UNIVERSITY OF CAMBRIDGE
Faculty of Economics

CAMBRIDGE WORKING PAPERS IN ECONOMICS

Market segmentation through information

Matthew Elliott
University Of Cambridge

Andrea Galeotti
London Business School

Andrew Koh
Massachusetts Institute of Technology

Wenhao Li
Pennsylvania State University

Abstract

An information designer has precise information about consumers' preferences over products sold by oligopolists. The designer chooses what information to reveal to differentiated firms who, then, compete on price by making personalized offers. We ask what market outcomes the designer can achieve. The information designer is a metaphor for an internet platform who collects data about users and sells it to firms who can, in turn, target discounts and promotions towards different consumers. Our analysis provides new benchmarks demonstrating the power that users' data can endow internet platforms with. These benchmarks speak directly to current regulatory debates.

Reference Details

CWPE	2105
Published	14 January 2021
Updated	13 December 2021
Key Words	Information design, market segmentation, price discrimination
JEL Codes	D43, D83, L13
Website	www.econ.cam.ac.uk/cwpe

Market segmentation through information

Matthew Elliott, Andrea Galeotti, Andrew Koh and Wenhao Li*

December 13, 2021

Abstract

An information designer has precise information about consumers' preferences over products sold by oligopolists. The designer chooses what information to reveal to differentiated firms who, then, compete on price by making personalized offers. We ask what market outcomes the designer can achieve. The information designer is a metaphor for an internet platform who collects data about users and sells it to firms who can, in turn, target discounts and promotions towards different consumers. Our analysis provides new benchmarks demonstrating the power that users' data can endow internet platforms with. These benchmarks speak directly to current regulatory debates.

1 Introduction

The last two decades have witnessed the emergence of a new internet based business model whose revenue stream emanates from collecting and using information about its users to target advertisements. Concerns about competition and users' privacy issues have attracted the attention of antitrust authorities around the world. To guide their thinking towards possible economic harms in downstream markets, antitrust authorities often lean on two benchmarks—complete information, in which all firms know all consumers' preferences, and no information, in which all firms know only the distribution of consumers' preferences.

When firms have complete information, Bertrand competition leads to an efficient pricing equilibrium in which each consumer buys the product for which she is willing to pay the biggest premium above its marginal cost, at a price that makes her indifferent between buying it and her next most preferred product, which is priced at its marginal cost. In comparison, with no information on individual consumers' preferences, gains from trade can be forgone through some consumers being excluded from trading and others buying the wrong product. Hence, comparing these cases reveals that the collection and use of information that permits price discrimination can be welfare enhancing and, in fact, can increase consumer surplus.[1] This provides a salient, cautionary note for regulations that limit the use of information about consumer preferences. For instance, the Council of Economic Advisors (CEA) report on big data and price discrimination observes that

*This paper subsumes "Using Information to Amplify Competition" by Wenhao Li. Elliott: Cambridge University, email mle30@cam.ac.uk; Galeotti: London Business School; Koh: Massachusetts Institute of Technology; Li: Pennsylvania State University. First version: August 2019. We wish to thank Nageeb Ali, Ben Golub, Konstantin Guryev, Navin Kartik, Stephen Morris and Ludovic Renou. Jörg Kalbfuss and Alastair Langtry provided excellent research assistance. Elliott acknowledges funding from the European Research Council under the grant EMBED #757229 and the JM Keynes Fellowships Fund. Galeotti acknowledges funding from the European Research Council under the grant #724356.

[1] While this debate been reopened by the possibility of data-driven price discrimination, its provenance dates back at least to Pigou (1920) and Robinson (1933).

"Economic reasoning suggests that differential pricing, whether online or offline, can benefit both buyers and sellers," and goes on to conclude that "we should be cautious about proposals to regulate online pricing." (Council of Economic Advisors, 2015).[2]

Although full information and no information are useful benchmarks, an intermediary who commands access to consumers' data has more options available than just choosing between either withholding all information or disclosing all information to all firms. Regulatory oversight should be aware of the full power information about consumers grants to intermediaries. For example, in response to privacy concerns, Google has announced that it will replace the use of third-party tracking cookies on its Chrome web browser with its "Privacy Sandbox". The "Privacy Sandbox" groups users into "cohorts" based on their browsing behaviour and targets firms' ads and promotions to these cohorts rather than to individuals. The aim of this paper is to provide a richer set of benchmarks that can shed new light on the effect on downstream competition of technologies like these.

We formulate an information design problem in which the information designer chooses what information about consumers to reveal to competing firms who, then, play a simultaneous pricing game. The information designer can be thought of as an intermediary and, depending on its revenue model, both producer and consumer surplus may be monetized. We consider an intermediary whose objective is increasing in consumer surplus and producer surplus, and the ratios of producer to consumer surplus such an intermediary can achieve though information design. Two cases of particular interest are when the intermediary seeks to maximize producer surplus, which we call producer-optimal outcome, or maximize consumer surplus, which we call consumer-optimal outcome.

Theorem 1 provides a necessary and sufficient condition under which the information designer can induce the producer-optimal outcome. This condition is easier to satisfy when consumers' preferences are polarised, i.e., consumers have a strong taste for their most preferred product (Proposition 1). Although firms compete, under the producer-optimal information design, each consumer buys their most preferred product at a price equal to her valuation for it. In effect, the information designer implements the fully collusive outcome between downstream competitors even in the absence of dynamic incentives.

The producer-optimal information design groups consumers with different valuations for a firm's product and only reveals to this firm which group a given consumer belongs to. Key elements of the information design are that (i) the groups may have to be different for different firms; (ii) a group for firm i consists of some consumers who value i's product most, and who all have the same value for it, and some consumers who value another product more, and with a distribution of lower values for i's product; (iii) each group for a firm i induces a demand function that makes it optimal for firm i to set a price that only the high value consumers—those who value its product most—are willing to pay.

Theorem 2 characterizes the consumer-optimal information design. Taking the prices set by the other firms as given, a firm can always ignore all information it receives and charge the profit maximizing uniform price. Assuming the other firms charge a price equal to

[2] Likewise, the Digital Competition Experts Panel report writes: "There are many reasons why consumers may wish to share their data with a third party. This might enable them to access more accurate price information, to better compare goods and services or to access more tailored advice or recommendations. It may also help support a more effective market, for example where consumers can make a conscious choice to share their data in exchange for some benefit, for example a monetary payment, price discount or free service." (Digital Competition Expert Panel, 2019).

their respective marginal costs therefore provides a theoretical lower bound on the profits each firm can achieve. By subtracting this lower bound on total profits from the total available surplus we obtain an upper bound on consumer surplus. We show that this bound is tight, and hence call it the consumer-optimal outcome. The consumer-optimal information design groups consumers such that: (i) firms receive the same information; (ii) each group contains consumers whose ideal product is the same; (iii) for all groups containing consumers whose ideal product is i, each firm $j \neq i$ charges a price of 0, thereby imposing strong constraints in the pricing for firm i; (iv) each group containing consumers who most prefer i's product is chosen to induce a demand function that makes it optimal for firm i to set a sufficiently low price that all consumers in this group buy product i.

The consumer-optimal information design is closely related to the consumer surplus maximizing information design when each firm is a monopolist (Bergemann, Brooks, and Morris, 2015). Each firm receives the same information about the consumers it is efficient for it to sell to as in the consumer-optimal design when it is a monopolist, but with the key difference that the distribution of consumer values is the difference between consumers' most and second-most preferred products. Unlike in the monopoly setting there is scope for information design to intensify competition. The consumer-optimal design does this through the information provided to *other* firms about the consumers who most prefer i's product. Thus the information provided drives producer surplus down.

Both the consumer-optimal and producer-optimal outcomes are efficient—in both cases all consumers buy the product they value most. We also consider information designs that achieve the other points along the efficient frontier and provide a sufficient condition under which all interior points can be achieved.

A comparison of the producer-optimal and consumer-optimal design is of particular interest to antitrust authorities mandated to protect consumer surplus. First, both the producer-optimal and consumer-optimal information designs are consistent with privacy enhancing technologies like Google Privacy Sandbox: they both pool consumers into flocks and transmit this coarsened information to competing firms. An intermediary who monetizes the firms' side may have strong incentives to develop privacy enhancing technologies that create groups like those in the producer-optimal information design. In this case, enhancing users' privacy in this way is no impediment to extracting consumer surplus—to the contrary, it facilitates it. To avoid this, regulators could formulate guidelines or rules of conduct that ensure such groups of consumers are formed in line with the principles characterizing the consumer-optimal design—i.e., only consumers with similar preferences (and hence the same most preferred product) should be grouped together.[3]

Second, the producer-optimal design often relies on private signals, whereas the consumer-optimal outcome can always be implemented with public signals. A more direct intervention is a non-discrimination requirement at the level of the information provided to firms by an intermediary. That is, regardless of the way in which the intermediary constructs flocks, the aggregated information should be public amongst firms, e.g., requiring Google to put the same consumers into the same flocks for competing firms. To understand the

[3]For instance, the regulators can formulate rules of conduct prescribing that machine learning techniques used to aggregate consumers in flocks to have the objective to group similar consumers together, as in the consumer-optimal design. For related issues about algorithms used directly by competing firms interacting with each other and softening competition (possibly inadvertently) see Calvano et al. (2020).

implication of such a policy, we take the worse case scenario in which the intermediary wishes to maximise producer surplus and we compare how much consumer surplus increases when the policy is introduced. We show that, in a duopoly, this problem can be formulated as a linear programming problem and this allows us to solve numerical examples relatively easily. We apply this numerical method to a canonical Hotelling duopoly model and show that when product differentiation is intermediate this policy increases consumer surplus significantly and, at the same time, does not compromise aggregate efficiency.

Our analysis is based on the assumption that firms will price discriminate if they can. But overt price discrimination can create severe negative publicity.[4] If firms expect consumers to become aware of differential pricing based on consumers' willingness to pay, the ensuing reputational damage may deter the implementation of these practices. There are, however, ways in which such pricing can be concealed. First, a 2019 report by the UK's Digital Competition Experts Panel writes that if firms can "send secret deals to consumers, for example by directly offering discounts via email, the price discrimination becomes entirely opaque." The use of discount codes is widespread and encouraged by internet intermediaries.[5] When firms attempt to conceal price discrimination from consumers in this way it will be relatively challenging to detect it empirically. A web-scraping 'robot,' used in experiments like that run by Cavallo and Rigobon (2016) to compare online and offline prices, does not have the same web-surfing or purchase history as real profiles. As such, firms do not have the opportunity to target them with discount codes (for instance, through social media feeds). Second, in industries where the cost of providing the service being sold depends on the characteristics of the individual (e.g., insurance and credit markets), and in industries that use dynamic demand-based pricing (e.g., flights and ride-hailing), it is hard for consumers to understand what underlies price differences.[6] Again, in such cases, it is challenging for empirical work using publicly available data to identify price discrimination.

A lack of strong evidence for widespread price discrimination does not necessarily imply that such practices are not taking place, albeit in more subtle ways. And this is why the possibility that consumer data are used to facilitate discriminatory pricing has drawn regulatory interest. China's new anti-monopoly guidelines—tailored exclusively to reigning in tech firms—explicitly outline the phenomena for data being used to 'achieve coordinated behaviour' (State Administration for Market Regulation, 2021).[7] In a similar vein, a recent report by the Competition and Markets Authority in the UK reported

[4]Examples of it back-firing include the online retailer Boohoo Group offering the same item of clothing at different prices via its different brands, see `https://www.bbc.com/news/business-56506859`, and Amazon charging different prices back in 2005, see `https://edition.cnn.com/2005/LAW/06/24/ramasastry.website.prices/`.

[5]See Google's marketer playbook and Facebook's webpage for small businesses. Targeted discounts are also ubiquitous in the grocery market. A supermarket is an intermediary which collects detailed data on consumers and price discriminates using coupons. There is also some explicit evidence for price discrimination. Hannak et al. (2014) compare the prices charged to real consumer profiles obtained via Amazon Mechanical Turk. They find evidence that Home Depot, Sears, many travel sites (e.g. Cheaptickets, Orbitz, Priceline etc.) price discriminate.

[6]For example, a 2018 report by the Competition and Markets Authority, the UK's competition regulator found that some home and motor insurance firms use complex and opaque pricing techniques to charge consumers with a higher willingness to pay markedly higher prices (Competition & Markets Authority, 2018).

[7]There is considerable anecdotal evidence for widespread price discrimination occurring in China. A

that 'even if there is limited evidence for personalized pricing, this could change quickly' (Competition & Markets Authority, 2021). Similar issues are highlighted in regulatory documents from the EU, US and Canada.[8]

Overall, our analysis highlights new benchmarks on the power of information in shaping market outcomes. We believe these will serve as useful additions to the toolkit of antitrust authorities. Our paper contributes to a recent literature studying how the informational environment interacts with consumer and producer surplus. Bergemann, Brooks, and Morris (2015) studies price discrimination when a monopolist obtains information about consumer valuations and show that every combination of consumer and producer surplus in the 'surplus triangle' can be obtained through information design. We extend the analysis to an oligopoly setting—the introduction of competition poses additional technical challenges, but also leads to new economic insights which can be related to contemporary regulatory debates.[9]

We investigate what outcomes an intermediary with exogenous consumer data can achieve by sharing the data with firms. Complementary to this, Ali, Lewis, and Vasserman (2020) consider a disclosure game in which a consumer chooses some verifiable information about her preferences to convey to firms. They show that the ability to reveal only partial information can play firms against each other and intensify competition. We focus on a setting in which firms are uncertain about consumer valuations, while Roesler and Szentes (2017) study the converse problem in which consumers, rather than a monopoly, have uncertain valuation; they characterise the signal structure which is best for consumers. Armstrong and Zhou (2019) extend this setting to the duopoly case, and characterise both firm-optimal and consumer-optimal signal structures.

Our paper also relates to a burgeoning literature on markets for information broadly conceived—the transaction, pricing, and design of information (see, e.g., Admati and Pfleiderer (1986), Armstrong and Vickers (2019) Lizzeri (1999), Taylor (2004) Calzolari and Pavan (2006), Bergemann and Bonatti (2015), Bergemann et al. (2018), Acemoglu et al. (2019), Bergemann et al. (2019), Fainmesser and Galeotti (2019), Kehoe et al. (2018) Montes et al. (2019), Jones and Tonetti (2020), Bounie, Dubus, and Waelbroeck (2020); also see Bergemann and Bonatti (2019) for a summary). Perhaps the closest paper to ours is Bounie, Dubus, and Waelbroeck (2018). Like us, they consider an intermediary choosing what information to reveal to firms about consumer valuations. A major focus in their paper is when the intermediary will share information to a single firm or both. They conduct their analysis within a Hotelling model with linear transportation costs and restrict the set of possible information structures that the intermediary can offer firms. We abstract away from the way industry profits are shared between firms and the intermediary and study our information design problem in a general oligopoly model with differentiated products and arbitrary information structures.

survey conducted in 2019 by the Beijing Consumer Association finds that 88% of consumers believe that the practice of big data-enabled price discrimination is significant, and 57% have personally experienced this.

[8]See European Comission (2019), Council of Economic Advisors (2015), Competition Bureau Canada (2018).

[9]A large literature has studied how firms choose which information to share when competing, see, among others, Novshek and Sonnenschein (1982), Vives (1988), Raith (1996). Our approach is very different. We explicitly model an information designer who has granular information at the level of individual valuations and chooses what information to reveal to each firm.

2 Model

2.1 Setup. There is a finite set of firms, indexed $\mathcal{N} = \{1, \ldots, n\}$ each of which produces a single product at zero cost. There is a continuum of consumers with unit mass each of whom demands a single unit inelastically.[10] A consumer of type $\boldsymbol{\theta} := (\theta_1, \theta_2, \ldots, \theta_n)$ obtains utility $\theta_i \in V$ from purchasing from firm i where $V = \{v_1, \ldots, v_K\}$, and $\{0 < v_1 < v_2 < \ldots < v_K < 1\}$, $K < \infty$. The distribution of consumers over V^n is given by the mass function $f : V^n \to [0, 1]$ such that $\sum_{\boldsymbol{\theta} \in V^n} f(\boldsymbol{\theta}) = 1$, which is common knowledge. We will primarily work with the support of the distribution, which we denote with $\Theta := \text{supp} f \subseteq V^n$.

We denote the consumer types that value product i the most by $E_i := \{\boldsymbol{\theta} \in \Theta : \theta_i > \theta_j \text{ for all } j \neq i\}$. We assume that all consumers have strict preferences so that there is no mass on the types $\{\boldsymbol{\theta} \in V^n : |\operatorname{argmax}_j \theta_j| \geq 2\}$. This implies that $\{E_i\}_{i \in \mathcal{N}}$ partitions Θ. We focus on discrete type distributions and assume preferences are strict just to streamline the exposition. All our main results extend to a continuous version of the model; see Online Appendix B4.

An *information designer*, knowing the valuation of each consumer for each product, chooses how to distribute information about consumer preferences across firms. For each firm, the information designer commits in advance to a signal structure it will provide about each type of consumer. Thus, the information designer chooses a mapping

$$\boldsymbol{\psi} : \Theta \to \Delta(\mathcal{M})$$

from consumer types to a joint probability distribution over messages $\Delta(\mathcal{M})$ where $\mathcal{M} = \prod_{i \in \mathcal{N}} M_i$, and $M_i = [0, 1]$ is the message space for firm i. Denote the set of message functions with $\boldsymbol{\Psi}$ and for $\boldsymbol{\psi} \in \boldsymbol{\Psi}$, $\psi_i(\boldsymbol{\theta})$ is the marginal distribution of messages firm i receives.

Call $m_i \in M_i$ a message realisation for firm i. Given the messages received for each consumer, firms then play a simultaneous move pricing game. A pure strategy for firm $i \in \{1, \ldots, n\}$ is $p_i : M_i \to [0, 1]$. A mixed strategy for firm i is $\sigma_i : M_i \to \Delta([0, 1])$.

The information designer can be thought of as an intermediary that has detailed information about consumer preferences, and chooses what market segmentation of consumers to present to each firm.[11] We assume the only role of the intermediary is to give better firm-specific information to different firms about consumers' preferences. In some situations, however, platforms might also be able to withhold some firms' access to certain consumers.[12] We discuss the implication of this at the end of Section 3, which is supported by Supplementary Appendix B.1.

[10] All results translate into an alternate setting with a single consumer of uncertain type.

[11] Bergemann and Morris (2013, 2016) consider many-player settings and examine how the informational environment maps to resultant equilibria. In the special case with a single receiver, Kamenica and Gentzkow (2011) show that concavification of the designer's payoff as a function of receiver's posteriors binds the designer's maximum attainable utility and characterises the optimal signal structure. However, there are well-known difficulties applying such techniques when the type space is large. A contribution of our analysis is to show that it can be helpful to reframe certain information design problems as matching problems.

[12] See Bergemann and Bonatti (2019) for a discussion of this distinction.

3 Producer-Optimal information Design

We first characterize conditions under which there exist information structures such that, in an equilibrium of the resultant subgame, the following property holds:

P **(Full Surplus Extraction)** each consumer of type $\boldsymbol{\theta} \in \Theta$ pays $\max_{i \in \mathcal{N}} \theta_i$.

Condition P characterizes the fully collusive outcome when transfers are possible. Note that an equilibrium that satisfies conditions P is efficient. Let $\Gamma(\boldsymbol{\psi})$ denote the pricing subgame induced by the message function $\boldsymbol{\psi}$, and let $\boldsymbol{\Gamma^*}$ denote the set of induced games in which there exists an equilibrium satisfying condition P, and let

$$\boldsymbol{\Psi}^* := \{\boldsymbol{\psi} : \Gamma(\boldsymbol{\psi}) \in \boldsymbol{\Gamma}^*\}$$

be the set of message functions that the information designer can use to fulfil condition P. We refer to $\boldsymbol{\psi} \in \boldsymbol{\Psi}^*$ as a producer-optimal information design and to the induced outcome as the producer-optimal outcome.

3.1 Simplifying the problem. In what follows we characterise producer-optimal information designs. In order to satisfy condition P all consumers must buy their most preferred product and each firm i must be able to perfectly separate the consumers in E_i with different values for its product so that it can charge all such consumers their valuation. Recalling that $\psi_i(\boldsymbol{\theta})$ is a probability distribution over the messages i can receive for a given type $\boldsymbol{\theta}$, we denote the support of this distribution by $\text{supp}\big(\psi_i(\boldsymbol{\theta})\big)$. Hence, $\boldsymbol{\psi} \in \boldsymbol{\Psi}^*$ only if for all $i \in \mathcal{N}$:

$$\text{supp}\big(\psi_i(\boldsymbol{\theta})\big) \cap \text{supp}\big(\psi_i(\boldsymbol{\theta}')\big) = \emptyset \text{ for all } \boldsymbol{\theta}, \boldsymbol{\theta}' \in E_i \text{ such that } \theta_i \neq \theta_i'.$$

Given that firm i must receive different signals for consumers in E_i with different values for its product, but must charge the same price to consumers in E_i with the same value for its product, it is without loss of generality to set the message received for such consumers to $m_i = \theta_i$. This yields the natural interpretation of messages as price (or discount) recommendations.

In the monopoly case the producer optimal design gives the monopoly perfect information about all the consumers for whom it is efficient for the monopolist to sell to. When firms compete, an analogous information structure is to give each firm perfect information about all the consumers it should sell to (i.e., those in E_i) and no further information about any other consumers, i.e., for consumers not in E_i firm i would receive the same null message. However, for such an information structure to induce a producer-optimal outcome, firm j must be extracting all surplus from its consumers in E_j. But if so, firm i could profitably steal some of these consumers by setting a sufficiently low price $p_i > 0$ upon receiving the null message. In fact, by the same argument, firm i can never receive a message different from one of the valuations a consumer in E_i has for i's product. If it did, it would know the consumer is not in E_i and it could, for example, set a price $p_i \in (0, v_1)$ to profitably steal the consumer from firm j.

In order to prevent business stealing like this the information design must make it costly for a firm i to sell to consumers in E_j for $j \neq i$. This requires firm i to be unable to distinguish a consumer in E_j from one with a higher valuation in E_i—that way, firm i trades off charging a high price and selling only to the consumer in E_i (niche strategy) versus charging a low price and selling to both the consumer in E_i and E_j (mass strategy). By controlling the probability that the consumer is in E_i versus E_j the information designer can make it incentive compatible for firm i to only sell to consumers in E_i.

Lemma 1 formalizes the above arguments. Let $M_i' \subseteq V$ be the set of valuations that consumers in E_i have for product i and let $\mathcal{M}' := \prod_{i \in \mathcal{N}} M_i'$.

Lemma 1. A producer-optimal information design exists when the set of available messages is $\mathcal{M}$ if and only if an information design exists when the set of available messages is $\mathcal{M}'$.

3.2 Restricting the messages. From now on, we restrict type space to $\mathcal{M}'$. We have shown that it is without loss of generality and we can let $\text{supp}\big(\psi_i(\boldsymbol{\theta})\big) = \{\theta_i\}$ for all $\boldsymbol{\theta} \in E_i$ and to let firm i receive a message $m_i = \theta_i$ for a consumer type $\boldsymbol{\theta} \in E_i$. This implies that we can think of the information design problem as one of assigning the consumers not in E_i to the same set of messages that firm i receives about consumers in E_i.

Recall that the message function $\psi_i(\boldsymbol{\theta})$ denotes the distribution over messages firm i receives about consumers of type $\boldsymbol{\theta}$. We let $\psi_i(m_i|\boldsymbol{\theta})$ denote the proportion of consumers of type $\boldsymbol{\theta}$ for whom i receives message m_i. It is helpful to define the functions $(\mu_i : (\Theta \setminus E_i) \times M_i' \to [0,1])_{i=1}^n$ which fulfil

$$\mu_i(\boldsymbol{\theta}, m_i) := \psi_i(m_i|\boldsymbol{\theta}) f(\boldsymbol{\theta}).$$

This gives the mass of consumers of type $\boldsymbol{\theta} \in \Theta \setminus E_i$ that firm i receives message m_i for. Recall that firm i also receives the message m_i for consumers in E_i with value m_i. This yields the interpretation of $\mu_i(\cdot,\cdot)$ as a matching function. Note that conditions imposed on $\boldsymbol{\mu}$ can be interpreted as conditions directly imposed on $\boldsymbol{\psi}$.

We now consider what restrictions must be satisfied for condition P to be satisfied in equilibrium. A first observation, as already argued, is that $\boldsymbol{\psi} \in \boldsymbol{\Psi}^*$ only if: for all firms i and all $\boldsymbol{\theta} \in E_i$ firm i receives the message $m_i = \theta_i$ with probability 1, i.e.,

$$\psi_i(m_i|\boldsymbol{\theta}) = \begin{cases} 1 & \text{if } m_i = \theta_i \\ 0 & \text{otherwise.} \end{cases} \qquad \textbf{(Separation)}$$

Furthermore, the total mass of consumers matched to some message $m_i \in M_i'$ must be consistent with the actual distribution of consumers. That is, $\boldsymbol{\psi} \in \boldsymbol{\Psi}^*$ only if: for all firms i, and all $\boldsymbol{\theta} \notin E_i$

$$\sum_{m_i \in M_i'} \mu_i(\boldsymbol{\theta}, m_i) = f(\boldsymbol{\theta}). \qquad \textbf{(Consistency)}$$

Next, in order to satisfy condition P, upon receiving the message m_i firm i must set a price $p_i = m_i$ while all consumers must buy the product they value highest, leaving consumers with no surplus. Consider a consumer of type $\boldsymbol{\theta}$ who most prefers j's product.

If firm i receives message m_i about this consumer it must set the consumer a price m_i. However, if m_i is less than the consumer's value for i's product, the consumer will deviate and buy from i instead of j, violating P. This gives us that $\boldsymbol{\psi} \in \boldsymbol{\Psi^*}$ only if for all $\boldsymbol{\theta} \notin E_i$,

$$\mu_i(\boldsymbol{\theta}, m_i) = 0 \text{ for all } \theta_i \geq m_i, \qquad \textbf{(Consumer IC)}$$

which ensures that when each firm i is pricing to extract all surplus from consumers in E_i, it is incentive compatible for these consumers to buy product i. **Consumer IC** also implies that if firm i receives a message m_i and sets a price $p_i > m_i$ it makes no sales.

We finally need to derive the conditions for firm i to not want to deviate to a price $\hat{p}_i < m_i$ upon receiving message m_i. By charging a price lower than m_i firm i might be able to capture extra consumers. To prevent this from being profitable the inframarginal losses this entails for consumers in E_i (now being charged a price less than their valuations) must be greater than the extra profits made via any additional sales. Formally, $\boldsymbol{\psi} \in \boldsymbol{\Psi^*}$ only if for all $m_i \in M_i'$ and for all $\hat{p}_i < m_i$,

$$\underbrace{(m_i - \hat{p}_i) \sum_{\boldsymbol{\theta'} \in E_i : \theta_i' = m_i} f(\boldsymbol{\theta'})}_{\text{Inframarginal losses}} \geq \underbrace{\hat{p}_i \sum_{\boldsymbol{\theta'} \in \Theta \setminus E_i : \theta_i' \geq \hat{p}_i} \mu(\boldsymbol{\theta'}, m_i)}_{\text{Business stealing gains}} \qquad \textbf{(Firm IC)}$$

We have argued that these conditions "must" be satisfied by a producer-optimal information design. Lemma 2 shows that satisfying these conditions is also sufficient for a producer-optimal information design to be achieved.

Lemma 2. A producer-optimal information design exists if and only if there exists an information design $\boldsymbol{\psi}$ (with induced matching functions $\boldsymbol{\mu}$) which, for all firms $i \in \mathcal{N}$, satisfies **Consumer IC**, **Firm IC**, **Separation** and **Consistency**.

3.3 Existence of producer-optimal information design. Our goal is to find a necessary and sufficient condition on the distribution of consumer types for there to exist a producer-optimal information design. Lemma 2 implies that we can restrict our attention to whether there exists an information structure satisfying **Separation**, **Consistency**, **Consumer IC** and **Firm IC**.

By **Separation** and **Consistency** we can frame the problem as one of matching: we need to match consumers of types that firm i should not sell to, to consumers of types that firm i should sell to. **Consumer IC** gives us the first restriction on such matchings: A consumer type $\boldsymbol{\theta} \notin E_i$ with value for firm i's product greater than v_i cannot be matched to a consumer type $\boldsymbol{\theta'} \in E_i$ with value for firm i's product equal to v_i.

We are then left to derive conditions on the distribution of consumer types under which it is possible to find a matching that also satisfies **Firm IC**. The construction has two steps.

In the first step we characterize the distribution of types not in E_i that can be matched to each message $m_i \in E_i$ firm i receives and that makes firm i indifferent between charging *any* price $\hat{p}_i \leq m_i$ upon receiving m_i. This matching causes every possible firm i's IC constraint to bind; for this to happen we must have that, for all $m_i \in E_i$

$$\sum_{\boldsymbol{\theta'} \in \Theta \setminus E_i : \theta_i' \geq \hat{p}_i} \mu(\boldsymbol{\theta'}, m_i) = \frac{(m_i - \hat{p}_i)}{\hat{p}_i} \sum_{\boldsymbol{\theta'} \in E_i : \theta_i' = m_i} f(\boldsymbol{\theta'}) \quad \text{for all } \hat{p}_i \leq m_i.$$

It is, then, helpful to define for each firm i

$$G_i(\hat{p}_i, m_i) := \begin{cases} \dfrac{(m_i - \hat{p}_i)}{\hat{p}_i} \displaystyle\sum_{\boldsymbol{\theta}' \in E_i : \theta'_i = m_i} f(\boldsymbol{\theta}') & \text{if } \hat{p}_i \leq m_i \\ 0 & \text{otherwise,} \end{cases}$$

which gives the exact mass of consumers not in E_i with valuation for i's product in $[\hat{p}_i, m_i)$ that can be matched to consumers in E_i with value m_i to bind firm i's incentive compatibility condition upon receiving the message m_i. The next example illustrates the construction of the function G.

Example 1. There are two firms and four possible consumer values for each product, which we set equal to 0.2, 0.4, 0.6 and 0.8. Panel (a) of Figure 1 identifies the consumer types for which it is efficient for firms 1 and 2 to sell to. Panel (b) illustrates the mass of consumers of each type.

Figure 1: The set E_1 and E_2, and distribution of consumers types in Example 1.

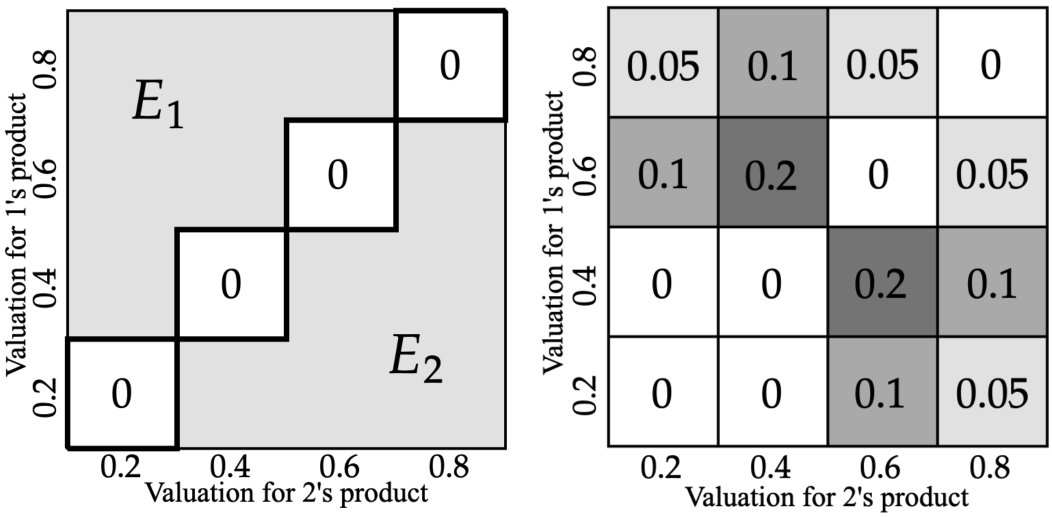

(a) (b)

Consider the 0.2 mass of consumers in E_1 with valuation 0.8 for 1's product. Firm 1 must receive the message $m_1 = 0.8$ for these consumers and offer them a price of $p_1 = 0.8$; this leads to a profit of 0.16. Suppose firm 1 deviates to a price of $\hat{p}_1 = 0.6$. The fraction of consumers in E_2 with a valuation for product 1 in $[0.6, 0.8)$ that can be matched to the message $m_1 = 0.8$ to make firm 1 indifferent between following the price recommendation of 0.8 and lowering the price to 0.6 is $G_1(0.6, 0.8) = 0.2/3$. Similarly, to make firm 1 indifferent about deviating to the price $\hat{p}_1 = 0.4$ we can have a mass of consumers in E_2 equal to $G_1(0.4, 0.8) = 0.2$ with value in $[0.4, 0.8)$ for 1's product. Finally, to make firm 1 indifferent about deviating to $\hat{p}_1 = 0.2$ we must have $G_1(0.2, 0.8) = 0.6$ consumers in E_2 with value in $[0.2, 0.8)$ for 1's product.

Next, consider the message $m_1 = 0.6$. There is a mass of consumers equal to 0.3 in E_1 for whom firm 1 must receive this message and charge a price of 0.6, obtaining a profit of 0.18. Note that by firm 1's IC constraint none of the 0.05 mass of consumers in E_2

with value 0.6 for 1's product can be assigned to this message—otherwise either firm 1 can profitably just undercut, or all these consumers in E_2 buy from firm 1, in which case consumer IC would be violated. Firm 1 is indifferent about deviating to $\hat{p}_1 = 0.4$ when an additional mass of $G_1(0.4, 0.6) = 0.15$ consumers in E_2 with a value for product 1 in $[0.4, 0.6)$ are matched to message $m_1 = 0.6$. Likewise, for firm 1 to be indifferent about deviating to $\hat{p}_1 = 0.2$, it must gain a mass of consumers in E_2 equal to $G_1(0.2, 0.6) = 0.6$.

Finally, as there are no consumers in E_1 with value 0.2 or 0.4 firm 1 never receives the message 0.2 or 0.4 for consumers in E_1 and so no consumers in E_2 can be assigned such messages either. ■

The second step of our construction is to sum $G_i(\hat{p}_i, m_i)$ across all consumers' types whose ideal product is i, i.e., across all messages $m_i \in M_i'$. Formally, we define the function

$$H_i(\hat{\theta}_i) := \sum_{m_i \in M_i'} G_i(\hat{\theta}_i, m_i) = \sum_{m_i > \hat{\theta}_i} G_i(\hat{\theta}_i, m_i),$$

which gives the total mass of consumers of types not in E_i with valuations for i's product more than or equal to $\hat{\theta}_i$ that will be matched when all firm i's IC constraints bind.

The value of $H_i(\hat{\theta}_i)$ is important because it gives us a maximum mass of types not in E_i with a value for product i larger than $\hat{\theta}_i$ that we can have if we want to construct a producer-optimal information design. Consider a value $\hat{\theta}_i$ for i's product. The mass of consumers not in E_i that have at least this value for i's product is

$$\sum_{\boldsymbol{\theta}' \in \Theta \setminus E_i : \theta_i' \geq \hat{\theta}_i} f(\boldsymbol{\theta}').$$

If this mass of consumers is greater than $H_i(\hat{\theta}_i)$ then, by construction of H_i, we know there is no way to assign all these consumers any message $m_i \in \{M_i' : m_i > \hat{\theta}_i\}$ without firm i sometimes having a profitable deviation to capture some of these consumers. On the other hand, if this mass of consumers is weakly less than $H_i(\hat{\theta}_i)$ for all $\hat{\theta}_i$, then there is a way of assigning the consumers not in E_i messages $m_i \in V$ such that firm i wants to follow the price recommendation $p_i = m_i$ upon receiving the message m_i. This leads us to the necessary and sufficient condition stated below in Theorem 1 for a producer-optimal information design to exist.

Theorem 1. A producer-optimal information design exists if and only if for all firms $i \in \mathcal{N}$ and all consumer valuations $\hat{\theta}_i \in V$,

$$H_i(\hat{\theta}_i) \geq \sum_{\boldsymbol{\theta}' \in \Theta \setminus E_i : \theta_i' \geq \hat{\theta}_i} f(\boldsymbol{\theta}').$$

To gain further intuitions for Theorem 1 we observe that the key condition can be rewritten as

$$\sum_{m_i > \hat{\theta}_i} (m_i - \hat{\theta}_i) \sum_{\boldsymbol{\theta}' \in E_i : \theta_i' = m_i} f(\boldsymbol{\theta}') \geq \hat{\theta}_i \sum_{\boldsymbol{\theta}' \in \Theta \setminus E_i : \theta_i' \geq \hat{\theta}_i} f(\boldsymbol{\theta}'),$$

where the left-hand side is firm i's aggregate infra-marginal losses from deviating to $\hat{\theta}_i$ after every price recommendation $m_i > \hat{\theta}_i$, and the right-hand side is the maximum business stealing profit that firm i can hope to obtain when setting a price $\hat{\theta}_i$.

Example 1 (continued). The condition in the Theorem 1 is satisfied for firm 1 because

$$\begin{aligned} H_1(0.2) &= 1.2 > f(0.2,0.6) + f(0.2,0.8) + f(0.4,0.6) + f(0.4,0.8) + f(0.6,0.8) = 0.5 \\ H_1(0.4) &= 0.35 = f(0.4,0.6) + f(0.4,0.8) + f(0.6,0.8) = 0.35 \\ H_1(0.6) &= 0.2/3 > f(0.6,0.8) = 0.05 \end{aligned}$$

The same can be verified for firm 2. Hence, there exists a producer-optimal information design.

Figure 2: A producer-optimal information design in Example 1. The proportion of each square shaded blue and red denotes the proportion of each type assigned the messages 0.8 and 0.6 provided to each firm respectively.

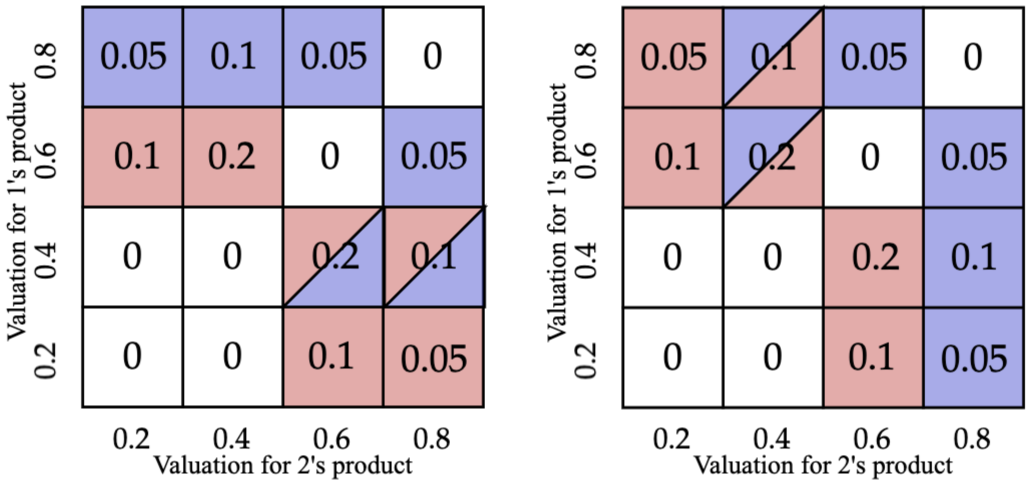

(a) Information for firm 1

(b) Information for firm 2

An example of a producer-optimal information design is shown in Panel (b) of Figure 2. The proportion of each square shaded blue and red denotes the proportion of each type assigned the messages 0.8 and 0.6 provided to firm 1, respectively. Upon receiving message 0.6 firm 1 learns this group contains 0.3 mass of consumers in E_1 with valuation 0.6 for product 1, 0.15 mass of consumers in E_2 with valuation 0.4 for 1's product, and 0.15 mass of consumers in E_2 with valuation 0.2 for product 1. Firm 1 is indifferent between charging 0.6 and 0.4 and strictly prefers to charge 0.6 instead of 0.2. Upon receiving message 0.8 firm 1 learns that this group contains all 0.2 mass of consumers in E_1 with valuation for firm 1's product 0.8, a mass of 0.05 consumers in E_2 with valuation of 0.6 for product 1, and a mass of 0.15 of consumers in E_2 with valuation 0.4 for product 1. With this information, firm 1 strictly prefers to charge 0.8 instead of charging 0.6 or 0.2 and is indifferent between charging 0.8 and 0.4. ■

We now investigate when the condition in Theorem 1 is more likely to be satisfied. In online Appendix B5 we consider some canonical distributions in the continuous version of our model. An interesting benchmark is the duopoly with anticorrelated values such that $v_2 = 1 - v_1$ for all consumers (like the Hotelling model). In this case, when consumers' valuations are uniformly distributed over the unit interval, the producer optimal outcome is feasible. Moving away from the anticorrelated setting a second benchmark is when valuations are uniformly distributed over the unit square. Again, the producer optimal outcome is feasible for this distribution.

A third case of interest, back in the Hotelling setting, is when valuations are drawn from a truncated normal distribution with mean 1/2 and variance σ^2. In this case the producer optimal outcome is feasible for all $\sigma > 0.15$. Note that as the variance increases consumers have stronger preferences for one product over the other. We now show that this holds generally; When "polarization" like this is sufficiently strong the producer optimal outcome is feasible.

Definition 1. Consumers' preferences are more polarized under distribution $\tilde{f}$ relative to distribution f whenever

(i) For all $i \in \mathcal{N}$,

$$\sum_{\boldsymbol{\theta}' \in E_i : \theta'_i \geq \hat{\theta}_i} f(\boldsymbol{\theta}') \leq \sum_{\boldsymbol{\theta}' \in E_i : \theta'_i \geq \hat{\theta}_i} \tilde{f}(\boldsymbol{\theta}') \quad \text{for all } \hat{\theta}_i \in V,$$

(ii) For all $i, j \in \mathcal{N}$,

$$\sum_{\boldsymbol{\theta}' \in E_j : \theta'_i \geq \hat{\theta}_i} f(\boldsymbol{\theta}') \geq \sum_{\boldsymbol{\theta}' \in E_j : \theta'_i \geq \hat{\theta}_i} \tilde{f}(\boldsymbol{\theta}') \quad \text{for all } \hat{\theta}_i \in V.$$

Condition (i) states that consumers who, under f, preferred firm i's product, under $\tilde{f}$, prefer it by more; condition (ii) states that consumers who, under f, preferred j's product now, under $\tilde{f}$, prefer i's product less. Intuitively, as consumers' preferences become more polarized this slackens constraints posed by **Firm IC** which in turn strengthens the designer's ability to implement a producer-optimal information design.

Proposition 1 (Polarization aids segmentation)**.** Assume consumers' preferences are more polarized under $\tilde{f}$ relative to f. If a producer-optimal information design exists under f then it also exists under $\tilde{f}$, i.e., $\boldsymbol{\Psi^*} \neq \emptyset$ under f then $\boldsymbol{\Psi^*} \neq \emptyset$ under $\tilde{f}$.

There are various ways in which consumers' preference can become more polarised. An obvious avenue is for firms to make their products more differentiated.[13] Beyond adjusting their product offerings firms can also affect the extent of product differentiation through advertising.[14] Proposition 1 shows that all these managerial options to increase polarization can help make the producer-optimal information structure feasible.

On the other hand, firm actions that uniformly increase the value consumers place on one product relative to another, thereby skewing the mass of consumer valuations towards a particular firm—and, in the process, shrinking the fraction of consumers it is efficient

[13]This includes vertical differentiation when the marginal cost of production is increasing in the quality of the product. Suppose, for example, that there are two firms. They produce products at quality q_1 and $q_2 = 1 - q_1$ with a marginal cost of production q_i, and consumer j has a willingness to pay equal to $\alpha + \gamma_j q_i$ where $\alpha > 0$ is a parameter and γ_j is drawn from a distribution with support $[0, 1]$. When $q'_1 > q_1$ consumers preferences for the available products will be more polarized.

[14]We refer to Johnson and Myatt (2006) for many examples on how firms can use product design and advertising to shape the distribution of consumers' preferences.

for the other firms to serve—can inhibit the ability to achieve the producer-optimal outcome. This is because the firm with a reduced consumer base has stronger incentives to undercut other firms. This implies that imbalanced competition in which some firms have a much smaller market share than others can in fact more severely inhibit an intermediary from implementing a producer-optimal information design than balanced competition, where market shares are more symmetrically distributed. Weak firms can pose sterner competitive constraints on pricing than strong ones.[15]

It is can also be seen that a merger weakens the condition in Theorem 1. Consider a merger between two firms i and j, with $E_i \neq \emptyset$ and $E_j \neq \emptyset$, into the firm k (with the same product offering). The new condition that must be satisfied for the producer optimal outcome will be weaker than either of the two conditions required prior to the merger. This is because $E_k \supset E_i$ and $E_k \supset E_j$ which slackens the condition under which the producer optimal outcome is feasible for the merged firm, without affecting the conditions for the other firms.

We conclude this section with three observations.

First, generically it is possible to make all incentive compatibility constraints for all firms slack.[16] Hence, the producer-optimal information design does not rely on making firms indifferent. An implication of this is that the information design can be made robust to slight misspecifications of the firms' incentives. For example, if firms maximize a convex combination of profits and revenues with weight $\alpha \in [0, 1]$ on profits, the producer-optimal information design based on $\alpha = 1$ generically works for values of $\alpha \in [\bar{\alpha}, 1]$ for some $\bar{\alpha} < 1$.

Second, we have thus far focused on the benchmark case in which the only role of the intermediary is to provide information about consumer preferences. In practice, however, the intermediary might, in addition to providing information about preferences, also control firms' access to certain consumers. As Bergemann and Bonatti (2019) emphasise, this distinction is key for understanding market outcomes. We show that, if in addition to designing information, the intermediary can restrict some firms' access to certain consumers, this weakens the conditions under which the producer-optimal outcome can be achieved. In this sense, access to consumers is complementary to information design.[17]

Third, in practice, in order to (approximately) implement the producer-optimal information design, a platform has to first estimate different consumers' values for different products. A standard starting point in the empirical industrial organization literature is to view products as bundles of characteristics over which consumers have preferences, and to then estimate these preferences. A large literature building on Berry et al. (1995) shows how the joint distribution of consumer values can be estimated in this way using information about product characteristics and demand data. This approach can be combined with detailed consumer data, like search data (Armona et al., 2021), to more accurately estimate consumer preferences. Once preferences over characteristics have been

[15]This is similar to the argument that, in the presence of switching costs, firms with a smaller customer base can be a stronger competitive constraint on market behaviour than more established firms (Klemperer, 1995). Based in part on this logic the UK antitrust authorities prohibited the acquisition of Abbey National by Lloyds TSB Group in 2001.

[16]In Appendix B.2 we give a necessary and sufficient condition for this to be possible which generically holds when the condition in Theorem 1 is satisfied.

[17]We formalize these claims in Proposition 3 in Online Appendix B.1.

estimated, it is straightforward to back out consumers' valuations for different products.

4 Characterization of Consumer-Optimal Design

We now characterize information structures that maximize equilibrium consumer surplus. Let us start by establishing an upper bound on consumer surplus. We do so by first establishing a lower bound on total producer surplus.

Allowing other firms to randomize their prices, let $\boldsymbol{P}_{-i} \in \Delta(\Theta \times [0,1]^{n-1})$ denote the joint distribution of types and other firms' prices which firm i takes as given. Letting $\Pi_i^*(\boldsymbol{P}_{-i})$ be firm i's profits from charging an optimal uniform price, this is minimized when all other firms charge all consumers a price of zero.[18] Thus, a lower bound on i's profits in any equilibrium is:

$$\underline{\Pi_i^*} = \max_{p_i} p_i \sum_{\substack{\boldsymbol{\theta}:\theta_i - p_i \geq \theta_j \\ \text{for all } j}} f(\boldsymbol{\theta}) = \max_{p_i} p_i \sum_{\substack{\boldsymbol{\theta}\in E_i:\theta_i - p_i \geq \theta_j \\ \text{for all } j}} f(\boldsymbol{\theta}).$$

Letting $S^* = \sum_{i=1}^{n} \sum_{\boldsymbol{\theta}\in E_i} f(\boldsymbol{\theta})\theta_i$ be the total surplus available in the economy, a corresponding upper bound on consumer surplus in all equilibria is

$$CS^* = S^* - \sum_{i=1}^{n} \underline{\Pi_i^*}.$$

It remains to show that this upper bound is tight i.e., there exists some information structure which induces this welfare outcome in the resultant subgame.

A possible candidate for a consumer-optimal information structure is one in which all firms publicly learn which consumers prefer product i the most, for all products i, and nothing else. Since we want to maximise consumer surplus, we assume that all firms j, other than i charges a price of 0 to the group of consumers E_i. Given this, the effective willingness to pay of a consumer of type $\boldsymbol{\theta} \in E_i$ for product i is $\theta_i - \max_{j\neq i}\theta_j$. We denote the distribution of effective valuations for product i by $d_i : [0,1] \to [0,1]$, with:

$$d_i(\hat{\theta}_i) := \sum_{\substack{\boldsymbol{\theta}\in E_i: \\ \theta_i - \max_{j\neq i}\theta_j = \hat{\theta}_i}} f(\boldsymbol{\theta}) \quad \text{for all } \hat{\theta}_i \in [0,1].$$

Firm i, by charging p_i to consumers in E_i faces demand $\sum_{\theta\geq p_i} d_i(\theta)$ and makes a profit of $\pi_i = p_i \sum_{\theta\geq p_i} d_i(\theta)$. Hence, firm i is, in effect, a monopolist facing this effective demand schedule and so will set a price $p_i^* = \operatorname{argmax}_{p_i} p_i \sum_{\theta\geq p_i} d_i(\theta)$ to achieve profits of $\underline{\Pi_i^*}$.

[18] Facing $\boldsymbol{P}_{-i}$, firm i makes

$$\Pi_i^*(\boldsymbol{P}_{-i}) = \max_{p_i} \quad p_i \sum_{\mathbf{p}_{-i}} \sum_{\boldsymbol{\theta}\in\Theta} \mathbb{1}(\theta_i - p_i \geq \max\{0, \max_j(\theta_j - p_j)\})\boldsymbol{P}_{-i}(\boldsymbol{\theta}, \mathbf{p}_{-i})$$

which is minimized when $\boldsymbol{P}_{-i}(\boldsymbol{\theta}, \mathbf{0}) = f(\boldsymbol{\theta})$ for all $\boldsymbol{\theta}$. The case without finite support extends straightforwardly. Further note that this expression assumes that consumers break ties in favour of firm i. However, this remains a lower bound on i's profits if consumers broke ties against firm i—firm i can undercut slightly to make profits arbitarily close to $\Pi_i^*(\boldsymbol{P}_{-i})$.

There are two possibilities. The first is that at p_i^* all consumers in E_i buy firm i's product.[19] In this case the information design is consumer-optimal. This is the case in Example 1.

Example 1 (continued). Panel (a) of Figure 3 reproduces the distribution of consumer values over V^2, while panel (b) illustrates the distribution of effective valuations d_i for each firm $i = 1, 2$.

Figure 3: A distribution of consumers' types under which all firms publicly learn only which consumers prefer product i the most is a consumer-optimal information design.

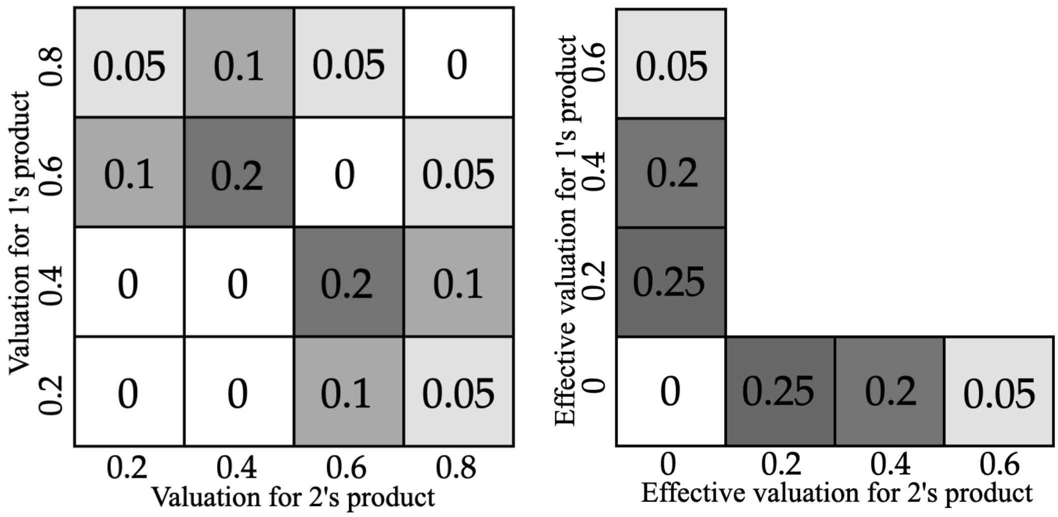

(a) (b)

Consider firm 1. Given that firm 2 charges a price of 0 to all E_1 consumers, there is a mass 0.05 of consumers in E_1 that are just willing to pay a price 0.6 for 1's product, a mass 0.2 that will pay a price 0.4 and a mass 0.25 that will pay a price of 0.2. Thus, firm 1 maximizes its profits by charging a price of 0.2. By symmetry, firm 2 maximizes its profits by charging a price 0.2 to consumers in E_2. Thus the information design in which both firms learn whether each consumer is in E_1 or E_2 induces an equilibrium in which firms obtain the lower bound on equilibrium profits and all consumers buy their most preferred product. Hence, this information design is consumer-optimal.■

The second possibility is that firm i would optimally chooses a uniform price to consumers in E_i higher than some such consumers are willing to pay (i.e., $p_i^* > \min_{\boldsymbol{\theta} \in E_i} \min_{j \neq i} \theta_i - \theta_j$). In this case, the induced allocation is neither efficient nor consumer-optimal. We illustrate this in the next example then show how to construct an information design which is consumer-optimal.

Example 1 (modified). Consider the distribution of consumer values shown in panel (a) of Figure 4 and, in panel (b), the distribution of effective valuations.

When firm 2 charges a price 0 to all consumers in E_1, there is a mass 0.1 of consumers in E_1 that are just willing to pay a price 0.6 for 1's product, a mass 0.3 that will pay a price 0.4 and a mass 0.1 that will pay a price of 0.2. Thus, firm 1's optimal price is $p_1^* = 0.4$.

[19]That is, $p_i^* = \min_{\boldsymbol{\theta} \in E_i} \min_{j \neq i} \theta_i - \theta_j$.

Figure 4: A distribution of consumers' types under which all firms publicly learn only which consumers prefer product i the most is not consumer-optimal.

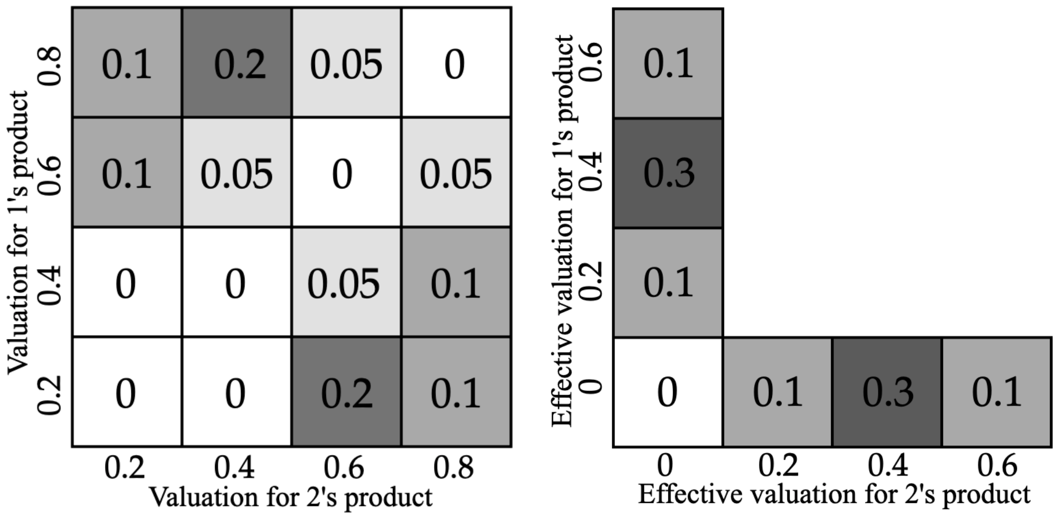

(a) (b)

and this generates a profit of $\underline{\Pi_1^*} = 0.16$. However, this outcome is inefficient because it excludes the 0.1 mass of consumers in E_1 with effective valuation 0.2 (the 0.05 mass of consumers in E_1 with valuations $\theta_1 = 0.6$ and $\theta_2 = 0.4$, and the 0.05 mass of consumers with valuations $v_1 = 0.8$ and $\theta_2 = 0.6$).

Next consider the following alternative information design. Firm 1 receives three messages. The first message is received by firm 1 for all consumers with effective valuation 0.2 for product 1 and a mass 0.1 of consumers with effective valuation 0.4. It is then optimal for firm 1 to charge a price of 0.2 to these consumers; the profit obtained is $0.04 = \underline{\Pi_1^*}/4$. The second message reveals to firm 1 a group of consumers of mass 0.15 with effective valuation of 0.4 for product 1, so firm 1 charges 0.4 to them and obtains a profit of $0.06 = 3\underline{\Pi_1^*}/8$. Finally, the third message reveals to firm 1 a group of consumers composed of a mass 0.05 of consumers with effective valuation 0.4 for product 1 and the mass 0.1 of consumers with effective valuation 0.6. It is optimal for firm 1 to charge a price 0.4 to this group and obtain a profit of $0.06 = 3\underline{\Pi_1^*}/8$.

All consumers in E_1 now buy from firm 1 and so this is efficient. Furthermore, firm 1 obtains a profit of $\underline{\Pi_1^*} = 0.16$ which is the same as the profits obtained when maximizing the profit given demand d_1. This is because, under the constructed information structure, when firm 1 receives message m she is indifferent between charging all effective valuations for product 1 of the consumers' types for which she receives message m and, in each group, there is a positive mass of consumers with effective valuation 0.4 which is equal to the optimal price p_1^* given demand d_1. Since $\underline{\Pi_1^*} = 0.16$ is a lower bound on firm 1's profit and the outcome is efficient, this information design is consumer-optimal.■

The consumer-optimal information structure that we have constructed in the above example can be generalized to arbitrary distributions over V^n. Key properties of this information design are: (i) consumers are partitioned according to their types into $\{E_1, E_2, \dots E_n\}$ where for each set of types E_i there is a dominant firm i; (ii) all consumers of types in E_i are further partitioned into groups with consumers in the same group assigned the same message; (iii) in each group there is a positive mass of consumers whose effective valuation equals p_i^* (the price that firm i charges when facing demand d_i); and (iv) if a group contains consumers of different effective valuations, then firm i is indifferent between charging any price equal to any of their effective valuations. Property (iv) implies that is incentive compatible for firm i to serve all consumers in each element of the partition of E_i, and property (iii) ensures that firm i makes the same profit as when pricing to all consumers in E_i according to demand d_i.

It is not obvious that conditions (ii)-(iv) can be met for every distribution of valuations. As it turns out, Bergemann, Brooks, and Morris (2015) establish that for an arbitrary distribution of valuations, it is always possible to partition up the mass of consumers in E_i into groups fulfilling these conditions.[20] For a group l, let d_i^l denote the distribution of effective valuation of consumers for i's product and let $\text{supp}(d_i^l)$ denote its support. We say that $\{d_i^1, d_i^2, \ldots, d_i^L\}$ is a partition of the consumers with effective valuation d_i if $\sum_{l=1}^{L} d_i^l(\hat{\theta}) = d_i(\hat{\theta})$ for all $\hat{\theta} \in [0,1]$. We obtain

Theorem 2. The consumer-optimal information structure takes the following form: For each $i \in \mathcal{N}$, consumers in E_i are partitioned into $\{d_i^1, d_i^2, \ldots, d_i^L\}$ where $L_i < \infty$. Each consumer in group $1 \leq l \leq L_i$ is assigned the same message. Furthermore,

$$\text{supp}(d_i^l) = \underset{p}{\text{argmax}}\, p \sum_{\theta \geq p} d_i^l(\theta)$$

and

$$p_i^* \in \text{supp}(d_i^l)$$

for some $p_i^* \in \text{argmax}_p\, p \sum_{\theta \geq p} d_i(\theta)$ which is optimal when firm i sets a uniform price facing the distribution of effective valuations d_i.

5 Efficient Information Structures

In the previous sections, we showed that whenever the condition in Theorem 1 is met, all surplus can be allocated to producers and the platform faces no efficiency trade-off when maximizing producer surplus. The same is true when the platform maximizes consumer surplus: although it is not possible to allocate all available surplus to consumers, allocating as much as possible still leads to an outcome in which all consumers buy their most preferred products.

While the consumer-optimal and producer-optimal outcomes are useful benchmarks, it is also informative to consider what other points on the efficient frontier an information designer can obtain.

A first insight is that providing full information to all firms about all consumers (full information design) always results in an efficient market outcome. Indeed, under full information there is an equilibrium in which each firm i sets a price 0 to all consumers in E_j for $j \neq i$, and charges each consumer in E_i her effective valuation. This is illustrated by point B in Figure 7. The consumer-optimal outcome, illustrated by point A in panel

[20] Bergemann, Brooks, and Morris (2015) call this partition uniform profit preserving extremal segmentation.

(a) of Figure 7, achieves a weakly higher proportion of consumer surplus. These are two points on the efficient frontier that can always be achieved. Hence, by assigning a fraction λ of consumers to the consumer-optimal information design and a fraction $1-\lambda$ to the full information design any outcome on the efficient frontier between points A and B can also be achieved. The next proposition gives tight conditions under which points A and B coincide exactly.

Proposition 2. The full information design is consumer-optimal if, and only if for all firms i, all consumers in E_i have the same effective valuation i.e., for each i and any pair $\boldsymbol{\theta}, \boldsymbol{\theta}' \in E_i$, $f(\boldsymbol{\theta}) > 0$, $f(\boldsymbol{\theta}') > 0$ implies $\theta_i - \max_{j\neq i} \theta_j = \theta'_i - \max_{j\neq i} \theta'_j$.

The proof of Proposition 2 is deferred to Appendix A.5. The condition requires that for all consumers in E_i, the differences between valuations for i and their second most favorite product must coincide. In the setting of V^2 considered in Example 1, this implies that for E_1, only a single 'diagonal' e.g., types $\{(0.4, 0.2), (0.6, 0.4), (0.8, 0.6)\}$, can have positive masses of consumers, and likewise for E_2. This is a restrictive condition that generically does not hold and we should therefore expect that the full-information design to be sub-optimal for consumers.

Figure 5: Efficient Information Structures

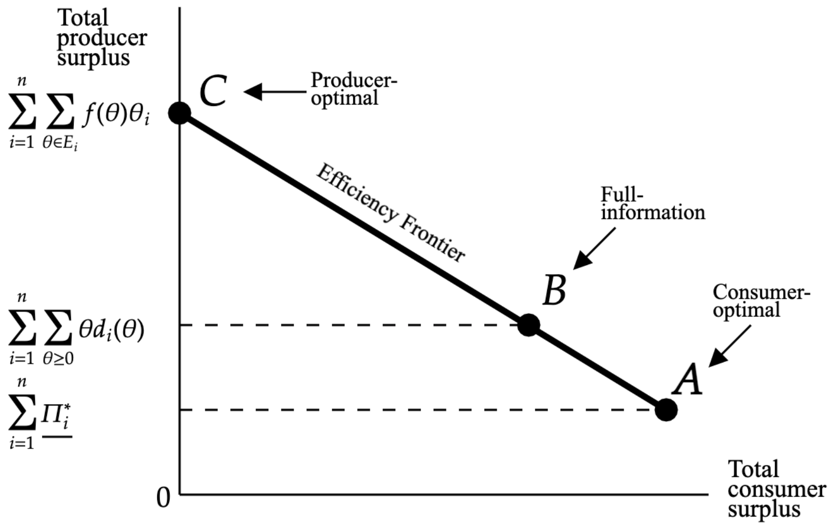

While it is always possible to reach any point in the efficient frontier between the full information and the consumer-optimal outcome, the producer-optimal outcome (point C in Figure 7) can be only achieved when the condition in Theorem 1 is met. When this condition holds, all points between point B and point C can also be obtained.

To see this, suppose we wish to obtain a point $D = \lambda B + (1-\lambda)C$ for some $\lambda \in (0, 1)$. We can partition the distribution of consumers f into $f_B(\boldsymbol{\theta}) = \lambda f(\boldsymbol{\theta})$ and $f_C(\boldsymbol{\theta}) = (1-\lambda)f(\boldsymbol{\theta})$ for each $\boldsymbol{\theta} \in \Theta$. We then apply the producer-optimal information design to f_C (which has mass λ) and the full information design to f_B (which has mass $1-\lambda$). Since the condition in Theorem 1 holds for f it also holds for f_C because this is simply a re-normalization of total mass.

The information design we constructed to obtain point D allocates each type randomly between the full information design and the producer-optimal information design. This construction works when the condition in Theorem 1 holds. It remains an open question to find necessary and sufficient conditions for the existence of an information design that obtains a given intermediate point between B and C.

As the condition in Theorem 1 is also sufficient for obtaining any point on the efficient frontier, the results of Proposition 1 apply and polarizing consumer tastes will continue to allow the designer to achieve all points on the frontier. If a given point on the efficient frontier cannot be achieved, then actions such as product differentiation and advertising that polarize consumer tastes might help.

6 Policy considerations

An intermediary may have incentives to induce, through information design, consumer surplus in downstream markets which is too low from the perspective of a regulator. In this section we use our framework to reflect on contemporary regulatory debates.

6.1 Creating flocks of consumers. In response to concerns over users' privacy, internet platforms are developing new technologies. A prominent example is Google's Privacy Sandbox which prevents information being conveyed to downstream firms about individual consumers. Instead, users are algorithmically grouped into flocks[21] based on the information the platform has collected about them. The platform then discloses only the cohort an individual user belongs to, which prevents her from being identified. An important question is whether these technologies will, in addition to preserving users' privacy, also enhance consumer surplus.

Information designs which implement both the producer-optimal outcome (Theorem 1) and the consumer-optimal outcome (Theorem 2) pool consumers into flocks and provide only this aggregated information to downstream firms. Hence, although such technologies may improve privacy, they need not constrain the platform from achieving the outcomes it desires in downstream markets. Indeed, an intermediary with a revenue model that just monetizes the surplus accruing to downstream firms may group consumers to soften competition in downstream markets, considerably reducing consumer surplus.

An important difference between the consumer-optimal and producer-optimal information structures is the types of consumers they each group together. In the consumer-optimal design, consumers with the same most preferred product are grouped together. By contrast, in the producer-optimal design consumers with different most preferred products are grouped together. Hence, regulators should pay attention to the principles that determine the ways in which privacy enhancing technologies pool consumers into flocks. Ideally, regulators should formulate legally binding rules of conduct that ensure that such groups are formed in line with the consumer-optimal information design.[22]

[21] Google terms this the Federated Learning of Online Cohorts (FLoC).

[22] The need of pro-active regime based on legally binding rules of conduct that shape the behaviour of powerful tech intermediaries is at the core of recent development among regulators, e.g., https://www.gov.uk/government/news/cma-advises-government-on-new-regulatory-regime-for-tech-giants.

6.2 Public information design. Rule of conducts on the way in which an intermediary should aggregate users' data may be difficult to enforce. Indeed, these groupings may be done by complex algorithms which are not open to scrutiny; at the same time, there is little transparency on the users' information that the intermediary accesses. A more direct intervention is to prescribe that, regardless of the way in which the intermediary constructs flocks, the aggregated information should be shared among firms. In terms of Google's flocks, for example, this implies restricting Google to put the same consumers into the same flocks for competing firms.

Note that to achieve the producer-optimal outcome the information designer might need to rely on private signals, whereas the consumer-optimal outcome can always be implemented with a public signal. Consider, for example, the producer-optimal information design illustrated in Figure 2. When firm 1 receives price recommendation 0.8, firm 1 knows that there is a mass of 0.2 consumers in E_1 with $v_1 = 0.8$, and a total mass of 0.2 consumers in E_2 among which a mass 0.1 has values $(\theta_1 = 0.4, \theta_2 = 0.8)$, a 0.05 has values $(\theta_1 = 0.4, \theta_2 = 0.8)$ and a mass 0.05 has values $(\theta_1 = 0.6, \theta_2 = 0.8)$. With this information firm 1 has incentives to follow the price recommendation. With private signals, the information designer can send firm 2 the message 0.8 corresponding to a mirrored group of consumers so that firm 2 similarly extracts all surplus from the consumers in E_2 with a valuation of 0.8 for product 2. Note here that the way consumers are grouped differs for the two firms.

If, instead, we constrain messages to be public, consumers must be grouped in the same way for firm 2 as they are for firm 1. If we take the group of consumers for which firm 1 receives the price recommendation of 0.8, and firm 2 also charges this group a price of 0.8, then the 0.1 mass of consumers in E_2 with value 0.6 for product 2 will not buy and, therefore, some potential producer surplus is lost. If firm 2 instead charges a price of 0.6 she will not extract all the surplus of the 0.1 mass of consumers in E_2 with value 0.8 for product 2.

To resolve this problem we must pool the consumers in E_1 with a value 0.8 for 1's product only with consumers in E_2 that all have the same value for 2's product (e.g., a value of 0.6 or 0.8, but not both). Suppose we try to create a group that pools all consumers in E_1 and E_2 with a valuation 0.8 for their most preferred product. Then both firms 1 and 2 will optimally follow the price recommendation of 0.8 and extract all surplus from those consumers. However, this implies that the intermediary will need to pool together all consumers in E_1 and E_2 with valuation 0.6 for their most preferred product and recommend a price of 0.6. If the two firms follow these price recommendations they each get a profit of 0.18. However, note that if firm 1 deviates and charges a price of 0.4 (or just under) it will steal a mass of 0.2 consumers from firm 2. Her profit will then increase from 0.18 to 0.2. In fact, one can check that in the distribution given in Example 1, the producer-optimal outcome in which all available surplus is allocated to producers cannot be achieved under public signals: in this case numerical calculations show that the producer-optimal design with private signals generates at least $\approx 7\%$ more producer surplus than can be achieved under any design implemented through public signals.

To better understand the more general implication of forcing signals to be public (i.e., mandating that the intermediary group consumers in the same way for different downstream firms), we consider the worse case scenario in which the intermediary wishes to maximise producer surplus and we compare the consumer surplus that is achieved when the information designer can use private signals and when it can only use public signals. We focus on the duopoly case and study pure strategy equilibria of the pricing subgame and provide two contributions.

First, we show that the problem of finding an information structure that maximises the producer surplus (even when the condition in Theorem 1 fails) can be formulated as a linear programming problem.[23]

Second, we apply this method to a canonical Hotelling duopoly model. Firm 1 is located at -1 and firm 2 is located at 1; there are six consumers' types uniformly distributed on $V := \{-1, -3/5, -1/5, 1/5, 3/5, 1\}$. Valuations are perfectly anti-correlated: the valuation of type $\theta \in V$ for product 1 is $1 - t(1 + \theta)$ and the valuation for product 2 is $1 - t(1 - \theta)$where $0 \leq t \leq 1/2$ controls the transportation cost. For each t we calculate the maximum producer surplus that can be extracted under no restriction on the set of information structures and under a regulation that only allows for public signals. Figure 6 plots the total producer surplus achievable under public and private signals alongside the corresponding total available surplus as transportation costs vary along the horizontal axis. We note that for both public and private optimal signals, the total surplus generated under the producer-optimal design typically coincides with the total available surplus, implying that these designs are typically efficient.[24] This implies that the corresponding consumer surplus under private and public signals is approximately the difference between the total available surplus and the maximum producer surplus.

When products are fairly homogeneous (e.g, small t), competition drives prices down to cost no matter how information is provided to firms. When product differentiation is high, consumers have strong preferences for their most preferred product. Regardless of the information firms have, the market will be segmented and firms will extract considerable surplus from their loyal consumers. It is when product differentiation is at an intermediate level that information design can shape market outcomes significantly. In particular, the restriction to public signals has a sizeable effect on the ability of the platform to appropriate consumer surplus. For example, for $t \approx 0.2$, the restriction to public signals assures a doubling of the consumer surplus. Note also that when moving from private to public signals the total gains from trade that are realised remain virtually unchanged. So the policy is effective in shifting surplus from producers to consumers without sacrificing aggregate efficiency.

Overall, this analysis suggests that regulations that constrain an intermediary's ability to provide information to firms privately may have a strong effect on limiting the power of platforms to extract consumer surplus without sacrificing aggregate efficiency. Moreover, although this analysis has focused on the scenario in which the platform seeks to maximize downstream producer surplus, it is worth emphasizing that if, instead, the platform sought to maximize downstream consumer surplus, the restriction to public signals would not inhibit their ability to do so: while the producer-optimal design we constructed relied on private signals, our consumer-optimal design can always be implemented through public ones.

[23] In fact this procedure holds for information structures that maximise any convex combination of consumer surplus and producer surplus. This formulation is useful as it allows us to numerically solve examples. The formal analysis is developed in Online Appendix B.3.

[24] When they do not coincide, the producer-optimal design under public and private signals achieves total surpluses of at least 99% of the total available surplus.

Figure 6: Producer surplus under the producer-optimal public and private signal

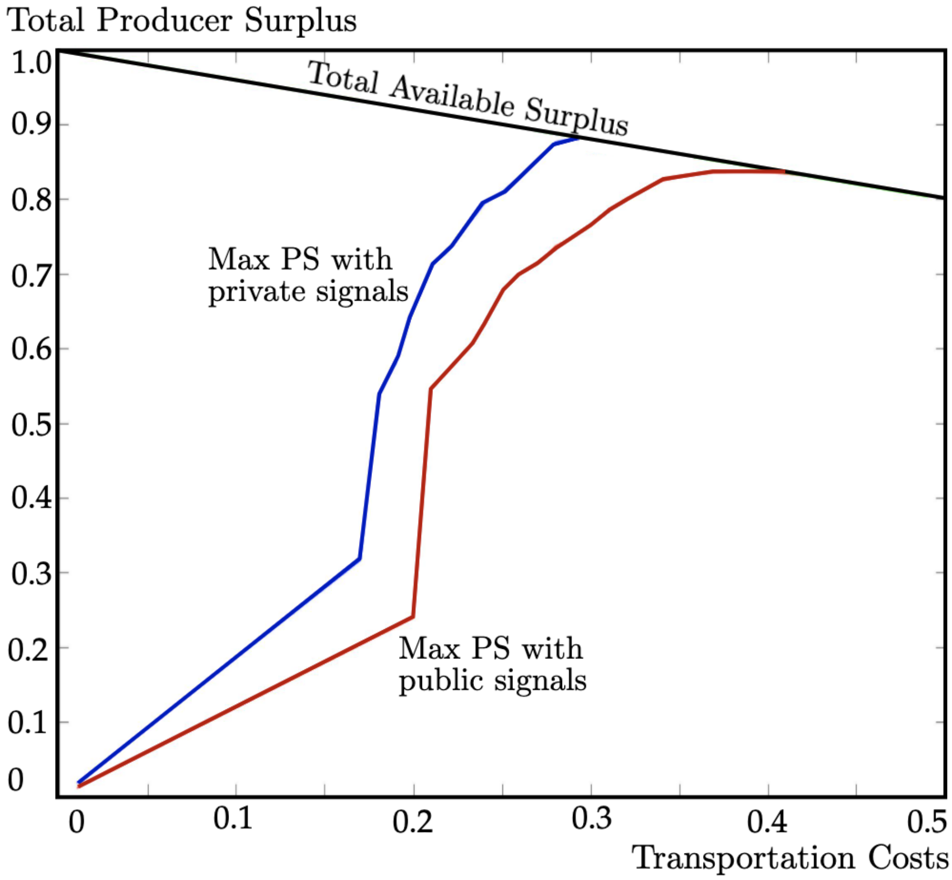

7 Conclusion

Internet platforms collect data about consumers and use that data to allow downstream firms to target advertisements at specific consumer groups. This raises new concerns for antitrust authorities. In this paper we provide a benchmark model to investigate and illustrate the power this information confers to data intermediaries in shaping market outcomes.

Through the lens of information design we consider an internet platform that seeks to maximize an increasing function of consumer surplus and producer surplus. How much influence can an internet platform, with perfect information about consumer preferences, wield over the ratio of producer versus consumer surplus obtained in downstream markets? We provide conditions under which the intermediary has absolute power—without sacrificing any surplus (so moving along the efficient frontier) the platform can choose any feasible ratio of consumer to producer surplus, including the perfectly collusive outcome under which consumer surplus is zero. Downstream firms can help the conditions be met by increasing the polarization of preferences through, for example, advertising and product design. Even when consumer preferences do not meet the required condition, the platform retains considerable power (many points on the frontier can be achieved).

From the perspective of an antitrust authority mandated with protecting consumer surplus, this raises a delicate problem. Outrightly preventing the use of information will typically sacrifice efficiency, and the platform might even use information to maximize consumer surplus, as in cases in which the platform prioritizes its reputation amongst users. At the same time, an intermediary with a revenue model based on monetizing consumer information can design an information structure that implements the perfectly collusive outcome in an otherwise highly competitive downstream market. Our analysis shows that the details of what information is provided to which firms matter for attendant market outcomes. The power of internet companies to extract consumer surplus can be curtailed (without restricting their ability to facilitate efficient trade) by forcing them to provide the same information about each group of consumers to all competing firms. This analysis allows us to speak to the modern debate on the regulation of data intermediaries: for instance, in the case of Google's recent flocks, the restriction to public signals would constrain Google to allocating the same consumers into the same flocks for all competing firms.

Our analysis also raises interesting questions for future work. While we consider a monopoly platform, in practice there are multiple internet companies, each with the ability to collect extensive proprietary data about consumers. This raises the prospect of competition among platforms. It would be interesting to study how this manifests, and what implications it has for the competitiveness of downstream markets. One possibility is that the internet companies compete for consumers and their information via their product offerings, while maintaining power over downstream market outcomes.

References

Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2019): "Too much data: Prices and inefficiencies in data markets," Tech. rep., National Bureau of Economic Research.

Admati, A. R. and P. Pfleiderer (1986): "A monopolistic market for information," *Journal of Economic Theory*, 39, 400–438.

Ali, S. N., G. Lewis, and S. Vasserman (2020): "Voluntary disclosure and personalized pricing," in *Proceedings of the 21st ACM Conference on Economics and Computation*, 537–538.

Armona, L., G. Lewis, and G. Zervas (2021): "Learning Product Characteristics and Consumer Preferences from Search Data," *Available at SSRN 3858377*.

Armstrong, M. and J. Vickers (2019): "Discriminating against captive customers," *American Economic Review: Insights*, 1, 257–72.

Armstrong, M. and J. Zhou (2019): "Consumer information and the limits to competition," .

Bergemann, D. and A. Bonatti (2015): "Selling cookies," *American Economic Journal: Microeconomics*, 7, 259–94.

——— (2019): "Markets for information: An introduction," *Annual Review of Economics*, 11, 85–107.

Bergemann, D., A. Bonatti, and T. Gan (2019): "The economics of social data," .

Bergemann, D., A. Bonatti, and A. Smolin (2018): "The design and price of information," *American economic review*, 108, 1–48.

Bergemann, D., B. Brooks, and S. Morris (2015): "The limits of price discrimination," *American Economic Review*, 105, 921–57.

Bergemann, D. and S. Morris (2013): "Robust predictions in games with incomplete information," *Econometrica*, 81, 1251–1308.

——— (2016): "Bayes correlated equilibrium and the comparison of information structures in games," *Theoretical Economics*, 11, 487–522.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile prices in market equilibrium," *Econometrica: Journal of the Econometric Society*, 841–890.

Bounie, D., A. Dubus, and P. Waelbroeck (2018): "Selling strategic information in digital competitive markets," .

——— (2020): "Market for Information and Selling Mechanisms," *Available at SSRN 3454193*.

Calvano, E., G. Calzolari, V. Denicolo, and S. Pastorello (2020): "Artificial intelligence, algorithmic pricing, and collusion," *American Economic Review*, 110, 3267–97.

Calzolari, G. and A. Pavan (2006): "On the optimality of privacy in sequential contracting," *Journal of Economic Theory*, 130, 168–204.

Cavallo, A. and R. Rigobon (2016): "The billion prices project: Using online prices for measurement and research," *Journal of Economic Perspectives*, 30, 151–78.

Competition & Markets Authority (2018): "Pricing algorithms: Economic working paper on the use of algorithms to facilitate collusion and personalised pricing," https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/746353/Algorithms_econ_report.pdf.

——— (2021): "Algorithms: How they can reduce competition and harm consumers," https://www.gov.uk/government/publications/algorithms-how-they-can-reduce-competition-and-harm-consumers/algorithms-how-they-can-reduce-competition-and-harm-consumers.

Competition Bureau Canada (2018): "Big data and innovation: key themes for competition policy in Canada," https://www.competitionbureau.gc.ca/eic/site/cb-bc.nsf/eng/04342.html.

Council of Economic Advisors (2015): "Big Data and Differential Pricing," https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/docs/Big_Data_Report_Nonembargo_v2.pdf.

Digital Competition Expert Panel (2019): "Unlocking digital competition," https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/785547/unlocking_digital_competition_furman_review_web.pdf.

European Comission (2019): "Competition policy for the digital era," https://ec.europa.eu/competition/publications/reports/kd0419345enn.pdf.

Fainmesser, I. P. and A. Galeotti (2019): "Pricing network effects: Competition," *Johns Hopkins Carey Business School Research Paper*.

Hannak, A., G. Soeller, D. Lazer, A. Mislove, and C. Wilson (2014): "Measuring price discrimination and steering on e-commerce web sites," in *Proceedings of the 2014 conference on internet measurement conference*, 305–318.

Johnson, J. P. and D. P. Myatt (2006): "On the simple economics of advertising, marketing, and product design," *American Economic Review*, 96, 756–784.

Jones, C. I. and C. Tonetti (2020): "Nonrivalry and the Economics of Data," *American Economic Review*, 110, 2819–58.

Kamenica, E. and M. Gentzkow (2011): "Bayesian persuasion," *American Eco-*

nomic Review, 101, 2590–2615.

KEHOE, P. J., B. J. LARSEN, AND E. PASTORINO (2018): "Dynamic competition in the era of big data," Tech. rep., Working paper, Stanford University and Federal Reserve Bank of Minneapolis.

KLEMPERER, P. (1995): "Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade," *The Review of Economic Studies*, 62, 515–539.

LIZZERI, A. (1999): "Information revelation and certification intermediaries," *The RAND Journal of Economics*, 214–231.

MONTES, R., W. SAND-ZANTMAN, AND T. VALLETTI (2019): "The value of personal information in online markets with endogenous privacy," *Management Science*, 65, 1342–1362.

NOVSHEK, W. AND H. SONNENSCHEIN (1982): "Fulfilled expectations Cournot duopoly with information acquisition and release," *The Bell Journal of Economics*, 214–218.

PIGOU, A. (1920): "The Economics of Welfare (London, 1920)," 12.

RAITH, M. (1996): "A general model of information sharing in oligopoly," *Journal of economic theory*, 71, 260–288.

ROBINSON, J. (1933): *The economics of imperfect competition*, London: Macmillan.

ROESLER, A.-K. AND B. SZENTES (2017): "Buyer-optimal learning and monopoly pricing," *American Economic Review*, 107, 2072–80.

STATE ADMINISTRATION FOR MARKET REGULATION (2021): "Anti-Monopoly Guidelines for the Platform Economy Sector," `http://gkml.samr.gov.cn/nsjg/fldj/202102/t20210207_325967.html`.

TAYLOR, C. R. (2004): "Consumer privacy and the market for customer information," *RAND Journal of Economics*, 631–650.

VIVES, X. (1988): "Aggregation of information in large Cournot markets," *Econometrica: Journal of the Econometric Society*, 851–876.

Appendices

A Omitted Results and Proofs

A.1 Proof of Lemma 1.

Proof. Since $\mathcal{M} \supset \mathcal{M}'$, there is nothing to prove for the if direction. It remains to show that there exists a producer optimal information structure under message space $\mathcal{M}$, then there also exists a producer optimal information structure under message space $\mathcal{M}'$. For any $\boldsymbol{\psi} \in \boldsymbol{\Psi}^*$ and for each i, first observe that P requires

$$\operatorname{supp}\big(\psi_i(\boldsymbol{\theta})\big) \cap \operatorname{supp}\big(\psi_i(\boldsymbol{\theta}')\big) = \emptyset \quad \text{for all } \boldsymbol{\theta}, \boldsymbol{\theta}^{'} \in E_i \text{ such that } \theta_i \neq \theta_i^{'}.$$

Further, P requires that upon receiving message $m_i \in \operatorname{supp}\big(\psi_i(\boldsymbol{\theta})\big)$ for $\boldsymbol{\theta} \in E_i$, firm i must find it optimal to charge price θ_i. P also requires that

$$\bigcup_{\boldsymbol{\theta} \in \Theta} \operatorname{supp}\big(\psi_i(\boldsymbol{\theta})\big) = \bigcup_{\boldsymbol{\theta} \in E_i} \operatorname{supp}\big(\psi_i(\boldsymbol{\theta})\big).$$

If this were not true, then must exist $\boldsymbol{\theta}' \in E_j, j \neq i$ such that

$$\operatorname{supp}\big(\psi_i(\boldsymbol{\theta}')\big) \setminus \bigcup_{\boldsymbol{\theta} \in E_i} \operatorname{supp}\big(\psi_i(\boldsymbol{\theta})\big) \neq \emptyset$$

and firm i must receive messages in this set for some strictly positive mass of consumers. Upon receipt of this message, firm i can infer that the associated consumers are not in E_i. Property P requires that firm i only sell to consumers in E_i. But then firm i has a strictly profitable deviation by charging a price $0 < \varepsilon < v_1$, a contradiction.

Now for each firm i, and each type $\boldsymbol{\theta} \in E_i$, relabel all messages in $\operatorname{supp}\big(\psi_i(\boldsymbol{\theta})\big)$ with the message $m_i = \theta_i$. By the above argument, all types $\boldsymbol{\theta} \in \Theta$ are then assigned some message in M_i', and this equilibrium continues to satisfy P.

□

A.2 Proof of Lemma 2.

Proof. We first show that if the information structure satisfies **Consumer IC**, **Firm IC**, **Consistency** and **Separation** then it is producer-optimal. Suppose there exists an information structure that satisfies **Consumer IC**, **Firm IC**, **Consistency** and **Separation**. Consider a strategy profile in which all firms set prices equal to the messages they receive. Given **Separation**, $p_i = m_i = \theta_i$ for all $\boldsymbol{\theta} \in E_i$. Given Consistency and **Consumer IC** a consumer type $\boldsymbol{\theta} \in E_i$ buys from firm i because that consumer will never receive a price for product j that is less than her value θ_j. Hence, the outcome of this pricing strategy is producer-optimal. To see that the pricing strategy is an equilibrium note that **Consumer IC** guarantees that it is unprofitable for a firm i to deviate and set a price above m_i, and **Firm IC** guarantees it is unprofitable for a firm i to deviate and set a price below m_i.

We now show that if an information structure supports an equilibrium which is producer-optimal then there must exist an information structure that satisfies **Consumer IC**, **Firm IC**, **Consistency** and **Separation**.

By Lemma 1, if a producer-optimal information structure exists, we can find an information structure $\boldsymbol{\psi}^*$ which only sends messages in $\mathcal{M}'$ and induces an equilibrium satisfying P. Further, as in the equilibrium constructed in the proof of Lemma 1, all firms charge the prices equal to the message they receive. We now show that such an information structure necessarily satisfies the above conditions.

First, suppose $\boldsymbol{\psi}^*$ violates **Separation**. Then for some firm i, there exists $\boldsymbol{\theta} \in E_i$ such that $\psi_i^*(m_i|\boldsymbol{\theta}) > 0$ for some $m_i \in M_i', m_i \neq \theta_i$. But since in equilibrium, firms charge prices equal to the messages they receive, some strictly positive proportion of type $\boldsymbol{\theta}$ do not buy from i at price θ_i which violates P.

Second, $\boldsymbol{\psi}^*$ only sends messages in $\mathcal{M}'$, Hence **Consistency** must hold.

Third, suppose $\boldsymbol{\psi}^*$ violates **Consumer IC**. Then there exists firms i, j, type $\boldsymbol{\theta} \in E_j$, and message $m_i \in M_i'$ such that $\mu_i(\boldsymbol{\theta}, m_i) > 0$ and $\theta_i \geq m_i$ i.e., although type $\boldsymbol{\theta} \in E_j$ has valuation for i's product larger than m_i, some positive mass of them are nonetheless assigned the message m_i. If $\theta_i > m_i$, then since firms must charge prices equal to the messages they receive, it violates P. If $\theta_i = m_i$, then since $\boldsymbol{\theta} \in E_j$, P requires that such consumers buy from firm j at price θ_j. But then a positive mass of such types are exactly indifferent between firms i and j so no matter how ties are broken, at least one firm does not sell to all such consumers—this firm has a strictly profitable deviation to a lower price, violating P.

Finally, suppose $\boldsymbol{\psi}^*$ violates **Firm IC**. In equilibrium firms charge prices equal to the messages they receive, and all surplus is extracted from all consumers. Hence, a violation of **Firm IC** implies at least one firm has a strictly profitable deviation to a lower price. But this delivers strictly positive surplus to some consumers, violating P.

□

A.3 Proof of Theorem 1.

Proof. When the condition in Theorem 1 is fulfilled, we show that we can find an information structure $\boldsymbol{\psi}$ which fulfils **Separation**, **Consumer IC**, **Firm IC**, and **Consistency**. By Lemma 2 this is sufficient.

Since messages are private, we proceed by first constructing the marginals $(\psi_i)_i$, then defining the joint $\psi(\boldsymbol{\theta}) := \Pi_i \psi_i(\boldsymbol{\theta})$.

For $\boldsymbol{\theta} \in E_i$, set

$$\psi_i(m_i|\boldsymbol{\theta}) := \begin{cases} 1 & \text{if } m_i = \theta_i \\ 0 & \text{otherwise} \end{cases}$$

and observe this fulfils **Separation**.

To facilitate constructing $\psi_i(\boldsymbol{\theta})$ for $\boldsymbol{\theta} \in \Theta \setminus E_i$, we introduce some additional notation. Recall that $M_i' := \{\theta_i : \boldsymbol{\theta} \in E_i\}$. This will be the messages received by firm i in our construction. For technical reasons it is helpful to also define $M_i'' := \{\theta_i : \boldsymbol{\theta} \in \Theta \setminus E_i\}$, the valuations for i's product held by *positive masses* of consumers of types not in E_i.

Observe that since the condition in Theorem 1 is fulfilled, for each firm i we have (i) $M_i' \neq \emptyset, M_i'' \neq \emptyset$; and (ii) $\max M_i'' < \max M_i'$.

The key to constructing the information structure is assigning messages to the types that firm i should not sell to. A starting point for this is given by the functions $G_i(\hat{\theta}_i, m_i)$ which measure that maximum mass of consumers not in E_i with valuations weakly above $\hat{\theta}_i$ that can be matched to the message m_i in an incentive compatible manner. Summing over these messages, we obtain our function $H_i(\hat{\theta}_i)$ which is the maximum mass of consumers not in E_i with valuations weakly greater than $\hat{\theta}_i$ which can be matched to some message in M_i' in an incentive compatible manner. We need to get from these "matching capacities" to the actual marginal distribution while preserving incentive compatibility. Moreover, as this neglects the valuations consumers have for other products, we need to pin this down and move from there to the full distribution over $\Theta \setminus E_i$.

Letting $|M_i''| = K_i$, it will be helpful to let $v_{l_i(k)}$ be the kth highest value in M_i''. Thus, $M_i'' = \{v_{l_i(1)}, \ldots, v_{l_i(K_i)}\}$. Notice that $v_{l_i(K_i)} < \max M_i' \leq v_K$.

Next, we define an adjustment to the G_i functions which will, by construction, correspond to the actual marginal distribution for values of i's product among consumers not in E_i.

$$G_i'(\hat{\theta}_i, m_i) := \begin{cases} \dfrac{G_i(\hat{\theta}_i, m_i)}{H_i(\hat{\theta}_i)} \displaystyle\sum_{\substack{\boldsymbol{\theta}' \in \Theta \setminus E_i: \\ \theta_i' \geq \hat{\theta}_i}} f(\boldsymbol{\theta}') & \text{if } \hat{\theta}_i < \max M_i'; \\ G_i(\hat{\theta}_i, m_i) & \text{otherwise.} \end{cases}$$

Further define

$$H_i'(\hat{\theta}_i) := \sum_{m_i \in M_i'} G_i'(\hat{\theta}_i, m_i) = \sum_{\substack{\boldsymbol{\theta}' \in \Theta \setminus E_i: \\ \theta_i' \geq \hat{\theta}_i}} f(\boldsymbol{\theta}'),$$

so this does indeed correspond to the marginal distribution of valuations for i's product. We now move to constructing the full distribution of messages.

For $\boldsymbol{\theta} \in \Theta \setminus E_i$, construct $\psi_i(\boldsymbol{\theta})$ as follows. Let $\theta_i = v_{l(k)}$ and set

$$\psi_i(m_i|\boldsymbol{\theta}) := \frac{G_i'(v_{l_i(k)}, m_i) - G_i'(v_{l_i(k+1)}, m_i)}{H_i'(v_{l_i(k)}) - H_i'(v_{l_i(k+1)})}$$

where we define $v_{l(K_i+1)} := v_K$. This is well-defined since by the definition of M_i'' there is a strictly positive mass of consumers in $\Theta \setminus E_i$ with value $v_{l_i(k)}$ for all k, the denominator is strictly positive.

Now we have constructed our message function, which satisfies **Separation**, so we just need to show that it aso satisfies the **Consistency**, **Consumer IC** and **Firm IC** conditions.

For $\boldsymbol{\theta} \in \Theta \setminus E_i$, observe

$$\sum_{m_i \in M_i'} \mu_i(\boldsymbol{\theta}, m_i) = \sum_{m_i \in M_i'} \psi_i(m_i|\boldsymbol{\theta}) f(\boldsymbol{\theta}) = f(\boldsymbol{\theta})$$

fulfilling **Consistency**.

Next, observe that

$$\mu_i(\boldsymbol{\theta}, m_i) = f(\boldsymbol{\theta})\frac{G_i'(v_{l_i(k)}, m_i) - G_i'(v_{l_i(k+1)}, m_i)}{H_i'(v_{l_i(k)}) - H_i'(v_{l_i(k+1)})} = 0 \quad \text{whenever } \theta_i = v_{l_i(k)} \geq m_i$$

since $G_i'(\theta_i, m_i) \leq G_i(\theta_i, m_i) = 0$ by the construction of G_i in the main text. Hence, this fulfils **Consumer IC**.

It remains to show **Firm IC** which requires that for all $m_i \in M_i'$,

$$m_i\Bigg(\sum_{\substack{\boldsymbol{\theta}' \in E_i: \\ \theta_i' = m_i}} f(\boldsymbol{\theta}')\Bigg) \geq \hat{\theta}_i\Bigg(\sum_{\substack{\boldsymbol{\theta}' \in E_i: \\ \theta_i' = m_i}} f(\boldsymbol{\theta}') + \sum_{\substack{\boldsymbol{\theta} \in \Theta \setminus E_i: \\ \theta_i \geq \hat{\theta}_i}} \mu_i(\boldsymbol{\theta}, m_i)\Bigg) \quad \text{for all } \hat{\theta}_i < m_i.$$

It is sufficient to check for deviations to $\{\hat{\theta}_i \in M_i'' : \hat{\theta}_i < m_i\}$ since deviating to any other price is dominated. To do so, fix $\hat{\theta}_i = v_{l_i(k)}, 1 \leq k \leq K_i$.

$$\begin{aligned}
\sum_{\boldsymbol{\theta} \in \Theta \setminus E_i : \theta_i \geq v_{l_i(k)}} \mu_i(\boldsymbol{\theta}, m_i) &= \sum_{k \leq t \leq K_i} \sum_{\boldsymbol{\theta} \in \Theta \setminus E_i : \theta_i = v_{l_i(t)}} \mu_i(\boldsymbol{\theta}, m_i) \\
&= \sum_{k \leq t \leq K_i} \left(\frac{G_i'(v_{l_i(t)}, m_i) - G_i'(v_{l_i(t+1)}, m_i)}{H_i'(v_{l_i(t)}) - H_i'((v_{l_i(t+1)})}\right) \sum_{\boldsymbol{\theta} \in \Theta \setminus E_i : \theta_i = v_{l_i(t)}} f(\boldsymbol{\theta}) \\
&= \sum_{k \leq t \leq K_i} \left(\frac{G_i'(v_{l_i(t)}, m_i) - G_i'(v_{l_i(t+1)}, m_i)}{H_i'(v_{l_i(t+1)}) - H_i'(v_{l_i(t+1)})}\right)\Big(H_i'(v_{l_i(t)}) - H_i'(v_{l_i(t+1)})\Big) \\
&= \sum_{k \leq t \leq K_i} \Big(G_i'(v_{l_i(t)}, m_i) - G_i'(v_{l_i(t+1)}, m_i)\Big) \\
&= G_i'(v_{l_i(k)}, m_i) \\
&\leq G_i(v_{l_i(k)}, m_i)
\end{aligned}$$

where the last equality is from the telescoping sum since $G_i'(v_{l_i(K_i+1)}, m_i) = 0$. But since G_i has been constructed so that **Firm IC** is tight everywhere, and we have shown that under our constructed information design, weakly fewer consumers with valuations above $\hat{\theta}$ are matched to message m_i, **Firm IC** holds.

We now show that if the condition in Theorem 1 is not fulfilled, $\boldsymbol{\Psi^*} = \emptyset$. Towards a contradiction suppose the condition is violated but $\boldsymbol{\Psi^*} \neq \emptyset$. By Lemma 2 there exists an information structure $\boldsymbol{\psi} \in \boldsymbol{\Psi^*}$ satisfying **Separation** with corresponding matchings $(\mu_i)_i$, satisfying **Consistency**, **Consumer IC** and **Firm IC**.

Given any such matching $(\mu_i)_i$, consider the distribution function it generates, $(\tilde{G}_i)_i$ which fulfils

$$\tilde{G}_i(\hat{\theta}_i, m_i) = \sum_{\substack{\boldsymbol{\theta} \in \Theta \setminus E_i: \\ \theta_i \geq \hat{\theta}_i}} \mu(\boldsymbol{\theta}, m_i) \quad \text{for all } m_i \in M_i',\, \hat{\theta}_i \in V.$$

As the condition in Theorem 1 does not hold, there must exist a firm i and alternative price $\tilde{\theta}_i$ that i can set such that

$$\sum_{\substack{\boldsymbol{\theta}' \in \Theta \setminus E_i: \\ \theta_i' \geq \tilde{\theta}_i}} f(\boldsymbol{\theta}') > H_i(\tilde{\theta}_i) = \sum_{m_i \in M_i'} G_i(\tilde{\theta}_i, m_i) \geq \sum_{m_i \in M_i'} \tilde{G}_i(\tilde{\theta}_i, m_i),$$

where the inequality follows from **Firm IC**. But by **Consistency**,

$$\sum_{\substack{\boldsymbol{\theta}'\in\Theta\setminus E_i:\\ \theta_i'\geq\tilde{\theta}_i}} f(\boldsymbol{\theta}') = \sum_{m_i\in M_i'} \tilde{G}_i(\tilde{\theta}_i, m_i),$$

a contradiction. □

A.4 Proof of Proposition 1.

Proof. Let H and $\tilde{H}$ be the corresponding functions defined in the main text for distribution f and $\tilde{f}$, respectively. Suppose $\boldsymbol{\Psi}^* \neq \emptyset$ under the distribution f and fix $\hat{\theta}_i \in (v_{k-1}, v_k]$ for some $1 \leq k \leq K$, where we set $v_0 = 0$ and $v_{K+1} = 1$. We have

$$\begin{aligned}
H_i(\hat{\theta}_i) &= \sum_{m_i\in M_i'} G_i(\hat{\theta}_i, m_i) \\
&= \sum_{\substack{m_i\in M_i':\\ m_i\geq\hat{\theta}_i}} \left(\frac{m_i-\hat{\theta}_i}{\hat{\theta}_i} \sum_{\substack{\boldsymbol{\theta}'\in E_i:\\ \theta_i'=m_i}} f(\boldsymbol{\theta}')\right) \\
&= \sum_{K\geq l\geq k} \left(\frac{v_l-\hat{\theta}_i}{\hat{\theta}_i} \sum_{\substack{\boldsymbol{\theta}'\in E_i:\\ \theta_i'=v_l}} f(\boldsymbol{\theta}')\right) \\
&= \sum_{K\geq l\geq k} \left(\frac{v_l-\hat{\theta}_i}{\hat{\theta}_i}\right)\Bigg(\underbrace{\sum_{\substack{\boldsymbol{\theta}'\in E_i:\\ \theta_i'\geq v_l}} f(\boldsymbol{\theta}')}_{:=a_l} - \underbrace{\sum_{\substack{\boldsymbol{\theta}'\in E_i:\\ \theta_i'\geq v_{l+1}}} f(\boldsymbol{\theta}')}_{:=a_{l+1}}\Bigg) \\
&= \frac{1}{\hat{\theta}_i}\sum_{K\geq l\geq k} v_l(a_l-a_{l+1}) - \sum_{K\geq l\geq k}(a_l-a_{l+1}) \\
&= \left(\frac{v_k}{\hat{\theta}_i}-1\right)a_k + \frac{1}{\hat{\theta}_i}\sum_{K\geq l\geq k}(v_{l+1}-v_l)a_{l+1} \\
&\leq \left(\frac{v_k}{\hat{\theta}_i}-1\right)b_k + \frac{1}{\hat{\theta}_i}\sum_{K\geq l\geq k}(v_{l+1}-v_l)b_{l+1} = \tilde{H}_i(\hat{\theta}_i)
\end{aligned}$$

where here

$$b_k := \sum_{\substack{\boldsymbol{\theta}'\in E_i:\\ \theta_i'\geq v_k}} \tilde{f}(\boldsymbol{\theta}') \geq \sum_{\substack{\boldsymbol{\theta}'\in E_i:\\ \theta_i'\geq v_k}} f(\boldsymbol{\theta}') := a_k$$

for all $1 \leq k \leq K+1$ by condition (i) of Proposition 1. The inequality follows because (i) $v_k/\hat{\theta}_i - 1 \geq 0$; (ii) $v_{l+1} - v_l > 0$; and (iii) $b_k \geq a_k$. The last equality follows from the same expansion of H_i, but replacing a with b.

Furthermore, condition (ii) of Proposition 1 implies that

$$\sum_{\substack{\boldsymbol{\theta}'\in\Theta\setminus E_i:\\ \theta_i'\geq\hat{\theta}_i}} f(\boldsymbol{\theta}') = \sum_{j\neq i}\sum_{\substack{\boldsymbol{\theta}'\in E_j:\\ \theta_i'\geq\hat{\theta}_i}} f(\boldsymbol{\theta}') \geq \sum_{j\neq i}\sum_{\substack{\boldsymbol{\theta}'\in E_j:\\ \theta_i'\geq\hat{\theta}_i}} \tilde{f}(\boldsymbol{\theta}') = \sum_{\substack{\boldsymbol{\theta}'\in\Theta\setminus E_i:\\ \theta_i'\geq\hat{\theta}_i}} \tilde{f}(\boldsymbol{\theta}').$$

Hence,

$$\tilde{H}_i(\hat{\theta}_i) \geq \sum_{\substack{\boldsymbol{\theta}' \in \Theta \setminus E_i: \\ \theta'_i \geq \hat{\theta}_i}} \tilde{f}(\boldsymbol{\theta}') \quad \text{for all } \hat{\theta}_i \in (0, v_K]$$

and so, by Theorem 1, $\boldsymbol{\Psi^*} \neq \emptyset$ under $\tilde{f}$. □

A.5 Proof of Proposition 2.

Proof. If: If the condition in Proposition 2 holds, then this implies that for each firm i, d_i has singleton support, so full information is both efficient and yields profits $\underline{\Pi}^*$ for each firm which implies that it is consumer-optimal.

Only if: Suppose there exists some pair $\boldsymbol{\theta}, \boldsymbol{\theta}' \in E_i$ such that

$$a = \theta_i - \max_{j \neq i} \theta_j > \theta'_i - \max_{j \neq i} \theta'_j = b.$$

This implies that the support of d_i includes at least a and b. It will suffice to restrict our attention to these two effective valuations.

Under the full information design, firm i makes $ad_i(a) + bd_i(b)$ from these two points. Now consider a modification of the full information design which continues to give full information about types with effective valuation not in $\{a, b\}$. For effective valuations a, b, we now group all consumers with effective valuations equal to b, as well as mass $\varepsilon > 0$ of consumers with effective valuations equal to a together. Firm i continues to find it optimal to set a price equal to b for this group since

$$b(d_i(b) + \varepsilon) \geq a\varepsilon$$

for sufficiently small ε. We group the remaining mass of consumers with effective valuations equal to a in a separate group. Now firm i makes profits

$$b(d_i(b) + \varepsilon) + a(d_i(a) - \varepsilon)$$

from the types in E_i with effective valuations a, b, and profits from all other types remain unchanged. As such, it makes $\varepsilon(a - b) > 0$ less than under the full information design. But since this equilibrium is efficient, total consumer surplus is strictly higher than under full information. □