

Pavan, Alessandro; Jullien, Bruno

Working Paper

Platform pricing under dispersed information

CSIO Working Paper, No. 0119

Provided in Cooperation with:

Department of Economics - Center for the Study of Industrial Organization (CSIO), Northwestern University

Suggested Citation: Pavan, Alessandro; Jullien, Bruno (2013) : Platform pricing under dispersed information, CSIO Working Paper, No. 0119, Northwestern University, Center for the Study of Industrial Organization (CSIO), Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/92478>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

THE CENTER FOR THE STUDY
OF INDUSTRIAL ORGANIZATION
AT NORTHWESTERN UNIVERSITY

Working Paper #0119

Platform Pricing under Dispersed
Information

By

Alessandro Pavan
Northwestern University

and

Bruno Jullien
Toulouse School of Economics
(CNRS-GREMAQ and IDEI)

August 2013

Platform Pricing under Dispersed Information^{*}

Bruno Jullien

Alessandro Pavan

Toulouse School of Economics (CNRS-GREMAQ and IDEI)

Northwestern University

August 2013

Abstract

We study monopoly and duopoly pricing in a two-sided market with dispersed information about users' preferences. We first show how the dispersion of information introduces idiosyncratic uncertainty about participation rates and how the latter shapes the elasticity of the demands and thereby the equilibrium prices. We then study informative advertising campaigns and product design affecting the agents' ability to estimate their own valuations and/or the distribution of valuations on the other side of the market.

JEL classification: D82

Keywords: two-sided markets, dispersed information, platform competition, global-games, informative advertising

^{*}An earlier version circulated under the title "Platform Competition under Dispersed Information." For useful comments and suggestions, we thank Mark Armstrong, Marco Ottaviani, Michael Riordan, Patrick Rey, Jean Tirole and seminar participants at Berkeley Haas, Bocconi University, the 2013 CSIO-IDEI Conference, Universidad Carlos III of Madrid, NYU Stern, and the 2013 SAET Conference. Pavan also thanks the National Science Foundation and Nokia for financial support and the University of California at Berkeley for hospitality during the 2011-1012 academic year. Emil Temnyalov provided excellent research assistance.

1 Introduction

Many markets feature platforms mediating the interactions among the various sides of the market. Examples include media outlets mediating the interactions between readers/viewers on one side and content providers and advertisers on the other side; video-game consoles mediating the interactions between gamers and video-game developers; operating systems mediating the interactions between end-users and software developers; e-commerce websites mediating the interactions between buyers and sellers; employment agencies mediating the interactions between employers and job seekers; and dating agencies mediating the search of partner-seekers.

The literature on two-sided markets has studied the role of prices in implementing such mediated interactions.¹ The assumption commonly made in this literature is that preferences on each side of the market are commonly known. This assumption implies that, given the prices set by the platforms, each agent from each side can perfectly predict the participation decisions of all other agents. In equilibrium, such predictions are accurate and coincide with the platforms' predictions.

While a convenient modelling shortcut, the assumption that preferences are commonly known does not square well with most markets. Preferences over the products and services of different platforms typically reflect personal traits, making it difficult for an agent to predict the behavior of other agents. This is especially so when the product offered by one (or multiple) platform is relatively new. Due to network effects, predicting how many agents from the opposite side will choose a given platform is key to an agent's own decision about which platform to join. (Think of a consumer trying to determine whether or not he needs an iPad at the time the latter was launched to the market; his willingness to purchase the new tablet increases with his expectation about the number of applications that will be developed).

In this paper, we develop a tractable, yet rich, model of platform pricing under dispersed information, where the distribution of preferences in the cross-section of the population is unknown to both the platforms and to each individual agent, and where each agent possibly has private information about his own preferences as well as about the distribution of preferences in the cross-section of the population. Part of the contribution is in showing how such dispersion of information, by introducing heterogeneity in the users' expectations about the participation rates, shapes the elasticity of demand on each side. We then use such a characterization to examine the effects of the dispersion of information on the equilibrium prices and on the network allocations that they induce. Finally, we examine the platforms' incentives to change the information available to each side through informative advertising and marketing campaigns, as well as their incentives to invest in product design so as to change the way their product is likely to be perceived relative to those offered by the competitors.

Model preview. We consider a market where two platforms compete on two sides. Each side is populated by a continuum of agents. Each agent derives a direct utility from each of the two

¹Early contributions include Caillaud and Jullien (2001,2003), Rochet and Tirole (2003, 2006), and Armstrong (2006). See Rysman (2009) and Weyl (2010) for excellent overviews and more recent developments.

platforms' products hereafter referred to as the agent's "*stand-alone valuation*".² In addition, each agent derives an indirect (reduced-form) utility from interacting with the other side of the market that is proportional to the number of agents from the other side who join the same platform; hereafter, we refer to this component of the agent's utility as "*network effect*".³ Each agent is uncertain about the distribution of stand-alone valuations in the cross section of the population at the time he must choose which platform to join.⁴ In addition, each agent may face uncertainty about his own stand-alone valuations, reflecting the idea that agents need not know which products and services serve best their needs.⁵

For simplicity, but also to isolate the novel effects brought in by the dispersion of information, we consider markets where all agents from the same side attach the same value to interacting with agents from the opposite side.⁶ However, because different agents hold different expectations about how many agents from the opposite side are likely to join each platform, the model de facto accommodates a particular form of heterogeneity in the *estimated* network effects.

We allow for the possibility that the network effects are negative but restrict attention to markets where they are positive on one side (for example, in the case of a media outlet competing for readers, or viewers, on one side and for advertisers on the other side, it is reasonable to assume that network effects are negative on the readers' side—most readers dislike advertisement—but positive on the advertisers' side). We also assume that stand-alone valuations are positively correlated between any two agents from the same side but possibly negatively correlated between two agents from opposite sides (think of the market for operating systems; a system that appeals to software developers need not necessarily appeal to end-users, for the latter typically value features of the operating system differently from the developers—e.g., they may value the simplicity of the key tasks more than the

²Other expressions favored in the literature are "intrinsic benefit" — e.g., Armstrong and Wright (2007) — and "membership benefit" — e.g., Weyl (2010).

³Other expressions favored in the literature are "usage value" — e.g., Rochet and Tirole (2006) — "cross-side externality" — e.g., Armstrong (2006) — and "interaction benefit" — e.g., Weyl, (2010).

⁴In the baseline version of the model we do not allow agents to multi-home (that is, to join both platforms). Later in the paper, however, we relax this assumption and show that multihoming does not obtain under reasonable parameter configurations if one assumes that platforms cannot set negative prices (see Armstrong and Wright (2006), Bergemann and Bonatti (2011), Athey, Calvano and Gans (2012), and Ambrus, Calvano and Reisingerx (2013) for models that allow for multihoming under complete information, and Amelio and Jullien (2010) for how the impossibility to set negative prices may lead to tying).

⁵We consider a one-shot interaction between the platforms. In future work it will be interesting to extend the analysis to a dynamic setting where the platforms affect the speed of individual and social learning through their pricing strategies. We see the static analysis in the present paper as a necessary first step towards the analysis of such richer settings (see Mitchell and Skrzypacz (2006), and Cabral (2011) for dynamic models of network competition focusing, however, on different issues).

⁶The role of the heterogeneity in the values assigned to the network effects under complete information is studied in Weyl (2010), Veiga and Weyl (2011), and White and Weyl (2012). These papers, however, do not consider the possibility that preferences are correlated across sides, which is one of the key forces behind the mechanism we consider in the present paper.

flexibility and sophistication of the code).

We build on the global-game literature⁷ by assuming that the cross-sectional distribution of the stand-alone valuations can be parametrized by the realization of a random vector hereafter referred to as the "aggregate state". Each agent receives a noisy signal about the aggregate state that he uses to estimate his own stand-alone valuations as well as to predict the participation decisions on the other side of the market. This inference problem introduces novel effects that are missing under complete information and that are at the core of our analysis.

Implications for equilibrium prices. As in most of the literature, we abstract from price discrimination and assume that platforms compete by setting access fees on each side of the market. By paying the fee, an agent is granted access to the platform's product and thereby also to the other side of the market. We also assume that the platforms do not possess any private information relative to the rest of the market. This permits us to abstract from the signaling role of prices and isolate the novel effects that emerge when agents extrapolate from their own preferences to estimate the participation decisions on the other side of the market.⁸

The advantage of casting the analysis within a global-game framework is twofold: (i) it facilitates the analysis of how the dispersion of information impacts the equilibrium prices; (ii) it also guarantees that the equilibrium demand functions are unique (thus avoiding the usual "chicken and egg" problem of many two-sided markets⁹). In particular, for any given vector of prices there is a unique distribution of users over the two platforms (note that this is true despite the fact that platforms, in our model, compete in simple access fees—as they do in most markets—which do not condition on participation rates from the opposite side¹⁰).¹¹

A key difference relative to complete information is that each agent's beliefs about the participation decisions on the opposite side depend on the agent's own estimated stand-alone valuations. As the platform changes its price on one side, the marginal agent's beliefs about the participation rate on the opposite side thus also change (the marginal agent is the one who is indifferent between joining one platform or the other). Dispersed information thus induces a specific form of correlation between the estimated stand-alone valuations and the estimated network effects. Importantly, the endogeneity of such correlation has important implications for the equilibrium prices that differ from

⁷See, among others, Carlsson and Van Damme (1993), and Morris and Shin (2003).

⁸We also abstract from within-side externalities and heterogeneity in users' attractiveness. See Damiano and Li (2007), Gomes and Pavan (2013), and Veiga and Weyl (2012) for models that combine certain forms of price discrimination with heterogeneity in attractiveness. See also Ambrus and Argenziano (2009) for a model with heterogeneous network effects where platforms discriminate by offering multiple networks.

⁹See, e.g., Caillaud and Jullien (2003).

¹⁰Weyl (2010) and White and Weyl (2012) study how multiplicity can be eliminated if platforms may offer tariffs where the price on each side is a function of the participation rate on the other side—also known as "insulating tariffs".

¹¹Another advantage of this modelization is that the unique equilibrium in the continuation game where users choose which platform to join, if any, coincides with the unique rationalizable strategy profile. It thus does not require a high ability to coordinate with other agents. This is appealing, especially in large markets, which are the focus of the paper.

those obtained by assuming an exogenous correlation structure within each side.¹²

Suppose, for example, that network effects are positive on both sides (meaning that all agents benefit from a higher participation rate on the opposite side) and that preferences are positively correlated between the two sides (so that a high stand-alone valuation is "good news" about participation from the opposite side). Then suppose that a platform were to raise its price on, say, side 1. Because the marginal agent from side 1 who is excluded is the most "pessimistic" about side 2's participation, among those who join the platform, the drop in expected demand is smaller than in a world where all agents share the same beliefs about the other side's participation (as is necessarily the case under complete information). In other words, when preferences are positively correlated between the two sides and network effects are positive on both sides, the dispersion of information contributes to a reduction in the own-price elasticity of the demand functions. As a result of this new effect, the equilibrium price on each side increases with the intensity of that side's network effects when preferences are positively correlated between the two sides, and decreases otherwise.¹³

Despite the complexity of the strategic interactions, the model yields an extremely simple formula for the equilibrium duopoly prices. Holding fixed the ex-ante distribution of estimated stand-alone valuations (which amounts to fixing the ex-ante degree of differentiation between the two platforms), the equilibrium prices depend on the distribution of information only through a *coefficient of mutual forecastability*. The latter is an increasing transformation of the correlation coefficient between the signals of any two agents from opposite sides.

The reason why the equilibrium prices respond to the dispersion of information in such a coarse way is that each side values its ability to forecast the distribution of preferences on the other side only insofar this permits it to forecast participation decisions on the other side. For example, suppose that side-1 agents possess high-quality information that permits them to predict well the needs and *true* preferences of the side-2 agents. In contrast, suppose that the side-2 agents possess low-quality information. Then, the side-1 agents will expect the side-2 agents to not respond much to variations in their true stand-alone valuations, making the information of the side-1 agents of limited value. As a result, the equilibrium prices on the two sides will not differ significantly from a situation where both sides possess low-quality information. The implications of the aforementioned result are most striking in the case of a market that is perfectly symmetric under complete information (meaning that the intensity of the network effects is the same across the two sides and so is the ex-ante distribution of stand-alone valuations). In this case, the equilibrium prices remain symmetric under dispersed information, despite possible asymmetries in the distribution of information.

Implications for advertising, marketing, and product design. The results described

¹²See, e.g., Weyl (2010).

¹³In contrast, under complete-information, irrespective of the degree of correlation of preferences between the two sides, the equilibrium price on each side is invariant in that side's intensity of network effects, when the importance that each agent assigns to such network effects is independent of his own stand-alone valuation (see, e.g., Armstrong, 2006, and Rochet and Tirole, 2006).

above have important implications for the platforms' incentives to change the information available to each side, possibly through advertising and marketing campaigns as well as various information disclosures aimed at affecting the agents' ability to estimate both their own valuations as well as the distribution of valuations on the other side of the market.

We show that campaigns that increase the agents' ability to estimate their own stand alone valuations always increase profits. This is because such campaigns, by increasing the sensitivity of individual demands to information increase the ex-ante degree of differentiation between the two platforms (equivalently, they reduce the elasticity of the residual demands), thus softening competition.¹⁴

In contrast, campaigns whose role is primarily to help the agents predict the participation decisions on the opposite side of the market increase profits if and only if the correlation of stand-alone valuations between the two sides is of the same sign as the sum of the intensity of the network effects. In particular, such campaigns increase profits when network effects are positive on both sides and stand-alone valuations are positively correlated (as is probably the case in the market for video-game consoles). On the contrary, they decrease profits when either (a) stand-alone valuations are negatively correlated and network effects are positive on each side (as is possibly the case for some operating systems), or (b) stand-alone valuations are (weakly) positively correlated but one side suffers from the presence of the other side more than the other side benefits from its presence (as in the case of certain media outlets).

To understand this last result, assume that stand-alone valuations are positively correlated between the two sides and consider a campaign that increases the ability of the side-1 agents to forecast the stand-alone valuations of the side-2 agents (think of a video-game company advertising that, thanks to recent innovation in its software, the cost to the programmers to develop new games is low). An increase in such ability reduces the own-price elasticity of demand on side 1 by making the marginal agent's beliefs more sensitive to his private information (As explained above, a higher sensitivity to private information implies a lower drop in demand in response to an increase in price due to the fact that the marginal agent is less optimistic about participation from the opposite side than any of the inframarginal agents). Interestingly, when preferences are positively correlated, an increase in the quality of the side-1 agents' information about the side-2 agents' valuations also reduces the own-price elasticity of the side-2 demand by making the behavior of the side-1 agents more predictable in the eyes of the side-2 agents. These effects unambiguously contribute to higher equilibrium prices on both sides. At the same time, more precise information on side 1 also implies a higher sensitivity of both sides to variations in prices on the opposite side, which contributes negatively to the equilibrium prices. While the net effect on the equilibrium price on each side depends on the relative importance that the two sides attach to interacting with one another, the net effect on total profits is always unambiguously positive when the sum of the network effects is

¹⁴A similar result is obtained in Anderson and Renault (2010) in the context of an ex-ante symmetric one-sided market.

positive (more generally, of the same sign as the correlation of preferences between the two sides). This is because, holding constant the ex-ante distributions of estimated stand-alone valuations, the equilibrium price on each side depends on the dispersion of information only through the coefficient of mutual forecastability, which is increasing in the quality of information on each of the two sides. When the sum of the network effects is positive, then any possible loss of revenues on one side must necessarily be more than compensated by an increase in revenues on the opposite side, making the equilibrium total profits unambiguously increase with each side's ability to forecast the preferences of the other side.

We also investigate how equilibrium profits change with variations in the prior distribution from which stand-alone valuations are drawn. These comparative statics, contrary to the ones pertaining the quality of information, are meant to shed light on a platform's incentives to differentiate its product from the competitor's, without knowing the exact distribution of preferences on either side of the market. For instance, we show that raising the similarity with the opponent's product always reduces the equilibrium profits by intensifying competition. On the other hand, aligning the preferences of the two sides by favoring product dimensions that are appealing to both sides increases profits for positive network effects but reduces them when the sum of the network effects is negative (that is, when one side suffers from the presence of the other side more than the other side benefits from its presence).

Outline. The rest of the paper is organized as follows. Below we wrap up the Introduction with a brief discussion of the contribution of the paper relative to the pertinent literature. Section 2 presents the model. Section 3 introduces some preliminary results concerning the ability of each side to forecast its own preferences and the cross-sectional distribution of preferences on the other side of the market, and discusses the benchmark case with no network effects. Section 4 characterizes optimal prices for a monopolistic platform. Section 5 contains the main results for the duopoly case. Section 6 contains implications for advertising and product design. Section 7 offers a few concluding remarks. All proofs are in the Appendix.

(Most) pertinent literature. The paper contributes to three lines of inquiry. The first is the one examined in the two-sided-market literature. This literature is too vast to be successfully summarized here. We refer the reader to Rysman (2009) and Rochet and Tirole (2006) for excellent overviews.¹⁵ The closest papers to ours are Armstrong (2006), Rochet and Tirole (2006), Weyl (2010), and White and Weyl (2012). The first two papers study monopoly and duopoly pricing in a market with differentiated products, assuming homogenous network effects. Our model is the incomplete-information analog of the model studied in these papers. Weyl (2010) extends this model by allowing for heterogenous network effects, focusing on a monopolistic platform, while White and Weyl (2012) extends the analysis in Weyl (2010) to a duopoly (see also Ambrus and Argente (2009) who were the first to introduce heterogenous network effects and show how the latter can lead

¹⁵The paper is also related to the literature on one-sided markets with network effects. See Katz and Shapiro (1985) for a pioneering contribution and Farrell and Klemperer (2006) for a recent overview.

to asymmetric networks under coalition-rationalizable strategies).

The key contribution of our paper relative to this literature is in uncovering the implications of dispersed information about participation decisions. We identify a new channel by which the dispersion of information affects the elasticity of the demands on the two sides and thereby the equilibrium prices. This in turn permits us to uncover novel effects. For example, under complete information, it is the discrepancy between the importance assigned to network effects by the marginal user and by the average user that is responsible for distortions in prices and in network allocations, along the lines of those identified in Spence (1975) (see Weyl, 2010). In contrast, under dispersed information, it is the discrepancy between the participation rates expected by the marginal user on each side and the participation rates expected by the two platforms that is responsible for novel distortions.

Our paper focuses on dispersed information at the subscription stage. In contrast, Halaburda and Yehezkel (2013) analyze a model where two homogenous platforms compete by offering access fees and menus of trades and where buyers and sellers privately learn their valuations and costs only after joining a platform but before transacting with the other side. While the two papers address very different questions, they both point to the importance of asymmetric information for platforms' pricing decisions.

The second line of inquiry the paper contributes to is the one considered in the literature on coordination under incomplete information, and in particular in the global-games literature. Our paper does not contribute to this literature in any theoretically significant way. However, to the best of our knowledge, it is the first paper to provide a global-game application where two distinct populations (the two sides) coordinate under dispersed information and where the outcome of such coordination is shaped by two competing "big players" (the platforms). The paper in this literature closest to ours is Argenziano (2008). That paper uses a global-game approach to study efficiency under product differentiation in the contest of a one-sided network duopoly. The questions addressed in that paper are fundamentally different from those addressed in the present paper which are largely motivated by the two-sidedness of the problem under examination. We also allow for a richer class of information structures which permits us to study the platforms' incentives to change the information available to the two sides via advertising and marketing campaigns, as well as product design.

The third line of inquiry is the one that studies informative advertising and marketing campaigns (for recent contributions see, e.g., Anderson and Renault (2006, 2009), Johnson and Myatt (2006) and the references therein). Our results about the effects on profits of campaigns that help the agents understand their own needs and preferences are in the same spirit of those established in this literature. The main contribution is in investigating the effects of campaigns that help the agents forecast the preferences and behavior of other agents from the opposite side, which is new and brings novel implications.

2 Model

Players. Two platforms, indexed by $k = A, B$, compete on two sides, $i = 1, 2$. Each side is populated by a measure-one continuum of agents, indexed by $l \in [0, 1]$.

Actions and payoffs. Each agent $l \in [0, 1]$ from each side $i = 1, 2$ must choose which platform to join, if any.¹⁶ The payoff U_{il}^k that agent l from side i derives from joining platform k is given by

$$U_{il}^k = u_{il}^k + \gamma_i m_j^k - p_i^k$$

where u_{il}^k is the idiosyncratic *stand-alone valuation* of joining platform k , $m_j^k \in [0, 1]$ is the mass of agents from side $j \neq i$ that join platform k , $\gamma_i \in \mathbb{R}$ is a parameter that controls for the intensity of the *network effects* on side i and p_i^k is the price (the access fee) charged by platform k to side i .

We assume that the network effects are positive on at least one of the two sides but allow them to be negative on the opposite side; that is, we assume that $\gamma_i > 0$ for some $i \in \{1, 2\}$.

The payoff that each agent $l \in [0, 1]$ from each side $i = 1, 2$ obtains from not joining any platform is assumed to be equal to zero.

Each platform's payoff Π^k is the total revenue from collecting the prices from the two sides.¹⁷

$$\Pi^k = p_1^k m_1^k + p_2^k m_2^k.$$

All players are risk-neutral expected-utility maximizers.

Horizontal differentiation and information. We assume that the stand-alone valuations are given by

$$u_{il}^A = s_i - \frac{1}{2}v_{il} \text{ and } u_{il}^B = s_i + \frac{1}{2}v_{il}$$

$i = 1, 2$, $k = A, B$, $l \in [0, 1]$, where $s_i \in \mathbb{R}$ is a known scalar whose role is to control for the agents' payoff relative to their outside options.¹⁸ The above specification is chosen so that the difference in stand-alone valuations is $v_{il} \equiv u_{il}^B - u_{il}^A$.

The "aggregate state" of the market corresponds to the joint distribution of stand-alone valuations and of the agents' information. We parametrize this state by a pair $\theta \equiv (\theta_1, \theta_2)$ and assume that θ is drawn from a bivariate Normal distribution with mean $(0, 0)$ and variance-covariance matrix

$$\Sigma_\theta = \begin{bmatrix} (\alpha_1)^{-1} & \frac{\rho_\theta}{\sqrt{\alpha_1 \alpha_2}} \\ \frac{\rho_\theta}{\sqrt{\alpha_1 \alpha_2}} & (\alpha_2)^{-1} \end{bmatrix}$$

where the parameter ρ_θ denotes the coefficient of linear correlation between $\tilde{\theta}_1$ and $\tilde{\theta}_2$.¹⁹

¹⁶Below we will also discuss the possibility that the agents may choose to join both platforms (multihoming).

¹⁷All results extend to the case where the platforms incur costs to provide access to the users. Because these costs do not play any role, we disregard them to facilitate the exposition.

¹⁸As it will become clear in a moment, s_i coincides with the unconditional average stand-alone valuation of each side- i 's agent for each of the two products.

¹⁹Throughout, we will use tildes " \sim " to denote random variables and denote their realization without the tildes.

Neither the platforms nor the agents observe θ . Furthermore, each agent may have an imperfect knowledge of his own valuations. We formalize this by assuming that each agent l from each side $i = 1, 2$ privately observes a signal x_{il} that is imperfectly correlated with both θ and v_{il} . More precisely we assume that

$$v_{il} = z_i (\theta_i + \varepsilon_{il}) \text{ and } x_{il} = \theta_i + \eta_{il}$$

where z_i is a non-negative scalar and where the variables $(\tilde{\varepsilon}_{il}, \tilde{\eta}_{il})$ are idiosyncratic terms drawn from a bivariate Normal distribution with mean $(0, 0)$ and variance-covariance matrix

$$\Sigma_i = \begin{bmatrix} (\beta_i^\varepsilon)^{-1} & \frac{\rho_i}{\sqrt{\beta_i^\varepsilon \cdot \beta_i^\eta}} \\ \frac{\rho_i}{\sqrt{\beta_i^\varepsilon \cdot \beta_i^\eta}} & (\beta_i^\eta)^{-1} \end{bmatrix}$$

with the parameter $\rho_i \geq 0$ denoting the coefficient of linear correlation between $\tilde{\varepsilon}_i$ and $\tilde{\eta}_i$. The pairs $(\varepsilon_{il}, \eta_{il})_{l \in [0,1], i \in \{0,1\}}$ are drawn independently across agents and independently from θ .

Timing.

- At stage 1, platforms simultaneously set prices on each side.
- At stage 2, after observing the prices $(p_i^k)_{i=1,2}^{k=A,B}$, and after receiving the information x_{il} , each agent $l \in [0, 1]$ from each side $i = 1, 2$, simultaneously chooses which platform to join, if any.
- Finally, at stage 3, payoffs are realized.

Comment. The above specification has the advantage of being tractable, while at the same time rich enough to capture a variety of situations. Thanks to Normality, the "aggregate state" (i.e., the cross-sectional distribution of preferences and information) is uniquely pinned down by the bivariate variable $\theta = (\theta_1, \theta_2)$. The information about θ is dispersed so that different agents have different beliefs about θ .

The pure common-value case where all agents from side i have identical preferences over the products of the two platforms but different information about the quality differential is captured as the limit in which $\beta_i^\varepsilon \rightarrow \infty$, in which case, almost surely, $v_{il} = z_i \theta_i$ all $l \in [0, 1]$. The parameter α_i is then a measure of differentiation between the two products, as perceived by side i . Letting $\alpha_1 = \alpha_2$ and $\rho_\theta = 1$ while allowing $\beta_1^\eta \neq \beta_2^\eta$ then permits us to capture situations where the quality differential between the two platforms is the same on both sides, but where one side may have superior information than the other. Letting $z_i = 0$ in turn permits us to capture situations where agents on side i do not care about the intrinsic quality differential between the two products but nonetheless possess information about the distribution of preferences on the opposite side (as in the case of advertisers who choose which media platform to place ads on entirely on the basis of their expectation of the platform's ability to attract readers and viewers from the opposite side).

More generally, allowing the correlation coefficient ρ_θ to be different from one permits us to capture situations where the quality differential between the two products differs across the two

sides (including situations where it is potentially negatively correlated), as well as situations where one side may be able to perfectly predict the behavior of each agent from that side but not the behavior of agents from the opposite side (which corresponds to the limit in which $\beta_i^\eta \rightarrow \infty$).

The model can also capture situations in which different users from the same side have different preferences for the two products. This amounts to letting the variance of $\tilde{\varepsilon}_{il}$ be strictly positive or, equivalently, $\beta_i^\varepsilon < \infty$. Depending on the degree of correlation ρ_i between $\tilde{\varepsilon}_{il}$ and $\tilde{\eta}_{il}$, agents may then possess more or less accurate information about their own stand-alone valuations. For example, the case where each agent perfectly knows his own valuations but is imperfectly informed about the valuations of other agents (from either side) is captured as the limit in which $\rho_i \rightarrow 1$. The extreme case of independent private values then corresponds to the limit in which $\alpha_i \rightarrow \infty$ and $\beta_i^\varepsilon < \infty$.

Also note that the scalars (z_1, z_2) only serve the purpose of parametrizing the quality of the agents' information about their own stand-alone valuations relative to the quality of their information about the distribution of stand-alone valuations on the other side of the market. These parameters are not crucial and could have been dispensed with by introducing two separate signals for each agent, one for $\tilde{\theta}_1$, the other for $\tilde{\theta}_2$. This, however, would have made the subsequent analysis significantly less tractable by essentially requiring that we describe the equilibrium strategies in terms of semi-planes as opposed to simple cut-off rules.

Finally note that, as mentioned already, the remaining parameters (s_1, s_2) play a role only for the agents' decision to opt out of the market by not joining any platform.

3 Preliminaries

Reduced form representation. The key determinant of the equilibrium allocations (prices and participation decisions) will be the agents' ability to forecast their own stand-alone valuations, as well as the distribution of such valuations on the other side of the market. As described above, the information of each agent l from each side i is encoded in a single signal x_{il} . This signal is drawn for a Normal distribution with zero mean and variance

$$\frac{1}{\beta_i^x} \equiv \text{var}(\tilde{x}_{il}) = \frac{\alpha_i + \beta_i^\eta}{\alpha_i \beta_i^\eta}. \quad (1)$$

Notice that the agents' signals are correlated both within sides and across sides. The important correlation is the one across sides. For any two agents l and l' from opposite sides, the coefficient of linear correlation of their signals is

$$\rho_x \equiv \frac{\text{cov}(\tilde{x}_{1l}, \tilde{x}_{2l'})}{\sqrt{\text{var}(\tilde{x}_{1l}) \text{var}(\tilde{x}_{2l'})}} = \rho_\theta \sqrt{\frac{\beta_1^\eta \beta_2^\eta}{(\alpha_1 + \beta_1^\eta)(\alpha_2 + \beta_2^\eta)}}. \quad (2)$$

Based on the signal x_{il} , each agent i then believes that the differential \tilde{v}_{il} in his stand-alone valuations is Normally distributed with mean

$$V_{il} \equiv \mathbb{E}[\tilde{v}_{il} \mid x_{il}] = \kappa_i x_{il} \text{ with } \kappa_i \equiv \frac{\text{cov}[\tilde{v}_{il}, \tilde{x}_{il}]}{\text{var}[\tilde{x}_{il}]} = z_i \frac{\beta_i^\eta + \rho_i \alpha_i \sqrt{\beta_i^\eta / \beta_i^\varepsilon}}{\beta_i^\eta + \alpha_i}. \quad (3)$$

Hereafter, we will refer to $V_{il} \equiv \mathbb{E}[\tilde{v}_{il} \mid x_{il}]$ as to the *estimated stand-alone differential*. Note that V_{il} uniquely pins down not only the differential but also the agent's estimated stand-alone valuations.

Next, consider the agents' ability to forecast the participation decisions on the other side of the market. Because each agent observes only a noisy signal of his valuations, the best an agent can do to predict participation decisions on the other side of the market is to use his own signal x_{il} to forecast the distribution of signals on the other side. Now observe that each agent l from each side i , after observing a signal x_{il} , believes that each agent l' from the opposite side received a signal $\tilde{x}_{jl'} = \tilde{\theta}_j + \tilde{\eta}_{jl'}$ drawn from a Normal distribution with mean

$$\mathbb{E}[\tilde{x}_{jl'} \mid x_{il}] = \rho_x \sqrt{\frac{\beta_i^x}{\beta_j^x}} x_{il} \quad (4)$$

and variance

$$\text{var}[\tilde{x}_{jl'} \mid x_{il}] = \frac{1 - \rho_x^2}{\beta_j^x}. \quad (5)$$

It is then easy to see that, by varying the coefficient ρ_i of correlation between the two idiosyncratic terms $(\tilde{\varepsilon}_{il}, \tilde{\eta}_{il})$ while keeping all other parameters fixed, one can capture variations in the agents' ability to estimate their own stand-alone valuations, holding fixed the agents' ability to estimate the participation decisions on the other side of the market. Likewise, by varying ρ_x (for example by varying ρ_θ) holding fixed all other parameters, one can capture variations in the agents' ability to estimate the participation decisions on the other side of the market, holding constant their ability to estimate their own stand-alone valuations.

For the first part of the paper, the key parameters of the model will be $(\beta_1^x, \beta_2^x, \rho_x)$, which parametrize the agents' information²⁰, and the parameters $(s_1, s_2, \kappa_1, \kappa_2)$ which define the individual estimated stand-alone valuations for given information.

In the second part of the paper, we will discuss how more structural parameters such as ρ_i or ρ_θ affect the equilibrium and how firms can modify them with marketing and advertising campaigns, as well as product design.

Benchmark: Absence of network effects. As a warm-up (but also as a useful step to fix ideas and introduce notation that will be used throughout the rest of the analysis), consider for a moment a market without network effects. In our framework this corresponds to setting $\gamma_1 = \gamma_2 = 0$. In this case the demand on each side is independent of the pricing and participation decisions on the other side of the market.

Consider first the case where platform A is a monopolist. Given the price p_i^A on side i , each agent l from side i buys only if his estimated stand-alone valuation for the platform's product is above the price; that is, only if $\mathbb{E}[\tilde{u}_{il}^A \mid x_{il}] - p_i^A \geq 0$. Using the fact that $\mathbb{E}[\tilde{u}_{il}^A \mid x_{il}] = s_i - \frac{1}{2}\kappa_i x_{il}$ we have that the agent buys only if his signal is low enough,

$$x_{il} < \hat{x}_i \equiv 2 \left(\frac{s_i - p_i^A}{\kappa_i} \right).$$

²⁰Formally, one should consider also the correlation of signals within sides, but this will play no role in the analysis.

Notice that, by choosing the price, the platform chooses the signal of the marginal consumer \hat{x}_i . The total demand m_i^A on side i then depends on the realization of $\tilde{\theta}_i$, which pins down the distribution of stand-alone valuations, and which is unknown to the platform at the time the platform sets its price. Letting Φ denote the c.d.f. of the standard Normal distribution and ϕ its density, we then have that the demand the platform expects on side i when it sets a price p_i^A (equivalently, when it chooses a marginal agent \hat{x}_i) is given by

$$Q_i^A = \mathbb{E}[\tilde{m}_i^A] = \Pr(x_{il} < \hat{x}_i) = \Phi(\sqrt{\beta_i^x} \hat{x}_i). \quad (6)$$

Now let

$$\mu_i(x_i) = -\frac{Q_i^A}{\frac{dQ_i^A}{dp_i^A}} = \frac{\kappa_i}{2} \frac{Q_i^A}{\frac{dQ_i^A}{dx_i}} = \frac{\kappa_i}{2} \frac{\Phi(\sqrt{\beta_i^x} x_i)}{\sqrt{\beta_i^x} \phi(\sqrt{\beta_i^x} x_i)} \quad (7)$$

denote the *inverse semi-elasticity* of the stand-alone demand evaluated at the price $p_i^A = s_i - \frac{1}{2}\kappa_i \hat{x}_i$.²¹ The monopoly price p_i^A is then given by the usual first-order condition

$$p_i^A = \mu_i(\hat{x}_i) \Leftrightarrow s_i - \frac{1}{2}\kappa_i \hat{x}_i = \mu_i(\hat{x}_i).$$

Next, consider a duopoly where platforms A and B set prices simultaneously on each side. Assuming full participation (that is, each agent who does not choose platform A chooses platform B), we then have that each agent l from side i buys from A if $\mathbb{E}[\tilde{u}_{il}^B - \tilde{u}_{il}^A \mid x_{il}] < p_i^B - p_i^A$ and from B if the inequality is inverted.²² Using the fact that $\mathbb{E}[\tilde{u}_{il}^B - \tilde{u}_{il}^A \mid x_{il}] = \mathbb{E}[\tilde{v}_{il} \mid x_{il}] = \kappa_i x_{il}$, we then have that the demand that platform A expects when the prices are p_i^A and p_i^B is given by

$$Q_i^A = \mathbb{E}[\tilde{m}_i^A] = \Phi\left(\sqrt{\beta_i^x} \hat{x}_i\right) = 1 - Q_i^B, \text{ with } \hat{x}_i = \frac{p_i^B - p_i^A}{\kappa_i}$$

where \hat{x}_i is the signal of the marginal agent (the agent who is indifferent between purchasing from A and purchasing from B). Now let

$$\mu_i^d(x) = -\frac{Q_i^A}{\left.\frac{dQ_i^A}{dp_i^A}\right|_{p_i^B=\text{const}}} = \kappa_i \frac{\Phi(\sqrt{\beta_i^x} x)}{\sqrt{\beta_i^x} \phi(\sqrt{\beta_i^x} x)} \quad (8)$$

denote the *inverse semi-elasticity* of the *residual demand curve* of platform A , evaluated at the price $p_i^A = p_i^B - \kappa_i x$. It is then easy to see that in the unique symmetric duopoly equilibrium each agent l from side i buys from platform A if $x_{il} < \hat{x}_i^d = 0$ and from platform B if $x_{il} > \hat{x}_i^d = 0$. In equilibrium, each firm serves half of the market (i.e., $Q_i^A = Q_i^B = 1/2$) and the equilibrium prices are given by

$$p_i^A = p_i^B = \mu_i^d(0). \quad (9)$$

²¹This semi-elasticity is referred to as *the market power* in Weyl (2010).

²²When $\mathbb{E}[\tilde{u}_j^A - \tilde{u}_j^B \mid x_{il}] = p_i^A - p_i^B$, the consumer is indifferent. Because this event has zero probability, the way such indifference is resolved is inconsequential for the choice of the optimal prices.

Using (3), note that the equilibrium semi-elasticity of the residual stand-alone demands is given by

$$\mu_i^d(0) = \frac{\kappa_i}{\sqrt{\beta_i^x} 2\phi(0)} = \frac{\sqrt{\text{var}[\tilde{V}_{il}]}}{2\phi(0)}. \quad (10)$$

where $\text{var}[\tilde{V}_{il}]$ is the ex-ante dispersion of the estimated stand-alone differentials $V_{il} = \mathbb{E}[\tilde{v}_{il} \mid x_{il}]$. Not surprisingly, a higher dispersion of estimated stand-alone differentials is isomorphic to a higher degree of differentiation between the two platforms, which lessens competition and thus results in higher equilibrium prices.

4 Monopoly

We now turn to the model with network effects. We start by considering the case of a monopolistic market, in which only platform A is active.

Given the prices (p_1^A, p_2^A) , each agent l from each side i finds it optimal to join the platform only if

$$\mathbb{E}[\tilde{u}_{il}^A \mid x_{il}] + \gamma_i \mathbb{E}[\tilde{m}_j^A \mid x_{il}] - p_i^A \geq 0. \quad (11)$$

Now let $\gamma_i^- \equiv \min\{\gamma_i; 0\}$ and $\gamma_i^+ \equiv \max\{\gamma_i; 0\}$. It is immediate to see that any agent whose expected stand-alone valuation $\mathbb{E}[\tilde{u}_{il}^A \mid x_{il}]$ is less than $(p_i^A - \gamma_i^+)$ finds it dominant not to join, whereas any agent whose expected stand-alone valuation $\mathbb{E}[\tilde{u}_{il}^A \mid x_{il}]$ is greater than $p_i^A - \gamma_i^-$ finds it dominant to join. Using $\mathbb{E}[\tilde{u}_{il}^A \mid x_{il}] = s_i - \kappa_i x_{il}/2$, we then have that iterated deletion of strictly dominated strategies leads to a pair of thresholds $\underline{x}_i = \underline{x}_i(p_1^A, p_2^A)$ and $\bar{x}_i = \bar{x}_i(p_1^A, p_2^A)$ on each side $i = 1, 2$ such that it is iteratively dominant for each agent l from each side i to join for $x_{il} < \underline{x}_i$ and not to join for $x_{il} > \bar{x}_i$. These observations also suggest existence of a continuation equilibrium in threshold strategies whereby each agent l from each side i joins if and only if $x_{il} \leq \hat{x}_i$. In any such continuation equilibrium, the participation rate on side j (i.e., the measure of agents from side j who join the platform) is given by

$$m_j^A = \Pr(\tilde{x}_{jl} \leq \hat{x}_j \mid \theta_j).$$

We refer to an allocation with this property as a threshold allocation (\hat{x}_1, \hat{x}_2) . Notice that m_j^A decreases with θ_j , since a higher θ_j means fewer agents with a high stand-alone valuation for the platform's product. Using (4) and (5), we then have that, from the perspective of agent l from side i , the expected participation rate on side $j \neq i$ is given by

$$\mathbb{E}[\tilde{m}_j^A \mid x_{il}] = \Pr(\tilde{x}_{jl} \leq \hat{x}_j \mid x_{il}) = \Phi \left(\sqrt{\frac{\beta_j^x}{1 - \rho_x^2}} \left(\hat{x}_j - \rho_x \sqrt{\frac{\beta_i^x}{\beta_j^x}} x_{il} \right) \right)$$

Now, for any $i, j \in \{1, 2\}$, $i \neq j$, any (\hat{x}_1, \hat{x}_2) , let $M_j^A(\hat{x}_1, \hat{x}_2) \equiv \mathbb{E}[\tilde{m}_j^A \mid x_{il} = \hat{x}_i]$ denote the expected participation rate on side j from the perspective of the *marginal* agent on side i (the one

with signal \hat{x}_i). Then

$$M_j^A(x_1, x_2) \equiv \Phi(X_{ji}(x_1, x_2)) \text{ where } X_{ji}(x_1, x_2) = \frac{\sqrt{\beta_j^x} x_j - \rho_x \sqrt{\beta_i^x} x_i}{\sqrt{1 - \rho_x^2}}$$

Letting

$$\Omega \equiv \frac{\rho_x}{\sqrt{1 - \rho_x^2}}, \quad (12)$$

we then have that the term X_{ji} can be expressed as follows

$$X_{ji}(x_1, x_2) = \sqrt{1 + \Omega^2} \sqrt{\beta_j^x} x_j - \Omega \sqrt{\beta_i^x} x_i. \quad (13)$$

Hereafter, we will refer to the term Ω as to the *coefficient of mutual forecastability*. Note that $|\Omega|$ is increasing in each side's ability to forecast the distribution of information on the opposite side. As anticipated in the Introduction, this term will play an important role in determining the equilibrium prices.

Using (11), we then have that, in any threshold equilibrium, the thresholds (\hat{x}_1, \hat{x}_2) must jointly solve the following system of conditions

$$G_i(\hat{x}_1, \hat{x}_2) = p_i^A \quad i = 1, 2 \quad (14)$$

where

$$G_i(x_1, x_2) = s_i - \kappa_i x_i / 2 + \gamma_i M_j^A(x_1, x_2). \quad (15)$$

Note that the function $G_i(x_1, x_2)$ represents the payoff, gross of payments, of joining platform A for an agent on side i whose signal is equal to the threshold signal x_i when he expects all users from side $j \neq i$ to join if and only if their signal is smaller than x_j . To ensure that, for any vector of prices, a continuation equilibrium in threshold strategies exists, we assume that the function G_i is decreasing in x_i . This is the case, for all x_i , if and only if the following condition holds, which we assume throughout:

Condition (M): *The parameters of the model are such that $2\mu_i(0) + \gamma_i\Omega > 0$.*

Note that the above condition imposes that, when side i values interacting with the other side—namely, when $\gamma_i > 0$, the preferences between the two sides be not too negatively correlated. Symmetrically, the condition requires the correlation between $\tilde{\theta}_1$ and $\tilde{\theta}_2$ to be sufficiently small when side i dislikes the presence of the other side, that is when $\gamma_i < 0$. This is intuitive. Consider the case where $\gamma_i > 0$; if $\tilde{\theta}_1$ and $\tilde{\theta}_2$ were strongly negatively correlated (relative to γ_i , of course), then an increase in the appreciation by agent l from side i of the platform's product could make the agent less willing to join if he expects a significant drop in the participation by agents from the opposite side.

We then have the following preliminary result:

Lemma 1 *For any vector of prices $p = (p_1^A, p_2^A)$, there exists at least one solution to the system of conditions given by (14), which implies that a threshold continuation equilibrium always exists.*

Now, to guarantee that the continuation equilibrium is unique, for all possible prices, we assume that the strength of the network effects is not too large, given the distribution of the stand-alone valuations, in the sense of Condition (Q) below, which we assume throughout the rest of the analysis.

Condition (Q). *The parameters of the model are such that*

$$\gamma_1 \gamma_2 < \frac{[2\mu_1(0) + \gamma_1 \Omega][2\mu_2(0) + \gamma_2 \Omega]}{\sqrt{(1 + \Omega^2)} + \Omega^2}$$

We then have the following result:

Lemma 2 *For any vector of prices (p_1^A, p_2^A) , the continuation equilibrium is unique.*

The proof in the Appendix first shows that, when conditions M and Q hold, then, for any vector of prices, there exists a unique pair of thresholds $\hat{x}_i = \hat{x}_i(p_1^A, p_2^A)$, $i = 1, 2$, that solve the system of equations defined by the indifference conditions (14). Standard arguments from the global-games literature based on iterated deletion of strictly dominated strategies then imply that the unique monotone equilibrium defined by the thresholds \hat{x}_i , $i = 1, 2$, is the unique equilibrium of the continuation game.

Notice that condition Q implies condition M if γ_1 and γ_2 have the same sign, while condition M implies condition Q if network effects have opposite sign.

The above result implies that there exists a unique pair of demand functions. For any vector of prices (p_1^A, p_2^A) , the demand on side i in state $\theta = (\theta_1, \theta_2)$ is given by $m_i^A = \Phi(\sqrt{\beta_i^\eta}(\hat{x}_i - \theta_i))$, while the unconditional expected demand is $Q_i^A = \Phi(\sqrt{\beta_i^x} \hat{x}_i)$, where the thresholds $\hat{x}_i = \hat{x}_i(p_1^A, p_2^A)$, $i = 1, 2$, are the unique solution to the system of equations given by (14).

Now consider the choice of prices by the monopolist. For any pair of prices (p_1^A, p_2^A) , the monopolist's profits are equal to

$$\Pi^A(p_1^A, p_2^A) = \sum_{i=1,2} p_i^A \Phi\left(\sqrt{\beta_i^x} \hat{x}_i(p_1^A, p_2^A)\right).$$

Notice that the system of demand equations (14) defines a bijective relationship between (p_1^A, p_2^A) and (\hat{x}_1, \hat{x}_2) . The monopolist's problem can thus also be seen as choosing a pair of thresholds (\hat{x}_1, \hat{x}_2) so as to maximize

$$\hat{\Pi}^A(\hat{x}_1, \hat{x}_2) \equiv \sum_{i=1,2} G_i(\hat{x}_1, \hat{x}_2) \Phi\left(\sqrt{\beta_i^x} \hat{x}_i\right) \quad (16)$$

where $G_i(\hat{x}_1, \hat{x}_2)$ (defined in (15)) is the expected gross surplus of the marginal agent on side i , whose signal is equal to the threshold \hat{x}_i .

Next, for $i = 1, 2$, let

$$G_i^-(x) \equiv \left[s_i - \frac{\kappa_i}{2}x + \gamma_i^-\right] \Phi\left(\sqrt{\beta_i^x}x\right)$$

where recall that $\gamma_i^- \equiv \min\{\gamma_i; 0\}$. Throughout, we will assume that the following condition also holds, which guarantees that the optimal prices will be interior.

Condition (W). *The parameters of the model are such that, for any $i, j = 1, 2, j \neq i$,²³ $\max_{x \in \mathbb{R}} G_i^-(x) > |\gamma_j|$.*

Note that Condition W is trivially satisfied when s_i are large enough. The condition simply guarantees that it is always optimal to induce a strictly positive participation rate on both sides, despite the possibility that one side may suffer from the presence of the other side. We then have the following result:

Lemma 3 *A vector of prices (p_1^A, p_2^A) maximizing firm A's profits always exists. Furthermore any such vector must satisfy $p_i^A = G_i(\hat{x}_1, \hat{x}_2)$, $i = 1, 2$, with (\hat{x}_1, \hat{x}_2) solving the system of conditions given by²⁴*

$$G_i(\hat{x}_1, \hat{x}_2) \sqrt{\beta_i^x} \phi\left(\sqrt{\beta_i^x} \hat{x}_i\right) + \frac{\partial G_i(\hat{x}_1, \hat{x}_2)}{\partial x_i} \Phi\left(\sqrt{\beta_i^x} \hat{x}_i\right) + \frac{\partial G_j(\hat{x}_1, \hat{x}_2)}{\partial x_i} \Phi\left(\sqrt{\beta_j^x} \hat{x}_j\right) = 0. \quad (17)$$

To shed light on what lies underneath the first-order conditions for the monopolist's profit-maximizing prices, note that the latter are equivalent to

$$p_i^A + \frac{dp_i^A}{dQ_i^A} \Big|_{Q_j^A = \text{const}} \cdot Q_i^A + \frac{dp_j^A}{dQ_i^A} \Big|_{Q_j^A = \text{const}} \cdot Q_j^A = 0 \quad (18)$$

where $Q_i^A = \mathbb{E}[\tilde{m}_i^A]$ is the demand on side i , as expected by the platform. These first-order conditions are the incomplete-information analogs of the familiar complete-information optimality conditions according to which, at the optimum, profits must not vary when the monopolist changes the price on side i and, at the same time, adjusts the price on side j so as to maintain the expected demand on side j constant.

Notice that, under complete information about $\theta = (\theta_1, \theta_2)$, the demand on each side $i = 1, 2$ expected by the platform coincides with the demand expected by the marginal agent from the opposite side (that is, $M_i^A = Q_i^A$ for $i = 1, 2$). This leads to the familiar optimality condition

$$p_i^A = \mu_i(\hat{x}_i) - \gamma_j Q_j^A,$$

according to which the monopolist's price on each side equals the usual one-sided inverse semi-elasticity adjusted by the effect of a variation in the side- i 's participation on side- j 's revenues (the second term)—see, for example, Weyl (2010).

²³That the function G_i^- has a maximum follows from the fact that it is continuous, positive for $\hat{x}_i < 2(s_i + \gamma_i^-)/\kappa_i$, negative for $\hat{x}_i > 2(s_i + \gamma_i^-)/\kappa_i$ and such that $\lim_{\hat{x}_i \rightarrow -\infty} g_i(\hat{x}_i) = 0$.

²⁴While we did not prove that a solution (\hat{x}_1, \hat{x}_2) to the system of equations given below is unique, we conjecture that this is the case. Importantly, our results are independent of whether or not such a solution is unique. What is important is that, for any vector of prices, the continuation equilibrium is unique. This is what permits us to establish the properties of the equilibrium prices described below.

What is interesting here is how incomplete information affects the slope of the demand functions on the two sides and thereby the prices. While, with complete information, these slopes are the same irrespective of whether they are computed by the platform or by any other agent, this is not the case with dispersed information. In particular, even if the platform adjusts the price on side j so as to maintain the threshold \hat{x}_j fixed (which amounts to maintaining the side- j 's demand Q_j^A constant, as perceived by the platform), from the perspective of the new marginal agent on side i , the expected side- j 's demand changes in response to variations in the side- i 's price. Formally,

$$\frac{\partial M_j^A(\hat{x}_1, \hat{x}_2)}{\partial \hat{x}_i} = -\Omega \sqrt{\beta_i^x} \phi(X_{ji}(\hat{x}_1, \hat{x}_2)) \neq 0 \text{ if } \Omega \neq 0. \quad (19)$$

This in turn affects the slope of the side- i 's (inverse) demand function. Indeed as the side- i demand $Q_i^A = \sqrt{\beta_i^x} \phi(\sqrt{\beta_i^x} \hat{x}_i)$ changes in response to a variation in the side- i 's price, the side- j 's participation expected by the side- i 's marginal consumer also changes according to the relationship:

$$\left. \frac{dM_j^A}{dQ_i^A} \right|_{Q_j^A = \text{const}} = \frac{\frac{\partial M_j^A(\hat{x}_1, \hat{x}_2)}{\partial \hat{x}_i}}{\frac{dQ_i^A}{d\hat{x}_i}} = -\Omega \frac{\phi(X_{ji}(\hat{x}_1, \hat{x}_2))}{\phi(\sqrt{\beta_i^x} \hat{x}_i)}, \quad (20)$$

where we use

$$\frac{dQ_i^A}{d\hat{x}_i} = \sqrt{\beta_i^x} \phi(\sqrt{\beta_i^x} \hat{x}_i). \quad (21)$$

The conditions above highlight a key difference with respect to complete information. Even if the platform adjusts the price on side j in response to a variation in the price on side i so as to maintain the expected demand Q_j^A on side j constant, the slope

$$\left. \frac{dp_i^A}{dQ_i^A} \right|_{Q_j^A = \text{const}} = \frac{\mu_i(\hat{x}_i)}{Q_i^A} - \gamma_i \Omega \frac{\phi(X_{ji}(\hat{x}_1, \hat{x}_2))}{\phi(\sqrt{\beta_i^x} \hat{x}_i)}$$

of the side- i 's (inverse) demand curve naturally depends on the intensity γ_i of the side- i 's *own* network effects (in contrast, under complete information, the slope is independent of γ_i). This new effect, of course, plays an important role for the equilibrium prices.

There is a second difference with respect to complete information. The variation in the side- i 's demand that the platform expects to trigger by changing the price p_i^A and then adjusting the price p_j^A to keep the expected side- j demand Q_j^A constant need not coincide with the variation in the side- i demand expected by the marginal agent on side j , which is given by

$$\frac{\partial M_i^A(\hat{x}_1, \hat{x}_2)}{\partial \hat{x}_i} = \sqrt{1 + \Omega^2} \sqrt{\beta_i^x} \phi(X_{ij}(\hat{x}_1, \hat{x}_2)), \quad (22)$$

Comparing (21) with (22), one can then see that the variation in the side- i 's demand expected by the marginal agent on side j differs from the variation expected by the platform when $\Omega \neq 0$. This effect in turn impacts the adjustment in the side- j price that the platform must undertake to maintain the side- j expected demand constant, as it can be observed from the following decomposition:

$$\left. \frac{dM_i^A}{dQ_i^A} \right|_{Q_j^A = \text{const}} = \frac{\frac{\partial M_i^A(\hat{x}_1, \hat{x}_2)}{\partial \hat{x}_i}}{\frac{dQ_i^A}{d\hat{x}_i}} = \sqrt{1 + \Omega^2} \frac{\phi(X_{ji}(\hat{x}_1, \hat{x}_2))}{\phi(\sqrt{\beta_i^x} \hat{x}_i)} \neq 1 \text{ if } \Omega \neq 0. \quad (23)$$

The above two effects, combined, lead to the following first-order condition

$$p_i^A + \left[\frac{\mu_i(\hat{x}_i)}{Q_i^A} - \gamma_i \frac{dM_j^A}{dQ_i^A} \Big|_{Q_j^A = \text{const}} \right] Q_i^A + \left[\gamma_j \frac{dM_i^A}{dQ_j^A} \Big|_{Q_j^A = \text{const}} \right] Q_j^A = 0 \quad (24)$$

where the first bracket term is the change in p_i^A for one unit of extra sale on side i , while the second bracket term is the change in p_j^A required to maintain the expected side- j demand constant. The following proposition combines the above observations into a formula for the monopolist's equilibrium prices that will turn useful when considering competition between the two platforms (the proof follows from the arguments above):

Proposition 1 *The monopolist's profit-maximizing prices, expressed as a function of the demand thresholds they induce, satisfy the following conditions:*

$$p_i^A = \mu_i(\hat{x}_i) - \gamma_i \left[\frac{\Omega \frac{\phi(X_{ji}(\hat{x}_1, \hat{x}_2))}{\phi(\sqrt{\beta_i^x} \hat{x}_i)}}{\phi(\sqrt{\beta_i^x} \hat{x}_i)} \right] Q_i^A - \gamma_j \left[\frac{\sqrt{1 + \Omega^2} \frac{\phi(X_{ij}(\hat{x}_1, \hat{x}_2))}{\phi(\sqrt{\beta_i^x} \hat{x}_i)}}{\phi(\sqrt{\beta_i^x} \hat{x}_i)} \right] Q_j^A \quad i = 1, 2 \quad (25)$$

where \hat{x}_1 and \hat{x}_2 are implicitly defined by the system of equations given by (14), $\mu_i(\hat{x}_i)$ denotes the inverse-semi-elasticity of the stand-alone demand curves, as defined in (7), and $M_j^A(\hat{x}_1, \hat{x}_2) \equiv \Phi(X_{ji}(\hat{x}_1, \hat{x}_2))$ denotes the side- j demand, as expected by the side- i marginal agent.

The first term in the price equation (25), which corresponds to the inverse semi-elasticity of the demand curves in the absence of network effects, expressed in terms of thresholds as opposed to prices, is completely standard and entirely driven by the distribution of the estimated stand-alone valuations. In our model it depends on the information structure only because the latter also affects the distribution of the estimated stand-alone valuations.

The third-term in (25) captures the familiar extra cost of raising prices in a two-sided market due to a reduction of demand (or equivalently of price) on the other side. When side j benefits from the presence of side i , that is, when $\gamma_j > 0$, this term is known to contribute negatively to the price charged by the monopolist (see e.g., Armstrong, 2006). As discussed above, the novelty relative to complete information comes from the fact that the variation in the side- i demand that the platform expects to trigger by raising p_i^A now differs from the variation expected by the marginal agent on side j . This novel effect is captured in the bracket in the third term, which measures the sensitivity of the beliefs of the marginal agent on side j to changes in the average demand on side i .

The second term in (25) is absent under complete information. As explained above, this term originates in the fact that a variation in the side- i demand now implies a variation in the side- i 's expectation about side- j 's participation (this despite the fact that, from the platform's perspective, the side- j expected demand does not change, given the adjustment in the side- j price). Whether this new term contributes positively or negatively to side- i own price elasticity (and thus ultimately to the monopolist's profit-maximizing price) depends on the interaction between (a) the sign of side- i

network effects, γ_i , and (b) the sign of the correlation between the two sides' preferences (Formally, the sign of this new term is the sign of $\gamma_i \rho_\theta$). For a given increase in expected demand Q_i^A , the extra adjustment in the side- i price that the platform must undertake due to this novel effect is given by $\gamma_i \left. dM_j^A / dQ_i^A \right|_{\hat{x}_j = \text{const}}$, which corresponds to the change in the network effects expected by the marginal consumer. To understand this, recall that, by lowering the price p_i^A , the monopolist raises the threshold \hat{x}_i . Equivalently, it lowers the estimated stand-alone valuation of the marginal agent who is just indifferent between joining and staying home. When valuations are positively correlated between the two sides, this means that the new marginal agent will also expect that fewer agents from the opposite side will like the platform and thus join. When side i values positively the participation of the side- j agents, this new effect thus reduces the elasticity of the demand on side i and thus contributes to a higher optimal price.

It is interesting to contrast our results with the analysis in Weyl (2010). In that paper, information is complete but consumers are heterogenous in the importance that they assign to network effects. This possibility can be captured in our model by letting α_1 and α_2 go to infinity, with $\rho_x = 0$, but then allowing the coefficient γ_{il} to vary across agents. To preserve the property that the heterogeneity among the agents is parametrized by x_{il} , then let $\gamma_i(x_{il}) = \mathbb{E}[\tilde{\gamma}_{il} \mid x_{il}]$ and assume that $\kappa_i x - 2\gamma_i(x)$ is increasing in x , so as to preserve the threshold property of the demand curves. Then, both in Weyl (2010) and in our model, the intensity of the network effects is correlated with the perceived stand-alone valuations:

$$\begin{aligned} \mathbb{E}[\tilde{u}_{il}^A + \tilde{\gamma}_{il} \tilde{m}_j^A \mid x_{il}] &= s_i - \frac{\kappa_i}{2} x_{il} + \gamma_i(x_{il}) Q_j^A \text{ with heterogenous network effects,} \\ \mathbb{E}[\tilde{u}_{il}^A + \tilde{\gamma}_{il} \tilde{m}_j^A \mid x_{il}] &= s_i - \frac{\kappa_i}{2} x_{il} + \gamma_i M_j^A(x_{il}, \hat{x}_j) \text{ with dispersed information.} \end{aligned}$$

The equilibrium prices with heterogenous network effects are then given by

$$p_i^A = \underbrace{\mu_i(\hat{x}_i) - \left[\gamma_i'(\hat{x}_i) \frac{d\hat{x}_i}{dQ_i^A} Q_j^A \right]}_{\mu^{\mathcal{I}}} Q_i^A - \hat{\gamma}_j Q_j^A$$

where $\hat{\gamma}_j = \gamma_j(\hat{x}_j)$ and where the term $\mu^{\mathcal{I}}$ corresponds to what in Weyl is called *classical market power*. Notice that the market power $\mu^{\mathcal{I}}$ differs from the usual stand-alone market power due to the correlation between the stand-alone valuations and the importance of the network effects. In our model, a similar formula obtains under dispersed information, but with different interpretations of $\mu^{\mathcal{I}}$ and $\hat{\gamma}_j$. In our model,

$$\mu^{\mathcal{I}} = \mu_i(\hat{x}_i) - \left[\gamma_i \left. \frac{dM_j^A}{dQ_i^A} \right|_{Q_j^A = \text{const}} Q_i^A \right]$$

differs from the usual stand-alone market power index $\mu_i(\hat{x}_i)$ because of the correlation between the stand-alone valuations across the two sides, as opposed to the correlation between stand-alone valuations and the importance of the network effects, within the same side. Interestingly, the sign

of $\mu^{\mathcal{I}} - \mu_i(\hat{x}_i)$ in our model depends on two primitive variables, the sign of the network effects, γ_i , and the sign of the correlation coefficient, ρ_θ .

Next, consider the term $\hat{\gamma}_j$. As pointed out in Weyl (2010), this term reflects the fact that the monopoly internalizes the effect of variations in side- i participation on the utility of the *marginal* consumer on the opposite side. Because $\hat{\gamma}_j$ differs from the average importance $\mathbb{E}\{\gamma_j(\hat{x}_j) \mid x_{jl} < \hat{x}_j\}$ assigned to the network effects by the participating agents on side j , the monopolist's optimal price exhibits a distortion along the lines of Spence (1975). In our model too

$$\hat{\gamma}_j = \gamma_j \left. \frac{dM_i^A}{dQ_i^A} \right|_{Q_j^A = \text{const}},$$

differs from the average value of the expected network effects among the side- j 's participants, which is constant at γ_j .²⁵ However, in our model, the reason for the distortion is very different from the one in Spence (1975). It originates in the fact that the monopolist accounts for the discrepancy between its own beliefs and the beliefs of the marginal agents about participation rates on each of the two sides of the market.

5 Competition

We now reintroduce platform B and examine the outcome of the duopoly game where the two platforms simultaneously compete in prices on each side, assuming full participation²⁶.

Consider the continuation game starting in stage 2 given the prices $(p_1^A, p_2^A, p_1^B, p_2^B)$. Each agent l from each side $i = 1, 2$ chooses platform A when

$$\mathbb{E}[\tilde{u}_{il}^A - \tilde{u}_{il}^B \mid x_{il}] + \gamma_i \mathbb{E}[\tilde{m}_j^A - \tilde{m}_j^B \mid x_{il}] > p_i^A - p_i^B \quad (26)$$

and platform B when the inequality is reversed. Using $m_i^A + m_i^B = 1$, $i = 1, 2$, and (3), Condition (26) can be rewritten as

$$-\kappa_i x_{il} + 2\gamma_i \mathbb{E}[\tilde{m}_j^A \mid x_{il}] - \gamma_i > p_i^A - p_i^B.$$

Now suppose that each agent l from side $j \neq i$ follows a *threshold* strategy according to which he chooses platform A if $x_{jl} < \hat{x}_j$ and B if $x_{jl} > \hat{x}_j$. When this is the case, the measure of agents from side j on platform A is a decreasing function of θ_j and is given by $m_j^A = \Pr(\tilde{x}_{jl} \leq \hat{x}_j \mid \theta_j)$. Given the expectation that each agent from side $j \neq i$ follows such a strategy, each agent l from side i then finds it optimal to choose platform A if

$$-\kappa_i x_{il} + 2\gamma_i \Pr(\tilde{x}_{jl} \leq \hat{x}_j \mid x_{il}) - \gamma_i > p_i^A - p_i^B. \quad (27)$$

²⁵This can be seen by applying the law of iterated expectations.

²⁶As it will become clear below, full participation can be justified by assuming that the stand-alone valuations are sufficiently high—see Proposition 3.

Under Condition (M), the left hand side in (27) is decreasing in x_{il} . Applying the same logic to each side, we then conclude that a monotone continuation equilibrium is characterized by a pair of thresholds (\hat{x}_1, \hat{x}_2) that jointly solve

$$-\kappa_i \hat{x}_i + 2\gamma_i M_j^A(\hat{x}_1, \hat{x}_2) - \gamma_i = p_i^A - p_i^B \quad i, j = 1, 2, j \neq i. \quad (28)$$

Note that the left-hand side of (28) is the gross payoff differential of joining platform A relative to platform B for the marginal agent \hat{x}_i on side i when users on both sides follow threshold strategies with respective cutoffs \hat{x}_1 and \hat{x}_2 .

Recognizing that

$$-\kappa_i \hat{x}_i + 2\gamma_i M_j^A(\hat{x}_1, \hat{x}_2) - \gamma_i = 2G_i(\hat{x}_1, \hat{x}_2) - 2s_i - \gamma_i$$

where G_i are the functions defined above for the monopolist, we then have that many of the properties identified above for the monopolist case carry over to the duopoly case. In particular, for any vector of prices $p = (p_1^A, p_2^A, p_1^B, p_2^B)$, there always exists a solution to the system of conditions given by (28), which implies that a threshold continuation equilibrium always exists. Furthermore, under Condition (Q), this continuation equilibrium is the unique continuation equilibrium, which implies that we can associate to any vector of prices a unique system of demands given, in each state $\theta = (\theta_1, \theta_2)$ by

$$m_i^A = \Phi(\sqrt{\beta_i^\eta}(\hat{x}_i - \theta_i)) = 1 - m_i^B \quad i = 1, 2.$$

Thus consider the choice of prices by the two platforms. For any $p = (p_1^A, p_2^A, p_1^B, p_2^B)$, we have $Q_i^A = \mathbb{E}[\tilde{m}_i^A] = \Phi(\sqrt{\beta_i^x} \hat{x}_i)$ and the two platforms' profits are equal to

$$\Pi^A(p_1^A, p_2^A, p_1^B, p_2^B) = \sum_{i=1,2} p_i^A \Phi(\sqrt{\beta_i^x} \hat{x}_i)$$

and

$$\Pi^B(p_1^A, p_2^A, p_1^B, p_2^B) = \sum_{i=1,2} p_i^B (1 - \Phi(\sqrt{\beta_i^x} \hat{x}_i))$$

with the thresholds (\hat{x}_1, \hat{x}_2) uniquely defined by the system (28).

Now fix (p_1^B, p_2^B) and consider the choice of prices by platform A . Given the bijective relationship between (p_1^A, p_2^A) and (\hat{x}_1, \hat{x}_2) given by

$$p_i^A = p_i^B + 2G_i(\hat{x}_1, \hat{x}_2) - 2s_i - \gamma_i$$

we have that the prices (p_1^A, p_2^A) constitute a best-response for platform A if and only if the corresponding thresholds (\hat{x}_1, \hat{x}_2) solve the following problem:

$$\max_{(\hat{x}_1, \hat{x}_2)} \hat{\Pi}^A(\hat{x}_1, \hat{x}_2) \equiv \sum_{i=1,2} [p_i^B + 2G_i(\hat{x}_1, \hat{x}_2) - 2s_i - \gamma_i] \Phi(\sqrt{\beta_i^x} \hat{x}_i) \quad (29)$$

Arguments similar to those for the monopolist case then easily permit us to verify that, under Condition (Q), for any vector of prices (p_1^B, p_2^B) , the prices (p_1^A, p_2^A) that maximize platform A 's profits must be a solution to the system of first-order conditions given by

$$\begin{aligned} [p_i^B + 2G_i(\hat{x}_1, \hat{x}_2) - 2s_i - \gamma_i] \sqrt{\beta_i^x} \phi\left(\sqrt{\beta_i^x} \hat{x}_i\right) + 2 \frac{\partial G_i(\hat{x}_1, \hat{x}_2)}{\partial x_i} \Phi\left(\sqrt{\beta_i^x} \hat{x}_i\right) \\ + 2 \frac{\partial G_j(\hat{x}_1, \hat{x}_2)}{\partial x_i} \Phi\left(\sqrt{\beta_j^x} \hat{x}_j\right) = 0. \end{aligned} \quad (30)$$

The above conditions are the duopoly analogs of the optimality conditions (18) for the monopoly case; they describe the relation between the profit-maximizing thresholds and the corresponding prices. Following steps similar to those in the previous section, we can then show that the combination of optimal prices and corresponding thresholds for platform A must satisfy the following conditions

$$p_i^A = \kappa_i \frac{Q_i^A}{\frac{dQ_i^A}{d\hat{x}_i}} - 2\gamma_i \left(\frac{\frac{\partial M_i^A(\hat{x}_1, \hat{x}_2)}{\partial \hat{x}_i}}{\frac{dQ_i^A}{d\hat{x}_i}} \right) Q_i^A + 2\gamma_j \left(\frac{\frac{\partial M_i^A(\hat{x}_1, \hat{x}_2)}{\partial \hat{x}_i}}{\frac{dQ_i^A}{d\hat{x}_i}} \right) Q_j^A \quad (31)$$

along with $p_i^A = p_i^B + 2G_i(\hat{x}_1, \hat{x}_2) - 2s_i - \gamma_i$, $i = 1, 2$. The advantage of the above representation is that it highlights the analogy with the monopolist's case (the only difference is that the optimality conditions now apply to the residual demands). It also permits us to identify the unique prices that are sustained in a symmetric equilibrium.

Proposition 2 *In the unique symmetric equilibrium, the prices that both platforms charge on each side are given by*

$$p_i^* = \mu_i^d(0) + \gamma_i \Omega - \gamma_j \sqrt{1 + \Omega^2} \quad i, j = 1, 2, \quad j \neq i, \quad (32)$$

where $\mu_i^d(0)$ is the inverse semi-elasticity of the stand-alone residual demand and where Ω is the coefficient of mutual forecastability between the two sides.

As in the monopolist's case, the first term in (32) is the inverse semi-elasticity of the component of the demand on side i that comes from the stand-alone valuations, accounting for the relation between information and estimated valuations. Notice that it coincides with the equilibrium price in the absence of network effects (see (9)).

The last two terms in (32) capture the interaction between the network effects and the dispersion of information. To appreciate the role of these terms observe that, under complete information, the equilibrium duopoly prices are given by

$$p_i^c = \mu_i^{d;c}(0) - \gamma_i, \quad i = 1, 2, \quad (33)$$

where $\mu_i^{d;c}(0)$ is the complete-information inverse semi-elasticity of the component of the side- i demand that comes from the stand-alone valuations.²⁷ The term $\gamma_i \Omega$ in (32) captures the effects of

²⁷Note that this formula is qualitative the same as in Armstrong (2006). It can be obtained in our setting by

dispersed information on side- i own-price elasticity. As in the monopolist's case, whether this term contributes positively or negatively to the equilibrium prices depends on the sign of the network effects γ_i on side i and on the correlation ρ_θ between the stand-alone valuations on the two sides (recall that $\text{sign}(\Omega) = \text{sign}(\rho_\theta)$). Finally, the third term in (32) captures the cost of increasing the price on side i due to the effect that this has on the platform's profits on the other side of the market. As in the case of complete-information, this effect contributes to a lower equilibrium price when side j benefits from the presence of side i , i.e., when $\gamma_j > 0$, and to a higher price when $\gamma_j < 0$. Contrary to complete information, though, the impact of this effect now depends on the ability of side- j agents to forecast variations in the side- i demand, which, as discussed above, depend on the coefficient of mutual forecastability Ω^2 (see also the discussion below).

We summarize the implications of the above result in the following corollary:

Corollary 1 *As in the complete-information case, equilibrium duopoly prices (i) increase with the inverse-semi-elasticity of the component of the demand that comes from the estimated stand-alone valuations and (ii) decrease with the intensity of the network effect from the opposite side. However, contrary to complete information, equilibrium prices under dispersed information (a) increase with the intensity of the own-side network effects when stand-alone valuations are positively correlated between the two sides, and (b) decrease when they are negatively correlated.*

A second important observation is that, despite the sophistication of the strategic effects at play, the formula for the equilibrium duopoly prices is extremely simple. In particular, fixing the ex-ante distribution of the estimated stand-alone valuations (the first term in the price equation (32)), the equilibrium price on each side depends on the properties of the information structure only through the *coefficient of mutual forecastability* Ω . Recall that

$$\Omega = \frac{\rho_x}{\sqrt{1 - \rho_x^2}}.$$

As anticipated above, the sign of Ω is what determines whether an agent becomes more or less optimistic about the other side's participation as his appreciation for the platform's product increases. As a result, the sign of Ω is what determines whether the equilibrium price p_i on each side increases or decreases with the intensity γ_i of that side's own network effects. In contrast, when it comes to the impact on equilibrium prices of the intensity of the network effects γ_j on the opposite side, what matters is only the square of Ω . To interpret this result, use the variance decomposition $\text{var}(\tilde{x}_{il}) -$

considering the limit where $\alpha_i \rightarrow \infty$, $i = 1, 2$, which suffices to eliminate any aggregate uncertainty. The precise value of

$$\mu_i^{d;c}(0) = \frac{\sqrt{\text{var}[\tilde{V}_{il}]}{2\phi(0)} = z_i \frac{\rho_i \sqrt{\beta_i^\eta / \beta_i^\varepsilon}}{\beta_i^\eta 2\phi(0)}$$

then depends on whether or not one assumes that the agents know their own stand-alone valuations (and if not, on the correlation between the taste-shocks $\tilde{\varepsilon}_i$ and the noise shocks $\tilde{\eta}_i$). Irrespective of what one assumes on $(z_i, \beta_i^\eta, \beta_i^\varepsilon)$, in the absence of aggregate uncertainty $\Omega = 0$, which gives the formula in (33).

$var \left[\tilde{x}_{il} - \rho_x \sqrt{\frac{\beta_j^x}{\beta_i^x}} \tilde{x}_{jl'} \right] = \rho_x^2 var(\tilde{x}_{il})$ to see that

$$\Omega^2 = \frac{var(\tilde{x}_{il})}{var[\tilde{x}_{il} - \mathbb{E}[\tilde{x}_{il}|\tilde{x}_{jl'}]]} - 1.$$

Hence Ω^2 measures the ability of side j to forecast variations in participation decisions on side i triggered by variations in prices.²⁸ It is then natural that the sensitivity of the equilibrium price on side i to the intensity γ_j of the network effects on the opposite side depends on Ω only through Ω^2 .

The above properties also suggest that equilibrium prices need not be too sensitive to the specific way the information is distributed across the two sides. Fixing again the ex-ante distribution of the estimated stand-alone valuations (equivalently, the inverse semi-elasticity of the component of the demands that comes from the stand-alone valuations), we have that any two information structures that result in the same coefficient of mutual forecastability yield the same equilibrium prices.

This observation is particularly sharp in the case of a market whose primitives are perfectly symmetric under complete information. That is, consider a market where both the intensity of the network effects and the inverse semi-elasticity of the stand-alone demand is the same across the two sides, i.e., $\gamma_1 = \gamma_2 = \gamma$ and $\mu_i^d = \mu^d$, $i = 1, 2$. Using (33), we then have that the complete-information equilibrium prices are given by

$$p_i^c = \mu^{d;c}(0) - \gamma, \quad i = 1, 2.$$

Not surprising, these prices are the same across the two sides. Perhaps more surprising, the equilibrium prices continue to be the same across the two sides, even when the distribution of information is not symmetric. This is because, holding fixed the distribution of the estimated stand-alone valuations, and assuming that the intensity of the network effect is the same across the two sides, a variation in the quality of information on side i has an identical effect on the elasticity of demand on each of the two sides.

To gauge some intuition, consider the case where preferences are perfectly correlated between the two sides so that $\tilde{\theta}_1 = \tilde{\theta}_2$ almost surely (in which case $\alpha_1 = \alpha_2$ and $\rho_\theta = 1$). Now suppose that information is very precise on side 1, while very imprecise on side 2, so that $\beta_1^\eta \rightarrow \infty$ while $\beta_2^\eta \rightarrow 0$. Because participation decisions on side 2 do not vary much with the state θ_2 , the value of the information held by the side-1 agents is pretty much the same as if side-1 was itself uninformed about the distribution of the side-2's valuations.

More generally, the result in Proposition 2 implies that shocks that affect the agents' ability to forecast the cross-sectional distribution of valuations in an asymmetric way across the two sides have nonetheless a symmetric effect on the equilibrium prices, as long as the intensity of the network effect is the same across the two sides. This is because, holding fixed the ex-ante distribution of

²⁸Note that Ω^2 is reminiscent of the "coefficient of fit" R^2 for the regression of \tilde{x}_i on \tilde{x}_j . The difference is in the denominator, which here is the variance of the residual, while it is the total variance in R^2 .

estimated stand-alone valuations, the value that each side assigns to being able to predict the *realized* distribution of estimated stand-alone valuations on the opposite side comes entirely from its ability to coordinate its participation decisions with those on the opposite side. When the importance of the network effects is the same across the two sides (that is, when $\gamma_1 = \gamma_2$), the two platforms then equalize the prices over the two sides, despite possibly asymmetries in the distribution of information.

We conclude this section with two results that show that, under plausible additional assumptions, the equilibrium prices characterized above (along with the participation decisions they induce) continue to remain equilibrium outcomes when agents can choose to "opt out" of the market, or to "multihome" by joining both platforms. These results should be interpreted as (minimal) robustness checks aimed at showing that the above results are not unduly driven by the choice of simplifying the analysis by abstracting from these possibilities. In future work, it would be interesting to extend the analysis to markets where multihoming and partial market-coverage occur in equilibrium.

We start with the following result that pertains our assumption of full market-coverage:

Proposition 3 *There exist finite $(\underline{s}_i)_{i=1,2}$ such that, for any $(s_i)_{i=1,2}$ with $s_i > \underline{s}_i$, $i = 1, 2$, the equilibrium in the game where agents must join one of the two platforms is also an equilibrium in the game where agents can "opt out" of the market by choosing not to join any platform.*

The reason why the equilibrium prices in the game with compulsory participation need not remain equilibrium prices in the game where agents can opt out of the market is the following. First, when platforms set the prices at the level of Proposition 2, some agents may experience a negative equilibrium payoff and hence prefer to opt out. Because the equilibrium prices p_i^* in Proposition 2 are independent of the levels of the stand-alone valuations (formally, of s_1 and s_2) this possibility can be ruled out by assuming that the marginal agents' equilibrium payoffs are positive, which amounts to assuming that $s_i + \gamma_i/2 \geq p_i^*$, $i = 1, 2$. Under these conditions, no agent finds it optimal to opt out, given that any agent's equilibrium payoff is at least as high as that of the marginal agents.

This condition, however, does not suffice. In fact, platforms may have an incentive to raise one of their prices above the equilibrium levels of Proposition 2 if they expect that, by inducing some agents to opt out, their demand will fall less than that of the other platform, relative to the case where participation is compulsory. Consider, for example, a deviation by platform A to a vector of prices (p_1^A, p_2^A) with $p_1^A > p_1^*$. Now suppose that, in the unique continuation equilibrium of the game where participation is compulsory, the payoff of the marginal agent $\hat{x}_1(p_1^A, p_2^A, p_1^*, p_2^*)$ on side 1 is negative (that is, below his outside option). This means that, in the game where participation is voluntarily, some agents in a neighborhood of $\hat{x}_1(p_1^A, p_2^A, p_1^*, p_2^*)$ may now decide to opt out. Note that some of these agents were joining platform B in the game with compulsory participation. When network effects are positive, this in turn implies that such a deviation may now be profitable for firm A if the measure of agents on side 1 who would have joined platform B in the game with compulsory participation and now decide to opt out is larger than the measure of agents who would have joined

platform A and now opt out. That is, when the platform expects a larger drop in the rival's demand than in its own (relative to the case where participation is compulsory), then a deviation that was not profitable in the game where participation is compulsory may now become profitable. For this to be the case, however, it must be that the intensity of the network effects is sufficiently strong to prevail on the direct effect coming from the stand-alone valuations. The proof in the Appendix shows that this is never the case when s_1 and s_2 are sufficiently large.

Next, consider the possibility that agents multihome by choosing to join both platforms. We assume that, by doing so, each agent l from each side i obtains a gross payoff equal to $(2 - \kappa)s_i + \gamma_i(m_j^A + \mu_j^B)$, where μ_j^B is the measure of agents from side j who join platform B without joining platform A (to avoid double counting), and where $\kappa \in [0, 1]$ denotes the loss of utility coming from combining the two products.²⁹

We then have the following result:

Proposition 4 *Consider the variant of the game where agents from each side of the market can multihome, as described above. For any vector of prices $(p_1^A, p_2^A, p_1^B, p_2^B)$ such that $p_i^A + p_i^B \geq \gamma_i + 2(1 - \kappa_i)s_i$, $i = 1, 2$, there exists a continuation equilibrium where each agent from each side singlehomes. Conversely, such a continuation equilibrium fails to exist for any vector of prices for which $p_i^A + p_i^B < \gamma_i + 2(1 - \kappa_i)s_i$, for some $i \in \{1, 2\}$.*

As we show in the Appendix, the condition in the proposition guarantees that *any* agent who expects all other agents to singlehome (according to the same threshold rule as in the game in which multihoming is not possible) prefers to join his most preferred platform to multihoming. As the proposition makes clear, the condition is also necessary, in the sense that, when it is violated, then in any continuation equilibrium some agents necessarily multihome. The following corollary is then an immediate implication of the above result:

Corollary 2 *Let (p_1^*, p_2^*) be the equilibrium prices in the game where multihoming is not possible, as defined in (32), and assume that $p_i^* \geq \gamma_i + 2(1 - \kappa_i)s_i$, $i = 1, 2$. Suppose that platforms cannot set negative prices. The equilibrium in the game where agents are not allowed to multihome then continues to be an equilibrium in the game where multihoming is possible.*

Because equilibrium prices are increasing in the ex-ante dispersion of the estimated stand-alone valuations (formally in $\mu_i^d(0)$) and because such dispersion measures the degree of differentiation between the two platforms, the result in Corollary 2 is consistent with the finding in Armstrong and Wright (2007) that strong product differentiation on both sides of the market implies that agents have no incentive to multihome when prices are restricted to be non-negative (As argued in that paper, and in other contexts as well, the assumption that prices must be non-negative can be justified by the fact that negative prices can create moral hazard and adverse selection problems).

²⁹Note that $(2 - \kappa_i)s_i + \gamma_i(m_j^A + \mu_j^B) = u_i^A + u_i^B - \kappa_i s_i + \gamma_i(m_j^A + \mu_j^B)$.

Together, the results in Proposition 3 and Corollary 2 imply that, when the stand-alone valuations of the marginal agents are neither too high nor too low (intermediate s_i) and when the two platforms are seen as sufficiently differentiated on both sides of the market (the ex-ante distribution of estimated stand-alone valuations is sufficiently diffuse), then the unique symmetric equilibrium of the baseline game is also an equilibrium in the more general game where agents can multihome and opt out of the market.

6 Implications for advertising and product design

We now turn to the effects on equilibrium profits of variations in (i) the quality of the agents' information and (ii) the prior distribution from which stand-alone valuations are drawn. These comparative statics results have implications for advertising and product design.

We start by showing how the equilibrium prices depend on the various structural parameters of the model. From Proposition 2, the relevant terms for the equilibrium prices are (a) the inverse semi-elasticities μ_i^d of the stand-alone demands and (b) the coefficient Ω of mutual forecastability. The inverse semi-elasticities of the stand-alone demands (evaluated at the equilibrium prices) are in turn proportional to the dispersion of the estimated stand-alone differentials (see (10)):

$$\text{var}[\tilde{V}_{il}] = z_i^2 \frac{\left(\beta_i^\eta + \rho_i \alpha_i \sqrt{\beta_i^\eta / \beta_i^\varepsilon}\right)^2}{(\alpha_i + \beta_i^\eta) \alpha_i \beta_i^\eta}. \quad (34)$$

As one can see from (34), $\text{var}[\tilde{V}_{il}]$ increases with the correlation ρ_i between the noise $\tilde{\eta}_{il}$ in the agents' information and the idiosyncratic taste shock $\tilde{\varepsilon}_{il}$ in the stand-alone differentials. It also increases with z_i , which parametrizes the overall sensitivity of the agents' stand-alone differentials to common and idiosyncratic shocks ($\tilde{\theta}_i$ and $\tilde{\varepsilon}_{il}$, respectively). Finally, it decreases with β_i^ε , for a higher β_i^ε implies a lower dispersion of idiosyncratic taste shocks.

On the other hand, $\text{var}[\tilde{V}_{il}]$ is typically non-monotone in α_i and in β_i^η . The non-monotonicity with respect to α_i (which parametrizes the precision of the prior about θ_i) reflects the fact that a higher α_i implies a lower dispersion of stand-alone differentials but also a higher precision of the agents' information. Because the latter effect makes the agents respond more to their information, it contributes to a higher dispersion of estimated differentials. The non-monotonicity with respect to the precision β_i^η of the agents' information in turn reflects the fact that, holding constant the correlation coefficient ρ_i , a higher β_i^η implies a lower covariance between the noise in the signals and the idiosyncratic taste shocks in the differentials. Because a lower covariance between the noise in the signals and the taste shock in turn contributes to a lower sensitivity of estimated differentials to the agents' signals, the net effect of a higher β_i^η on $\text{var}[\tilde{V}_{il}]$ is typically non-monotone.

Next, consider the coefficient Ω of mutual forecastability. As illustrated above, Ω is an increasing transformation of the coefficient ρ_x of correlation between signals from the two sides, which in turn determines the two sides' ability to forecast each other. To be precise, we measure the ability of

side i to forecast the information on side j by the (inverse of the) variance of the forecast errors $\tilde{x}_{j\nu} - \mathbb{E}[\tilde{x}_{j\nu}|\tilde{x}_{il}]$, which can be decomposed as follows

$$var[\tilde{x}_{j\nu} - \mathbb{E}[\tilde{x}_{j\nu}|\tilde{x}_{il}]] = var[\tilde{\theta}_j - \mathbb{E}[\tilde{\theta}_j|\tilde{x}_{il}]] + \frac{1}{\beta_j^\eta}. \quad (35)$$

Clearly, the ability of side i to forecast the information (and hence the valuations) on side j increases as the noise in the side- j 's signals decreases (that is, as β_j^η increases). It also increases with its ability to forecast the correlated taste shock $\tilde{\theta}_j$ in the side- j signals, which is inversely proportional to

$$var[\tilde{\theta}_j - \mathbb{E}[\tilde{\theta}_j|\tilde{x}_{il}]] = \left(1 - \rho_\theta^2 \frac{\beta_i^\eta}{\alpha_i + \beta_i^\eta}\right) \frac{1}{\alpha_j} \quad (36)$$

Not surprisingly, the ability of side i to forecast $\tilde{\theta}_j$ increases with $|\rho_\theta|$ and β_i^η , and decreases with α_i .

Building on these observations, we now investigate the firms' incentives to take actions that affect either (i) the agents' ability to estimate their own preferences as well as those of the agents from the opposite side (e.g., through informative advertising and marketing campaigns as well as through personalized disclosures aimed at allowing consumers to learn of their personal match with the product's characteristics), or (ii) the distributions from which the agents' true preferences are drawn (e.g., through product design). We examine each of the two channels separately.

6.1 Advertising campaigns

Think of a software firm entering the market with a new operating system. The firm must decide how much information to disclose to the public about the various features of its product. We think of these disclosures as affecting both the developers' and the end-users' ability to estimate their own stand-alone valuations (both in absolute value and relative to the operating system produced by a rival incumbent firm), as well as their ability to forecast the distribution of valuations on the other side of the market.

Formally, we think of these disclosure and advertising campaigns as affecting the information available to the two sides of the market, for fixed distribution of true stand-alone valuations. That is, fix the parameters $(\alpha_1, \alpha_2, \rho_\theta, \beta_1^\varepsilon, \beta_2^\varepsilon, z_1, z_2)$ defining the prior distribution from which individual stand-alone valuations are drawn and consider the effects on profits of variations in (i) the agents' ability to estimate their own stand-alone valuations (as measured by the inverse of the volatility of the forecast error $var[\tilde{v}_{il} - \tilde{V}_{il}]$), and (ii) their ability to forecast the distribution of stand-alone valuations on the other side of the market (as measured by the inverse of (35)). Hereafter, we isolate the effects of the variations in (i) by looking at changes in the coefficient ρ_i of correlation between the noise $\tilde{\eta}_{il}$ in the signals and the idiosyncratic taste shock $\tilde{\varepsilon}_{il}$. We then isolate the effects of the variations in (ii) by looking at joint changes in $(\beta_i^\eta, \rho_i)_{i=1,2}$ that leave $var[\tilde{v}_{il} - \tilde{V}_{il}]$, $i = 1, 2$ fixed.

We then have the following result:

Proposition 5 *Informative advertising and marketing campaigns that increase the agents' ability to estimate their own stand-alone valuations without affecting their ability to forecast the distribution of (true or estimated) stand-alone valuations on the other side of the market always increase profits.*

Conversely, campaigns that increase the agents' ability to forecast the distribution of (true or estimated) stand-alone valuations on the other side of the market without affecting their ability to estimate their own valuations increase profits if $\rho_\theta(\gamma_1 + \gamma_2) > 0$ and reduce profits otherwise.

The first part of the result is very intuitive. Campaigns that increase the agents' ability to understand their own needs and preferences, without affecting their ability to forecast other agents' preferences, make agents more responsive to their own idiosyncrasies. As such, these campaigns increase the ex-ante dispersion of estimated stand-alone valuations, thus reducing the semi-price elasticity of the part of the demand on each side that comes from the stand-alone valuations. These campaigns are thus similar to those that increase the degree of differentiation between the two platforms under complete information (e.g., Johnson and Myatt, 2006). By reducing the intensity of the competition between the two platforms, such campaigns unambiguously contribute to higher prices and hence to higher equilibrium profits.

Next, consider campaigns whose primary effect is to make agents more informed about what is likely to be "hip" on the other side of the market (formally, that help agents predict the other side's cross-sectional distribution of preferences). As we show in the Appendix, these campaigns impact the coefficient of mutual forecastability Ω , without affecting the ex-ante distribution of estimated stand-alone valuations $\text{var}[\tilde{V}_i]$. From the equilibrium price equation (32), one can then see that, depending on the intensity of the network effects, such campaigns may either increase or decrease the equilibrium prices. Their total effect on equilibrium profits, which in a symmetric equilibrium are given by

$$\Pi^* = \frac{1}{2}(p_1^* + p_2^*) = \frac{1}{2} \left\{ \mu_1^d(0) + \mu_2^d(0) + (\gamma_1 + \gamma_2) \left(\Omega - \sqrt{1 + \Omega^2} \right) \right\}, \quad (37)$$

is then determined by (i) the sign of the total network effects $\gamma_1 + \gamma_2$ and (ii) whether increasing the agents' ability to forecast the distribution of preferences on the other side (which, by (35), corresponds to an increase in the precisions β_i^η , $i = 1, 2$, of the agents' information) increases or decreases the coefficient of mutual forecastability Ω . Because the latter is increasing in the quality of the agents' information β_1^η and β_2^η if and only if preferences are positively correlated between the two sides (that is, if and only if $\rho_\theta > 0$), we then have that the effect of such campaigns on profits is positive if and only if the correlation of tastes between the two sides is of the same sign as the sum of the intensity of the network effects (that is if and only if $\rho_\theta(\gamma_1 + \gamma_2) > 0$).

To better understand this result, recall that the term $\gamma_i\Omega$ in the price equation captures the effect of the dispersion of information on side- i 's own-price elasticity. From the discussion in the previous section, when network effects are positive and preferences are positively correlated between the two sides, then $\gamma_i\Omega$ increases in either of the two sides' quality of information (that is in either β_1^η

and β_2^η). This effect comes from the fact that more precise information on side i makes the marginal agent on both sides more responsive to his private information. When preferences are positively correlated and network effects are positive, this effect in turn contributes to a higher equilibrium price on each side by making each side's demand less elastic.

At the same time, more precise information also implies a higher sensitivity of both demands to variations in prices on the opposite side. These effects, which are captured by the terms $\gamma_j\sqrt{1+\Omega^2}$ in the price equations, contribute negatively to the equilibrium prices. While the net effect on the equilibrium prices on each side then depends on the relative strengths of the network effects γ_1 and γ_2 , the net effect on total profits is unambiguously positive when the sum of the network effects is positive (more generally, when it is of the same sign as the correlation of preferences between the two sides). This is because any loss of profits on one side is more than compensated by an increase in profits on the opposite side, as one can see from (37).

What is interesting about the results in the proposition is that they identify two fairly general channels through which information affects profits, without specifying the particular mechanics by which advertising and marketing campaigns operate. In reality, most campaigns operate through both channels. That is, they impact both the agents' ability to understand their own preferences and their ability to understand what other agents are likely to find attractive. The results in the proposition then indicate that such campaigns unambiguously increase profits in markets where (i) preferences are positively correlated between the two sides and (ii) the sum of the network effects is positive (which is always the case when each side benefits from the presence of the other side). In contrast, in markets where the sum of the network effects is positive but where preferences are negatively correlated between the two sides (or, vice versa), profits may decrease with the agents' ability to forecast other agents' preferences and platforms may find it optimal to conceal part of the information they have.

Note that the above results refer to informative campaigns. They do not apply to campaigns that distort the average perception the agents have about the quality differential between the platforms' products. These campaigns could be modelled in our framework by allowing the platforms to manipulate the mean of the distributions from which the signals are drawn. However, because in our environment platforms do not possess any private information and the agents are fully rational, the effect of such campaigns on profits is unambiguously negative. This is because each agent can always "undo" the manipulation by adjusting the interpretation of the information he receives. As discussed in the "signal-jamming" literature (e.g., Fudenberg and Tirole (1986)), platforms may then be trapped into a situation where they have to invest resources in such campaigns, despite the fact that, in equilibrium, such campaigns have no effect on the agents' decisions.

6.2 Product design

We conclude by considering the effects on profits of changes in the distribution from which the true stand-alone valuations are drawn. As anticipated in the Introduction, these effects—formally captured by variations in the parameters $(\alpha_1, \alpha_2, \rho_\theta, \beta_1^\varepsilon, \beta_2^\varepsilon, z_1, z_2)$ —should be interpreted as the result of product design. For example, an increase in α_1 and α_2 should be interpreted as the choice to enter the market with a product that is more similar to the one provided by the rival incumbent firm. An increase in ρ_θ , instead, should be interpreted as the choice to favor product characteristics that are expected to appeal to both sides. We then have the following result:

Proposition 6 *Fix the quality of the information on each side of the market (that is, fix (β_i^η, ρ_i) , $i = 1, 2$). An increase in the similarity between the two products (as captured by an increase in (α_1, α_2)) always reduces the equilibrium profits. The same is true for a reduction in the cross-sectional heterogeneity of individual preferences (as captured by an increase in $(\beta_1^\varepsilon, \beta_2^\varepsilon)$).*

Conversely, an increase in the alignment of preferences between the two sides (as captured by an increase in ρ_θ) increases profits if $\gamma_1 + \gamma_2 > 0$ and reduces them if $\gamma_1 + \gamma_2 < 0$.

That both a higher similarity in the products and a smaller relevance of dimensions that are responsible for idiosyncratic appreciations contribute negatively to profits is immediate, for they both contribute to fiercer competition on prices.

The result pertaining the decision to favor dimensions that appeal to both sides is less obvious. Observe from the price equation (32) that an increase in the alignment of preferences (which amounts to an increase in the coefficient Ω of mutual forecastability) may increase prices on one side while decreasing prices on the other side. This is true even if each side benefits from the participation of the other side. The net effect of profits is however always positive if the sum of the network effects is positive, while it is negative otherwise. For example, in a market for media outlets, favoring dimensions that appeal to both viewers and advertisers is likely to be profit-enhancing if the viewers' tolerance towards advertisement is high, whereas it may reduce profits otherwise.

7 Conclusions

We examined the effects of dispersed information on market outcomes in a simple, yet flexible, model of platform competition with differentiated products. Dispersed information naturally introduces heterogeneity in the users' expectations about participation rates on both sides of the market.

The analysis identified a novel channel through which the dispersion of information interacts with the network effects in determining the elasticity of the demand on each side and thereby the equilibrium prices. We then showed how equilibrium profits are affected by variations in (i) the prior distribution from which valuations are drawn and (ii) the quality of information available to the market participants. Finally, we used these results to shed light on the platforms' incentives to

invest in product design to align the preferences of the two sides and/or to engage in advertising and marketing campaigns affecting the agents' ability to understand their own needs and/or the distribution of preferences on the other side of the market.

In future work, it would be interesting to extend the analysis in a few directions. First, one may want to consider equilibria without full coverage and/or where some agents multihome. While in the paper we identified conditions that guarantee that, in equilibrium, all agents participate and singlehome, it seems interesting to relax these conditions and investigate the implications for equilibrium prices and market coverage. In the same vein, it would be interesting to construct asymmetric equilibria whereby the two platforms set their prices differently on one or both sides.

Second, one may want to extend the analysis to a dynamic setting and investigate the platforms' incentives to price aggressively at the early stages so as to build a user base as a barrier to entry and to future competition. Such an extension could also be used to shed light on how the platforms' pricing strategies affect the dynamics of individual and social learning and thereby the speed of technology adoption.

Third, it would be interesting to consider the possibility of price discrimination, whereby the platforms grant differential access to the participating population from the opposite side.³⁰

Lastly, interesting new effects are expected to emerge by introducing within-side network effects (e.g., congestion or other negative externalities among software developers and/or collaboration or other positive externalities among software adopters), thus introducing the possibility that participation decisions be determined by the agents' ability to forecast the joint distribution of preferences on the two sides of the market.

References

- Ambrus, A., Calvano E. and M. Reisinger (2013): "Either or Both Competition: A Two-sided Theory of Advertising with Overlapping Viewerships," Working Paper, Duke University.
- Ambrus, A. and R. Argenziano (2009): "Asymmetric Networks in Two-Sided Markets," *American Economic Journal: Microeconomics*, 1(1), 17-52.
- Amelio A. and B. Jullien (2012): "Tying and Freebies in Two-sided Markets," *International Journal of Industrial Organization*, 30(5), 436-446.
- Anderson, S. and R. Renault (2006): "Advertising Content," *American Economic Review* 96,(1), 93-113.
- Anderson, S. and R. Renault (2009): "Comparative Advertising: Disclosing Horizontal Match Information," *RAND Journal of Economics*, 40,(3), 558-581.
- Argenziano, R. (2008): "Differentiated Networks: Equilibrium and Efficiency" *Rand Journal of Economics* 39, 747-769.
- Armstrong, M. (2006): "Competition in Two-Sided Markets," *Rand Journal of Economics* 37, 668-691.

³⁰See Gomes and Pavan (2013) for an analysis in the context of a monopolistic market.

- Armstrong, M. and J. Wright (2007): "Two-Sided Markets, Competitive Bottlenecks and Exclusive Contracts," *Economic Theory* 32, 353-380.
- Athey, S., E. Calvano and J.S. Gans (2012): "The Impact of the Internet on Advertising Markets for News Media," Working Paper, Harvard University.
- Bergemann, D. and A. Bonatti (2011): "Targeting in Advertising Markets: Implications for Offline versus Online Media," *RAND Journal of Economics*, 42, 417-443.
- Cabral L. (2011): "Dynamic Price Competition with Network Effects", *Review of Economic Studies* 78 (1): 83-111.
- Caillaud, B. and B. Jullien (2001): "Competing Cybermediaries," *European Economic Review (Papers & Proceedings)* 45, 797-808.
- Caillaud, B. and B. Jullien (2003): "Chicken & Egg: Competition Among Intermediation Service Providers," *Rand Journal of Economics*, 34, 309-328.
- Carlsson, H., and E. Van Damme (1993): "Global Games and Equilibrium Selection," *Econometrica* 61, 989-1018.
- Damiano, E. and H. Li (2007): "Price Discrimination and Efficient Matching," *Economic Theory*, Vol. 30, 243-263.
- Farrell, J. and P Klemperer (2006): "Coordination and Lock-In: Competition with Switching Costs and Network Effects", in Armstrong, M. and Porter, R. (eds) *Handbook of Industrial Organization*, Vol. 3 (North-Holland), 1967–2072.
- Fudenberg D. and J. Tirole (1986): "A Signal-jamming Theory of Predation," *The RAND Journal of Economics* 17, 366-376.
- Gomes, R. and A. Pavan (2012): "Cross-Subsidization and Matching Design," mimeo Northwestern University.
- Halaburda, H. and Y. Yehezkel (2013): "Platform Competition under Asymmetric Information," *American Economic Journal: Microeconomics*, 5(3): 22-68.
- Hagiu, A. (2006): "Optimal Pricing and Commitment in Two-Sided Markets," *The RAND Journal of Economics* 37(3), 720–737.
- Hagiu, A. (2009): "Two-Sided Platforms: Product Variety and Pricing Structures," *Journal of Economics and Management Strategy* 18, 1011-1043.
- Johnson, J. and D. Myatt (2006): "On the Simple Economics of Advertising, Marketing, and Product Design" *American Economic Review* 96,(3), 756-784
- Jullien, B. (2007): "Two-Sided Markets and Electronic Intermediation," in G. Illing and M. Peitz (eds.), *Industrial Organization and the Digital Economy*, MIT Press.
- Katz M., and C.Shapiro (1985): "Network Externalities, Competition and Compatibility", *American Economic Review*, 75, 424–440.
- Mitchell, M. and A. Skrzypacz (2006): "Network Externalities and Long-Run Market Share", *Economic Theory*, 29, 621–648.
- Morris, S. and H.S. Shin (2003): "Global Games—Theory and Applications," in *Advances in Eco-*

- nomics and Econometrics*, 8th World Congress of the Econometric Society (M. Dewatripont, L. Hansen, and S. Turnovsky, eds.), Cambridge University Press.
- Rochet, J.C. and J. Tirole (2003): "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association* 1, 990-1029.
- Rochet, J.C. and J. Tirole (2006): "Two-Sided markets: a Progress Report," *The RAND Journal of Economics* 37, 645-667.
- Rysman, M. (2009): "The Economics of Two-Sided Markets," *Journal of Economic Perspectives* 23(3), 125-143.
- Spence, A.M. (1975): "Monopoly, Quality, and Regulation," *Bell Journal of Economics*, 6(2), 417-29.
- Veiga, A. and E. G. Weyl (2011): "Multidimensional Product Design," mimeo, the University of Chicago.
- Weyl, E. G. (2010): "A Price Theory of Multi-sided Platforms," *American Economic Review*, 100(4): 1642-72.
- Whinston, M. (1990): "Tying, Foreclosure, and Exclusion," *American Economic Review* 80(4), 837-859.
- White, A. and E. G. Weyl (2012) "Insulated Platform Competition", mimeo University of Chicago.

A Appendix

Proof of Lemma 1. Fix (p_1^A, p_2^A) . Under Assumption M, $G_i(x_1, x_2)$ is a continuous decreasing³¹ function onto \mathbb{R} of \hat{x}_i . Thus for any x_2 there exists a unique value $x_1 = \xi_1(x_2)$ that solves $G_1(\xi_1(x_2), x_2) = p_1^A$. Thus consider the function

$$F(x_2) \equiv G_2(\xi_1(x_2), x_2) - p_2^A.$$

This is a continuous function, positive for x_2 small enough and negative for x_2 large enough. Thus a solution to $F(x_2) = 0$ always exists, which establishes the result. ■

Proof of Lemma 2. To fix ideas, we assume here that $\gamma_1 \geq 0$. The proof for the case where $\gamma_1 < 0 \leq \gamma_2$ is symmetric to the one for the case where $\gamma_2 < 0 \leq \gamma_1$ which is covered below. Consider again the function $F(x_2) \equiv G_2(\xi_1(x_2), x_2)$ introduced in the proof of Lemma 1, where $\xi_1(x_2)$ is the unique solution to $G_1(\xi_1(x_2), x_2) = p_1^A$. From the implicit function theorem, and given that $\partial G_i(x_1, x_2) / \partial x_i < 0$, we have that the sign of $\frac{dF(x_2)}{dx_2}$ is the sign of

$$\frac{\partial G_2(\xi_1(x_2), x_2)}{\partial x_1} \frac{\partial G_1(\xi_1(x_2), x_2)}{\partial x_2} - \frac{\partial G_2(\xi_1(x_2), x_2)}{\partial x_2} \frac{\partial G_1(\xi_1(x_2), x_2)}{\partial x_1}.$$

³¹To see this note that $\frac{\partial G_i(x_1, x_2)}{\partial x_i} = -\kappa_i/2 - \gamma_i \Omega \sqrt{\frac{\alpha_i \beta_i^x}{\alpha_i + \beta_i^x}} \phi(X_{ji}(x_1, x_2))$. Hence, when $\gamma_i \Omega \geq 0$, $\frac{\partial G_i(x_1, x_2)}{\partial x_i} < 0$ while for $\gamma_i \Omega < 0$, $\frac{\partial G_i(x_1, x_2)}{\partial x_i} \leq -\kappa_i/2 - \gamma_i \Omega \sqrt{\frac{\alpha_i \beta_i^x}{\alpha_i + \beta_i^x}} \phi(0)$ which is again negative by Assumption M.

Using

$$\begin{aligned}\frac{\partial G_i(x_1, x_2)}{\partial x_i} &= -\kappa_i/2 - \gamma_i \Omega \sqrt{\beta_i^x} \phi(X_{ji}(x_1, x_2)) \\ \frac{\partial G_i(x_1, x_2)}{\partial x_j} &= \gamma_i \sqrt{1 + \Omega^2} \sqrt{\beta_j^x} \phi(X_{ji}(x_1, x_2))\end{aligned}$$

after some algebra, we obtain that

$$\begin{aligned}& \frac{\partial G_2(x_1, x_2)}{\partial x_1} \frac{\partial G_1(x_1, x_2)}{\partial x_2} - \frac{\partial G_2(x_1, x_2)}{\partial x_2} \frac{\partial G_1(x_1, x_2)}{\partial x_1} \\ &= \left(\gamma_1 \gamma_2 \sqrt{1 + \Omega^2} \sqrt{\beta_1^x \beta_2^x} \phi(X_{12}(x_1, x_2)) - \frac{\kappa_2}{2} \gamma_1 \Omega \sqrt{\beta_1^x} \right) \phi(X_{21}(x_1, x_2)) \\ & \quad - \frac{\kappa_1}{2} \gamma_2 \Omega \sqrt{\beta_2^x} \phi(X_{12}(x_1, x_2)) - \frac{\kappa_1 \kappa_2}{4}.\end{aligned}\tag{38}$$

Now we claim that, under Condition Q, the expression in (38) is strictly negative for any (x_1, x_2) . To see this, suppose, on the contrary, that there exists (x_1, x_2) for which the sign of the expression in (38) is nonnegative. Consider first the case where $\gamma_1, \gamma_2, \Omega \geq 0$. Then for the expression in (38) to be nonnegative, it must be that

$$\gamma_1 \gamma_2 \frac{\Omega^2 \sqrt{\beta_1^x \beta_2^x}}{\rho_x^2} \phi(X_{12}(x_1, x_2)) - \frac{\kappa_2}{2} \gamma_1 \Omega \sqrt{\beta_1^x} > 0$$

which in turn implies that

$$\begin{aligned}& \frac{\partial G_2(x_1, x_2)}{\partial x_1} \frac{\partial G_1(x_1, x_2)}{\partial x_2} - \frac{\partial G_2(x_1, x_2)}{\partial x_2} \frac{\partial G_1(x_1, x_2)}{\partial x_1} \\ & \leq \left(\gamma_1 \gamma_2 \sqrt{1 + \Omega^2} \sqrt{\beta_1^x \beta_2^x} \phi(X_{12}(x_1, x_2)) - \frac{\kappa_2}{2} \gamma_1 \Omega \sqrt{\beta_1^x} \right) \phi(0) \\ & \quad - \frac{\kappa_1}{2} \gamma_2 \Omega \sqrt{\beta_2^x} \phi(X_{12}(x_1, x_2)) - \frac{\kappa_1 \kappa_2}{4}.\end{aligned}\tag{39}$$

Because the right-hand side of (39) can also be rewritten as

$$\left(\gamma_1 \gamma_2 \sqrt{1 + \Omega^2} \sqrt{\beta_1^x \beta_2^x} \phi(0) - \frac{\kappa_1}{2} \gamma_2 \Omega \sqrt{\beta_2^x} \right) \phi(X_{12}(x_1, x_2)) - \frac{\kappa_2}{2} \gamma_1 \Omega \sqrt{\beta_1^x} \phi(0) - \frac{\kappa_1 \kappa_2}{4}\tag{40}$$

for the sign of the expression in (40) to be nonnegative, by the same reasoning as above, it must be that the sign of the first term in (40) is also strictly positive. It must then be that

$$\left(\gamma_1 \gamma_2 \sqrt{1 + \Omega^2} \sqrt{\beta_1^x \beta_2^x} \phi(0) - \frac{\kappa_1}{2} \gamma_2 \Omega \sqrt{\beta_2^x} \right) \phi(0) - \frac{\kappa_2}{2} \gamma_1 \Omega \sqrt{\beta_1^x} \phi(0) - \frac{\kappa_1 \kappa_2}{4} \geq 0\tag{41}$$

which is impossible when Condition Q holds.

Next assume that $\gamma_1, \gamma_2 \geq 0 > \Omega$. Then, by the same arguments as above, the existence of a pair (\hat{x}_1, \hat{x}_2) for which the sign of the expression in (38) is nonnegative contradicts the assumption that Condition Q holds.

Next, assume that $\gamma_1, \Omega \geq 0 > \gamma_2$. It follows that

$$\frac{\partial G_2(x_1, x_2)}{\partial x_1} \frac{\partial G_1(x_1, x_2)}{\partial x_2} - \frac{\partial G_2(x_1, x_2)}{\partial x_2} \frac{\partial G_1(x_1, x_2)}{\partial x_1} \leq -\frac{\kappa_1}{2} \gamma_2 \Omega \sqrt{\beta_2^x} \phi(X_{12}(x_1, x_2)) - \frac{\kappa_1 \kappa_2}{4}\tag{42}$$

For the expression in the right-hand-side of (42) to be nonnegative, it must then be that

$$-\gamma_2 \Omega \sqrt{\beta_2^x} \phi(0) - \frac{\kappa_2}{2} \geq 0$$

which is impossible under condition (M).

Next consider the case where $\gamma_1 \geq 0 > \Omega, \gamma_2$. We then have that

$$\frac{\partial G_2(x_1, x_2)}{\partial x_1} \frac{\partial G_1(x_1, x_2)}{\partial x_2} - \frac{\partial G_2(x_1, x_2)}{\partial x_2} \frac{\partial G_1(x_1, x_2)}{\partial x_1} \leq -\frac{\kappa_2}{2} \gamma_1 \Omega \sqrt{\beta_1^x} \phi(0) - \frac{\kappa_1 \kappa_2}{4} < 0$$

where the last inequality is again by Condition (M).

We conclude that the function $F(\cdot)$ is strictly decreasing which implies that the threshold continuation equilibrium of Lemma 1 is unique. Standard global-game arguments then imply that there do not exist continuation equilibria other than the threshold one, which establishes the result.

■

Proof of Lemma 3. Existence of a maximizer. Because of the bijective relation between (p_1^A, p_2^A) and (\hat{x}_1, \hat{x}_2) it suffices to show that there exists a vector of thresholds (\hat{x}_1, \hat{x}_2) that maximize (16). To see this, note that, for any pair (\hat{x}_1, \hat{x}_2) ,

$$\hat{\Pi}^A(\hat{x}_1, \hat{x}_2) \equiv \sum_{i=1,2} \left[s_i - \frac{\kappa_i}{2} \hat{x}_i + \gamma_i M_i^A(\hat{x}_1, \hat{x}_2) \right] \Phi\left(\sqrt{\beta_i^x} \hat{x}_i\right)$$

which means that

$$\sum_{i=1,2} \left[s_i - \frac{\kappa_i}{2} \hat{x}_i + \gamma_i^- \right] \Phi\left(\sqrt{\beta_i^x} \hat{x}_i\right) \leq \hat{\Pi}^A(\hat{x}_1, \hat{x}_2) \leq \sum_{i=1,2} \left[s_i - \frac{\kappa_i}{2} \hat{x}_i + \gamma_i^+ \right] \Phi\left(\sqrt{\beta_i^x} \hat{x}_i\right) \quad (43)$$

Next, consider the function

$$G_1^+(x_i) \equiv \left[s_i - \frac{\kappa_i}{2} x_i + \gamma_i^+ \right] \Phi\left(\sqrt{\beta_i^x} x_i\right)$$

and note that this function is bounded from above but not from below.³² By looking at the right-hand side of (43), it is then immediate that, for any $i = 1, 2$, there exists a finite \bar{x}_i such that $\hat{\Pi}^A(\hat{x}_1, \hat{x}_2) < 0$ for any (\hat{x}_1, \hat{x}_2) such that $\hat{x}_i \geq \bar{x}_i$. Because the platform can always guarantee itself zero profits by setting prices equal to zero, this means that, to find a maximizer of $\hat{\Pi}^A(\hat{x}_1, \hat{x}_2)$, one can restrict attention to pairs (\hat{x}_1, \hat{x}_2) such that $\hat{x}_i \leq \bar{x}_i$, $i = 1, 2$.

Next, note that $\lim_{x_i \rightarrow -\infty} G_1^+(x_i) = 0$. This means that for any $i = 1, 2$, $j \neq i$ and $\varepsilon > 0$ arbitrarily small, there exists a finite \underline{x}_i such that, for any (\hat{x}_1, \hat{x}_2) with $\hat{x}_i \leq \underline{x}_i$,

$$\hat{\Pi}^A(\hat{x}_1, \hat{x}_2) \leq \varepsilon + \left[s_j - \frac{\kappa_j}{2} \hat{x}_j + \gamma_j^+ \right] \Phi\left(\sqrt{\beta_j^x} \hat{x}_j\right) \quad (44)$$

Now take any $\hat{x}_i^\# \in \arg \max_x G_i^-(x)$ and note that any such $\hat{x}_i^\#$ is such that $\hat{x}_i^\# > \underline{x}_i$. This means, for any (\hat{x}_1, \hat{x}_2) with $\hat{x}_i \leq \underline{x}_i$, the inequality in (44) holds whereas the following inequality

$$\hat{\Pi}^A(\hat{x}_1, \hat{x}_2) > G_i^-(\hat{x}_i^\#) + \left[s_j - \frac{\kappa_j}{2} \hat{x}_j + \gamma_j^- \right] \Phi\left(\sqrt{\beta_j^x} \hat{x}_j\right) \quad (45)$$

³²This follows from the fact that the standard Normal distribution satisfies the property that $\lim_{x \rightarrow -\infty} x \Phi(x) = 0$.

holds for $(\hat{x}_i^\#, \hat{x}_j)$. By Condition (W), we then have that, for any $i = 1, 2$, any pair (\hat{x}_1, \hat{x}_2) with $\hat{x}_i \leq \underline{x}_i$, there exists a pair (\hat{x}'_1, \hat{x}'_2) with $\hat{x}'_i = \hat{x}_i^\#$ and $\hat{x}'_j = \hat{x}_j$ such that $\hat{\Pi}^A(\hat{x}'_1, \hat{x}'_2) > \hat{\Pi}^A(\hat{x}_1, \hat{x}_2)$. Together with the result above, this means that, when looking for maximizers of $\hat{\Pi}^A(\hat{x}_1, \hat{x}_2)$ one can restrict attention to pairs (\hat{x}_1, \hat{x}_2) such that $\underline{x}_i \leq \hat{x}_i \leq \bar{x}_i$, $i = 1, 2$. Because the above is a compact set, and because the function $\hat{\Pi}^A(\hat{x}_1, \hat{x}_2)$ is continuous and differentiable, this proves that a maximizer to $\hat{\Pi}^A(\hat{x}_1, \hat{x}_2)$ always exists.

Necessity of the first order conditions. By construction of the intervals $[\underline{x}_i, \bar{x}_i]$, any maximizer of $\hat{\Pi}^A(\hat{x}_1, \hat{x}_2)$ is necessarily interior to the rectangular $[\underline{x}_1, \bar{x}_1] \times [\underline{x}_2, \bar{x}_2]$ and thus must satisfy the first-order conditions (17). ■

Proof of Proposition (1). Instead of proving equivalence with (17), we rewrite condition 18 as ,

$$p_i^A \cdot \frac{dQ_i^A}{d\hat{x}_i} \frac{d\hat{x}_i}{dp_i^A} \Big|_{\hat{x}_j=const} + Q_i^A + \gamma_j \frac{\partial M_i^A(\hat{x}_1, \hat{x}_2)}{\partial \hat{x}_i} \frac{d\hat{x}_i}{dp_i^A} \Big|_{\hat{x}_j=const} Q_j^A = 0.$$

Then using $\frac{d\hat{x}_i}{dp_i^A} \Big|_{\hat{x}_j=const} = \frac{1}{-\kappa_i/2 + \gamma_i \frac{\partial M_j^A(\hat{x}_1, \hat{x}_2)}{\partial \hat{x}_i}}$ we have that

$$p_i^A = \frac{\kappa_i}{2} \frac{Q_i^A}{\frac{dQ_i^A}{d\hat{x}_i}} - \gamma_i \frac{\partial M_j^A(\hat{x}_1, \hat{x}_2)}{\partial \hat{x}_i} \frac{Q_i^A}{\frac{dQ_i^A}{d\hat{x}_i}} - \gamma_j \frac{\frac{\partial M_i^A(\hat{x}_1, \hat{x}_2)}{\partial \hat{x}_i}}{\frac{dQ_i^A}{d\hat{x}_i}} Q_j^A = 0$$

which gives the result. ■

Proof of Proposition 2. By definition, in a symmetric equilibrium, $p_i^A = p_i^B$, $i = 1, 2$. Under Conditions (M), (Q) and (W), the unique continuation equilibrium is then a threshold equilibrium with thresholds $\hat{x}_1 = \hat{x}_2 = 0$ and expected demands $Q_i^A = \mathbb{E}[\tilde{m}_i^A] = 1/2$, $i = 1, 2$. Substituting $\hat{x}_i = 0$ and $Q_i^A = 1/2$, $i = 1, 2$, into the the formulas for $dQ_i^A/d\hat{x}_i$, $dM_j^A/d\hat{x}_i$, and $dM_i^A/d\hat{x}_i$ (as given by (21), (19) and (22), respectively) and replacing these formulas into the optimality conditions (31), we then have that the equilibrium prices are given by

$$p_i^* = \frac{\kappa_i}{2\sqrt{\beta_i^x \phi(0)}} + \gamma_i \Omega - \gamma_j \sqrt{1 + \Omega^2}$$

Noticing that $\frac{\kappa_i}{2\sqrt{\beta_i^x \phi(0)}} = \mu_i^d(0)$ then gives the result. ■

Proof of Proposition 3. First note that, when $\underline{s}_i > p_i^* - \gamma_i^-$, in the proposed equilibrium where participation to one of the two platforms is compulsory, each agent obtains more than his outside option (normalized to zero). Now suppose that platform B offers the equilibrium prices and consider the problem faced by platform A (the problem faced by platform B is symmetric). Clearly, for any deviation entailing a reduction in the price offered to each side, one can construct a continuation equilibrium where each agent behaves exactly as in the game where participation is compulsory, in which case the deviation is unprofitable. Next, for any $i = 1, 2$, let $x_i^\#$ be implicitly defined by

$$s_i + \frac{1}{2} \kappa_i x_i^\# + \gamma_i^- = p_i^*$$

and observe that, no agent from side i receiving a signal $x_i > x_i^\#$ will ever opt out, irrespective of the prices charged by platform A , for, irrespective of the other agents' decisions, he can obtain a positive surplus by joining platform B .

Now observe that the equilibrium prices p_i^* , $i = 1, 2$, are independent of s_i and that $x_i^\#$ is strictly decreasing in s_i , going to $-\infty$ as s_i goes to $+\infty$. Suppose now that there exists a vector of prices (p_1^A, p_2^A) such that, in any of the continuation equilibria that follow the selection of the prices $(p_1^A, p_2^A, p_1^*, p_2^*)$, platform A is strictly better off than under the monotone equilibrium that follows the selection of the equilibrium prices $(p_1^*, p_2^*, p_1^*, p_2^*)$. Clearly, for this to be possible, there must exist $i \in \{1, 2\}$ such that $\hat{x}_i(p_1^A, p_2^A, p_1^*, p_2^*) \leq x_i^\#$, where $\hat{x}_i(p_1^A, p_2^A, p_1^*, p_2^*)_{i=1,2}$ are the thresholds defined by (28) in the game where participation is compulsory. Finally, let $x_i^+(p_1^A, p_2^A, p_1^*, p_2^*)$ be implicitly defined by

$$s_i - \frac{1}{2}\kappa_i x_i^+ + \gamma_i^+ = p_i^A$$

and observe that no agent from side i with signal $x_i > x_i^+(p_1^A, p_2^A, p_1^*, p_2^*)$ will ever join platform A , irrespective of his beliefs about the other agents' participation decisions. Now, letting side i be the one for which $\hat{x}_i(p_1^A, p_2^A, p_1^*, p_2^*) \leq x_i^\#$, observe that, necessarily,

$$x_i^+(p_1^A, p_2^A, p_1^*, p_2^*) < \hat{x}_i(p_1^A, p_2^A, p_1^*, p_2^*) + 2|\gamma_i|/\kappa_i. \quad (46)$$

To see this, let $q(\cdot)$ and $r(\cdot)$ be the function defined by

$$\begin{aligned} q(x_i) &\equiv s_i - \frac{1}{2}\kappa_i x_i + \gamma_i^+ - p_i^A \text{ and} \\ r(x_i) &\equiv s_i - \frac{1}{2}\kappa_i x_i + \gamma_i \Phi \left(\sqrt{\frac{\beta_j^x}{1 - \rho_x^2}} \left(\hat{x}_j - \rho_x \sqrt{\frac{\beta_i^x}{\beta_j^x}} x_{il} \right) \right) - p_i^A \end{aligned}$$

where, again, $\hat{x}_i(p_1^A, p_2^A, p_1^*, p_2^*)_{i=1,2}$ are the thresholds defined by (28) in the game where participation is compulsory. Note that, for any x_i , $0 \leq q(x_i) - r(x_i) \leq |\gamma_i|$. Because $r(\hat{x}_i) < 0$, it follows that $q(x_i) \leq |\gamma_i|$. Given the linearity of $q(\cdot)$ in x_i , we then have that the unique solution x_i^+ to $q(x_i^+) = 0$ must necessarily satisfy (46).

Having established that $x_i^\#, x_i^+, \hat{x}_i$ all converge (uniformly) to $-\infty$ as $s_i \rightarrow +\infty$, we then have that, in the limit as $s_i \rightarrow +\infty$, $m_i^A(p_1^A, p_2^A, p_1^*, p_2^*) \rightarrow 0$ and $m_i^B(p_1^A, p_2^A, p_1^*, p_2^*) \rightarrow 1$, exactly as in the game where participation is compulsory. This means that, when s_i goes to infinity, $i = 1, 2$, platform A 's payoff given the prices $(p_1^A, p_2^A, p_1^*, p_2^*)$ under any continuation equilibrium in the game where participation is voluntarily must converge to its' payoff in the unique continuation equilibrium of the game where participation is compulsory. Because the latter is necessarily less than the platform's payoff under the equilibrium prices, and because, by quasi-concavity of payoffs, there exists $K, M > 0$ such that, in the game where participation is compulsory

$$\Pi^A(p_1^*, p_2^*, p_1^*, p_2^*) - \Pi^A(p_1^A, p_2^A, p_1^*, p_2^*) > K$$

for any $(p_1^A, p_2^A, p_1^*, p_2^*)$ for which there exists $i \in \{1, 2\}$ such that $p_i^A > M$, we conclude that, no matter the selected continuation equilibrium, any deviation resulting in partial participation is necessarily unprofitable. This completes the proof. ■

Proof of Proposition 4. Recall that each agent l from each side i prefers joining platform A to joining platform B if and only if

$$\mathbb{E} [z_i(\tilde{\theta}_i + \tilde{\varepsilon}_{il}) \mid x_{il}] + \gamma_i \mathbb{E} [\tilde{m}_j^B - \tilde{m}_j^A \mid x_{il}] \leq p_i^B - p_i^A. \quad (47)$$

The same agent then prefers joining platform A to multihoming if and only if

$$(1 - \kappa_i)s_i + \frac{1}{2}\mathbb{E} [z_i(\tilde{\theta}_i + \tilde{\varepsilon}_{il}) \mid x_{il}] + \gamma_i \mathbb{E} [\tilde{\mu}_j^B \mid x_{il}] - p_i^B \leq 0. \quad (48)$$

Note that Condition (48) is implied by Condition (47) if and only if

$$2(1 - \kappa_i)s_i + 2\gamma_i \mathbb{E} [\tilde{\mu}_j^B \mid x_{il}] - \gamma_i \mathbb{E} [\tilde{m}_j^B - \tilde{m}_j^A \mid x_{il}] \leq p_i^A + p_i^B \quad (49)$$

In any continuation equilibrium where all agents singlehome $m_j^B = \mu_j^B = 1 - m_j^A$, in which case the inequality in (49) becomes equivalent to $\gamma_i + 2(1 - \kappa_i)s_i \leq p_i^A + p_i^B$. The same conclusion applies to those agents that prefer platform B to platform A . From the results above, we know that the game where multihoming is not possible always admits a continuation equilibrium. We then conclude that, when $p_i^A + p_i^B \geq \gamma_i + 2(1 - \kappa_i)s_i$ such a continuation equilibrium is also a continuation equilibrium in the game where agents can multihome.

Conversely, when $p_i^A + p_i^B < \gamma_i + 2(1 - \kappa_i)s_i$, there exists no continuation equilibrium where all agents singlehome, for, if such equilibrium existed, then it would satisfy $m_j^B = \mu_j^B = 1 - m_j^A$. Inverting the inequalities above, we would then have that some agent from side i would necessarily prefer to multihome. ■

Proof of Proposition 5. Recall that agent l 's (from side i) ability to forecast his own stand-alone valuations is measured by the inverse of

$$\text{var}[\tilde{v}_{il} - \tilde{V}_{il}] = z_i^2 \frac{\alpha_i + \beta_i^\varepsilon}{\alpha_i \beta_i^\varepsilon} - z_i^2 \frac{(\beta_i^\eta + \rho_i \alpha_i \sqrt{\beta_i^\eta / \beta_i^\varepsilon})^2}{(\alpha_i + \beta_i^\eta) \alpha_i \beta_i^\eta} \quad (50)$$

Likewise, the agent's ability to forecast the distribution of true stand-alone valuations on the other side of the market is measured by the inverse of

$$\text{var}[\tilde{\theta}_j - \mathbb{E}[\tilde{\theta}_j \mid \tilde{x}_{il}]] = \left(1 - \rho_\theta^2 \frac{\beta_i^\eta}{\alpha_i + \beta_i^\eta}\right) \frac{1}{\alpha_j}.$$

Finally, the agent's ability to forecast the estimated valuations of any agent l' from side j is given by the inverse of

$$\text{var}[\tilde{x}_{jl'} - \mathbb{E}[\tilde{x}_{jl'} \mid \tilde{x}_{il}]] = \text{var}[\tilde{\theta}_j - \mathbb{E}[\tilde{\theta}_j \mid \tilde{x}_{il}]] + \frac{1}{\beta_j^\eta}.$$

Finally, recall that the ex-ante distribution of estimated stand-alone valuations on each side $i = 1, 2$ of the market is Normal with zero mean and variance

$$var[\tilde{V}_i] = z_i^2 \frac{(\beta_i^\eta + \rho_i \alpha_i \sqrt{\beta_i^\eta / \beta_i^\varepsilon})^2}{(\alpha_i + \beta_i^\eta) \alpha_i \beta_i^\eta} \quad (51)$$

Now observe that the equilibrium profits are given by $\Pi^A = \Pi^B = \Pi^* \equiv \frac{1}{2}(p_1^* + p_2^*)$ with

$$p_i^* = \frac{\sqrt{var[\tilde{V}_i]}}{2\phi(0)} + \gamma_i \Omega - \gamma_j \sqrt{1 + \Omega^2}$$

and

$$\Omega \equiv \rho_\theta \sqrt{\frac{\beta_1^\eta \beta_2^\eta}{\alpha_1 \alpha_2 + \beta_1^\eta \alpha_2 + \beta_2^\eta \alpha_1 + (1 - \rho_\theta^2) \beta_1^\eta \beta_2^\eta}}.$$

Because the prior distribution is fixed, so are the parameters $(\alpha_1, \alpha_2, \rho_\theta, \beta_1^\varepsilon, \beta_2^\varepsilon, z_1, z_2)$. It is then immediate from (50) and (51) that campaigns that increase the agents' ability to forecast their own stand-alone valuations increase the ex-ante dispersion of estimated stand-alone valuations. From the formula for the equilibrium prices, it is then easy to see that, when such campaigns do not affect the agents' ability to forecast the distribution of true (and estimated) stand-alone valuations on the other side of the market (that is, when they leave β_1^η and β_2^η unchanged), they necessarily increase equilibrium prices and hence equilibrium profits.

Next consider campaigns that leave unchanged the agents' ability to forecast their own stand-alone valuations (and hence the ex-ante dispersion of estimated stand-alone valuations). Then such campaigns increase profits if and only if they increase $(\gamma_1 + \gamma_2) \left(\Omega - \sqrt{1 + \Omega^2} \right)$ which is the case if and only if

$$\frac{\partial \Omega}{\partial \beta_i^\eta} (\gamma_1 + \gamma_2) \geq 0 \quad i = 1, 2.$$

Using the fact that Ω is increasing in β_1^η and β_2^η if and only if $\rho_\theta \geq 0$, we then have that such campaigns increase profits if and only if $\rho_\theta (\gamma_1 + \gamma_2) \geq 0$, thus establishing the result. ■

Proof of Proposition 6. The results concerning the comparative statics with respect to $(\alpha_1, \alpha_2, \beta_1^\varepsilon, \beta_2^\varepsilon)$ follow directly from inspecting the formula for the equilibrium prices and observing that the ex-ante dispersion of estimated stand alone-valuations $var[\tilde{V}_{il}]$ on each $i = 1, 2$ decreases with $(\alpha_i, \beta_i^\varepsilon)$ and is independent of $(\alpha_j, \beta_j^\varepsilon)$, whereas the coefficient of mutual forecastability Ω is independent of $(\alpha_1, \alpha_2, \beta_1^\varepsilon, \beta_2^\varepsilon)$.

Next, consider the comparative statics with respect to the coefficient of correlation ρ_θ . The result then follows from observing that

$$\frac{\partial \Pi^*}{\partial \rho_\theta} = \frac{1}{2} (\gamma_1 + \gamma_2) \frac{\partial \Omega}{\partial \rho_\theta} \left\{ 1 - \frac{\Omega}{\sqrt{1 + \Omega^2}} \right\}$$

which is positive if and only if $\gamma_1 + \gamma_2 \geq 0$. ■