Two-sided Markets and Matching Design*

Renato Gomes† Alessandro Pavan‡

August 21, 2021

1 Introduction

Two-sided markets are markets in which agents match through a platform, which designs and prices matching opportunities. Typical examples include ad-exchanges matching advertisers with publishers; media outlets matching readers/viewers with content providers and advertisers; video-game consoles matching gamers with game developers; operating systems matching end-users with software developers; e-commerce websites matching buyers with sellers; business-to-business platforms matching procurers with service providers; and employment agencies matching employers with job seekers.

In the last few years, platform markets have gained a prominent role in the organization of business activities. As a result, a conspicuous literature has flourished examining various aspects of such markets, ranging from pricing to platform design. In this chapter, we focus on monopolistic pricing and its connection to matching design. Section 2 contains a flexible model of platform-mediated matching with transfers. Section 3 reviews some of the classical results on monopolistic pricing in two-sided markets. Section 4 extends some of these results to markets in which both the platform and the agents face uncertainty over the distribution of preferences over the two sides of the market and hence over the eventual participation decisions. Section 5 considers markets in which the platform engages in discriminatory practices matching different agents to different subsets of the participating agents from the other side of the market. It first considers the case of one-to-one matching and then the case of many-to-many matching. Throughout the entire chapter, special attention is given to the distortions in the provision of matching services that emerge in two-sided markets when the platform enjoys significant market power.

*This chapter will appear in "Online and Matching-Based Market Design," Federico Echenique, Nicole Immorlica and Vijay V. Vazirani, Editors, Cambridge University Press, 2021. For comments and suggestions, we thank the Editors as well as Atila Abdulkadiroglu, Guillaume Haeringer, Hanna Halaburda, Bruno Jullien, Philipp Kircher, Ravi Jagadeesan, and Paulo Somaini.

†Toulouse School of Economics, CNRS, University of Toulouse Capitole: renato.gomes@tse-fr.eu.

‡Northwestern University and CEPR: alepavan@northwestern.edu.

2 General Setup

Consider a "large" two-sided market in which the impact of each individual agent in isolation on the platform's profits is small. We capture such a situation by assuming that each side is populated by a unit-mass continuum of agents. This is a point of departure with respect to what assumed in the other chapters in this book (see, however, the chapter on large matching markets in this book for other benefits of assuming a large market). The assumption of a continuum of agents permits us to illustrate in the simplest possible way the distortions that arise when agents are privately informed and the platform has market power. It also permits us to bridge the analysis of matching design in this chapter with the literature on two-sided markets in Industrial Organization where demands are smooth.

To capture the platform's market power in the starkest possible terms, we consider a situation where a single platform matches agents from the two sides of the market.

Each agent from each side $k = a, b$ has a type ω_k drawn from a distribution F_k with support Ω_k, independently across agents (from either side of the market). Agents privately know their types.

A matching mechanism consists of a *matching rule* and a *payment rule* for each side of the market. By the Revelation Principle, it is without loss of generality to focus on direct-revelation mechanisms. Accordingly, we define a matching mechanism as $M := \{\mathbf{s}_k(\cdot), \mathbf{p}_k(\cdot)\}_{k=a,b}$, where, for each $k = a, b$, and each $\omega_k \in \Omega_k$, $\mathbf{s}_k(\omega_k) \subset \Omega_{-k}$ is the set of types from side $-k$ (that is, from the opposite side of the market) that type ω_k is matched to, whereas $\mathbf{p}_k(\omega_k)$ is the payment asked/given to the agent. Formally, $\mathbf{p}_k : \Omega_k \to \mathbb{R}$ (both positive and negative payments are allowed), while $\mathbf{s}_k : \Omega_k \to \mathcal{C}(\Omega_{-k})$, where $\mathcal{C}(\Omega_{-k})$ is a subset of the power set of Ω_{-k} describing the collection of *admissible matching sets*. This set captures technological or institutional constraints on the shape matching sets. For instance, in the case of one-to-one matching, $\mathcal{C}(\Omega_{-k})$ is the collection of singleton sets $\{\omega_{-k}\}$.

A matching rule $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ is *feasible* if and only if the following reciprocity condition holds for all $\omega_k \in \Omega_k$, $k = a, b$:

$$\omega_{-k} \in \mathbf{s}_k(\omega_k) \Rightarrow \omega_k \in \mathbf{s}_{-k}(\omega_{-k}). \tag{1}$$

We assume that agents' preferences are quasi-linear. Namely, the utility that a side-k agent of type ω_k derives from the matching set $\tilde{s}_k \in \mathcal{C}(\Omega_{-k})$ when making a payment $\tilde{p}_k \in \mathbb{R}$ to the platform is equal to $u_k(\tilde{s}_k|\omega_k) - \tilde{p}_k$, with the real-valued function $u_k(s_k|\omega_k)$ describing the agent's gross payoff. Therefore, the payoff that type ω_k obtains when reporting type ω'_k under the mechanism M is given by

$$\hat{U}_k(\omega_k, \omega'_k; M) := u_k(\mathbf{s}_k(\omega'_k)|\omega_k) - \mathbf{p}_k(\omega'_k),$$

whereas the payoff from truthful reporting his type is equal to $U_k(\omega_k; M) := \hat{U}_k(\omega_k, \omega_k; M)$. A mechanism M is *individually rational* (IR) if $U_k(\omega_k; M) \geq 0$

for all $\omega_k \in \Omega_k$, $k = a, b$, and is *incentive compatible* (IC) if $U_k(\omega_k; M) \geq \hat{U}_k(\omega_k, \omega_k'; M)$ for all $\omega_k, \omega_k' \in \Omega_k$, $k = a, b$. The definition of incentive compatibility is the same as in the mechanism design literature; it is de facto equivalent to the notion of *strategy proofness* in the introductory chapter and used in most of the matching literature.

A feasible matching rule is *implementable* if there exists a payment rule $\{\mathbf{p}_k(\cdot)\}_{k=a,b}$ such that the mechanism $M = \{\mathbf{s}_k(\cdot), \mathbf{p}_k(\cdot)\}_{k=a,b}$ is individually rational and incentive compatible.

Example 1. Let $\Omega_k = [\underline{\omega}_k, \bar{\omega}_k] \subset \mathbb{R}_+$ and $u_k(s_k|\omega_k) = \omega_k|s_k|$, $k = a, b$, with $|s_k| := \int_{s_k} dF_{-k}(\omega_{-k})$ denoting the measure of the set s_k. A feasible matching rule $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ is implementable if and only if $|\mathbf{s}_k(\cdot)|$ is non-decreasing, $k = a, b$. To see this, note that $u_k(s_k|\omega_k)$ is supermodular. Hence, if type ω_k weakly prefers the pair $(\mathbf{s}_k(\omega_k), \mathbf{p}_k(\omega_k))$ to the pair $(\mathbf{s}_k(\tilde{\omega}_k), \mathbf{p}_k(\tilde{\omega}_k))$, and $|\mathbf{s}_k(\omega_k)| > |\mathbf{s}_k(\tilde{\omega}_k)|$, then any type $\omega_k' > \omega_k$ strictly prefers the pair $(\mathbf{s}_k(\omega_k), \mathbf{p}_k(\omega_k))$ to the pair $(\mathbf{s}_k(\tilde{\omega}_k), \mathbf{p}_k(\tilde{\omega}_k))$, which implies that the mechanism $M = \{\mathbf{s}_k(\cdot), \mathbf{p}_k(\cdot)\}_{k=a,b}$ is incentive compatible only if $|\mathbf{s}_k(\cdot)|$ is non-decreasing, $k = a, b$. Lastly, to see that any feasible rule $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ such that $|\mathbf{s}_k(\cdot)|$ is non-decreasing, $k = a, b$, is implementable, consider the mechanism M in which the payment rule is given by $\mathbf{p}_k(\omega_k) = \omega_k \mathbf{s}_k(\omega_k) - \int_{\underline{\omega}_k}^{\omega_k} |\mathbf{s}_k(\tilde{\omega}_k)| d\tilde{\omega}_k$, all $\omega_k \in \Omega_k$, $k = a, b$. It is then easy to see that, because $|\mathbf{s}_k(\cdot)|$ is non-decreasing,

$$U_k(\omega_k; M) = \int_{\underline{\omega}_k}^{\omega_k} |\mathbf{s}_k(\tilde{\omega}_k)| d\tilde{\omega}_k$$
$$\geq \int_{\underline{\omega}_k}^{\omega_k'} |\mathbf{s}_k(\tilde{\omega}_k)| d\tilde{\omega}_k + (\omega_k - \omega_k')|\mathbf{s}_k(\omega_k')| = \hat{U}_k(\omega_k, \omega_k'; M),$$

which implies that the mechanism M is incentive compatible. That the same mechanism is also individually rational follows from the fact that $U_k(\omega_k; M) = \int_{\underline{\omega}_k}^{\omega_k} |\mathbf{s}_k(\tilde{\omega}_k)| d\tilde{\omega}_k \geq 0$ for all ω_k, $k = a, b$. Hence the monotone rule $\mathbf{s}_k(\cdot)\}_{k=a,b}$ is implementable. ♢

In what follows, we specialize the above formulation to capture specific aspects of pricing and matching design in two-sided markets.

3 Pricing in Two-sided Markets

The study of pricing in two-sided markets originally focused on environments characterized by the absence of discrimination (across agents from the same side), and the presence of cross-side network effects.

The first property implies that all side-k agents that join the platform are assigned the same matching set. This restriction is motivated by the inability of many two-sided platforms to customize matching opportunities.[1] Accordingly, for $k = a, b$, there are sets $\{\hat{\Omega}_k\}_{k=a,b}$, with $\hat{\Omega}_k \in \mathcal{C}(\Omega_{-k}) \neq \emptyset$, such that, for any $\omega_k \in \Omega_k$, either $\mathbf{s}_k(\omega_k) = \emptyset$, meaning that type ω_k is excluded, or $\mathbf{s}_k(\omega_k) = \hat{\Omega}_{-k}$,

[1]For instance, in a shopping mall or fair, it is impossible or impractical to prevent all participating buyers and sellers to freely interact.

where $\hat{\Omega}_{-k}$ is the set of participating types from side $-k$. Feasibility obviously requires that $\mathbf{s}_k(\omega_k) = \hat{\Omega}_{-k}$ if and only if $\omega_k \in \hat{\Omega}_k$ for $k = a, b$. We say that such a matching rule induces *the single network* of participating agents $(\hat{\Omega}_a, \hat{\Omega}_b)$.

The presence of cross-side network effects is captured by the following assumptions on the agents' types and preferences. The type of each agent is a two-dimensional vector $\omega_k = (\omega_k^s, \omega_k^i) \in \mathbb{R}^2$, where ω_k^s denotes the agent's "*stand-alone value*," that is, the benefit the agent derives from all products and services the platform provides in addition to matching agents from the two sides of the market, whereas ω_k^i is the agent's "*interaction benefit*," that is, the value the agent derives from interacting with agents from the other side of the market. The gross utility of a type-ω_k agent from being matched to a set $s_k \subseteq \Omega_{-k}$ of types from the other side of the market takes the form

$$u_k(s_k|\omega_k) := \omega_k^s + \omega_k^i |s_k|, \tag{2}$$

where $|s_k| := \int_{s_k} dF_{-k}(\omega_{-k})$ is the measure of the set s_k. Accordingly, the presence of more agents from the opposite side enhances the utility of a side-k agent of type ω_k if and only if $\omega_k^i > 0$. In advertising markets, for instance, it is typically assumed that advertisers (on side a) have positive interaction benefits, $\omega_a^i > 0$, whereas consumers (on side b) have negative interaction benefits $\omega_b^i < 0$ (that is, dislike advertising). In this example, a consumer's (positive) stand-alone value is the utility she derives from the content provided by the platform (e.g., news or services), whereas an advertiser's (negative) stand-alone value is cost of producing the advertisement.

The next lemma relates the matching rules inducing a single network to the transfer rules that implement them. The proof is straightforward and hence omitted. To simplify, assume that, whenever indifferent, agents join the platform.

Lemma 2. *A matching rule $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ inducing a single network $(\hat{\Omega}_a, \hat{\Omega}_b)$ is implementable if and only if there exist access prices P_a and P_b such that, for all $k = a, b$,*

$$\hat{\Omega}_k = \left\{ \omega_k \in \Omega_k : \omega_k^s + \omega_k^i |\hat{\Omega}_{-k}| \geq P_k \right\}. \tag{3}$$

In light of Lemma 2, consider the game in which the side-k agents are offered to join the platform at the price P_k, and in which the agents' participation decisions are simultaneous. There exists an equilibrium of this game in which each side-k agent joins the platform if and only if $\omega_k \in \hat{\Omega}_k$. However, the equilibrium need not be unique. For example, when the agents' stand-alone values are identically equal to zero, that is, $\omega_k^s = 0$, and the interaction benefits are homogenous within and across sides, that is, $\omega_k^i = 1$ for $k = a, b$, any price vector (P_a, P_b) with $0 \leq P_a, P_b \leq 1$, implements the single complete network $\hat{\Omega}_k = \Omega_k$, $k = a, b$. The implementation is however partial, in that the game also admits a continuation equilibrium in which none of the agents participates.

In the rest of this section, consistently with the mechanism design literature, we shall disregard the multiplicity issue and describe the platform's problem as choosing a pair of participation values (N_a, N_b), with $N_k = |\hat{\Omega}_k|$ denoting the mass of agents from side k joining the platform. By virtue of Lemma 2, such a pair of participation values can be supported in equilibrium if and only if there exist sets $(\hat{\Omega}_a, \hat{\Omega}_b)$ and prices (P_a, P_b) satisfying (3) and $N_k = |\hat{\Omega}_k|$, $k = a, b$.

3.1 Profit-Maximizing Prices

We consider three scenarios, corresponding to different specifications of the agents' preferences. To this end, denote by F_k^s and F_k^i the marginal distributions of the joint cdf F_k, and by f_k (alternatively, f_k^s, f_k^i) the density of F_k (alternatively, of F_k^s, F_k^i) if it exists. We let $\mathbf{1}\{A\}$ be the indicator function taking value 1 if statement A is true, and zero otherwise.

Scenario 1. *Agents are heterogeneous in their stand-alone values, but homogenous in their interaction benefits, with the latter equal to* $\mathring{\omega}_k^i > 0$. *In this case,* $F_k(\omega_k) = \mathbf{1}\left\{\omega_k^i \geq \mathring{\omega}_k^i\right\} F_k^s(\omega_k^s)$*, with the marginal* F_k^s *absolutely continuous over* $\mathbb{R}$.

Scenario 2. *Agents are heterogeneous in their interaction benefits, but homogenous in their stand-alone values, which we normalize to zero with no loss of generality:* $\mathring{\omega}_k^s = 0$. *In this case,* $F_k(\omega_k) = \mathbf{1}\{\omega_k^s \geq 0\} F_k^i(\omega_k^i)$*, with the marginal* F_k^i *absolutely continuous over* $\mathbb{R}$.

Scenario 3. *Agents are heterogeneous in both their stand-alone values and their interaction benefits. In this case,* F_k *is absolutely continuous over* $\mathbb{R}^2$.

In light of Lemma 2, we can formulate the platform's problem in terms of the measure of agents that join from each side of the market. To do so, fix $(N_a, N_b) \in (0, 1]^2$ and, for each $k \in \{a, b\}$, define $P_k(N_a, N_b)$ as the unique solution to the system of equations given by

$$N_k = \int_{\left\{\omega_k \in \Omega_k : \omega_k^s + \omega_k^i N_{-k} \geq P_k\right\}} dF_k(\omega_k), \qquad k = a, b. \tag{4}$$

Intuitively, the two equations in (4) identify the access prices that implement a single network in which the measure of participating agents from each side $k = a, b$ is equal to $N_k > 0$. In either of the three scenarios considered above, the inverse demand functions $P_k(N_a, N_b)$, $k = a, b$, are differentiable. We then define the *own-price demand elasticity on side* k as

$$\varepsilon_k(N_a, N_b) := -\frac{P_k(N_a, N_b)}{N_k} \left(\frac{\partial P_k}{\partial N_k}(N_a, N_b)\right)^{-1}.$$

Fixing the measure N_{-k} of participating agents from side $-k$, $1/\varepsilon_k$ captures the sensitivity of the side-k inverse demand with respect to variations in the side-k participation N_k. Equivalently, given the prices $(P_a, P_b) = \{P_k(N_a, N_b)\}_{k=a,b}$ implementing the participation vector (N_a, N_b), $\varepsilon_k(N_a, N_b)$ is the elasticity of the side-k direct demand with respect to variations in the side-k price, for fixed participation on side $-k$.

Next, assume that the platform incurs a participation cost c^s_k for each side-k agent it brings on board, and an interaction cost c^i for every interaction between the two sides it induces. For any (N_a, N_b), its profit is then equal to

$$\Pi(N_a, N_b) := \sum_{k=a,b} N_k \left(P_k(N_a, N_b) - c^s_k\right) - c^i N_a N_b.$$

Proposition 3. *Consider a profit-maximizing platform designing a single network, and let the agents' preferences be given by (2). The profit-maximizing prices (P^*_a, P^*_b), along with the participation profile (N^*_a, N^*_b) they induce, solve*

$$\frac{P^*_k - \left[c^s_k + N^*_{-k}\left(c^i - \tilde{\omega}^i_{-k}(N^*_a, N^*_b)\right)\right]}{P^*_k} = \frac{1}{\varepsilon_k(N^*_a, N^*_b)} \tag{5}$$

for $k = a, b$, where $P^*_k = P_k(N^*_a, N^*_b)$, with $P_k(N_a, N_b)$ given by (4), and where

$$\tilde{\omega}^i_{-k}(N^*_a, N^*_b) := \mathbb{E}\left[\omega^i_{-k} | \omega^s_{-k} + \omega^i_{-k} N^*_k = P^*_{-k}\right]$$

is the average interaction benefit of those agents from side $-k$ who are just indifferent between participating and not participating, under the profile (N^*_a, N^*_b).

Proof. The result is obtained by differentiating the objective function $\Pi(N_a, N_b)$ with respect to N_k, $k = a, b$, and then noting that $\frac{\partial P_k(N_a,N_b)}{\partial N_k} \frac{N_k}{P_k(N_a,N_b)} = -\frac{1}{\varepsilon_k(N_a,N_b)}$ and that $\frac{\partial P_{-k}(N_a,N_b)}{\partial N_k} = \tilde{\omega}^i_{-k}(N_a, N_b)$, where the last property follows from the implicit function theorem applied to (4). □

Condition (5) is instrumental to understanding how pricing in two-sided markets differs from pricing in more traditional one-sided environments. Relative to the classic Lerner formula for monopoly pricing, two differences stand out. The first is that, when choosing its side-k price, the platform must account for the fact that its effective side-k marginal cost is endogenous and depends on the set of participating agents on side $-k$. Indeed, when it increases the participation on side k, to hold constant the participation on side $-k$, the platform must adjust its price on side $-k$ by $\tilde{\omega}^i_{-k}(N_a, N_b)$, which is the average interaction benefit among all agents on side $-k$ who are just indifferent between joining the platform and staying out. When this term is positive, this effect contributes to a reduction in the side-k marginal cost, whereas the opposite is true when $\tilde{\omega}^i_{-k}(N_a, N_b) < 0$.

The second difference is that the demand elasticities are a function of the entire profile (N_a, N_b) of participating agents. Because of cross-side network effects, the participation on side $-k$ affects the willingness to pay of the participating agents on side k, and hence the elasticity of the side-k demand with respect to the side-k price.

The "two-sided" Lerner formula in (5) highlights that price skewness is a fundamental feature of markets with cross-side network effects. If, for instance, the demand elasticity is low on side a, but high on side b, the platform tends to set high prices on the former side, and low (potentially negative) prices on the latter side. In light of such considerations, prices on each side should not be taken *in vacuo* for competition policy purposes: Neither a low price on side b is a sign of predation (below-cost pricing) or a high-price on side a is a sign of abuse of market power (high markup). Rather, the welfare effects of monopolistic pricing should be evaluated by comparing prices on each side to their efficient counterparts, as we show next.

3.2 Welfare-maximizing Pricing

Let social welfare include all agents' utilities and the platform's profit. Accordingly, the welfare induced by a single network with participation profile (N_a, N_b) is given by

$$W(N_a, N_b) := \sum_{k\in\{a,b\}} \int_{\left\{\omega_k : \omega_k^s + \omega_k^i N_{-k} \geq P_k(N_a,N_b)\right\}} \left(\omega_k^s + \omega_k^i N_{-k} - c_k^s\right) dF_k(\omega_k) - c^i N_a N_b.$$

The term inside the integrals is the total gross surplus of the matches enabled by the platform (the sum of the agents' utilities net of all participation costs), while the last term is the platform's total interaction cost. The next proposition derives the welfare-maximizing participation profile, which we hereafter refer to as the *efficient participation profile.*

Proposition 4. *Consider a welfare-maximizing platform designing a single network, and let the agents' preferences be given by (2). The welfare-maximizing prices* (P_a^e, P_b^e) *along with the efficient participation profile* (N_a^e, N_b^e) *they induce solve*

$$P_k^e = c_k^s + N_{-k}^e \left(c^i - \bar{\omega}_{-k}^i(N_a^e, N_b^e)\right) \tag{6}$$

for $k = a, b$*, where* $P_k^e = P_k\left(N_a^e, N_b^e\right)$*, with* $P_k(N_a, N_b)$ *given by (4), and where*

$$\bar{\omega}_{-k}^i(N_a^e, N_b^e) := \mathbb{E}\left[\omega_{-k}^i | \omega_{-k}^s + \omega_{-k}^i N_k^e \geq P_{-k}^e\right]$$

is the average interaction benefit of the participating agents from side $-k$*.*

Proof. The result is obtained by differentiating the objective function $W(N_a, N_b)$ with respect to N_k, $k = a, b$, and then applying the implicit function theorem to (4) to obtain the formula for $\frac{\partial P_k(N_a,N_b)}{\partial N_k}$. □

Condition (6) is the two-sided incarnation of the Pigouvian precept according to which, to achieve efficiency, agents should be charged (or remunerated) for the externalities they impose on other market participants (as in the VCG mechanism). The price on each side is equal to the total marginal cost that the platform incurs to bring a marginal agent on board, adjusted by the network externality that the marginal agent exerts on the participating agents from the other side of the market. This externality equals the measure of agents on side $-k$ multiplied by the average interaction benefit among *all the participating agents* on side $-k$. Note the contrast to the pricing practiced by a profit-maximizing monopolist, whereby only the externality imposed on the *marginal* agents from the opposite side is accounted for.

3.3 Distortions

To understand how the profit-maximizing price profile compares to its efficient counterpart, let us take the difference between (5) and (6). This leads to the following decomposition:

$$P_k^* - P_k^e = \underbrace{\frac{P_k^*}{\varepsilon_k(P_a^*, P_b^*)}}_{\text{markup}} + \underbrace{N_{-k}^e\left(\bar{\omega}_{-k}^i(N_a^e, N_b^e) - \tilde{\omega}_{-k}^i(N_a^e, N_b^e)\right)}_{\text{Spence distortion}}$$
$$+ \underbrace{N_{-k}^e\left(\tilde{\omega}_{-k}^i(N_a^e, N_b^e) - \tilde{\omega}_{-k}^i(N_a^*, N_b^*)\right)}_{\text{displacement distortion}} + \underbrace{\left(N_{-k}^e - N_{-k}^*\right)\left(\tilde{\omega}_{-k}^i(N_a^*, N_b^*) - c^i\right)}_{\text{scale distortion}}. \tag{7}$$

The usual markup distortion reflects the market power enjoyed by the monopolistic platform. The Spence distortion captures the fact that a profit-maximizing monopolist, when setting the price on side k, internalizes the effect of expanding the side-k participation on the marginal agent from side $-k$, rather than on all participating agents from side $-k$. The displacement distortion accounts for the difference between the benefits that the marginal agents on side $-k$ derive from the expansion of the side-k participation under the profit-maximizing and the efficient allocation, respectively. In turn, the scale distortion reflects the difference between the participation on side $-k$ induced by the profit-maximizing monopolist and a welfare-maximizing platform: the net average benefit $\tilde{\omega}_{-k}^i - c^i$ that the marginal agents on side $-k$ derive from the expansion of the side-k participation (net of the platform's interaction cost) applies to a measure of agents equal to N_{-k}^* under profit maximization, whereas it applies to N_{-k}^e agents under welfare maximization.

In general, it is not possible to sign the net effect of these four distortions. As a result, the profit-maximizing prices can be either higher or lower than their efficient counterparts in either one or both sides of the market. To obtain further insights, it is useful to express prices on a *per-unit* basis, that is, by normalizing the side-k price by the size of the participation on side $-k$. Further assume that the participation costs are equal to zero on each side so that $c_a^s = c_b^s = 0$. Letting $p_k^* := \frac{P_k^*}{N_{-k}^*}$ and $p_k^e = \frac{P_k^e}{N_{-k}^e}$ denote the side-k per-unit price under profit and welfare maximization, respectively, we then have that

$$p_k^* - p_k^e = \underbrace{\frac{p_k^*}{\varepsilon_k(N_a^*, N_b^*)}}_{\text{markup}} + \underbrace{\left(\bar{\omega}_{-k}^i(N_a^e, N_b^e) - \tilde{\omega}_{-k}^i(N_a^e, N_b^e)\right)}_{\text{Spence distortion}}$$
$$+ \underbrace{\left(\tilde{\omega}_{-k}^i(N_a^e, N_b^e) - \tilde{\omega}_{-k}^i(N_a^*, N_b^*)\right)}_{\text{displacement distortion}}, \tag{8}$$

Corollary 5. *Consider a monopolistic platform designing a single network, and let the agents' preferences be given by (2).*

1. *Under Scenario* 1*, the displacement and Spence distortions are nil. The per-unit price on each side is higher under profit than welfare maximization:* $p_k^* > p_k^e$*,* $k = a, b$*.*

2. *Under Scenario 2, the Spence distortion is always positive on both sides, whereas the displacement distortion is negative on at least one side. Moreover, the sum of the per-unit prices is higher under profit than welfare maximization:* $p_a^* + p_b^* > p_a^e + p_b^e$*.*

Proof. Part 1 follows directly from (8). For part 2, note that $\bar{\omega}_{-k}^i(N_a^e, N_b^e) - \tilde{\omega}_{-k}^i(N_a^e, N_b^e) > 0$ because, by definition, $\tilde{\omega}_{-k}^i(N_a^e, N_b^e) = p_{-k}^e$ whereas $\bar{\omega}_{-k}^i(N_a^e, N_b^e)$ is the expectation over all ω_{-k}^i satisfying $\omega_{-k}^i \geq p_{-k}^e$. Hence, the Spence distortion is always positive. Because, $\tilde{\omega}_{-k}^i(N_a^e, N_b^e) = p_{-k}^e$ and $\tilde{\omega}_{-k}^i(N_a^*, N_b^*) = p_{-k}^*$, Condition (8) can be rewritten as

$$p_a^* + p_b^* = p_a^e + p_b^e + \underbrace{\frac{p_k^*}{\varepsilon_k(N_a^*, N_b^*)}}_{\text{markup}} + \underbrace{\left(\bar{\omega}_{-k}^i(N_a^e, N_b^e) - \tilde{\omega}_{-k}^i(N_a^e, N_b^e)\right)}_{\text{Spence distortion}}$$

for $k = a, b$. Because the markup and the Spence distortions are both positive, we have that $p_a^* + p_b^* > p_a^e + p_b^e$. Therefore, for some k, $p_k^* > p_k^e$. Because the displacement distortion is equal to $-(p_{-k}^* - p_{-k}^e)$, it follows that the displacement distortion is negative on at least one side of the market. □

Corollary 5 shows that, under scenarios 1 or 2, total per-unit prices are excessively high when set by a profit-maximizing monopolist. Interestingly, in Scenario 2, this does not rule out the possibility that the profit-maximizing price on one side of the market is lower than its efficient counterpart. As we shall see in Section 5, this possibility stems from the platform's inability to discriminate among agents from the same side of the market. When discrimination is possible, all agents from both sides face higher prices under profit than under welfare maximization.

4 Unknown preference distribution

We now consider markets in which the joint distribution of preferences in the population is *unknown*, both to the platforms and to each agent from either side of the market. This dimension is important because it introduces uncertainty over the size of the network externalities.

Specifically, suppose that preferences are consistent with the specification in Scenario 1. The uncertain "aggregate state" of the world is thus given by the pair of distributions $F^s = (F_a^s, F_b^s)$ from which the agents' stand-alone valuations are drawn. As in the previous section, each ω_k^s is drawn from F_k^s independently across agents. To keep things simple, further assume that each agent's stand-alone valuation ω_k^s parametrizes both the agent's preferences and the agent's beliefs over the aggregate state. For simplicity, the platforms are assumed not to possess any private information.

For any $\omega_k^s \in \mathbb{R}$, $k = 1, 2$, then denote by $Q_k^s(\omega_k^s)$ the measure of agents from side k that the platform believes to have stand-alone valuations no smaller than ω_k^s.

Next, consider the agents. For any $k = a, b$, and any $(\omega_k^s, \omega_{-k}^s) \in \mathbb{R}^2$, let $M_{-k}^s\left(\omega_{-k}^s|\omega_k^s\right)$ denote the measure of agents from side $-k$ with stand-alone valuation no smaller than ω_{-k}^s, as expected by any agent from side k with stand-alone valuation equal to ω_k^s. These functions thus reflect the agents' beliefs over the cross-sectional distribution of preferences on the other side of the market. They may capture, for example, how consumers use their own appreciation for the features of a new platform's product (e.g., its operating system, interface, and the like) to form beliefs over the number of applications that will be developed for the new product. Importantly, such beliefs need not coincide with the platforms' beliefs.

For any ω_k^s, we assume that $M_{-k}^s\left(\omega_{-k}^s|\omega_k^s\right)$ is strictly decreasing in ω_{-k}^s, and differentiable in each argument.

Definition 6. Preferences are *aligned* if, for all ω_{-k}^s, $M_{-k}^s\left(\omega_{-k}^s|\omega_k^s\right)$ is increasing in ω_k^s, $k = a, b$. They are *misaligned* if, for all ω_{-k}^s, $M_{-k}^s\left(\omega_{-k}^s|\omega_k^s\right)$ is decreasing in ω_k^s, $k = a, b$.

When preferences are aligned, agents with a higher appreciation for a platform's product also expect a higher appreciation by agents from the opposite side, whereas the opposite occurs when preferences are misaligned. Importantly, the definition does not presume that stand-alone valuations are drawn from a common prior. It simply establishes a monotone relationship between beliefs and stand-alone valuations.

The special case of a common prior corresponds to the case in which all players commonly believe that $F^s = (F_a^s, F_b^s)$ is drawn from a set of distributions $\mathscr{F}$ according to a distribution $\mathbf{F}$. In this case, the platforms' and the agents' beliefs are given by

$$Q_k^s(\omega_k^s) = \mathbb{E}_{\mathbf{F}}[1-F_k^s(\omega_k^s)], \quad \text{and} \quad M_{-k}^s\left(\omega_{-k}^s|\omega_k^s\right) = \frac{\mathbb{E}_{\mathbf{F}}\left[\left(1 - F_{-k}^s(\omega_{-k}^s)\right) f_k^s\left(\omega_k^s\right)\right]}{\mathbb{E}_{\mathbf{F}}\left[f_k^s\left(\omega_k^s\right)\right]},$$

where, as in the baseline model, f_k^s denotes the density of F_k^s, and all expectations are computed by integrating over $\mathscr{F}$ under the common prior $\mathbf{F}$.

A strategy profile for the agents then constitutes a *continuation (Bayes-Nash) equilibrium* in the game that starts after the platform announces its access prices $P = (P_a, P_b)$ if each agent's participation decision is a best response to all other agents' equilibrium strategies.

Each agent from side k with stand-alone valuation ω_k^s then joins the platform if and only if

$$\omega_k^s + \mathring{\omega}_k^i \mathbb{E}\left[N_{-k}|\omega_k^s\right] \geq P_k, \tag{9}$$

where $\mathbb{E}\left[N_{-k}|\omega_k^s\right]$ is the participation on side $-k$ expected by the agent. Provided the interaction benefits $(\mathring{\omega}_a^i, \mathring{\omega}_b^i)$ are not too large, we then have that, for any vector of prices (P_a, P_b), the demand expected by the platform on each side $k = a, b$ is given by $Q_k^s(\hat{\omega}_k^s)$, where $(\hat{\omega}_a^s, \hat{\omega}_b^s)$ is the unique solution to the system of equations given by

$$\hat{\omega}_k^s + \hat{\omega}_k^i M_{-k}^s \left(\hat{\omega}_{-k}^s | \hat{\omega}_k^s\right) = P_k,\ k = a, b. \tag{10}$$

Now, suppose the platform aims at getting on board N_a agents from side a and N_b agents from side b. Because the platform does not know the exact distribution of preferences, N_a and N_b must be interpreted as the participation expected by the platform, where the expectation is taken over all possible distributions F^s using the platform's own beliefs. Given (10), the platform should set prices (P_a, P_b) such that the thresholds $(\hat{\omega}_a^s, \hat{\omega}_b^s)$ satisfy $Q_k(\hat{\omega}_k^s) = N_k,\ k = a, b$. The key difference with respect to the case of complete information is the following. When the platform adjusts its prices so as to change the participation it expects from side k while keeping constant the participation it expects from side $-k$, it need not be able to keep constant the side-k's marginal agent's beliefs over the participation of side $-k$. This is because uncertainty over the distribution of preferences in the population de facto introduces statistical dependence between the agents' beliefs over the other side's participation and their own preferences, reflected in the fact that

$$\frac{\partial \mathbb{E}[N_{-k} | \hat{\omega}_k^s]}{\partial \hat{\omega}_k^s} = \frac{\partial M_{-k}^s \left(\hat{\omega}_{-k}^s | \hat{\omega}_k^s\right)}{\partial \hat{\omega}_k^s} \neq 0. \tag{11}$$

In particular, when preferences are aligned between the two sides, $\mathbb{E}[N_{-k} | \hat{\omega}_k^s]$ is increasing in $\hat{\omega}_k^s$, whereas the opposite is true when preferences are misaligned. In the first case, this novel effect contributes to *steeper* inverse demand curves, whereas in the second case to *flatter* inverse demands. When preferences are aligned, the new marginal agent that the platform attracts by lowering its price on side k is more pessimistic about the participation of the other side than any of the infra-marginal agents who are already on board (those with a higher stand-alone valuation). To get the new marginal agent on board, the platform must thus cut its side-k price more than what it would have done under complete information. Importantly, this novel effect is present even if the platform adjusts its price on side $-k$ so as to maintain its expectation of that side's participation constant (which amounts to maintaining $\hat{\omega}_{-k}^s$ constant).

The above novel effects play an important role for how platforms price access to their network on each side of the market. Using (10), we can reformulate the platform's objective in terms of the participation thresholds $(\hat{\omega}_a^s, \hat{\omega}_b^s)$ rather than the prices (P_a, P_b) that induce these thresholds. Accordingly, the platform chooses $(\hat{\omega}_a^s, \hat{\omega}_b^s)$ to maximize

$$\sum_{k=a,b} \left\{\hat{\omega}_k^s + \hat{\omega}_k^i M_{-k}^s \left(\hat{\omega}_{-k}^p | \hat{\omega}_k^s\right) - c_k^s\right\} Q_k^s(\hat{\omega}_k^s) - c^i Q_a^s(\hat{\omega}_a^s) Q_b^s(\hat{\omega}_b^s). \tag{12}$$

We then have the following result:

Proposition 7. *Suppose that preferences are as in Scenario 1 and that the distribution of preferences over the two sides is unknown to the agents and the platform. The profit-maximizing prices* (P_a^*, P_b^*)*, along with the stand-alone thresholds* $(\hat{\omega}_a^{s*}, \hat{\omega}_b^{s*})$ *they induce, satisfy the following optimality conditions*

$$P_k^* = c_k^s + c^i Q_{-k}^s(\hat{\omega}_{-k}^{s*}) + \left[1 + \mathring{\omega}_k^i \frac{\partial M_{-k}^s\left(\hat{\omega}_{-k}^{s*}|\hat{\omega}_k^{s*}\right)}{\partial \hat{\omega}_k^s}\right] \frac{Q_k^s(\hat{\omega}_k^{s*})}{|dQ_k^s(\hat{\omega}_k^{s*})/d\hat{\omega}_k^s|}$$
$$+\mathring{\omega}_{-k}^i \frac{\partial M_k^s\left(\hat{\omega}_k^{s*}|\hat{\omega}_{-k}^{s*}\right)}{\partial \hat{\omega}_k^s} \frac{Q_{-k}^s(\hat{\omega}_{-k}^{s*})}{|dQ_k^s(\hat{\omega}_k^{s*})/d\hat{\omega}_k^s|}, \tag{13}$$

with $P_k^* = \hat{\omega}_k^{s*} + \mathring{\omega}_k^i M_{-k}^s\left(\hat{\omega}_{-k}^{s*}|\hat{\omega}_k^{s*}\right), \quad k = a, b.$

Proof. The result follows directly from differentiating the profit function in (12) and then using (10). □

Note that the price formula in (13) is the incomplete-information analog of the corresponding complete-information formula

$$P_k^* = c_k^s + c^i N_{-k}^* - \frac{\partial P_k(N_a^*, N_b^*)}{\partial N_k} N_k^* - \frac{\partial P_{-k}(N_a^*, N_b^*)}{\partial N_k} N_{-k}^* \tag{14}$$

derived above. It requires that profit does not change when the platform increases the participation it expects from side k (given its own beliefs), while adjusting the price on side $-k$ to maintain the participation it expects from side $-k$ constant.

In particular, the last term in the right-hand side of (13) is the benefit of cutting the price on side k due to the possibility of raising the price on side $-k$, typical of two-sided markets (see (5)).[2] Note, however, an important difference with respect to complete information. The measure of additional agents from side k that the platform expects to bring on board by cutting its price on side k now differs from the measure of agents expected by the marginal agent on side $-k$ (the one with signal $\hat{\omega}_{-k}^{s*}$ who is just indifferent between joining and not joining). This novel effect is captured by the term

$$\left.\frac{\partial \mathbb{E}[N_k \mid \hat{\omega}_{-k}^{s*}]}{\partial N_k}\right|_{\hat{\omega}_{-k}^{s*}=const} = -\frac{\partial M_k^s\left(\hat{\omega}_k^{s*} \mid \hat{\omega}_{-k}^{s*}\right)}{\partial \hat{\omega}_k^s} \frac{1}{|dQ_k^s(\hat{\omega}_k^{s*})/d\hat{\omega}_k^s|}$$

in (13).[3] Irrespective of whether preferences are aligned or misaligned between the two sides, this term is always positive, thus contributing to a lower price on side k.

[2] Observe that $\partial M_k\left(\hat{\omega}_k^{s*} \mid \hat{\omega}_{-k}^{s*}\right)/\partial \hat{\omega}_k^s < 0$, irrespective of whether preferences are aligned or misaligned.

[3] Under complete information, $\frac{\partial M_k^s\left(\hat{\omega}_k^s|\hat{\omega}_{-k}^p\right)}{\partial \hat{\omega}_k^s} = \frac{1}{|dQ_k^s(\hat{\omega}_k^s)/d\hat{\omega}_k^s|}$ in which case the second term in (13) reduces to $\mathring{\omega}_{-k}^i N_{-k}$, as discussed above.

The term in square brackets in the right-hand side of (13) captures the adjustment in the side-k price necessary to expand the side-k demand. Interestingly, this term accounts for how a variation in the participation from side k comes with a variation in the beliefs of the side-k's marginal agent about the participation from side $-k$ (the second term in the square bracket). Such variation occurs even when the platform adjusts its price on side $-k$ to maintain the identity of the marginal agent on that side unchanged (thus maintaining the participation expected by the platform from that side constant). As indicated above, when preferences are aligned across sides, this effect contributes to a steeper inverse demand on each side and hence, other things equal, to higher prices. The opposite is true in markets in which preferences are misaligned. The latter effect has no counterpart under complete information.

We conclude by comparing the profit-maximizing prices to their efficient counterparts. To this purpose, consider the problem of a planner that shares the same beliefs as the platform (which is always the case when stand-alone values are drawn from a common prior). The planner's problem then consists in choosing $(\hat{\omega}_a^s, \hat{\omega}_b^s)$ so as to maximize

$$\hat{W}\left(\hat{\omega}_a^s, \hat{\omega}_b^s\right) := \sum_{k=a,b} \int_{\left\{\omega_k^s \geq \hat{\omega}_k^s\right\}} \left(\omega_k^s + \mathring{\omega}_k^i M_{-k}^s\left(\hat{\omega}_{-k}^s \mid \omega_k^s\right) - c_k^s\right) d[1 - Q_k^s(\omega_k)]$$
$$-c^i Q_a^s(\hat{\omega}_a^s) Q_b^s(\hat{\omega}_b^s).$$

Proposition 8. *Consider a welfare-maximizing platform designing a single network, and let the agents' preferences be as in Scenario 1. The welfare-maximizing prices* (P_a^e, P_b^e)*, along with the efficient stand-alone thresholds* $(\hat{\omega}_a^{se}, \hat{\omega}_b^{se})$ *they induce solve*

$$P_k^e = c_k^s + c^i Q_{-k}^s(\hat{\omega}_{-k}^{se}) + \mathring{\omega}_{-k}^i \frac{\int_{\left\{\omega_{-k}^s \geq \hat{\omega}_{-k}^{se}\right\}} \left(\frac{\partial M_k^s\left(\hat{\omega}_k^{se} \mid \omega_{-k}^s\right)}{\partial \hat{\omega}_k^s}\right) d[1 - Q_{-k}^s(\omega_{-k}^s)]}{\left|dQ_k^s(\hat{\omega}_k^{se})/d\hat{\omega}_k^{se}\right|} \tag{15}$$

where $P_k^e = \hat{\omega}_k^{se} + \mathring{\omega}_k^i M_{-k}^s\left(\hat{\omega}_{-k}^{se} \mid \hat{\omega}_k^{se}\right)$*, for* $k = a, b$.

Proof. The result follows directly from differentiating the welfare function in (12) and then using (10) to relate the stand-alone thresholds to the prices that induce them. □

When combined with Proposition 7, the result in the previous proposition identifies the distortions due to market power. First, as under complete information, a profit-maximizing platform accounts for the effect that a reduction in the side-k price has on the profit collected from all infra-marginal agents from the same side. This effect is captured by the third term in the right-hand side of (13). This is the same markup distortion discussed above. The novelty relative to complete information is that the adjustment in the side-k price must account for the difference between the platform's beliefs and those of the side-k's marginal agents. Other things equal, such a difference contributes to larger distortion when preferences are aligned and to a smaller one when they are misaligned. Second, a profit-maximizing platform internalizes the effect of expanding the side-k participation by looking at the externality exerted on the marginal agent from side $-k$ instead of all participating agents from side $-k$. This effect is the analog of the sum of the Spence and the displacement distortions discussed above and is captured by the difference between the last term in (13) and the last term in (15). Other things equal, whether such distortions are amplified or mitigated by dispersed information depends to a large extent on the modularity of the beliefs, that is, on whether $|\partial M_k^s \left(\hat{\omega}_k^{se} \mid \omega_{-k}^s\right) / \partial \hat{\omega}_k^s|$ is increasing or decreasing in ω_{-k}^s. The last factor contributing to a discrepancy between the profit-maximizing and the welfare-maximizing prices is a scale distortion analogous to the one under complete information, but again adjusted for the difference in beliefs between the marginal and the infra-marginal agents.

5 Matching Design

Matching design relaxes one of the the key restrictions in the analysis of pricing in two-sided markets; namely the absence of discrimination (within agents from the same side). This opens the door to customized matching rules, where the matching set of each participating agent depends on his type.

5.1 One-to-One Matching

We first consider markets in which matching is one-to-one, capturing situations or rivalry or of severe capacity constraints on the part of agents. Recall that, in this case, the set $\mathcal{C}(\Omega_{-k})$ is the collection of all singletons $\{\omega_{-k}\}$, with $\omega_{-k} \in \Omega_{-k}$.

We identify the type of each agent with a vertical characteristic which we refer to as *quality*. We let $\Omega_k = [\underline{\omega}_k, \bar{\omega}_k] \subset \mathbb{R}_{++}$ denote the set of types from side k, and assume that the type of each side-k agent is an independent draw from the distribution F_k.

Agents' preferences take the following form: for each $\omega_k \in \Omega_k$, the gross utility that each side-k agent with type ω_k derives from the matching set $s_k = \{\omega_{-k}\}$ is given by

$$u_k(s_k|\omega_k) := \phi_k(\omega_k, \omega_{-k}), \tag{16}$$

where the function ϕ_k is differentiable, equi-Lipschitz continuous, strictly increasing in both arguments, and supermodular. Accordingly, the surplus of each match increases with each of the involved agents' quality, and the gain from a better-quality partner is higher for agents of higher quality. A simple example satisfying these assumptions is the multiplicative surplus function $\phi_a(\omega_a, \omega_b) = \omega_a \omega_b$ often assumed in applications.

The cost the platform incurs for each match it induces is $c \in \mathbb{R}_+$. To make things simple (but interesting), assume that $\sum_k \phi_k(\underline{\omega}_a, \underline{\omega}_b) \leq c \leq \sum_k \phi_k(\bar{\omega}_k, \bar{\omega}_k)$.

To properly describe the platform's matching design problem, we first need to amend the definition of matching rules introduced in Section 2. Namely, it is necessary to allow for stochastic matching rules that assign to each type a distribution over the type of the matching partner from the opposite side.

To do so formally, let $\hat{\Omega}_k$ be the (Lebesgue-measurable) set of participating types from side k. For any $\omega_k \in \Omega_k \setminus \hat{\Omega}_k$, $\mathbf{s}_k(\omega_k) = \emptyset$. For $\omega_k \in \hat{\Omega}_k$, instead, $\mathbf{s}_k(\omega_k) \in \Delta(\hat{\Omega}_{-k})$ is a probability measure over $\hat{\Omega}_{-k}$ describing the likelihood that each side-k agent of type ω_k is matched with any of the participating types from the other side of the market. We denote by $G_{\mathbf{s}_k}(\cdot|\omega_k)$ the cdf associated with the measure $\mathbf{s}_k(\omega_k)$.

Feasibility then dictates that the sets of participating agents from the two sides have the same measure, $|\hat{\Omega}_a| = |\hat{\Omega}_b|$, where, as before, $|\hat{\Omega}_k|$ is the F_k-measure of the set $\hat{\Omega}_k$. This requirement is self-explanatory, as no one-to-one matching rule can be constructed when the mass of participating agents is unequal across sides. In addition, feasibility also requires that the measure of types from side $-k$ matched to any subset $\tilde{\Omega}_k \subseteq \hat{\Omega}_k$ of participating types from side k have the same measure as $\tilde{\Omega}_k$. That is, for any $\tilde{\Omega}_k \subseteq \hat{\Omega}_k$,

$$|\tilde{\Omega}_k| := \int_{\tilde{\Omega}_k} dF_k(\omega_k) = \int_{\hat{\Omega}_{-k}} \int_{\tilde{\Omega}_k} dG_{\mathbf{s}_{-k}}(\omega_k|\omega_{-k}) dF_{-k}(\omega_{-k}),$$

where the double integral in the right-hand side is the total measure of agents from side $-k$ that are matched to those agents from side k whose type is in $\tilde{\Omega}_k$. Because the equality above has to hold for any measurable set $\tilde{\Omega}_k$, we can define the joint distribution $\mathbf{F}$ according to $d\mathbf{F}(\omega_k, \omega_{-k}) := dG_{\mathbf{s}_{-k}}(\omega_k|\omega_{-k})dF_{-k}(\omega_{-k})$. Note that, by construction, this joint distribution couples the marginals $|\hat{\Omega}_a|^{-1}F_a$ and $|\hat{\Omega}_b|^{-1}F_b$.

Accordingly, feasibility requires that *there exists* a joint distribution $\mathbf{F}$, with support $\hat{\Omega}_a \times \hat{\Omega}_b$ and marginals $|\hat{\Omega}_a|^{-1}F_a$ and $|\hat{\Omega}_b|^{-1}F_b$, such that, for each $\omega_k \in \hat{\Omega}_k$, $G_{\mathbf{s}_k}(\cdot|\omega_k)$ is the conditional distribution of ω_{-k} given ω_k, as induced by the joint cdf $\mathbf{F}$. Intuitively, this requirement guarantees that *any* profile of *realized* matches satisfies the reciprocity condition (1). Heuristically, the function $\mathbf{F}$ describes the distribution of matched pairs under the rule $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$.

For instance, *random matching* (with full participation, that is, with $\hat{\Omega}_k = \Omega_k$, $k = a, b$) corresponds to the joint cdf $\mathbf{F}$ defined by $\mathbf{F}(\omega_a, \omega_b) = F_a(\omega_a)F_b(\omega_b)$, all $(\omega_a, \omega_b) \in \Omega$, and the associated stochastic matching rule has cdf's $G_{\mathbf{s}_a}(\cdot|\omega_a) = F_b$ and $G_{\mathbf{s}_b}(\cdot|\omega_b) = F_a$, for all $(\omega_a, \omega_b) \in \Omega$. In this example, the partner of each side-k agent is drawn (independently across agents) from the marginal distribution F_{-k}, irrespectively of the agent's own type.

The case of a *deterministic* matching rule corresponds to the case where each type from each side $k = a, b$ is assigned a single type from the opposite side. We shall abuse notation and write the corresponding matching rule as $\mathbf{s}_k(\omega_k) = \omega_{-k}$, with the understanding that, in this case, $G_{\mathbf{s}_k}(\cdot|\omega_k)$ is a degenerate Dirac delta assigning probability one to type $\{\omega_{-k}\}$. Formally, this corresponds to an endogenous joint distribution $\mathbf{F}$ whose conditional distribution $\mathbf{F}_{-k|k}$ specifies a collection of Dirac deltas, one for each ω_k from sides $k = a, b$.

A deterministic matching rule of special interest is the (truncated*) positive assortative* one. In order to define it formally, we first need to introduce the following:

Definition 9. A pair of absolutely continuous random variables (X, Y), with cdf's F_X and F_Y, respectively, is co-monotone if there is a random variable U uniformly distributed over $[0, 1]$ such that $X = F_X^{-1}(U)$ and $Y = F_Y^{-1}(U)$, where $F_k^{-1}(U) := \inf\{\omega_k \in \Omega_k : F_k(\omega_k) \geq U\}$.

As it is well known, any random variable can be represented as being generated from a draw from a uniform distribution (probability integral transform theorem). The definition imposes that the draw be the same across the two random variables, implying that the two variables are related by the identity $Y = F_Y^{-1}(F_X(X))$. Co-monotone random variables are intimately related to positive assortative matching:

Definition 10. A matching rule $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ is truncated positive assortative if, for each $k \in \{a, b\}$, there exists a participation threshold $\hat{\omega}_k \in \Omega_k$ such that

$$\mathbf{s}_k(\omega_k) = \begin{cases} (F_{-k})^{-1}(F_k(\omega_k)) & \text{if} \quad \omega_k \geq \hat{\omega}_k \\ \emptyset & \text{if} \quad \omega_k < \hat{\omega}_k. \end{cases}$$

Therefore, a positive assortative matching rule renders the types of matched agents co-monotone random variables (also note that, by feasibility, $F_a(\hat{\omega}_a) = F_b(\hat{\omega}_b)$).

5.1.1 Efficient Matching Design

Equipped with the notation introduced above, the platform's welfare-maximization problem consists of choosing a pair of sets $(\hat{\Omega}_a, \hat{\Omega}_b)$ and a joint distribution $\mathbf{F}$ over $\hat{\Omega}_a \times \hat{\Omega}_b$ that couples the marginals $|\hat{\Omega}_a|^{-1}F_a$ and $|\hat{\Omega}_b|^{-1}F_b$ to maximize[4]

$$\hat{W}(\mathbf{F}; \hat{\Omega}_a, \hat{\Omega}_b) := |\hat{\Omega}_a| \int_{\hat{\Omega}_a \times \hat{\Omega}_b} (\phi_a(\omega_a, \omega_b) + \phi_b(\omega_b, \omega_a) - c)\, d\mathbf{F}(\omega_a, \omega_b).$$

Proposition 11. *Consider a welfare-maximizing platform designing one-to-one matches, and let the agents' preferences be given by (16). The efficient matching rule is truncated positive assortative with participation thresholds* $(\hat{\omega}_a^e, \hat{\omega}_b^e)$ *satisfying* $\hat{\omega}_b^e = F_b^{-1}(F_a(\hat{\omega}_a^e))$ *and*

$$\phi_a(\hat{\omega}_a^e, \hat{\omega}_b^e) + \phi_b(\hat{\omega}_a^e, \hat{\omega}_b^e) = c. \tag{17}$$

Proof. First, consider the case in which $c = 0$. It is evident that, in this case, it is never optimal to exclude any type, implying that $(\hat{\Omega}_a, \hat{\Omega}_b) = (\Omega_a, \Omega_b)$. Next, let the the match surplus function be step-function with coefficients $(\alpha, \beta) \in \mathbb{R}^2$:

$$\phi_k(\omega_a, \omega_b) = \phi^{\alpha,\beta}(\omega_a, \omega_b) := \mathbf{1}\{\omega_a \geq \alpha\}\, \mathbf{1}\{\omega_b \geq \beta\}$$

[4]Because of one-to-one matching, the measure of matches induced by the platform is $|\hat{\Omega}_a| = |\hat{\Omega}_b|$.

Observe that, in this case, the welfare-maximizing matching rule is untruncated positive assortative (in that $\hat{\omega}_k = \underline{\omega}_k$ for each k).

For general match surplus functions, let $\phi(\omega_a, \omega_b) := \sum_{k=a,b} \phi_k(\omega_k, \omega_{-k})$, and consider the function $\Phi(\omega_a, \omega_b)$ defined by

$$\Phi(\omega_a, \omega_b) := \frac{\phi(\omega_a, \omega_b) - \phi(\underline{\omega}_a, \omega_b) - \phi(\omega_a, \underline{\omega}_b) + \phi(\underline{\omega}_a, \underline{\omega}_b)}{\phi(\bar{\omega}_a, \bar{\omega}_b) - \phi(\underline{\omega}_a, \bar{\omega}_b) - \phi(\bar{\omega}_a, \underline{\omega}_b) + \phi(\underline{\omega}_a, \underline{\omega}_b)}$$

for all $(\omega_a, \omega_b) \in \Omega$. Because ϕ is supermodular, Φ is a cdf with support Ω. Obviously,

$$\Phi(\omega_a, \omega_b) = \int_\Omega \phi^{\alpha,\beta}(\omega_a, \omega_b) d\Phi(\alpha, \beta).$$

Now consider the objective

$$\tilde{W}(\mathbf{F}) := \int_\Omega \Phi(\omega_a, \omega_b) d\mathbf{F}(\omega_a, \omega_b) = \int_\Omega \int_\Omega \phi^{\alpha,\beta}(\omega_a, \omega_b) d\mathbf{F}(\omega_a, \omega_b) d\Phi(\alpha, \beta),$$

where the equality follows from Fubini's Theorem. From the arguments above, among all joint cdf's that couple F_a and F_b, the one that maximizes the integral $\int_\Omega \phi^{\alpha,\beta}(\omega_a, \omega_b) d\mathbf{F}(\omega_a, \omega_b)$ is the cdf $\mathbf{F}^*$ that renders ω_a and ω_b co-monotone random variables. Because this is true for all $(\alpha, \beta) \in \Omega$, we conclude that $\tilde{W}(\mathbf{F})$ is maximized by $\mathbf{F}^*$. Because Φ has the form $\Phi(\omega_a, \omega_b) = \delta\phi(\omega_a, \omega_b) + \sum_k \gamma_k \varphi_k(\omega_k) + K$, where δ, γ_a, γ_b and K are constants, and where each function $\varphi_k(\omega_k)$ is invariant in ω_{-k}, it is then easy to see that, when $c = 0$, $\mathbf{F}^*$ also maximizes $\hat{W}(\mathbf{F}; \Omega_a, \Omega_b)$ over all joint cdf's $\mathbf{F}$ that couple the marginals F_a and F_b. We conclude that the un-truncated positive assortative matching rule is optimal when $c = 0$.

Next consider the case in which $c > 0$. Because each $\phi_k(\omega_a, \omega_b)$ is strictly increasing in both arguments, it is easy to see that the welfare maximizing participation sets $\hat{\Omega}_k^e$ have the form $\hat{\Omega}_k^e = [\hat{\omega}_k^e, \bar{\omega}_k]$, for some $\hat{\omega}_k^e \geq \underline{\omega}_k$, $k = a, b$. Because it is never efficient to match any pair of types $(\omega_a, \omega_b) \in \Omega$ for which $\sum_k \phi_k(\omega_a, \omega_b) - c < 0$, we then conclude that the welfare-maximizing participation thresholds $\hat{\omega}_a^e$ and $\hat{\omega}_b^e$ satisfy (17). Applying the arguments above to the truncated marginal $|\hat{\Omega}_a^e|^{-1} F_a$ and $|\hat{\Omega}_b^e|^{-1} F_b$ then yields the result that the function $\hat{W}(\mathbf{F}; \hat{\Omega}_a, \hat{\Omega}_b)$ is maximized by the truncated positive assortative rule with participation sets $(\hat{\Omega}_a^e, \hat{\Omega}_b^e)$.

The next lemma establishes the implementability of truncated positive assortative rules.

Lemma 12. *Let $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ be a truncated positive assortative matching rule. Then $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ can be implementable by the following payment rule:*

$$\mathbf{p}_k(\omega_k) = \phi_k(\omega_k, \mathbf{s}_k(\omega_k)) - \int_{\hat{\omega}_k}^{\omega_k} \frac{\partial \phi_k}{\partial \omega_k}(\tilde{\omega}_k, \mathbf{s}_k(\tilde{\omega}_k)) d\tilde{\omega}_k,$$

all $\omega_k \in \Omega_k$, $k = a, b$.

Proof. Consider the mechanism defined by the matching rule $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ along with the payment rule $\{\mathbf{p}_k(\cdot)\}_{k=a,b}$ in the lemma. Each agent with type ω_k then chooses a report ω'_k so as to maximize

$$\phi_k(\omega_k, \mathbf{s}_k(\omega'_k)) - \phi_k(\omega'_k, \mathbf{s}_k(\omega'_k)) + \int_{\hat{\omega}_k}^{\omega'_k} \frac{\partial \phi_k}{\partial \omega_k}(\tilde{\omega}_k, \mathbf{s}_k(\tilde{\omega}_k))d\tilde{\omega}_k.$$

Because $\mathbf{s}_k(\cdot)$ is monotone and ϕ_k is supermodular, it is then easy to see that truthful reporting uniquely maximizes the agent's payoff. □

The result in the proposition then follows from the arguments above along with Lemma 12. □

It is easy to verify that the efficient rule is essentially unique (that is, up to zero-measure perturbations). This implies that random matching is never optimal provided the total match surplus ϕ is supermodular.

5.1.2 Profit-Maximizing Matching Design

Having derived the efficient one-to-one matching rule, we now turn to its profit-maximizing counterpart. The platform's profit under any incentive compatible mechanism $\{\mathbf{s}_k(\cdot), \mathbf{p}_k(\cdot)\}_{k=a,b}$ is equal to

$$\Pi := \sum_{k=a,b} \int_{\hat{\Omega}_k} \mathbf{p}_k(\omega_k)dF_k(\omega_k) - c|\hat{\Omega}_a|.$$

The next lemma expresses the platform's profit solely in terms of the endogenous joint distribution $\mathbf{F}$ induced by the matching rule $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$.

Lemma 13. *Suppose that $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ is an implementable matching rule with participation sets $\hat{\Omega}_k = [\hat{\omega}_k, \overline{\omega}_k]$, $\hat{\omega}_k \in \Omega_k$, $k = a, b$, and let $\mathbf{F}$ be the endogenous joint distribution over $\hat{\Omega}_a \times \hat{\Omega}_b$ induced by $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$. The platform's maximal profit under such a rule are given by*

$$\Pi(\mathbf{F}; \hat{\Omega}_a, \hat{\Omega}_b) = |\hat{\Omega}_a| \int_{\hat{\Omega}_a \times \hat{\Omega}_b} \left(\widehat{\phi}_a(\omega_a, \omega_b) + \widehat{\phi}_b(\omega_b, \omega_a) - c\right) d\mathbf{F}(\omega_a, \omega_b),$$

where

$$\widehat{\phi}_k(\omega_k, \omega_{-k}) := \phi_k(\omega_k, \omega_{-k}) - \left(\frac{1 - F_k(\omega_k)}{f_k(\omega_k)}\right) \frac{\partial \phi_k}{\partial \omega_k}(\omega_k, \omega_{-k})$$

is the side-k virtual match surplus, $k = a, b$.

Proof. Let $\{\mathbf{p}_k(\cdot)\}_{k=a,b}$ be any payment rule implementing $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ and $M = \{\mathbf{s}_k(\cdot), \mathbf{p}_k(\cdot)\}_{k=a,b}$ be the mechanism defined by $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ and $\{\mathbf{p}_k(\cdot)\}_{k=a,b}$. For any participating type $\omega_k \in \hat{\Omega}_k$, the equilibrium payoff (that is, the payoff under truth-telling) is equal to

$$U_k(\omega_k; M) = \int_{\hat{\Omega}_{-k}} \phi_k(\omega_k, \omega_{-k})dG_{\mathbf{s}_k}(\omega_{-k}|\omega_k) - \mathbf{p}_k(\omega_k)$$

The envelope theorem implies that

$$U_k(\omega_k; M) = U_k(\hat{\omega}_k; M) + \int_{\hat{\omega}_k}^{\omega_k} \int_{\hat{\Omega}_{-k}} \frac{\partial \phi_k}{\partial \omega_k}(\tilde{\omega}_k, \omega_{-k}) dG_{\mathbf{s}_k}(\omega_{-k}|\omega_k) d\tilde{\omega}_k.$$

Hence,

$$\mathbf{p}_k(\omega_k) = \textstyle\int_{\hat{\Omega}_{-k}} \phi_k(\omega_k, \omega_{-k}) dG_{\mathbf{s}_k}(\omega_{-k}|\omega_k)$$

$$- \textstyle\int_{\hat{\omega}_k}^{\omega_k} \int_{\hat{\Omega}_{-k}} \frac{\partial \phi_k}{\partial \omega_k}(\tilde{\omega}_k, \omega_{-k}) dG_{\mathbf{s}_k}(\omega_{-k}|\omega_k) d\tilde{\omega}_k - U_k(\hat{\omega}_k; M).$$

Replacing the above expression into the profit function Π and integrating by parts we have that

$$\Pi = \sum_k \int_{\hat{\Omega}_k} \int_{\hat{\Omega}_{-k}} \left(\phi_k(\omega_k, \omega_{-k}) - \left(\frac{1 - F_k(\omega_k)}{f_k(\omega_k)} \right) \frac{\partial \phi_k}{\partial \omega_k}(\omega_k, \omega_{-k}) \right) dG_{\mathbf{s}_k}(\omega_{-k}|\omega_k) dF_k(\omega_k)$$

$$- \sum_k |\hat{\Omega}_k| U_k(\hat{\omega}_k; M) - c|\hat{\Omega}_a|. \tag{18}$$

Clearly, any profit-maximizing transfer scheme implementing the matching rule $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ must satisfy $U_k(\hat{\omega}_k; M) = 0$, $k = a, b$. Feasibility requires that $dG_{\mathbf{s}_k}(\omega_{-k}|\omega_k) dF_k(\omega_k) = |\hat{\Omega}_k| d\mathbf{F}(\omega_k, \omega_{-k})$. Combining the above properties then yields the result. □

Next, observe that, in any incentive compatible mechanism M, if $\mathbf{s}_k(\omega_k) \neq \emptyset$, then necessarily $\mathbf{s}_k(\omega'_k) \neq \emptyset$ for all $\omega'_k > \omega_k$, $k = a, b$. Hence, under any profit-maximizing mechanism, the participating set on each side has the interval structure $\hat{\Omega}_k = [\hat{\omega}_k, \overline{\omega}_k]$, for some $\hat{\omega}_k \in \Omega_k$. In light of this observation and the result in Lemma 13, we have that the problem of a profit-maximizing platform is identical to that of a welfare-maximizing one, after one replaces the match surplus function by its virtual counterpart. Accordingly, the next result follows from Proposition 11.

Proposition 14. *Consider a profit-maximizing platform designing one-to-one matches, and let the agents' preferences be given by (16). Suppose that, for each $k \in \{a, b\}$, the virtual match surplus function $\widehat{\phi}_k(\omega_k, \omega_{-k})$ is supermodular. Then, the profit-maximizing matching rule is a truncated positive assortative rule with participation thresholds $(\hat{\omega}^*_a, \hat{\omega}^*_b)$ satisfying $\hat{\omega}^*_b = F_b^{-1}(F_a(\hat{\omega}^*_a))$ and*

$$\widehat{\phi}_a(\hat{\omega}^*_a, \hat{\omega}^*_b) + \widehat{\phi}_b(\hat{\omega}^*_a, \hat{\omega}^*_b) = c.$$

Hence, when the virtual match surplus functions $\widehat{\phi}_k$ satisfy the same properties as their primitive counterparts ϕ_k, $k = a, b$, the matching partner of any participating agent is the same as under welfare maximization. Profit maximization, however, introduces extensive-margin distortions, in that it excludes an inefficiently large set of agents from both sides.

How strong is the requirement that virtual match surpluses are supermodular? For an illustration, suppose that $\phi_k(\omega_k, \omega_{-k}) = \omega_k \omega_{-k}$. Then $\widehat{\phi}_k(\omega_k, \omega_{-k})$ being supermodular is equivalent to the usual regularity condition from mechanism design requiring that virtual values

$$\omega_k - \left(\frac{1 - F_k(\omega_k)}{f_k(\omega_k)} \right)$$

being strictly increasing. For other match value functions, this supermodularity condition is harder to satisfy. When the condition is not satisfied, profit-maximization, in addition to excluding too many agents, leads to inefficient matching of the participating agents.[5]

5.2 Many-to-many matching design

Now suppose the platform can engage in many-to-many matching, therefore customizing the matching set that each agent receives. Formally, $\mathcal{C}(\Omega_{-k})$ is now the collection of all (measurable) subsets of Ω_{-k}. As in the previous subsection, continue to denote each agent's type by the unidimensional characteristics $\omega_k \in \Omega_k := [\underline{\omega}_k, \overline{\omega}_k] \subseteq \mathbb{R}$ but allow now the latter to take on negative values. An agent's type continues to parametrize both the agent's preferences and the utility she brings to agents from the other side. In particular, let $\sigma_k(\omega_k) \in \mathbb{R}_+$ denote the "salience" of each agent from side k with type ω_k. Such salience contributes positively to the utility of those agents from side $-k$ who like interacting with the side-k agents and negatively to those who dislike it. To make things simple, assume that the absolutely continuous distribution F_k (with density f_k) from which each ω_k is drawn is "regular", meaning that the function $\omega_k - [1 - F_k(\omega_k)]/f_k(\omega_k)$ is nondecreasing.

Next, assume that the gross payoff $u_k(\mathbf{s}|\omega_k)$ that an agent from side k obtains by interacting with a set of types $\mathbf{s} \in \mathcal{C}(\Omega_{-k})$ from the opposite side takes the form

$$u_k(\mathbf{s}|\omega_k) = \omega_k g_k \left(|\mathbf{s}|_{-k} \right), \tag{19}$$

where $g_k(\cdot)$ is a positive, strictly increasing, and continuously differentiable function satisfying $g_k(0) = 0$, and where

$$|\mathbf{s}|_{-k} := \int_{\omega_{-k} \in \mathbf{s}} \sigma_{-k}(\omega_{-k}) dF_{-k}(\omega_{-k}) \tag{20}$$

is the aggregate *salience* of the set $\mathbf{s}$.

An agent from side k with a negative ω_k is thus one who dislikes interacting with agents from the opposite side. To avoid trivial cases, assume that $\bar{\omega}_k > 0$ for some $k \in \{a, b\}$. The functions $g_k(\cdot)$, $k = a, b$, capture increasing (alternatively, decreasing) marginal utility (alternatively, disutility) for matching intensity.

Finally, assume that all costs are equal to zero, so as to simplify the analysis. Following standard arguments from mechanism design, it is easy to verify that

[5] Indeed, it can be easily shown that the converse to Proposition 14 is true; namely, that positive assortative matching is optimal only if the total virtual match surplus is supermodular.

a mechanism M is individually rational and incentive compatible *if and only if* the following conditions jointly hold for each side $k = a, b$:[6]

(i) the matching intensity of the set $\mathbf{s}_k(\omega_k)$ is nondecreasing;

(ii) the payoff $U_k(\underline{\omega}_k; M)$ of those agents with the lowest type is non-negative;

(iii) the pricing rule satisfies the envelope formula

$$\mathbf{p}_k(\omega_k) = \omega_k g_k\left(|\mathbf{s}_k(\omega_k)|_l\right) - \int_{\underline{\omega}_k}^{\omega_k} g_k\left(|\mathbf{s}_k(x)|_{-k}\right) dx - U_k(\underline{\omega}_k; M). \tag{21}$$

It is also easy to see that, in any mechanism that maximizes the platform's profit, the IR constraints of those agents with the lowest types bind, that is, $U_k(\underline{v}_k; M) = 0$, $k = a, b$. As already shown above, the problem of maximizing the platform's profit is then analogous to that of maximizing welfare in a fictitious environment in which the agents' types are equal to their virtual analogs. To economize on notation, for any $k = a, b$, any $\omega_k \in \Omega_k$, let $\varphi_k^W(\omega_k) := \omega_k$ and $\varphi_k^P(\omega_k) := \omega_k - [1 - F_k(\omega_k)]/f_k(\omega_k)$. The platform's problem thus consists in finding a pair of matching rules $\{\mathbf{s}_k(\cdot)\}_{k=a,b}$ that maximize

$$\sum_{k=a,a} \int_{\Omega_k} \varphi_k^h(\omega_k) g_k\left(|\mathbf{s}_k(\omega_k)|_l\right) dF_k(\omega_k) \tag{22}$$

among all rules that satisfy the monotonicity constraint (i) and the reciprocity condition (1). Hereafter, we will say that a matching rule $\{\mathbf{s}_k^h(\cdot)\}_{k=a,b}$ is h-optimal if it solves the above h-problem, with the understanding that, when $h = W$, this means that the rule is efficient, that is, welfare-maximizing, whereas when $h = P$, the rule is profit-maximizing.

For future reference, for both $h = W, P$, we also define the *reservation value* $r_k^h := \inf\{\omega_k \in \Omega_k : \varphi_k^h(\omega_k) \geq 0\}$ when $\{\omega_k \in \Omega_k : \varphi_k^h(\omega_k) \geq 0\} \neq \emptyset$.

5.2.1 Threshold Rules

Definition 15. A matching rule is a *threshold rule* if there exits a pair of weakly decreasing functions $t_k : \Omega_k \to \Omega_{-k} \cup \{\emptyset\}$ along with threshold types $\hat{\omega}_k \in \Omega_k$ such that, for any $\omega_k \in \Omega_k$, $k = a, b$,

$$\mathbf{s}_k(w_k) = \begin{cases} [t_k(\omega_k), \overline{\omega}_{-k}] & \text{if } \omega_k \geq \hat{\omega}_k \\ \emptyset & \text{otherwise,} \end{cases}$$

and, for any $\omega_k \in [\hat{\omega}_k, \overline{\omega}_k]$,

$$t_k(\omega_k) = \min\{\omega_{-k} : t_{-k}(\omega_{-k}) \leq \omega_k\}. \tag{23}$$

The interpretation is that any type below $\hat{\omega}_k$ is excluded, while a type $\omega_k > \hat{\omega}_k$ is matched to any agent from the other side whose type is above the threshold $t_k(\omega_k)$. To satisfy the reciprocity condition (1), the threshold functions $\{t_k(\cdot)\}_{k=a,b}$ have to satisfy the property in (23).

[6] See also Example 1.

Note that threshold rules are always implementable because matching intensity is nondecreasing under such rules. However, many other implementable matching rules do not have a threshold structure.

Proposition 16. *Assume that one of the following two sets of conditions holds:*

(a) the functions $g_k(\cdot)$ are weakly concave, and the functions $\sigma_k(\cdot)$ are weakly increasing, for both $k = a, b$;

(b) the functions $g_k(\cdot)$ are weakly convex, and the functions $\sigma_k(\cdot)$ are weakly decreasing, for both $k = a, b$.

Then both the profit-maximizing and the welfare-maximizing matching rules are threshold rules.

Proof sketch. Consider an agent for whom $\varphi_k^h(\omega_k) \geq 0$. Ignoring the monotonicity constraints, it is easy to see that it is always optimal to assign to this agent a matching set that includes all agents from the other side whose φ_l^h-value is non-negative. This is because (i) these latter agents contribute positively to type ω_k's payoff and, (ii) these latter agents have non-negative φ_l^h-values, which implies that adding type ω_k to these latter agents' matching sets (as required by reciprocity) never reduces the platform's payoff.

Next, consider an agent for whom $\varphi_k^h(\omega_k) < 0$. It is also easy to see that it is never optimal to assign to this agent a matching set that contains agents from the opposite side whose φ_{-k}^h-values are also negative. This is because matching two agents with negative h-valuations decreases the platform's payoff.

Now suppose that $g_k(\cdot)$ is weakly concave and $\sigma_k(\cdot)$ is weakly increasing, on both sides. Pick an agent from side k with $\varphi_k^h(\omega_k) > 0$ and suppose that the platform wants to assign to this agent a matching set whose intensity

$$q = |\mathbf{s}|_{-k} > \int_{[r_{-k}^h, \bar{\omega}_{-k}]} \sigma_{-k}(\omega_{-k}) dF_{-k}(\omega_{-k})$$

exceeds the aggregate matching intensity of those agents from side $-k$ with non-negative φ_l^h-values (that is, for whom $\omega_l \geq r_l^h$). That $g_k(\cdot)$ is weakly concave and $\sigma_k(\cdot)$ is weakly increasing, along with the fact that types are private information, implies that the *least costly* way to deliver such matching intensity is to match the agent to all agents from the opposite side whose $\varphi_{-k}^h(\omega_{-k})$ is the least negative. This is because (a) these latter agents are the most attractive ones, and (b) by virtue of g_{-k} being concave, using the same agents from side $-k$ with a negative φ_{-k}^h-valuations *intensively* is less costly than using different agents with negative φ_l^h-valuations. Threshold rules thus *minimize the costs of cross-subsidization* by delivering to those agents who play the role of consumers (that is, whose φ_k^h-valuation is nonnegative) matching sets of high quality in the most economical way.

Next, suppose that $g_k(\cdot)$ are weakly convex, and $\sigma_k(\cdot)$ are weakly decreasing, on both sides. The combination of the above properties with the fact that types are private information, implies that the most profitable way of using any type ω_k for whom $\varphi_k^h(\omega_k) < 0$ is to match him to those agents from side $-k$ with the highest positive φ_{-k}^h-valuations. This is because (a) these latter types are the ones that benefit the most from interacting with type[7] ω_k and (b) these latter types are the least salient ones and hence exert the lowest negative externalities on type ω_k. A threshold structure is thus optimal in this case as well. □

The matching allocations induced by threshold rules are consistent with the practice followed by many media platforms (e.g., newspapers) of exposing all readers to premium ads (displayed in all versions of the newspaper), but only those readers with high tolerance to advertising to discount ads (displayed only in the tabloid or printed version).

5.2.2 Distortions

We conclude by discussing the distortions in the provision of matching services due to market power.

Proposition 17. *Assume that the conditions for the optimality of threshold rules in Proposition 16 hold and that, in addition, the functions* $\psi_k^h : \Omega_k \to \mathbb{R}$ *defined by*

$$\psi_k^h(\omega_k) := \frac{\varphi_k^h(\omega_k)}{g'_{-k}\left(|[\omega_k, \bar{\omega}_k]|_k\right) \cdot \sigma_k(\omega_k)}$$

are strictly increasing, $k = a, b$, $h = W, P$. *Then, relative to the welfare-maximizing rule, the profit-maximizing rule (a) completely excludes a larger group of agents, that is,* $\hat{\omega}_k^P \geq \hat{\omega}_k^W$, $k = a, b$, *and (b) matches each agent who is not excluded to a subset of his efficient matching set, that is,* $\mathbf{s}_k^P(\omega_k) \subseteq \mathbf{s}_k^W(\omega_k)$, *for all* $\omega_k \geq \hat{\omega}_k^P$, $k = a, b$.

Proof sketch. Let $\hat{g}_k : \Omega_{-k} \to \mathbb{R}_+$ be the function defined by

$$\hat{g}_k(\omega_{-k}) := g_k\left(|[\omega_{-k}, \bar{\omega}_{-k}]|_{-k}\right) = g_k\left(\int_{\omega_{-k}}^{\overline{\omega}_{-k}} \sigma_{-k}(x) dF_{-k}(x)\right),$$

$k = a, b$. The utility that an agent with type ω_k obtains under a threshold rule from the matching set $[t_k(\omega_k), \overline{\omega}_{-k}]$ is then equal to $\omega_k \hat{g}_k\left(t_k(\omega_k)\right)$. Then let $\triangle_k^h : \Omega_k \times \Omega_{-k} \to \mathbb{R}$ be the function defined by

$$\triangle_k^h(\omega_k, \omega_{-k}) := -\hat{g}'_k(\omega_{-k})\varphi_k^h(\omega_k) f_k(\omega_k) - \hat{g}'_{-l}(\omega_k)\varphi_{-k}^h(\omega_{-k}) f_{-k}(\omega_{-k}), \quad (24)$$

for $k = a, b$. Note that $\triangle_a^h(\omega_a, \omega_b) = \triangle_b^h(\omega_b, \omega_a)$ represents the marginal effect on the platform's h-objective of decreasing the threshold $t_k^h(\omega_k)$ below ω_{-k} while also reducing the threshold $t_{-k}^h(\omega_{-k})$ below ω_k by reciprocity. One can then show that, under the conditions in the proposition, when $\triangle_k^h(\underline{\omega}_k, \underline{\omega}_{-k}) \geq 0$, the h-optimal matching rule is such that $\mathbf{s}_k^h(\omega_k) = \Omega_l$ for all $\omega_k \in \Omega_k$, $k = a, b$.[8] In this case, the platform induces all agents on board from either side of the market.

[7] Indeed such types have matching sets with the highest matching intensity – as required by incentive compatibility – and, because of the convexity of $g_{-k}(\cdot)$, the highest marginal utility for meeting additional agents.

[8] Equivalently, $t_k^h(\omega_k) = \underline{\omega}_{-k}$, for all $\omega_k \in \Omega_k$, $k = a, b$.

When, instead, $\triangle_k^h(\underline{\omega}_k, \underline{\omega}_{-k}) < 0$, the h-optimal matching rule has the following structure: (a) If $\triangle_k^h(\bar{\omega}_k, \underline{\omega}_{-k}) > 0$, the optimal rule induces bunching at the top on side k and no exclusion at the bottom on side $-k$ (that is, $\hat{\omega}_{-k} = \underline{\omega}_{-k}$ with $t_{-k}^h(\underline{\omega}_{-k})$ given by the unique solution to $\triangle_k^h(t_{-k}^h(\underline{\omega}_{-k}), \underline{\omega}_{-k}) = 0$); (b) If $\triangle_k^h(\bar{\omega}_k, \underline{\omega}_{-k}) < 0$, the optimal policy induces exclusion at the bottom on side $-k$ and no bunching at the top on of side k ($\hat{\omega}_{-k} = t_k^h(\bar{\omega}_k)$ with $t_k^h(\bar{\omega}_k)$ given by the unique solution to $\triangle_k^h(\bar{\omega}_k, t_k^h(\bar{\omega}_k)) = 0$. Intuitively, wherever possible, the platform balances the marginal gains of expanding the matching set of each agent with the corresponding marginal cost, taking into account the constraint imposed by reciprocity and the optimality of using a threshold structure.[9] The results in the proposition then follow from the above properties along with the fact that $\varphi_k^P(\omega_k) \leq \varphi_k^W(\omega_k)$ for all $\omega_k \in \Omega_k$, $k = a, b$. □

The role of the extra condition in the proposition is to guarantee that, under the h-optimal rule, bunching occurs only at the bottom of the distribution where it takes the form of exclusion ($\mathbf{s}_k^h(\omega_k) = \emptyset$ for all $\omega_k < \hat{\omega}_k^h$) or at the very top of the distribution where agents are matched to all agents on board from the other side of the market (that is, $\mathbf{s}_k^h(\omega_k) = [\hat{\omega}_{-k}^h, \bar{\omega}_k]$ for all $\omega_k > t_{-k}(\hat{\omega}_{-k}^h)$).[10] To interpret the condition, take the case of profit-maximization, $h = P$. The numerator in $\psi_k^h(\omega_k)$ is the agent's "virtual type". This term captures the effect on the platform's profit of expanding the intensity of the matching set of each side-k individual with type ω_k, taking into account that the expansion requires increasing the rents of all side-k agents with higher types. In other words, it captures the marginal value of a type-ω_k agent as a *consumer*. The denominator, instead, captures the effect on the platform's profit of adding such an agent to the matching set of any agent from the opposite side whose matching set is $[\omega_k, \bar{\omega}_k]$ (that is, whose threshold $t_{-k}^h(\omega_{-k}) = \omega_k$). In other words, it captures the marginal value of a type-ω_k agent as an *input*, under a threshold rule. The condition then requires that the contribution of an agent as a consumer increases faster than his contribution as an input.

The intuition for the result in the proposition is the following. Under profit-maximization, the platform only internalizes the effects of cross-subsidization on marginal revenues (which are proportional to the virtual valuations), rather than their effects on welfare (which are proportional to the true valuations). Contrary to other mechanism design problems, inefficiencies do not necessarily vanish as agents' types approach the "top" of the distribution (that is, the highest valuation for matching intensity). The reason is that, although virtual

[9]The formal proof is tedious and requires adapting some calculus-of-variation results to the non-standard reciprocity constrained given by (23). However, the intuition is fairly straightforward and is well captured by the property of the marginal net benefit function Δ_k^h above.

[10]Note that the condition is equivalent to the property that the marginal benefit function $\triangle_k^h$ satisfies the following single-crossing property: whenever $\triangle_k^h(\omega_k, \omega_{-k}) \geq 0$, then $\triangle_k^h(\omega_k, \omega'_{-k}) > 0$ for all $\omega'_{-k} > \omega_{-k}$ and $\triangle_k^h(\omega'_k, \omega_{-k}) > 0$ for all $\omega'_k > \omega_k$. This single-crossing property guarantees that, whenever the Euler condition $\triangle_k^h(\omega_k, t_k^h(\omega_k)) = 0$ admits an interior solution $\underline{\omega}_{-k} < t_k^h(\omega_k) < \bar{\omega}_{-k}$, the threshold $t_k^h(\omega_k)$ is strictly decreasing. The condition in the proposition is the "weakest" regularity condition that rules out non-monotonicities (or bunching) in the matching rule.

valuations converge to the true valuations as agents' types approach the top of the distribution, the cost of cross-subsidizing these types remains strictly higher under profit maximization than under welfare maximization, due to the infra-marginal losses implied by reciprocity on the opposite side.

6 Conclusions

The organization of the chapter reflects, to a large extent, the evolution of many platform markets from a format where all agents on board are matched to all participating agents from the opposite side of the market (the case considered in most of the earlier literature) to one where platforms engage in more sophisticated design, matching participating agents in a customized manner. Strong market power in such markets may result in the complete exclusion of many agents from either side of the market, and/or to inefficient matching whereby those agents on board are either matched to the wrong partners (in the case of one-to-one matching) or to a subset of their efficient matching set (in the case of many-to-many matching). Such distortions call for regulation and government interventions, an area that is receiving growing attention in recent years.

The analysis in this chapter has confined attention to markets dominated by a single platform. An important part of the literature studies competition between platforms in multi-sided markets, both in the case in which agents can multi-home (that is, join multiple platforms) and in the case where they single-home (that is, join at most one platform). Another area that has started receiving attention recently is dynamics, whereby agents strategically time their joining of the platform, experience shocks to their preferences over time, and learn the attractiveness of potential partners by interacting with them. The literature is also now considering richer specifications of the agents' preferences by allowing for different combination of vertical and horizontal differentiation that permit one to investigate the effects of targeting, a form of third-degree price discrimination often encountered in mediated matching markets. Finally, recent work studies matching markets where agents are asked to submit match-specific bids and where the selection of the matches is done through auctions that use match-specific scores to control for the agents' prominence in the market, a practice employed by many ad-exchanges and sponsored-search engines.

7 Bibliographical Notes

Section 3 is based on Rochet and Tirole (2003), Armstrong (2006), Rochet and Tirole (2006), Weyl (2010), and Tan and Wright (2018). In particular, we considered three scenarios regarding the agents' preferences: Scenario 1 is the one studied by Armstrong (2006), Scenario 2 the one in Rochet and Tirole (2003), and Scenario 3 the one in Rochet and Tirole (2006) and Weyl (2010). Section 4 is based on Jullien and Pavan (2019). Section 5 is based on Damiano and Li (2007), Jonhson (2013), Galichon (2016), and Gomes and Pavan (2016).

The literature on competing platforms mentioned in Section 6 is very broad. See Caillaud and Jullien (2001, 2003), Armstrong (2006), Rochet and Tirole (2006), and Armstrong and Wright (2007) for earlier contributions and Tan and Zhou (2020) for recent developments. See also Belleflamme and Peitz (2021) and Jullien et al. (2021) for an overview of this literature.

The literature on platform pricing in dynamic settings mentioned in Section 6 includes Cabral (2011), Halaburda et al. (2020), and Biglaiser and Crémer (2020). The case where agents learn about the attractiveness of potential partners over time and submit bids for specific interactions mentioned in Section 6 is examined in Fershtman and Pavan (2017, 2020). The case of platforms practicing a combination of second- and third-degree price discrimination mentioned in Section 6 is examined in Belleflamme and Peitz (2020) and Gomes and Pavan (2020).

References

[1] Armstrong, M., 2006, "Competition in two-sided markets." *Rand Journal of Economics*, Vol. 37, No. 3, pp. 668-691.

[2] Armstrong, M., and J. Wright, 2007, "Two-sided Markets, Competitive Bottlenecks and Exclusive Contracts." *Economic Theory*, Vol. 32(2), pp. 353-380.

[3] Belleflamme, P., and M. Peitz, 2020, "Network goods, Price Discrimination, and Two-sided Platforms." Mimeo, University of Bonn and University of Mannheim.

[4] Belleflamme, P., and M. Peitz, 2021, "The Economics of Platforms: Concepts and Strategy." Cambridge University Press, forthcoming.

[5] Biglaiser, G., and J. Crémer, 2020, "The Value of Incumbency in Heterogeneous Platforms." *American Economic Journal: Micro*, 12(4), pp. 229-69.

[6] Caillaud, B., and B. Jullien, 2001, "Competing Cybermediaries." *European Economic Review*, Vol. 45(4-6), pp. 797-808.

[7] Caillaud, B., and B. Jullien, 2003, "Chicken & Egg: Competition among Intermediation Service Providers", *Rand Journal of Economics*, Vol. 34(2), pp. 309-328.

[8] Damiano, E. and H. Li, 2007, "Price Discrimination and Efficient Matching." *Economic Theory*, Vol. 30, pp. 243-263.

[9] Fershtman D. and A. Pavan, 2017, "Pandora's Auctions: Dynamic Matching with Unknown Preferences,." *American Economic Review*, 107(5), pp. 186-90.

[10] Fershtman D. and A. Pavan, 2020, "Matching Auctions." mimeo, Northwestern University and Tel Aviv University.

[11] Galichon, A., 2016, "Optimal Transport Methods in Economics." Princeton University Press.

[12] Gomes, R., and A. Pavan, 2016, "Many-to-Many Matching and Price Discrimination." *Theoretical Economics*, Vol. 11, pp. 1005-1052.

[13] Gomes, R. and A. Pavan, 2020, "Price Customization and Targeting in Matching Markets." Mimeo Northwestern University and TSE.

[14] Halaburda, H., Jullien, B., and Y. Yehezkel, 2020, "Dynamic Competition with Network Effects: Why History Matters", *Rand Journal of Economics*, 51(1), pp. 3-31.

[15] Johnson, T., 2013, "Matching Through Position Auctions." *Journal of Economic Theory*, Vol. 148, pp. 1700-1713.

[16] Jullien, B. and A. Pavan, 2019, "Information Management and Pricing in Platform Markets." *Review of Economic Studies*, Vol. 86(4), pp. 1666-1703.

[17] Jullien, B., Pavan, A. and M. Rysman, 2021, "Two-sided Markets and Network Effects." Handbook of Industrial Organization, forthcoming.

[18] Rochet, J.C., and J. Tirole, 2003, "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association*, Vol. 33(4): 549-570.

[19] Rochet, J.C., and J. Tirole, 2006, "Two-sided markets: a progress report." *Rand Journal of Economics*, Vol. 37, No. 3, pp. 645-667.

[20] Tan, H., and J. Wright, 2018, "A price theory of multi-sided platforms: Comment." *American Economic Review*, 2018, Vol. 108(9): 2761-62.

[21] Tan, G. and J. Zhou, 2021, "Price Competition in Multi-Sided Markets." *Review of Economic Studies,* forthcoming.

[22] Weyl, G., 2010, "A Price Theory of Two-Sided Markets." *American Economic Review*, Vol. 100(4), pp. 1642–72.