The Economics of Data Externalities

Shota Ichihashi*

May 28, 2021

Abstract

A firm buys data from consumers to learn about some uncertain state of the world. There are data externalities, whereby data of some consumers reveal information about other consumers' data. I characterize data externalities that maximize or minimize consumer surplus and the firm's profit. I use the result to solve an information design problem in which the firm chooses what information to buy from consumers, balancing the value and price of information. The firm collects no less information than the efficient amount. In some cases we can solve the firm's data collection problem with a two-step concavification method.

Keywords: externalities; data markets; privacy; information design

^{*}Bank of Canada. Email: shotaichihashi@gmail.com. For valuable comments and suggestions, I thank the editor, the anonymous referee, and seminar participants at the Bank of Canada, Canadian Economic Theory Conference 2021, and the 6th Annual Conference on Network Science and Economics. The opinions expressed in this article are the author's own and do not reflect the views of the Bank of Canada.

1 Introduction

Digital platforms, such as Facebook and Google, collect data from users, learn their characteristics, and personalize services and advertising. Biotechnology companies collect genetic information to assess people's health risks. Carmakers collect driving data through vehicles to predict the behavior of a human driver. What is common in these examples is that a firm collects data from consumers to learn about some uncertain state of the world.

Motivated by the examples, I study a model of data collection: A firm wants to learn about an uncertain state of the world. There are *n* consumers, and each consumer *i* has data represented by a Blackwell experiment μ_i about the state. The game consists of two stages. First, the firm offers prices to buy data. Second, each consumer *i* decides whether to sell her data μ_i without observing the state or the signal drawn from μ_i . The payoffs of the firm and consumers depend on monetary transfers and what information the firm learns about the state—e.g., the firm may use data to learn about consumers' tastes and customize its product.

The main question is how data collection affects the firm and consumers. The key idea is that how the firm and consumers divide the surplus created by data depends on data externalities—i.e., what information each consumer's data reveal about other consumers' data. As a result, the impact of data collection depends on what kind of data externalities consumers impose on each other.

To highlight the idea, I begin with the following problem: Fix any experiment μ_0 about the state, which represents the aggregate data. A profile of *n* experiments (μ_1, \ldots, μ_n) , which I call the allocation of data, summarizes what data the consumers hold and how they are correlated. The allocation of data is feasible if it contains the same information about the state as aggregate data μ_0 . For any fixed μ_0 , I ask which feasible allocation of data maximizes or minimizes the equilibrium consumer surplus and the firm's profit.

Section 4 characterizes the best and worst allocations of data under a few different assumptions on consumer preferences. First, I assume that all consumers are worse off if the firm learns more about the state. Consumer welfare is minimized and the firm's profit is maximized if consumers hold "substitutable" data, where the data of any n - 1 consumers perfectly reveal information about the remaining consumer's data. In such a case, from each consumer *i*'s perspective, other consumers' data already reveal what *i* would prefer to hide. As a result, the firm can collect all data at a price of zero. In contrast, consumer surplus is maximized and the firm's profit is minimized if consumers hold "complementary" data, where the data of any n-1 consumers are totally uninformative, but the data of all n consumers jointly reveal the same information as μ_0 . The firm then compensates consumers for collecting data because each consumer's data contribution is pivotal. For any given μ_0 I construct feasible allocations of data that are substitutable or complementary. Second, if all consumers benefit from the firm learning the state, we obtain the opposite result—e.g., consumer welfare is maximized when they hold substitutable data. Finally, I extend the consumer-worst outcome to any consumer preferences.

Section 5 allows the firm to choose what information (μ_1, \ldots, μ_n) to buy from consumers without the feasibility constraint. The firm can buy data at lower prices by collecting more data and suitably designing data externalities between consumers. For any consumer preferences, the firm collects no less information about the state than the efficient amount, and the profit-maximizing data collection makes all consumers worse off compared to no data collection. Under a certain condition, we can solve the firm's problem with a concavification method.

Section 6 applies the results to a setting in which the firm uses data to price discriminate consumers in a product market. In the spirit of Bergemann et al. (2015b), I characterize all pairs of the firm's profits and consumer surpluses across all allocations of data. Data externalities drastically expand the set of possible outcomes.

The paper relates to recent work that studies the welfare impacts of data externalities (Easley et al., 2018; Acemoglu et al., 2019; Bergemann et al., 2019; Choi et al., 2019).¹ In particular, Acemoglu et al. (2019) and Bergemann et al. (2019) study a firm that buys information from consumers to learn about their types. There are two main differences between these papers and my paper. First, these papers assume that consumers' types and signals follow normal distributions. In contrast, I allow any information structure to study new questions, such as the firm's data collection problem under arbitrary consumer preferences and the characterization of the consumer-optimal outcome. Second, Acemoglu et al. (2019) and Bergemann et al. (2019) mainly consider data collection that harms consumers, but I allow more general consumer preferences. The relaxation reveals that a certain data externality protects consumers from the firm's monopsony power in the

¹Earlier works that consider externalities in information sharing include MacCarthy (2010) and Fairfield and Engel (2015).

data market.

The paper also relates to the broad literature on information markets. One branch of the literature studies the optimal collection and sales of personal data (e.g., Admati and Pfleiderer 1986, Taylor 2004, Calzolari and Pavan 2006, Eső and Szentes 2007, Babaioff et al. 2012, Bergemann et al. 2015a, Hörner and Skrzypacz 2016, Bergemann et al. 2018, Agarwal et al. 2019, Ichihashi forthcoming). Another branch studies the optimal use of data such as price discrimination and targeting (e.g., Conitzer et al. 2012, De Corniere and De Nijs 2016, Ali et al. 2019, Madio et al. 2019, Montes et al. 2019, Bonatti and Cisternas 2020, De Corniere and Taylor 2020). Relative to the literature I simplify how a firm uses data but focus on general data externalities.

2 Model

The set of consumers is $\mathcal{N} = \{1, 2, ..., n\}$. A firm wants to learn about the state of the world, $X \in \mathcal{X}$. The set \mathcal{X} is finite, and all players share a common prior belief about X. Given any finite set \mathcal{S} of realizations, I call any function $\mu : \mathcal{X} \to \Delta \mathcal{S}$ an *experiment*.² Let Σ denote the set of all experiments with finite realization spaces. Given any $\mu \in \Sigma$, let $\langle \mu \rangle \in \Delta \Delta \mathcal{X}$ denote the distribution of posteriors induced by the prior and μ . We say that μ is more informative than μ' if $\langle \mu \rangle$ is a mean preserving spread of $\langle \mu' \rangle$, and write it as $\mu \succeq \mu'$ (Blackwell, 1953).

The aggregate data is an experiment μ_0 . An allocation of data is a profile of n experiments $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n) : \mathcal{X} \to \Delta S^{\mathcal{N}}$, where μ_i represents consumer *i*'s data. Given $X \in \mathcal{X}$, realizations from $(\mu_1(X), \ldots, \mu_n(X))$ may not be independent. An allocation of data $\boldsymbol{\mu}$ is *feasible* (with respect to μ_0) if $\langle \boldsymbol{\mu} \rangle = \langle \mu_0 \rangle$. For any $\mu_0 \in \boldsymbol{\Sigma}$, let $\mathcal{F}(\mu_0)$ denote the set of all allocations of data that are feasible with respect to μ_0 . I describe the game by taking allocation $\boldsymbol{\mu}$ as exogenous, then study how the equilibrium depends on $\boldsymbol{\mu}$.

The game consists of two stages. In the first stage the firm chooses a price vector $\mathbf{p} = (p_1, \ldots, p_n) \in \mathbb{R}^n$, where p_i is the price offer to consumer *i*. Each consumer *i* privately observes p_i . A negative price $p_i < 0$ is a transfer from consumer *i* to the firm. In the second stage, all consumers simultaneously decide whether to sell their data. Specifically, let $a_i \in \{0, 1\}$ denote the data-sharing decision of consumer *i* with $a_i = 1$ corresponding to sharing. Denote the profile

 $^{^{2}\}Delta S$ denotes the set of all probability distributions over S.

of sharing decisions by $\mathbf{a} = (a_1, \dots, a_n)$. Let $\mathcal{N}_{\mathbf{a}} = \{i \in \mathcal{N} : a_i = 1\}$ denote the set of consumers who sell their data under \mathbf{a} . Given $\mathcal{N}_{\mathbf{a}}$ the firm acquires experiment $\boldsymbol{\mu}_{\mathbf{a}} = (\mu_i)_{i \in \mathcal{N}_{\mathbf{a}}} : \mathcal{X} \to \Delta \mathcal{S}^{\mathcal{N}_{\mathbf{a}}}$. All players move before X is realized and have no private information about X.

A profile of data-sharing decisions other than consumer *i* is denoted by $\mathbf{a}_{-i} \in \{0, 1\}^{n-1}$, and $\mathbf{1}_{-i}$ denotes \mathbf{a}_{-i} such that all consumers but *i* sell their data. For $a \in \{0, 1\}$, (a, \mathbf{a}_{-i}) denotes the profile of data-sharing actions such that consumer *i* chooses *a* and other consumers choose \mathbf{a}_{-i} . Finally, $\boldsymbol{\mu}_{-i}$ denotes $\boldsymbol{\mu}_{(0,1-i)}$.

All players maximize their expected payoffs, and ex post payoffs are as follows. If the firm collects data $\mu_{\mathbf{a}}$, it obtains a payoff of $\pi(\langle \mu_{\mathbf{a}} \rangle) - \sum_{i \in \mathcal{N}} a_i p_i$, and each consumer *i* obtains a payoff of $u_i(\langle \mu_{\mathbf{a}} \rangle) + a_i p_i$. The functions $\pi(\cdot)$ and $(u_i(\cdot))_{i \in \mathcal{N}}$ are defined on $\Delta \Delta \mathcal{X}$. For simplicity, write $\pi(\langle \mu \rangle)$ and $u_i(\langle \mu \rangle)$ as $\pi(\mu)$ and $u_i(\mu)$. We normalize $\pi(\mu_{\emptyset}) = u_i(\mu_{\emptyset}) = 0$, where μ_{\emptyset} is an uninformative experiment (i.e., $\langle \mu_{\emptyset} \rangle$ is degenerate at the prior).

The firm prefers more information: If $\mu \succeq \mu'$, then $\pi(\mu) \ge \pi(\mu')$. For example, suppose that the firm collects data, learns about the state, then chooses some payoff-relevant action. In such a case, $\pi(\mu)$ is the firm's expected payoff in the (unmodeled) decision problem when it acts optimally based on information μ .

The solution concept ("equilibrium") is perfect Bayesian equilibrium (PBE) in which consumers hold passive beliefs—i.e., each consumer *i*'s belief about \mathbf{p}_{-i} does not depend on what price the firm offers to *i*, on and off the equilibrium paths.³ Section 4.4 discusses the role of passive beliefs and the robustness of the results.

The following notions simplify exposition. Consumer surplus refers to the sum of the expected payoffs of all consumers. Given an allocation of data μ , let $\mathcal{E}(\mu)$ denote the set of all equilibria.

Definition 1. Fix any experiment $\mu_0 \in \Sigma$. An allocation of data μ^* maximizes (respectively, minimizes) consumer surplus with respect to μ_0 if $\mu^* \in \mathcal{F}(\mu_0)$, and there is an equilibrium $E^* \in \mathcal{E}(\mu^*)$ such that for any $\mu \in \mathcal{F}(\mu_0)$ and any $E \in \mathcal{E}(\mu)$, consumer surplus at E^* is weakly greater (respectively, smaller) than the one at E.

Analogously, I define an allocation of data that maximizes or minimizes the firm's expected payoff. Definition 1 fixes μ_0 , so we can compare two markets that have the same aggregate data

³I allow mixed strategy PBE, but the assumption of passive beliefs excludes mixed strategy equilibrium in which the firm offers prices that are not independent across consumers.

and different data externalities between consumers.

According to Definition 1, $\mathcal{E}(\mu^*)$ may contain multiple equilibria, and we select an equilibrium E that maximizes or minimizes consumer surplus. Section 4.4 discusses when the results do not depend on how we choose an equilibrium from $\mathcal{E}(\mu^*)$.

If there is a single consumer, the equilibrium is efficient and the firm extracts full surplus. The result follows from the standard argument of monopoly pricing with inelastic demand.

Claim 1. Suppose n = 1 and the consumer holds data μ_0 . In any equilibrium, consumer surplus is zero and the firm obtains a payoff of max $\{0, \pi(\mu_0) + u_1(\mu_0)\}$.

3 Substitutable and Complementary Allocations of Data

We now turn to the market with multiple consumers $(n \ge 2)$. I introduce two allocations of data that are useful for describing the results.

Definition 2. An allocation of data μ is *perfectly substitutable* if for any $i \in \mathcal{N}$, $\langle \mu \rangle = \langle \mu_{-i} \rangle$.

Definition 3. An allocation of data μ is *perfectly complementary* if for any $i \in \mathcal{N}$, $\langle \mu_{-i} \rangle = \langle \mu_{\emptyset} \rangle$, where μ_{\emptyset} is an uninformative experiment.

An allocation of data is perfectly substitutable if the marginal value of individual data is zero. It captures a situation in which a firm can perfectly learn about one consumer from the data of other consumers. A perfectly complementary allocation of data is such that the marginal value of individual data equals the value of the entire dataset. In such a case, the dataset is valueless if the data of any single consumer is missing. Perfect complementarity captures increasing returns to scale, whereby the data of some consumers increase the marginal value of data on other consumers.⁴ If n = 2, the definitions satisfy the complementarity and substitutability of experiments in Börgers et al. (2013). For any aggregate data μ_0 we can find a feasible allocation of data that is perfectly substitutable or complementary.

Lemma 1. Suppose $n \ge 2$, and take any experiment $\mu_0 \in \Sigma$ as the aggregate data.

⁴Arrieta-Ibarra et al. (2018) offer an insightful discussion on when data may exhibit increasing or decreasing returns to scale.

1. There is a feasible and perfectly substitutable allocation of data.

2. There is a feasible and perfectly complementary allocation of data.

Proof. Take any experiment $\mu_0 : \mathcal{X} \to \Delta \mathcal{Y}$. For Point 1, take an allocation of data $(\mu_1^*, \dots, \mu_n^*)$ such that $\mu_i^* = \mu_0$ for all $i \in \mathcal{N}$. We have $\langle \boldsymbol{\mu}^* \rangle = \langle \boldsymbol{\mu}_{-i}^* \rangle = \langle \mu_0 \rangle$ because the firm observes the same realization $Y \in \mathcal{Y}$ across all μ_i^* 's with probability 1.

For Point 2, we use the secret sharing algorithm of Shamir (1979). It provides a set S and a function $\boldsymbol{\nu} : \mathcal{Y} \to \Delta S^n$ such that for any distribution over $\mathcal{Y}, \boldsymbol{\nu}_{-i}$ is uninformative about a realized $Y \in \mathcal{Y}$ for any $i \in \mathcal{N}$, but $\boldsymbol{\nu}$ is perfectly informative about Y (here, $\boldsymbol{\nu}_{-i}$ is the experiment created from $\boldsymbol{\nu}$ by omitting the *i*-th experiment). Define $\boldsymbol{\mu}^* : \mathcal{X} \to \Delta S^n$ as a composite of μ_0 and $\boldsymbol{\nu}$: For any $X \in \mathcal{X}, \boldsymbol{\mu}^*(X)$ draws Y according to $\mu_0(X) \in \Delta \mathcal{Y}$, then draws $(S_1, \ldots, S_n) \in S^n$ according to $\boldsymbol{\nu}(Y) \in \Delta S^n$. Consumer *i*'s data μ_i^* reveals S_i . The experiment $\boldsymbol{\mu}^*$ is perfectly complementary and satisfies $\langle \boldsymbol{\mu}^* \rangle = \langle \mu_0 \rangle$.

If n = 2, we can construct complementary signals as follows.⁵ Take any $\mu_0 : \mathcal{X} \to \Delta \mathcal{Y}$ with $\mathcal{Y} = \{1, 2, ..., m\}$. We decompose μ_0 into two experiments, μ_0^1 and μ_0^2 : First, independently of the state X, μ_0^1 draws a permutation α of 1, 2, ..., m uniformly randomly from the set of all the permutations. Second, given the realized permutation α , μ_0^2 draws signal realization $\alpha(Y)$ whenever μ_0 draws $Y \in \mathcal{Y}$. The allocation (μ_0^1, μ_0^2) is perfectly complementary and contains the same information as μ_0 .⁶ Intuitively, μ_0^1 provides a dictionary that interprets the signal drawn from μ_0^2 , so we need μ_0^1 to learn about the state from μ_0^2 .

An economic example of a complementary allocation is as follows: Suppose consumer 1 is a seller, whose type θ_S is -1 or 1. Consumer 2 is a buyer, whose type θ_B is -1 or 1. For example, θ_S is the horizontal characteristics of the seller's product, and θ_B is the buyer's taste. A firm, such as a platform, tries to learn the match quality, $X = \theta_S \cdot \theta_B$. Suppose θ_S and θ_B are independently distributed with the uniform prior $\mathbf{P}(\theta_S = 1) = \mathbf{P}(\theta_B = 1) = 0.5$. If the buyer and the seller have signals that respectively reveal θ_B and θ_S , it is a perfectly complementary allocation: Knowing θ_B or θ_S alone provides no information about the match quality.

⁵Shamir's algorithm uses the polynomial interpolation.

⁶Suppose we only have μ_0^2 and observe realization $Y \in \mathcal{Y}$. The conditional probability of X given Y is $\mathbf{P}^{\mathcal{X}}(X|Y) = \frac{\mathbf{P}^{\mathcal{X}}(X) \cdot \left[\frac{1}{m} \mu^0(1|X) + \dots + \frac{1}{m} \mu^0(m|X)\right]}{\frac{1}{m} \mu^0(1) + \dots + \frac{1}{m} \mu^0(m)} = \mathbf{P}^{\mathcal{X}}(X)$. Here, $\mu^0(Y)$ is the ex ante probability of Y under μ^0 , and $\mathbf{P}^{\mathcal{X}}(X)$ is the prior probability of X.

Finally, the following example, motivated by Bergemann et al. (2019), illustrates substitutability and complementarity of data. Suppose $X = (\theta_1, \ldots, \theta_n)$, where θ_i denotes consumer *i*'s type the firm wishes to learn. Each consumer has data that reveal a noisy signal of her type: $s_i = \theta_i + \epsilon_i$.⁷ First, suppose that consumers have a common type and the data contain idiosyncratic noise: $\theta_i = \theta$ for all i, and ϵ_i is independent across i. For example, θ represents a demand parameter for the firm's product. The purchase history of each consumer serves as a signal of θ but contains noise such as her idiosyncratic taste. For a large n the firm can use the data of n-1 consumers to accurately estimate θ , which approximates perfect substitutability. Second, suppose consumers have idiosyncratic types and the data contain a common noise: θ_i is independent across *i*, and $\epsilon_i = \epsilon$ for all *i*. For example, consumers have independent values for some product, and the firm can learn about the values from purchase histories. The data contain noise such as common traits or trends that influence consumers' behavior but are orthogonal to their product values. If the noise has a large variance, the data of consumer i is nearly uninformative about θ_i , and the data of consumers -ireveal nothing about θ_i . However, for a large n the firm can combine all data to estimate ϵ and then calculate $\theta_i = s_i - \varepsilon$. Thus the data of consumer *i* and that of consumers -i are complementary in learning about θ_i .⁸

4 Welfare Implications of Data Externalities

Lemma 1 helps us find an allocation of data that maximizes or minimizes consumer surplus and the firm's profit. I first assume monotone consumer preferences, then generalize a part of the results. Throughout the section we assume $n \ge 2$. All omitted proofs are in the Appendix.

4.1 Harmful Data Collection

Assume that data collection harms consumers: For any experiments $\mu, \mu' \in \Sigma$ such that $\mu \succeq \mu'$, $u_i(\mu) \leq u_i(\mu') \leq 0$ for all $i \in \mathcal{N}$.

⁷Assume θ_i and ε_i are independent. Bergemann et al. (2019) use a Gaussian information structure.

⁸The example does not formally capture Definition 3 in the following sense: Perfect complementarity requires that the firm learns nothing about $X = (\theta_1, \ldots, \theta_n)$ once the data of any consumer *i* is missing. In the current example, if *n* is large, the data of consumers -i still provide information about θ_{-i} . Nonetheless, the example captures complementarity between the data of consumer *i* and that of consumers -i upon learning about θ_i .

Proposition 1. For any aggregate data $\mu_0 \in \Sigma$, a perfectly substitutable allocation of data μ^* minimizes consumer surplus and maximizes the firm's profit with respect to μ_0 . Given μ^* , the firm's profit is $\pi(\mu_0)$, consumer surplus is $\sum_{i \in \mathcal{N}} u_i(\mu_0)$, and the prices of data are zero.

Suppose that the allocation of data is substitutable and n - 1 consumers sell their data. The remaining consumer *i* is willing to give up her data for free, because she correctly believes that her data contribution has no marginal effect on what the firm learns about the state. As a result the firm buys data from every consumer for free, even though data collection harms consumers.

Conversely, if the allocation of data is complementary and n - 1 consumers sell their data, the remaining consumer *i* faces a private cost of $u_i(\mu_0) < 0$ from selling her data. The firm then compensates each consumer according to their loss of data collection, leading to zero consumer surplus. The outcome is best for consumers and worst for the firm, because the firm never leaves a positive surplus to consumers when data collection harms them.

Proposition 2. For any aggregate data μ_0 such that $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0) \ge 0$, a perfectly complementary allocation of data μ^* maximizes consumer surplus and minimizes the firm's profit with respect to μ_0 . Under μ^* , the firm pays $-u_i(\mu_0) \ge 0$ to each consumer *i*. Consumer surplus is zero and the firm's profit is $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0)$.

Section 4.4 discusses the case of $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0) < 0.$

4.2 Beneficial Data Collection

Assume now that consumers are better off if the firm has more data: For any $\mu, \mu' \in \Sigma$ such that $\mu \succeq \mu', u_i(\mu) \ge u_i(\mu') \ge 0$ for each $i \in \mathcal{N}$. The firm may now charge a negative price to extract surplus from each consumer, but she can always retain her data to secure a non-negative payoff. The best and worst allocations of data are the mirror images of those under harmful data collection.

Proposition 3. For any aggregate data $\mu_0 \in \Sigma$, a perfectly complementary allocation of data μ^* minimizes consumer surplus and maximizes the firm's profit with respect to μ_0 . Under μ^* , each consumer *i* pays $u_i(\mu_0) \ge 0$ and the firm extracts full surplus $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0)$.

Proposition 4. For any aggregate data $\mu_0 \in \Sigma$, a perfectly substitutable allocation of data μ^* maximizes consumer surplus and minimizes the firm's profit. Under μ^* , the firm collects data at a price of zero. Consumer surplus is $\sum_{i \in \mathcal{N}} u_i(\mu_0)$ and the firm's profit is $\pi(\mu_0)$. Proposition 4 resembles the free-rider problem. When the allocation of data is substitutable, consumers have low willingness to pay for having their data collected, provided other consumers sell their data. If prices were exogenous, the incentive to free-ride would inefficiently lower the level of data provision. However, prices are now endogenous, so the free-riding prevents the firm from setting negative prices for collecting data. As a result, the data externality protects consumers from the firm's monopsony power.⁹

Remark 1 (The set of possible payoffs). The above results characterize the highest and lowest payoffs of each player across all feasible allocations of data. A natural question is whether we can use the results to characterize the entire set of possible payoffs. The answer is trivially yes if the allocation of data can depend on some public randomization device. In such a case we can use a random allocation of data to attain any payoff between the highest and lowest payoffs for each player.¹⁰ Appendix D further shows that public randomization is unnecessary: For each player, we can construct an allocation of data to attain any payoff between the bounds. However the results do not extend to the joint characterization of possible *payoff vectors*. To characterize the payoff of each player separately, we could focus on equilibria in which the firm collects data from all consumers. In contrast, to characterize the set of payoff vectors, we may need to study an equilibrium in which the firm does not collect data from some consumers. Appendix D shows an example in which we can attain some payoff profile only when the firm does not collect data from a consumer who incurs a high cost of selling data. In general the firm's incentive to collect data depends on the shape of $(u_i(\cdot))_{i\in\mathcal{N}}$ and $\pi(\cdot)$ even for monotone preferences, and the characterization of possible payoff profiles is left for future research.

Remark 2 (Non-monotone, separable payoffs). The results under monotone preferences extend to the following setting. Suppose we can write the state as $X = (X_G, X_B)$, where X_G and X_B are independent. Consumers prefer the firm to learn about X_G but not about X_B . Formally, if the firm acquires information μ_G and μ_B about X_G and X_B , respectively, then consumer *i* receives a gross

⁹The logic is similar to how free-riding by shareholders prevents the raider from capturing surplus in corporate takeover (cf. Tirole 2010).

¹⁰As an example, suppose data collection is harmful. Suppose that the allocation of data is perfectly substitutable with probability α and complementary with probability $1 - \alpha$ and that this realization is observable before the firm makes an offer. By Propositions 1 and 2, the expected payoff of the firm becomes $\alpha \pi(\mu_0) + (1 - \alpha) \{\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0)\} = \pi(\mu_0) + (1 - \alpha) \sum_{i \in \mathcal{N}} u_i(\mu_0)$. By varying α between 0 and 1, we can attain any payoff of the firm between the best outcome and the worst outcome.

payoff of $u_i^G(\mu_G) + u_i^B(\mu_B)$. For each $i, u_i^G(\mu_G)$ increases in the informativeness of μ_G and $u_i^B(\mu_B)$ decreases in that of μ_B . Propositions 1 - 4 extend to such a setting. For example, the consumeroptimal allocation is the one in which consumers hold information about X_G as a substitutable allocation and information about X_B as a complementary allocation. Similarly, the firm-optimal allocation is the one in which consumers hold information about X_G as a complementary allocation and X_B as a substitutable allocation. Appendix E provides details.

4.3 The Consumer-Worst Outcome Under General Preferences

I extend the results on the consumer-worst outcome to arbitrary preferences.

Proposition 5. Take any $(u_i(\cdot))_{i\in\mathcal{N}}$ and any aggregate data μ_0 . There is an allocation of data μ^* that minimizes consumer surplus with respect to μ_0 . Under μ^* , the firm collects μ^* and each consumer *i* receives a payoff of $\min_{\mu \leq \mu_0} u_i(\mu) \leq 0$.

Proof. In any equilibrium, consumer i can refuse to sell data and secure a payoff of $\mathbb{E}[u_i(\mu_{-i})]$, where μ_{-i} is an experiment the firm collects from other consumers. The expectation operator is relevant for a mixed strategy equilibrium in which μ_{-i} is random. Because μ_{-i} is less informative than μ_0 , we have $\mathbb{E}[u_i(\mu_{-i})] \geq \min_{\mu \leq \mu_0} u_i(\mu)$. As a result, $\min_{\mu \leq \mu_0} u_i(\mu)$ is a lower bound of consumer i's payoff across all feasible allocations of data and all equilibria. I construct μ^* that achieves $\min_{\mu \prec \mu_0} u_i(\mu)$ as an equilibrium payoff of *i*. For each $i \in \mathcal{N}$, pick an experiment $\mu_i^{MIN} \in$ $\arg\min_{\mu \leq \mu_0} u_i(\mu)$.¹¹ Let $\boldsymbol{\nu}_{-i}^i = (\nu_i^i)_{j \in \mathcal{N} \setminus \{i\}}$ denote a perfectly complementary allocation of data for consumers in $\mathcal{N} \setminus \{i\}$ such that $\langle \boldsymbol{\nu}_{-i}^i \rangle = \langle \mu_i^{MIN} \rangle$. Let $\boldsymbol{\nu}^*$ denote a perfectly complementary allocation for n consumers such that $\langle \boldsymbol{\nu}^* \rangle = \langle \mu_0 \rangle$. Consider the allocation of data $\boldsymbol{\mu}^*$ such that each consumer j has $(\nu_j^i)_{i \in \mathcal{N} \setminus \{j\}}$ and ν_j^* . Consider the strategy profile in which the firm offers $p_i^* := -u_i(\mu_0) + u_i(\mu_i^{MIN})$ to each *i*, and all consumers sell their data. It is optimal for consumer *i* to sell her data: If *i* does not sell data, her payoff is $u_i(\mu_i^{MIN})$ because other n-1 consumers sell data and the firm obtains ν_{-i}^i . If i sells data, her gross utility is $u_i(\mu_0)$. Thus, p_i^* is the maximum amount that i is willing to pay. The firm has no profitable deviation, because $p_i^* \leq 0$ holds for all *i*, and the firm cannot lower the price.

¹¹Remark 4 provides a condition under which $\arg \min_{\mu \preceq \mu_0} u_i(\mu)$ exists.

The worst outcome for consumer *i* is that other consumers hold information $\min_{\mu \leq \mu_0} u_i(\mu)$, which minimizes *i*'s payoff when she refuses to sell her data. The low outside option enables the firm to collect her data at a low price. The proof constructs a feasible allocation of data that simultaneously minimizes the outside options of all consumers. The result implies that regardless of the social value of data collection, it can harm all consumers under a certain data externality.

4.4 Discussion on the Multiplicity of Equilibria and Passive Beliefs

Definition 1 selects an equilibrium $E^* \in \mathcal{E}(\mu^*)$ that attains a higher consumer surplus than any equilibrium under any feasible allocations. The definition leaves the possibility that $\mathcal{E}(\mu^*)$ contains another equilibrium with a low consumer surplus. For example, Proposition 4 selects an equilibrium in which the firm sets $p_i = 0$ for all *i*. In another equilibrium the firm charges a negative price of $-u_i(\mu_0)$ to consumer *i* and a price of zero to others, leading to a strictly lower consumer surplus than the candidate equilibrium.

However the results are not sensitive to the equilibrium selection in the following sense. If data collection is beneficial, we can construct a sequence $(\mu^k)_{k\in\mathbb{N}}$ of feasible allocations such that it converges to μ^* (in Proposition 3 or Proposition 4) and a unique equilibrium exists under each μ^k . Thus we can approximate the consumer or the firm-optimal outcome with an allocation of data that has a unique equilibrium (see Appendix B for the proof).

If data collection is harmful and $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0) > 0$, the equilibrium is unique under the consumer-optimal allocation of data in Proposition 2. Under the consumer-worst allocation in Proposition 1, the equilibrium is unique if $\pi(\mu_0) + u_i(\mu_0) > 0$ for some *i*. Given the inequality, the firm collects data from at least one consumer. Other consumers are then willing to sell their data for free, given the perfectly substitutable allocation. As a result the firm collects data at a price of zero in any equilibrium.

Another restriction is that I compare perfect Bayesian equilibria with passive beliefs. Nonpassive beliefs can introduce other equilibria. For example, consumers can sustain high prices with a belief system such that if the firm deviates to a low price, each consumer believes that the firm collects data from other consumers in a way that increases her cost of selling data. Nonetheless except for the consumer-optimal outcome under harmful data collection, all the results remain the same even if we allow non-passive beliefs (see Appendix C).

Beyond the settings we examined above, allowing non-passive belief is sometimes necessary to ensure the existence of equilibrium. Consider the case of $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0) < 0$ in Proposition 2. There is no equilibrium in which the firm collects data μ_0 , because it will then pay $-u_i(\mu_0)$ to each consumer *i* and earn a negative profit. However, it cannot be an equilibrium that all consumers refuse to sell data; given the passive beliefs, the firm will then deviate and collect the data of each consumer at a small positive price. Indeed, under a mild condition there is no equilibrium with passive beliefs under a complementary allocation of data. At the same time, once we relax the passive belief assumption, we can find an equilibrium with no data collection:

Claim 2. Assume that data collection is harmful and the allocation of data is perfectly complementary. Suppose $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0) < 0$. There is a perfect Bayesian equilibrium in which consumers do not have passive beliefs, the firm collects no data, and all players obtain zero payoffs. In addition, if $\pi(\mu_0) + \min_{i \in \mathcal{N}} u_i(\mu_0) > 0$, there is no equilibrium with passive beliefs.

Remark 3 (Restriction on the firm's strategy). We have assumed that the firm chooses whether to collect each consumer's data μ_i but it cannot request part of the data (i.e., a garbling of μ_i). Depending on the allocation of data the firm may benefit from doing so: For example, suppose n = 1 and the consumer holds μ_0 such that $\pi(\mu_0) + u_1(\mu_0) < 0$, but there is some μ that is less informative than μ_0 and satisfies $\pi(\mu) + u_1(\mu) > 0$. The current assumption implies that the firm does not collect any data, but it could earn a positive profit by requesting μ . At the same time, most of the above results are robust to the setting in which the firm can request any $\mu \leq \mu_i$ from each consumer *i*.¹² For example, suppose data collection is harmful and consumers hold substitutable data. In the equilibrium of Proposition 1 the firm collects all data at the lowest possible price of zero. In such an equilibrium the firm cannot increase profit by requesting part of the data. The equilibrium maximizes the firm's profit and minimizes consumer surplus, even though the firm could collect partial information under other allocations of data. Similarly, Propositions 3, 4, and 5 extend. In contrast, Proposition 2 might change if the firm could request partial data.

¹²Formally, the firm can offer to collect any $\hat{\mu}_i \preceq \mu_i$ from consumer *i* in exchange for p_i , and consumers have passive beliefs regarding the firm's offers. The informativeness $\hat{\mu}_i \preceq \mu_i$ means that we can obtain $\hat{\mu}_i$ by garbling signal realizations of μ_i .

equal to $\pi(\mu_0) + \sum_i u_i(\mu_0)$, which is the total surplus from collecting μ_0 . If $\pi(\mu) + \sum_i u_i(\mu)$ is maximized at some μ' that is less informative than μ_0 , the firm would be better off by committing to garble the information it collects from consumers.

5 The Firm-Optimal Data Collection

I now study the firm's data collection problem, in which it can request any experiments from any consumers. Although it is a strong assumption that the firm can potentially source any information, we may view the problem as the firm's first-best benchmark. The problem is equivalent to finding an allocation of data that maximizes a profit without any feasibility constraint.

Definition 4. An allocation of data μ^* globally maximizes the firm's profit if there is an equilibrium $E^* \in \mathcal{E}(\mu^*)$ such that for any allocation of data μ and any equilibrium $E \in \mathcal{E}(\mu)$, the firm's expected payoff in E^* is weakly higher than the one in E.

If there is a single consumer, any welfare-maximizing experiment $\mu^* \in \arg \max_{\mu \in \Sigma} \pi(\mu) + u_1(\mu)$ globally maximizes the firm's profit (see Claim 1). For $n \ge 2$, the following result characterizes the firm-optimal outcome (I assume that the relevant optimization problems have solutions; Remark 4 provides sufficient conditions for it).

Proposition 6. Suppose $n \ge 2$ and take any $(u_i(\cdot))_{i\in\mathcal{N}}$. Let $\mu_0^* \in \Sigma$ solve

$$\max_{\mu \in \mathbf{\Sigma}} \left(\pi(\mu) + \sum_{i \in \mathcal{N}} u_i(\mu) - \sum_{i \in \mathcal{N}} \min_{\mu_i \preceq \mu} u_i(\mu_i) \right).$$
(1)

There is an allocation μ^* that satisfies $\langle \mu^* \rangle = \langle \mu_0^* \rangle$ and globally maximizes the firm's profit. The maximum equilibrium payoff of the firm under μ^* is (1).

Proof. In any equilibrium at which the firm acquires $\mu \in \Sigma$, the firm's payoff is total surplus from μ minus consumer surplus. Proposition 5 implies that across all such situations, the lowest consumer surplus is $\sum_{i\in\mathcal{N}} \min_{\mu_i \preceq \mu} u_i(\mu_i)$. The firm's maximum profit conditional on collecting μ is then $\pi(\mu) + \sum_{i\in\mathcal{N}} u_i(\mu) - \sum_{i\in\mathcal{N}} \min_{\mu_i \preceq \mu} u_i(\mu_i)$. The firm can optimize across all $\mu \in \Sigma$, leading to the optimal profit (1). Compared to the welfare-maximizing social planner, the firm's profit (1) contains an extra term $-\sum_{i\in\mathcal{N}}\min_{\mu_i\preceq\mu}u_i(\mu_i)$, which is increasing in the informativeness of μ . It captures the firm's inefficient incentive to collect data: By collecting more data and designing the signal structure properly, the firm can lower the payoff of consumer *i* from refusing to sell data. The firm can then collect *i*'s data at a lower price. As a result, the firm tends to collect too much information, and the firm-optimal data collection never benefits consumers.

Corollary 1. The firm-optimal allocation of data in *Proposition 6* has the following properties.

- 1. There is no efficient experiment $\mu_E^* \in \arg \max_{\mu \in \Sigma} \pi(\mu) + \sum_{i \in \mathcal{N}} u_i(\mu)$ that is strictly more informative than all firm-optimal allocations of data.
- 2. Compared to no data collection, the firm-optimal data collection weakly decreases the payoffs of all consumers.

Proof. For Point 1, suppose that some firm-optimal allocation μ^* is strictly less informative than μ_E^* . Because μ_E^* maximizes total surplus, we have

$$\pi(\boldsymbol{\mu}^*) + \sum_{i \in \mathcal{N}} u_i(\boldsymbol{\mu}^*) \le \pi(\mu_E^*) + \sum_{i \in \mathcal{N}} u_i(\mu_E^*).$$

Because $\mu_E^* \succeq \mu^*$, we have $\min_{\mu' \preceq \mu_E^*} u_i(\mu') \leq \min_{\mu' \preceq \mu^*} u_i(\mu')$, so

$$\pi(\boldsymbol{\mu}^*) + \sum_{i \in \mathcal{N}} u_i(\boldsymbol{\mu}^*) - \sum_{i \in \mathcal{N}} \min_{\mu' \preceq \boldsymbol{\mu}^*} u_i(\mu') \leq \pi(\mu_E^*) + \sum_{i \in \mathcal{N}} u_i(\mu_E^*) - \sum_{i \in \mathcal{N}} \min_{\mu' \preceq \mu_E^*} u_i(\mu').$$

The inequality implies that μ_E^* is also a firm-optimal information, which completes the proof. Point 2 holds because under the allocation of data that globally maximizes the firm's payoff, each consumer obtains an equilibrium payoff of $\min_{\mu_i \leq \mu^*} u_i(\mu_i) \leq 0$.

I now turn to solving the firm's problem (1). First, if data collection harms some consumers and benefits others, the firm collects full information:

Corollary 2. If each $u_i(\cdot)$ is either increasing or decreasing in Blackwell's ordering \succeq , a solution to the firm's problem (1) is a fully informative experiment.

Proof. If $u_i(\cdot)$ is increasing in \succeq , $\min_{\mu_i \leq \mu} u_i(\mu_i) = u_i(\mu_{\emptyset}) = 0$. If $u_i(\cdot)$ is decreasing in \succeq , $\min_{\mu_i \leq \mu} u_i(\mu_i) = u_i(\mu)$. As a result the maximand in (1) is $\Pi(\mu) := \pi(\mu) + \sum_{i \in \mathcal{N}_+} u_i(\mu)$, where \mathcal{N}_+ is the set of *i*'s such that $u_i(\cdot)$ is increasing. Because $\Pi(\cdot)$ is increasing in \succeq , the fully informative experiment solves the problem (1).

Second, if all consumers have the same preferences, we can solve the firm's problem (1) using concavification (e.g., Aumann and Maschler 1995 and Kamenica and Gentzkow 2011). To state the result, we assume that there are functions $\hat{\pi} : \Delta \mathcal{X} \to \mathbb{R}$ and $\hat{u} : \Delta \mathcal{X} \to \mathbb{R}$ such that for each $\mu \in \Sigma$, we have $\pi(\mu) = \int_{\Delta \mathcal{X}} \hat{\pi}(b) d\langle \mu \rangle(b)$ and $u(\mu) = \int_{\Delta \mathcal{X}} \hat{u}(b) d\langle \mu \rangle(b)$.¹³ For simplicity, we identify π and u with $\hat{\pi}$ and \hat{u} . For any function $f : \Delta \mathcal{X} \to \mathbb{R}$, let $\mathcal{V}[f]$ denote the concavification of f, and let $\mathcal{V}[f](b)$ denote the concavification evaluated at $b \in \Delta \mathcal{X}$.¹⁴

Corollary 3. Assume all of the $n \ge 2$ consumers have the same utility function, $\frac{1}{n}u(\cdot)$. Given the common prior $b_0 \in \Delta \mathcal{X}$, the firm's payoff (1) under the optimal allocation of data is

$$\mathcal{V}\left[\mathcal{V}[\pi+u]-u\right](b_0).\tag{2}$$

Proof. Given the common preferences, we can write the firm's problem (1) as

$$\max_{\nu \in \Sigma} \left\{ \max_{\mu \succeq \nu} \left(\pi(\mu) + u(\mu) - u(\nu) \right) \right\}.$$
(3)

The term $\max_{\mu \succeq \nu} (\pi(\mu) + u(\mu) - u(\nu))$ is an information design problem in which the designer chooses μ that is more informative than ν to maximize $\pi(\mu) + u(\mu) - u(\nu)$. Appendix F shows

$$\max_{\mu \succeq \nu} \left(\pi(\mu) + u(\mu) - u(\nu) \right) = \int_{\Delta \mathcal{X}} \left\{ \mathcal{V}[\pi + u](b) - u(b) \right\} d\langle \nu \rangle(b).$$
(4)

As a result, (3) is written as

$$\max_{\nu \in \Sigma} \int_{\Delta \mathcal{X}} \left\{ \mathcal{V}[\pi + u](b) - u(b) \right\} d\langle \nu \rangle(b),$$
(5)

¹³Functions $\hat{\pi}$ and \hat{u} exist, for example, if the firm chooses some action a after learning about X from the information collected, and the expost payoff of each player depends only on (a, X).

 $^{^{14}}$ A concavification of f is the smallest concave function that is everywhere weakly greater than f. For details, see Kamenica and Gentzkow (2011).

which equals (2).

Under common preferences, we can obtain the firm's optimal profit by (i) concavifying total surplus, (ii) subtracting consumer surplus from (i), then (iii) concavifying (ii) again and evaluating it at the prior. Step (iii) identifies the information the firm collects from all but one consumer, which determines each consumer's outside option. Step (i) then identifies the information the firm collects on the equilibrium path, where all consumers sell their data. Generally the solution to Step (iii) depends on the identity i of the consumer who refuses to sell data, but the assumption of common preferences enables the firm to use the same information as the outside option of all consumers. As a result the firm's problem of choosing n offers becomes information design that determines what information the firm will acquire if all consumers sell their data and if all but one consumer do so. We apply Corollary 3 to an example.

Example 1. The state space is binary, i.e., $\mathcal{X} = \{0, 1\}$. We identify $\Delta \mathcal{X}$ with [0, 1], where $b \in [0, 1]$ is the probability of X = 1. Assume $\pi(\cdot) \equiv 0$, and all consumers have $\frac{1}{n}u(b)$. In Figure 1, the black and red solid lines depict u and $\mathcal{V}[u]$ evaluated at each $b \in [0, 1]$. Figure 2 depicts $\mathcal{V}[u](b) - u(b)$ and its concavification, $\mathcal{V}[\mathcal{V}[u] - u](b)$.



Figure 1: Consumer surplus u and its concavification $\mathcal{V}[u]$.



Figure 2: Function $\mathcal{V}[u] - u$ and its concavification.

Corollary 3 states that across all allocations of data and equilibria, the maximum payoff of the firm at prior b_0 is $\mathcal{V}[\mathcal{V}[u] - u](b_0)$ (see Figure 2). We can also derive the firm-optimal signal (i.e., μ_0^* in Proposition 6) from the two-step concavification. First, the concavification in Figure 2, which concavifies $\mathcal{V}[u] - u$, splits the prior b_0 into b_L and b_H . Second, Figure 1, which concavifies u, splits b_L into 0 and β_L , and splits b_H into β_H and 1. As a result, the firm-optimal signal generates four posteriors, $0, \beta_L, \beta_H$, and 1. Under the firm-optimal allocation, the data of any n - 1 consumers induce posteriors b_L and b_H . These posteriors do not arise on the equilibrium path, but decrease the outside option of a consumer from refusing to sell her data. Finally, the firm-optimal signal is strictly more informative than any signal that maximizes total surplus, which induces posteriors in $[\beta_L, \beta_H]$. The observation conforms to Corollary 1.

Remark 4. Appendix G provides conditions under which the firm's problem (1) has a solution. In particular, the appendix studies a setting in which the firm learns about the state from collected information, then takes a payoff-relevant action. In such a setting, $\pi(\mu)$ and $(u_i(\mu))_{i\in\mathcal{N}}$ are the equilibrium payoffs of the subgame in which the firm chooses an action based on the collected information μ . The problem (1) has a solution, for example, if the firm has finitely many actions and chooses an action according to the following tie-breaking rule: First, the firm breaks ties to maximize the sum of the payoffs of all players on the equilibrium path. Second, it breaks ties to minimize the payoff of consumer *i* whenever she unilaterally deviates and refuses to sell the data. Such a tie-breaking rule ensures that a version of (1) has a solution, which globally maximizes the firm's profit.

6 Application: Monopoly Price Discrimination

I now apply the results to a setting in which the firm uses data for price discrimination. The firm sells a good to consumers, each of whom demands one unit. The production cost is zero. The consumers have a common value of X to the good.¹⁵ The product value X has a finite support and is positive with probability 1.

Given the allocation of data μ about X, the firm and consumers play the following game: First, in the *data market* the firm chooses a price vector **p**, and each consumer *i* decides whether to sell data μ_i . Second, in the *product market* the firm updates its belief about X based on the information collected in the data market, then sets a product price t. Finally, consumers observe X and make (identical) purchase decisions. It is without loss of generality that the firm sets the same product price across all consumers.

Suppose that the firm pays the total amount of P for data and sets a product price of t, and m consumers buy goods. In ex post terms, the firm's profit and consumer surplus are mt - P and m(X - t) + P. The average firm profit and the average consumer surplus are $\frac{1}{n}(mt - P)$ and $\frac{1}{n}[m(X - t) + P]$.

Let $\bar{w} := \mathbb{E}[X]$ denote the average total surplus under the efficient outcome. Let u_{\emptyset} and π_{\emptyset} denote the average expected consumer surplus and the average firm profit, when the firm buys no information and all players behave optimally in the product market. For simplicity, assume that the optimal product price given no information is unique, so that u_{\emptyset} is unique.

I characterize all possible outcomes across all allocations of data. If there is a single consumer the firm can buy data at a price that makes her indifferent between selling and not selling data. If she does not sell data, her payoff in the product market is u_{\emptyset} . As a result the consumer's net equilibrium payoff is always u_{\emptyset} ; the data affect only the firm's profit.

Claim 3. Suppose n = 1. The following two conditions are equivalent.

- 1. There is an allocation of data (which is equal to the aggregate data) such that the equilibrium payoffs of the firm and the consumer are π^* and u^* , respectively.
- 2. $u^* = u_{\emptyset}$ and $\pi_{\emptyset} \leq \pi^* \leq \bar{w} u_{\emptyset}$.

¹⁵The common value assumption simplifies exposition. The same result holds for independent and private values, provided we allow allocations of data such that each consumer has information about the values of other consumers.

Data externalities between multiple consumers expand the set of possible outcomes. To state the main result, define the surplus triangle as follows.

$$\Delta := \left\{ (\pi, u) \in \mathbb{R}^2 : \pi + u \le \bar{w}, u \ge 0, \pi \ge \pi_{\emptyset} \right\}.$$
(6)

If there are multiple consumers, any outcome in Δ can arise.

Proposition 7. Suppose $n \ge 2$. A pair (π, u) of the average profit and the average consumer surplus can arise in some equilibrium given some allocation of data if and only if $(\pi, u) \in \Delta$.

Proof. Suppose $n \ge 2$. To show the "if" part, take any $(\pi^*, u^*) \in \Delta$. Bergemann et al. (2015b) construct a $\mu^* \in \Sigma$ such that if the firm has μ^* , the resulting average outcome in the product market is (π^*, u^*) . Suppose $u^* < u_{\emptyset}$ (resp. $u^* \ge u_{\emptyset}$). Proposition 1 (resp. Proposition 4) provides an allocation of data such that the firm collects μ^* at a price of zero. In the equilibrium, (π^*, u^*) arises as the net average payoffs of the firm and consumers. The "only if" part holds because consumers can secure zero payoffs by selling no data and buying nothing, and the firm can secure π_{\emptyset} by obtaining no data and set an optimal price given the prior.

Figure 3 depicts the possible outcomes for n = 1 and $n \ge 2$. The surplus triangle Δ is *AEC*. The segment *EC* represents the firm's profit from no data, and *AC* describes the total surplus from the efficient allocation. All the values are in terms of the average across consumers. If the market consists of a single consumer, the possible outcomes correspond to the segment *BD*, so the consumer never benefits from data. In contrast if the market consists of multiple consumers, any outcome in *AEC* can arise for some allocation of data.

The social planner who cares about consumer surplus should consider not only what inference the firm can make from the aggregate data, but also how the data are initially allocated across consumers. To see this, compare the following two scenarios. First, suppose the aggregate data enable the firm to perfectly price discriminate. In such a case, consumer surplus is zero in the product market, but the net surplus can be positive, because the firm compensates consumers for providing data if they hold complementary data. Second, suppose that the aggregate data correspond to a consumer-optimal segmentation in Bergemann et al. (2015b), leading to a high consumer surplus in the product market. However, if consumer data are complementary, the firm will charge a fee in the data market to extract the surplus that accrues to consumers in the product market.



Figure 3: The set of possible outcomes for n = 1 (blue segment, BD) and $n \ge 2$ (gray triangle, AEC).

7 Conclusion

The paper has studied a stylized model in which a firm collects data from consumers. The welfare implication of data collection depends on data externalities—i.e., what information each consumer's data reveal about other consumers' data. Data externalities can render data collection beneficial or harmful to consumers, regardless of its social value. If the firm can flexibly design what information to collect, it chooses a data externality that reduces each consumer's outside option from refusing to share data. Such a policy leads to an inefficiently high level of data collection and harms consumers.

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar (2019), "Too much data: Prices and inefficiencies in data markets." Working Paper 26296, National Bureau of Economic Research.

- Admati, Anat R and Paul Pfleiderer (1986), "A monopolistic market for information." *Journal of Economic Theory*, 39, 400–438.
- Agarwal, Anish, Munther Dahleh, and Tuhin Sarkar (2019), "A marketplace for data: An algorithmic solution." In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 701–726.
- Ali, S Nageeb, Gregory Lewis, and Shoshana Vasserman (2019), "Voluntary disclosure and personalized pricing." Technical report, National Bureau of Economic Research.
- Aliprantis, Charalambos D and Kim Border (2006), *Infinite dimensional analysis: a hitchhiker's guide*. Springer Science & Business Media.
- Arrieta-Ibarra, Imanol, Leonard Goff, Diego Jiménez-Hernández, Jaron Lanier, and E Glen Weyl (2018), "Should we treat data as labor? Moving beyond "Free"." In AEA Papers and Proceedings, volume 108, 38–42.
- Aumann, Robert J and Michael Maschler (1995), *Repeated games with incomplete information*. MIT press.
- Babaioff, Moshe, Robert Kleinberg, and Renato Paes Leme (2012), "Optimal mechanisms for selling information." In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 92– 109, ACM.
- Bergemann, Dirk, Alessandro Bonatti, and Tan Gan (2019), "The economics of social data." Cowles foundation discussion paper 2203, Yale University.
- Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin (2015a), "Selling experiments." Technical report, Mimeo, MIT and Yale.
- Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin (2018), "The design and price of information." *American Economic Review*, 108, 1–48.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2015b), "The limits of price discrimination." *The American Economic Review*, 105, 921–957.

- Blackwell, David (1953), "Equivalent comparisons of experiments." *The annals of mathematical statistics*, 24, 265–272.
- Bonatti, Alessandro and Gonzalo Cisternas (2020), "Consumer scores and price discrimination." *The Review of Economic Studies*, 87, 750–791.
- Börgers, Tilman, Angel Hernando-Veciana, and Daniel Krähmer (2013), "When are signals complements or substitutes?" *Journal of Economic Theory*, 148, 165–195.
- Calzolari, Giacomo and Alessandro Pavan (2006), "On the optimality of privacy in sequential contracting." *Journal of Economic theory*, 130, 168–204.
- Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim (2019), "Privacy and personal data collection with information externalities." *Journal of Public Economics*, 173, 113–124.
- Conitzer, Vincent, Curtis R Taylor, and Liad Wagman (2012), "Hide and seek: Costly consumer privacy in a market with repeat purchases." *Marketing Science*, 31, 277–292.
- De Corniere, Alexandre and Romain De Nijs (2016), "Online advertising and privacy." *The RAND Journal of Economics*, 47, 48–72.
- De Corniere, Alexandre and Greg Taylor (2020), "Data and competition: a general framework with applications to mergers, market structure, and privacy policy."
- Easley, David, Shiyang Huang, Liyan Yang, and Zhuo Zhong (2018), "The economics of data." *Available at SSRN 3252870.*
- Eső, Péter and Balazs Szentes (2007), "Optimal information disclosure in auctions and the handicap auction." *The Review of Economic Studies*, 74, 705–731.
- Fairfield, Joshua AT and Christoph Engel (2015), "Privacy as a public good." Duke LJ, 65, 385.
- Hörner, Johannes and Andrzej Skrzypacz (2016), "Selling information." Journal of Political Economy, 124, 1515–1562.
- Ichihashi, Shota (forthcoming), "Competing data intermediaries." RAND Journal of Economics.

- Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian persuasion." *American Economic Review*, 101, 2590–2615.
- MacCarthy, Mark (2010), "New directions in privacy: Disclosure, unfairness and externalities." *ISJLP*, 6, 425.
- Madio, Leonardo, Yiquan Gu, and Carlo Reggiani (2019), "Exclusive data, price manipulation and market leadership."
- Montes, Rodrigo, Wilfried Sand-Zantman, and Tommaso Valletti (2019), "The value of personal information in online markets with endogenous privacy." *Management Science*, 65, 1342–1362.
- Shamir, Adi (1979), "How to share a secret." Communications of the ACM, 22, 612–613.
- Taylor, Curtis R (2004), "Consumer privacy and the market for customer information." *RAND Journal of Economics*, 631–650.
- Tirole, Jean (2010), The theory of corporate finance. Princeton University Press.

Appendix

A Omitted Proofs for Section 4

Proof of Proposition 1. The existence of a feasible and substitutable μ^* follows from Lemma 1. Suppose that the allocation of data is μ^* and the firm chooses $p_i = 0$ for all *i*. There is an equilibrium in which all consumers sell their data: Because μ^* is perfectly substitutable, each consumer is indifferent between selling and not selling her data, whenever all other consumers sell their data. The firm does not benefit from decreasing p_i because consumer *i* will reject it. This equilibrium leads to the firm's profit $\pi(\mu_0)$ and consumer surplus $\sum_{i \in \mathcal{N}} u_i(\mu_0)$.

To show the welfare implications, take any feasible allocation and any (possibly mixed-strategy) equilibrium. The firm's profit is at most $\pi(\mu_0)$ because it cannot charge negative prices when data collection harms consumers. Also, the equilibrium payoff of each consumer *i* is at least $u_i(\mu_0)$, because she can refuse to sell her data. A substitutable allocation μ^* attains these bounds.

Proof of Proposition 2. The existence of a feasible and complementary μ^* follows from Lemma 1. Suppose that the firm offers each consumer *i* a price of $-u_i(\mu_0)$ and all consumers sell their data. Because of complementarity, price $-u_i(\mu_0)$ is consumer *i*'s loss of selling her data conditional on that other consumers sell their data. A similar argument as Proposition 1 implies this is an equilibrium. All consumers receive zero payoffs, and the firm obtains $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0)$. The outcome maximizes the payoff of each consumer, who never obtains a positive payoff.¹⁶ It also minimizes the firm's payoff because for any $\mu \in \mathcal{F}(\mu_0)$, the firm can offer each consumer a price of $-u_i(\mu_0)$ to secure $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0)$. As a result μ^* maximizes consumer surplus and minimizes the firm's profit.

Proof of Proposition 3. Under μ^* , there is an equilibrium in which the firm sets $p_i^* = -u_i(\mu_0) \le 0$ for all $i \in \mathcal{N}$ and all consumers share their data. Given \mathbf{p}^* , each consumer is indifferent between selling and not selling her data, conditional on that all other consumers share their data. This is an equilibrium that maximizes the firm's profit, because the firm extracts the efficient total surplus while giving consumers the lowest possible payoff of zero.

Proof of Proposition 4. There is an equilibrium in which the firm sets $p_i^* = 0$ for all *i* and all consumers sell their data. In particular, consumer *i* refuses to give data whenever the firm deviates to a lower (i.e., negative) price, because *i* believes that other consumers already give their data. The equilibrium maximizes consumer surplus and minimizes the firm's profit because there is no equilibrium in which the firm pays a positive price when data collection is beneficial.¹⁷

B Proofs for Section 4.4: Multiplicity of Equilibria

I impose the following restriction on consumer preferences.

Assumption 1. For any $i \in \mathcal{N}$, any $\mu, \mu' \in \Sigma$, and any $\alpha \in [0, 1]$, $u_i(\alpha \langle \mu \rangle + (1 - \alpha) \langle \mu' \rangle) = \alpha u_i(\langle \mu \rangle) + (1 - \alpha) u_i(\langle \mu' \rangle).$

Assumption 1 holds if each consumer *i* has some underlying payoff $u_i(a, X)$ that depends on the firm's (unmodeled) action *a* and a realized state *X*. Denoting the firm's action at a posterior

¹⁶Suppose to the contrary that there is an equilibrium in which consumer i obtains a positive payoff. It means that the firm pays a positive price, and consumer i strictly prefers to share her data. Because she holds a passive belief, the firm can lower the price to buy the same data, which is a contradiction.

¹⁷This observation follows from the same argument as Footnote 16.

 $b \in \Delta \mathcal{X}$ by a(b), we can write $u_i(\mu) = \int_{\Delta \mathcal{X}} \int_{\mathcal{X}} u_i(a(b), X) db(X) d\langle \mu \rangle(b)$. Because $u_i(\mu)$ is linear, it satisfies Assumption 1.

B.1 Approximation Result for Proposition 3

Proposition 8. Fix any aggregate data $\mu_0 \in \Sigma$ such that $\pi(\mu_0) > 0$. There is a sequence of feasible allocations of data $(\boldsymbol{\mu}^k)_{k \in \mathbb{N}}$ that satisfies the following.

- 1. For each μ^k , there is a unique equilibrium.
- 2. As $k \to +\infty$, the equilibrium consumer surplus converges to 0, and the firm's profit converges to $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0)$.
- 3. For each $i \in \mathcal{N}$, as $k \to +\infty$, $\langle \mu_i^k \rangle$ weakly converges to $\langle \mu_{\emptyset} \rangle$.

Proof. Take any perfectly complementary allocation of data $\mu^C = (\mu_1^C, \dots, \mu_n^C)$ such that $\langle \mu^C \rangle = \mu_0$. For each $k \in \mathbb{N}$, consider the following allocation $(\mu_i^k)_{i \in \mathcal{N}}$: (i) with probability $1 - \frac{1}{k}$, $(\mu_i^k)_{i \in \mathcal{N}} = \mu^C$, and (ii) with probability $\frac{1}{k}$, $\mu_i^k = \mu_0$ or $\mu_i^k = \mu_{\emptyset}$ holds. Specifically, conditional on (ii), one of *n* consumers (say *i*) is randomly picked with probability $\frac{1}{n}$, and $\mu_i^k = \mu_0$ holds. For any other consumer j, $\mu_j^k = \mu_{\emptyset}$ holds. Thus, given (ii), there is exactly one consumer with $\mu_i^k = \mu_0$.

Take any equilibrium, and suppose to the contrary that the firm does not collect data from (say) consumer 1. Suppose that the firm deviates and offers $p_1 = \varepsilon > 0$. Consumer 1 then strictly prefers to sell data because she earns ε and benefits from the firm's learning. Consider the positive probability event in which (ii) is realized and $\mu_1^k = \mu_0$. On this event, the firm's profit strictly increases by collecting data μ_1^k . Thus, the above deviation is profitable for a small $\varepsilon > 0$, leading to a contradiction. Given that the firm collects all data, the maximum price the firm can charge to i is $p_i^k = -u_i(\mu_{1,1-i}^k) + u_i(\mu_{0,1-i}^k) = -(1 - \frac{1}{k} + \frac{1}{nk}) u_i(\mu_0)$. In the unique equilibrium, the firm charges each consumer i a non-positive price of p_i^k , which converges to $-u_i(\mu_0)$ as $k \to +\infty$. Points 2 and 3 directly follows from these observations.

B.2 Approximation Result for Proposition 4

Proposition 9. Fix any aggregate data $\mu_0 \in \Sigma$ such that $\pi(\mu_0) > 0$. There is a sequence of feasible allocations of data $(\boldsymbol{\mu}^k)_{k \in \mathbb{N}}$ that satisfies the following.

- 1. For each μ^k , there is a unique equilibrium.
- 2. As $k \to +\infty$, the equilibrium consumer surplus converges to $\sum_{i \in \mathcal{N}} u_i(\mu_0)$, and the firm's profit converges to $\pi(\mu_0)$.
- 3. For each $i \in \mathcal{N}$, as $k \to +\infty$, $\langle \mu_i^k \rangle$ weakly converges to $\langle \mu_0 \rangle$.

Proof. For each $k \in \mathbb{N}$, consider the following allocation of data $(\mu_i^k)_{i \in \mathcal{N}}$: For each realized $X \in \mathcal{X}$, (i) with probability $1 - \frac{1}{k}$, $\mu_i^k = \mu_0$ for all $i \in \mathcal{N}$; (ii) with probability $\frac{1}{k}$, only one of n consumers, say i, has $\mu_i^k = \mu_0$, and any other consumer j has $\mu_j^k = \mu_0$ (i.e., (ii) is the same as the one in the proof of Proposition 8). By the same argument as the proof of Proposition 8, the firm collects data from all consumers in any equilibrium. Given that the firm collects all data, the maximum price the firm can charge to consumer i is $p_i^k = -u_i(\mu_{1,1-i}^k) + u_i(\mu_{0,1-i}^k) = -\frac{1}{nk}u_i(\mu_0)$. In the unique equilibrium, the firm charges each consumer i a price of p_i^k , which converges to 0 as $k \to +\infty$. Points 2 and 3 follow from the above observations.

C Equilibria with Non-Passive Beliefs

This section consists of four parts. First, I present an example in which an equilibrium with nonpassive beliefs exists under an extreme allocation of data identified in the main analysis. Second, I present an example in which non-passive beliefs can increase the surplus bound for the consumeroptimal outcome under harmful data collection. Third, I show that other surplus bounds do not change even if we allow non-passive beliefs. Finally we prove Claim 2.

C.1 Example of Equilibrium with Non-Passive Beliefs

Consider harmful data collection with a perfectly substitutable allocation of data. For simplicity, suppose all consumers have the same $u_i(\cdot)$. Proposition 1 shows that there is an equilibrium in which each consumer *i* receives a net payoff of $u_i(\mu_0)$, which minimizes consumer surplus.

There is also a PBE in which one consumer (say 1) receives a payoff of zero, and other consumers receive $u_i(\mu_0)$. To see this, consider the following strategy profile with non-passive beliefs: The firm offers price $-u_1(\mu_0)$ to consumer 1, who sells her data, and it offers a price of zero to other consumers, who do not sell their data. Whenever the firm deviates and offers a different price to consumer *i*, she believes that the firm offers negative prices to all consumers $j \neq i$, who will refuse to sell their data. If $\pi(\mu_0) + u_1(\mu_0) \ge 0$, this is an equilibrium. In particular, unlike the case of passive beliefs, the firm can no longer collect data for free from consumer $j \neq 1$, because any consumer who detects a deviation believes that her data provision is pivotal. In this equilibrium, consumer surplus is greater than the equilibrium of Proposition 1, in which all consumers provide their data for free. The example implies that the uniqueness of equilibrium discussed in Section 4.4 could fail if we allow non-passive beliefs.

C.2 Non-Passive Beliefs: Consumer-Optimal Allocation Under Harmful Data Collection

I show that the result on the consumer-optimal outcome under harmful data collection (i.e., part of Proposition 2) fails once we allow passive beliefs. To see this, suppose the state is two-dimensional: $X = (X_1, X_2)$, where X_1 and X_2 are independent at the prior. For each $j \in \{1, 2\}$, let $\mu(j)$ denote the experiment that fully reveals the realization of X_j and is uninformative about the other dimension. Also, let $\mu(12)$ denote the experiment that fully reveals the state.

Consider the following payoffs: $u(\mu(1)) = u(\mu(2)) = -1$, $u(\mu(12)) = -4$, and $\pi(\mu(1)) = \pi(\mu(2)) = \pi(\mu(12)) = 9$. Suppose there are two consumers, and consumers 1 and 2 hold $\mu(1)$ and $\mu(2)$, respectively. We set $\mu_0 = \mu(12)$, which satisfies the assumption in Proposition 2. Consider the following equilibrium: The firm collects data only from consumer 1 at $p_1 = 3$ and offers a price of zero to consumer 2. Consumer 2 has the passive belief. Consumer 1, in contrast, believes that if the firm deviates, it collects data from consumer 2 at price (say) $+\infty$. The firm has no profitable deviation: If the firm decreases the price to consumer 1, she refuses to sell her data, because she believes that consumer 2 sells her data, so consumer 1's loss of selling her data is now $u(\mu(12)) - u(\mu(2)) = -3$ as opposed to -1. Also, the firm does not benefit from buying data from consumer 2, because she requires a price of at least 3, and the marginal value of the second unit of data is zero to the firm. In this equilibrium, consumer surplus is $2u(\mu(1)) + 3 = 1$, which is strictly greater than the maximum value of zero under passive beliefs.

C.3 Non-Passive Beliefs: The Robustness of Other Surplus Bounds

All other results continue to hold even if we allow any belief systems. First, the consumer-worst outcomes remain the same, because I derive the results based on the observations that consumers can secure a payoff of $u_i(\mu_0)$ (under harmful data collection) or 0 (under beneficial data collection) by refusing to sell data. The payoff is a lower bound of a consumer's outside option for any beliefs, so non-passive beliefs do not change the consumer-worst outcome. The same argument implies that the general consumer-worst outcome in Proposition 5 and the firm-optimal data collection in Proposition 6 remain the same. The rest of the results continue to hold, because they only use the upper and lower bounds of possible equilibrium prices that do not depend on consumers' beliefs: The firm-best and the firm-worst outcomes under harmful data collection remain the same, because the results only use the fact that the firm cannot charge negative prices but can always pay $-u_i(\mu_0)$ to collect data from each consumer *i*. Also, the firm-worst and the consumer-best outcome under beneficial data collection remain the same, because the firm can always collect data at a small positive price. Finally, the firm-best outcome under beneficial data collection remains the same, because the firm cannot charge more than $u_i(\mu_0)$ on each consumer *i*.

C.4 Proof of Claim 2

Proof. Assume $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0) < 0$. For the first part, consider the following strategy profile: On the equilibrium path, the firm offers a price of zero to all consumers, who refuse to sell their data. If the firm deviates and increases a price to consumer *i*, then *i* believes that the firm offers a greater price than $-u_j(\mu_0)$ and buys data from all other consumers $j \neq i$. This is an equilibrium: Given the consumers' non-passive beliefs, the firm can acquire data μ_0 only by paying $-u_i(\mu_0)$ or more to each consumer *i*, but such an offer will lead to a negative profit. As a result the firm has no incentive to deviate. Consumers have no incentive deviate. In particular, each consumer is indifferent between selling and not selling data on the equilibrium path.

For the second part, suppose to the contrary that there is a (possibly mixed) perfect Bayesian equilibrium with passive beliefs. The rest of the proof consists of three steps.

Step 1: *There is no positive probability event where the firm collects* μ_0 . Suppose to the contrary that there is a positive probability event in which all consumers sell their data. Take any *i* with

 $u_i(\mu_0) < 0$. Consumer *i* incurs a loss if the firm collects μ_0 , so she has to receive a positive price with a positive probability. On such an event, the firm collects data from other consumers -i for sure; otherwise, it would collect an uninformative signal and make a negative profit from *i*, but it could then profitably deviate by offering a negative price to *i*. As a result, consumer *i* correctly anticipates that the firm will collect data from all other consumers. To sum up, whenever all consumers sell their data, the firm has to pay a price of at least $-u_i(\mu_0)$ to each consumer *i*. Because $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0) < 0$ the firm will earn a negative profit on such an event, which is a contradiction.

Step 2: *There is some consumer j who never sells data*. Passive belief implies that the prices offered to different consumers are independent, which also implies that whether consumer *i* sells her data is independent across *i*. Thus if each consumer sells data with a positive probability, all consumers simultaneously sell their data with a positive probability, which contradicts Step 1.

Step 3: *The firm has a profitable deviation.* Suppose consumer j never sells data in equilibrium. The firm can deviate and buy data from consumer j at a price of $-u_j(\mu_0)$ and data from other consumers for free. Indeed, any consumer $i \neq j$ believes that the firm does not collect data from consumer j and thus i's data does not affect the firm's inference. The deviation is profitable because the firm's profit increases from zero to $\pi(\mu_0) + u_j(\mu_0) \geq \pi(\mu_0) + \min_{i \in \mathcal{N}} u_i(\mu_0) > 0$, where the strict inequality is by assumption. We obtain a contradiction and conclude there is no equilibrium with passive beliefs.

D Appendix for **Remark 1**: The Set of Equilibrium Payoffs

This appendix provides results on the set of all equilibrium payoffs across all feasible allocations of data. Throughout the appendix, we fix the aggregate data, μ_0 .

Claim 4. Consider harmful data collection with $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0) \ge 0$. For the firm and consumers, any payoff between the worst and the best payoffs can arise under some data allocation: For any $\pi \in [\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0), \pi(\mu_0)]$, there is some feasible allocation of data and the corresponding equilibrium in which the firm receives a payoff of π . Similarly, for any $U \in [\sum_{i \in \mathcal{N}} u_i(\mu_0), 0]$, there is some feasible allocation of data and the corresponding equilibrium in which the firm receives a payoff of π .

Proof. Take any $\alpha \in [0, 1]$. Take a perfectly complementary allocation of data $\mu^C = (\mu_1^C, \dots, \mu_n^C)$ such that $\langle \mu^C \rangle = \langle \mu_0 \rangle$. Consider the following allocation $(\mu_i^*)_{i \in \mathcal{N}}$: (i) with probability α , $(\mu_i^*)_{i \in \mathcal{N}} = \mu^C$ and (ii) with probability $1 - \alpha$, $\mu_i^* = \mu_0$ for all $i \in \mathcal{N}$. Conditional on that all consumers $j \neq i$ sell their data, consumer *i*'s utility from refusing to sell data is $(1 - \alpha)u_i(\mu_0)$. If she sells her data, her utility is $u_i(\mu_0)$. As a result, the minimum price consumer *i* is willing to accept to sell her data is $p_i^* = (1 - \alpha)u_i(\mu_0) - u_i(\mu_0) = -\alpha u_i(\mu_0) \geq 0$. Now, there is an equilibrium in which the firm offers price $-\alpha u_i(\mu_0)$ to each consumer *i* and all consumers sell their data. In particular, if the firm deviates and buys data only from consumer *i*, then relative to collecting all data, the firm would lose the gross revenue of $\alpha \pi(\mu_0)$ and saves the total price of $-\alpha \sum_{j \neq i} u_j(\mu_0)$. Because $\pi(\mu_0) + \sum_{j \neq i} u_j(\mu_0) \geq \pi(\mu_0) + \sum_{j \in \mathcal{N}} u_j(\mu_0) \geq 0$, the firm does not benefit from such a deviation. A similar argument implies that the firm has no other profitable deviation, and neither do consumers. In this equilibrium, the firm's payoff is $\pi(\mu_0) + \alpha \sum_{i \in \mathcal{N}} u_i(\mu_0)$, and the consumer surplus is $(1 - \alpha) \sum_{i \in \mathcal{N}} u_i(\mu_0)$. By moving α from 0 to 1, we obtain the result.

By the same argument as above, we obtain the following result.

Claim 5. Consider beneficial data collection. For the firm and consumers, any payoff between the worst and the best payoffs can arise under some data allocation: For any $\pi \in [\pi(\mu_0), \pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i(\mu_0)]$, there is some feasible allocation of data and the corresponding equilibrium in which the firm receives a payoff of π . Similarly, for any $U \in [0, \sum_{i \in \mathcal{N}} u_i(\mu_0)]$, there is some feasible allocation of data and the corresponding equilibrium in which the consumer surplus is U.

The above claims characterize the set of possible payoffs for each player, separately. To do so, we have focused on equilibria in which the firm collects data from all consumers. The following example shows that to characterize the set of payoff *vectors*, we need to study equilibrium in which the firm does not collect data from some consumers. Thus our results, which are based on equilibria with full data collection, do not immediately apply to the full set characterization.

Example 2. The state is two-dimensional, $(X_Y, X_Z) \in \mathcal{X}$. Let Y and Z denote the fully informative signals for X_Y and X_Z . There are three consumers, who have the same $u(\cdot)$. Let u(Y) denote a consumer's gross payoff when the firm obtains data Y; analogously, define u(Z), u(YZ), $\pi(Y)$, $\pi(Z)$, and $\pi(YZ)$. Data collection is harmful: u(Y) = u(Z) = -0.5, u(YZ) = -2,

and $\pi(Y) = \pi(Z) = \pi(YZ) = 7$. Suppose consumers 1 and 2 hold Y as a complementary allocation, and consumer 3 holds Z. In one equilibrium, the firm collects data from consumers 1 and 2. The firm does not collect data from consumer 3, who demands compensation of at least $u(Y) - u(YZ) = 1.5 > 0 = \pi(YZ) - \pi(Y)$. In this equilibrium the firm earns 6, consumers 1 and 2 earn zero, and consumer 3 earns -0.5. Claim 4 implies that the payoff of each player—say payoff 6 for the firm—can arise given another allocation of data under which the firm collects data from all consumers. However, the payoff vector (6, 0, 0, -0.5) leads to total surplus 5.5, which is strictly greater than the one when the firm collects all data, i.e., $7 - 2 \cdot 3 = 1$. Thus the payoff vector (6, 0, 0, -0.5) arises only if the firm does not collect data from some consumers.

E Appendix for Remark 2: Non-Monotone, Separable Payoffs

We impose the following structure on the baseline model. The state space is written as $\mathcal{X} = \mathcal{X}_G \times \mathcal{X}_B$, where the two dimensions of the state, $X_G \in \mathcal{X}_G$ and $X_B \in \mathcal{X}_B$, are independent according to the prior. The aggregate data $\mu_0 : \mathcal{X} \to \Delta(\mathcal{S}_G \times \mathcal{S}_B)$ satisfies

$$\mu_0(S_G, S_B | X_G, X_B) = \mu_0^G(S_G | X_G) \cdot \mu_0^B(S_B | X_B), \forall (S_G, S_B, X_G, X_B) \in \mathcal{S}_G \times \mathcal{S}_B \times \mathcal{X}_G \times \mathcal{X}_B.$$
(7)

The condition (7) restricts the kind of aggregate data we consider. However, given such μ_0 , we still consider all feasible allocations of data that may not satisfy the separability condition like (7).

Take any information $\mu : \mathcal{X} \to \Delta \mathcal{S}$ the firm acquires, and let $\mu_G \in \Delta \Delta \mathcal{X}_G$ and $\mu_B \in \Delta \Delta \mathcal{X}_B$ denote the distribution of posteriors induced by μ about X_G and X_B , respectively. The payoff of each consumer *i* is $u_i^G(\mu_G) + u_i^B(\mu_B)$, where $u_i^G(\mu_G)$ is increasing in the informativeness of μ_G , and $u_i^G(\mu_B)$ is decreasing in the informativeness of μ_B . We normalize $u_i^G(\cdot)$ and $u_i^B(\cdot)$ to be zero when they are evaluated at the uninformative experiments. The combination of Propositions 1 - 4 in the main text leads to the following result.

Claim 6. Across all feasible allocations of data, consumer surplus is minimized and the firm's profit is maximized if consumers hold μ_0^G as complementary data and μ_0^B as substitutable data. If $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i^B(\mu_0^B) \ge \pi(\mu_0^G)$, consumer surplus is maximized and the firm's profit is minimized if consumers hold μ_0^G as substitutable data and μ_0^B as complementary data.

Proof. We show the second part, as it is more complicated. Suppose consumers hold μ_0^G as substitutable data and μ_0^B as complementary data. Formally, consumers hold $(\mu_i^G, \mu_i^B)_{i \in \mathcal{N}}$ such that $\langle (\mu_i^G)_{i \in \mathcal{N}} \rangle = \langle (\mu_i^G)_{i \neq j} \rangle = \langle \mu_0^G \rangle$ and $\langle (\mu_i^B)_{i \in \mathcal{N}} \rangle = \langle \mu_0^B \rangle$, but $\langle (\mu_i^B)_{i \neq j} \rangle$ is totally uninformative for each *j*. Given such an allocation, there is an equilibrium in which the firm collects data from each consumer *i* with price $p_i = -u_i^B(\mu_0^B)$. Suppose the firm offers $(p_i)_{i \in \mathcal{N}}$ and all consumers sell their data. If consumer *i* unilaterally deviates and refuses to sell her data, her gross payoff increases by $-u_i^B(\mu_0^B)$ but she cannot receive p_i . Thus she is indifferent between selling and not selling data. The firm has no profitable deviation either. If the firm follows the above strategy, it receives a payoff of $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i^B(\mu_0^B) \ge \pi(\mu_0^G) \ge 0$. If the firm deviates and collects no data, it receives a payoff of zero. If it deviates and collects data from the set \mathcal{N}' of consumers such that even after the deviation, the firm has to pay $-u_i^B(\mu_0^B)$ to collect data of each consumer $i \in \mathcal{N}'$.) In either case, the firm's payoff is lower than $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i^B(\mu_0^B)$.

The above equilibrium (say E) maximizes consumer surplus and minimizes the firm's profit. To see this, take any feasible allocation of data and any equilibrium. If consumer i sells her data, she is indifferent between selling and not selling her data given our passive belief assumption. Thus regardless of whether consumer i sells her data, her equilibrium payoff is the one from not selling data, which is at most $u_i^G(\mu_0^G)$, i.e., her payoff in E. Thus equilibrium E maximizes consumer surplus. The firm can always pay $-u_i^B(\mu_0^B)$ to each consumer i to collect all data, so $\pi(\mu_0) + \sum_{i \in \mathcal{N}} u_i^B(\mu_0^B)$ is a lower bound of the firm's profit. Thus equilibrium E, which attains the lower bound, minimizes the firm's profit.

By the similar argument we can prove the first part: If consumers hold μ_0^G as complementary data and μ_0^B as substitutable data, there is an equilibrium in which the firm collects all data and consumer *i* pays $u_i^G(\mu_0^G)$. Such an equilibrium minimizes consumer surplus because each consumer obtains the lowest possible payoff of $u_i^B(\mu_0^B)$. The equilibrium maximizes the firm's profit because it cannot charge more than $u_i^G(\mu_0^G)$ for collecting data.

F Appendix for the Proof of Corollary 3

We identify μ with $\langle \mu \rangle$. We prove equation (4), which is equivalent to

$$\max_{\mu \succeq \nu} \int_{\Delta \mathcal{X}} f(b) d\mu(b) = \int_{\Delta \mathcal{X}} \mathcal{V}[f](b) d\nu(b), \tag{8}$$

where $f = \pi + u$. First, by the definition of $\mathcal{V}[f]$, we can find a mean-preserving transition kernel $Q: b \mapsto Q(\cdot|b) \in \Delta \Delta \mathcal{X}$ such that¹⁸

$$\int_{\Delta \mathcal{X}} \mathcal{V}[f](b) d\nu(b) = \int_{\Delta \mathcal{X}} \int_{\Delta \mathcal{X}} f(\beta) dQ(\beta|b) d\nu(b) = \int_{\Delta \mathcal{X}} f(b) d\mu'(b),$$

where $\mu'(\cdot):=\int_{\Delta\mathcal{X}}Q(\cdot|b)d\nu(b).$ Because $\mu'\succeq\nu,$ we have

$$\max_{\mu \succeq \nu} \int_{\Delta \mathcal{X}} f(b) d\mu(b) \ge \int_{\Delta \mathcal{X}} \mathcal{V}[f](b) d\nu(b).$$
(9)

To show the converse, take any $\mu \succeq \nu$. Given a mean-preserving transition kernel Q that satisfies $\mu(\cdot) = \int_{\Delta \mathcal{X}} Q(\cdot|b) d\nu(b)$, we have

$$\int_{\Delta \mathcal{X}} f(b) d\mu(b) = \int_{\Delta \mathcal{X}} \int_{\Delta \mathcal{X}} f(\beta) dQ(\beta|b) d\nu(b) \leq \int_{\Delta \mathcal{X}} \mathcal{V}[f](b) d\nu(b) = \int_{\Delta \mathcal{X}} f(\beta) d\mu(b) d\nu(b) = \int_{\Delta \mathcal{X}} \int_{\Delta \mathcal{X}} f(\beta) d\mu(b) d\nu(b) = \int_{\Delta \mathcal{X}} \int_{\Delta \mathcal{X}} f(\beta) d\mu(b) d\nu(b) d\nu(b) = \int_{\Delta \mathcal{X}} \int_{\Delta \mathcal{X}} \int_{\Delta \mathcal{X}} f(\beta) d\mu(b) d\nu(b) d\nu(b) d\nu(b) = \int_{\Delta \mathcal{X}} \int_{\Delta \mathcal{X}} \int_{\Delta \mathcal{X}} f(\beta) d\mu(b) d\nu(b) d\nu(b)$$

which implies (8).

G The Existence of a Solution to Problem (1)

I provide conditions under which the firm's problem (1), which is written as

$$\max_{\mu \in \mathbf{\Sigma}} \left\{ \pi(\mu) + \sum_{i \in \mathcal{N}} u_i(\mu) + \sum_{i \in \mathcal{N}} \max_{\mu' \in G(\mu)} \left(-u_i(\mu') \right) \right\}, \quad \text{where}$$
(10)

$$G(\mu) = \{\mu' \in \Sigma : \mu' \preceq \mu\},\tag{11}$$

has a solution. I identify $\mu \in \Sigma$ with a Bayes plausible element of $\Delta \Delta \mathcal{X}$ and use weak* topology on it. The first existence result uses the standard argument based on Berge Maximum Theorem.

¹⁸A function $Q: \Delta \mathcal{X} \to \Delta \Delta \mathcal{X}$ is a mean-preserving transition kernel if for all $b \in \Delta(\mathcal{X})$, $\int_{\Delta \mathcal{X}} \beta Q(d\beta|b) = b$.

Claim 7. If $\pi(\cdot)$ is upper semicontinuous and $(u_i(\cdot))_{i\in\mathcal{N}}$ is continuous, the problem (1) has a solution.

Proof. We show $G(\cdot)$ is compact-valued and upper hemicontinuous. Let $\Phi \subset \mathbb{R}^{\Delta \mathcal{X}}$ denote the set of all continuous convex functions. By Theorem 7 of Blackwell (1953), we can write $G(\mu)$ as

$$\begin{split} G(\mu) &= \bigcap_{\phi \in \Phi} \left\{ \mu' \in \mathbf{\Sigma} : \int_{\Delta \mathcal{X}} \phi d\mu \geq \int_{\Delta \mathcal{X}} \phi d\mu' \right\} \\ &= \bigcap_{\phi \in \Phi} G(\phi, \mu), \end{split}$$

where $G(\phi, \mu) = \{\mu' \in \Sigma : \int_{\Delta \mathcal{X}} \phi d\mu \ge \int_{\Delta \mathcal{X}} \phi d\mu'\}$. The set $G(\phi, \mu)$ is compact, because $\Sigma \subset \Delta \Delta \mathcal{X}$ is compact and $G(\phi, \mu)$ is a closed subset of it. Indeed, if $(\mu_n)_n \subset G(\phi, \mu)$ and $\mu_n \to \mu_0$, then $\int \phi \mu_n \to \int \phi \mu_0$ for any $\phi \in \Phi$, because ϕ is a continuous function defined on a compact set $\Delta \mathcal{X}$ (and thus bounded). Thus, if $\mu_n \to \mu_0$ and $\int \phi d\mu \ge \int \phi d\mu_n$, then $\int \phi d\mu \ge \int \phi d\mu_0$. This implies $\mu_0 \in G(\phi, \mu)$, so $G(\phi, \mu)$ is closed.

The upper hemicontinuity of $G(\phi, \cdot)$ is shown as follows. Take $\hat{\mu}_n \to \hat{\mu}$ and $\mu_n \to \mu$ such that $\int \phi d\hat{\mu}_n \geq \int \phi d\mu_n$ for any n. By the similar argument as above, we get $\int \phi d\hat{\mu} \geq \int \phi d\mu$, and thus $\mu \in G(\phi, \hat{\mu})$. Since $G(\phi, \cdot)$ has a closed graph and is closed-valued, it is upper hemicontinuous. Because $G(\mu)$ is the intersection of compact-valued upper hemicontinuous correspondences, Point 2 of Theorem 17.25 of Aliprantis and Border (2006) implies that it is upper hemicontinuous.

Because the correspondence $G(\cdot)$ is compact-valued and upper hemicontinuous, Lemma 17.30 of Aliprantis and Border (2006) implies that $\sum_{i \in \mathcal{N}} \max_{\mu' \in G(\mu)} (-u_i(\mu'))$ is upper semicontinuous in μ . The maximand of (10) then becomes upper semicontinuous, so Theorem 2.43 of Aliprantis and Border (2006) implies a solution exists.

Claim 7 provides a condition for existence, but it may not be easy to check. For example, suppose the firm uses a signal to learn about X, then takes a payoff-relevant action. In such a case, $\pi(\mu)$ and $(u_i(\cdot))_{i\in\mathcal{N}}$ are not the primitives but the payoffs from the firm's optimal behavior.

We now provide a condition for existence when the firm uses information to choose an action. Consider the following three-stage game: First, the firm sets prices to buy data. Second, consumers decide whether to sell their data. Finally, the firm chooses an action a from a subset \mathcal{A} of \mathbb{R}^m . Let $\pi(a, X)$ and $u_i(a, X)$ denote the ex post gross payoffs of the firm and consumer i if the firm chooses $a \in A$ and the realized state is X. The solution concept is perfect Bayesian equilibrium. Abusing terminology, we say that an allocation of data μ^* globally maximizes the firm's profit if μ^* satisfies Definition 4, where "equilibrium" there now refers to equilibrium of this extended game.

We prepare some notations. For each posterior $b \in \Delta \mathcal{X}$, let $a^i(b)$ denote the firm's best response that breaks ties to minimize consumer *i*'s expected payoff. Also, let $a^{\mathcal{N}}(b)$ denote the firm's best response that breaks ties to maximize the sum of the expected payoffs of all players. (The assumptions stated in the following results ensure that these best responses exist). Define $\pi^{\mathcal{N}}(\mu) := \int_{\Delta \mathcal{X}} \pi(a^{\mathcal{N}}(b), b) d\mu, u_i^{\mathcal{N}}(\mu) := \int_{\Delta \mathcal{X}} u_i(a^{\mathcal{N}}(b), b) d\mu$, and $u_i^i(\mu) := \int_{\Delta \mathcal{X}} u_i(a^i(b), b) d\mu$. Given the tie-breaking rules, the problem similar to (10) has a solution and provides the firmoptimal data collection.

Claim 8. If the firm's action space A is finite, then

$$\max_{\mu \in \mathbf{\Sigma}} \left\{ \pi^{\mathcal{N}}(\mu) + \sum_{i \in \mathcal{N}} u_i^{\mathcal{N}}(\mu) + \sum_{i \in \mathcal{N}} \max_{\mu' \in G(\mu)} \left(-u_i^i(\mu') \right) \right\}$$
(12)

has a solution, μ_0^* . There is an allocation of data μ^* such that $\langle \mu^* \rangle = \langle \mu_0^* \rangle$ and μ^* globally maximizes the firm's profit.

The result is a special case of the following.

Claim 9. Suppose the firm's action space $\mathcal{A} \subset \mathbb{R}^m$ is compact, and $\pi(\cdot, X)$ and $u_i(\cdot, X)$ are continuous on \mathcal{A} for any given X. The problem (12) has a solution, μ_0^* . There is an allocation of data μ^* such that $\langle \mu^* \rangle = \langle \mu_0^* \rangle$ and μ^* globally maximizes the firm's profit.

Proof. First, we show the following result: Take any function v(a, X). Let $a^v(b)$ denote the firm's best response at each posterior $b \in \Delta \mathcal{X}$ such that it breaks ties to maximize $v(a, b) := \int_{\mathcal{X}} v(a, X) db(X)$. Then, $v(a^v(b), b)$ is upper semicontinuous. To show this result, suppose to the contrary that there is an $\varepsilon > 0$ and $\{b_n\} \subset \Delta \mathcal{X}$ such that $b_n \to b$ but $v(a^v(b_n), b_n) \ge v(a^v(b), b) + \varepsilon$ for all $n \in \mathbb{N}$. Because \mathcal{A} is compact, we can find a subsequence of $(a^v(b_n))_n$ that converges to (say) a'. We have

$$v(a^{v}(b), b) \ge v(a', b) \ge v(a^{v}(b), b) + \varepsilon,$$
(13)

which is a contradiction. Thus, $v(a^v(b), b)$ is upper semi-continuous on $\Delta \mathcal{X}$.

The above observation implies that $\pi^{\mathcal{N}}(\mu) + \sum_{i \in \mathcal{N}} u_i^{\mathcal{N}}(\mu)$ is upper semicontinuous in μ . Also, because the tie breaking that minimizes consumer *i*'s payoff is equivalent to the one that maximizes $-u_i, -u_i^i(\mu')$ is upper semicontinuous in μ' . Lemma 17.30 of Aliprantis and Border (2006) then implies that $\max_{\mu' \in G(\mu)} (-u_i^i(\mu'))$ is upper semicontinuous in μ . Therefore, the maximand of (12) is upper semicontinuous, and the problem has a solution.

By the construction of $\pi^{\mathcal{N}}(\mu)$ and $(u_i^{\mathcal{N}}(\mu), u_i^i(\mu))_{i \in \mathcal{N}}$ and the same argument as Proposition 6, any solution of (12) has the associated allocation of data μ^* that gives the firm a payoff equal to (12). In the (firm-optimal) equilibrium under μ^* , the firm breaks ties to minimize consumer *i*'s payoff whenever she refuses to sell her data, and the firm breaks ties to maximize the total surplus on the equilibrium path.

The allocation μ^* globally maximizes the firm's profit. To see this, take any allocation μ and equilibrium $E \in \mathcal{E}(\mu)$. Without loss of generality, assume that all consumers sell their (possibly empty) data in equilibrium. We modify E so that upon choosing a, the firm breaks ties to minimize the payoff of a consumer who has refused to sell data; if all consumers sell their data, the firm breaks ties to maximize the sum of the payoffs of all players. Correspondingly the firm offers a price that makes each consumer indifferent between selling and not selling her data. The modification increases total surplus and decreases consumer surplus, so it creates a new strategy profile that gives the firm a weakly greater payoff. The firm's payoff under the new strategy profile is the maximand of (12) evaluated at μ . Therefore, the firm is better off under an equilibrium of μ^* than any $E \in \mathcal{E}(\mu)$.

H The Proof of Claim 3

Proof. Suppose Point 1 holds. $u^* \ge u_{\emptyset}$ holds because the consumer can secure u_{\emptyset} by not sharing data. Suppose to the contrary that $u^* > u_{\emptyset}$, which means that the consumer shares data (say) μ^* . Let u_{μ^*} denote the consumer's payoff in the product market given data μ^* . Let p_1^* denote the equilibrium transfer from the firm to the consumer for data. Since $u^* = u_{\mu^*} + p_1^* > u_{\emptyset}$, the firm can slightly lower p_1^* to strictly increase its profit while collecting μ^* . This is a contradiction, and thus $u^* = u_{\emptyset}$. The sum of the payoffs of the consumer and the firm is at most \bar{w} , and the firm can

always choose to not collect data. Thus, $\pi_{\emptyset} \leq \pi^* \leq \bar{w} - u_{\emptyset}$.

Suppose Point 2 holds. I write π_{μ} for the firm profit in the product market given $\mu \in \Sigma$. Bergemann et al. (2015b) shows that there is a $\mu^* \in \Sigma$ such that $u_{\mu^*} = u_{\emptyset}$ and $\pi_{\mu^*} = \pi^*$. Consider the following strategy profile: The seller sets $p_1 = 0$ and the consumer sells data. Regardless of whether the consumer sells data, the firm sets a price optimally to achieve (π^*, u_{\emptyset}) in the product market. This consists of an equilibrium. In particular, the consumer is willing to share data because doing so does not change her payoff in the product market. \Box