Strategic Release of Information in Platforms: Entry, Competition, and Welfare

Kostas Bimpikis Giacomo Mantegazza *

Abstract

Two-sided platforms play an important role in reducing frictions and facilitating trade, and in doing so they increasingly engage in collecting and processing data about supply and demand. This paper establishes that platforms have an incentive to strategically disclose (coarse) information about demand to the supply side as this can considerably boost their profits. However, this practice may also adversely affect the welfare of consumers. By optimally designing its information disclosure policy, a platform can influence the entry and pricing decisions of its potential suppliers. In general, it is optimal for the platform to disclose its information only partially to either "nudge" entry when it is a priori costly for suppliers to join or, conversely, discourage it when suppliers do not have access to attractive outside options. On the other hand, consumers may end up being worse off as they have access to fewer trading options and/or face higher prices compared to when the platform refrains from sharing any demand information to its potential suppliers.

Keywords: Two-sided platforms, information design, competition, consumers welfare

1 Introduction

Two-sided platforms are increasingly becoming ubiquitous in everyday life as they have already transformed a number of industries ranging from short-term accommodation, e.g., Airbnb and Vrbo, to transportation, e.g., Uber and Lyft, freelance service provision, e.g., Upwork, and e-commerce, e.g., eBay and Etsy. Typically, the role of such platforms involves providing the infrastructure, i.e., an online marketplace, and a number of decision support tools to facilitate the exchange of goods or services between two sets of agents, the "buyers" and "sellers", who would otherwise find it challenging to transact.

*Graduate School of Business, Stanford University. Email: {kostasb,giacomom}@stanford.edu.
We are grateful to Kimon Drakopoulos, Avi Goldfarb, Yoni Gur, Bar Light, Haim Mendelson, and Ilan Morgenstern for very helpful suggestions.

Although the market environments in which platforms have found success are diverse, they share a number of key features. First, they are quite fragmented both in terms of the ownership of trade assets, e.g., accommodation for Airbnb and Vrbo or labor for Upwork, and in terms of information, i.e., market participants have only a limited view of market conditions. Second, they involve a great deal of uncertainty both in terms of the underlying demand (as in traditional marketplaces) but also in terms of the availability and willingness of the supply-side to participate, i.e., suppliers can freely choose whether, when, and how long to remain active on the platform depending on the terms and volume of trade they expect.

Thus, besides their primary function of providing a venue and the infrastructure for buyers and sellers to transact, platforms increasingly use their size, resources, and analytical capabilities to collect, process, and share information with market participants to best facilitate the match between supply and demand. Such information sharing takes a number of forms depending on the idiosyncratic features of the environment in which the platform operates. For example, Airbnb recently introduced a price-suggesting tool for hosts, which can be seen as an (indirect) way of providing them with information about local demand and supply conditions. Similarly, Guda and Subramanian (2019) report that Uber frequently shares information with its drivers about places where demand for rides may be high. Finally, Etsy, an online marketplace focused on handmade items and craft supplies, frequently shares "marketplace insights" and "seasonal tips" that can be viewed as demand forecasts for products sold on the platform.[1]

Despite their growing adoption in online marketplaces, the role of information provision tools on profits for the platform and welfare for consumers is not yet resoundingly clear. Conventional wisdom would suggest that more information typically leads to better outcomes for everyone involved. However, this is not necessarily as intuitive in the settings we focus on, given that information about demand has first-order implications on suppliers' entry and pricing decisions. This is precisely the goal of the present paper: we aim to shed light on (i) whether and how profit-maximizing platforms may choose to disclose information about the market environment, e.g., demand forecasts, and (ii) what the implications are of optimally disclosing such forecasts for profits and consumer welfare.

Towards this end, we develop a model of a two-sided platform that facilitates transactions between sellers and buyers of a homogeneous good. Besides acting as an intermediary between the two sides, the platform commits to and announces an information disclosure policy, i.e., a mapping from the information about demand it obtains, e.g., its demand forecast, to a set of messages that it sends to sellers. After committing to its policy, the platform observes whether demand will be "high" or "low" and sends a message to sellers, who, in turn, decide whether to join the platform and forgo their outside option. Sellers active on the platform compete in prices for buyers who have

[1] For example, refer to `https://www.etsy.com/seller-handbook/category/seasonal-tips`.

independent, private valuations for the good and arrive over two periods, i.e., a buyer arrives in the first period whereas a second buyer may arrive in the second period if the demand is "high". If a seller completes a transaction, she exits the platform, i.e., sellers are endowed with a single unit of the good. When setting their prices, sellers take into account both the likelihood that demand is high in the second period and that they may face less competition in case there is a transaction in the first one. Finally, the platform appropriates a fixed share of the value of each transaction that takes place and designs its disclosure policy to maximize its profits.

Upon receiving a message from the platform, potential sellers update their belief about the demand being high according to the platform's information disclosure policy. All else equal, the prices sellers set and subsequent profits are increasing in their beliefs about demand whereas, conversely, their profits decrease as the number of competitors that join the platform increases. Thus, in determining its disclosure policy, the platform takes into account the following two main (competing) drivers that affect equilibrium outcomes.

(i) *Modifying posterior beliefs to affect prices:* Since the price a seller sets for her product is increasing in her posterior belief, the platform has an incentive to induce high beliefs. On the other hand though, high prices imply a low probability of transactions taking place as consumers are price sensitive.

(ii) *Modifying posterior beliefs to affect entry*: Since, all else equal, the sellers' profits increase in their posterior beliefs, inducing higher beliefs makes joining the platform a more attractive option for sellers. However, the implications of having more active sellers on the platform are not obvious for the platform's profits, given that the likelihood of transactions taking place increases, but their respective prices decrease.

The interaction between information disclosure and the ensuing competition among sellers is quite complex: we establish that, in general, the platform finds it optimal to "nudge" or discourage entry depending on the prior and the value of the sellers' outside option; it achieves this by disclosing a coarser version of the information it has at its disposal. Optimal information disclosure considerably increases profits for the platform, but consumers may end up being significantly worse off as a result, as they face higher prices relative to the case when the platform does not disclose any information. For example, we establish that when potential sellers do not have access to attractive outside options, the optimal information disclosure policy leads to higher profits for the platform by disincentivizing entry; however, this comes at the expense of consumers who bear the brunt of the change in the competitive environment. Thus, we conclude that information disclosure is a powerful tool for a platform to increase its profitability, but its use may often be detrimental to consumers.

The rest of the paper is organized as follows. In Section 2 we present our benchmark model whereas in Section 3 we derive the equilibrium prices posted by sellers. Then, in Section 4 we provide a characterization of the optimal information disclosure policy and in Section 5 we discuss its welfare implications. Finally, in Section 6 we report on a number of extensions that support our main findings and we conclude in Section 7. All proofs can be found in the Electronic Companion.

1.1 Literature Review

The paper contributes to the emerging literature exploring applications of information design to settings of operational interest.[2] For example, Lingenbrink and Iyer (2019) determine the optimal way to disclose information about queue length to customers deciding whether to join a service system and Papanastasiou, Bimpikis, and Savva (2018) explore information provision by a service provider aiming to induce exploration of available alternatives by short-lived agents. Furthermore, Candogan and Drakopoulos (2020) extend the study of information design to social networks and Anunrojwong, Iyer, and Manshadi (2020b) discuss how the framework can be applied to improve the allocation of social services. Lingenbrink and Iyer (2020) and Drakopoulos, Jain, and Randhawa (2021) show that selective disclosure of inventory availability may lead to higher profits for a monopolist selling a fixed inventory of goods. Kostami (2020) focuses on the comparison between static and dynamic lead-time information provision in inventory systems whereas Hu, Wang, and Feng (2020) explore information disclosure and pricing policies for network goods. Finally, Küçükgül, Özer, and Wang (2019) show how to optimally design information disclosure in the context of time-locked sales campaigns whereas Alizamir, de Véricourt, and Wang (2020) and de Véricourt, Gurkan, and Wang (2020) explore the use of information design by public agencies on how to best release information to the public to mitigate potential disasters.

Closer to our work is the recent contribution by Bimpikis, Papanastasiou, and Zhang (2020), which explores information design in a two-sided service platform that connects quality-differentiated sellers with customers who are heterogeneous in how they value the quality of service. They establish that the platform may find it optimal to understate the quality of its best providers as a way to boost its volume of transactions and increase its profits. Our modeling framework differs significantly from theirs: first, in our setting, information about the size of the demand is conveyed to potential suppliers as opposed to information about sellers quality being provided to consumers and, second, in our setting, sellers are informed before they join the platform, which in turn affects their entry decisions, while in Bimpikis et al. (2020) they receive a quality label after entry. In the

[2] For the theory of information design, refer to contributions by Kamenica and Gentzkow (2011) and Rayo and Segal (2010) and, more recently, Arieli and Babichenko (2019), Anunrojwong, Iyer, and Lingenbrink (2020a), and Candogan and Strack (2021). Bimpikis and Papanastasiou (2019) and Candogan (2020) provide comprehensive surveys of work related to applications of information design to operations.

same spirit, Johari, Light, and Weintraub (2019) also consider a platform, which is better informed about the quality of its sellers and decides how to optimally communicate this information to buyers. They derive conditions under which barring a subset of sellers from joining and disclosing no information about the remaining ones to buyers is revenue-optimal. As also mentioned above, a key difference with our paper is that in Johari et al. (2019) the platform has information about the quality of potential sellers as opposed to information about future demand as in our case. Consequently, the mechanisms by which platforms increase their profits through information design are quite different from the ones we identify. Finally, their setting is not tailored to pursue a welfare analysis, which is instead a central piece of our work.

Finally, Gur, Macnamara, Morgenstern, and Saban (2022) study the problem of a platform that is informed about buyers' types and discloses such information to sellers. In their setting, only one seller performs Bayesian updates, while its competitors follow a given pricing policy. Our setting differs from theirs along multiple dimensions: in our case entry is a strategic decision and we focus on how the platform can use information disclosure to affect the level of competition it induces, while in Gur et al. (2022) the pool of sellers is given and the focus is on understanding the interplay between information disclosure and promotion policies that "spotlight" only one of the sellers. By and large, our work contributes in this line of work by highlighting how a platform can disclose its information on customer demand to induce the optimal level of entry and affect the intensity of competition among sellers. Importantly, besides illustrating that information design can lead to a significant increase in platform profits, we explore the welfare implications of this practice and find that, unlike the platform, consumers may be significantly worse off.

Our focus on consumer welfare also complements recent work that examines the value of information sharing in the specific context of ride-hailing platforms. In particular, Guda and Subramanian (2019) and Hu, Hu, and Zhu (2021) show that ride-hailing platforms can boost profits through communicating information about demand and appropriately pricing their service as drivers are incentivized to move to locations with excess demand. Besides mostly looking at platform profits, the aforementioned work does not adopt an information design perspective, i.e., demand signals are assumed to be fully disclosed.[3]

Finally, our paper also relates to the earlier and influential work on information sharing in the context of supply chain management. In particular, a set of seminal contributions, e.g., Lee, Padmanabhan, and Whang (1997), Lee and Whang (2000), and Cachon and Fisher (2000), identify the downsides of information distortion along the tiers of a supply chain and suggest how information

[3]We also contribute to the broader literature on the design and operations of two-sided platforms. Recent examples of this literature include Manshadi and Rodilitz (2020), Kanoria and Saban (2020), and Birge, Candogan, Chen, and Saban (2020), that study the optimal design of platforms with focus on efficient crowdsourcing of volunteer labor, matching, and optimal commission and subscriptions fees, respectively.

sharing, e.g., demand forecasts or sales data, may generate benefits resulting from lower inventory or shortage costs. On the other hand, Lee and Whang (2000) argue that one of the biggest challenges to supply chain wide information sharing may be confidentiality in the presence of competition. Motivated by this, Li (2002) among others explores both the direct and indirect effects, e.g., due to "leakage", of information sharing in a vertical supply chain that features competition and shows when downstream retailers may be discouraged from information transparency. Taken together, this early work illustrates that the implications of information sharing within a supply chain are nuanced and suggests that market participants may have the incentive to strategically determine when and what information to disclose. This is precisely the view we adopt in our study: we focus squarely on the disclosure of information as another strategic lever for a firm (a two-sided platform) and use tools from information design to derive optimal disclosure policies.[4]

2 Model

We consider a market for a homogeneous good, which can be sold from a set of *sellers* to *buyers* through an intermediary, the *platform*. The interaction between sellers, buyers, and the platform is modeled as a dynamic game with four time periods, $t \in \{-1, 0, 1, 2\}$. In particular, first, the platform commits to an *information disclosure* policy (at $t = -1$), which provides a mapping from the platform's information about demand (its *signal*) to a message to sellers.[5] Importantly, the platform determines and announces its information disclosure policy *before* it observes the demand signal. Next, at $t = 0$, the platform observes its signal and sends a message to sellers according to the disclosure policy. Subsequently, sellers decide whether they want to join and list their products on the platform. If they join, they forgo an opportunity cost, e.g., associated with selling their product through another channel, and they compete in prices. We assume that each seller can list at most one unit of their product on the platform, i.e., sellers have limited inventory.

Buyers arrive sequentially to the platform and decide whether to make a purchase based on the prices quoted by sellers. Each buyer wishes to acquire at most a single unit of the good. For most of the paper, we restrict attention to an economy with two sellers and two potential buyers arriving at period $t = 1$ and (possibly) at period $t = 2$. This setting allows for a tractable and transparent analysis—we discuss a general framework with a continuum of buyers and multiple sellers in Section 6.3, where we illustrate the robustness of our main qualitative insights.

Finally, the platform facilitates trade and generates its profits by appropriating a share $\alpha \in (0, 1)$

[4] At a higher level, our finding that strategic information disclosure may prevent entry on online markets also relates to the extensive literature in Industrial Organization that studies market concentration (e.g., Sutton (1991)). Specifically, the increasing relevance of two-sided platforms has sparked further research on the potential for anti-competitive effects in two-sided markets (e.g., Evans and Schmalensee (2005)).

[5] For example, the signal can be interpreted as a forecast about the size of the demand in future periods.

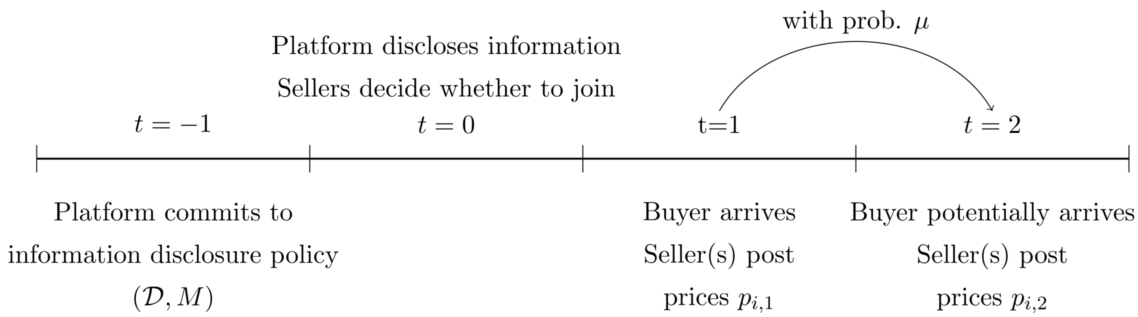

Figure 1: Timeline of the game

of the price of each transaction that takes place. We assume α is fixed and given (we revisit this in Section 6.2). This assumption is motivated by industry practice, where typically commission rates are not dynamically adjusted at the same rate at which new information about demand/supply is obtained.[6] In the remainder of this section, we provide a more detailed description of the interaction between sellers, buyers, and the platform, and formally introduce the equilibrium concept we adopt in our analysis. Figure 1 provides a comprehensive illustration of the timeline of the game.

To further facilitate the formal definition of the dynamic game among sellers, buyers, and the platform, we let $\mathcal{H}$ denote the set of all possible histories in the game and h_t indicate a generic history at time period t. For the first period, $t = -1$, we have $h_{-1} = \varnothing$. Furthermore, for every other t, a history h_t is a tuple containing the information disclosure policy, the message sent by the platform to the sellers, and all actions taken by time t. Finally, we use the notation σ_i^S and σ_i^B to denote the strategies of sellers and buyers, which are maps from the set of histories to the set of (distributions over) available actions.

Buyers Much of the paper focuses on a setting with two sellers and two potential buyers. In particular, demand for the good is initially unknown: although it is certain that a buyer will request the good in period $t = 1$, a second buyer may arrive in $t = 2$ with some probability. We can formally model this by introducing a set of states of the world $\Omega = \{0, 1\}$, where state $\omega = 1$ denotes that a second buyer will arrive in period $t = 2$. We assume that there exists a commonly held prior belief about the state of the world, i.e.,

$$\mathbb{P}(\omega = 1) = \mu \in [0, 1].$$

The arrival of buyers is exogenous, i.e., buyers cannot strategically time their arrival.

Buyers wish to acquire a single unit of the good for which they have a private valuation v. A

[6]Furthermore, platforms operate in multiple markets that may differ considerably in their demand and supply profiles. By and large, commissions/revenue shares are not tailored to each individual market.

buyer's valuation is independently drawn from a (known) distribution with cumulative distribution function F and support on (a, b), with $0 \leq a < b \leq \infty$. We make the following assumption.

Assumption 1. *The cumulative distribution function F of the buyers' valuation has a differentiable density f and satisfies the increasing failure rate property, i.e., function h defined below is increasing over the support of the distribution*

$$h(x) = \frac{f(x)}{1 - F(x)}.$$

We also introduce the notation $\overline{F}(x) = 1 - F(x)$, and let $g(x) = xh(x)$ denote the generalized failure rate introduced in Lariviere (2006).

Since the first buyer arrives at the platform at time $t = 1$ and the second may arrive at $t = 2$, we can index buyers with the label of the time period they arrive at. Consider buyer t, and let S_t and $(p_{i,t})_{i \in S_t}$ denote the set of active sellers in period t and their prices, respectively. The buyer's expected utility is given by the following expression,

$$U_t^B \left(\sigma_t^B; S_t, (p_{i,t})_{i \in S_t} \right) = \mathbb{E}_{B_t \sim \sigma_t^B} \left[\sum_{i \in S_t} \mathbf{1}\{B_t = i\}(v_t - p_{i,t}) \right], \tag{1}$$

where v_t is the buyer's private valuation. Here, B_t is associated with the purchase decision taken by buyer t, i.e., $B_t = i$ for some $i \in S_t$ implies that the buyer engaged in a transaction with seller i, whereas $B_t = 0$ denotes that buyer t decided not to purchase. Expectation is taken with respect to the possible randomization in buyer t's decision. Notice that not transacting with anyone gives utility zero. We interpret this as the normalized value of the outside option available to buyers. Thus, if the buyer does not purchase the good, she leaves the platform and takes her outside option. Finally, note that we are also assuming that buyers cannot decide to delay their purchase: in particular, either buyer 1 purchases in $t = 1$, or she does not purchase at all.

The optimal decision of a buyer can be summarized in a straightforward manner as follows: the buyer engages in a transaction with the seller who quotes the lowest price, provided that it does not exceed her valuation (she randomizes between sellers when they post the same price).[7] Specifically,

$$B_t^* = \begin{cases} 0 & \text{if } p_{i,t} > v_t \quad \text{for all } i \in S_t, \\ \arg\max\limits_{i \in S_t} \{v_t - p_{i,t}\} & \text{if there exists } i \in S_t \text{ such that } \; p_{i,t} \leq v_t \text{ and } p_{1,t} \neq p_{2,t}. \\ i \in S_t \text{ with probability } \frac{1}{2} & \text{if } p_{1,t} = p_{2,t} \leq v_t \end{cases} \tag{2}$$

[7] The assumption that if $v_t = p_{i,t}$ for some seller i, buyer t will make a purchase is without loss of generality, because under Assumption 1 this event has measure zero, and is consistent with the Sender-preferred equilibrium we consider in Definition 1. Note also that, as long as customer t buys from either seller with strictly positive probability when posted prices are equal and less than her valuation, our results continue to hold. To keep the analysis simple, we make the assumption that the buyer picks one of the sellers uniformly at random.

Since buyers do not take any other decision in the course of the game, in what follows, we abstract away from their decision making and we take it as given, i.e., we assume it is prescribed by Expression (2). In turn, this simplifies the exposition considerably and allows us to streamline the definition of the equilibrium.

Sellers Sellers are endowed with a single, indivisible good, which they can sell by joining the platform. They also have access to an outside option valued at $c > 0$, which they forgo if they join. When a seller transacts with a buyer, she exits the market. Otherwise, if a seller fails to make a sale in period $t = 1$, she remains active on the platform in period $t = 2$, as well. They make forward-looking pricing decisions, and cannot commit *ex-ante* to a sequence of prices. Finally, we assume that a seller can observe whether her competitor decided to join the platform (which, in turn, may affect her pricing strategy) and whether her competitor made a sale in time period $t = 1$ and, thus, exited the market.

As far as sellers' payoffs are concerned, we define them using backwards induction. We assume that sellers set their prices at time period $t = 2$ before they observe the realization of the period's demand, i.e., whether a second buyer arrived at the platform. In particular, seller i's expected continuation payoff is given by

$$\Pi^S_{2,i}\left(\sigma^S_i; h_2\right) = (1-\alpha)\mathbb{E}_\rho\left[\mathbb{E}_{p_{i,2}\sim\sigma^S_i,\, B^*_2}\left[p_{i,2}\mathbf{1}\{B^*_2 = i\}\mathbf{1}\{\omega = 1\} \mid \omega\right]\right], \tag{3}$$

where we recall that α denotes the share of the transaction fee that the platform appropriates. In principle, a seller can randomize its price although in equilibrium this is never the case. Here, the event $\{\omega = 1\}$ denotes the state of the world where a second buyer arrives at $t = 2$ and ρ is the seller's belief about ω given her information set (we discuss this at greater length in what follows).

Next, consider a history h_1 at time period $t = 1$. Seller i's expected continuation payoff is given by the following expression

$$\Pi^S_{1,i}\left(\sigma^S_i; h_1\right) = \mathbb{E}_{p_{i,1}\sim\sigma^S_i,\, B^*_1}\left[(1-\alpha)p_{i,1}\mathbf{1}\{B^*_1 = i\} + \Pi^S_{2,i}\left(\sigma^S_i, \langle h_1, B^*_1\rangle\right)\mathbf{1}\{B^*_1 \neq i\}\right], \tag{4}$$

given the arrival of a buyer in time period $= 1$. Finally, consider time period $t = 0$, when the seller decides whether to join the platform based on the platform's information disclosure policy. We can state seller i's *ex-ante* payoff as

$$\Pi^S_{0,i}\left(\sigma^S_i; h_0\right) = \mathbb{E}_{E_i\sim\sigma^S_i}\left[\mathbf{1}\{E_i = 1\}\Pi^S_{1,i}\left(\sigma^S_i, \langle h_0, E_i = 1, E_j\rangle\right) + \mathbf{1}\{E_i = 0\}c\right], \tag{5}$$

where E_i denotes the decision to join the platform. Note that we explicitly incorporated the opponent's entry decision to determine the history at time period $t = 1$ for the sake of clarity. In summary, a seller's strategy is a function

$$\sigma_i^S : h_t \in \mathcal{H} \mapsto \Delta\left(\mathcal{S}_A(h_t)\right),$$

where $\mathcal{S}_A(\cdot)$ is the map giving the action space available to sellers at history h_t.

Platform and information disclosure As mentioned before, at $t = -1$ demand is unknown and platform and sellers share a common prior μ about its state. It is assumed that the platform is endowed with a signal about future demand, which is informative about state ω. This signal is realized at time $t = 0$. After observing the realization of the signal, the platform sends a "message" m to the sellers, which may be informative about the demand. The sellers observe the platform's message and decide whether to join and forgo their outside option.

Importantly, the platform's message to sellers is generated according to an *information disclosure* policy, which the platform commits to and communicates to the sellers before it observes the signal's realization. For simplicity, we assume that the signal is perfectly informative, i.e., that it equals ω (we relax this in Section 6.1, further analyzed in Section EC 2 of the Electronic Companion). Formally, then, the information disclosure policy is a mapping from the set of states of the world to the set of (distributions over) potential messages,

$$\mathcal{D} : \Omega \to \Delta(M),$$

i.e., we allow the platform to follow a randomized messaging strategy. The platform designs its information disclosure policy in order to maximize its expected profits. Intuitively, communicating information about demand can be seen as a tool to adjust to market conditions relatively quickly compared, for example, to adjusting the platform's commission structure. After observing the platform's message, sellers update their beliefs about demand and determine whether to enter and how to price. Thus, when determining its policy, the platform has to take into account the equilibrium responses it induces from sellers.

Formally, the platform's strategy consists of a pair $\sigma^P = (\mathcal{D}, M)$. We pose no restrictions on the mapping $\mathcal{D}$ or the message space. Recall that the platform generates revenues by appropriating a share α of every transaction. Thus, its payoff can be written as

$$\Pi^P\left((\mathcal{D}, M); h_{-1}\right) = \frac{\alpha}{1-\alpha}\mathbb{E}_{\omega\sim\mu}\left[\mathbb{E}_{m\sim\mathcal{D}(\omega)}\left[\sum_{i=1}^{2}\mathbf{1}\{E_i = 1\}\Pi_{0,i}^S\left(\sigma_i^S; \langle h_{-1}, (\mathcal{D}, M), m\rangle\right) \mid \omega\right]\right]. \quad (6)$$

2.1 Equilibrium definition

Our choice of equilibrium concept is that of Perfect Bayesian Equilibrium (PBE). Following the literature on information design, we consider a refinement of PBE, i.e., a Sender-preferred PBE (refer to Kamenica and Gentzkow (2011) for additional details). Intuitively, a Sender-preferred Perfect Bayesian Equilibrium in our context is such that sellers always take the action that maximizes the platform's profits when they are indifferent between multiple actions.

Formally, a PBE consists of a collection of strategies and beliefs about the state of the world, one for each of the players of the game; we refer to them as strategy-belief pairs.[8] We introduce the notation $\gamma_k \mid h$ to indicate the posterior belief of agent k after history h, and we let $\gamma_k = (\gamma_k \mid h)_{h\in\mathcal{H}}$ denote the entire belief system of agent k. Recall that S denotes the set of sellers, and we let P denote the platform. We focus on equilibria in pure strategies although the definition remains valid for mixed strategies as well.

Definition 1 (Sender-preferred Perfect Bayesian Equilibrium)**.** *A collection of strategy-belief pairs $(\sigma_k, \gamma_k)_{k\in S\cup\{P\}}$ is a Sender-preferred PBE if the following conditions are satisfied:*

(1) For every seller $i \in S$, every time period $t \in \{0, 1, 2\}$, and every history $h_t \in \mathcal{H}$, for any feasible strategy $\sigma_i'^S$ of seller i we have

$$\Pi^S_{t,i}\left(\sigma^S_i; h_t\right) \geq \Pi^S_{t,i}\left(\sigma'^S_i; h_t\right).$$

(2) For every agent $k \in S \cup \{P\}$, $\gamma_k \mid \varnothing = \mu$, i.e., both sellers and the platform share a common prior μ. Moreover,

(i) For $i \in S$, $\gamma_i \mid h$ is determined by Bayes' rule after history $h \in \mathcal{H}$, whenever possible.

(ii) For the platform, $\gamma_P \mid h = \delta_{\{\omega\}}$ for all $h \neq \varnothing$, where ω is the realization of the state of the world.

(3) For any fixed $(\mathcal{D}, M)$, whenever there exist multiple collections of strategy-belief pairs for the sellers such that all of the previous conditions hold, then $(\sigma_i, \gamma_i)_{i\in S}$ yields the highest payoff for the platform.

(4) Finally, $(\mathcal{D}, M)$ is the information disclosure policy that maximizes the platform's profits

$$\Pi^P\left((\mathcal{D}, M); h_{-1}\right) \geq \Pi^P\left((\mathcal{D}, M)'; h_{-1}\right),$$

assuming that sellers follow the strategies prescribed by the equilibrium.

[8] These are commonly referred to as *assessments* in the literature.

Let us briefly parse through Definition 1: points (1) and (2) require sequential rationality and Bayes consistency of the strategy-belief pairs. In particular, sequential rationality demands that agent i's strategy σ_i is a best response to the strategies of the rest of the agents; point (3) selects the collection of strategy-belief pairs that gives the highest profit to the platform among all collections that are such that each strategy-belief pair is Bayes consistent and it is a best-response to the other agents' strategies. Thus, as defined, the equilibrium is "Sender-preferred" because sellers, when indifferent, follow the strategies that maximizes the platform's profits.

In the next section, we proceed by showing that an equilibrium exists. In principle, existence is not straightforward since the agents' payoff functions can be discontinuous. The proof is constructive and proceeds by backwards induction.

3 Equilibrium Analysis

As a first step in our analysis, we consider any subgame that results after sellers observe the platform's message and update their beliefs about demand according to the platform's information disclosure policy. In particular, for the discussion that follows, we assume that sellers share a common belief $\rho \in [0, 1]$ regarding ω and describe their equilibrium actions (entry and pricing decisions) as a function of ρ and the rest of the modeling primitives. In Section 4, we turn our attention to the platform's optimization problem, i.e., how to design the information disclosure policy to induce the set of sellers' beliefs that lead to the highest expected payoff for the platform.

3.1 Pricing games

Taking the entry decisions of sellers at time $t = 0$ as given, there are two cases to consider depending on whether one or both sellers decided to join the platform (the case when neither seller joins is straightforward as the platform does not generate any revenues).

Single entrant First, we consider the case that a single seller, i.e., seller i, decides to join the platform when the induced belief is equal to ρ. The seller's optimal pricing strategy can be summarized in the following proposition.

Proposition 1 (Single Entrant)**.** *For every history where only seller i joins the platform, there exist unique prices $\hat{p}^M(\rho)$ and p^M such that it is optimal to set $p_{i,1} = \hat{p}^M(\rho)$ and $p_{i,2} = p^M$. Moreover, $\hat{p}^M > p^M$ if and only if $\rho > 0$.*

Here, note that $\hat{p}^M(\rho)$ is a function of the belief ρ while p^M does not depend on ρ. This captures the fact that the seller sets a higher price for her good at $t = 1$, given the opportunity cost of selling the good at $t = 2$ if a second buyer arrives at the platform. Using Proposition 1, the expected optimal profit for a single entrant is given by

$$\Pi^S_{0,i}\left(E_i = 1, E_{-i} = 0; \rho\right) = W^M(\rho) = (1-\alpha)\left[\hat{p}^M(\rho)\overline{F}\left(\hat{p}^M(\rho)\right) + \rho\pi^M F\left(\hat{p}^M(\rho)\right)\right],$$

whereas the expected profits for the platform are equal to

$$\hat{V}^M(\rho) = \alpha\left[\hat{p}^M(\rho)\overline{F}\left(\hat{p}^M(\rho)\right) + \rho\pi^M F\left(\hat{p}^M(\rho)\right)\right],$$

where we introduced the notation $\pi^M = p^M\overline{F}\left(p^M\right)$; we refer to π^M and p^M as one-shot monopoly profit and price, respectively. We conclude the discussion on a single entrant by noting that her and the platform's profits are increasing and convex in the induced belief ρ.

Corollary 1. *The functions W^M and $\hat{V}^M$ are increasing and strictly convex in ρ. Furthermore, $\hat{p}^M$ is increasing in ρ.*

Intuitively, a seller who is monopolist has an incentive to price higher in $t = 1$ than the one-shot monopoly price p^M: the reason is that the possibility of a second buyer in $t = 2$ decreases the marginal cost of high prices in $t = 1$ by providing a "second chance" to sell the good. Similarly, profits are convex in the belief since as the probability of a second buyer increases, both the price at $t = 1$ and the likelihood that the good will be sold to the second customer increase.

Both sellers join Next, we consider the case when both sellers decide to enter the platform at belief ρ. In turn, we distinguish between when there was a sale at $t = 1$ (thus, there is a single seller at $t = 2$) or when there was no sale. In the latter case, taking seller j's price $p_{j,2}$ as given, seller i's continuation value at $t = 2$ is given by

$$\Pi^S_{2,i}(p_{i,2}; \rho) = \begin{cases} 0 & \text{if } p_{i,2} > p_{j,2}, \\ \rho(1-\alpha)p_{i,2}\overline{F}\left(p_{i,2}\right) & \text{if } p_{i,2} < p_{j,2}, \\ \frac{1}{2}\rho(1-\alpha)p_{i,2}\overline{F}\left(p_{i,2}\right) & \text{if } p_{i,2} = p_{j,2}. \end{cases} \tag{7}$$

This is a standard Bertrand game where sellers have zero marginal costs. Thus, the only equilibrium of this subgame is such that they both set price equal to zero and, thus, make zero profits (for example, see Tirole (1988)). On the other hand, if seller j engaged in a transaction at $t = 1$, then seller i sets its price at p^M at $t = 2$, as described above.

Next, we consider the interaction between the two sellers at time $t = 1$. The expected payoff of seller i as a function of j's price is given as

$$\Pi^S_{1,i}(p^1_i;\rho) = \begin{cases} \overline{F}\left(p_{j,1}\right)\left[(1-\alpha)\rho\pi^M\right] & \text{if } p_{i,1} > p_{j,1}, \\ (1-\alpha)p_{i,1}\overline{F}(p_{i,1}) & \text{if } p_{i,1} < p_{j,1}, \\ (1-\alpha)\left[\frac{1}{2}\overline{F}(p_{i,1})p_{i,1} + \frac{1}{2}\overline{F}(p_{j,1})\rho\pi^M\right] & \text{if } p_{i,1} = p_{j,1}. \end{cases} \tag{8}$$

If seller i is undercut, she expects to earn $\rho\pi^M$ from period $t = 2$, provided her opponent succeeds in selling, which happens with probability $\overline{F}\left(p_{j,1}\right)$; if instead she undercuts the opponent, her expected profit is $p_{i,1}\overline{F}(p_{i,1})$. When they post the same price, the expected profit is an average of the two. The proposition below summarizes equilibrium play when both sellers join the platform.

Proposition 2 (Both sellers join)**.** *At every history when both sellers join the platform, the unique equilibrium prices are given as* $p_{i,1} = p_{j,1} = \rho\pi^M$ *and*

$$p_{i,2} = \begin{cases} p^M & \textit{if } j \textit{ makes a sale at } t = 1, \\ 0 & \textit{otherwise.} \end{cases}$$

To obtain some intuition, note that both sellers expect to be able make a profit of $\rho\pi^M$ if they are undercut by the opponent. This is equivalent to having an alternative valued at $\rho\pi^M$, i.e, a positive marginal cost. It is then reasonable that price competition leads to equilibrium prices equal to $\rho\pi^M$. Substituting equilibrium prices into (5), we get that the expected profit in this case when the induced belief is ρ is given by

$$\Pi^S_{0,i}(E_i = 1, E_i = 1;\rho) = W^D(\rho) = (1-\alpha)\rho\pi^M\overline{F}\left(\rho\pi^M\right), \tag{9}$$

while the value to the platform of inducing the entry of two sellers at belief ρ is

$$\hat{V}^D(\rho) = 2\alpha\rho\pi^M\overline{F}\left(\rho\pi^M\right). \tag{10}$$

Contrary to the case of a single entrant, sellers' and the platform's profits are concave in the induced belief ρ.

Corollary 2. *The functions* W^D *and* $\hat{V}^D$ *as defined in* (9) *and* (10) *are increasing and concave in belief* ρ*.*

Intuitively, expected profits in duopoly are concave in the belief because, while a higher probability of a second customer increases the value of the alternative of each entrant if they are undercut, higher prices in $t = 1$ make it more likely that the first customer will not buy, so that in $t = 2$ the sellers will make zero profits. Thus, the greater profits due to higher prices in $t = 1$ are dampened by the prospect of competition in $t = 2$.

3.2 Entry decisions

Next, we turn our attention to the sellers' decision whether to join the platform and forgo their outside option. Naturally, this decision hinges upon their expectations about the competitive environment they will face on the platform, which in turn are shaped by the platform's information disclosure policy. In particular, it is straightforward to obtain that

$$\min_{\rho} W^M = W^M(0) = \pi^M > \pi^M \overline{F}\left(\pi^M\right) = W^D(1) = \max_{\rho} W^D(\rho).$$

In other words, the lowest possible expected profits for a seller when she faces no competition on the platform are strictly higher than the highest possible profits when both sellers join. We provide a characterization of the sellers' entry decisions under a condition, which we introduce in Assumption 2 below.

Assumption 2. *The cumulative distribution F of the buyers' valuation satisfies the following*

$$\overline{F}\left(\pi^M\right)\left(1 - g\left(\pi^M\right)\right) \geq \frac{F(p^M)}{2} \qquad \text{and} \qquad 2\overline{F}\left(\pi^M\right) \leq 1 + F(p^M). \tag{11}$$

Formally, this assumption introduces an additional "wedge" between the profit functions $W^M(\cdot)$ and $W^D(\cdot)$, which in turn makes it possible to compare the revenue of the platform in both competitive regimes. Together with the properties of the profit functions W^M and W^D derived in Section 3.1, it constitutes an important ingredient of our theoretical analysis; for this reason, we further discuss Assumption 2 at the end of this section to provide a justification for requiring it.

The next proposition highlights that there are essentially three regimes with regard to entry, depending on the size of the sellers' outside option and the belief about demand induced by the platform.

Proposition 3 (Entry equilibrium)**.** *Let ρ denote the sellers' common posterior belief about demand after they obtain the platform's message. Then, the following hold along the equilibrium path:*

(i) Suppose that $c > W^M(0)$. Then, there exists $\rho^M(c, \alpha)$ such that if $\rho \geq \rho^M(c, \alpha)$ a single seller joins the platform and, otherwise, there is no entry. In addition, when $c \geq W^M(1)$, we have $\rho^M(c, \alpha) \geq 1$, i.e., entry is too costly for sellers irrespective of their beliefs about demand.

(ii) Suppose that $W^D(1) \leq c \leq W^M(0)$. Then, a single seller joins irrespective of the platform's information disclosure policy.

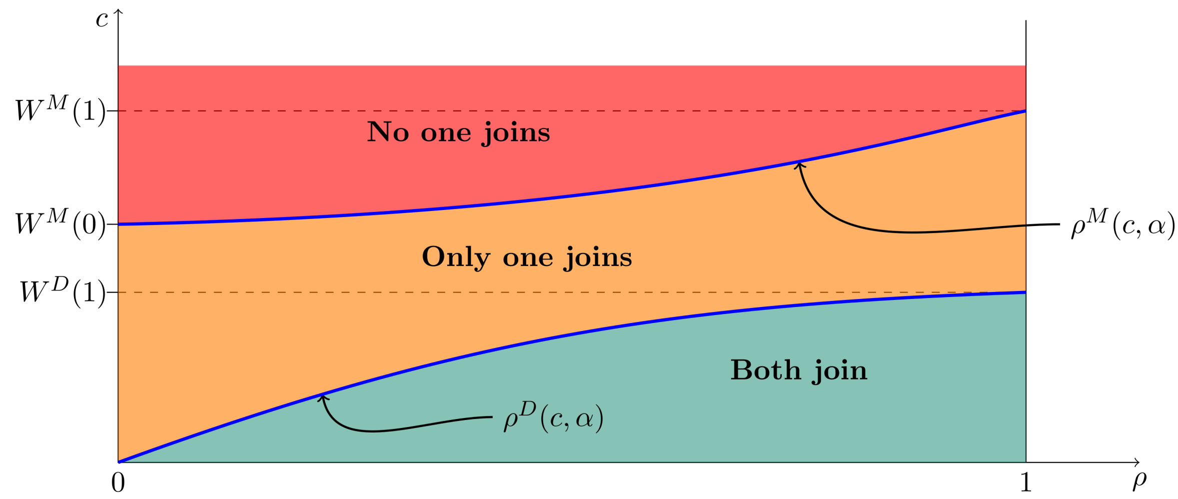

Figure 2: Equilibrium entry on the platform

(iii) Suppose that $c < W^D(1)$. Then, there exists $\rho^D(c, \alpha)$ such that if $\rho > \rho^D(c, \alpha)$ both sellers join the platform and, otherwise, there is a single entrant.

Figure 2 illustrates the regimes described in Proposition 3 and completes our equilibrium characterization. Essentially, since profits are increasing in the belief, for a fixed value of the outside option there exists a threshold belief such that joining the platform is more profitable than taking the outside option for all beliefs above the threshold. The thresholds ρ^M and ρ^D are defined as those beliefs such that $W^M(\rho^M) = c$ and $W^D(\rho^D) = c$.

Discussion of Assumption 2 Under Assumption 2, the Sender-preferred equilibrium whenever sellers are indifferent between joining and not features entry by a single buyer, as detailed in Proposition 3. Without the assumption, the pattern of entry in the Sender-preferred equilibrium at ρ^D might change as the rest of the parameters of the model vary. Moreover, as we shall see, Assumption 2 makes it possible to derive analytically the optimal disclosure policy for the platform; outside its scope, we must either resort to numerical analysis, or introduce other types of restrictions. In Appendix A we explore numerically the effect of relaxing Assumption 2 on the results of Sections 4 and 5: while the nature of the Sender-preferred equilibrium and the disclosure policies may be substantially different, the implications of information disclosure on the welfare of market participants remain (qualitatively) similar. Concretely, Assumption 2 restricts the focus on distributions of the willingness to pay that make buyers' sensitivity to prices neither too high nor too low. Distributions that satisfy the conditions include the exponential, uniform, and, in general, $Beta(a, b)$ and $Gamma(a, b)$ with $b \gg a$. We refer the reader to Appendix A for a more exhaustive robustness analysis, as well as for a more detailed interpretation in economic terms.

4 Optimal Information Disclosure

Armed with the equilibrium characterization from Section 3, we turn our attention to the problem of identifying the optimal information disclosure policy for the platform. Note that given an information disclosure policy and a prior belief μ, a message m sent by the platform to its potential sellers induces a posterior belief ρ on the demand being high. Furthermore, since the state of the world is random, the induced posterior is itself random, with a distribution τ that depends on both the message m and the information disclosure policy. Therefore, if ρ follows distribution τ, the platform attaches value

$$\Pi^P\left((\mathcal{D}, M); h_{-1}\right) = \mathbb{E}_{\rho\sim\tau}\left[\hat{V}(\rho)\right]$$

to employing policy $(\mathcal{D}, M)$.[9] Theorem 1 below establishes that the platform need only communicate two messages and the optimal mapping between the platform's information to its message to sellers is in most cases stochastic, i.e., the platform finds it optimal to disclose its information only partially.

Theorem 1 (Platform's optimal policy)**.** *Suppose that Assumptions 1 and 2 hold. Then, there exists an optimal policy $(\mathcal{D}^*, M^*)$ that involves sending one of two messages, i.e., $M^* = \{Y, N\}$. Moreover, if we let $\mathcal{D}^*(\omega)$ denote the probability that message Y is sent when the state of the world is ω under information disclosure policy $(\mathcal{D}^*, \{Y, N\})$, an optimal policy for the platform takes the following form:*

(i) When $c > W^M(0)$, then

$$\mathcal{D}^*(1) = \begin{cases} 1 & \text{for } \mu < \rho^M(c,\alpha) \\ q_u^M & \text{for } \mu \geq \rho^M(c,\alpha) \end{cases} \qquad \text{and} \qquad \mathcal{D}^*(0) = \begin{cases} q_l^M & \text{for } \mu < \rho^M(c,\alpha) \\ 0 & \text{for } \mu \geq \rho^M(c,\alpha) \end{cases},$$

where the disclosure probabilities are equal to

$$q_l^M = \frac{\mu(1-\rho^M(c,\alpha))}{\rho^M(c,\alpha)(1-\mu)} \quad \text{and} \quad q_u^M = \frac{\mu - \rho^M(c,\alpha)}{\mu(1-\rho^M(c,\alpha))};$$

[9] Although there can be multiple policies that induce the same distribution over beliefs, the expected revenues for the platform depend only on the posterior distribution τ. Thus, in what follows, we explicitly consider that the platform aims to optimize the posterior τ induced by its information disclosure policy, following the example of Kamenica and Gentzkow (2011).

(ii) When $W^D(1) \leq c \leq W^M(0)$*, full-disclosure is optimal, i.e.,*

$$\mathcal{D}^*(\omega) = \omega;$$

(iii) Finally, when $c < W^D(1)$*,*

$$\mathcal{D}^*(1) = \begin{cases} 1 & \textit{for } \mu \leq \rho^D(c,\alpha) \\ q_u^D & \textit{for } \mu > \rho^D(c,\alpha) \end{cases} \qquad \textit{and} \qquad \mathcal{D}^*(0) = \begin{cases} q_l^D & \textit{for } \mu \leq \rho^M(c,\alpha) \\ 0 & \textit{for } \mu > \rho^D(c,\alpha) \end{cases},$$

where the disclosure probabilities are equal to

$$q_l^D = \frac{\mu(1-\rho^D(c,\alpha))}{\rho^D(c,\alpha)(1-\mu)} \quad \textit{and} \quad q_u^D = \frac{\mu-\rho^D(c,\alpha)}{\mu(1-\rho^D(c,\alpha))}.$$

There are a number of points to note regarding the optimal information disclosure policy described in Theorem 1. First, although we did not restrict the space of possible messages, it turns out that there exists an optimal policy that involves sending just one of two messages, i.e., Y or N suggesting that the demand is high or low, respectively.[10] Sellers translate the message they receive using Bayes rule to a posterior belief about the demand being high according to mapping $\mathcal{D}^*(\omega)$. Interestingly, the platform finds it optimal to fully disclose its information *only when* the sellers' outside option takes intermediate values, i.e., when $W^D(1) \leq c \leq W^M(0)$. Otherwise, i.e., in cases (i) and (iii), the mapping is stochastic and depending on the prior and the state of the world ω, it may induce a posterior belief about the demand being high strictly in interval $(0,1)$.

Figure 3 represents the revenue function $\hat{V}$ and the optimal value function V^* obtained as the concave envelope of $\hat{V}$ for the cases of Theorem 1 where full-disclosure is not optimal. The threshold beliefs ρ^M and ρ^D correspond to posteriors that, when induced, affect the level of entry in the platform (from zero to one seller in the left panel and from one to two in the right one).

On a higher level, the platform's and the sellers' incentives are misaligned, i.e., the platform does not fully internalize the costs borne by sellers when they join and compete on the platform. As such, the platform uses its information disclosure policy as a lever to mainly affect the seller's

[10] On a technical level, Theorem 1 cannot be directly derived from Kamenica and Gentzkow (2011), because our setting involves multiple receivers who are engaged in a dynamic game. Our approach, inspired by Arieli and Babichenko (2019), involves defining an auxiliary game with a Sender and a single Receiver with three actions, whose payoffs correspond to the profits of the entry and pricing games of Section 3. However, these functions are non-linear in the belief about the state of the world, which prevents us from using revelation principle-type arguments to reduce the size of message space. Instead, we use an intermediate lemma from Anunrojwong et al. (2020a), which allows us to take as message space the interval $[0,1]$, and then retrieve the optimal value function as the concavification of $\hat{V}$, as in Kamenica and Gentzkow (2011) and Aumann and Maschler (1995). In sum, the finding that a binary message space suffices follows from the structure of the optimal policy and not directly from Kamenica and Gentzkow (2011).

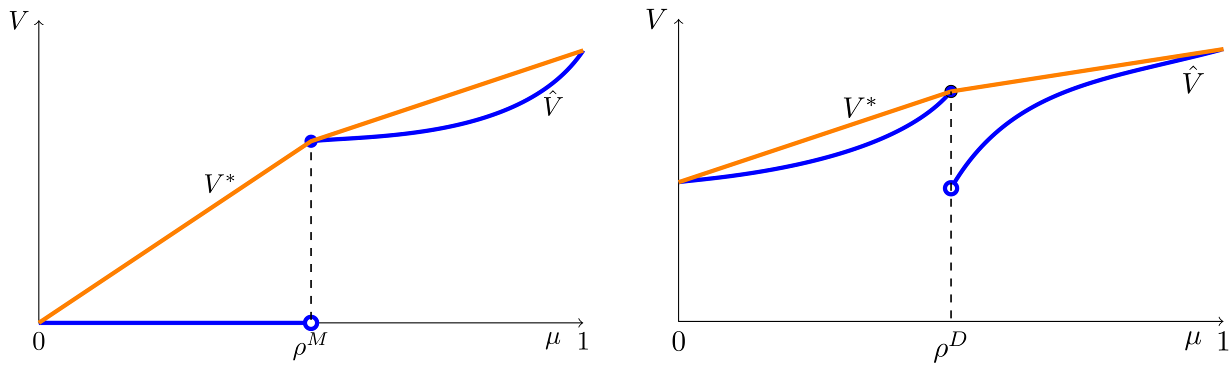

(a) High outside option: $c \geq W^M(0)$

(b) Low outside option: $c < W^D(1)$

Figure 3: Optimal value function V^* obtained as a concavification of $\hat{V}$.

entry decisions, i.e., nudge one of them to join when the prior would not justify entry, or preclude both from joining and driving prices down when the prior is sufficiently high. It is through affecting entry and the resulting competition among sellers that the platform increases its profits: if sellers had no outside option and found it always optimal to both join the platform irrespective of their beliefs on demand, information design would not result in any gains to the platform, i.e., it would have been optimal to simply employ no-disclosure.

4.1 Platform gains to information disclosure

We conclude this section by comparing the profits obtained under optimal disclosure with other benchmark policies the platform can employ. This comparison highlights that the benefits associated with strategically disclosing information can be substantial for platforms.

We compare the platform's profits under the optimal policy with two benchmarks: (i) no information, and (ii) full information disclosure. The former can be taken as representing a platform that does not engage in data collection and forecasting, but rather only facilitates trade, while the latter represents a platform that collects and discloses demand information, albeit not in a strategic fashion. Thus, we can interpret the relative profit gains associated with optimal disclosure as follows: profit gains relative to the no-disclosure benchmark represent the maximum return for the platform from collecting and processing demand-side data. On the other hand, additional gains relative to full-disclosure can be interpreted as the premium of optimizing information disclosure, i.e., the gains associated with *strategic* behavior from the platform.

Corollary 3. *The following hold true:*

(i) The optimal information disclosure policy always yields strictly higher profits for the platform

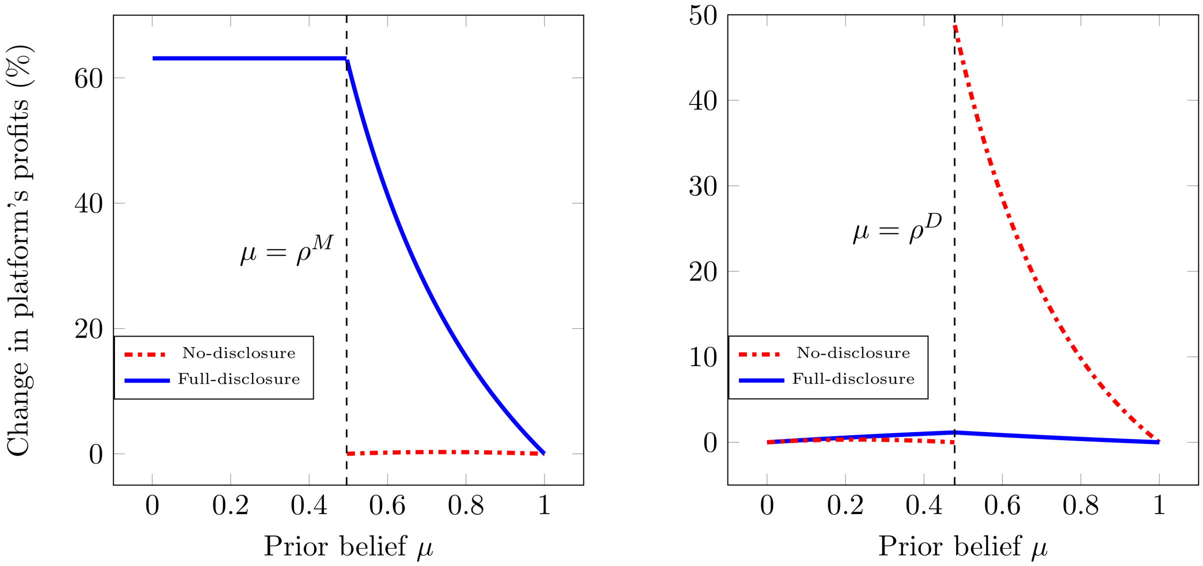

Figure 4: Percentage *increase* in the equilibrium profits for the platform from using the optimal policy instead of the full- and no-disclosure benchmarks. In the left panel, the profits under no-disclosure are zero for $\mu < \rho^M$, and therefore the relative increase from employing the optimal policy is infinite (hence, not shown). Results depicted were obtained with $v \sim U[0,1]$ and $\alpha = 5\%$, $c = 0.1$ for the left panel and $c = 0.3$ for the right panel.

than no-disclosure. In addition, it strictly outperforms full-disclosure too, unless $W^D(1) \leq c \leq W^M(0)$.

(ii) No-disclosure yields strictly higher profits than full-disclosure if (a) $c > W^M(0)$ *and* $\mu \geq \rho^M$*, or (b)* $c < W^D(1)$ *and* $\mu \leq \rho^D$.

Figure 4 illustrates the findings described in Corollary 3: the gains over both benchmarks can be significant. For example, the left panel shows an increase of more than 60% in profits from using optimal disclosure instead of full-disclosure, if the outside option for sellers is relatively high. Similarly, the right panel displays a 50% increase in profits by switching to optimal disclosure from no-disclosure, when the outside option is low. Thus, the simulation study reinforces the intuition we built from our analytical results: optimal information disclosure benefits the platform mostly by affecting the sellers' entry decisions. Finally, when the optimal policy does not affect entry, full-disclosure considerably lowers the platform's profits even compared to no-disclosure.

5 Welfare Effects

The platform's optimal disclosure policy increases profits by affecting the competitive structure of the market. As such, it also impacts the welfare of all market participants. In this section, we turn our attention from platform's profits to consumer surplus, sellers' profits, and aggregate welfare induced by the optimal policy. Before we proceed, we formally define expected consumer surplus and social welfare in our setting.

Definition 2 (Single seller)**.** *When a single seller joins the platform and the prior on the demand being high is μ, the expected social welfare is given by*

$$SW^M(\mu) = \int_{\hat{p}^M(\mu)}^{\infty} vdF + \underbrace{\mu F\left(\hat{p}^M(\mu)\right)}_{\substack{\text{Probability a second buyer arrives} \\ \text{and product is available}}} \int_{p^M}^{\infty} vdF \underbrace{+c}_{\text{Payoff of second seller not joining}}$$

and the consumer surplus is given by

$$CS^M(\mu) = \int_{\hat{p}^M(\mu)}^{\infty} \left(v - \hat{p}^M(\mu)\right) dF + \mu F\left(\hat{p}^M(\mu)\right) \int_{p^M}^{\infty} (v - p^M) dF.$$

Similarly, we have

Definition 3 (Two sellers)**.** *When both sellers join the platform and the prior on the demand being high is μ, the expected social welfare is*

$$SW^D(\mu) = \int_{\mu\pi^M}^{\infty} vdF + \mu\left(\overline{F}\left(\mu\pi^M\right) \int_{p^M}^{\infty} wdF + F\left(\mu\pi^M\right) \int_0^{\infty} wdF\right)$$

and the consumer surplus is given by

$$CS^D(\mu) = \int_{\mu\pi^M}^{\infty} \left(v - \mu\pi^M\right) dF + \mu\left(\overline{F}\left(\mu\pi^M\right) \int_{p^M}^{\infty} \left(w - p^M\right) dF + F\left(\mu\pi^M\right) \int_0^{\infty} wdF\right).$$

Finally, when none of the sellers joins the platform, the aggregate welfare is equal to $2c$, i.e., the sellers' opportunity costs.[11] The analytical results that follow compare the optimal policy with an uninformative one (no-disclosure) as the benchmark. Optimal information disclosure affects the welfare of market participants differently: the supply side of the market, i.e., sellers, are never worse-off by the optimal policy; on the other hand, when communication by the platform discourages entry, consumer surplus decreases, while, conversely, when the optimal policy encourages entry consumers are better-off. Finally, changes in social welfare always align with those to consumer surplus. The

[11]For the sake of simplicity, we assume that buyers' valuations are uniformly distributed on the unit interval.

following results formalize these observations.

First, we consider the case when joining the platform is relatively attractive for sellers, because the value of their outside option is low. Arguably, two-sided platforms are most successful in settings where market participants do not have access to high-value alternatives, i.e., they thrive in fragmented markets where there exist frictions to trade. As such, this represents the most natural setting for our welfare analysis. Moreover, it is also when the platform can gain the most by resorting to strategic information disclosure (as seen in Figure 4).

Theorem 2. *Suppose that $c < W^D(1)$. Then, if $\mu \leq \rho^D(c, \alpha)$, social welfare, consumer surplus and profits for the sellers increase under the optimal policy. On the other hand, if $\mu > \rho^D(c, \alpha)$, then profits for sellers increase; however, both consumer surplus and aggregate welfare decrease.*

We distinguish between two cases depending on the prior that the demand is high. In particular, when $\mu \leq \rho^D$ a single seller would find it optimal to join irrespective of the platform's message. In this case, the effect of information disclosure is to induce the seller to set different prices without affecting the competitive structure. Given that the expected profit of a monopolist seller is convex in the posterior belief, the optimal policy increases sellers' profits in expectation. We can show that consumer surplus and social welfare are also convex in the posterior belief; thus, they both increase under the optimal policy.[12] Intuitively, this follows since as the belief increases, price $\hat{p}^M(\mu)$ increases as well. Therefore, it is more likely that the product will not be sold in the first time period (as the quoted price is high) and, thus, it will be available for a potential second customer. The decrease in the first customer's surplus due to the higher price is compensated by the higher likelihood that the product will be sold at a lower price to the second customer.

The second case, when $\mu > \rho^D$, illustrates the adverse effects of a platform strategically disclosing the information at its disposal. In particular, the platform finds it optimal to "nudge" sellers *against* entry as a way of reducing competition in the market, increasing scarcity and, thus, resulting in higher prices. In particular, when the prior is relatively high, both sellers would join the platform in the benchmark case of no-disclosure. That is no longer the case under the platform's optimal policy: there is positive probability, i.e., $1 - \tau_u^D$, that only one of the sellers joins and consumers are worse-off as a result due to higher prices.[13] Moreover, even when the platform announces that a second customer will arrive (thus, inducing both sellers to join), sellers post higher prices than without any information disclosure. Hence in both cases consumers face higher prices as a result of the platform's strategic behavior.

[12] The convexity of the consumer surplus with respect to the belief holds under our assumption that the buyers' valuation for the product is uniform, but it also true for other commonly used distributions, e.g., Beta and Exponential.

[13] $1 - \tau_u^D = \mu(1 - q_u^D) + (1 - \mu)$ is the probability that a posterior belief is induced which prevents entry of one of the sellers.

Figure 5 provides an illustration of the magnitude of the welfare decrease under the conditions of Theorem 2. In particular, when the prior belief exceeds the ρ^D threshold, consumer surplus can be as much as 70% lower under the platform's optimal disclosure than under no-disclosure, and social welfare as a whole decreases by more than 15%.[14] At this point, it is worth emphasizing that the conditions under which Theorem 2 holds are not particularly restrictive: they just describe an environment where the presence of a two-sided platform facilitating trade would induce competition among potential sellers. In fact, in Section 6.3, we revisit this finding in the context of a considerably more involved environment with many potential sellers and buyers and confirm this intuition: a platform may strategically design its policy of disclosing demand information so as to create (artificial) scarcity of sellers and induce high prices to the detriment of consumers.

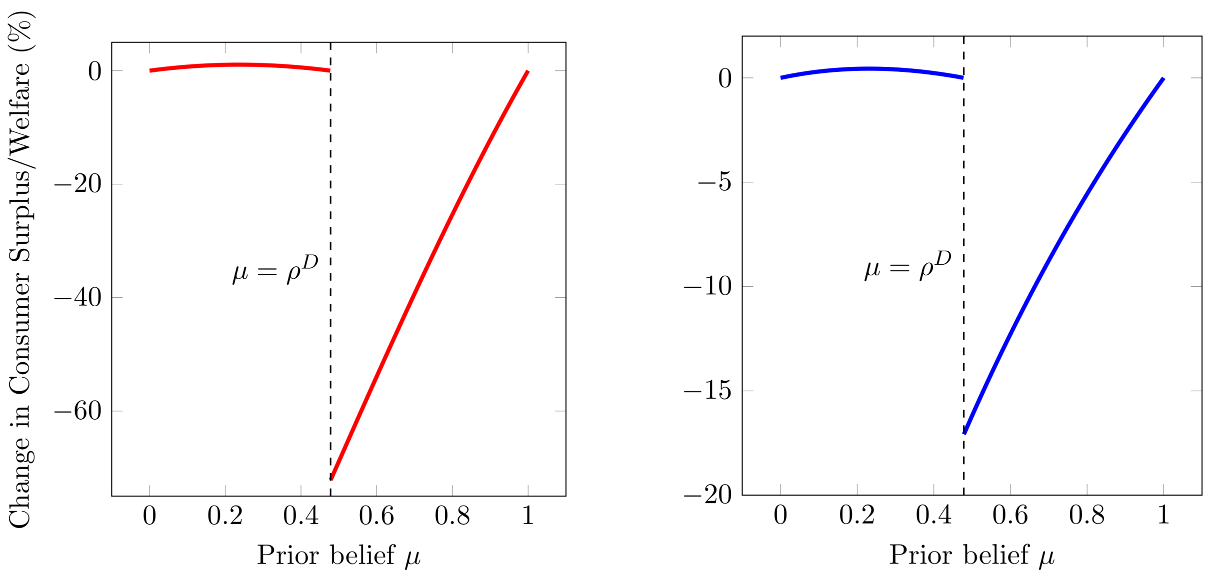

Figure 5: Percentage *decrease* in consumer surplus (left) and aggregate social welfare (right) due to the platform employing the optimal disclosure policy instead of no-disclosure. Results depicted were obtained with $v \sim U[0,1]$, $\alpha = 5\%$, and $c = 0.1$ (which satisfy the assumptions of Theorem 2). Aggregate profits for sellers are higher under the optimal disclosure policy (thus, they mitigate the losses in social welfare).

Next, we consider the case when the platform's strategic disclosure of information leads to higher social welfare.

Theorem 3. *When $c > W^M(0)$, both social welfare and consumer surplus increase under the optimal policy. In addition, when $\mu \geq \rho^M(c,\alpha)$, the profits for sellers increase, as well (otherwise, they remain unchanged).*

[14]The same set of simulations also shows that the profit of sellers can almost double under optimal disclosure.

Theorem 3 concerns the case when joining the platform is relatively costly for sellers, in that their outside option is sufficiently high. As before, we distinguish between two cases depending on the prior: in particular, when $\mu \geq \rho^M$ a single seller would find it optimal to join irrespective of the platform's message. In this case, we recover the same intuition as in Theorem 2 when $\mu \leq \rho^D$.

Instead, when $\mu < \rho^M$, no seller would have entered the market in the absence of any information from the platform. In this case, the role of information disclosure for the platform is to "nudge" a seller to join. Therefore, the resulting surplus is higher in expectation, given that a market is "created" for consumers. In sum, when joining the platform is relatively costly, the optimal policy induces a Pareto improvement over no-disclosure.

In a similar vein, the following proposition, which follows directly from Theorem 3, establishes that optimal information disclosure benefits all market participants when $W^D(1) \leq c \leq W^M(0)$, as well.

Proposition 4. *When $W^D(1) \leq c \leq W^M(0)$, social welfare, consumer surplus and profits for the sellers increase under the optimal policy.*

In summary, this section establishes that strategic information disclosure has nuanced effects on the welfare of market participants. While sellers tend to benefit, in the setting of most interest, i.e., when the value of the sellers' outside option is low, consumers end up being worse off.

6 Extensions

In this section, we briefly discuss a number of extensions to our basic model. For the sake of brevity, we summarize the main points for each extension, while we formalize the discussion in the Electronic Companion.

6.1 Noisy Signals

One of the simplifying assumptions of Section 2 is that the platform obtains a perfectly informative signal about the state of the demand. In the Electronic Companion, Section EC 2 we relax this assumption and study a setting where the signal received is only accurate with probability a, with $a < 1$.

The model we consider is identical to that of Section 2, with the difference that in this case the platform's information disclosure policy is conditional on the realization of the signal and not on the realization of the demand. In particular, this implies that the platform cannot induce any posterior belief using its policy, because its signal is accurate only with some probability. Formally, given a prior μ the range of possible posterior beliefs lies in an interval (i.e., $[\mu_{\min}(\mu, a), \mu_{\max}(\mu, a)] \subset [0, 1]$).

This introduces an additional constraint to the information design problem and makes it of a "non-standard" form: we characterize its solution as a function of the concave envelope of $\hat{V}$ that obtains when its domain is restricted to the range of admissible posteriors. At a high level, this characterization implies that, whenever the threshold beliefs ρ^M and ρ^D of Section 3.2 lie within the range of admissible posteriors, the optimal information disclosure policy either nudges or discourages entry as suggested by Theorem 1; in all other cases information disclosure does not affect entry.

Finally, we turn our attention on the marginal return for the platform of increasing the accuracy of its signal a, e.g., by investing in its data processing capabilities. In general, increasing a, i.e., making the signal more precise, simply amounts to expanding the range of posterior beliefs that the platform can induce. However, if by increasing a the range expands enough so that it includes (a previously excluded) threshold belief, then profits for the platform jump upwards: this confirms the intuition that the greatest benefit from information disclosure occurs exactly when platforms can use this tool to affect the composition of the market.

6.2 Revenue Share

When laying out the assumptions of the basic model, we argued that the platform may find it challenging to adjust its revenue share to quickly evolving market conditions or use different revenue shares in different sub-markets, e.g., corresponding to different geographies. These considerations provided justification behind our assumption that α is fixed and does not depend on other market primitives, e.g., the prior belief μ. In addition, they motivated us to focus our attention solely on how to design the platform's information disclosure policy. Here, we consider the case when the platform could also optimize over the revenue share (a more detailed discussion and analytical details are presented in Section EC 3 of the Electronic Companion).

In particular, we study a variant of the model in Section 2, where the platform optimizes over both the share of revenue it retains from each transaction and its information disclosure policy, with the goal of maximizing its profits. We establish that in the equilibrium induced by the platform's optimal revenue share α^*, at most a single seller joins the platform. The intuition behind this is that a single seller can generate higher aggregate revenues in expectation for the platform than a duopoly. In other words, the higher prices that can be sustained by a single seller benefit the platform more than the potential for more transactions in a duopoly. In addition, when the platform sets the revenue share optimally, there is no loss by employing either full- or no-disclosure. To an extent, strategically disclosing demand signals and optimizing over the revenue share are substitutes: when the revenue share cannot be tailored to the prevailing and potentially rapidly evolving market conditions, information disclosure can be employed to increase profits; if instead, the revenue share can be optimized for each set of primitives, optimal information disclosure takes a simpler form.

Finally, turning our attention to aggregate welfare, we establish that a platform optimizing its revenue share (and using the corresponding information disclosure policy) results in lower consumer surplus compared to a benchmark where the revenue share takes a fixed value.

6.3 Market with Many Sellers

Our basic model features two sellers and buyers. This allows us to characterize the optimal information disclosure policy as function of the market primitives and establish that obfuscation generates higher profits for the platform, but may lead to considerably lower surplus for consumers. In Section EC 4 of the Electronic Companion, we consider a significantly richer environment, where an arbitrary number of sellers can potentially join the platform. Furthermore, we allow product differentiation among sellers and model price competition by assuming that the demand for seller i is decreasing in the price it sets and increasing in the average price of the rest of the entrants, i.e.,

$$Q_i = \max\left\{0, \theta + (\phi - 1)P - \phi p_i\right\},$$

where $\theta > 0$, $P = \frac{1}{|S|}\sum_{j \in S} p_j$ is the average price of all participating sellers, and $\phi > 1$ denotes the degree of substitutability among products. Uncertainty about demand is captured by θ, which is assumed to be random. The platform observes the realization of θ (its signal about demand) and, as before, sends a message to sellers. To maintain tractability, we assume that the market clears in a single period, i.e., we do not consider a dynamic interaction among sellers.[15]

Although the complexity introduced by the fact that there exist multiple sellers precludes us from providing an analytical characterization of the optimal information disclosure policy, we show that there always exists a range of prior beliefs such that consumer surplus under the optimal policy is lower than under no-disclosure. This result confirms the main of our previous conclusions, i.e., strategically disclosing information may induce a considerable loss in consumer welfare. We complement this result with an example for which we derive the optimal disclosure policy numerically. The policy, which involves information obfuscation, generates higher profits for the platform and lower surplus for consumer by either inducing higher equilibrium prices (without altering the competitive structure compared with no-disclosure) or by discouraging entry. This mechanism ensures a profits growth of 30% for the platform, while incurring a loss of almost 50% of consumer surplus.

In summary, we show that the qualitative nature of our results holds in a significantly richer environment (that does not allow for deriving the optimal information disclosure policy). In par-

[15]This demand system can arise from the interaction of a continuum of buyers and has been micro-founded in Myatt and Wallace (2015).

ticular, we illustrate that the strategic disclosure of demand information yields significant benefits to a platform. Secondly, we highlight that the fact that consumers may be worse off under optimal disclosure policies is indeed a structural feature of the interaction, i.e., the platform generates higher profits in the expense of lower consumer surplus.

7 Concluding Remarks

This paper explores strategic information disclosure in an environment where a two-sided platform that has private information of future demand interacts with sellers who decide whether to join the platform and forgo their outside option. We demonstrate that the platform can increase its profits by appropriately disclosing its information on demand, potentially though at the expense of consumers. In particular, the platform uses its information strategically to change the sellers' beliefs about demand and affect their entry decisions to either "nudge" entry, when it is a priori costly for them to join, or discourage it, when sellers find it optimal to join the platform on the basis of their priors.

We also discuss a richer environment, which features multiple potential sellers, and confirm that our main results continue to hold qualitatively. In fact, we prove that consumers can be worse off if the platform engages in optimal information disclosure. We recover in an extensive numerical analysis that, when entry is relatively attractive for sellers, the platform may find it optimal to use its demand information strategically as a means of discouraging sellers to join the platform and, thus, induce higher prices. We also confirm that this intuition carries over to the case where the platform has only an imperfect knowledge of the level of demand. Finally, we argue that revenue shares/commissions and strategic information disclosure can be, to some extent, thought of as substitutes and that also optimizing commissions goes to the detriment of consumers.

We view our paper as contributing to the emerging and growing literature on platform design and, in particular, the recent line of work that focuses on the use of non-monetary levers to affect market outcomes. Apart from extending our model by relaxing the assumptions we made to ensure tractability, e.g., sellers are homogeneous and have the same outside option, we hope that our model and results will motivate more work on exploring the role of (non-monetary) platform levers on affecting the composition and size of both the demand- and the supply-side of the market, and the implications on welfare of the employment of such tools.

References

Alizamir, Saed, Francis de Véricourt, Shouqiang Wang. 2020. Warning against recurring risks: An information design approach. *Management Science* **66**(10) 4612–4629.

Anunrojwong, Jerry, Krishnamurthy Iyer, David Lingenbrink. 2020a. Persuading risk-conscious agents: A geometric approach. *Available at SSRN 3386273* .

Anunrojwong, Jerry, Krishnamurthy Iyer, Vahideh Manshadi. 2020b. Information design for congested social services: Optimal need-based persuasion. *arXiv preprint arXiv:2005.07253* .

Arieli, Itai, Yakov Babichenko. 2019. Private Bayesian persuasion. *Journal of Economic Theory* **182** 185–217. doi:10.1016/j.jet.2019.04.008.

Aumann, Robert J, Michael Maschler. 1995. *Repeated games with incomplete information*. MIT press.

Bimpikis, Kostas, Yiangos Papanastasiou. 2019. Inducing exploration in service platforms. *Sharing Economy*. Springer, 193–216.

Bimpikis, Kostas, Yiangos Papanastasiou, Wenchang Zhang. 2020. Information provision in two-sided platforms: Optimizing for supply. *Available at SSRN* .

Birge, John, Ozan Candogan, Hongfan Chen, Daniela Saban. 2020. Optimal commissions and subscriptions in networked markets. *Manufacturing & Service Operations Management* .

Cachon, Gerard P, Marshall Fisher. 2000. Supply chain inventory management and the value of shared information. *Management science* **46**(8) 1032–1048.

Candogan, Ozan. 2020. Information design in operations. *Pushing the Boundaries: Frontiers in Impactful OR/OM Research*. INFORMS, 176–201.

Candogan, Ozan, Kimon Drakopoulos. 2020. Optimal signaling of content accuracy: Engagement vs. misinformation. *Operations Research* **68**(2) 497–515.

Candogan, Ozan, Philipp Strack. 2021. Optimal disclosure of information to a privately informed receiver. *Working paper* .

de Véricourt, Francis, Huseyin Gurkan, Shouqiang Wang. 2020. Informing the public about a pandemic. *Forthcoming in Management Science* .

Drakopoulos, Kimon, Shobhit Jain, Ramandeep Randhawa. 2021. Persuading customers to buy early: The value of personalized information provisioning. *Management Science* **67**(2) 828–853.

Evans, David S, Richard Schmalensee. 2005. The industrial organization of markets with two-sided platforms. Working Paper 11603, National Bureau of Economic Research. doi:10.3386/w11603. URL http://www.nber.org/papers/w11603.

Guda, Harish, Upender Subramanian. 2019. Your uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives. *Management Science* **65**(5) 1995–2014.

Gur, Yonatan, Gregory Macnamara, Ilan Morgenstern, Daniela Saban. 2022. On the disclosure of promotion value in platforms with learning sellers. *Available at SSRN 3468674* .

Hu, Bin, Ming Hu, Han Zhu. 2021. Surge pricing and two-sided temporal responses in ride hailing. *Manufacturing & Service Operations Management* .

Hu, Ming, Zizhuo Wang, Yinbo Feng. 2020. Information disclosure and pricing policies for sales of network goods. *Operations Research* **68**(4) 1162–1177.

Johari, Ramesh, Bar Light, Gabriel Weintraub. 2019. Quality selection in two-sided markets: A constrained price discrimination approach. *arXiv preprint arXiv:1912.02251* .

Kamenica, Emir, Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* **101**(6) 2590–2615. doi:10.1257/aer.101.6.2590.

Kanoria, Yash, Daniela Saban. 2020. Facilitating the search for partners on matching platforms. *Management Science (to appear)* .

Kostami, Vasiliki. 2020. Price and lead time disclosure strategies in inventory systems. *Production and Operations Management* .

Küçükgül, Can, Özalp Özer, Shouqiang Wang. 2019. Engineering social learning: Information design of time-locked sales campaigns for online platforms. *Available at SSRN 3493744* .

Lariviere, Martin A. 2006. A note on probability distributions with increasing generalized failure rates. *Operations Research* **54**(3) 602–604. doi:10.1287/opre.1060.0282.

Lee, Hau L, Venkata Padmanabhan, Seungjin Whang. 1997. Information distortion in a supply chain: The bullwhip effect. *Management science* **43**(4) 546–558.

Lee, Hau L, Seungjin Whang. 2000. Information sharing in a supply chain. *International Journal of Manufacturing Technology and Management* **1**(1) 79–93.

Li, Lode. 2002. Information sharing in a supply chain with horizontal competition. *Management science* **48**(9) 1196–1212.

Lingenbrink, David, Krishnamurthy Iyer. 2019. Optimal signaling mechanisms in unobservable queues. *Operations research* **67**(5) 1397–1416.

Lingenbrink, David, Krishnamurthy Iyer. 2020. Signaling in online retail: Efficacy of public signals. *Available at SSRN: https://ssrn.com/abstract=3179262* .

Manshadi, Vahideh, Scott Rodilitz. 2020. Online policies for efficient volunteer crowdsourcing.

Myatt, David P, Chris Wallace. 2015. Cournot competition and the social value of information. *Journal of Economic Theory* **158** 466–506.

Papanastasiou, Yiangos, Kostas Bimpikis, Nicos Savva. 2018. Crowdsourcing exploration. *Management Science* **64**(4) 1727–1746.

Rayo, Luis, Ilya Segal. 2010. Optimal information disclosure. *Journal of Political Economy* **118**(5) 949–987.

Sutton, John. 1991. *Sunk costs and market structure: Price competition, advertising, and the evolution of concentration*. MIT press.

Tirole, Jean. 1988. *The Theory of Industrial Organization*. The MIT Press, Cambridge, Massachusetts.

Appendices

A Discussion of Assumption 2

In Section 3.2 we introduced Assumption 2, that restricts the set of distributions of the buyers' willingness to pay to which our results apply. In this appendix we clarify the need for the conditions we assume to hold and provide some intuition into their meaning.

The definition of equilibrium given in Definition 1 requires that whenever sellers are indifferent between equilibrium strategies, we concentrate on the profile that yields the highest payoff to the platform. The pricing subgames have a unique equilibrium in pure strategies, and therefore no selection criterion is necessary to pin down the equilibrium actions we consider. However, there exist multiple equilibrium profiles of actions at the entry stage of the game, and we use the refinement of Definition 1 to select the relevant equilibrium. In fact, notice that if either (i) $W^M(0) > c$ and $\rho = \rho^M$ or (ii) $c < W^D(1)$ and $\rho = \rho^D$, sellers are indifferent between joining or not. In the former case, taking as given that the opponent will stay out, one of the sellers receives the same expected payoff from joining and not; since the platform is strictly better off if she joins, we select as equilibrium that for $\rho = \rho^M$ one seller enters.

Suppose now that $c < W^D(1)$ and $\rho = \rho^D$: taking as given that the opponent will join, the other seller is indifferent between joining or nor. According to Definition 1, she will join if $\hat{V}^D(\rho^D) > \hat{V}^M(\rho^D)$, and not otherwise. Assumption 2 implies that $\hat{V}^M(\rho) \geq \hat{V}^D(\rho)$ for all $\rho \in [0,1]$, as given in the following lemma.

Lemma 1. *Suppose F satisfies these conditions simultaneously:*

$$\overline{F}\left(\pi^M\right)\left(1 - g(\pi^M)\right) \geq \frac{F(p^M)}{2} \tag{A.1}$$

$$2\overline{F}\left(\pi^M\right) \leq 1 + F(p^M). \tag{A.2}$$

Then

$$2\rho^D\pi^M\overline{F}\left(\rho^D\pi^M\right) \leq \hat{p}^M\left(\rho^D\right)\overline{F}\left(\hat{p}^M\left(\rho^D\right)\right) + \rho^D\pi^M F\left(\hat{p}^M\left(\rho^D\right)\right). \tag{A.3}$$

for all $\rho \in [0,1]$.

We can then conclude that at ρ^D only one seller join. Moreover, this makes sure that the pattern of entry at ρ^D remains the same for every $c < W^D(1)$.[16] The assumption is satisfied by $Uniform$ and $Exponential(\lambda)$ distributions, while some numerical studies suggest that for $Beta(a,b)$ and

[16] If $\hat{V}^M$ and $\hat{V}^D$ were not globally ordered, depending on the value of the other parameters of the model there might be entry of one or both sellers.

$Gamma(a, b)$ it is verified when $b \geq r(a)$ for some increasing function $r(\cdot)$.

Interpretation Towards a more intuitive understanding of Assumption 2, let us recall that p^M is the price that a monopolist would set in a one-shot game, π^M is the profit that it obtains and $\overline{F}(p^M)$ is the demand, i.e., the probability that a random customer would accept to buy at price p^M. By definition, p^M is the price at which the elasticity of demand to price equals 1, $g(p^M) = 1$: for prices $p < p^M$ the elasticity is less than 1, i.e., demand is inelastic, and for prices $p > p^M$ the elasticity is larger than 1, i.e., demand is elastic.[17] Note that condition (A.2) implies $\overline{F}(p^M) \leq \frac{2}{3}$; in other words, if $\overline{F}(p^M) > \frac{2}{3}$ Assumption 2 is violated. Since we know that $\overline{F}(p) \in [0, 1]$ for every p, $\overline{F}(p^M)$ measures the length of the elastic part of the demand schedule on a normalized scale: the larger $\overline{F}(p^M)$, the longer this is, and so the more elastic demand is "on average".[18] Thus, condition (A.2) rules out those distributions for which demand is too elastic on average. On the other hand, condition (A.1) implies that $\overline{F}(p^M)$ cannot be too small, either. Hence, Assumption 2 effectively restricts attention to those distributions that induce neither very elastic nor very inelastic demand schedules. At a high level then, our model better fits those real-world settings where the sensitivity of customers' demand to prices lies on a middle ground.

Relaxing the assumption Outside the scope of Assumption 2, the optimal information disclosure policy can only be computed numerically. Therefore, we conduct a simulation study to assess the effect on consumer surplus of optimally revealing information to sellers, when the distribution of the willingness to pay does not satisfy Assumption 2. We focus on the $Beta(a, b)$ distribution and consider a grid $(a, b) \in [1, 10] \times (0, 20]$ with steps of 0.2 for the parameters, and a grid $\Gamma = \{0, 0.02, 0.04, \ldots, 1\}$ for the beliefs. For any pair (a, b), we consider all possible threshold beliefs $\rho^D \in \Gamma$ and, for fixed ρ^D, all possible priors $\mu \in \Gamma$. First, we derive the optimal policy in every one of these cases. Figure 6 shows the proportion of threshold beliefs such that for all priors $\mu > \rho^D$ consumers' welfare decreases under the optimal policy compared to no-disclosure. As we expect, consumers are worse off under the optimal policy when Assumption 2 holds (the yellow-shaded area in Figure 6, color online); however, consumers' welfare decreases also when the assumption does not hold, and in particular when the platform finds it optimal to fully reveal its information. Figure 6 clarifies that, on average, consumers would be worse off at least 50% of the time under the optimal information disclosure policy relative to no information disclosure if the prior belief was randomly sampled from $[0, 1]$.

In summary, Assumption 2 is a technical condition that simplifies the analysis and allows for a

[17] As remarked in Lariviere (2006), the generalized failure rate of a distribution, denoted by $g(\cdot)$ in the paper, is the (absolute value of the) elasticity of demand. By the IFR assumption, $g(\pi^M) < 1$.

[18] In fact, demand at p^M is large if and only if the price is low.

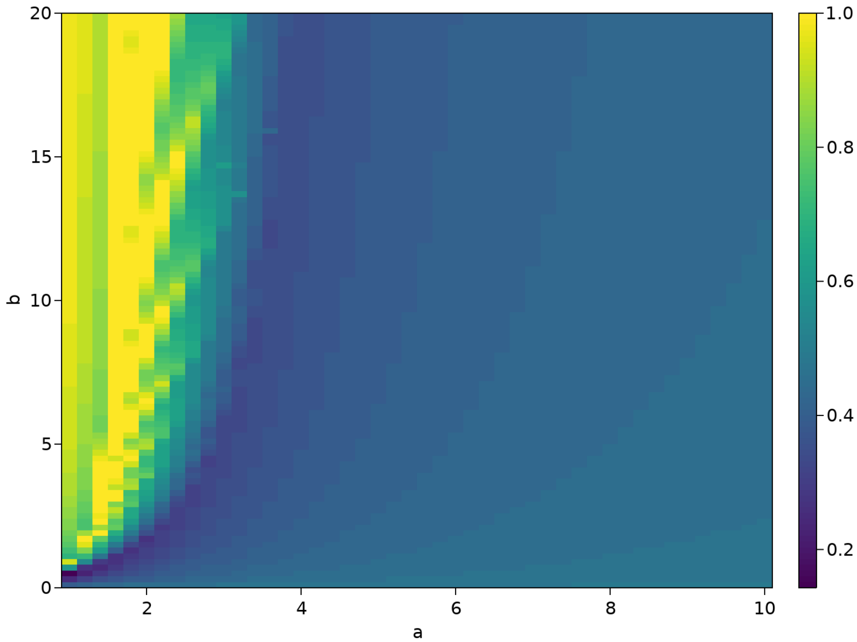

Figure 6: Proportion of threshold beliefs ρ^D such that for all $\mu > \rho^D$ the welfare of consumers decreases under the optimal policy as compared to no-disclosure.

transparent exposition. However, even when the assumption does not hold, the main qualitative insight of our work still holds, i.e., that a platform by appropriately disclosing the information it has at its disposal often leads to lower consumer welfare relative to no information disclosure.

Electronic Companion

Strategic Release of Information in Platforms:
Entry, Competition, and Welfare

Kostas Bimpikis Giacomo Mantegazza *

EC 1 Proofs of results from the main text

Proposition 1 *For every history where only seller i joins the platform, there exist unique prices $\hat{p}^M(\rho)$ and p^M such that it is optimal to set $p_{i,1} = \hat{p}^M(\rho)$ and $p_{i,2} = p^M$. Moreover, $\hat{p}^M > p^M$ if and only if $\rho > 0$.*

Proof. The proof proceeds backwards. At period $t = 2$, it is optimal to set price $p_{i,2}$ such that

$$\max_{p_{i,2}} \quad \rho(1-\alpha)p_{i,2}\overline{F}(p_{i,2}) + (1-\rho)\times 0. \tag{EC.1}$$

Notice that the maximand reproduces equation (3), since $\mathbb{E}_\rho[\mathbf{1}\{\omega = 1\}] = \rho$ and $\mathbb{E}[\{\mathbf{1}\{B_2^* = i\}] = \overline{F}(p_{i,2})$. This is a one-shot monopoly pricing problem, as analyzed in Lariviere (2006). Recalling our notation, $g(p)$ is the generalized failure rate of F and is $h(p)$ the failure rate, for which we have

$$g(p) = ph(p) = p\frac{f(p)}{1-F(p)}.$$

The First Order Condition (FOC) for (EC.1) is

$$\overline{F}(p)[1-g(p)] = 0, \tag{EC.2}$$

and since we assume IFR, [1] a solution to (EC.2) exists, is unique and is interior; we denote it by p^M and write $\pi^M = p^M\overline{F}\left(p^M\right)$.

*Graduate School of Business, Stanford University. Email: {kostasb,giacomom}@stanford.edu

[1] Implicitly we assume $\lim_{p\to a} g(p) < 1$ and $\lim_{p\to b} g(p) > 1$, which rule out distributions such that it would always be optimal to set $p^M = a$ or $p^M = b$. We write IFR for the increasing failure rate property and IGFR for the increasing generalized failure rate property.

Using this result and plugging back into equation (4) for period $t = 1$, seller i's optimal price in this period must solve

$$\max_{p_{i,1}} \quad (1-\alpha)\left[p_{i,1}\overline{F}\left(p_{i,1}\right) + F\left(p_{i,1}\right)\rho\pi^M\right]. \tag{EC.3}$$

The first summand, $p_{i,1}\overline{F}\left(p_{i,1}\right)$, represent the gross expected profit from selling the good to the first customer; with probability $F\left(p_{i,1}\right)$ the first buyer has too low a valuation for the good, and the sale does not happen, so that the unitary good of seller i is still available for purchase, should the second customer arrive.

The FOC associated to (EC.3) can be written as

$$\overline{F}(p)\left[1 - g(p) + \rho\pi^M h(p)\right] = 0. \tag{EC.4}$$

Notice that $h(p) > 0$ for all $p \in (a, b)$, so that

$$\lim_{p\to a}\left[1 - g(p) + \rho\pi^M h(p)\right] = 1 - \lim_{p\to a} g(p) + \rho\pi^M \lim_{p\to a} h(p) > 0,$$

because $\lim_{p\to a} g(p) < 1$. Also

$$\lim_{p\to b}\left[1 - g(p) + \rho\pi^M h(p)\right] = 1 + \lim_{p\to b}\left[\left(\frac{\rho\pi^M}{p} - 1\right) g(p)\right].$$

If $b = \infty$ the limit becomes $1 - \lim_{p\to\infty} g(p)$, which is negative by assumption. If b is finite, then notice that

$$\frac{\rho\pi^M}{b} = \underbrace{\rho\overline{F}\left(p^M\right)}_{<1}\underbrace{\frac{p^M}{b}}_{<1} < 1.$$

Together with $\lim_{p\to b} g(p) > 1$, this implies

$$\lim_{p\to b}\left[1 - g(p) + \rho\pi^M h(p)\right] < 0,$$

which shows that a solution to (EC.4) exists and is in the interior. To show uniqueness of the maximizer, consider the second derivative of the objective function, which is well-defined since f is assumed to be differentiable, given by

$$-2f(p) + f'(p)(\rho\pi^M - p).$$

Using (EC.4), we can rewrite it as

$$-2f(p^*) - f'(p^*)\frac{\overline{F}(p^*)}{f(p^*)},$$

where p^* denotes that it is evaluated at the stationary point. Then the Second Order Condition is

$$\begin{aligned}-2f(p^*) - f'(p^*)\frac{\overline{F}(p^*)}{f(p^*)} < 0 &\iff 2f(p^*) + f'(p^*)\frac{\overline{F}(p^*)}{f(p^*)} > 0 \\ &\iff f(p)^2 + \overline{F}(p)^2 h'(p) > 0.\end{aligned}$$

This condition is satisfied at every stationary point because the distribution is IFR, i.e. $h'(p) \geq 0$. We showed that the objective function in (EC.3) is strictly concave at any stationary point, which we are assured to exist. Since it is a continuously differentiable function (because F necessarily has continuous density), this implies that it cannot have more than one stationary point, which is also the unique global maximizer; denote it by $\hat{p}^M$.

Finally, plugging p^M in (EC.4) and using (EC.2), we have

$$\overline{F}\left(p^M\right)\left[1 - g\left(p^M\right) + \rho\pi^M h\left(p^M\right)\right] = \overline{F}\left(p^M\right)\rho\pi^M h\left(p^M\right) \geq 0,$$

with equality if and only if $\rho = 0$; therefore $\hat{p}^M \geq p^M$. ■

Corollary 1 *The functions W^M and $\hat{V}^M$ are increasing and strictly convex in ρ. Furthermore, $\hat{p}^M$ is increasing in ρ.*

Proof. Setting aside α, which does not affect monotonicity and curvature, we consider the function

$$V^M(\rho) = \hat{p}^M(\rho)\overline{F}\left(\hat{p}^M(\rho)\right) + \rho\pi^M F\left(\hat{p}^M(\rho)\right),$$

where we highlight that $\hat{p}^M$ is a function of ρ, too. The derivative with respect to ρ is

$$\begin{aligned}\frac{\partial V^M}{\partial \rho} &= \frac{\partial \hat{p}^M}{\partial \rho}\left[\overline{F}\left(\hat{p}^M\right) - \hat{p}^M f\left(\hat{p}^M\right) + \rho\pi^M f\left(\hat{p}^M\right)\right] + \pi^M F\left(\hat{p}^M\right) \\ &= \frac{\partial \hat{p}^M}{\partial \rho}\overline{F}\left(\hat{p}^M\right)\left[1 - g\left(\hat{p}^M\right) + \rho\pi^M h\left(\hat{p}^M\right)\right] + \pi^M F\left(\hat{p}^M\right) \\ &= \pi^M F\left(\hat{p}^M\right) > 0,\end{aligned}$$

where the last equality holds because $\hat{p}^M$ must satisfy (EC.4) (equivalently, by the envelope theorem). For the second derivative,

$$\frac{\partial^2 V^M}{\partial \rho^2} = \frac{\partial}{\partial \rho}\left[\pi^M F\left(\hat{p}^M\right)\right] = \pi^M f\left(\hat{p}^M\right)\frac{\partial \hat{p}^M}{\partial \rho} > 0,$$

since, as we will show, $\hat{p}^M$ is strictly increasing in ρ.

Applying the Implicit Function Theorem to (EC.4), one obtains

$$\frac{\partial \hat{p}^M}{\partial \rho} = \frac{-\pi^M h\left(\hat{p}^M\right)}{\rho \pi^M h'\left(\hat{p}^M\right) - g'\left(\hat{p}^M\right)}. \tag{EC.5}$$

While the numerator is negative, we can rewrite the denominator as

$$\begin{aligned}&\rho\pi^M h'\left(\hat{p}^M\right) - g'\left(\hat{p}^M\right) = \left(\rho\pi^M - \hat{p}^M\right) h'\left(\hat{p}^M\right) - h\left(\hat{p}^M\right)\\ &= -\frac{\overline{F}\left(\hat{p}^M\right)}{f\left(\hat{p}^M\right)} h'\left(\hat{p}^M\right) - h\left(\hat{p}^M\right) < 0,\end{aligned}$$

because $h(p) > 0$ on all the support. Comparing this with the numerator in (EC.5) we get $\partial\hat{p}^M/\partial\rho > 0$, which justifies the above claim about V^M. ■

Proposition 2 *At every history when both sellers join the platform, the unique equilibrium prices are given as* $p_{i,1} = p_{j,1} = \rho\pi^M$ *and*

$$p_{i,2} = \begin{cases} p^M & \text{if } j \text{ makes a sale at } t = 1, \\ 0 & \text{otherwise.}\end{cases}$$

Proof. Recalling the payoff function for seller i from equation (8), notice that when $p_j^1 > p^M$, then also $p_j^1 > \rho\pi^M$ for every $\rho \in [0,1]$. This means that

$$\rho\pi^M \overline{F}\left(p_j^1\right) < p_j^1 \overline{F}\left(p_j^1\right) < \pi^M,$$

so that the optimal price is $p_i^1 = p^M$ if $p_j^1 > p^M$. On the other hand, when $p_j^1 \leq p^M$, $\Pi_{1,i}^S(p_i^1;\rho)$ is strictly increasing for $p_i^1 < p_j^1$ and constant thereafter; therefore, the optimal action is some $p_i^1 > p_j^1$ whenever $\rho\pi^M > p_j^1$ and some $p_i^1 \geq p_j^1$ when $\rho\pi^M = p_j^1$. Finally, when $\rho\pi^M < p_j^1$, there is no well-defined best reply: for every $p_i^1 < p_j^1$ there exists $\varepsilon > 0$ such that

$$(p_i^1 + \varepsilon)\overline{F}\left(p_i^1 + \varepsilon\right) > p_i^1 \overline{F}\left(p_i^1\right),$$

and $p_i^1 + \varepsilon < p_j^1$, while

$$\Pi_{1,i}^S \left(p_i^1 = p_j^1; \rho\right) < \lim_{p \uparrow p_j^1} p\overline{F}(P).$$

Therefore the best-response correspondence for seller i is

$$BR_i(p_j^1) = \begin{cases} \left(p_j^1, \infty\right), & p_j^1 < \rho\pi^M \\ \left[p_j^1, \infty\right), & p_j^1 = \rho\pi^M \\ \varnothing, & \rho\pi^M < p_j^1 \leq p^M \\ p^M, & p_j^1 > p^M \end{cases}$$

Now, it is easy to see that the only pair (p_i^1, p_j^1) for which the equilibrium condition

$$\begin{cases} \{p_i^1\} \subseteq BR_i(p_j^1) \\ \{p_j^1\} \subseteq BR_j(p_i^1) \end{cases}$$

is satisfied is exactly $(\rho\pi^M, \rho\pi^M)$. ■

Corollary 2 *The functions W^D and $\hat{V}^D$ as defined in (9) and (10) are increasing and concave in belief ρ.*

Proof. Let us consider only the function $V^D(\rho) = \rho\pi^M\overline{F}\left(\rho\pi^M\right)$. Then we have

$$\frac{\partial V^D}{\partial \rho} = \pi^M \left[\overline{F}\left(\rho\pi^M\right) - \rho\pi^M f\left(\rho\pi^M\right)\right] \geq 0$$
$$\Longleftrightarrow \overline{F}\left(\rho\pi^M\right) \geq \rho\pi^M f\left(\rho\pi^M\right) \Longleftrightarrow 1 \geq g\left(\rho\pi^M\right).$$

But remember $\rho\pi^M = \rho p^M \overline{F}(p^M) < p^M$ and by IGFR property it must be $1 = g(p^M) \geq g\left(\rho\pi^M\right)$; therefore V^D is increasing in ρ.

Moreover, it can be checked that

$$\frac{\partial^2 V^D}{\partial \rho^2} = -\left(\pi^M\right)^2 \left[f\left(\rho\pi^M\right)\left(1 - g\left(\rho\pi^M\right)\right) + \overline{F}\left(\rho\pi^M\right) g'\left(\rho\pi^M\right)\right] \leq 0,$$

again by IGFR and what we proved above. ■

Proposition 3 *Let ρ denote the sellers' common posterior belief about demand after they obtain the platform's message. Then, the following hold along the equilibrium path:*

(i) Suppose that $c > W^M(0)$. Then, there exists $\rho^M(c,\alpha)$ such that if $\rho \geq \rho^M(c,\alpha)$ a single seller joins the platform and, otherwise, there is no entry. In addition, when $c \geq W^M(1)$, we have $\rho^M(c,\alpha) \geq 1$, i.e., entry is too costly for sellers irrespective of their beliefs about demand.

(ii) Suppose that $W^D(1) \leq c \leq W^M(0)$. Then, a single seller joins irrespective of the platform's information disclosure policy.

(iii) Suppose that $c < W^D(1)$. Then, there exists $\rho^D(c,\alpha)$ such that if $\rho > \rho^D(c,\alpha)$ both sellers join the platform and, otherwise, there is a single entrant.

Proof. As further discussed in Appendix A, under Assumption 2 the platform's expected profits are always strictly larger with just one entrant than when both sellers join. This is formalized in the following lemma.

Lemma 1 *Suppose F satisfies these conditions simultaneously:*

$$\overline{F}\left(\pi^M\right)\left(1-g(\pi^M)\right) \geq \frac{F(p^M)}{2}$$
$$2\overline{F}\left(\pi^M\right) \leq 1 + F(p^M),$$

Then

$$2\rho^D\pi^M\overline{F}\left(\rho^D\pi^M\right) \leq \hat{p}^M\left(\rho^D\right)\overline{F}\left(\hat{p}^M\left(\rho^D\right)\right) + \rho^D\pi^M F\left(\hat{p}^M\left(\rho^D\right)\right).$$

for all $\rho \in [0,1]$

For $c > W^M(1)$, since $W^M(1)$ is the maximal amount seller i can expect to earn under any market condition, there will be no entrance; this is equivalent to $\rho^M(c,\alpha) > 1$, with ρ^M defined as below.

For $W^M(0) < c \leq W^M(1)$, since the function W^M is strictly increasing in the belief, there exists a unique solution in $[0,1]$ to the equation $W^M(\rho) = c$; denote it with $\rho^M(c,\alpha)$. When $\rho < \rho^M$, then $W^D(\rho) < W^M(\rho) < c$; hence, no seller has incentive to enter. If instead $\rho > \rho^M$, then $W^M(\rho) > c > W^D(\rho)$. Suppose seller i decides to enter: it follows that seller j is better off by staying out. Vice versa, if j joins, then i should stay out to maximize profits. Therefore, at equilibrium only one seller joins, and the other stays out. Finally, suppose $\rho = \rho^M$, so that $W^M(\rho) = c > W^D(\rho)$: if seller j stays out, seller i is indifferent between joining and not joining. Since the platform's expected payoff is higher when i enters than when she does not, the Sender-preferred equilibrium requires that i enters at belief ρ^M.

Take now $W^D(1) \leq c \leq W^M(0)$. For every $\rho \in (0,1)$, $W^M(\rho) > c > W^D(\rho)$, and therefore the same argument of before applies. For $\rho = 0$, if $c < W^M(0)$, once again we are in the same case as above; if $c = W^M(0)$, then at indifference the equilibrium definition requires one seller to join. Finally, for $\rho = 1$, if $c > W^D(1)$, then entry is never profitable for both sellers, and hence just one joins; if instead $c = W^D(1)$, if seller i joins, then seller j is indifferent between entering and not. Since the platform is better off with one seller, our equilibrium requires that j decides to stay out.

When $c < W^D(1)$, then since W^D is strictly increasing there exists a unique solution in $[0, 1]$ to the equation $W^D(\rho) = c$; denote this solution by $\rho^D(c, \alpha)$. Whenever $\rho < \rho^D$, then $W^M(\rho) > c > W^D(\rho)$ and thus only one joins in equilibrium. If $\rho > \rho^D$, then $W^D(\rho) > c$; thus, if seller i joins, also seller j is better off by joining, and notice that seller i does not have incentive to change its behaviour, because by opting out she would earn less than $W^D(\rho)$. Hence, it is equilibrium that both enter. Finally, for $\rho = \rho^D$, if seller i enters then j is indifferent between joining and not; again by the Sender-preferred equilibrium definition, it must be that she stays out. Hence, at belief ρ^D only one seller enters. This concludes the proof. ■

Lemma 1 *Suppose F satisfies these conditions simultaneously:*

$$\overline{F}\left(\pi^M\right)\left(1 - g(\pi^M)\right) \geq \frac{F(p^M)}{2}$$
$$2\overline{F}\left(\pi^M\right) \leq 1 + F(p^M),$$

Then

$$2\rho^D \pi^M \overline{F}\left(\rho^D \pi^M\right) \leq \hat{p}^M\left(\rho^D\right)\overline{F}\left(\hat{p}^M\left(\rho^D\right)\right) + \rho^D \pi^M F\left(\hat{p}^M\left(\rho^D\right)\right).$$

for all $\rho \in [0, 1]$

Proof. In Corollary 1 we proved that $\hat{V}^M(\rho)$ is strictly convex. Therefore, (A.3) is implied by

$$2\rho\pi^M \overline{F}\left(\rho\pi^M\right) \leq \pi^M + \rho\pi^M F(p^M), \tag{EC.6}$$

where the RHS is the first order approximation of $\hat{p}^M(\rho)\overline{F}\left(\hat{p}^M(\rho)\right) + \rho\pi^M F\left(\hat{p}^M(\rho)\right)$ at $\rho = 0$. Notice that (EC.6) is trivially satisfied as a strict inequality at $\rho = 0$. To prove that inequalities (A.1)-(A.2) imply (EC.6), we show that under Assumption 2 the function l, defined by

$$l(\rho) = 1 + \rho F(p^M) - 2\rho\overline{F}\left(\rho\pi^M\right),$$

is decreasing in ρ and satisfies $l(1) \geq 0$. Indeed

$$l(1) = 1 + F(p^M) - 2\overline{F}\left(\pi^M\right) \geq 0$$

is precisely condition (A.2). We can also notice that, as proved in Corollary 2, $2\rho\overline{F}\left(\rho\pi^M\right)$ is concave, so that $l(\rho)$ is convex. Hence, to show that it is decreasing it is sufficient (and necessary) that $\left.\frac{\partial l(\rho)}{\partial\rho}\right|_{\rho=1} \leq 0$. Therefore, since

$$\frac{\partial l(\rho)}{\partial\rho} = F(p^M) - 2\overline{F}\left(\rho\pi^M\right)\left(1 - g\left(\rho\pi^M\right)\right)$$

the condition (A.1) in Lemma 1 implies that also $\left.\frac{\partial l(\rho)}{\partial\rho}\right|_{\rho=1} \leq 0$ is satisfied. ■

Theorem 1 *Suppose that Assumptions 1 and 2 hold. Then, there exists an optimal policy $(\mathcal{D}^*, M^*)$ that involves sending one of two messages, i.e., $M^* = \{Y, N\}$. Moreover, if we let $\mathcal{D}^*(\omega)$ denote the probability that message Y is sent when the state of the world is ω under information disclosure policy $(\mathcal{D}^*, \{Y, N\})$, an optimal policy for the platform takes the following form:*

(i) When $c > W^M(0)$, then

$$\mathcal{D}^*(1) = \begin{cases} 1 & \text{for } \mu < \rho^M(c,\alpha) \\ q_u^M & \text{for } \mu \geq \rho^M(c,\alpha) \end{cases} \quad \text{and} \quad \mathcal{D}^*(0) = \begin{cases} q_l^M & \text{for } \mu < \rho^M(c,\alpha) \\ 0 & \text{for } \mu \geq \rho^M(c,\alpha) \end{cases},$$

where the disclosure probabilities are equal to

$$q_l^M = \frac{\mu(1-\rho^M(c,\alpha))}{\rho^M(c,\alpha)(1-\mu)} \quad \text{and} \quad q_u^M = \frac{\mu - \rho^M(c,\alpha)}{\mu(1-\rho^M(c,\alpha))};$$

(ii) When $W^D(1) \leq c \leq W^M(0)$, full-disclosure is optimal, i.e.,

$$\mathcal{D}^*(\omega) = \omega;$$

(iii) Finally, when $c < W^D(1)$,

$$\mathcal{D}^*(1) = \begin{cases} 1 & \text{for } \mu \leq \rho^D(c,\alpha) \\ q_u^D & \text{for } \mu > \rho^D(c,\alpha) \end{cases} \quad \text{and} \quad \mathcal{D}^*(0) = \begin{cases} q_l^D & \text{for } \mu \leq \rho^M(c,\alpha) \\ 0 & \text{for } \mu > \rho^D(c,\alpha) \end{cases},$$

where the disclosure probabilities are equal to

$$q_l^D = \frac{\mu(1-\rho^D(c,\alpha))}{\rho^D(c,\alpha)(1-\mu)} \quad \text{and} \quad q_u^D = \frac{\mu - \rho^D(c,\alpha)}{\mu(1-\rho^D(c,\alpha))}.$$

Proof. We first show that the problem of the platform is equivalent to one with just one receiver.

Consider the following game between two players: one, which we call P', is an *alter ego* of the platform, while the other one, which we call S, represents the system of sellers. S's action space is the set $\{00, 01, 11\}$. P' and S engage in a two-stage game in which first P' observes the state of the world and sends a message according to an information disclosure policy; subsequently S takes an action, which affects its own and the payoff of P'. The expected utility of S is given by the following function,[2]

$$u(I, \mu) = \begin{cases} c & I = 00 \\ W^M(\mu) & I = 01 \\ W^D(\mu)\mathbf{1}\{W^D(\mu) \leq c\} + \left[W^M(\mu) + \varepsilon\right] \mathbf{1}\{W^D(\mu) > c\} & I = 11 \end{cases}$$

for some $\varepsilon > 0$. Let us interpret action 00 being *No one joins*, 01 being *Only one joins* and 11 being *Both join*. S's optimal action given belief μ about the state of the world is the same as the entry decision that realizes in our model. The expected utility of P' is instead

$$\hat{u}(I, \mu) = \begin{cases} 0 & I = 00 \\ \hat{V}^M(\mu) & I = 01 \\ \hat{V}^D(\mu) & I = 11 \end{cases}$$

The information design problem of P' is equivalent to that of the platform in the original model, in the sense that every optimal solution for P' is optimal for the platform and vice-versa. This setting satisfies the hypotheses of Anunrojwong, Iyer, and Lingenbrink (2020), so we can use one of their preliminary results and take the message space to be $M = \Delta(\Omega) = \Delta(\{0, 1\}) = [0, 1]$ without loss of generality. The problem can then be stated as

$$\begin{aligned} &\max_{\tau \in \Delta([0,1])} \mathbb{E}_{\rho \sim \tau}\left[\hat{u}(J, \rho)\right] \quad \text{s.t.} \\ &J \in \underset{I \in \{00,01,11\}}{\arg\max}\; u(I, \rho)\ \forall \rho \in [0, 1] \\ &\mathbb{E}_\tau[\rho] = \mu \end{aligned}$$

Let us define the function $\hat{V} : [0, 1] \to \mathbb{R}_+$, given by $\hat{V}(\rho) = 0$ for $c > W^M(1)$, and further as

$$\hat{V}(\rho) = \begin{cases} 0 & \text{if } \rho < \rho^M(c, \alpha) \\ \hat{V}^M(\rho) & \text{if } \rho \geq \rho^M(c, \alpha) \end{cases}$$

[2] Expectation is taken with respect to μ.

for $W^M(0) < c \leq W^M(1)$. As $\hat{V}(\rho) = \hat{V}^M(\rho)$ for $W^D(1) \leq c \leq W^M(0)$. Finally, as

$$\hat{V}(\rho) = \begin{cases} \hat{V}^M(\rho) & \text{if } \rho \leq \rho^D(c, \alpha) \\ \hat{V}^D(\rho) & \text{if } \rho > \rho^D(c, \alpha) \end{cases}$$

for $c < W^D(1)$. Since $\hat{V}(\rho) = \hat{u}(J, \rho)$ if $J \in \arg\max_{I \in \{00,01,11\}} u(I, \rho)$ $\forall \rho \in [0, 1]$, the original problem can be further restated as

$$\begin{aligned} &\max_{\tau \in \Delta([0,1])} \mathbb{E}_{\rho \sim \tau} \left[\hat{V}(\rho) \right] \quad \text{s.t.} \\ &\mathbb{E}_\tau[\rho] = \mu \end{aligned}$$

This is now a problem similar to those analyzed by, e.g., Kamenica and Gentzkow (2011) and Aumann and Maschler (1995), and we identify the optimal value function for the platform with the "concavification" of $\hat{V}$, i.e.

$$V^*(\mu) = \sup \left\{ y : \ (y, \mu) \in \mathtt{cov}(\hat{V}) \right\},$$

where $\mathtt{cov}(\hat{V})$ denotes the convex hull of the graph of $\hat{V}$. In other terms, this is the concave closure of $\hat{V}$. We now turn to characterizing V^* in each of the cases we have for $\hat{V}$. To avoid trivialities, suppose $c \leq W^M(1)$.

When $W^M(0) < c \leq W^M(1)$, $\hat{V}$ stays at zero for $\mu < \rho^M$ and then jumps to a positive convex function. Notice that full-disclosure is not optimal, unless $c = W^M(0)$: in fact, the value function corresponding to such policy is equal to $\hat{V}^M(1)\mu$. However, it has

$$\begin{aligned} \hat{V}^M(\mu) &> \hat{V}^M(1) + \alpha \mu \pi^M F\left(\hat{p}^M(1)\right)(\mu - 1) \\ &\geq \hat{V}^M(1)\mu, \end{aligned}$$

where the first equality holds because the RHS is a linear approximation to $\hat{V}^M$ at $\mu = 1$ and the second by algebra. Thus, the concavification of $\hat{V}$ is the piecewise-linear function

$$V^*(\mu) = \begin{cases} \dfrac{1}{\rho^M} \dfrac{\alpha c}{1 - \alpha} \mu & \mu < \rho^M \\ \dfrac{\mu - \rho^M}{1 - \rho^M} \left(\hat{V}^M(1) - \dfrac{\alpha c}{1 - \alpha} \right) + \dfrac{\alpha c}{1 - \alpha} & \mu \geq \rho^M \end{cases}$$

since $\hat{V}^M(\rho^M) = \frac{\alpha c}{1-\alpha}$. Thus, we infer that the optimal strategy for the platform is to alternatively induce beliefs 0 and ρ^M when $\mu < \rho^M$, and beliefs ρ^M and 1 when $\mu \geq \rho^M$. To determine the optimal probability with which each belief is induced, we employ the constraint on τ given by $\mathbb{E}_\tau[\rho] = \mu$: in this case τ should put mass only on two points in $[0,1]$ and their mean must be μ. Hence: for $\mu < \rho^M$ we need $0 \times (1-\tau) + \tau \times \rho^M = \mu$ so that $\tau = \frac{\mu}{\rho^M}$. An information disclosure policy that sends message ρ^M with probability q_l^M (and 0 with complementary probability) when $\omega = 0$, and with probability 1 when $\omega = 1$, induces this distribution of posterior beliefs; for $\mu \geq \rho^M$, the constraint requires $\rho^M \times (1-\tau) + 1 \times \tau = \mu$, which implies $\tau = \frac{\mu - \rho^M}{1-\rho^M}$. A policy sending message ρ^M with probability q_u^M (and 0 with complementary probability) when $\omega = 1$, and always message 0 when $\omega = 0$, induces this distribution.

Suppose now $W^D(1) \leq c \leq W^M(0)$: since $\hat{V}(\mu) = \hat{V}^M(\mu)$ and the latter is strictly convex, the concavification V^* is the straight line joining $(0, \pi^M)$ and $(1, \hat{V}^M(1))$ in the (μ, v) plane. Consequently, one easily reads that the platform optimal strategy is to fully reveal all the information it has. This is achieved, for example, by the policy that sends message 0 when $\omega = 0$ and message 1 when $\omega = 1$.

Finally, assume $c < W^D(1)$. To find the optimal policy for this case we use Assumption 2. Full transparency is not optimal because the value function of this policy would be $\widetilde{V}(\mu) = \alpha\left[\pi^M + \mu\pi^M\left(2\overline{F}(\pi^M) - 1\right)\right]$, and we have for $\mu > 0$

$$\begin{aligned}\hat{V}^M(\mu) &> \alpha\left[\pi^M + \mu\pi^M F(p^M)\right] \\ &\geq \alpha\left[\pi^M + \mu\pi^M\left(2\overline{F}(\pi^M) - 1\right)\right] = \widetilde{V}(\mu),\end{aligned}$$

where the first inequality follows because the RHS is the linear approximation of $\hat{V}^M$ at $\mu = 0$ and the second reduces to the second condition in (11). In addition, and always thanks to conditions (11), $\hat{V}^D$ always lies below a linear approximation to $\hat{V}^M$ at $\mu = 0$: this implies that for any threshold ρ^D the straight line joining $(\rho^D, \hat{V}^M(\rho^D))$ with $(1, 2\pi^M\overline{F}(\pi^M))$, which is the value of a policy that induces beliefs ρ^D and 1, always dominates any other policy that induces different beliefs when $\rho^D < \mu \leq 1$.[3] Hence, also in this case the optimal value function is a piecewise-linear function, given by

$$V^*(\mu) = \begin{cases} \alpha\pi^M + \left(\hat{V}^M(\rho^D) - \alpha\pi^M\right)\dfrac{\mu}{\rho^D} & \mu \leq \rho^D \\ \hat{V}^M(\rho^D) + \left(\alpha 2\pi^M\overline{F}(\pi^M) - \hat{V}^M(\rho^D)\right)\dfrac{\mu - \rho^D}{1-\rho^D} & \mu > \rho^D \end{cases}$$

As before, this value function tells that the optimal policy is to alternatively induce beliefs 0 and ρ^D when $\mu \leq \rho^D$ and beliefs ρ^D and 1 when $\mu > \rho^D$. Following the same reasoning of the case

[3] In particular, the first condition in (11) ensures that it is not optimal to induce beliefs ρ^D and $\eta < 1$ for $\rho^D < \mu \leq \eta$ and do nothing for $\mu > \eta$.

above where $W^M(0) < c \leq W^M(1)$, one can check that an optimal policy is to send message ρ^D with probability 1 when $\omega = 1$ and with probability q_l^D (and 0 with complementary probability) when $\omega = 0$, if $\mu \leq \rho^D$; and to send message 1 with probability q_u^D (and ρ^D with complementary probability) when $\omega = 1$ and ρ^D with probability 1 when $\omega = 0$, if $\mu > \rho^D$.

Finally, let us observe that the policy we have described always sends one of two possible messages for every combination of c and α. Therefore, the message space can be restricted to $\{Y, N\}$, so that we recover the policy given in the statement of the proposition. ■

Corollary 3 *The following hold true:*

(i) The optimal information disclosure policy always yields strictly higher profits for the platform than no-disclosure. In addition, it strictly outperforms full-disclosure too, unless $W^D(1) \leq c \leq W^M(0)$.

(ii) No-disclosure yields strictly higher profits than full-disclosure if (a) $c > W^M(0)$ and $\mu \geq \rho^M$, or (b) $c < W^D(1)$ and $\mu \leq \rho^D$.

Proof. Since $\mu = 0$ and $\mu = 1$ are absorbing cases in which the disclosure problem is trivial (in fact, the prior belief cannot be modified), without loss of generality assume $\mu \neq 0, 1$.

Notice that the first part of the corollary follows directly from Theorem 1. For the second part, suppose $W^M(0) < c \leq W^M(1)$ and $\mu < \rho^M$: the no-disclosure policy gives 0 profit in this case, while full-disclosure yields $\mu \hat{V}^M(1) > 0$, so full-disclosure strictly dominates no-disclosure. If $\mu > \rho^M$, then no-disclosure strictly dominates full-disclosure if and only if

$$\mu \hat{V}^M(1) < \hat{V}^M(\mu) \iff \frac{\hat{V}^M(1)}{1} < \frac{\hat{V}^M(\mu)}{\mu}.$$

$\frac{\hat{V}^M(\mu)}{\mu}$ is a strictly decreasing function of μ, since

$$\frac{\partial}{\partial \mu}\left(\frac{\hat{V}^M(\mu)}{\mu}\right) = \alpha \frac{\mu F\left(\hat{p}^M(\mu)\right)\pi^M - \hat{V}^M(\mu)}{\mu^2} < 0,$$

where the inequality follows by equation (3.1), and thus claim (a) holds.

When $W^D(1) \leq c \leq W^M(0)$ full-disclosure yields larger profits, because $\hat{V}^M$ is strictly convex. So, assume $c < W^D(1)$: full-disclosure gives $\mu \hat{V}^D(1) + (1-\mu)\hat{V}^M(0)$. If $\mu \leq \rho^D$, no-disclosure earns $\hat{V}^M(\mu)$. As in the proof of Theorem 1, conditions (11) imply that the linear approximation of $\hat{V}^M$ at $\mu = 0$ is always strictly larger than the value of full-disclosure. Recalling the convexity of $\hat{V}^M$ yields the result. Finally, for $\mu > \rho^D$, full-disclosure yields strictly larger profits if the linear approximation to $\hat{V}^D$ at $\mu = 1$ is always strictly smaller than the value of full-disclosure. Notice the two are equal at $\mu = 1$ by construction, so it is enough (by concavity of $\hat{V}^D$) to show that $\frac{\partial \hat{V}^D}{\partial \mu} > \hat{V}^D(1) - \hat{V}^M(0)$ at $\mu = 1$. From Corollary 2, we have

$$\begin{aligned}\frac{\partial \hat{V}^D(1)}{\partial \mu} = 2\pi^M \left[\overline{F}(\pi^M) - \pi^M f(\pi^M)\right] &> \pi^M \left[2\overline{F}(\pi^M) - 1\right] \\ &\iff 1 > 2\overline{F}\left(\pi^M\right) g(\pi^M).\end{aligned}$$

By the first inequality in (11) we have $2\overline{F}\left(\pi^M\right) - F\left(p^M\right) \geq 2\overline{F}\left(\pi^M\right) g(\pi^M)$, and the LHS is smaller than 1 by the second inequality in (11), so also (b) holds. ■

Theorem 2 *Suppose that $c < W^D(1)$. Then, if $\mu \leq \rho^D(c, \alpha)$, social welfare, consumer surplus and profits for the sellers increase under the optimal policy. On the other hand, if $\mu > \rho^D(c, \alpha)$, then profits for sellers increase; however, both consumer surplus and aggregate welfare decrease.*

Proof. Suppose first that $\mu \leq \rho^D$: under the optimal policy, if $\omega = 1$ message Y is always sent, which induces belief ρ^D; if $\omega = 0$, message Y is sent with probability q_l^D, and message N with complementary probability, so that after N belief 0 obtains . The variation in social welfare is then

$$\begin{aligned}\Delta SW = \mu \Bigg[&\int_{\hat{p}^M(\rho^D)}^{\infty} v dF + F\left(\hat{p}^M(\rho^D)\right) \int_{p^M}^{\infty} v dF + c - \int_{\hat{p}^M(\mu)}^{\infty} v dF - F\left(\hat{p}^M(\mu)\right) \int_{p^M}^{\infty} v dF - c \Bigg] \\ &+ (1-\mu) \left[q_l^D \left(\int_{\hat{p}^M(\rho^D)}^{\infty} v dF + c \right) + \left(1 - q_l^D\right) \left(\int_{\hat{p}^M(0)}^{\infty} v dF + c \right) - \int_{\hat{p}^M(\mu)}^{\infty} v dF - c \right] \\ &= \tau_l^D SW^M(\rho^D) + (1 - \tau_l^D) SW^M(0) - SW^M(\mu),\end{aligned}$$

recalling Definition 2 and that $\tau_l^M = \frac{\mu}{\rho^D}$ from the proof of Theorem 1. Since we know that $\mu = \tau_l^D \rho^D + (1 - \tau_l^D) \times 0$, whether social welfare increases or not depends on whether the function $SW^M(\mu)$ is convex or concave. Moreover, it is easy to see that in this case $CS^M(\mu) = SW^M(\mu) - V^M(\mu) - c$, where $V^M = \hat{V}^M + W^M$, and therefore

$$\begin{aligned}\Delta CS &= \tau_l^D CS^M(\rho^D) + \left(1 - \tau_l^D\right) CS^M(0) - CS^M(\mu) \\ &= \tau_l^D \left[SW^M(\rho^D) - V^M(\rho^D) - c\right] + (1 - \tau_l^D) \left[SW^M(0) - V^M(0) - c\right] - SW^M(\mu) + V^M(\mu) + c.\end{aligned}$$

Since V^M is convex, if $CS^M(\mu)$ is convex also $SW^M(\mu)$ is: under the uniform assumption one obtains

$$CS(\mu) = \frac{1}{2}\left(\frac{1}{2} + \frac{\mu}{8}\right)^2 + \frac{\mu}{8}\left(\frac{\mu}{8} - \frac{1}{2}\right),$$

which is strictly convex. Hence both social welfare and consumer surplus increase by adopting the optimal policy. Finally, since also W^M is convex, the profit of the seller that decides to join increases; since the expected profit of the other seller remains unchanged, information disclosure also increases aggregate sellers' profits in this case.

Consider now the case $\mu > \rho^D$. We begin with sellers' profits: after message Y both sellers join, while after N, which induces ρ^D, only one enters; accordingly the variation in aggregate profits is

$$\begin{aligned}
\Delta \sum_{i=1}^{2} \Pi_{0,i}^S =& \mu \Bigg[q_u^D (1-\alpha) 2\pi^M \overline{F}(\pi^M) + \left(1-q_u^D\right)(1-\alpha)\left(\hat{p}^M(\rho^D)\overline{F}\left(\hat{p}^M(\rho^D)\right) + F\left(\hat{p}^M(\rho^D)\right)\pi^M + c\right) \\
& -(1-\alpha)\overline{F}\left(\mu\pi^M\right)\pi^M(1+\mu)\Bigg] + (1-\mu)\Bigg[(1-\alpha)\hat{p}^M(\rho^D)\overline{F}\left(\hat{p}^M(\rho^D)\right) + c - (1-\alpha)\mu\pi^M\overline{F}(\mu\pi^M)\Bigg] \\
=&(1-\alpha)\Big[\tau_u^D 2\pi^M\overline{F}(\pi^M) + (1-\tau_u^D)\left(\hat{p}^M(\rho^D)\overline{F}\left(\hat{p}^M(\rho^D)\right) + F\left(\hat{p}^M(\rho^D)\right)\pi^M\right) \\
& - 2\mu\pi^M\overline{F}(\mu\pi^M)\Big] + c(1-\tau_u^D) \\
=&\frac{1-\alpha}{\alpha}\left[\tau_u^D \hat{V}^D(1) + (1-\tau_u^D)\hat{V}^M(\rho^D) - \hat{V}^D(\mu)\right] + c(1-\tau_u^D) \\
=&\frac{1-\alpha}{\alpha}\left[V^*(\mu) - \hat{V}^D(\mu)\right] + c(1-\tau_u^D) \geq 0,
\end{aligned}$$

where $\tau_u^D = \frac{\mu-\rho^D}{1-\rho^D}$ and the last inequality follows by definition of V^*. The inequality is strict for all $\mu \neq \rho^D, 1$. Let us now consider consumer surplus, whose variation is

$$\Delta CS = \tau_u^D CS^D(1) - (1-\tau_u^D)CS^M(\rho^D) - CS^D(\mu).$$

We cannot immediately deduce the sign of the change in consumer surplus, because with the entry of a second seller the curvature changes. Using the uniformity hypothesis we have that $\rho^D = 2\left(1-\sqrt{1-\frac{4c}{1-\alpha}}\right)$, $W^D(1) = \frac{3}{16}(1-\alpha)$ so that

$$\Delta CS = \frac{\tau_u^D}{2} + (1-\tau_u^D)\left[\frac{1}{2}\left(\frac{1}{2}+\frac{\rho^D}{8}\right)^2 + \frac{\rho^D}{8}\left(\frac{\rho^D}{8}-\frac{1}{2}\right)\right] - \frac{\mu^2}{8} + \frac{\mu}{8} - \frac{1}{2} < 0,$$

for all $c < \frac{3}{16}(1-\alpha)$, $\alpha \in [0,1)$ and $\mu > \rho^D$. Finally, the variation in social welfare is

$$\begin{aligned}\Delta SW =&\tau_u^D SW^D(1) - (1-\tau_u^D)SW^M(\rho^D) - SW^D(\mu)\\ =&\tau_u^D\left(\frac{7}{8}\right) + (1-\tau_u^D)\left(\frac{5}{128}(\rho^D)^2 + \frac{\rho^D}{8} + \frac{3}{8} + c\right) - \frac{3}{8}\mu - \frac{1}{2}\\ =&(1-\mu)\left[\frac{27}{128} + \frac{c}{1-\rho^D} - \frac{5\rho^D}{128} - \frac{43}{128(1-\rho^D)}\right].\end{aligned}$$

Further algebra shows that $\Delta SW(\mu) < 0$ always. ■

Theorem 3 *When $c > W^M(0)$, both social welfare and consumer surplus increase under the optimal policy. In addition, when $\mu \geq \rho^M(c,\alpha)$, the profits for sellers increase, as well (otherwise, they remain unchanged).*

Proof. Notice first that the claim for the case $\mu \geq \rho^M$ follows the same argument based on convexity of $CS^M(\cdot)$ that we employed for the proof of Theorem 2.

Assume then that $\mu < \rho^M$: when $\omega = 1$ the platform sends message Y and belief ρ^M is induced, whereby one seller joins the market; when $\omega = 0$, with probability q_l^M message Y is sent and N otherwise, so that after N the induced belief is 0 and no seller enters. Under this the variation in social welfare can be written as

$$\begin{aligned}\Delta SW =&\mu\left[\int_{\hat{p}^M(\rho^M)}^{\infty} vdF + F\left(\hat{p}^M(\rho^M)\right)\int_{p^M}^{\infty} vdF + c - 2c\right]\\ &+ (1-\mu)\left[q_l^M\left(\int_{\hat{p}^M(\rho^M)}^{\infty} vdF + c - 2c\right) + (1-q_l^M)(2c-2c)\right]\\ =&\tau_l^M\left[\int_{\hat{p}^M(\rho^M)}^{\infty} vdF + \rho^M F\left(\hat{p}^M(\rho^M)\right)\int_{p^M}^{\infty} vdF - c\right]\\ =&\tau_l^M\left[\underbrace{\int_{\hat{p}^M(\rho^M)}^{\infty}\left(v-\hat{p}^M(\rho^M)\right)dF + \rho^M F\left(\hat{p}^M(\rho^M)\right)\int_{p^M}^{\infty}(v-p^M)dF}_{\text{consumer surplus}} + \hat{V}^M(\rho^M)\right] \geq 0,\end{aligned}$$

where $\tau_l^M = \frac{\mu}{\rho^M}$ from the proof of Theorem 1. Therefore, social welfare always increases.

Let us now consider aggregate sellers' profits: the variation is

$$\begin{aligned}\Delta\sum_{i=1}^{2}\Pi_{0,i} =&\mu\left[(1-\alpha)\hat{p}^M\left(\rho^M\right)\overline{F}\left(\hat{p}^M(\rho^M)\right)+(1-\alpha)F\left(\hat{p}^M(\rho^M)\right)\pi^M+c-2c\right]\\ &+(1-\mu)\left[q_l^M\left((1-\alpha)\hat{p}^M\left(\rho^M\right)\overline{F}\left(\hat{p}^M(\rho^M)\right)+c-2c\right)+\left(1-q_l^M\right)(2c-2c)\right]\\ =&\tau_l^M\left[W^M\left(\rho^M\right)-c\right]=0,\end{aligned}$$

where the last equality follows by definition of ρ^M. So aggregate sellers' profits remain the same for all prior beliefs. Finally, notice that without disclosure no seller would join, and therefore consumer surplus would be zero; under the optimal policy one seller enters and, since $\hat{p}^M(\rho^M)$ is smaller than the upper bound of the support of F, there is a non null probability that a sale will realize in the first period, so that consumer surplus under optimal disclosure is positive. Hence, the optimal policy always increases consumer surplus. We then conclude that the optimal policy increases all our metrics of welfare; notice that this conclusion holds even without the uniformity assumption. ■

EC 2 Noisy Signals

The setting we consider assumes that the platform is able to directly observe *in advance* the realized state of the world, and then condition the messages it sends to sellers on this observation. Our interpretation of this assumption is that the platform can generate a "perfectly" informative signal about the state of the demand, i.e., one that is it always correct, and then condition its messaging policy on it. In this appendix, we extend our model to allow for "noisy" signals and explore also how the accuracy of the platform's signal affects the policy it chooses to implement.

EC 2.1 Model

We consider a model akin to that of Section 2, but more general as far as the platform's signal is concerned: as before, there are two sellers who evaluate joining the platform to sell the single unit of a homogeneous good they are endowed with, and an uncertain number of buyers that have random valuations for the good. All assumptions set forth in Sections 2 and 3 are taken to hold also in the current setting, and the sequence of play after the platform's message is the same as well.

Let there be a set of states of the world $\Omega = \{0, 1\}$, with $\omega = 1$ denoting that the demand is high, i.e., that a second customer will look for the good on the platform in $t = 2$. The commonly shared prior probability on Ω is that $\mathbb{P}(\omega = 1) = \mu$. The platform cannot directly observe the realized ω; rather, it observes a partially accurate signal $\varphi \in \mathscr{F}$ about the state of the world and conditions its messaging policy on it. Formally, define the random variable φ taking values in $\{0, 1\}$ and such that

$$\varphi \mid \omega = \begin{cases} \omega & \text{with probability } a \\ 1 - \omega & \text{with probability } 1 - a \end{cases}$$

In words, conditional on the realization of the state of the world, φ is equal to the state of the world only with probability a (i.e., the signal is "accurate" only with some probability, which is taken to be an exogenous parameter and reflects the platform's ability to forecast its demand). Throughout this section we assume that $a \geq \frac{1}{2}$, with $a = \frac{1}{2}$ denoting a situation in which signals are not informative at all about the true state of the world, and $a = 1$ describing perfect ones; thus, this enriched model encompasses the one described in Section 2. Importantly, the accuracy level is common knowledge among the players of the game. The main difference we introduce is in the definition of *information disclosure* policies available to the platform, which can now only be conditioned on the signal. Formally, we denote an information disclosure policy as the mapping between the space of φ and the space of distributions over messages,

$$\mathcal{D} : \varphi \mapsto q(m \mid \varphi) \in \Delta(M).$$

The platform commits to a disclosure policy before observing the realization of the signal, and sellers are assumed to know that the messages are sent conditional on the signal and not the true demand.

Equilibrium The payoff functions of the agents participating in the game and the sequence of play are the same as in Section 2, and therefore we can apply the same equilibrium given in Definition 1, with one minor modification. In fact, for any history h_t with $t \geq 0$, the platform's belief is the posterior belief determined by Bayes' rule after observing the realization of the signal, while for sellers it is the posterior obtained by Bayes' rule after observing the message sent by the platform, but not the signal. In other words, φ is the platform's private information.

Similarly to before, the message sent by the platform to sellers induces the sellers to hold a new belief about the state of the demand, based on which they decide whether to join the platform. Hence, all results from Section 3 apply also in this context, and in particular those relating to the shape of the function $\hat{V}(\rho)$ that gives the expected profit of the platform at belief ρ. Therefore, we only need to characterize the optimal disclosure policy of the platform to be able to compare this setting with the one analyzed before.

EC 2.2 Optimal Information Disclosure

Suppose the platform adopts some information disclosure policy with message space M. Since the sellers know that each message $m \in M$ is sent to them according to the signal obtained by the platform, which is inaccurate in general, when they form their posterior belief they take into account this additional layer of uncertainty. As a result, not all beliefs in $[0, 1]$ can be induced, as the following lemma shows.

Lemma 1. *Take any information disclosure policy $(\mathcal{D}, M)$ and suppose the prior belief of demand being high is μ. Then the maximal posterior belief that can be induced is*

$$\mu_{\max}(\mu, a) = \frac{a\mu}{1 - \mu + a\,(2\mu - 1)}, \tag{EC.7}$$

and the minimal is

$$\mu_{\min}(\mu, a) = \frac{\mu(1 - a)}{a\,(1 - 2\mu) + \mu}. \tag{EC.8}$$

When $a = \frac{1}{2}$, i.e., the signal is not informative about the state of the world, Lemma 1 clarifies that the platform can never induce a posterior belief different from the prior, because the sellers know that any message based on uninformative signals cannot provide more information about the state of the world than what they already have; symmetrically, when $a = 1$ and foresight is perfect, $\mu_{\max} = 1$ and $\mu_{\min} = 0$, so that any belief can be reached as a posterior; finally, for information that is partially accurate, $\mu_{\min} > 0$ and $\mu_{\max} < 1$. Thus, Lemma 1 restricts the range of distributions over posteriors that the platform can induce, and their support: in particular, for any given prior μ the support of any distribution τ over posteriors must be a subset of the interval $[\mu_{\min}, \mu_{\max}]$, that depends on the value of the prior.

The restriction on the range of posterior beliefs brought by imperfect signals makes the analysis of the optimal disclosure policy more challenging, because the standard theory of information design does not allow for the possibility of imperfect observation.[4] Nevertheless, we can prove a characterization of the optimal disclosure policy similar to that of Kamenica and Gentzkow (2011), but specific to our setting.

Theorem 1 (Optimal policy characterization)**.** *Let there be given a prior μ and accuracy level a. The platform's optimal profit $V^*(\mu)$ can always be achieved with $|M| = 2$. Moreover,*

$$V^*(\mu) = \sup\left\{z : \ (\mu, z) \in \boldsymbol{cov}_{ltd}\left(\hat{V};\mu\right)\right\},$$

where $\boldsymbol{cov}_{ltd}\left(\hat{V};\mu\right) = \boldsymbol{cov}\left\{(\rho, v) \in \Gamma\left(\hat{V}\right) : \ \rho \in [\mu_{\min}(\mu, a),\ \mu_{\max}(\mu, a)]\right\}$*;* $\boldsymbol{cov}$ *denotes the convex hull of a set and* $\Gamma(\cdot)$ *is the graph of a function.*

More intuitively, Theorem 1 first makes clear that a binary message space is sufficient to achieve the optimal value, even in this more complicated setting; secondly, it characterizes the optimal profit of the platform in terms of what we call *sliding* concavification of $\hat{V}$. Essentially, the sliding concavification of $\hat{V}$ assigns to μ the value that one would obtain by computing the concavification of $\hat{V}$ at μ, with domain restricted to $[\mu_{\min}(\mu, a),\ \mu_{\max}(\mu, a)]$. Further applying this machinery, we obtain the following corollary that details the optimal information disclosure policy.

Corollary 1 (Platform's optimal policy)**.** *There exists an optimal policy $(\mathcal{D}^*, M^*)$ with $M^* = \{Y, N\}$. Assume that $a > \frac{1}{2}$ and let $q(m \mid \varphi)$ denote the probability that message m is sent after signal φ (we suppress the dependence on μ for readability). Then an optimal policy for the platform takes the following form:*

[4] There is little work that has analyzed the case of a privately informed sender, e.g. Hedlund (2017), whose framework does not apply to our setting.

(i) When $c > W^M(0)$, then

$$(q(Y \mid 1), q(Y \mid 0)) = \begin{cases} (1,0) & \text{for } \mu \text{ s.t. } \mu_{\max} \leq \rho^M(c,\alpha) \\ \left(q^1_{\mu_{\min},\rho^M}, q^0_{\mu_{\min},\rho^M}\right) & \text{for } \mu < \rho^M(c,\alpha) \text{ s.t. } \mu_{\min} < \rho^M < \mu_{\max} \\ (1,1) & \text{for } \mu = \rho^M(c,\alpha) \\ \left(q^1_{\rho^M,\mu_{\max}}, q^0_{\rho^M,\mu_{\max}}\right) & \text{for } \mu > \rho^M(c,\alpha) \text{ s.t. } \mu_{\min} < \rho^M < \mu_{\max} \\ (1,0) & \text{for } \mu \text{ s.t. } \mu_{\min} \geq \rho^M(c,\alpha) \end{cases}$$

(ii) When $W^D(1) \leq c \leq W^M(0)$, then

$$(q(Y \mid 1), q(Y \mid 0)) = \left(q^1_{\mu_{\min},\mu_{\max}}, q^0_{\mu_{\min},\mu_{\max}}\right)$$

(iii) Finally, when $c < W^D(1)$

$$(q(Y \mid 1), q(Y \mid 0)) = \begin{cases} (1,0) & \text{for } \mu \text{ s.t. } \mu_{\max} \leq \rho^D(c,\alpha) \\ \left(q^1_{\mu_{\min},\rho^D}, q^0_{\mu_{\min},\rho^D}\right) & \text{for } \mu < \rho^D(c,\alpha) \text{ s.t. } \mu_{\min} < \rho^D < \mu_{\max} \\ (1,1) & \text{for } \mu = \rho^D(c,\alpha) \\ \left(q^1_{\rho^D,\mu_{\max}}, q^0_{\rho^M,\mu_{\max}}\right) & \text{for } \mu > \rho^D(c,\alpha) \text{ s.t. } \mu_{\min} \leq \rho^D < \mu_{\max} \\ (1,1) & \text{for } \mu \text{ s.t. } \mu_{\min} > \rho^D(c,\alpha) \end{cases}$$

In the above, $\left(q^1_{x,y}, q^0_{x,y}\right)$ for $x < \mu < y$ are the unique solutions to the system

$$\begin{cases} \dfrac{\left[q^1_{x,y}a + q^0_{x,y}(1-a)\right]\mu}{\left[q^1_{x,y}a + q^0_{x,y}(1-a)\right]\mu + \left[q^1_{x,y}(1-a) + q^0_{x,y}a\right](1-\mu)} = y \\ \\ \dfrac{\left[\left(1-q^1_{x,y}\right)a + \left(1-q^0_{x,y}\right)(1-a)\right]\mu}{\left[\left(1-q^1_{x,y}\right)a + \left(1-q^0_{x,y}\right)(1-a)\right]\mu + \left[\left(1-q^1_{x,y}\right)(1-a) + \left(1-q^0_{x,y}\right)a\right](1-\mu)} = x \end{cases}$$

It is worthwhile to compare the policy of Corollary 1 with that of Theorem 1 in Section 4. With direct observability of the state of demand the platform is able to induce the competitive scenario it prefers, by discouraging or enticing entry. In the current setting this is no longer possible, because not all posterior beliefs can be obtained; as a result, taking as given the accuracy of the signals, there only exists a restricted range of prior beliefs for which the platform can alter the level competition compared to what would result without information disclosure. In greater detail, this occurs when the outside option is large ($c > W^M(0)$) and the prior belief is such that

$\mu_{\min} < \mu < \rho^M < \mu_{\max}$; and when the cost of entry is low ($c < W^D(1)$) and the prior belief is such that $\mu_{\min} < \rho^D < \mu < \mu_{\max}$. Intuitively, these are the cases where the range of feasible posterior beliefs includes the entry threshold ρ^M or ρ^D. However, in all these cases the platform's optimal information disclosure induces the type of competition that results from optimal disclosure in the case of perfectly informative signals: one of the sellers is nudged to join when none of them would not, and is discouraged from entry when both would join. Conveniently, in the limit where $a \to 1$, we recover exactly the policy from Theorem 1. Finally, notice that when the value of the outside option is low and the prior is such that $\mu_{\min} > \rho^D$, the platform's optimal policy is to send an uninformative message, which is never optimal under perfect foresight of the state of demand.

Platform's value for more accurate signals The discussion above stressed that the level of accuracy of the platform's signals is an important driver in its ability to induce a wider range of posterior beliefs, and thereby achieve higher profits. While it is clear that producing uninformative signals ($a = \frac{1}{2}$) is equivalent to not engaging in information disclosure at all, and that perfect foresight ($a = 1$) makes it possible to achieve the first-best of Section 4, it is also interesting to assess the marginal return from improving accuracy. In fact, this is the relevant metric for platforms to ascertain the return on investments in analytic capabilities able to generate the signals. Figure 1 depicts by how much would the platform's profits increase if it had a-precise signals instead of uninformative ones (or, equivalently, if it received no signal at all) , and we consider three values of accuracy ranging from $a = 0.6$ to $a = 0.9$. We concentrate on the case of low outside option because of its greater relevance for applications. We make a number of observations from the figure, which we can reconcile with the theory from Section 4.

First, for those prior beliefs for which information disclosure does not modify the number of entrants, we see quite a small gain from increasing the precision of the signal: this again points to the relevance of information disclosure first, and foremost, as a tool to modify the degree of competition on the markets. This intuition is further reinforced if one considers those prior beliefs for which information disclosure actually prevents entry of one of the sellers, at which the increase in profits that occurs can be as high as 50%.

Second, we observe that, contrary to Figure 4, the increase in profits presents discontinuities in the beliefs, and that these discontinuities occur at different beliefs for different accuracy levels. The jumps happen at those beliefs for which $\mu_{\min} = \rho^D$: intuitively, for all those priors for which the minimum attainable posterior is larger than the threshold ρ^D, the platform can no longer prevent entry using information disclosure, and therefore needs to settle for its second-best (which is not to disclose anything). This provides a new insight: increasing the precision of signals is particularly profitable for platforms that would otherwise be unable to persuade some of the sellers not to join.

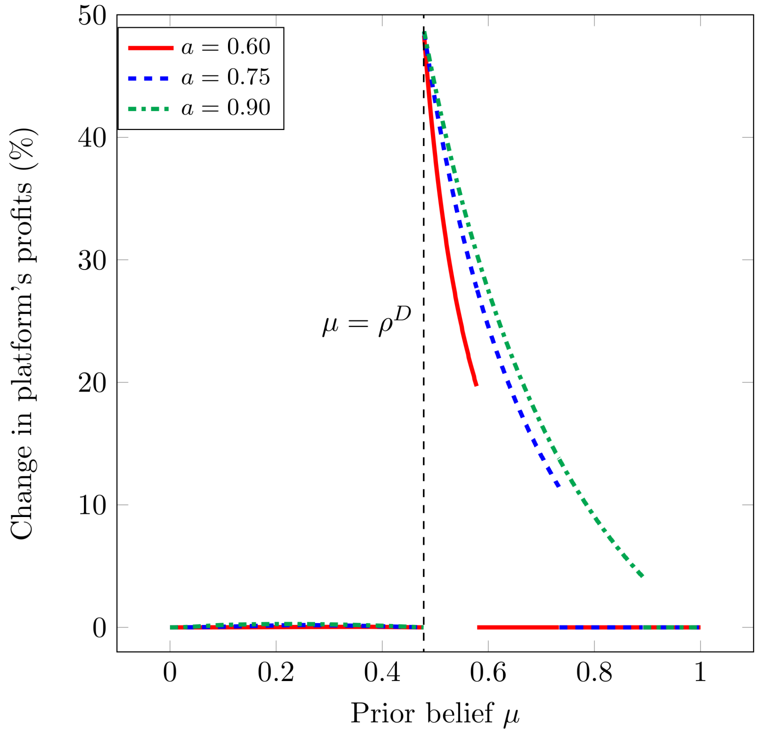

Figure 1: Percentage *increase* in platform's optimal profits from having a-precise signals instead of completely uninformative ones ($a = .5$). Results obtained with $v \sim U[0,1]$, $\alpha = 5\%$ and $c = 0.1$.

Finally, if we consider those prior beliefs for which increasing the precision of signals does not fundamentally alter the disclosure policy, we note that there are decreasing marginal returns from accuracy. This last intuition is made formal in the following lemma (for which we assume that the buyers' willingness to pay is uniformly distributed, for the sake of simplicity).

Lemma 2. *Suppose $c < W^D(1)$ and fix: a prior μ, and accuracy level $\overline{a}$ such that there exists $\epsilon > 0$ for which $\mu_{\min}(\mu, a) < \rho^D < \mu < \mu_{\max}$ holds for every $a \in (\overline{a} - \epsilon, \overline{a} + \epsilon)$. Let $V^*(a; \mu)$ denote the optimal platform's profit at prior μ when the accuracy level is a. Then*

$$\frac{\partial^2}{\partial a^2} V^*\left(\overline{a}; \mu\right) < 0,$$

i.e., the optimal profit is (locally) concave in the accuracy of the signals.

EC 2.3 Proofs EC 2

Lemma 1

Proof. Take a disclosure policy $(\mathcal{D}, M)$ with some arbitrary message space M, and without loss of generality assume $|M| \geq 2$; if not, it is clear that the only posterior consistent with Bayes' rule is the prior μ. Denote by $\mathscr{R}(\mathcal{D}, M)$ the range of posterior beliefs that can be achieved with this policy.

Take $m \in M$ and write $q(m \mid \varphi)$ for the probability that message m is sent conditional on the realized signal being φ; also without loss of generality, assume $q(m \mid \varphi) > 0$ for at least one realization of the signal, since otherwise the message has probability zero of being sent and the posterior belief cannot be determined. The sellers' posterior belief that $\{\omega = 1\}$ is then[5]

$$\begin{aligned}&\mathbb{P}\left(\omega = 1 \mid m\right)\\ &\quad= \frac{\left[q(m \mid \varphi = 1)a + q(m \mid \varphi = 0)(1-a)\right]\mu}{\left[q(m \mid \varphi = 1)a + q(m \mid \varphi = 0)(1-a)\right]\mu + \left[q(m \mid \varphi = 1)(1-a) + q(m \mid \varphi = 0)a\right](1-\mu)}.\end{aligned}$$

The posterior probability as a function of $\left(q(m \mid \varphi)\right)_{\varphi=0,1}$ is both quasi-convex and quasi-concave, because sub-(super-)level sets are half-spaces. In particular, it can be verified that it is maximized at $q(m \mid \varphi = 0) = 0$ and $q(m \mid \varphi = 1) \in (0, 1]$, and it achieves a value of

$$\frac{a\mu}{1 - \mu + a\left(2\mu - 1\right)}.$$

Similarly, the minimum is

$$\frac{\mu(1-a)}{a\left(1 - 2\mu\right) + \mu},$$

and is achieved at $q(m \mid \varphi = 0) \in (0, 1]$ and $q(m \mid \varphi = 1) = 0$. Finally, since $\mathbb{P}(\omega = 1 \mid m)$ is a continuous function, every posterior in $[\mu_{\min}, \mu_{\max}]$ can be achieved. This shows that $\mathscr{R}(\mathcal{D}, M) \subseteq [\mu_{\min}, \mu_{\max}]$ for any disclosure policy.

Take now a policy with $|M| = 2$, $M = \{m_1, m_2\}$. One immediately verifies that $\mathbb{P}(\omega = 1 \mid m_2)$ is minimized for $q(m_1 \mid \varphi = 1) = 1$ and $q(m_1 \mid \varphi = 0) \in [0, 1)$, and that the minimum is exactly $\mu_{\min}$. Likewise, $\mathbb{P}(\omega = 1 \mid m_1)$ attains its maximum of $\mu_{\max}$ at $q(m_1 \mid \varphi = 1) \in (0, 1]$ and $q(m_1 \mid \varphi = 0) = 0$. It follows that at $q(m_1 \mid \varphi = 1) = 1$ and $q(m_1 \mid \varphi = 0) = 0$ the posterior after m_1 is maximized and that after m_2 is minimized. Hence, for any policy with $|M| = 2$, $\mathscr{R}(\mathcal{D}, M) = [\mu_{\min}, \mu_{\max}]$, which also proves that for any policy with arbitrary message space the range of feasible beliefs must equal the interval $[\mu_{\min}, \mu_{\max}]$. ■

Theorem 1

Proof. Let $\mu \mid m$ be shorthand for $\mathbb{P}(\omega = 1 \mid m)$. The platform's optimization problem can be

[5] Implicitly assuming that $\mu \neq 0, 1$. These prior beliefs cannot be modified by any form of persuasion.

stated as:

$$\max_{M,\,(q(m|\varphi=1),\,q(m|\varphi=0))_{m\in M}} \mu \left[a \sum_{m\in M} q(m \mid \varphi=1)\hat{V}(\mu \mid m) + (1-a) \sum_{m\in M} q(m \mid \varphi=0)\hat{V}(\mu \mid m)\right] + (1-\mu)\left[a \sum_{m\in M} q(m \mid \varphi=0)\hat{V}(\mu \mid m) + (1-a) \sum_{m\in M} q(m \mid \varphi=1)\hat{V}(\mu \mid m)\right]$$

$$\text{s.t.} \sum_{m\in M} q(m \mid \varphi=i) = 1 \quad \forall i=0,1$$

$$q(m \mid \varphi=i) \geq 0 \quad \forall m \in M, \forall i=0,1$$

Since the objective is upper semi-continuous, there exists an optimal value $V^*(\mu)$ which is attained in the feasible region. Then there exist M^* and $(q^*(m \mid \varphi=i))_{i\in\{0,1\},m\in M}$ such that

$$V^*(\mu) = \mu \left[a \sum_{m\in M} q^*(m \mid \varphi=1)\hat{V}(\mu \mid m) + (1-a) \sum_{m\in M} q^*(m \mid \varphi=0)\hat{V}(\mu \mid m)\right] + (1-\mu)\left[a \sum_{m\in M} q^*(m \mid \varphi=0)\hat{V}(\mu \mid m) + (1-a) \sum_{m\in M} q^*(m \mid \varphi=1)\hat{V}(\mu \mid m)\right].$$

Denote by τ^* the distribution over posteriors induced by this policy. This distribution is Bayes plausible and such that $V^*(\mu) = \mathbb{E}_{\rho\sim\tau^*}\left[\hat{V}(\rho)\right]$. Moreover, by Lemma 1, $V^*(\mu) \in \mathtt{cov}\hat{V}\left([\mu_{\min}, \mu_{\max}]\right)$, the convex hull of the image of $[\mu_{\min}, \mu_{\max}]$ through $\hat{V}$. Hence, $(\mu, V^*(\mu)) \in \mathtt{cov}\left(hyp\left(\hat{V}\right)|_{[\mu_{\min},\mu_{\max}]}\right)$, the convex hull of the hypograph of $\hat{V}$ restricted to the interval of feasible posteriors. This is a connected subset of $\mathbb{R}^2$, and therefore by the Fenchel-Bunt theorem there exists $\overline{\tau}$ that satisfies the following: (i) $(\mu, V^*(\mu)) = \mathbb{E}_{\rho\sim\overline{\tau}}\left[(\rho, z(\rho))\right]$, with $(\rho, z(\rho)) \in \mathtt{cov}\left(hyp\left(\hat{V}\right)|_{[\mu_{\min},\mu_{\max}]}\right)$ so that $\overline{\tau}$ is Bayes plausible too; (ii) $supp(\overline{\tau}) \in [\mu_{\min}, \mu_{\max}]$ and $|supp(\overline{\tau})| \leq 2$. We now state a Lemma from the appendix of a working version of Kamenica and Gentzkow (2011), specialized to our setting, whose proof we also report for ease of reference.[6]

Lemma 3 (Kamenica and Gentzkow (2009))**.** *Given μ and $S \subset hyp(\hat{V}|_{[\mu_{\min},\mu_{\max}]})$, if $(\mu, V^*(\mu))$ is in the convex hull of S, it is also in the convex hull of the intersection of S and graph of $\hat{V}$ restricted to the feasible set of posteriors.*

Proof. We restrict to the case where $S = \{(\rho_1, z_1), (\rho_2, z_2)\}$ and suppose $(\mu, V^*(\mu)) = \gamma(\rho_1, z_1) + (1-\gamma)(\rho_2, z_2)$ for some $\gamma \in [0,1]$. Towards a contradiction, assume $z_1 < \hat{V}(\rho_1)$. Then we have

[6]This paper is available at https://www.wallis.rochester.edu/assets/pdf/wallisseminarseries/bayesianPersuasion.pdf

$V^*(\mu) = \gamma z_1 + (1-\gamma)z_2 < \gamma \hat{V}(\rho_1) + (1-\gamma)z_2$. But then $V^*(\mu)$ cannot be the optimal value because, as we will prove shortly, any Bayes plausible distribution with binary support in $[\mu_{\min}, \mu_{\max}]$ can be obtained by a disclosure policy. This is a contradiction and therefore $z_i = \hat{V}(\rho_i)$ for $i = 1, 2$. ■

It then follows that $V^*(\mu) = \mathbb{E}_{\rho\sim\overline{\tau}}\left[\hat{V}(\rho)\right]$, which proves that a distribution over posteriors with at most binary support is sufficient to achieve the optimum.

We now show that any Bayes plausible distribution τ with $supp(\tau) \in [\mu_{\min}, \mu_{\max}]$ and $|supp(\tau)| \leq 2$ can be obtained from some disclosure policy with $|M| = 2$. First, notice that in the case $supp(\tau) = \{\mu\}$ the claim is trivially true. Suppose then $supp(\tau) = \{\rho_1, \rho_2\}$, and without loss of generality $\rho_1 < \mu < \rho_2$ (since otherwise τ cannot be Bayes plausible). Let $M = \{m_1, m_2\}$. The linear system of equations

$$
\begin{cases}
q(m_1 \mid \varphi = 1)a + q(m_1 \mid \varphi = 0)(1-a) \\
= \dfrac{\rho_1}{\mu}\left\{\left[q(m_1 \mid \varphi = 1)a + q(m_1 \mid \varphi = 0)(1-a)\right]\mu + \left[q(m_1 \mid \varphi = 0)a + q(m_1 \mid \varphi = 1)(1-a)\right](1-\mu)\right\} \\
\\
q(m_2 \mid \varphi = 1)a + q(m_2 \mid \varphi = 0)(1-a) \\
= \dfrac{\rho_1}{\mu}\left\{\left[q(m_2 \mid \varphi = 1)a + q(m_2 \mid \varphi = 0)(1-a)\right]\mu + \left[q(m_2 \mid \varphi = 0)a + q(m_2 \mid \varphi = 1)(1-a)\right](1-\mu)\right\} \\
\\
q(m_1 \mid \varphi = i) + q(m_2 \mid \varphi = i) = 1 \quad \forall i = 0, 1 \\
q(m \mid \varphi = i) \geq 0 \quad \forall m \in M, \forall i = 0, 1
\end{cases}
$$

in the unknowns $q(m \mid \varphi = i)$ always has a solution, and therefore defines a Bayes plausible distribution with support $\{\rho_1, \rho_2\}$. However, there exists only one such distribution, so it must be that τ equals the distribution over posterior induced by the solutions to the system. In conclusion then, every Bayes plausible distribution over posteriors with binary support included in $[\mu_{\min}, \mu_{\max}]$ can be obtained by some disclosure policy with binary message space.

Given the equivalence between Bayes plausible distributions and disclosure policy we proved before, we can restate the platform's optimization problem as

$$
\begin{aligned}
&\max_{\tau\in\Delta([0,1])} \mathbb{E}_{\rho\sim\tau}\left[\hat{V}(\rho)\right] \\
&\text{s.t.} \quad \mathbb{E}_{\tau}[\rho] = \mu \\
&\qquad supp(\tau) = \{\mu_1, \mu_2\} \subset [\mu_{\min}, \mu_{\max}]
\end{aligned}
$$

Define now the convex hull of the graph of $\hat{V}$ restricted to the domain $[\mu_{\min}, \mu_{\max}]$ at μ as

$$\texttt{cov}_{\texttt{ltd}}\left(\hat{V};\mu\right) = \texttt{cov}\left\{(\rho, v) \in \Gamma\left(\hat{V}\right) : \rho \in [\mu_{\min}(\mu, a),\ \mu_{\max}(\mu, a)]\right\}$$

and notice that from all the previous parts of the proofs it follows that

$$V^*(\mu) = \sup\left\{z : \ (\mu, z) \in \texttt{cov}_{\texttt{ltd}}\left(\hat{V};\mu\right)\right\}$$

which finally proves the last claim and concludes the proof. ■

Corollary 1

Proof. The proof of the statement follows the same logic of the proof of Theorem 1 in Section 4, and entails verifying that the proposed policy gives the same value as the sliding concavification. As before, assume $c \leq W^M(1)$, since otherwise the platform's profit is always zero.

Suppose first that $c > W^M(0)$. When the prior is such that $\mu_{\max} < \rho^M$, any policy is optimal because irrespective of persuasion no seller join. If instead $\mu_{\max} = \rho^M$, it is optimal that the largest of the posteriors induced be ρ^M, since otherwise there would be no entry; let μ_2 be the smallest posterior belief induced: the value of such policy is

$$\frac{\mu - \mu_2}{\mu_{\max} - \mu_2}\hat{V}(\rho^M),$$

which is decreasing in μ_2, and therefore it is optimal to induce $\mu_{\min}$ as smallest posterior beliefs. One then verifies that the policy proposed induces exactly these two posteriors. When $\mu < \rho^M$ but $\mu_{\min} < \rho^M < \mu_{\max}$, notice that for any posteriors μ_1 and μ_2 it must be $\mu_2 < \mu$ and $\mu_1 > \rho^M > \mu$ (since otherwise the profits would be zero). Thus the value of this policy is

$$\frac{\mu - \mu_2}{\mu_1 - \mu_2}\hat{V}(\mu_1),$$

which is decreasing in μ_2, so that it is optimal to induce $\mu_{\min}$; moreover, this value is decreasing in μ_1 as long as $\mu_1 \geq \rho^M$, because

$$(\mu_{\min} - \mu_1)\frac{\partial \hat{V}(\mu_1)}{\partial \mu_1} + \hat{V}(\mu_1) < 0,$$

and therefore it is optimal to set $\mu_1 = \rho^M$. When $\mu = \rho^M$, it cannot be optimal to set the lowest posterior strictly below ρ^M, because otherwise (by the same reasoning as before) the optimum would be to set the largest posterior equal to ρ^M, which does not satisfy Bayes plausibility; hence, it is optimal to leave the prior unchanged, which is achieved by a policy that sends the same message with probability one irrespective of the signal realized. Finally, for $\mu > \rho^M$ the platform maximizes its profits by inducing the beliefs μ_1 and μ_2 such that $|\mu_1 - \mu_2|$ is maximal and there is entry at both: this follows from convexity of $\hat{V}$. Notice that the same reasoning carries out for $W^D(1) \leq c \leq W^M(0)$.

Suppose now that $c < W^D(1)$. As long as the prior is such that $\mu_{\max} \leq \rho^D$, convexity of $\hat{V}$ gives that the policy is the same as for $W^D(1) \leq c \leq W^M(0)$. For $\mu < \rho^D$ but $\mu_{\min} < \rho^D < \mu_{\max}$, it cannot be optimal to induce a posterior larger than ρ^D: this follows from Assumption 2 as in the proof of Theorem 1, where we ruled out full disclosure policies (which necessarily dominate any policy that induces posterior belief larger than ρ^D in this case). By convexity then, it is optimal to induce beliefs $\mu_{\min}$ and ρ^D. Always owing to Assumption 2, for $\mu = \rho^D$ it is optimal to leave the posterior unchanged. When $\mu > \rho^D$ but $\mu_{\min} \leq \rho^D$, it is clearly never optimal to induce a posterior less than ρ^D, which implies that the lower posterior belief must equal ρ^D; then Assumption 2 implies that the optimal upper posterior must be equal to $\mu_{\max}$. Finally, when $\mu_{\min} > \rho^D$, concavity of $\hat{V}$ yields that it is optimal to leave the prior belief unchanged. ■

Lemma 2

Proof. Under the hypotheses of the Lemma, the optimal disclosure policy is to induce belief $\mu_{\max}$ with probability $\frac{\mu-\rho^D}{\mu_{\max}-\rho^D}$ and ρ^D with complementary probability. Therefore,

$$V^*(\overline{a};\mu) = \frac{\mu-\rho^D}{\mu_{\max}-\rho^D}\hat{V}(\mu_{\max}) + \frac{\mu_{\max}-\mu}{\mu_{\max}-\rho^D}\hat{V}(\rho^D).$$

Recall that when the customers' willingness to pay is uniformly distributed it has

$$\rho^D = 2\left(1-\sqrt{1-\frac{4c}{1-\alpha}}\right)$$

and

$$\hat{V}(\rho^D) = \alpha\left(\frac{1}{2}+\frac{\rho^D}{8}\right)^2$$
$$\hat{V}(\mu_{\max}) = \alpha\frac{\mu_{\max}}{8}(4-\mu_{\max})$$

Moreover, it is assumed that there exists $\epsilon > 0$ for which $\mu_{\min}(\mu, a) < \rho^D < \mu < \mu_{\max}$ holds for every $a \in (\overline{a}-\epsilon, \overline{a}+\epsilon)$, which implies that the function $V^*(a;\mu)$ is twice differentiable in a neighbourhood of $\overline{a}$. But then it is a matter of algebra to show that

$$\frac{\partial^2 V^*(\overline{a};\mu)}{\partial a^2} = \alpha \frac{\partial^2}{\partial a^2}\left[\frac{\mu-\rho^D}{\mu_{\max}-\rho^D}\frac{\mu_{\max}}{8}(4-\mu_{\max}) + \frac{\mu_{\max}-\mu}{\mu_{\max}-\rho^D}\left(\frac{1}{2}+\frac{\rho^D}{8}\right)^2\right] < 0$$

■

EC 3 Revenue Share

In the description of our model we argued that the platform may find it impossible to tailor the revenue share it retains to different market conditions because of practical constraints; this justified taking α exogenous. In this appendix we explore the outcomes that would obtain if the platform could optimize also this quantity. In fact, a potential threat to our main results is, that leaving the platform the ability to choose α could lead to a "balancing" between the negative effects of optimal information disclosure and the positive effect of a low enough revenue share.

Suppose that the same model of Section 2 holds,[7] with the only difference that the platform chooses both the share of revenue it wants to retain and the information disclosure policy; sellers have the same action space as before. For the sake of tractability, we also assume that the buyers' willingness to pay is uniformly distributed on the unit interval. Formally, since for fixed α and $(\mathcal{D}, M)$ the game is the same as before and the sellers' optimal strategies do not change, we can write the platform's problem as

$$\max_{\alpha\in[0,1],\,(\mathcal{D},M)} \mathbb{E}_{\rho\sim\tau}\left[\hat{V}(\rho)\right],$$

where $\hat{V}(\cdot)$ is the function giving the expected profit of the platform when the induced belief is ρ; the expectation is taken with respect to the distribution τ over beliefs induced by the mechanism $(\mathcal{D}, M)$.[8] Notice that the optimal information disclosure policy derived in Section 4 is parametrized by $\alpha \in [0, 1]$, and so for each α we can identify the optimal $(\mathcal{D}, M)$; thus, we can write

$$\max_{\alpha\in[0,1],\,(\mathcal{D},M)} \mathbb{E}\left[\hat{V}(\rho)\right] = \max_{\alpha\in[0,1]} \max_{(\mathcal{D},M)} \mathbb{E}\left[\hat{V}(\rho)\right] = \max_{\alpha\in[0,1]} V^*(\mu),$$

where V^* is the concavification of $\hat{V}$, i.e. the optimal expected value of the policy to the platform given c and α. Without loss of generality we will assume $c \leq W^M(1)$; if it were larger, no one would ever join and the problem would be vacuous.

Notice that the platform faces a basic trade-off in this optimization problem. Since α concurs to determine how many sellers will join the market in equilibrium, as the opportunity cost decreases the platform should be able to retain an always larger revenue share, for fixed number of entrants. At the same time, we established that it is always more profitable to have a single seller on the market rather than two. The choice of the revenue share must then trade-off the incentive to appropriate a larger share of each transaction and the effects on competition and volumes of a larger α. The following proposition shows how this tension is resolved.

[7] We also revert to assuming that signals about demand are perfectly informative.

[8] A more formal definition can also be found in the proof of Theorem 1 in Appendix EC 1.

Proposition 1. *The optimal revenue share α^* takes value in the set*

$$\left\{1-\frac{64}{25}c,\, 1-\frac{c}{\left(\frac{1}{2}+\frac{\mu}{8}\right)^2},\, 1-4c,\, \overline{\alpha}\right\},$$

where $\overline{\alpha}$ solves

$$4\sqrt{\frac{c}{1-\alpha}}(1-\mu)(2-\alpha)=5(\mu+4)(1-\alpha).$$

The proof provides a more extensive description of the (rather involved) cases where each of these values are optimal. Notice, however, that the case where $\alpha^*=\overline{\alpha}$ is a "residual" one, in the sense that if c and μ were randomly and independently chosen from their respective ranges, the probability that this case would obtain is close to 0. Therefore, in what follows we will concentrate on the other, more relevant, instances. We highlight two main facts that follow from Proposition 1: (i) if α^* is chosen, there always occurs at most one entry; (ii) the optimal revenue share is decreasing in the opportunity cost. Hence, our previous intuition about the relation between α^* and the opportunity cost is confirmed.

The exact value of α^* depends both on the value of the opportunity cost and on the prior belief, and therefore requires the same precise knowledge of the market conditions of the optimal information disclosure policy of Section 4. However, optimally choosing the revenue share makes the disclosure policy straightforward to implement, because it degenerates to either full-disclosure or no-disclosure; Table 1 details this. The intuition is that, once it is established that it is optimal for the platform to have just one seller, the optimal α is the one that selects the best threshold belief $\rho^M(c,\alpha)$. This is because in the optimal disclosure with just one entrant, ρ^M is one of the induced beliefs. When the opportunity cost is large, it is better for the platform to retain a share that sets $\rho^M=1$, which lets one seller join only if two customers come, rather than trying to lower ρ^M so much that there would always be entry: the intuition is that to have the latter the revenue share would need to become too small. As the value of the outside option decreases, α is chosen (first) for ρ^M to exactly match the prior, and then to have $\rho^M=0$. In both cases one seller always joins, and the information disclosure policy moves from no-disclosure to full-disclosure. The simplicity of the ensuing information disclosure policy is a salient feature of this model, and points towards a direction we have hinted at in Section 2: adjusting the revenue share can be (partially) substituted with optimal information disclosure. When α is taken as given, the disclosure policy (i.e. the probability with which the true state is revealed) is made contingent on the market primitives; however, if the revenue share is allowed to be adjusted, then its optimal value subsumes also the effects of strategic information disclosure, so that the latter becomes straightforward.

The consequences on welfare of giving the platform an additional profit lever depend on which setting is taken as a benchmark. In fact, we briefly note that when α^* is chosen, the hypotheses of

Optimal α	**Induced policy**	**Entry**
$1-\dfrac{64}{25}c$	Full disclosure	$\rho^M=1 \Rightarrow$ One seller iff $\omega=1$
$1-\dfrac{c}{\left(\frac{1}{2}+\frac{\mu}{8}\right)^2}$	Uninformative	$\rho^M=\mu \Rightarrow$ One seller always in
$1-4c$	Full disclosure	$\rho^M=0 \Rightarrow$ One seller always in

Table 1: Optimal information disclosure when α^* is chosen.

Theorem 3 are satisfied; in turn, these imply that letting the platform apply the optimal information disclosure weakly increases consumer surplus, compared to the case where nothing is disclosed, but the same α^* is employed.[9] However, a more interesting benchmark is the setting in which the the platform sets neither the revenue share nor information disclosure, i.e. takes some α as given and reveals nothing. In this case, letting the platform use the additional lever makes buyers considerably worse off.

Theorem 2. *Let c, α and μ be given such that, with no-disclosure, at least one seller would join the platform. If $c>\frac{1}{36}$, or $c\leq\frac{1}{36}$ and $\mu\notin\left(\dfrac{1-36c}{14c+1},\overline{\mu}\right)$, then consumer surplus decreases when the platform optimally chooses both the revenue share and the disclosure policy.*

The belief $\overline{\mu}$ is defined in the proof of Proposition 1, and the the interval $\left(\dfrac{1-36c}{14c+1},\overline{\mu}\right)$ is the range of values of the prior for which, if the opportunity cost is very small, the optimal revenue share is $\overline{\alpha}$. Since no analytic results can be obtained for this case, one should interpret Theorem 2 as sufficient condition; indeed, numerical examples suggest that the claim also holds for most of the cases where $\alpha^*=\overline{\alpha}$.[10] The intuition behind the stark result is that when the platform is choosing α optimally, it makes sure that just one seller joins the market (and so makes buyers worse off compared to all cases where two would enter) and that she has all the information needed to set the highest price possible. As an illustrative example, consider what happens when $c>\dfrac{1}{4}$: in this case $\alpha^*=1-\frac{64}{25}c$ and the policy is of full disclosure, with a seller joining only when two customers are expected; under this, the first customer is either faced with the highest possible monopoly price or does not find a seller at all.

Overall, this section improved our understanding of the platform's market making abilities. Firstly, we can conclude that relaxing the assumption that α is exogenous does not alter, but strengthens, our insights from the previous sections: the platform can potentially severely harm the

[9] Notice that when α^* induces a no-disclosure policy, there is actually no change in consumer surplus.
[10] In any case, as remarked before, the range of values for which the theorem does not apply is very small.

buyers, and letting it choose the optimal revenue share only improves its ability to extract surplus to increase its profits. The fundamental mechanism through which the platform achieves this is its ability to modify the sellers' incentives to join the market. By selecting the most profitable competitive structure, the optimal revenue share basically incorporates all the gains that would come from optimal information disclosure, which is the reason why the optimal policy is so simple under α^*. In turn, this highlights that optimally choosing the revenue share and the information disclosure are, in a sense, substitutes. This reinforces our intuition that, provided α cannot be adapted to ever changing market conditions, information disclosure can be used in spite, thus constituting a flexible tool for managing market thickness.

EC 3.1 Proofs EC 3

Proposition 1

Proof. Before proceeding, let us specify the functional form many quantities of interest take in the uniform case. It has:

$$\rho^M(c,\alpha) = 4\left(2\sqrt{\frac{c}{1-\alpha}} - 1\right) \qquad \rho^D(c,\alpha) = 2\left(1 - \sqrt{1 - \frac{4c}{1-\alpha}}\right).$$

One can also check that $W^M(1) = (1-\alpha)\frac{25}{64}$, $W^M(0) = (1-\alpha)\frac{1}{4}$ and $W^D(1) = (1-\alpha)\frac{3}{16}$. Since α cannot be larger than 1, without loss of generality $c \leq \frac{25}{64}$. Thus, we have the following cases for V^*: when $\alpha > 1 - \frac{64}{25}c$, $V^*(\mu) = 0$ for all beliefs and we call it (for short) case (a); when $\alpha \in \left(1-4c, 1-\frac{64}{25}c\right] \cap [0,1]$, then

$$V^*(\mu) = \begin{cases} \dfrac{\alpha c}{1-\alpha}\dfrac{\mu}{\rho^M} & \text{if } \mu < \rho^M \\ \dfrac{\mu - \rho^M}{1-\rho^M}\left(\dfrac{25}{64}\alpha - \dfrac{\alpha c}{1-\alpha}\right) + \dfrac{\alpha c}{1-\alpha} & \text{if } \mu \geq \rho^M \end{cases}$$

and we denote it case (b); when $\alpha \in \left(1-\frac{16}{3}c, 1-4c\right] \cap [0,1]$, it obtains

$$V^*(\mu) = \frac{\alpha}{4} + \frac{9}{64}\alpha\mu,$$

termed case (c); finally, when $\alpha \in \left[0, 1-\frac{16}{3}c\right] \cap [0,1]$, we have

$$V^*(\mu) = \begin{cases} \dfrac{\alpha}{4} + \alpha\left[\left(\dfrac{1}{2} + \dfrac{\rho^D}{8}\right)^2 - \dfrac{1}{4}\right]\dfrac{\mu}{\rho^D} & \text{if } \mu \leq \rho^D \\ \alpha\left(\dfrac{1}{2} + \dfrac{\rho^D}{8}\right)^2 + \alpha\left[\dfrac{3}{8} - \left(\dfrac{1}{2} + \dfrac{\rho^D}{8}\right)^2\right]\dfrac{\mu - \rho^D}{1-\rho^D} & \text{if } \mu > \rho^D \end{cases}$$

which is denoted as case (d). When any of the intersections above is empty, V^* is not defined on that interval. Notice also that V^* as just defined is also a continuous function of α, so a maximum in $[0,1]$ exists.

Note that not all of the cases are feasible for every c. Indeed, when $\frac{1}{4} < c \leq \frac{25}{64}$, only case (b) is (partially) feasible: since $1-4c<0$, α ranges between 0 and $1-\frac{64}{25}c$; consequently, the smallest ρ^M that can be achieved is $8\sqrt{c}-4>0$. When $\frac{3}{16} < c \leq \frac{1}{4}$ case (b) is feasible, and case (c) is partially feasible. Finally, for $c \leq \frac{3}{16}$ both (b) and (c) are entirely feasible and (d) is partially feasible (it becomes entirely feasible only when $c=0$).

Let us notice that when $c=\frac{25}{64}$ then the only possible choice is $\alpha=0$ and therefore the problem is trivial. Without loss of generality, assume then $c<\frac{25}{64}$. We first show that, if feasible, it is never optimal for the platform to set α so low that case (d) would obtain. In fact, suppose $c \leq \frac{3}{16}$ and denote by $\alpha(c)$ some $\alpha \in \left(1-\frac{16}{3}c, 1-4c\right]$ and by $\alpha(d)$ some $\alpha \in \left[0, 1-\frac{16}{3}c\right]$. For fixed c, $\alpha(d)$ determines a threshold ρ^D. The expected profit for the platform without information disclosure under case (c) is

$$\hat{V}_{(c)}(\mu) = \alpha(c)\left[\hat{p}^M(\mu)\overline{F}\left(\hat{p}^M(\mu)\right) + \mu\pi^M F\left(\hat{p}^M(\mu)\right)\right],$$

and for case (d) with $\mu \leq \rho^D$ it has

$$\hat{V}_{(d)}(\mu) = \alpha(d)\left[\hat{p}^M(\mu)\overline{F}\left(\hat{p}^M(\mu)\right) + \mu\pi^M F\left(\hat{p}^M(\mu)\right)\right].$$

Since $\alpha(c) > \alpha(d)$, $\hat{V}_{(c)}(\mu) > \hat{V}_{(c)}(\mu)$ for $\mu \leq \rho^D$. Moreover, when $\mu > \rho^D$ so that two sellers join, we know that the platform gets strictly higher payoff when just one enters; therefore:

$$\alpha(d)\left[\hat{p}^M(\mu)\overline{F}\left(\hat{p}^M(\mu)\right) + \mu\pi^M F\left(\hat{p}^M(\mu)\right)\right] > \alpha(d) 2\mu\pi^M \overline{F}\left(\mu\pi^M\right) = \hat{V}_{(d)}(\mu),$$

and hence $\hat{V}_{(d)}(\mu) < \hat{V}_{(c)}(\mu)$ for $\mu > \rho^D$ as well. Recalling the definition of V^*, it follows $V^*_{(c)} \geq V^*_{(d)}$ for every $\alpha(c)$, $\alpha(d)$, $c \leq \frac{3}{16}$ and μ. Finally, with $c > \frac{3}{16}$ case (d) is not feasible and therefore we conclude that $\alpha \in \left[0, 1-\frac{16}{3}c\right]$ can never be optimal. We will now analyze each remaining "piece" of V^*, and then derive the optimal α. This is accomplished by first looking for the local maximum of each bit, and then using monotonicity to identify the global maximum.

Consider now the function

$$gb_l : \alpha \longmapsto \frac{\alpha c}{1-\alpha}\frac{\mu}{\rho^M}.$$

Since $c < \frac{25}{64}$, $\alpha \neq 1$ always. Moreover, this is the value of the revenue share α when $\mu < \rho^M$; using the definition of ρ^M one obtains

$$\mu < \rho^M \Longleftrightarrow \alpha > 1 - \frac{c}{\left(\frac{1}{2} + \frac{\mu}{8}\right)^2}.$$

For all μ in the unit interval we have

$$1 - 4c \leq 1 - \frac{c}{\left(\frac{1}{2} + \frac{\mu}{8}\right)^2},$$

and therefore $\rho^M \neq 0$ always, so that the function is always well defined. It can be checked that its first derivative is positive if and only if $c > \frac{1-\alpha}{(2-\alpha)^2}$, where the RHS reaches a maximum of $\frac{1}{4}$ at $\alpha = 0$. Hence for $c > \frac{1}{4}$ the function is strictly increasing on all its domain and its maximum is attained at $\alpha = 1 - \frac{64}{25}c$. Moreover, for $c \leq \frac{1}{4}$ the function is convex for $\alpha > 1 - 4c$, and consequently the maximum is found at either extreme of the domain. It has

$$gb_l\left(1 - \frac{c}{\left(\frac{1}{2} + \frac{\mu}{8}\right)^2}\right) \leq gb_l\left(1 - \frac{64}{25}c\right),$$

if and only if $c \leq \frac{15}{64}$, or $c > \frac{15}{64}$ and $\mu \geq 16(1 - 4c)$.

We now study the function

$$gb_u : \alpha \longmapsto \frac{\mu - \rho^M}{1 - \rho^M}\left(\frac{25}{64}\alpha + \frac{\alpha c}{1 - \alpha}\right) + \frac{\alpha c}{1 - \alpha}.$$

It can be checked that it is strictly concave for all $\mu \neq 1$. Since this is the value of α when $\mu \geq \rho^M$, its domain is the interval (provided c is small enough)

$$\left(1 - 4c, 1 - \frac{c}{\left(\frac{1}{2} + \frac{\mu}{8}\right)^2}\right].$$

The FOC for it to be maximized is

$$4\sqrt{\frac{c}{1 - \alpha}}(1 - \mu)(2 - \alpha) = 5(\mu + 4)(1 - \alpha). \tag{EC 3.1}$$

The first derivative evaluated at $\alpha = 1 - \frac{c}{\left(\frac{1}{2} + \frac{\mu}{8}\right)^2}$ is positive for all μ when $c > \frac{1}{36}$; when $c \leq \frac{1}{36}$ there exists $\overline{\mu}$ such that gb_u' is positive if $\mu \geq \overline{\mu}$. Such $\overline{\mu}$ is the unique root in $[0, 1]$ of the polynomial

$$\mu^3 + 7\mu^2 + (8 + 64c)\mu + 576c - 16,$$

with variable μ and parameter c. It follows that $gb'_u(1-4c) > 0$ for $c > \frac{1}{36}$, and when $c \leq \frac{1}{36}$ we have $gb'_u(1-4c) > 0$ if $\mu < \frac{1-36c}{14c+1} < \overline{\mu}$. Therefore, we have three possible cases: (i) when $gb'_u(1-4c) \leq 0$ the maximum is achieved at $\alpha = 1-4c$; (ii) if $gb'_u(1-4c) > 0$ and

$$gb'_u\left(1 - \frac{c}{\left(\frac{1}{2}+\frac{\mu}{8}\right)^2}\right) < 0,$$

then the function is maximized at the solution of (EC 3.1); finally (iii), if

$$gb'_u\left(1 - \frac{c}{\left(\frac{1}{2}+\frac{\mu}{8}\right)^2}\right) \geq 0,$$

the maximum is found at $\alpha = 1 - \frac{c}{\left(\frac{1}{2}+\frac{\mu}{8}\right)^2}$.

Finally, let us consider

$$gc : \alpha \longmapsto \alpha\left(\frac{1}{4} + \frac{9}{64}\mu\right),$$

which is maximized at $\alpha = 1-4c$, since it is linear in α.

Recall that the objective function we want to maximize is continuous in α. Putting together all previous observations, we arrive at the following cases.

1. $\frac{1}{4} < c \leq \frac{25}{64}$: only case (b) is feasible, and therefore from our discussion it must be that the platform wants $\rho^M > \mu$, with $\alpha^* = 1 - \dfrac{64}{25}c$.
2. $\frac{15}{64} < c \leq \frac{1}{4}$: only case (b) is feasible, but compared to the previous point gb_l is no longer monotone, while gb_u is increasing. Thus it has:
   (i) $\mu > 16(1-4c)$: $\alpha^* = 1 - \dfrac{64}{25}c$, because gb_L is convex and the maximum is attained at the right end of the domain.
   (ii) $\mu = 16(1-4c)$: $\alpha^* \in \left\{1 - \dfrac{64}{25}c, 1 - \dfrac{c}{\left(\frac{1}{2}+\frac{\mu}{8}\right)^2}\right\}$, because gb_l attains the same value at both ends of the domain.
   (iii) $\mu < 16(1-4c)$: $\alpha^* = 1 - \dfrac{c}{\left(\frac{1}{2}+\frac{\mu}{8}\right)^2}$, because gb_l is convex and the maximum is attained at the left end of the domain.
3. $\dfrac{1}{36} < c \leq \dfrac{15}{64}$: gb_u is increasing and gb_l is always smaller than gb_u, so $\alpha^* = 1 - \dfrac{c}{\left(\frac{1}{2}+\frac{\mu}{8}\right)^2}$.
4. $0 < c \leq \dfrac{1}{36}$: where the maximum is attained depends on the monotonicity of gb_u. From before, it has:

(i) $\mu \geq \overline{\mu}$: gb_u is increasing, and thus $\alpha^* = 1 - \dfrac{c}{\left(\frac{1}{2}+\frac{\mu}{8}\right)^2}$.

(ii) $\dfrac{1-36c}{14c+1} < \mu < \overline{\mu}$: gb_u admits an interior maximum, attained at the α^* solving equation (EC 3.1).

(iii) $\mu \leq \dfrac{1-36c}{14c+1}$: gb_u is decreasing and therefore $\alpha^* = 1-4c$.

This concludes the proof. ■

Theorem 2

Proof. We will start by recollecting from the previous proofs the value of expected consumer welfare when no optimization takes place; denote by α an exogenously given revenue share. If no seller joins, which happens for $\alpha > 1 - \frac{64}{25}c$, or $1-4c < \alpha \leq 1 - \frac{64}{25}c$ and $\mu < \rho^M(c,\alpha)$, the consumer surplus is 0. When only one seller enters, which occurs if $1-4c < \alpha \leq 1-\frac{64}{25}c$ and $\mu \geq \rho^M(c,\alpha)$, or $1-\frac{16}{3}c < c \leq 1-4c$, or $c \leq 1-\frac{16}{3}c$ and $\mu \geq \rho^D(c,\alpha)$, the expected consumer surplus is

$$\int_{\frac{1}{2}+\frac{\mu}{8}}^{1}\left(v-\frac{1}{2}-\frac{\mu}{8}\right)dv + \mu\left(\frac{1}{2}+\frac{\mu}{8}\right)\int_{\frac{1}{2}}^{1}\left(v-\frac{1}{2}\right)dv = \frac{1}{2}\left(\frac{1}{2}+\frac{\mu}{8}\right)^2 + \frac{\mu}{8}\left(\frac{\mu}{8}-\frac{1}{2}\right).$$

Finally, both sellers join only when $c \leq 1 - \dfrac{16}{3}c$ and $\mu > \rho^D(c,\alpha)$; the consumer surplus is given by

$$\int_{\frac{\mu}{4}}^{\infty}\left(v-\frac{\mu}{4}\right)dv + \mu\left(\left(1-\frac{\mu}{4}\right)\int_{\frac{1}{2}}^{\infty}\left(v-\frac{1}{2}\right)dv + \frac{\mu}{4}\int_0^{\infty} v dv\right) = \frac{\mu^2}{8} - \frac{\mu}{8} + \frac{1}{2}.$$

Let us first suppose that $\alpha^* = 1 - \frac{64}{25}c$, so that the induced policy is full-disclosure and in particular $\rho^M(c,\alpha^*) = 1$; this implies that one seller will enter when $\omega = 1$, and she is certain that a second customer comes, while no one joins when $\omega = 0$. Hence the expected consumer surplus induced by this policy is

$$\mu\left[\int_{\frac{5}{8}}^{1}\left(v-\frac{5}{8}\right)dv + \frac{5}{8}\int_{\frac{1}{2}}^{1}\left(v-\frac{1}{2}\right)dv\right] + (1-\mu)\times 0 = \frac{19}{128}\mu.$$

Algebra shows that

$$\frac{19}{128}\mu - \frac{1}{2}\left(\frac{1}{2}+\frac{\mu}{8}\right)^2 - \frac{\mu}{8}\left(\frac{\mu}{8}-\frac{1}{2}\right) \leq 0,$$

and

$$\frac{19}{128}\mu - \frac{\mu^2}{8} + \frac{\mu}{8} - \frac{1}{2} \leq 0,$$

for all $\mu \in [0,1]$.

Suppose now $\alpha^* = 1 - \frac{c}{\left(\frac{1}{2}+\frac{\mu}{8}\right)^2}$, which implies $\rho^M(c, \alpha^*) = \mu$ and therefore no-disclosure. Under this policy, one seller always enters, with a posterior belief equal to her prior. Thus consumer surplus is equal to

$$\frac{1}{2}\left(\frac{1}{2} + \frac{\mu}{8}\right)^2 + \frac{\mu}{8}\left(\frac{\mu}{8} - \frac{1}{2}\right) \geq 0,$$

and once again one verifies that

$$\frac{1}{2}\left(\frac{1}{2} + \frac{\mu}{8}\right)^2 + \frac{\mu}{8}\left(\frac{\mu}{8} - \frac{1}{2}\right) - \frac{1}{2}\left(\frac{1}{2} + \frac{\mu}{8}\right)^2 - \frac{\mu}{8}\left(\frac{\mu}{8} - \frac{1}{2}\right) = 0,$$

and

$$\frac{1}{2}\left(\frac{1}{2} + \frac{\mu}{8}\right)^2 + \frac{\mu}{8}\left(\frac{\mu}{8} - \frac{1}{2}\right) - \frac{\mu^2}{8} + \frac{\mu}{8} - \frac{1}{2} \leq 0,$$

for all $\mu \in [0, 1]$.

Finally, assume $\alpha^* = 1 - 4c$, so that $\rho^M(c, \alpha^*) = 0$ and the policy is again of full-disclosure. One seller always joins, which implies that the consumer surplus is

$$\mu\left[\int_{\frac{5}{8}}^{1}\left(v - \frac{5}{8}\right)dv + \frac{5}{8}\int_{\frac{1}{2}}^{1}\left(v - \frac{1}{2}\right)dv\right] + (1-\mu)\int_{\frac{1}{2}}^{1}\left(v - \frac{1}{2}\right)dv = \frac{3}{128}\mu + \frac{1}{8}.$$

Further algebra gives that

$$\frac{3}{128}\mu + \frac{1}{8} - \frac{1}{2}\left(\frac{1}{2} + \frac{\mu}{8}\right)^2 - \frac{\mu}{8}\left(\frac{\mu}{8} - \frac{1}{2}\right) \leq 0,$$

and

$$\frac{3}{128}\mu + \frac{1}{8} - \frac{\mu^2}{8} + \frac{\mu}{8} - \frac{1}{2} \leq 0,$$

for all $\mu \in [0, 1]$, which proves the claim. ■

EC 4 Market with Many Sellers

The model employed until Section 5 is stylized, but thanks to its simplicity we could neatly identify the essential aspects of our problem. We now seek to understand whether the insights developed also hold in a richer setting. To this end, we develop a model that allows for an arbitrary number of sellers and buyers to join the platform: our analysis will show that the previous intuitions carry over. However, this comes at the cost of a more contrived model, which will require some additional assumptions to be made tractable. We now describe its details, and postpone the discussion of the differences with the basic model to the end of the section.

EC 4.1 Model and equilibrium

The game we consider is dynamic, with three time periods, $t \in \{-1, 0, 1\}$: in period $t = -1$ the platform commits to the information disclosure policy; in $t = 0$, it observes the realization of the state of demand and sends a message to a countably infinite pool $\mathcal{S}$ of potentially differentiated sellers, who, in turn, make an entry decision. Finally, in $t = 1$, those of them who decided to join compete on prices. If sellers do not join the platform, they have an outside option valued at c. Analogously to the basic model, the platform generates its profits by retaining a fixed share of the value of the sales made by each participating seller. As in Section 2, we denote by h_t a generic history of the game in period t, with the convention that $h_{-1} = \varnothing$.

Buyers As before, we abstract away the behaviour of buyers, and take as primitive the demand function that arises from their (implicit) optimal decision. We use a demand system employed in Bimpikis, Crapis, and Tahbaz-Salehi (2019), which is in turn micro-founded in Myatt and Wallace (2015); this results from a continuum of buyer . When a subset $S \subset \mathcal{S}$ of sellers decide to join the platform, and each posts price $(p_j)_{j \in S}$, seller i faces a demand Q_i for her product given by

$$Q_i = \max\{0, \theta + (\phi - 1)P - \phi p_i\}, \tag{EC 4.1}$$

where $\theta > 0$, $P = \dfrac{1}{|S|}\sum_{j \in S} p_j$ is the average price posted by all participating sellers and $\phi > 1$. This demand system allows us to generally and economically model a market where many sellers sell (possibly) differentiated products. The parameter ϕ captures the degree of substitutability between goods supplied by different sellers: higher values of ϕ indicate more homogeneous markets. To model the idea that, when the number of participants in the market increases, the level of differentiation decreases[11], we assume that ϕ depends on the number $|S| = N$ of sellers that decide

[11] Instances of models where this phenomenon occurs are the classic Hotelling (linear city) or Salop (1979) competition games.

to join the market, so that $\phi = d(N) = N^\varepsilon$, $\varepsilon \geq 2$.

As discussed later, we include uncertainty about demand by assuming that θ is a random variable. A large realization of θ models an instance of "high" demand, while with a small realization we instantiate low demand in the market.

Sellers There is a countably infinite set of potential entrants, denoted by $\mathcal{S} = \{1, 2, \ldots\}$. We define the sellers' actions and payoffs backwards as in Section 2, with p_i the price posted by seller i and $E_i \in \{0, 1\}$ denoting her entry decision. Starting from $t = 1$, at any history h_1 such that subset $S \subset \mathcal{S}$ of seller decided to join the platform, for any $i \in S$ the payoff is

$$\Pi^S_{1,i}\left(\sigma^S_i; h_1\right) = (1-\alpha)\mathbb{E}_\rho\left[\mathbb{E}_{p_i \sim \sigma^S_i}\left[\max\left\{0, \theta + (\phi - 1)P - \phi p_i\right\} p_i \mid \theta\right]\right],$$

where ρ is the posterior belief at h_1. It is assumed that sellers get to observe how many of their competitors joined the platform before deciding on prices. Moreover, we suppose that each seller has an infinite inventory of her product, so that she can always entirely satisfy demand Q_i. Based on this continuation payoff, one defines also the utility of seller $i \in \mathcal{S}$ at time $t = 0$.

$$\Pi^S_{0,i}\left(\sigma^S_i; h_0\right) = \mathbb{E}_{E_i \sim \sigma^S_i}\left[\mathbf{1}\{E_i = 1\}\Pi^S_{1,i}\left(\sigma^S_i; \langle h_0, E_i = 1, E_{-i}\rangle\right) + c\mathbf{1}\{E_i = 0\}\right].$$

The entry decision is taken after having observed the message sent by the platform, which affects the expectation about θ and therefore influences $\Pi^S_{1,i}$.

Platform and information structure Demand is unknown at period $t = -1$, and this is modeled by assuming that θ in equation (EC 4.1) is a random variable, taking values in $\Theta = \{\theta_L, \theta_H\}$, with $\theta_H > \theta_L$. We assume $\theta_L > \frac{\theta_H}{2}$, i.e. that the realizations of demand are sufficiently similar; this is just to simplify exposition. The commonly shared prior probability that $\{\theta = \theta_H\}$ is denoted by μ as before. Consistently with Section 2, an information disclosure policy is a pair $(\mathcal{D}, M)$ such that

$$\mathcal{D} : \Theta \to \Delta(M).$$

Since the platform only moves at period $t = -1$, its payoff is[12]

$$\Pi^P\left((\mathcal{D}, M); h_{-1}\right) = \frac{\alpha}{1-\alpha}\mathbb{E}_{\theta\sim\mu}\left[\mathbb{E}_{m\sim\mathcal{D}(\theta)}\left[\sum_{i\in\mathcal{S}}\mathbf{1}\{E_i = 1\}\Pi^S_{0,i}\left(\sigma^S_i; \langle h_{-1}, (\mathcal{D}, M), m\rangle\right) \mid \theta\right]\right].$$

[12]Similarly to before, we can interpret this as saying that the platform receives perfectly informative signals about θ.

Under the rules of the game described in the previous paragraphs, the series above converges for all $c > 0$, since only finitely many of the sellers will join in equilibrium.

There are some important differences between this model and the one presented in Section 2. Besides the number of agents taking part in the transactions, these concern the timing of the game and the representation of demand; the two issues will be dealt together. Notice that simply expanding the number of buyers and sellers in the basic model, while leaving unchanged its dynamic pricing, gives rise to a much more difficult game, where existence and payoff equivalence of equilibria cannot be assured.[13] Therefore, in order to increase the number of market participants, and still retain a tractable model, we need to compress the time dimension. This also necessarily implies the different way in which we capture variability in demand.

EC 4.1.1 Equilibrium definition

The equilibrium concept we consider is completely analogous to that of the basic model: Sender-preferred perfect Bayesian equilibrium in pure strategies. Hence, the formal definition is very similar to that in Section 2.1.

Definition 1. *A collection of strategy-belief pairs $(\sigma_k, \gamma_k)_{k \in \mathcal{S} \cup \{P\}}$ is a Sender-preferred PBE if the following conditions are satisfied:*

(1) For every seller $i \in \mathcal{S}$, every time period $t \in \{0, 1\}$, and every history h_t, we have

$$\Pi^S_{t,i}\left(\sigma^S_i; h_t\right) \geq \Pi^S_{t,i}\left(\sigma'^S_i; h^t\right),$$

where σ'^S_i denotes a feasible strategy for seller i. Moreover, $\sigma^S_i(h_1) = \sigma^S_j(h_1)$ for all i, j such that $E_i = E_j = 1$ at $t = 0$

(2) For every agent k, $\gamma_k \mid \varnothing = \mu$, i.e., both sellers and the platform share a common prior μ. Moreover,

(i) For $i \in S$, $\gamma_i \mid h$ is determined by Bayes' rule after history h.

(ii) For the platform, $\gamma_P \mid h = \delta_{\{\theta\}}$ for all $h \neq \varnothing$, where θ is the realization of the state of the world.

(3) For fixed $(\mathcal{D}, M)$, whenever there exist multiple assessments such that all of the previous conditions hold, then $(\sigma_i, \gamma_i)_{i \in \mathcal{S}}$ yields the highest payoff for the platform.

[13] E.g. Martínez-De-Albéniz and Talluri (2011) conclude that with two sellers endowed with finite inventory, and facing multiple customers, there exists no equilibrium whenever buyers have random willingness to pay.

(4) Finally, $(\mathcal{D}, M)$ is the information disclosure policy that maximizes the platform's profits

$$\Pi^P\left((\mathcal{D}, M); h_{-1}\right) \geq \Pi^P\left((\mathcal{D}, M)'; h_{-1}\right),$$

assuming that sellers follow the strategies prescribed by the equilibrium.

The only substantial difference with the equilibrium in Section 2.1 is that here we impose that all entrants post the same price. This is required for simplicity reasons, and it allows to concentrate on the effects that information disclosure has on competition, since it eliminates possible sources of heterogeneity. Furthermore, it facilitates the comparison with the outcomes of the basic model.

EC 4.2 Equilibrium analysis

The identification of equilibrium decisions proceeds backwards, and therefore we begin with equilibrium play at period $t = 1$.

Proposition 2 (Pricing game)**.** *At every history h_1 when a subset S of sellers decide to join, and the belief is ρ, the only symmetric equilibrium in pure strategies gives*

$$p^* = \frac{\mathbb{E}_\rho[\theta]N}{N[d(N)+1]+1-d(N)}.$$

Each seller expects to earn

$$W_i(\rho, N, \alpha) = (1-\alpha)\left(\mathbb{E}_\rho[\theta]\right)^2 \frac{N\left[d(N)(N-1)+1\right]}{\left[d(N)(N-1)+N+1\right]^2},$$

and the platform's expected profits are

$$\hat{V}(\rho, N, \alpha) = \alpha\left(\mathbb{E}_\rho[\theta]\right)^2 \frac{N^2\left[d(N)(N-1)+1\right]}{\left[d(N)(N-1)+N+1\right]^2},$$

where $|S| = N$.

Further manipulations show the following facts: (i) holding ρ and α fixed, seller i's profits are strictly decreasing in N, and so are equilibrium prices; instead, (ii) $W_i(\rho, N, \alpha)$ increases in the belief for fixed N. We also note that $\hat{V}$ is convex in the belief ρ.

As an ideal counterpart of Proposition 3, the following result describes equilibrium entry decisions in this game.

Proposition 3 (Entry equilibrium)**.** *Given outside option value c and platform's commission α,* *define* $N_{max}(c,\alpha)$ *and* $N_{min}(c,\alpha)$ *as*

$$N_{max}(c,\alpha) = \max\{N : W_i(1,N,\alpha) \geq c\},$$
$$N_{min}(c,\alpha) = \max\{N : W_i(0,N,\alpha) \geq c\}.$$

For each $k \in \{N_{min}(c,\alpha)+1,\ldots,N_{max}(c,\alpha)\}$ *define the threshold beliefs* ρ_k *such that*

$$W_i(\rho_k,k,\alpha) = c.$$

Then at any history h_0 *such that the posterior belief is* ρ*:*

(i) if $\rho \leq \rho_{N_{min}+1}$, $N_{min}(c,\alpha)$ *sellers join;*

(ii) if $\rho_k < \rho \leq \rho_{k+1}$, k *sellers join (with* $k \in \{N_{min}(c,\alpha)+1,\ldots,N_{max}(c,\alpha)\}$*);*

(iii) if $\rho > \rho_{N_{max}}$, $N_{max}(c,\alpha)$ *sellers join.*

Proposition 3 shows that the pattern of entry at equilibrium in this model is not substantially different from that of the setting of Section 2: for a fixed value for the outside option and fees, as the belief about demand being high increases more sellers join. The main difference with Proposition 3 is that in this case there are infinitely many potential sellers, and therefore more than one of them may enter as the belief increases. The result is that the unit interval is partitioned in sub-intervals $(\rho_k,\rho_{k+1}]$, and for posterior beliefs within each interval, k sellers join the market. Finally, $N_{min}(c,\alpha)$ and $N_{max}(c,\alpha)$ describe the minimum and maximum number of entrants, respectively, that the value of the outside option and α allow in the market.

EC 4.3 Information disclosure

The information design problem of the platform exhibits now many more degrees of freedom, and the optimal information disclosure policy cannot be determined analytically for every combination of the parameters. The reason for the increased difficulty of the problem resides with the more complex nature of the entry pattern. In fact, in general it has $N_{max}(c,\alpha) - N_{min}(c,\alpha) > 1$, which implies that in principle all possible combinations of threshold beliefs ρ^k as potential posteriors to be induced should be evaluated. However, it is still possible to obtain a basic prediction of our basic model: optimal disclosure harms consumers, when taking as benchmark a policy of no-disclosure.

Theorem 3. *Suppose* $N_{min}(c,\alpha) < N_{max}(c,\alpha)$*. There exists* $\underline{N}$ *such that if* $N_{min}(c,\alpha) \geq \underline{N} \geq 2$ *and* $\hat{V}(1,N_{max}(c,\alpha)) < \hat{V}\left(\rho^{N_{max}},N_{max}(c,\alpha)-1\right)$*, then consumer surplus decreases under the optimal information disclosure policy for at least all* $\mu \geq \rho^{N_{max}}$.

It should be noted that the reduction in the welfare of consumers is caused by the same mechanism at work in the basic model. For $\mu \geq \rho^{N_{max}}$ the platform's optimal policy will always be one that induces beliefs 1 or $\rho^{N_{max}}$: when the latter obtains, the number of sellers in the market decreases compared to the benchmark, and this drives prices up and volumes down. Thus, we recover another base prediction of our simpler model, that the platform alters the competitive structure to increase its profits, and this damages consumers.

Theorem 3 provides a sufficient condition for buyers to be harmed by the disclosure from the platform, but it is not necessary. As such, it gives a lower bound on the instances where consumer surplus decreases. Indeed, it is not difficult to find examples where the reduction is consumer surplus is much more widespread; however, this can only be ascertained by first determining the optimal information disclosure policy, which in turn requires fixing some values for the parameters. Next we explore a numerical example where consumer surplus decreases for almost all prior beliefs.

EC 4.3.1 Numerical example

The example described here is intended to show that the implications of employing optimal information disclosure can be extended beyond the lower bound presented in Theorem 3. At the same time, the choice of parameters instantiates a fairly "standard" setting, whose main takeaways can be recovered also with other parameter values.

We fix the following values for the primitives: $\varepsilon = 2$, $\theta_H = 10$, $\theta_L = 5$, $c = 2.9$ and $\alpha = 5\%$. These imply that $N_{min} = 2$ and $N_{max} = 6$. The optimal information disclosure that ensues takes the following form.[14]

Optimal policy An optimal policy has message space $M = \{Y, N\}$. Write $\mathcal{D}(\theta)$ for the probability that message Y is sent when the state of the world is θ. Then,

$$\mathcal{D}(\theta_H) = \begin{cases} 1 & \text{for } \mu \leq \rho^{N_{min}+1} \\ q_k^H & \text{for } \rho^k < \mu \leq \rho^{k+1},\ k \in \{N_{min}+1, \ldots, N_{max}\} \\ \dfrac{\mu - \rho^{N_{max}}}{\mu\,(1-\rho^{N_{max}})} & \text{for } \mu > \rho^{N_{max}} \end{cases}$$

[14] A formal justification of why the policy takes this form can be found at the end of the proofs section of this appendix.

and

$$\mathcal{D}(\theta_L) = \begin{cases} \dfrac{\mu\left(1-\rho^{N_{min}+1}\right)}{\rho^{N_{min}+1}(1-\mu)} & \text{for } \mu \leq \rho^{N_{min}+1} \\ q_k^L & \text{for } \rho^k < \mu \leq \rho^{k+1},\ k \in \{N_{min}+1, \dots, N_{max}\} \\ 0 & \text{for } \mu > \rho^{N_{max}} \end{cases}$$

where q_k^L and q_k^H are the unique solutions to the systems

$$\begin{cases} \dfrac{q_k^H \mu}{q_k^H \mu + q_k^L (1-\mu)} = \rho^{k+1} \\ \dfrac{\left(1-q_k^H\right)\mu}{\left(1-q_k^H\right)\mu + \left(1-q_k^L\right)(1-\mu)} = \rho^k \end{cases}$$

for $k \in \{N_{min}+1, \dots, N_{max}\}$.

A verbal description of the policy clarifies that the optimal behaviour for the platform is to identify the threshold beliefs which are immediately larger and smaller than the prior μ, and to obfuscate information to induce either of these two as posterior beliefs. Since at the lower belief the incentives to entry are different, the outcome of the policy is to sometimes inducing one of the sellers that would have joined to stay out of the market. When the prior belief is below the threshold that would let $N_{min}+1$ sellers join, the possible posterior beliefs are 0 and $\rho^{N_{min}+1}$, from which follows that the number of participating sellers is not affected. Drawing a comparison with the optimal policy in Section 4, the case of $\mu \leq \rho^{N_{min}+1}$ closely mirrors the case where the value of the outside option is low, i.e., $c < W^D(1)$, with $\mu \leq \rho^D$. On the other hand, the outcomes of the policy when $\mu > \rho^{N_{min}+1}$ can be likened what occurs in our baseline model with low value of outside option and the prior exceeds the threshold ρ^D.

Figure 2 represents the welfare effects of employing the optimal information disclosure policy. From Figure 2a one observes that the optimal policy can lower consumer surplus by as much as 50%. At the same time, it is worth noting the similarity between Figure 2a and the left panel of Figure 5, which strengthens our previous parallelism with the baseline model. The increase in the platform's profits is large as well, as shown in Figure 2b, which also includes a comparison with the full-disclosure policy. The platform earns its additional profits at the expense of buyers: it exploits the fact that it can inflate prices by alternatively increasing the belief held by sellers or by restricting the number of entrants. The comparison with full-disclosure serves to confirm our intuition from Section 4, that fully revealing information may yield even worse outcomes than no-disclosure. Finally, we also obtain that social welfare is lower when the platform employs the optimal information disclosure policy, as given in Figure 2d.

In general, the results of this numerical exercise showed that it possible to retrieve our predictions

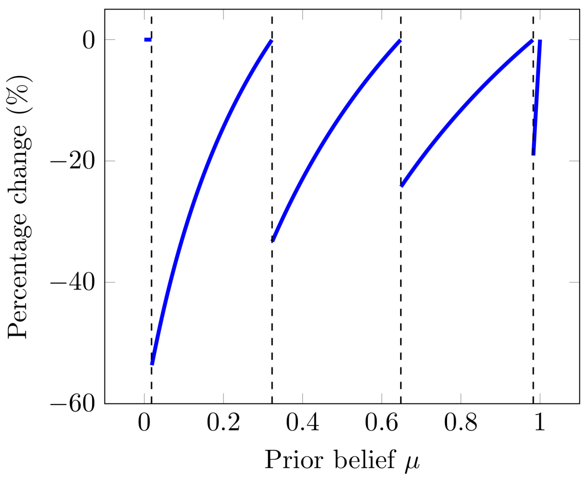

(a) Consumer surplus

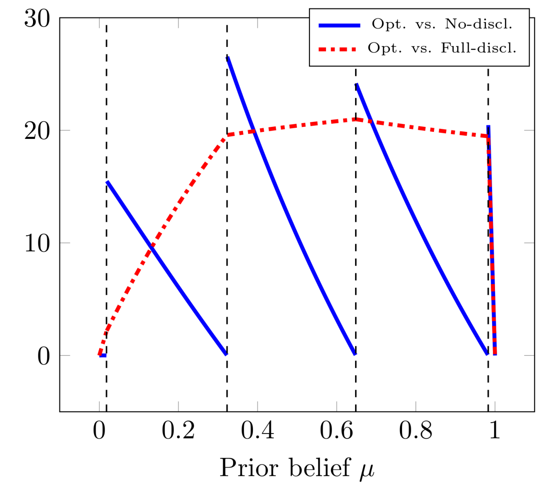

(b) Platform's profits

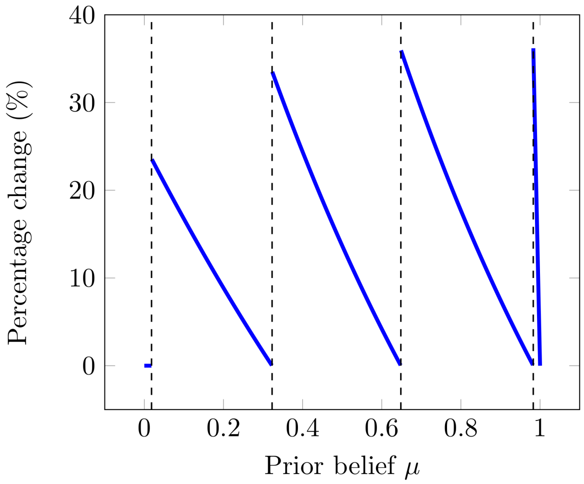

(c) Aggregate profits for sellers

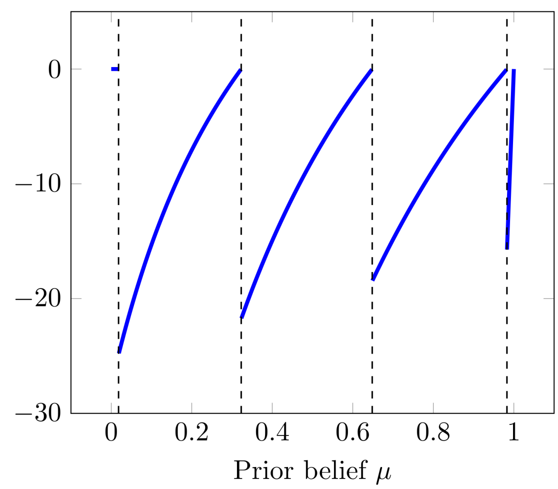

(d) Social welfare

Figure 2: Percentage change due to the platform employing the optimal disclosure policy instead of no-disclosure. Additionally, 2b illustrates the increase in profits with respect to full-disclosure policy. Dashed lines are drawn at the threshold beliefs ρ^3, ρ^4, ρ^5 and ρ^6. Benchmark profits and social welfare are derived with a total number of sellers equal to 10.

also from a more realistic model, that allows for additional customers and sellers. Results of the same type can be obtained by choosing different values for the parameters and selecting different functional forms for $\phi = d(N)$.

EC 4.4 Proofs EC 4

Proposition 2

Proof. Suppose first $N = 1$, in which case the demand faced by the single entrant reduces to $Q = \theta - p_i$, since $P = p_i$. The seller wishes to maximize expected profit and therefore the equilibrium price must solve

$$\max_p (1-\alpha)\mathbb{E}_\rho[p(\theta - p)].$$

Taking the FOC gives that $p^* = \frac{\mathbb{E}_\rho[\theta]}{2}$ and, since we are assuming $\theta_L > \frac{\theta_H}{2}$, the Seller faces positive demand even when the realized intercept is θ_L. Hence the expected profit for Seller i is $\pi_i(\rho, 1, \alpha) = \frac{(\mathbb{E}_\rho[\theta])^2}{4}$.

Suppose now $N > 1$, so that the optimization problem of seller i becomes

$$\max_{p_i} \quad (1-\alpha)\mathbb{E}_\rho\left[p_i\left(\theta + (d(N)-1)\frac{1}{N}\sum_{j=1}^{N} p_j - d(N)p_i\right)\right]. \tag{EC 4.2}$$

Taking the FOC gives

$$\mathbb{E}_\rho[\theta] - 2p_i\frac{(N-1)d(N)+1}{N} + \frac{d(N)-1}{N}\sum_{j\neq i} p_j = 0.$$

Imposing $p_j = p_i$ for every i, j and solving the equation yields

$$p_i^* = \frac{\mathbb{E}_\rho[\theta]N}{N(d(N)+1)+1-d(N)}.$$

At this point, simple substitution of p_i^* into (EC 4.2) gives the expected profit of a single seller. Finally, we can obtain the platform's profit by summing over the N expected profits of the sellers multiplied by the revenue share of the platform. ■

Proposition 3

Proof. Under our assumptions, $W_i(\rho, N, \alpha) \to 0$ for all ρ as $N \to \infty$. Also recall that W_i is increasing in ρ, which implies that the maximum possible profit for a seller obtains when $\rho = 1$ and $N = 1$, while the minimum is zero. Without loss of generality, assume $W_i(1, 1, \alpha) > c$, since otherwise there is never entry. Then $N_{max}(c,\alpha)$ and $N_{min}(c,\alpha)$ are well-defined as

$$N_{max}(c,\alpha) = \max\{N: \ W_i(1,N,\alpha) \geq c\},$$
$$N_{min}(c,\alpha) = \max\{N: \ W_i(0,N,\alpha) \geq c\}.$$

For each $k \in \{N_{min}(c,\alpha)+1,\dots,N_{max}(c,\alpha)\}$ define beliefs ρ_k such that

$$W_i(\rho_k,k,\alpha) = c.$$

Under our assumption, each of these equations has a unique solution: it represents the minimum belief such that k sellers would join the platform. In fact, take posterior belief ρ such that $\rho_k < \rho < \rho_{k+1}$: if $k-1$ sellers joined, then one of the sellers that stays out would have incentive to deviate and enter, because $W_i(\rho,k,\alpha) > c$; similarly, if $k+1$ joined then one of the seller in the market would deviate and stay out, because $W_i(\rho,k+1,\alpha) < c$. Hence, exactly k seller join at equilibrium. To establish entry for $\rho = \rho_k$, we evaluate $\hat{V}(\rho_k,k,\alpha)$ and $\hat{V}(\rho_k,k-1,\alpha)$. It has

$$\hat{V}(\rho_k,k-1,\alpha) \geq \hat{V}(\rho_k,k,\alpha \iff \frac{(k-1)^2\,[d(k-1)(k-2)+1]}{[d(k)(k-2)+k]^2} \geq \frac{k^2\,[d(k)(k-1)+1]}{[d(k)(k-1)+k+1]^2}, \tag{EC 4.3}$$

where we remind that $d(k) = k^\varepsilon$ for some $\varepsilon \geq 2$. Additional algebra shows that (EC 4.3) is always satisfied for every $k > 2$. Hence, we conclude that for $k = 1, 2$, exactly k sellers join at ρ_k, while for $k \geq 3$, $k-1$ sellers join at ρ_k. ■

Theorem 3

Proof. We start by noting that for any optimal policy, the only posterior beliefs that are induced can be $\rho = 0$, $\rho = 1$, or the threshold beliefs ρ^k, for some $k \in \left\{\rho^{N_{min}+1},\dots,\rho^{N_{max}}\right\}$ (and possibly all of them). This is because the function $\hat{V}$ is locally increasing and convex in the belief, which by the usual concavification argument implies that it is always possible to improve over a policy that induces posterior $\rho \neq 0,\, 1,\, \rho^k$. Observe now that at belief ρ^k exactly $k-1$ sellers join by the Sender-preferred condition of equilibrium. Therefore, at ρ^k the platform's profits reach a local maximum, and then jump downwards. We have

$$\hat{V}(\rho^k,k-1,\alpha) = \alpha\mathbb{E}_{\rho^k}[\theta]^2\frac{k-1}{h(k-1)} = c\frac{\alpha}{1-\alpha}(k-1)\frac{h(k)}{h(k-1)},$$

where $h(k) = \dfrac{[d(k)(k-1)+k+1]^2}{k\,[d(k)(k-1)+1]}$. It can be shown that $(k-1)\dfrac{h(k)}{h(k-1)}$ is eventually increasing in the number of entrants; denote by $\underline{N}$ the number, which depends on ε only, such that for all $k \geq \underline{N}$, $(k-1)\dfrac{h(k)}{h(k-1)}$ is increasing. It follows that, under our hypotheses, there is a sequence of increasing local maxima at ρ^k, with $\hat{V}\left(\rho^{N_{max}}, N_{max}-1, \alpha\right)$ being also the global maximum of $\hat{V}$ in $[0,1]$. Hence, it must be that at an optimal disclosure policy, for prior $\mu \in \left(\rho^{N_{max}}, 1\right]$, the platform mixes between beliefs $\rho^{N_{max}}$ and 1.

Consider now expected consumer surplus: since at equilibrium all sellers quote the same price, the expected consumer surplus of the customers acquiring the product from one seller amounts to $\frac{1}{2}\mathbb{E}_\rho\left[(\theta - p^*)^2\right]$. Hence, when there are k of them on the platform, the total expected consumer surplus is

$$CS(\rho, k) = \frac{k}{2}\mathbb{E}_\rho\left[\left(\theta - \frac{\mathbb{E}_\rho[\theta]k}{d(k)(k-1)+k+1}\right)^2\right].$$

All else equal, increasing the number of entrants increases consumer welfare, both because the price paid decreases and because a larger demand can be satisfied. Furthermore, algebra proves that $CS(\rho, k)$ is concave in ρ for k fixed, and increasing. Hence, expected consumer surplus is an increasing, piecewise concave function over $[0,1]$, with discontinuities at each threshold belief ρ^k. In particular, it has

$$CS(\rho^{N_{max}}, N_{max}-1) < \lim_{\rho \downarrow \rho^{N_{max}}} CS(\rho, N_{max}) \leq CS(1, N_{max}).$$

Since at any optimal disclosure policy the platform mixes between beliefs $\rho^{N_{max}}$ and 1, under the optimal policy the expected consumer surplus is

$$\tau CS\left(1, N_{max}\right) + (1-\tau)CS\left(\rho^{N_{max}}, N_{max}-1\right),$$

with τ determined by the optimal policy. Concavity of consumer surplus and the previous inequality then yield that under the optimal policy consumer welfare must be lower than under the no-disclosure benchmark. ■

Derivation of the optimal policy in the numerical example As in the proof of Theorem 1, we want to identify the concavification V^* of $\hat{V}$ and then deduce the optimal policy from it. The concavification of $\hat{V}$ can also be defined as the smallest concave function that is everywhere weakly larger than $\hat{V}$.

Notice that the beliefs $\{0, \rho^{N_{min}+1}, \ldots, \rho^{N_{max}}, 1\}$ partition the unit interval and, within each cell of the partition, $\hat{V}$ is a convex function of the prior belief μ. Thus, consider the set B given by

$$B = \left\{\left(\rho^k, \hat{V}(\rho^k, k-1, \alpha)\right) : \, k \in \{N_{min}+1, \ldots, N_{max}\}\right\} \cup \left\{\left(0, \hat{V}(0, N_{min}, \alpha)\right)\right\} \cup \left\{\left(1, \hat{V}(1, N_{max}, \alpha)\right)\right\}$$

and the piecewise affine function defined as

$$\widetilde{V}(\mu) = \begin{cases} \dfrac{\hat{V}(\rho^{N_{min}+1}, N_{min}, \alpha) - \hat{V}(0, N_{min}, \alpha)}{\rho^{N_{min}+1}}\mu + \hat{V}(0, N_{min}, \alpha) & \mu \leq \rho^{N_{min}+1} \\ \dfrac{\hat{V}(\rho^{k+1}, k, \alpha) - \hat{V}(\rho^{k}, k-1, \alpha)}{\rho^{k+1} - \rho^{k}}\mu + \hat{V}(\rho^{k}, k-1, \alpha) & \rho^{k} < \mu \leq \rho^{k+1} \\ \dfrac{\hat{V}(1, N_{max}, \alpha) - \hat{V}(\rho^{N_{max}}, N_{max}-1, \alpha)}{1 - \rho^{N_{max}}}\mu + \hat{V}(\rho^{N_{max}}, N_{max}-1, \alpha) & \mu > \rho^{N_{max}} \end{cases}$$

$k \in \{N_{min}+1, \ldots, N_{max}\}$, whose graph passes through every point in B. $\widetilde{V}$ is concave because under our choice of parameters the slope of each affine bit decreases as μ moves from 0 to 1; hence, B is a subset of the set of extreme points of the convex hull of the graph of $\hat{V}$. As a consequence, $V^*(\mu) = \hat{V}(\mu, N(\mu), \alpha)$ whenever $\left(\mu, \hat{V}(\mu, N(\mu), \alpha)\right) \in B$, by definition of extreme point.[15] Take now any other function l satisfying this requirement and such that $l(\mu) < \widetilde{V}(\mu)$ for some μ. Clearly $\mu \notin \{0, N_{min}+1, \ldots, N_{max}, 1\}$, but since $\widetilde{V}$ is affine, l must be non-concave in the cell of the partition containing μ. Thus, we cannot find another concave function that is smaller than $\widetilde{V}$ and that passes through all the points in B, which implies that $\widetilde{V} = V^*$. At this point, following arguments identical to those employed in Theorem 1 it follows that the given policy delivers a value equal to V^*.

[15] Here $N(\mu)$ denotes the number of entrants when the belief is μ, determined according to our equilibrium rule.

References

Anunrojwong, Jerry, Krishnamurthy Iyer, David Lingenbrink. 2020. Persuading risk-conscious agents: A geometric approach. *Available at SSRN 3386273* .

Aumann, Robert J, Michael Maschler. 1995. *Repeated games with incomplete information*. MIT press.

Bimpikis, Kostas, Davide Crapis, Alireza Tahbaz-Salehi. 2019. Information sale and competition. *Management Science* **65**(6) 2646–2664.

Hedlund, Jonas. 2017. Bayesian persuasion by a privately informed sender. *Journal of Economic Theory* **167** 229–268.

Kamenica, Emir, Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* **101**(6) 2590–2615. doi:10.1257/aer.101.6.2590.

Lariviere, Martin A. 2006. A note on probability distributions with increasing generalized failure rates. *Operations Research* **54**(3) 602–604. doi:10.1287/opre.1060.0282.

Martínez-De-Albéniz, Victor, Kalyan Talluri. 2011. Dynamic Price Competition with Fixed Capacities. *Management Science* **57**(6) 1078–1093. doi:10.1287/mnsc.lll0.1337.

Myatt, David P, Chris Wallace. 2015. Cournot competition and the social value of information. *Journal of Economic Theory* **158** 466–506.

Salop, Steven C. 1979. Monopolistic Competition with Outside Goods. *The Bell Journal of Economics* **10**(1) 141. doi:10.2307/3003323.