

Price Customization and Targeting in Many-to-Many Matching Markets*

Renato Gomes

Toulouse School of Economics
renato.gomes@tse-fr.eu

Alessandro Pavan

Northwestern University and CEPR
alepavan@northwestern.edu

January 2018

Abstract

Recent technologies permit matching intermediaries to engage in unprecedented levels of *targeting*. Yet, regulators fear that the welfare gains of such targeting be hindered by the high degree of *price customization* practiced by matching intermediaries, whereby prices finely depend on the characteristics of the matching partners. To shed light on this debate, we develop a matching model in which agents' preferences are both vertically and horizontally differentiated. Mirroring current practices, we show how platforms maximize profits by offering menus of matching plans defined by (a) a baseline configuration, (b) a baseline price, and (c) a collection of nonlinear tariffs for customization. We illustrate how, under such plans, prices are linked to structural elasticities, and derive primitive conditions under which market power distortions increase with the targeting level of a match. We then study the effects on targeting and consumer welfare of *uniform-pricing regulation* mandating that the price charged to the side- i agents be invariant in that side's observable characteristics (e.g., the requirement that the price charged to advertisers be invariant in the ads' content). Finally, we examine the transition of matching markets from a centralized structure to a decentralized one where sellers post prices and matching is unmediated. The analysis has implications for ad-exchanges, media platforms, cable TV, business-to-business platforms, and large online retailers.

JEL classification: D82

Keywords: Many-to-many matching, networks, asymmetric information, platforms, incentives, price discrimination

*For useful comments and suggestions, we thank Igal Hendel, Stephan Laueremann, Jean Tirole and seminar participants at the 2017 Barcelona Conference on the Economics of Platforms and various workshops where the paper was presented. Pavan also thanks the NSF (Research Grant SES 1530798) for financial support, and Bocconi University for its hospitality during the 2017-2018 academic year.

1 Introduction

Over the last two decades, new technologies have permitted the development of matching intermediaries of unprecedented scale engaging in unparalleled level of targeting. Notable examples include (a) ad exchanges, matching publishers with advertisers, (b) business-to-business platforms, matching firms with mutually beneficial commercial interests, and (c) dating websites, matching agents with common passions. The same advances in technology that favored high levels of targeting also enabled greater price customization, whereby the price of a match finely depend on observable characteristics of the matching partners. For instance, in advertising exchanges, the reservation prices demanded by publishers are allowed to vary with the identity of the advertisers.¹ Furthermore, the bids (and payments) by the advertisers depend on scores that summarize how compatible they are with each publisher’s content. A similar trend can be found in other markets, not traditionally analyzed through the lens of matching. In media markets, for instance, satellite TV providers use sophisticated pricing strategies that condition subscribers’ payments on the entire bundle of channels selected by the subscribers. In turn, because channels’ revenues often depend on the demographics of the subscribers (as this determines the value of advertising), the prices negotiated by the TV provider with the channels often depend in the profile of users subscribing the different packages in which a channel is included.

In some cases, such pricing practices are easy to enforce, as “horizontal” characteristics (i.e., those that determine targeting accuracy) are observable (for instance, the profile of an advertiser is typically observable). In other cases, instead, the characteristics relevant for targeting have to be indirectly elicited, and this may require bundling.²

While having a long history in the policy debate,³ price-customizing practices have attracted renewed attention due to the two-sided nature of matching intermediaries, and the amount of information now available for pricing.⁴ The concern is that, by leveraging the platforms’ market power, price customization hinders the efficiency gains permitted by better targeting technologies. The aim of this paper is to understand how targeting and customized pricing shape the matching plans offered by platforms and study the impact on targeting and consumer welfare of uniform-pricing regulation (whereby the payments charged by the platform to each agent cannot depend on the agent’s own profile).

To examine these issues, we develop a model where agents’ preferences exhibit elements of both vertical and horizontal differentiation. Certain agents value being matched with agents from the

¹See for example https://support.google.com/adxseller/answer/2913506?hl=en&ref_topic=3376095. Moreover, advertiser-specific reserve prices can be easily automated using proxy-bidding tools.

²Ad exchanges have recently developed new contractual arrangements allowing for bundling of advertising (see, for example, Mirrokni and Nazerzadeh (2017)).

³In the US, the first law to regulate price discrimination is the Robinson-Patman Act of 1936.

⁴In the case of media markets, the Federal Communications Commission (FCC) has published two reports analyzing the potential harm of price customization through bundling (FCC 2004, 2006). In the case of online retailing, the UK Office of Fair Trading issued in 2010 an eponymous report on online targeting of advertising and prices.

other side of the market uniformly more than others (vertical differentiation). At the same time, agents from the same side may disagree on the relative attractiveness of any two agents from the opposite side (horizontal differentiation). We capture the two dimensions of differentiation by letting agents' types be located on a cylinder, where the height represents the vertical dimension, while the radial position determines the horizontal preference. Each agent's utility from interacting with any other agent from the opposite side increases uniformly with the agent's vertical dimension. Fixing the vertical dimension, each agent's utility is single-peaked with respect to the horizontal dimension. More specifically, we identify each agent's radial position with his bliss point. Accordingly, each agent's utility for interacting with any other agent from the opposite side decreases with the distance between the agent's bliss point (his radial position) and the partner's location (the partner's radial position). Such preference structure, in addition to its analytical convenience, mirrors the one in the "ideal-point" models typically used in the empirical literature on media and advertising markets (see, for example, Goettler and Shachar 2001).

A key element of our analysis is the focus on *matching tariffs*, which describe how the payments asked by the platform vary with the matching set demanded by an agent. A tariff exhibits *uniform pricing* if all agents from a given side face the same price schedule for different quantities of matches from a given location of the opposite side. Formally, uniform tariffs are tariffs that do not condition an agent's payment to the platform on the agent's own radial position (i.e, the horizontal dimension of the agent's preferences). A particularly simple type of uniform pricing often proposed as a potential regulatory remedy to the market power enjoyed by media platforms is stand-alone linear pricing (for a discussion, see Crawford and Yurukoglu 2012).

Absent any regulation, platforms typically use *price-customizing* tariffs, whereby agents from each given side of the market are offered *matching plans*. Each plan is defined by its baseline configuration (i.e., a baseline set of partners from the opposite side), a baseline price, and a collection of prices describing the cost to the subscriber of customizing the plan by adding extra matches. The cost of the customization is typically non-linear in the volume of matches of any given type added to the plan (second-degree price discrimination). Importantly, the cost of the customization is also a function of the baseline plan selected by the subscriber. Because different plans are targeted to agents from different "locations," price-customizing tariffs thus also display a form of third-degree price discrimination. As the analysis reveals, when neither the vertical nor the horizontal dimensions defining the agents' preferences are observable by the platform, information about such dimensions is elicited through a careful design of the baseline plans and the associated price-customizing schedules.

Our first main result derives conditions under which the profit-maximizing tariffs exhibit price customization. Equipped with these conditions, we then offer a convenient representation of the optimal price schedules. The representation yields a formula describing the price each agent from each given location has to pay to include in his matching set any feasible amount of matches from any location on the opposite side of the market. The formula links location- and volume-specific prices to the various local elasticities of the demands on the two sides of the market. In this sense it constitutes

the analog in a matching market of the familiar Lerner-Wilson formula of optimal monopolistic pricing. The formula differs from the traditional one in that it accounts for (a) the reciprocity of the matching, and (b) the fact that the platform combines second-degree price discrimination (higher vertical types self-select into larger matching plans) with third-degree price discrimination (the total price paid by each subscriber depends on the horizontal dimension of the subscriber’s preferences, both when the latter are observed by the platform and when they are indirectly elicited through self-selection).

Our characterization of the profit-maximizing matching tariffs reveals interesting patterns of cross-subsidization, unique to matching environments. Namely, optimal tariffs induce a form of *negative assortative matching at the margin*: at any given location, agents with a low value for matching (a lower vertical dimension) are matched only to those agents from the other side whose value for matching is sufficiently high. This form of negative assortativeness naturally takes into account the agents’ mutual attractiveness, as determined by their joint locations. As a result, the matching sets of any two agents from the same side are nested only if the two agents share the same dimension of horizontal preferences.

We also leverage on this characterization to study the interplay between targeting and market power. Specifically, we derive conditions on primitives (on the agents’ utilities and distributions of their preferences) under which the under-provision of matches under profit maximization (relative to the efficient level) is either magnified or alleviated as the distance of a partner’s location from an agent’s bliss point increases. This analysis has no parallel in the screening literature, and can be brought to data using structural techniques along the lines, for instance, of those in Kang and You (2016).

Our second set of results investigates the effects on prices, the composition of the matching demands, and consumer welfare of regulation that imposes uniform pricing on a given side of the market (like the one proposed in recent years for consumers in media markets, or advertisers in digital platforms). Analogously to the generalized Lerner-Wilson formula discussed above, we provide a novel representation of the optimal price schedules that uses local elasticities to describe the prices agents on each side have to pay per quantity of matches from each location on the opposite side. Relative to the case of customized pricing, this new pricing formula identifies the relevant aggregate elasticities in environments where location-specific pricing is not possible. The typical marginal revenue and marginal cost terms (which determine the optimal cross-subsidization pattern) are now averages that take into account not only the uniform-pricing aspect of regulation, but also how the procurement costs of matches are affected by the horizontal component of preferences. From a more theoretical perspective, the characterization contributes to the mechanism design literature by developing a novel technique to handle constraints on the transfer rule employed by the principal (as opposed to constraints on the allocation rules, which are typically easier to analyze using standard techniques).

We then put this characterization to work, revealing how uniform-pricing regulation affects tar-

getting and welfare. Intuition might suggest that uniform pricing should increase targeting (by preventing platform from charging higher prices for the matches involving the most preferred partners). This simple intuition, however, may fail to account for the fact that platforms re-optimize their entire price schedules to respond to aggregate elasticities. Perhaps surprisingly, uniform-pricing regulation can either decrease or increase the equilibrium level of targeting, depending on how match-demand elasticities vary with location. We derive sufficient conditions in terms of primitives (on the modularity and convexity of the agents’ utilities, as well as the distributions of vertical preferences) under which targeting is higher (alternatively, lower) under uniform pricing (alternatively, customized pricing). We then use such conditions to look into the welfare effects of uniform-pricing regulation. Exploiting a novel connection between uniform pricing in matching markets and the literature on third-degree price discrimination, we show how to adapt the elegant analysis in Aguirre et al. (2010) to the matching markets under examination, and identify sufficient conditions for uniform-pricing regulation to increase consumer surplus in the side where it is mandated. These results, once combined with appropriate empirical work, can guide the design of regulatory interventions in platform markets where price customization is a concern.

Lastly, we show how the above results can also be used to study the transition from a centralized to a decentralized market. The same technological progress of the last few years that has facilitated the growth of matching markets is now expected to favor a gradual transition of such markets from a centralized structure where matching is controlled by platforms to a more decentralized structure where one side (typically, the “seller” side) posts stand-alone prices, following which agents from the other side (typically, the “buyer” side) then construct their matching sets. For example, in the market for media content, several analysts believe the increase in the speed of fiber-optic and broadband internet connection will favor a gradual transition of the market to a structure whereby viewers will pay directly the content producers, bypassing the intermediation of current TV providers. It remains unclear whether such developments will boost viewers’ surplus. We show that the same analysis that permits us to uncover the welfare effects of uniform-pricing regulation can be adapted to shed light on the welfare effects of a transition to such a decentralized structure, thus contributing a novel angle to the policy debate over whether such transition should be encouraged or slowed down.

Outline of the Paper. The rest of the paper is organized as follows. Below, we close the introduction by briefly reviewing the pertinent literature. Section 2 presents the model. Section 3 identifies conditions under which the profit-maximizing matching tariffs are price-customizing, and derives properties of the associated pricing schedules. Section 4 studies the effects of uniform-pricing regulation and of the transition from a centralized to a decentralized structure. Section 5 concludes. All proofs are in the Appendix at the end of the document.

Related Literature

This paper studies many-to-many matching (with monetary transfers) in markets where the agents' preferences are both vertically and horizontally differentiated. Particularly related are Jeon, Kim and Menicucci (2017) and Gomes and Pavan (2016). The first paper studies the provision of quality by a platform in a setting where quality provision enhances match values. The second paper studies the inefficiency of profit maximization in many-to-many matching markets. Both papers abstract from the possibility that agents' preferences be horizontally differentiated (in tastes and match values), thus ignoring the issues of targeting and price customization that are the heart of the analysis in the present paper. Fershtman and Pavan (2017) considers many-to-many matching in a model with a rich preference structure similar to the one in the present paper, combining elements of vertical and horizontal differentiation. Contrary to the present paper, however, it focuses on dynamic markets in which agents learn the attractiveness of their partners and experience shocks to their preferences over time. The structure of the matching sets as well as the focus of the analysis (learning and experimentation) are different from the one in the present paper.

Related are also Damiano and Li (2007), Johnson (2013), Jullien and Pavan (2017), and Tan and Zhou (2017). The first two papers study price discrimination in markets where matching is one-to-one and agents' preferences are differentiated only along a vertical dimension. The third paper studies platform competition in markets where agents' preferences for the different platforms are heterogenous but where agents value homogeneously the interactions with agents from the opposite side of the market. The fourth paper studies price competition in a general model where multiple platforms compete by offering differentiated services to the various sides of the market and where agents' preferences are heterogenous with both within-side and across-sides network effects. Contrary to the first two papers, however, the last two papers abstract from price discrimination.

More broadly, markets where agents purchase access to other agents are the focus of the broad literature on two-sided markets (see Belleflamme and Peitz (2017) for the most up-to-date overview). This literature, however, restricts attention to a single network, or to mutually exclusive networks, and abstracts from horizontal differentiation. By allowing for more flexible matching rules, and by considering a richer preferences structure, the present paper contributes to this literature by studying targeting and price customization in such markets.

The study of price customization is related to the literature on price discrimination. In the case of second-degree price discrimination, Mussa and Rosen (1978), Maskin and Riley (1983), and Wilson (1997) study the provision of quality/quantity in markets where agents possess private information about a vertical dimension of their preference. Our study of price customization in many-to-many matching markets introduces two novel features relative to the standard monopolistic screening problem. First, the platform's feasibility constraint (namely, the reciprocity of the matching rule) has no equivalent in markets for commodities. Second, agents' preferences are differentiated along both a vertical and a horizontal dimension. This richer preferences structure calls for a combination

of second- and third-degree price discrimination and leads to cross-subsidization patterns novel to the literature.⁵

The paper also contributes to the literature on third-degree price discrimination. In addition to the paper by Aguirre et al. (2010) mentioned above, see Bergemann, Brooks, and Morris (2015) for an excellent overview and recent developments. That paper characterizes all combinations of producer and consumer surplus that arise from different information structures about the buyers’ willingness-to-pay (alternatively, from different market segmentations). Our setup differs from theirs in many respects. First, the informational costs incurred by the intermediary are endogenous, and depend on its price-customization strategies on both sides of the market (the information structure is exogenous in Bergemann, Brooks, and Morris (2015)). Second, the preferences structure is different, reflecting specific features of many-to-many matching environments.

Related is also the literature on bundling (see, among others, Armstrong (2013), Hart and Reny (2015), and the references therein). The present work differs from that literature in two important aspects. First, while preferences are multi-dimensional both in the present paper and in that literature, in our setting, preferences can be orthogonally decomposed into a vertical and a horizontal dimension. The bundling literature, instead, assumes a more general preference structure, which, however, hinders the characterization of the optimal price schedules, except in special cases with specific distributions and only two goods. Second, reflecting the practices of many-to-many matching intermediaries, we assume that sales are monitored, so that prices can condition on the entire matching set of each agent. The bundling literature, by contrast, typically assumes that purchases are anonymous.

Lastly, the paper contributes to the literature on targeting in advertising markets (see, for example, Bergemann and Bonatti (2011, 2015) and Cox et al. (2017) and the references therein). Our work contributes to this literature by introducing a richer class of (non-linear) pricing strategies and by comparing the matching outcomes that emerge under a decentralized structure to their counterparts in platform markets where the matching between the advertisers and the publishers (or content providers) is mediated. Contrary to some of the papers in this literature, however, we abstain from platform competition and focus on a market with a single platform.

2 Model

A monopolistic platform matches agents from two sides of a market. Each side $k \in \{a, b\}$ is populated by a unit-mass continuum of agents. Each agent from each side $k \in \{a, b\}$ has a bi-dimensional type $\theta_k = (v_k, x_k) \in \Theta_k \equiv V_k \times X_k$ which parametrizes both the agent’s preferences and the agent’s

⁵Related is also Balestrieri and Izmalkov (2015). That paper studies price discrimination in a market with horizontally differentiated preferences by an informed seller who possesses private information about its product’s quality (equivalently, about the “position” of its good in the horizontal spectrum of agents’ preferences). The focus of that paper is information disclosure, while the focus of the present paper is matching, targeting, and price customization.

attractiveness.

The parameter x_k , which describes the location of the agent, captures horizontal differentiation in preferences. For convenience, we assume that agents are located on a circle of perimeter one, in which case $X_k = [0, 1]$, $k = a, b$. The parameter $v_k \in V_k \equiv [\underline{v}_k, \bar{v}_k] \subseteq \mathbb{R} \cup \{+\infty\}$, on the other hand, captures heterogeneity in preferences along a vertical dimension. It controls for the intensity of an agent's matching utility across locations (i.e., the overall utility the agent derives from interacting with a generic agent from the opposite side, before doing any profiling). Hereafter, we let $Int[V_k]$ denote the interior of the set V_k .

We assume that the vertical dimensions v_k are the agents' private information. As for the horizontal dimensions, i.e., the locations x_k , we consider both cases where they are publicly observable, as well as cases where they are the agents' private information. In particular, we consider the following four scenarios:

- *Scenario (i)*: Locations are publicly observable on both sides;
- *Scenario (ii)*: Locations are private information on side a and publicly observable on side b ;
- *Scenario (iii)*: Locations are publicly observable on side a and private information on side b ;
- *Scenario (iv)*: locations are private information on both sides.

Agents derive higher utility from being matched to agents whose locations are closer to their own. Their utility also increases, over all locations, with their vertical type. We assume the utility that an agent from side $k \in \{a, b\}$ with type $\theta_k = (v_k, x_k)$ obtains from being matched to an agent from side $l \neq k$ with type $\theta_l = (v_l, x_l)$ is represented by the function

$$u_k(v_k, |x_k - x_l|),$$

where $|x_k - x_l|$ is the distance between the two agents' locations. The function u_k is Lipschitz continuous, bounded, strictly increasing in v_k , and weakly decreasing in $|x_k - x_l|$. To make things interesting, we assume u_k is strictly decreasing in $|x_k - x_l|$ on at least one side. The following examples illustrate the type of preferences covered by the aforementioned specification.

Example 1. (*ad exchange*) The platform is an ad exchange matching advertisers on side a to publishers on side b . The expected payoff that an advertiser with type $\theta_a = (v_a, x_a)$ obtains from an impression at the website of a publisher with type $\theta_b = (v_b, x_b)$ is given by

$$u_a(v_a, |x_a - x_b|) = v_a \cdot \phi(|x_a - x_b|),$$

where the strictly decreasing function $\phi : [0, .5] \rightarrow [0, 1]$ describes how the probability of a conversion (i.e., the probability the ad view turns into a sale) varies with the distance between the publisher's profile, x_b , and the advertiser's ideal audience, x_a . By contrast, publishers can be viewed (to a first approximation) as indifferent with respect to the kind of advertisement displayed at their websites.

The matching (dis)utility of a publisher reflects the opportunity cost of not using the advertisement space to sell its own products, or from not selling the ad slot outside of the platform; given the above notation, this is captured by letting $u_b(x_b, |x_a - x_b|) = v_b \leq 0$, all $x_a, x_b \in [0, 1]$.

In the context of Example 1, it seems plausible that both an advertiser's ideal type of audience and a publisher's profile be observable by the platform, which corresponds to Scenario (i) above.

Example 2. (*media platform*) The platform is a media outlet, e.g., a cable TV provider, matching viewers on side a with content providers on side b . The utility that a viewer with type $\theta_a = (v_a, x_a)$ derives from having access to a content provider with type $\theta_b = (v_b, x_b)$ is given by the constant-elasticity-of-substitution (CES) function

$$u_a(v_a, |x_a - x_b|) = \left[\alpha \cdot (v_a)^\delta + (1 - \alpha) \cdot \phi(|x_a - x_b|)^\delta \right]^{\frac{1}{\delta}},$$

where $\alpha \in [0, 1]$ captures the relative importance of a viewer's vertical and horizontal preferences, and $V_a \subset \mathbb{R}_+$. In turn, the strictly decreasing function $\phi : [0, .5] \rightarrow \mathbb{R}_+$ describes how the viewer's utility varies with the distance between the viewer's ideal type of content, x_a , and the provider's profile, x_b . Finally, $\delta > 0$ measures the elasticity of substitution between the vertical and horizontal dimensions. By contrast, content providers might be viewed as indifferent with respect to the profile of the viewers that access their content. The matching (dis)utility of a content provider reflects the extra revenue from advertisers (as advertisers typically pay more to content providers with a higher exposure to viewers), or the expenses from broadcasting rights paid to third parties (which are proportional to the audience reached). In this case, for any $v_b \in V_b$, any $x_a, x_b \in [0, 1]$, $u_b(v_b, |x_b - x_a|) = v_b$, with $v_b > 0$ under the first interpretation, and $v_b < 0$ under the second interpretation.

In Example 2, each viewer's ideal type of content is likely to be his own private information, whereas each content provider's profile is likely to be publicly observable, which corresponds to Scenario (ii) above.

The type $\theta_k = (v_k, x_k)$ of each agent from each side $k \in \{a, b\}$ is an independent draw from the absolutely continuous distribution F_k with support Θ_k . We denote by F_k^v (alternatively, F_k^x) the marginal distribution of F_k with respect to v_k (alternatively, x_k), and by $F_k^{v|x}$ the distribution of v_k conditional on x_k . We denote by f_k^v the density of F_k^v and by $\lambda_k^v \equiv f_k^v / [1 - F_k^v]$ its hazard rate. We use analogous notation for the densities and hazard rates of the conditional distributions $F_k^{v|x}$.

Let $\Sigma(\Theta_l)$ be the Borel sigma algebra associated with the set Θ_l . The total payoff that an agent from side $k \in \{a, b\}$ with type $\theta_k = (v_k, x_k)$ obtains from being matched, at a price p , to a set of types $\mathbf{s}_k \in \Sigma(\Theta_l)$ from side $l \neq k$ is given by⁶

$$\pi_k(\mathbf{s}_k, p; \theta_k) = \int_{\mathbf{s}_k} u_k(v_k, |x_k - x_l|) dF_l(\theta_l) - p, \quad (1)$$

⁶The representation in (1) assumes the agent is matched to all agents from side $l \neq k$ whose type is in \mathbf{s}_k , which is the case under both profit-maximizing and welfare-maximizing mechanisms.

whereas the payoff that the same agent obtains from not interacting with the platform (in which case the agent is matched with no agent from side $l \neq k$ and makes no payment to the platform) is equal to zero.

Tariffs and Matching Demands

The platform offers matching tariffs to each side $k \in \{a, b\}$. A *matching tariff* T_k specifies the (possibly negative) total payment $T_k(\mathbf{s}_k)$ that each agent from each side k must pay to the platform for being matched to the set of types $\mathbf{s}_k \in \Sigma(\Theta_l)$ from side $l \neq k$.

Given the tariff T_k , the *matching demand* of each agent from side k with type $\theta_k = (v_k, x_k)$ is given by the set

$$\hat{\mathbf{s}}_k(\theta_k; T_k) \in \arg \max_{\mathbf{s}_k \in \Sigma(\Theta_l)} \left\{ \int_{\Theta_l} u_k(v_k, |x_k - x_l|) dF_l(\theta_l) - T_k(\mathbf{s}_k) \right\}. \quad (2)$$

To guarantee participation by all agents, we require that $T_k(\mathbf{s}_k) = 0$ if $\mathbf{s}_k = \emptyset$.

Definition 1. The tariffs T_k , $k = a, b$, are *feasible* if, for all $(\theta_k, \theta_l) \in \Theta_k \times \Theta_l$, $k, l \in \{a, b\}$, $l \neq k$,

$$\theta_l \in \hat{\mathbf{s}}_k(\theta_k; T_k) \iff \theta_k \in \hat{\mathbf{s}}_l(\theta_l; T_l). \quad (3)$$

A pair of feasible matching tariffs thus induces *reciprocal* demands. That is, if an agent from side k with type θ_k finds it optimal to be matched to all agents from side $l \neq k$ with type θ_l , then all agents from side l with type θ_l find it optimal to be matched to all agents from side k with type θ_k .

Given any matching set $\mathbf{s}_k \in \Sigma(\Theta_l)$, any location x_l , we denote by $q_{x_l}(\mathbf{s}_k)$ the “mass” of agents from side l located at x_l included in the matching set \mathbf{s}_k .⁷

Definition 2. The tariff T_k offered by the platform to the side- k agents, $k \in \{a, b\}$, is consistent with *uniform pricing* if there exists a collection of *price schedules* $p_k : [0, 1]^2 \rightarrow \mathbb{R}$, one for each location $x_l \in [0, 1]$, such that the total payment asked by the platform for each matching set $\mathbf{s}_k \in \Sigma(\Theta_l)$ is given by

$$T_k(\mathbf{s}_k) = \int_0^1 p_k(q_{x_l}(\mathbf{s}_k)|x_l) dx_l. \quad (4)$$

Conversely, the tariff T_k is *discriminatory* if there are no price schedules such that T_k is consistent with Condition (4), for all \mathbf{s}_k .

Hence, under uniform pricing, the tariff offered by the platform to the side- k agents consists of a collection of non-linear price schedules, $(p_k(\cdot|x_l))_{x_l \in [0,1]}$, with each schedule $p_k(q|x_l)$ specifying the total price each side- k agent has to pay to be matched to q agents from side $l \neq k$ located at $x_l \in [0, 1]$. Importantly, the price $p_k(q|x_l)$ is independent of the mass of agents from other locations included in the matching set \mathbf{s}_k .

⁷Hereafter, we abuse of terminology by referring to the density of agents of a certain type as the “mass” of agents of that type.

The next definition describes a type of discriminatory tariffs that plays an important role in the analysis below.

Definition 3. The tariff T_k is *customized* if there exists a collection of matching plans

$$\{(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathbf{S}_k(x_k)) : x_k \in [0, 1]\},$$

one for each side- k location $x_k \in [0, 1]$, such that each side- k agent selecting the basic plan $\underline{\mathbf{s}}_k(x_k)$ and then selecting the customization $\mathbf{s}_k \in \mathbf{S}_k(x_k) \subseteq \Sigma(\Theta_l)$ is asked to make a total payment equal to⁸

$$\underline{T}_k(x_k) + \int_0^1 \rho_k(q_{x_l}(\mathbf{s}_k)|x_l; x_k) dx_l, \quad (5)$$

with $\rho_k(q_{x_l}(\underline{\mathbf{s}}_k(x_k))|x_l; x_k) = 0$ all $x_l \in [0, 1]$.

Under customized tariffs, the platforms thus offer to the side- k agent a menu of matching plans, one for each side- k location x_k . Each plan specifies (a) a baseline configuration, formally captured by the default set of types $\underline{\mathbf{s}}_k(x_k) \subseteq \Theta_l$ from side $l \neq k$ included in the package, (b) a baseline price $\underline{T}_k(x_k)$, (c) a collection of possible customizations $\mathbf{S}_k(x_k) \subseteq \Sigma(\Theta_l)$, and (d) a collection of non-linear schedules $\rho_k(q|x_l; x_k)$, one for each location $x_l \in [0, 1]$, that jointly define the cost of customizing the plan. As in the case of uniform pricing, each non-linear schedule $\rho_k(q|x_l; x_k)$ specifies the price charged to the side- k agents for being matched to q agents from side $l \neq k$ located at x_l . Contrary to the case of non-linear pricing, though, the price depends on the plan selected by the side- k agent, which is conveniently indexed by the side- k locations, x_k . A menu of customized tariffs thus combines elements of second-degree price discrimination (each price function $\rho_k(q|x_l; x_k)$ is possibly non-linear in q) with elements of third-degree price discrimination (each non-linear price function $\rho_k(q|x_l; x_k)$ depends on the plan, and hence the location, of the side- k agents). That $\rho_k(q_{x_l}(\underline{\mathbf{s}}_k(x_k))|x_l; x_k) = 0$, all $x_l \in [0, 1]$, in turn means that an agent making no changes to a baseline plan is asked to make no further payments to the platform beyond $\underline{T}_k(x_k)$. Customized tariffs capture important features of real-world matching plans offered by platforms such as cable TV providers, Ad Exchanges, and stores.

The platform's problem consists in choosing a pair of feasible matching tariffs T_k , $k = a, b$, that maximizes its profits, which are given by

$$\sum_{k=a,b} \int_{\Theta_k} T_k(\hat{\mathbf{s}}_k(\theta_k; T_k)) dF_k(\theta_k). \quad (6)$$

⁸The payment specified by the tariff for any matching set $\mathbf{s}_k \notin \{\cup \mathbf{S}_k(x_k) : x_k \in [0, 1]\}$ can be taken to be arbitrarily large to guarantee that no type finds it optimal to select any such set. The existence of such payments is guaranteed by the assumption that u_k is bounded, $k = a, b$. Furthermore, in case locations are private information on side k , the collection of matching plans is required to have the property that for any set $\mathbf{s}_k \in \mathbf{S}_k(x_k) \cap \mathbf{S}_k(x'_k)$, the total payment associated with \mathbf{s}_k is the same no matter whether the set is obtained by selecting the plan x_k or the plan x'_k . When, instead, locations are public, the collection of matching plans $\{(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathbf{S}_k(x_k)) : x_k \in [0, 1]\}$ may entail multiple prices for the same matching set \mathbf{s}_k . However, because each agent is constrained to choosing the plan designed for his location, de facto each agent faces a tariff specifying a single price for each set.

Hereafter, we denote by T_k^* the profit-maximizing tariffs and by \mathbf{s}_k^* the induced matching sets, $k = a, b$.

3 Customized Tariffs

We start by studying the platform’s problem when no restrictions are imposed on the matching tariffs it can offer. Before describing the results, we introduce three conditions that play an important role in the analysis below.

Condition 1. [R] Regularity: For any $k, l \in \{a, b\}$, $l \neq k$, $(\theta_k, \theta_l) \in \Theta_k \times \Theta_l$, the virtual values

$$\varphi_k(\theta_k, \theta_l) \equiv u_k(v_k, |x_k - x_l|) - \frac{1 - F_k^{v|x}(v_k|x_k)}{f_k^{v|x}(v_k|x_k)} \cdot \frac{\partial u_k}{\partial v}(v_k, |x_k - x_l|)$$

are continuous and non-decreasing in v_k .

Condition R imposes that “virtual values” $\varphi_k(\theta_k, \theta_l)$ be non-decreasing in the agents’ vertical parameters. This assumption is the natural analog of standard regularity conditions (e.g., Myerson (1981)) in matching environments.

Condition 2. [VD] Virtual values decreasing in distance: For any $k, l \in \{a, b\}$, $l \neq k$, $\theta_k \in \Theta_k$, $v_l \in V_l$, the virtual values $\varphi_k(\theta_k, (v_l, x_l))$ are non-increasing in $|x_k - x_l|$.

Condition VD is another monotonicity condition, similar to R, requiring virtual values to be non-increasing in the distance between locations. Note that, because true values $u_k(v_k, |x_k - x_l|)$ are decreasing in $|x_k - x_l|$, a sufficient, albeit not necessary, condition for VD is that the match value functions u_k are supermodular in $(v, |\cdot|)$.

Condition 3. [I_k] Independence on side k : for any $\theta_k = (v_k, x_k) \in \Theta_k$, $F_k(\theta_k) = F_k^x(x_k)F_k^v(v_k)$.

Condition I_k requires the vertical parameters v_k to be drawn independently from the locations x_k . In the cable TV application, this condition implies that knowing a viewer’s “bliss point”, i.e., his preferred type of broadcasting, carries no information about the overall importance the viewer assigns to cable TV.

Condition 4. [S_k] Symmetry on side k : for any $\theta_k = (v_k, x_k) \in \Theta_k$, $F_k(\theta_k) = x_k F_k^v(v_k)$.

Condition S_k strengthens the independence condition by further requiring that locations be uniformly distributed over $X_k = [0, 1]$, as typically assumed in models of horizontal differentiation.⁹ As shown below, this assumption disciplines the agents’ matching demands on side k , when locations are private information on side $l \neq k$.

We then have the following result:

⁹Similar assumptions are typically made also in the targeting literature; see, for example, Bergemann and Bonatti (2011, 2015), and Cox et al. (2017).

Proposition 1. (properties of the optimum) *In addition to Condition R, suppose the environment satisfies the properties of one of the following four cases: Scenario (i); Scenario (ii) along with Conditions VD, I_a and S_b ; Scenario (iii) along with Conditions VD, S_a and I_b ; Scenario (iv) along with Conditions VD, S_a and S_b . Then, under the profit-maximizing tariffs, for any $k \in \{a, b\}$:*

1. *the matching tariff T_k^* is customized;*
2. *the matching sets \mathbf{s}_k^* exhibit negative assortativeness at the margin: there exist functions $t_k^* : \Theta_k \times [0, 1] \rightarrow V_l$ such that*

$$\mathbf{s}_k^*(\theta_k) = \{(v_l, x_l) \in \Theta_l : v_l > t_k^*(\theta_k, x_l)\},$$

with the threshold function t_k^ non-increasing in v_k . When condition VD holds, fixing $\theta_k = (v_k, x_k)$, the threshold functions $t_k^*(\theta_k, x_l)$ are non-decreasing in $|x_k - x_l|$. Finally, when locations are public on side $k \in \{a, b\}$, without loss of optimality, the side- k customized tariffs do not need to restrict the set of possible customizations, i.e., for each $x_k \in [0, 1]$, $\mathbf{S}_k(x_k) = \Sigma(\Theta_l)$.*

The conditions in Proposition 1 guarantee that the platform can price discriminate along the agents' locations, without leaving the agents rents for the private information the agents may possess regarding their locations. That is, in Scenarios (ii)-(iv), these conditions guarantee that the platform achieves the same profits as when locations are publicly observable, as in Scenario (i). Consider first Scenario (ii). Under Conditions I_a and S_b , the platform's pricing problem on side a is symmetric across any two locations. This is because of two reasons. First, the location of any agent from side a provides no information about the agent's vertical preferences (this is guaranteed by Condition I_a). Second, when the platform offers the same tariffs as in Scenario (i), the gross utility that each type $\theta_k = (v_a, x_a)$ obtains from the matching set $\mathbf{s}_a^*(\theta_k)$ coincides with the gross utility obtained by type $(v_a, x_a + d)$ from choosing the matching set $\mathbf{s}_a^*(v_a, x_a + d)$, $d \in [0, 1/2]$. Furthermore, the matching set $\mathbf{s}_a^*(v_a, x_a + d)$ is a *parallel translation* of the matching set $\mathbf{s}_a^*(v_a, x_a)$ by d units of distance, along the horizontal dimension (this is guaranteed by Condition S_b). As a result, when, in Scenario (ii), the platform offers the same profit-maximizing tariffs as in Scenario (i), the matching sets demanded by any two agents with types (v_a, x_a) and $(v_a, x_a + d)$ are parallel translations of one another, and are priced identically. The above properties imply that, when the platform offers the same tariffs as in Scenario (i), agents on side a continue to find it optimal to select the same matching sets as in Scenario (i). A symmetric situation applies to Scenario (iii).

In Scenario (iv), instead, locations are private information on both sides. In this case, when the platform offers the same profit-maximizing tariffs as in Scenario (i), agents continue to choose the same matching sets, provided that Condition S_k holds on both sides of the market.

Importantly, under the optimal tariffs of Proposition 1, for any given location x_k , the matching sets demanded by those agents with higher vertical types are supersets of those demanded by agents with lower vertical types. In this sense, the induced matching sets \mathbf{s}_k^* exhibit negative assortativeness

at the margin. Side- l agents located at x_l with a low vertical type v_l are included in the matching sets of the side- k agents located at x_k only if the latter’s vertical types v_k are large enough. To understand, consider Scenario (i), bearing in mind that the same conclusions apply to Scenarios (ii)-(iv) under the additional conditions in the proposition. Because locations are observable, the marginal profits the platform obtains by matching type $\theta_l = (v_l, x_l)$ from side l to type $\theta_k = (v_k, x_k)$ from side k are positive if, and only if,

$$\varphi_k(\theta_k, \theta_l) + \varphi_l(\theta_l, \theta_k) \geq 0. \tag{7}$$

Echoing Bulow and Roberts (1989), the above condition can be interpreted as stating that two agents are matched if, and only if, their *joint* marginal revenue to the platform is weakly positive (we elaborate on this point further in the next subsection). Condition R guarantees that $\varphi_l(\theta_k, \theta_l)$ is non-decreasing in v_l . This implies existence of a threshold $t_k^*(\theta_k, x_l)$ such that Condition (7) is satisfied if, and only if, $v_l \geq t_k^*(\theta_k, x_l)$. Moreover, the threshold $t_k^*(\theta_k, x_l)$ is non-increasing in v_k and, when virtual valuations are non-increasing in distance, $t_k^*(\theta_k, x_l)$ is also non-decreasing in the distance $|x_k - x_l|$. This means that, as v_k increases, the matching set of type θ_k expands to include new agents with lower vertical types. Moreover, when virtual valuations are non-increasing in distance, as v_k increases, the vertical type of the marginal agents located at x_l added to the matching set $\mathbf{s}_k^*(v_k, x_k)$ are higher the “farther” away the location x_l is from x_k . These assortativeness properties of matching demands can be tested empirically (see for, example, the recent work by Kang and You (2016)). Figure 1 illustrates the above properties by depicting the matching set of a representative agent from side a .




Figure 1: Matching sets under profit-maximizing tariffs. The shaded area in the figure describes the matching set for an agent from side a located at $x_a = 1/2$.

Remark 1. Proposition 1 describes symmetry conditions under which customized tariffs are offered by profit-maximizing platforms under imperfect information about the agents’ locations. These conditions rule out the optimality of bunching and other complications (such as lotteries over tariffs

and matching sets), while enabling us to highlight the role of price customization. These conditions can be relaxed by considering a model with a finite number of locations equidistant in the unit-circle. In a discrete setup, the agents' incentive to lie over the horizontal dimensions are diminished. We opted, here, for a continuum of locations for its greater analytical convenience, especially for pricing, as we show below.

3.1 Lerner-Wilson formula for matching schedules

We now derive further properties of the customized tariffs that maximize the platform's profits. To facilitate the exposition, hereafter, we assume that locations are public on both sides (that is, the environment satisfies the conditions of Scenario (i)). As explained in the proof of Proposition 1, in this case, because each side- k agent located at x_k is constrained to choosing the plan $(\underline{\mathbf{S}}_k(x_k), \underline{\mathbf{T}}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathbf{S}_k(x_k))$, designed for him, there is no need for the platform to restrict the set of possible customizations $\mathbf{S}_k(x_k)$ each x_k -agent can choose from; that is, without loss of optimality, the platform can set $\mathbf{S}_k(x_k) = \Sigma(\Theta_l)$, for all $x_k \in [0, 1]$. To ease the notation, hereafter, we then drop the sets $\mathbf{S}_k(x_k)$ from the specification of the matching plans.

Next, consider the problem of a side- k agent of type $\theta_k = (v_k, x_k)$ under the plan

$$(\underline{\mathbf{S}}_k(x_k), \underline{\mathbf{T}}_k(x_k), \rho_k(\cdot|\cdot; x_k)).$$

The mass of agents located at x_l demanded by type θ_k is given by

$$\hat{q}_{x_l}(\theta_k) \in \arg \max_{q \in [0, f_l^x(x_l)]} \{u_k(v_k, |x_k - x_l|) \cdot q - \rho_k(q|x_l; x_k)\}.$$

Assuming the price schedule $\rho_k(\cdot|x_l; x_k)$ is differentiable and convex, we then have that, whenever $\hat{q}_{x_l}(\theta_k) \in (0, f_l^x(x_l))$, $\hat{q}_{x_l}(\theta_k)$ is a solution to the following first-order condition:¹⁰

$$u_k(v_k, |x_k - x_l|) = \frac{d\rho_k}{dq}(\hat{q}_{x_l}(\theta_k)|x_l; x_k). \quad (8)$$

For future reference, for any pair of locations $x_k, x_l \in [0, 1]$, any marginal price $\frac{d\rho_k}{dq} \in [u_k(\underline{v}_k, |x_k - x_l|), u_k(\bar{v}_k, |x_k - x_l|)]$ let $\hat{v}_{x_l} \left(\frac{d\rho_k}{dq} |x_k \right)$ denote the value of v_k that solves the equation $u_k(v_k, |x_k - x_l|) = \frac{d\rho_k}{dq}$, where, to ease the notation, we dropped the arguments $(q|x_l; x_k)$ of the marginal price.

Note that, because the price function $\rho_k(\cdot|x_l; x_k)$ is strictly convex over the range of quantities purchased in equilibrium, the marginal price $d\rho_k/dq$ uniquely identifies the quantity q . Furthermore, because agents with higher vertical types demand larger matching sets, the *demand* for the q -th unit

¹⁰The strict convexity of the price function $\rho_k(\cdot|x_l; x_k)$ over the set of quantities purchased in equilibrium is a direct implication of the supermodularity of the agents' payoffs $u_k(v_k, |x_k - x_l|) \cdot q$ in (v_k, q) .

of the x_l -agents by the x_k -agents, at the *marginal price* $\frac{d\rho_k}{dq}$, is given by¹¹

$$D_k \left(\frac{d\rho_k}{dq} | x_l; x_k \right) \equiv \left[1 - F_k^{v|x} \left(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq} | x_k \right) | x_k \right) \right] f_k^x(x_k), \quad (9)$$

where, as above, we dropped the arguments $(q|x_l; x)$ of the marginal price to lighten the notation. Accordingly, $D_k \left(\frac{d\rho_k}{dq} | x_l; x_k \right)$ coincides with the mass of agents from side k located at x_k whose vertical type exceeds $\hat{v}_{x_l} \left(\frac{d\rho_k}{dq} | x_k \right)$.

Using (9), we then define the elasticity of the demand by the side- k agents located at x_k (in short, the x_k -demand) for the q -th unit of the x_l -agents with respect to its marginal price $\frac{d\rho_k}{dq}$ by (once again, the arguments of the marginal price $\frac{d\rho_k}{dq}$ are dropped to ease the notation)

$$\varepsilon_k \left(\frac{d\rho_k}{dq} | x_l; x_k \right) \equiv - \frac{\partial D_k \left(\frac{d\rho_k}{dq} | x_l, x_k \right)}{\partial \left(\frac{d\rho_k}{dq} \right)} \cdot \frac{\frac{d\rho_k}{dq}}{D_k \left(\frac{d\rho_k}{dq} | x_l, x_k \right)}. \quad (10)$$

Note that, when $\frac{d\rho_k}{dq} (q|x_l; x_k) \in [u_k(\underline{v}_k, |x_k - x_l|), u_k(\bar{v}_k, |x_k - x_l|)]$,

$$\varepsilon_k \left(\frac{d\rho_k}{dq} | x_l; x_k \right) = \lambda_k^{v|x} \left(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq} | x_k \right) | x_k \right) \cdot \left[\frac{\partial u_k}{\partial v} \left(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq} | x_k \right), |x_k - x_l| \right) \right]^{-1} \cdot \frac{d\rho_k}{dq}, \quad (11)$$

where recall that $\lambda_k^{v|x} (v_k|x_k) \equiv f_k^{v|x} (v_k|x_k) / [1 - F_k^{v|x} (v_k|x_k)]$ is the hazard-rate of the conditional distribution $F_k^{v|x}$. The next proposition characterizes the price schedules associated with the profit-maximizing customized tariffs of Proposition 1 in terms of the location-specific elasticities of the demands on both sides of the market.

Proposition 2. (Lerner-Wilson price schedules) *In addition to Condition R, suppose the environment satisfies the properties of Scenario (i). The price schedules $\rho_k^*(\cdot|x_l; x_k)$ associated with the profit-maximizing customized tariffs T_k^* are differentiable and convex over the equilibrium range $[q_{x_l}(\mathbf{s}_k(\underline{v}_k, x_k)), q_{x_l}(\mathbf{s}_k(\bar{v}_k, x_k))]$, $k = a, b$. Moreover, ρ_a^* and ρ_b^* jointly satisfy the following Lerner-Wilson formulas for all $(x_a, x_b) \in [0, 1]$, all $(q_a, q_b) \in [0, 1]$ such that $q_a = D_b \left(\frac{d\rho_b^*}{dq} (q_b|x_a; x_b) | x_a; x_b \right)$ and $q_b = D_a \left(\frac{d\rho_a^*}{dq} (q_a|x_b; x_a) | x_b; x_a \right)$:*

$$\underbrace{\frac{d\rho_a^*}{dq} (q_a|x_b; x_a) \left(1 - \frac{1}{\varepsilon_a \left(\frac{d\rho_a^*}{dq} (q_a|x_b; x_a) | x_b; x_a \right)} \right)}_{\text{net effect on side-a profits}} + \underbrace{\frac{d\rho_b^*}{dq} (q_b|x_a; x_b) \left(1 - \frac{1}{\varepsilon_b \left(\frac{d\rho_b^*}{dq} (q_b|x_a; x_b) | x_a; x_b \right)} \right)}_{\text{net effect on side-b profits}} = 0. \quad (12)$$

¹¹By the *demand* for the q -th unit of the x_l -agents by the x_k -agents we mean the mass of agents from side k located at x_k who demand at least q matches with the x_l -agents. Also, hereafter, we extend the definition in (9) to marginal prices $\frac{d\rho_k}{dq} (q|x_l; x_k) \notin [u_k(\underline{v}_k, |x_k - x_l|), u_k(\bar{v}_k, |x_k - x_l|)]$ by letting $\hat{v}_{x_l} \left(\frac{d\rho_k}{dq} | x_k \right) = \underline{v}_k$ for all $\frac{d\rho_k}{dq} (q|x_l; x_k) < u_k(\underline{v}_k, |x_k - x_l|)$, and $\hat{v}_{x_l} \left(\frac{d\rho_k}{dq} | x_k \right) = \bar{v}_k$ for all $\frac{d\rho_k}{dq} (q|x_l; x_k) > u_k(\bar{v}_k, |x_k - x_l|)$.

The Lerner-Wilson formulas (12) jointly determine the price schedules on both sides of the market. Intuitively, these formulas require that the marginal contribution to profits from adding to the matching sets of the x_k -agents the q_k -th unit of the x_l -agents coincide with the marginal contribution to profits from adding to the matching sets of the x_l -agents the q_l -th unit of the x_k -agents, where q_k and q_l are jointly related by the reciprocity condition in the Proposition (that is, $q_a = D_b \left(\frac{d\rho_b^*}{dq} (q_b|x_a; x_b) |x_a; x_b \right)$ and $q_b = D_a \left(\frac{d\rho_a^*}{dq} (q_a|x_b; x_a) |x_b; x_a \right)$). As for the standard Lerner-Wilson formula for monopoly/monopsony pricing, on each side, the marginal contribution to profits of such an adjustment has two terms: the term $\frac{d\rho_k^*}{dq} (q_k|x_l, x_k)$ captures the marginal benefit from adding the extra agents, whereas the semi-inverse-elasticity term $\frac{d\rho_k^*}{dq} (q_k|x_l, x_k) \left[\varepsilon_k \left(\frac{d\rho_k^*}{dq} (q_k|x_l; x_k) |x_l; x_k \right) \right]^{-1}$ capture its associated infra-marginal losses.

Importantly, as anticipated above, the quantities q_k and q_l at which the conditional price schedules are evaluated have to clear the market, as required by the reciprocity condition (3). The result in the proposition uses the fact that the demands under the optimal tariffs satisfy the threshold structure in Proposition 1 to establish that the mass of x_k -agents that, at the marginal price $\frac{d\rho_k^*}{dq} (q_k|x_l; x_k)$, demand q_k agents or more of type x_l coincide with the mass $D_k \left(\frac{d\rho_k^*}{dq} (q_k|x_l; x_k) |x_l; x_k \right)$ of x_k -agents with vertical type above $\hat{v}_{x_l} \left(\frac{d\rho_k^*}{dq} |x_k \right)$. Together with reciprocity, Proposition 1 then also implies that the mass q_k of x_l -agents that, at the marginal price $\frac{d\rho_l^*}{dq} (q_l|x_k, x_l)$, demand $q_l = D_k \left(\frac{d\rho_k^*}{dq} (q_k|x_l; x_k) |x_l; x_k \right)$ or more of the x_k -agents coincides with the mass of x_l -agents with vertical type above $\hat{v}_{x_k} \left(\frac{d\rho_l^*}{dq} |x_l \right)$.

Finally, that the price schedules $\rho_k^*(q_k|x_l; x_k)$ are convex in q_k reflects the fact that the matching demands of the x_k -agents for the x_l -agents are increasing in the vertical types v_k . As a result, the marginal price $\frac{d\rho_k^*}{dq} (q_k|x_l, x_k)$ for the q_k -unit of the x_l -agents has to increase with q_k .

The formulas in (12) also reveal how profit-maximizing platforms optimally cross-subsidize interactions among agents from multiple sides of the market while accounting for heterogeneity in preferences along both vertical and horizontal dimensions. In particular, the price schedules offered at any two locations x_k and x_l are a function of the *location-specific demand elasticities* $\varepsilon_k(\cdot|x_l; x_k)$ at these locations. This reflects the fact that, at the optimum, platforms make use of information about horizontal preferences to offer matching tariffs that extract as much surplus as possible from agents from both sides. As we show below, the ability to tailor price schedules to locations (a form of third-degree price discrimination) has important implications for the composition of the demands prevailing under optimal tariffs.

The result in Proposition 2 can also be used in empirical work, as it provides a system of structural equations that allows the econometrician to recover the distribution of the agents' preferences from price schedules and match volumes. The work by Kahn and You (2016) follows a related approach in the matching market for lobbying, but abstracting from horizontal differentiation. Proposition 2 might help in extending their empirical analysis to markets where horizontal differentiation is expected to play an important role.

3.2 Distortions and Horizontal Differentiation

We now investigate how distortions in the supply of matching opportunities (relative to efficiency) due to market power vary with the agents' horizontal preferences. In particular, we are interested in whether distortions increase or decrease as one considers locations farther away from an agent's bliss point. The analysis in this section has important implications for how policy makers should regulate mediated matching markets.

Let

$$\sum_{k=a,b} \int_{\Theta_k} \int_{\hat{s}_k(\theta_k; T_k)} u_k(v_k, |x_k - x_l|) dF_j(\theta_j) dF_k(\theta_k), \quad (13)$$

denote the welfare associated with a feasible pair of matching tariffs T_k , $k = a, b$. It is straightforward to see that a pair of tariffs T_k^e , $k = a, b$, maximizes social welfare if, and only if, the induced matching demands satisfy the following property: For any two agents with types $\theta_k = (v_k, x_k)$ and $\theta_l = (v_l, x_l)$,

$$\theta_l \in \hat{s}_k(\theta_k; T_k^e) \iff u_k(v_k, |x_k - x_l|) + u_l(v_l, |x_k - x_l|) \geq 0.$$

In turn, this means that, given any pair of welfare-maximizing tariffs T_k^e , $k = a, b$, there must exist threshold functions $t_k^e(\theta_k, x_l)$ such that $\theta_l \in \hat{s}_k(\theta_k; T_k^e)$ if, and only if, $v_l \geq t_k^e(\theta_k, x_l)$. Arguments similar to those establishing Proposition 1 and Proposition 2 then imply that the welfare-maximizing tariffs T_a^e and T_b^e are customized, and their associated price schedules jointly solve

$$\frac{d\rho_a^e}{dq}(q_a|x_b; x_a) + \frac{d\rho_b^e}{dq}(q_b|x_a; x_b) = 0,$$

at any pair q_a and q_b such that $q_a = D_b\left(\frac{d\rho_b^e}{dq}(q_b|x_a; x_b)|x_a; x_b\right)$ and $q_b = D_a\left(\frac{d\rho_a^e}{dq}(q_a|x_b; x_a)|x_b; x_a\right)$.

The following example illustrates the differences between the matching sets sustained under welfare maximization and their counterparts under profit maximization, when preferences are as in Example 2.

Example 3. (constant distortions) In addition to Condition R, suppose the environment satisfies the properties of Scenario (i) and Condition I_k holds for $k = a, b$. Further assume that agents' preferences are as in Example 2 with $\delta = 1$, v_b drawn from a shifted exponential distribution with parameters $\tilde{\lambda}_b > 0$ and $K < 0$, and v_a drawn from a distribution F_a^v such that $\varphi_a(\theta_a, \theta_b) > 0$ all (θ_a, θ_b) .¹² Assume $|K|$ is large so that, for any $(\theta_a, x_b) \in \Theta_a \times [0, 1]$, $t_a^*(\theta_a, x_b), t_a^e(\theta_a, x_b) \in \text{Int}[V_b]$. The distortions brought in by profit maximization are then captured by the discrepancy

$$t_a^*(\theta_a, x_b) - t_a^e(\theta_a, x_b) = \frac{1}{\tilde{\lambda}_b} + \frac{\alpha}{\lambda_a^v(v_a)}$$

between the thresholds defining the matching sets under profit maximization and welfare maximization, respectively. Under the specification of this example, such discrepancy is invariant in the distance $|x_a - x_b|$ between any two pair of locations, as illustrated in Figure 2.

¹²That is, $F_b^v(v_b) = 1 - \exp\{-\tilde{\lambda}_b(v_b - K)\}$, with $\tilde{\lambda}_b > 0$ and $K < 0$. Its support is then $[K, +\infty)$.




Figure 2: The welfare-maximizing demand threshold $t_a^e(\theta_a, x_b)$ (dashed blue curve) and profit-maximizing demand threshold $t_a^*(\theta_a, x_b)$ (solid black curve) for agents on side a located at $x_a = .5$ under the preferences of Example 2, when $\delta = 1$, $\phi(|x_a - x_b|) = \exp\{-|x_a - x_b|\}$, and vertical dimensions v_b are drawn from a shifted exponential distribution. Horizontal dimensions are drawn uniformly on both sides, independently of the vertical dimensions.

As the next proposition reveals, the above example is a knife-edge case, in that the match utility of the side- a agents is modular (i.e., the cross-derivative of u_a in v_a and $\mu = |x_a - x_b|$ is zero), the match utility of the side- b agents is modular and linear in v_b , and the distribution F_b^v of the vertical parameter v_b has a constant hazard rate.

To state the formal result, we first need to introduce a more general definition:

Definition 4. (distortions and distance) Distortions on side $k \in \{a, b\}$ decrease (alternatively, increase) with distance if, and only if,

$$u_l(t_k^*(\theta_k, x_l), |x_l - x_k|) - u_l(t_k^e(\theta_k, x_l), |x_l - x_k|)$$

decreases (alternatively, increases) with $|x_k - x_l|$.

Hence, fixing the type $\theta_k = (v_k, x_k)$ of a side- k agent, distortions increase with distance when the difference between the minimal utility asked by a profit-maximizing platform and a welfare-maximizing platform to each x_l -agent to be matched with type θ_k increases with the distance between the two agents' locations. Note that the difference in utilities $u_l(t_k^*(\theta_k, x_l), |x_l - x_k|) - u_l(t_k^e(\theta_k, x_l), |x_l - x_k|)$ reduces to the difference in the thresholds $t_k^*(\theta_k, x_l) - t_k^e(\theta_k, x_l)$ when the side- l 's preferences are invariant in the locations, as in Examples 1, 2, and 3 above.

Proposition 3. (distortions under customized pricing) In addition to Condition R, suppose the environment satisfies the properties of Scenario (i), and Conditions VD and I_k hold, $k = a, b$.

Consider any $(\theta_k, x_l) \in \Theta_k \times [0, 1]$ for which $t_k^*(\theta_k, x_l), t_k^e(\theta_k, x_l) \in \text{Int}[V_l]$.¹³ The following statements are true for $k, l = a, b, l \neq k$:

1. If u_k is submodular, u_l is submodular and concave in v_l , and F_l^v has an increasing hazard rate, distortions on side k decrease with distance (strictly, if at least one of the conditions is strict);
2. If u_k is supermodular, u_l is supermodular and convex in v_l , and F_l^v has a decreasing hazard rate, distortions on side k increase with distance (strictly, if at least one of the conditions is strict).

Consider again the specification in Example 3 but assume $\delta < 1$, in which case, the preferences of the side- a agents are strictly submodular. In this case, distortions decrease with distance. This means that, under profit maximization, the matching sets of the side- a agents are distorted by excluding primarily those agents from side b that the side- a agents like the most. Conversely, when $\delta > 1$, the preferences of the side- a agents are supermodular and distortions increase with distance, meaning that the side- b agents that are inefficiently excluded from the matching sets of the side- a agents are primarily those that the side- a agents regard as least attractive.¹⁴ Similar conclusions obtain when preferences are modular but F_b^v has a strictly monotone hazard rate (increasing or decreasing).

Next, consider the ad exchange specification of Example 1. Note that the agents' preferences under this specification are consistent with the conditions of part 1 in Proposition 3 (u_a is strictly submodular, whereas u_b is linear in v_b and invariant in $|x_a - x_b|$). The results in Proposition 3 then imply that, when F_b^v has an increasing hazard rate, under profit maximization, advertisers (on side a) are more often matched (relative to efficiency) to those publishers whose profile is more distant from their ideal audience. Figure 3 illustrates this situation.

¹³The results for the case in which either $t_k^*(\theta_k, x_l)$, or $t_k^e(\theta_k, x_l)$, coincide with the extreme points of V_l are not particularly interesting. In this case, the monotonicity of the difference in the thresholds $t_k^*(\theta_k, x_l) - t_k^e(\theta_k, x_l)$ does not depend on the hazard rate of the distributions, nor on the modularity of the match values. It follows directly from the monotonicity of the match values u_k in v_k and $|\cdot|$, and from Conditions R and VD. When $t_k^e(\theta_k, x_l) = \bar{v}_l$, $t_k^*(\theta_k, x_l) = \bar{v}_l$. In this case, distortions are invariant in the distance. If, instead, $t_k^e(\theta_k, x_l) = \underline{v}_l$, then distortions are weakly increasing in the distance. Finally, if $t_k^*(\theta_k, x_l) = \bar{v}_l$ then clearly distortions are weakly decreasing in distance.

¹⁴To see how the results follow from parts 1 and 2 in the proposition, recall that, in this example, the preferences of the side- b agents are linear in v_b and invariant in $|x_a - x_b|$; hence they are both weakly submodular and concave in v_b , and weakly supermodular and convex in v_b . Furthermore, F_l^v is exponential and hence it has a hazard rate that is both weakly increasing and weakly decreasing.




Figure 3: The welfare-maximizing demand threshold $t_a^e(\theta_a, x_b)$ (dashed blue curve) and profit-maximizing demand threshold $t_a^*(\theta_a, x_b)$ (solid black curve) for agents on side a located at $x_a = .5$ under the preferences specification of Example 1.

Proposition 3 identifies natural conditions on primitives (match utilities and distributions) that lead to location-specific elasticities that increase or decrease with distance. To understand the result, consider part 1 first. When u_a is submodular, the effect of vertical differentiation on match utility declines with the distance in locations. This implies that the location-specific semi-elasticity $\varepsilon_a \left(\frac{d\rho_a}{dq} | x_b; x_a \right) / \frac{d\rho_a}{dq}$ of the side- a demand, evaluated along the matching demand of type θ_a , is higher at locations x_b that are farther from x_a (as implied by equation (11)). Holding fixed $\frac{d\rho_b}{dq}$, this property of elasticities contributes to marginal prices on side a that are lower (relative to their efficient counterparts) at locations x_b that are farther from x_a , as implied by the Lerner-Wilson formulas (12). Likewise, when F_b^v has an increasing hazard rate and u_b is submodular and concave in v_l , the side- b location-specific semi-elasticity $\varepsilon_b \left(\frac{d\rho_b}{dq} | x_a; x_b \right) / \frac{d\rho_b}{dq}$, evaluated along the matching demand of type θ_a , is also higher at locations x_b that are farther from x_a (as implied by equation (11)). Holding fixed $\frac{d\rho_a}{dq}$, this property of elasticities thus contributes to marginal prices on side b that are lower (relative to their efficient counterparts) at locations x_b that are farther from x_a (as implied by the Lerner-Wilson formulas (12)). Because matchings is reciprocal, the side- a marginal prices then also decrease (relative to their efficient counterparts) as distances increase.

The two effects described above point in the same direction: Namely, the discrepancy between the side- a profit-maximizing marginal prices and their efficient counterparts are relatively lower at locations x_b that are farther from x_a . When this is the case, distortions decrease with distance. Mutatis mutandis, a similar explanation applies to the environment considered in part 2 in the proposition.

4 Uniform Pricing Regulation

In response to concerns about the detrimental welfare effects of price customization, policies advocating for the imposition of uniform-pricing regulations have recently gained prominence in the policy debate. In this section, we study how platforms optimally respond to the imposition of uniform pricing.

Uniform Pricing and Aggregate Demand Elasticities

Suppose the platform is forced to adopt a uniform price schedule $p_a(\cdot|x_b)$ on side a . Recall that, for each location $x_b \in [0, 1]$, and each $q \in [0, f_b^x(x_b)]$, such a schedule specifies the price that the side- a agents have to pay to be matched to q agents from side b located at x_b . Under such schedule, the aggregate demand (over all locations x_a) for the q -th unit of the x_b -agents at the marginal price $\frac{dp_a}{dq}(q|x_b)$ is equal to

$$\bar{D}_a \left(\frac{dp_a}{dq} | x_b \right) \equiv \int_0^1 D_a \left(\frac{dp_a}{dq} | x_b; x_a \right) dx_a = \int_0^1 \left\{ 1 - F_a^{v|x} \left(\hat{v}_{x_b} \left(\frac{dp_a}{dq} | x_a \right) | x_a \right) \right\} f_a^x(x_a) dx_a,$$

where, as in the previous section, $D_a \left(\frac{dp_a}{dq} | x_b; x_a \right)$ denotes the mass of agents located at x_a that demand q units or more of the x_b -agents, and where, as in the previous section, the arguments $(q|x_b)$ of the marginal prices $\frac{dp_a}{dq}(q|x_b)$ have been dropped, to ease the exposition.

The elasticity of the aggregate demand for the q -th unit of the x_b -agents with respect to its marginal price $\frac{dp_a}{dq}(q|x_b)$ is then equal to

$$\bar{\varepsilon}_a \left(\frac{dp_a}{dq} | x_b \right) \equiv - \frac{\partial \bar{D}_a \left(\frac{dp_a}{dq} | x_b \right)}{\partial \left(\frac{dp_a}{dq} \right)} \cdot \frac{\frac{dp_a}{dq}}{\bar{D}_a \left(\frac{dp_a}{dq} | x_b \right)} = \mathbb{E}_{\bar{H}(\tilde{x}_a | x_b, \frac{dp_a}{dq})} \left[\varepsilon_a \left(\frac{dp_a}{dq} | x_b; \tilde{x}_a \right) \right],$$

where the expectation is over $X_a = [0, 1]$, under the distribution $\bar{H}(\cdot | x_b, \frac{dp_a}{dq})$ whose density is equal to

$$\bar{h} \left(x_a | x_b, \frac{dp_a}{dq} \right) \equiv \frac{D_a \left(\frac{dp_a}{dq} | x_b, x_a \right)}{\int_0^1 D_a \left(\frac{dp_a}{dq} | x_b; x'_a \right) dx'_a}$$

with $\varepsilon_a \left(\frac{dp_a}{dq} | x_b; x_a \right)$ as defined in (10).

Hereafter, we refer to $\bar{\varepsilon}_a(\cdot|x_b)$ as to the *aggregate elasticity* of the side- a demand for the x_b -matches. This elasticity measures the percentage variation in the mass of agents from side a that demand at least q matches with the side- b agents located at x_b in response to a percentage change in the marginal price for the q -th unit of the x_b -agents. It is also equal to the average (over the side- a locations) elasticity of the x_a -demands for the q -th unit of the x_b -agents with respect to the marginal price $\frac{dp_a}{dq}$, where the average is under a distribution that assigns to each location x_a a weight equal to the mass of agents $D_a \left(\frac{dp_a}{dq} | x_b, x_a \right)$ located at x_a demanding q units, or more, of the x_b -agents.

The next proposition derives the optimal tariffs employed by a platform that is constrained not to engage in third-degree price discrimination on side a (equivalently, to price uniformly on side a). Consistently with the previous two propositions, we ease the exposition by assuming that locations are public information on both sides (that is, the environment satisfies the properties of Scenario (i)) and then denote by T_a^u and T_b^u the profit-maximizing tariffs on sides a and b , respectively, when the platform is constrained to price uniformly on side a .

Proposition 4. (Uniform Pricing) *In addition to Condition R, suppose the environment satisfies the properties of Scenario (i) and that the platform is constrained to price uniformly on side a (but is free to offer a customized tariff on side b). The profit-maximizing price schedules $p_a^u(\cdot|x_b)$ and $\rho_b^u(\cdot|x_a;x_b)$ are differentiable and convex over the equilibrium ranges $[q_{x_b}(\mathbf{s}_a(\underline{v}_a, x_b + .5)), q_{x_b}(\mathbf{s}_a(\bar{v}_a, x_b))]$, and $[q_{x_a}(\mathbf{s}_b(\underline{v}_b, x_b)), q_{x_a}(\mathbf{s}_a(\bar{v}_b, x_b))]$, and jointly satisfy the following optimality conditions for all $x_b \in [0, 1]$, all $q \leq f_b^x(x_b)$,*

$$\underbrace{\frac{dp_a^u}{dq}(q|x_b) \left(1 - \frac{1}{\bar{\varepsilon}_a \left(\frac{dp_a^u}{dq}(q|x_b) \right)} \right)}_{\text{net effect on side-a profits}} + \mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a}{dq})} \left[\underbrace{\frac{d\rho_b^u}{dq}(\hat{q}_b(q; \tilde{x}_a; x_b)|\tilde{x}_a; x_b) \left(1 - \frac{1}{\varepsilon_b \left(\frac{d\rho_b^u}{dq}(\hat{q}_b(q; \tilde{x}_a; x_b)|\tilde{x}_a; x_b)|\tilde{x}_a; x_b \right)} \right)}_{\text{net effect on side-b profits}} \right] = 0, \quad (14)$$

where $H\left(x_a|x_b, \frac{dp_a}{dq}\right)$ is the distribution over $X_a = [0, 1]$ whose density is given by

$$h\left(x_a|x_b, \frac{dp_a}{dq}\right) \equiv \frac{\frac{\partial D_a\left(\frac{dp_a}{dq}(q|x_b)|x_b; x_a\right)}{\partial\left(\frac{dp_a}{dq}\right)}}{\frac{\partial \bar{D}_a\left(\frac{dp_a}{dq}(q|x_b)|x_b\right)}{\partial\left(\frac{dp_a}{dq}\right)}},$$

and where $\hat{q}_b(q; x_a; x_b) \equiv D_a\left(\frac{dp_a}{dq}(q|x_b)|x_b; x_a\right)$ is the mass of side- a agents located at x_a that demand q or more matches with the side- b agents located at x_b when the (uniform) marginal price for the q -th unit of the x_b -agents is equal to $\frac{dp_a}{dq}(q|x_b)$.

The result in the proposition provides structural equations similar to those corresponding to the Lerner-Wilson formulas in (12), but adapted to account for the imposition of uniform pricing on side a . Such structural conditions jointly determine the price schedules on both sides of the market. Under uniform pricing, the price schedule on side a for the sale of the side- b matches cannot condition on the location of the side- a agents. As a result, the markup for the sale of the q -th unit of the x_b -matches is constant across all side- a locations x_a . The relevant elasticity for determining this markup

is then the aggregate elasticity $\bar{\varepsilon}_a(\cdot|x_b)$, rather than the location-specific elasticities $\varepsilon_a(\cdot|x_b; x_a)$ in the Lerner-Wilson formula (12). Interestingly, even if the platform can price discriminate on side b by offering a customized tariff to the side- b agents, when it is constrained to price uniformly on side a , the cost of procuring the side- b agents is the average (mark-up augmented) price

$$\mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a}{dq})} \left[\underbrace{\frac{d\rho_b^u}{dq}(\hat{q}_b(q; \tilde{x}_a; x_b)|\tilde{x}_a; x_b)}_{\text{net effect on side-}b \text{ profits}} \left(1 - \frac{1}{\varepsilon_b \left(\frac{d\rho_b^u}{dq}(\hat{q}_b(q; \tilde{x}_a; x_b)|\tilde{x}_a; x_b)|\tilde{x}_a; x_b \right)} \right) \right]$$

charged to the x_b -agents for their interactions with the various x_a -agents demanding q , or more, x_b -matches.

Also note that, by virtue of the reciprocity condition (3), the quantities q_a and q_b at which the conditional price schedules are evaluated have to clear the market for any pair of locations (x_a, x_b) . For this to be possible, it is important that the platform be able to employ a customized tariff on side b , as the latter ensures that the platform has enough price instruments to procure the side- b matches demanded by the side- a agents, while respecting reciprocity.

Finally, as in the case where price customization is allowed on both sides, the convexity of the price schedules $p_a^u(\cdot|x_b)$ and $\rho_b^u(\cdot|x_a; x_b)$ in q reflects the fact that the matching demands of those agents with a higher vertical type are supersets of those with a lower vertical type.

As revealed by the pricing formulas (12) and (14), the effects of the imposition of uniform pricing on side a on the composition of the matching sets on both sides hinge on the comparison between the aggregate inverse-elasticity $1/\bar{\varepsilon}_a$ and the location-specific inverse-elasticities $1/\varepsilon_a(\cdot|x_b, x_a)$ on side a , as well as the comparison between the average inverse-elasticity $\mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a}{dq})} [1/\varepsilon_b(\cdot|\tilde{x}_a; x_b)|\tilde{x}_a; x_b]$ of the x_b -demands for the various x_a -matches and the inverse-elasticities $1/\varepsilon_b(\cdot|\tilde{x}_a; x_b)|\tilde{x}_a; x_b)$ of the same demands for the specific matches. In turn, such comparisons naturally reflect how the average virtual valuations on both sides compare to their location-specific counterparts. To see this, first note that

$$\frac{1}{\bar{\varepsilon}_a \left(\frac{dp_a^u}{dq}(q|x_b) \right)} = \mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a}{dq})} \left[\frac{1}{\varepsilon_a \left(\frac{dp_a^u}{dq}(q|x_b)|x_b; \tilde{x}_a \right)} \right]. \quad (15)$$

That is, the inverse aggregate elasticity of the side- a demand for the q -th unit of the x_b -matches is equal to the average of the various location-specific inverse elasticities of the side- a agents for the same unit of the same x_b -matches, where the average is under the same measure $H \left(x_a|x_b, \frac{dp_a}{dq} \right)$ introduced in the proposition. Letting $v_b'(q; x_b) = \left(F_b^{v|x} \right)^{-1} (1 - q|x_b)$ denote the value of the vertical dimension of the x_b -agents such that the mass of x_b -agents with a vertical type higher than $v_b'(q; x_b)$

is equal to q , we then have that the optimality condition (14) can be re-written as

$$\begin{aligned} & \underbrace{\mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a}{dq})} \left[\varphi_a \left(\left(\hat{v}_{x_b} \left(\frac{dp_a}{dq} | \tilde{x}_a \right), \tilde{x}_a \right), (v'_b(q; x_b), x_b) \right) \right]}_{\text{net effect on side-}a \text{ profits}} \\ & \underbrace{\mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a}{dq})} \left[\varphi_b \left((v'_b(q; x_b), x_b), \left(\hat{v}_{x_b} \left(\frac{dp_a}{dq} | \tilde{x}_a \right), \tilde{x}_a \right) \right) \right]}_{\text{net effect on side-}B \text{ profits}} = 0. \end{aligned} \quad (16)$$

Hereafter, we assume that the left-hand side of (16) is monotone in the marginal price dp_a/dq , which amounts to quasi-concavity of the platform's profit function with respect to the marginal price, after accounting for the cost of procuring the x_b -agents, as explained in the proof of Proposition 4. The above property implies that the necessary condition in (16) is also sufficient for optimality.

Now recall that, under price customization (on both sides), the platform matches each pair of agents $\theta_a = (v_a, x_a)$ and $\theta_b = (v_b, x_b)$ if, and only if, type- θ_a 's virtual value for interacting with type θ_b is large enough to compensate for the virtual value that type θ_b derives from interacting with type θ_a (formally, θ_a and θ_b are matched if, and only if, $\varphi_a(\theta_a, \theta_b) + \varphi_b(\theta_b, \theta_a) \geq 0$). Under uniform pricing (on side a), instead, the platform matches the above pair of agents if, and only if, the following is true: if all side- a agents with the same *true value* for interacting with type θ_b as type θ_a were to be matched to type θ_b , the *average virtual value* among such agents for the match with type θ_b would compensate for the *average virtual value* that type θ_b derives from being matched with all such agents. Formally, under uniform pricing (on side a), types θ_a and θ_b are matched if, and only if,

$$\begin{aligned} & \mathbb{E}_{H(\tilde{x}_a|x_b, u_a(v_a, |x_b-x_a))} [\varphi_a((\hat{v}_{x_b}(u_a(v_a, |x_b-x_a)|\tilde{x}_a), \tilde{x}_a), \theta_b)] \\ & + \mathbb{E}_{H(\tilde{x}_a|x_b, u_a(v_a, |x_b-x_a))} [\varphi_b(\theta_b, (\hat{v}_{x_b}(u_a(v_a, |x_b-x_a)|\tilde{x}_a), \tilde{x}_a))] \geq 0. \end{aligned}$$

This observation plays an important role in determining how targeting and welfare are affected by the imposition of uniform pricing, as we show below.

4.1 Targeting under Uniform and Customized Pricing

Digital technology is often praised for its ability to increase match precision (or targeting) in a variety of markets. Yet, technology alone is no guarantee of large targeting gains, as the matching demands enjoyed by agents obviously depend on the pricing practices followed by platforms. Price customization allows a platform to charge agents prices that directly depend on their horizontal preferences. To the extent that agents value the most those matches of higher proximity, one might expect price-customization to hinder targeting, as it permits platforms to set higher prices for those matches the agents like the most. Without further inquiry, this observation seems to lend support to policies that impose uniform pricing. Indeed, recent proposals, requiring stand-alone pricing for

media content, or anonymous pricing for advertising slots, appear to follow this line of reasoning. This intuition is, however, incomplete, as it ignores the (endogenous) changes in prices that the platform undertakes in response to regulations mandating uniform pricing. The analysis below provides some guidelines when it comes to the effects of customized pricing on targeting.

Definition 5. (targeting) Customized pricing (on both sides) leads to more targeting than uniform pricing (on side a) if, for each $\theta_a = (v_a, x_a)$, there exists $d(\theta_a) \in (0, \frac{1}{2})$ such that

$$t_a^*(\theta_a, x_b) - t_a^u(\theta_a, x_b) \begin{cases} < 0 & \text{if } |x_a - x_b| < d(\theta_a) \\ > 0 & \text{if } |x_a - x_b| > d(\theta_a). \end{cases}$$

Conversely, uniform pricing on side a leads to more targeting than customized pricing on both sides if, for each $\theta_a = (v_a, x_a)$, there exists $d(\theta_a) \in (0, \frac{1}{2})$ such that

$$t_a^*(\theta_a, x_b) - t_a^u(\theta_a, x_b) \begin{cases} > 0 & \text{if } |x_a - x_b| < d(\theta_a) \\ < 0 & \text{if } |x_a - x_b| > d(\theta_a). \end{cases}$$

Intuitively, customized pricing (on both sides) leads to more targeting than uniform pricing (on side a) if, under the profit-maximizing customized tariffs, agents demand more matches close to their ideal points, and less matches far from their ideal points, relative to what they do under uniform pricing. Accordingly, the threshold function $t_a^*(\theta_a, x_b)$ under customized pricing is below the corresponding threshold function $t_a^u(\theta_a, x_b)$ for nearby matches (i.e., for locations x_b such that $|x_a - x_b| < d(\theta_a)$), and above $t_a^u(\theta_a, x_b)$ for more distant matches (for which $|x_a - x_b| > d(\theta_a)$). Figure 4 illustrates the situation captured by the above definition.




Figure 4: The threshold function $t_a^*(\theta_a, x_b)$ under customized pricing (solid black curve) and uniform pricing $t_a^u(\theta_a, x_b)$ (dashed blue curve) when customized pricing (on both sides) leads to more targeting than uniform pricing (on side a).

Note that, because matching is reciprocal, the above definition has an analogous implication for side b . Namely, when customized pricing (on both sides) leads to more targeting than uniform pricing

(on side a), then the side- b threshold function under customized pricing, $t_b^*(\theta_b, x_a)$, also single-crosses its counterpart under uniform pricing, $t_b^u(\theta_b, x_a)$, only once, and from below, as a function of the distance $|x_a - x_b|$.

The next proposition identifies primitive conditions on the side- a 's preferences under which uniform pricing on that side leads to more targeting than customized pricing. Let $t_k^*(\cdot)$ be the threshold functions describing the matching sets optimally induced by the platform when it can offer customized tariffs on both sides of the market. Let $L : [0, 1] \times \Theta_b \rightarrow \mathbb{R}$ be the function defined by

$$L(x_a|\theta_b) \equiv \varphi_b(\theta_b, (t_b^*(\theta_b, x_a), x_a)) - \frac{1 - F_a^{v|x}(t_b^*(\theta_b, x_a)|x_a)}{f_a^{v|x}(t_b^*(\theta_b, x_a)|x_a)} \cdot \frac{\partial u_a}{\partial v}(t_b^*(\theta_b, x_a), |x_a - x_b|)$$

all $(x_a, \theta_b) \in [0, 1] \times \Theta_b$.

Proposition 5. (comparison: targeting) *In addition to Condition R, suppose the environment satisfies the properties of Scenario (i). The following statements are true.*

1. *Suppose that, for any θ_b , the function $L(\cdot|\theta_b)$ is nondecreasing in $|x_b - x_a|$. Then uniform pricing (on side a) leads to more targeting than customized pricing (on both sides). The following conditions suffice for $L(x_a|\theta_b)$ to be nondecreasing in $|x_b - x_a|$: in addition to Condition VD and I_a , the match value function u_a is submodular and concave in v_a , F_a^v has an increasing hazard rate, and, for any θ_b , the virtual values $\varphi_b(\theta_b, \theta_a)$ are invariant in $|x_b - x_a|$.*

2. *Suppose that, for any θ_b , the function $L(\cdot|\theta_b)$ is nonincreasing in $|x_b - x_a|$. Then customized pricing (on both sides) leads to more targeting than uniform pricing (on side a). The following conditions suffice for $L(x_a|\theta_b)$ to be nonincreasing in $|x_b - x_a|$: in addition to Condition VD and I_a , the match value function u_a is supermodular and convex in v_a , and F_a^v has a decreasing hazard rate.*

The primitive conditions in Proposition 5 have implications for how location-specific elasticities on both sides of the market compare to the average elasticities. For simplicity, consider a market in which preferences on side b are location-invariant, that is, for all $x_a, x_b \in [0, 1]$, all $v_b \in V_b$, $u_b(v_b, |x_b - x_a|) = v_b$, as in Examples 1 and 2.

Consider first Part 1 and fix the side- b location x_b . Let $\hat{\rho} \equiv \frac{d\rho_a}{dq}(q|x_b; x_a)$ be the marginal price for the q -th unit of the x_b -agents charged to the side- a agents located at x_a under customized pricing. To make things interesting, assume $\hat{\rho} \in (u_a(\underline{v}_a, |x_b - x_a|), u_a(\bar{v}_a, |x_b - x_a|))$. We then have that, under customized pricing, the semi-price elasticity¹⁵

$$\varepsilon_a(\hat{\rho}|x_b; x_a)(\hat{\rho})^{-1} = \lambda_a^v(\hat{v}_{x_b}(\hat{\rho}|x_a)) \cdot \left[\frac{\partial u_a}{\partial v}(\hat{v}_{x_b}(\hat{\rho}|x_a), |x_b - x_a|) \right]^{-1}$$

of the demand for the q -th unit of the x_b -agents by the side- a agents is increasing in the distance $|x_b - x_a|$. By contrast, under uniform pricing, the aggregate counterpart of the above semi-elasticity, which is given by

$$\bar{\varepsilon}_a \left(\frac{dp_a^u}{dq}(q|x_b)|x_b \right) \left(\frac{dp_a^u}{dq}(q|x_b) \right)^{-1},$$

¹⁵Recall that, in this case, $\hat{v}_{x_b}(\hat{\rho}|x_a)$ is defined by the unique solution to $u_a(\hat{v}_{x_b}(\hat{\rho}|x_a), |x_b - x_a|) = \hat{\rho}$.

is constant in x_a (and therefore in the distance $|x_b - x_a|$), as the marginal price $\frac{dp_a^u}{dq}(q|x_b)$ for the q -th unit of x_b -matches is the same for all x_a -locations. As a consequence of the above property, under the assumptions in Part 1, the relevant semi-elasticity in the customized-pricing regime is lower (alternatively, higher) than in the uniform-pricing regime when the distance $|x_b - x_a|$ is small (alternatively, large). This implies that the marginal price $\frac{dp_a^*}{dq}(q|x_b; x_a)$ for the q -th unit of x_b -matches charged to the x_a -agents under the customized-pricing regime is lower than the corresponding price $\frac{dp_a^u}{dq}(q|x_b)$ under the uniform-pricing regime when locations are far apart, whereas the opposite is true at nearby locations. Accordingly, there is more targeting under uniform pricing than under customized pricing.

Mutatis mutandis, the assumptions in Part 2 in the proposition imply that, under customized pricing, the semi-price elasticity of the demand for the q -th unit of the x_b -agents by the side- a agents is strictly decreasing in $|x_b - x_a|$, in which case customized pricing leads to more targeting than uniform pricing.

Proposition 5 helps us understand the effects of price customization in a variety of markets. In the ad exchange setting of Example 1, price customization generates less targeting than uniform pricing (as the match function u_a is submodular and linear in v_a) when F_a^v has an increasing hazard rate (e.g., a uniform or exponential cdf). In the media markets setting of Example 2, instead, price customization leads to more targeting than uniform pricing when the elasticity of substitution is high (namely, when $\delta > 1$) and F_a^v has a decreasing hazard rate (e.g., a Pareto cdf).

When preferences on side b are also location-specific, the additional conditions in the proposition guarantee that the above comparisons remain valid after accounting for differences between average and location-specific elasticities on side b , which are responsible for the procurement costs.

4.2 Welfare under Uniform and Customized Pricing

The result in Proposition 5 can also be used to study the welfare implications of price customization. To see this, suppose the market satisfies the conditions in part 1 in the proposition. Then, under uniform pricing on side- a , the side- a agents face lower marginal prices $\frac{dp_a}{dq}(q|x_b)$ for the x_b -agents they like the most and higher marginal prices for those side- b agents whose location is far from their bliss point.

The above findings permit us to adapt results from the third-degree price discrimination literature to the matching environment under consideration here to identify conditions under which welfare of the side- a agents increases with the imposition of uniform pricing on side a . Formally, recall that, under uniform pricing, the demand by the x_a -agents for each q -th unit of the x_b -matches is given by

$$D_a \left(\frac{dp_a}{dq} | x_b; x_a \right) = \left[1 - F_a^{v|x} \left(\hat{v}_{x_b} \left(\frac{dp_a^u}{dq} | x_a \right) | x_a \right) \right] f_a^x(x_a)$$

where, to ease the notation, we dropped $(q|x_b)$ from the arguments of the the marginal price $\frac{dp_a}{dq}(q|x_b)$.

Now let

$$CD_a \left(\frac{dp_a}{dq} | x_b; x_a \right) = - \frac{\partial^2 D_a \left(\frac{dp_a^u}{dq} | x_b; x_a \right)}{\partial \left(\frac{dp_a^u}{dq} \right)^2} \frac{\frac{dp_a^u}{dq}}{\frac{\partial D_a \left(\frac{dp_a^u}{dq} | x_b; x_a \right)}{\partial \left(\frac{dp_a^u}{dq} \right)}}$$

denote the convexity of the demand by the x_a -agents for the q -th unit of the x_b -agents with respect to the marginal price $\frac{dp_a}{dq}$.¹⁶

Condition 5. [IR] Increasing Ratio: For any $(x_a, x_b) \in [0, 1]^2$, any $q \in [0, f_b^x(x_b)]$, the function

$$z_a \left(\frac{dp_a}{dq} | x_b; x_a \right) \equiv \frac{\frac{dp_a}{dq}}{2 - CD_a \left(\frac{dp_a}{dq} | x_b; x_a \right)}$$

is nondecreasing in the marginal price $\frac{dp_a}{dq}$ for the q -th unit of the x_b -agents.

We then have the following result:

Proposition 6. (comparison: welfare) *In addition to Condition R, suppose the environment satisfies the properties of Scenario (i) and Condition IR holds. The following statements are true.*

1. *Suppose the assumptions in part 1 of Proposition 5 hold and, for any dp_a/dq and x_b , the convexity $CD_a \left(\frac{dp_a}{dq} | x_b; x_a \right)$ of the demands by the x_a -agents for the q -th unit of the x_b -agents declines with the distance $|x_a - x_b|$. Then welfare of the side-a agents is higher under uniform pricing on side a than under customized pricing on both sides.*

2. *Suppose the assumptions in part 2 of Proposition 5 hold and, for any dp_a/dq and x_b , the convexity $CD_a \left(\frac{dp_a}{dq} | x_b; x_a \right)$ of the demands by the x_a -agents for the q -th unit of the x_b -agents increases with the distance $|x_a - x_b|$. Then welfare of the side-a agents is higher under uniform pricing on side a than under customized pricing on both sides.*

Condition IR, as well as the convexity properties of the demand functions in Proposition 6, parallel those in Aguirre et al (2010). The value of the proposition is in showing how our results about the connection between targeting and customized pricing also permit us to apply to the environment under examination here the welfare results from the third-degree price discrimination literature. Note that Proposition 5 is key to the result in Proposition 6. It permits us to identify “stronger markets,” in the sense of Aguirre et al. (2010), with those for matches involving agents from closer locations (part 1) or more distant locations (part 2). Once the connection between targeting and price customization is at hand, the welfare implications of customized pricing then naturally parallel those in the third-degree price discrimination literature.

Also note that the result in Proposition 6 is just an illustration of the type of welfare results that Proposition 5 permits. Paralleling the analysis in Proposition 2 in Aguirre et al. (2010), for example, we can also identify primitive conditions under which welfare is higher under price customization

¹⁶Note that $CD_a \left(\frac{dp_a}{dq} | x_b; x_a \right)$ is also the elasticity of the marginal demand $\partial D_a \left(\frac{dp_a}{dq} | x_b; x_a \right) / \partial \left(\frac{dp_a}{dq} \right)$ with respect to the marginal price $\frac{dp_a}{dq}$.

than under uniform pricing, as well as conditions under which price customization impacts negatively one side of the market and positively the other.

4.3 Centralized vs Decentralized Markets

The results in Propositions 5 and 6 also permit us to compare matching allocations and welfare in centralized markets (where matching is controlled by a platform) to their counterparts in decentralized markets (where matching is governed by agents on one side of the market posting prices for their matches with the other side).

To illustrate, suppose that side a is the “buyers” side and side b is the “sellers” side. In a decentralized market, sellers post prices for their products, or services, and their locations are observable, whereas the buyers’ locations are the buyers’ private information. Note that this amounts to assuming that the characteristics of the sellers’ products are publicly observable, whereas the buyers’ preferences for the sellers’ products are the buyers’ private information, which appears the most relevant scenario for many markets of interest. Consistently with the analysis in the rest of the paper, buyers have unit demands for the product of each seller (that is, each buyer purchases at most one unit from each seller) and the demand for each seller’s product is invariant in the price set by other sellers (that is, preferences are additively separable, as in the rest of the paper). Each seller is a local monopolist for its product; the products provided by different sellers from the same location are distinct in the eyes of the buyers, although they provide the same utility.

Now observe that the sellers’ inability to identify the buyers’ locations, along with the fact that each seller sells only one product, implies that, in a decentralized market, sellers are unable to engage in price customization on side a . Paralleling the analysis preceding Proposition 5, one should then expect that whether targeting is higher under a decentralized or a centralized market depends on the supermodularity/submodularity and concavity/convexity of the utility functions as well as on the marginal distributions of the vertical dimensions on both sides of the market. Indeed, the comparison is very similar to the one in Proposition 5 when one replaces the assumption that the environment satisfies the properties of Scenario (i) with the assumption that it satisfies the properties of Scenario (ii) (i.e., locations are private information on side a but public on side b), Condition I_k , $k = a, b$, with Conditions I_a and S_b (recall that these are the conditions that guarantee that, in a centralized market, demands under the profit-maximizing plans have a threshold structure), and finally the assumption that the vertical dimensions are private information on both sides with the assumption that they are private on side a but public on side b . In this case, results analogous to those in Propositions 5 and 6 apply to the comparison between centralized and decentralized markets. To see this, it suffices to note that the allocations and prices in decentralized markets coincide with those implemented in a centralized market when the platform is constrained to price uniformly on side a .

When, instead, the vertical dimensions are private information on both sides of the market, decentralization also brings a reduction in the prices set by the sellers, due to the elimination of

the monopsony mark-up applied by the platform under asymmetric information. Combining this extra effect with the ones identified in Propositions 5 and 6 permits us to establish the result in the following proposition.

Proposition 7. (*centralized vs decentralized markets*) *Suppose that, in addition to Condition R, the environment satisfies the properties of Scenario (ii) and Conditions VD, I_a , S_b , and IR hold. The following statements are true.*

1. *Suppose the assumptions in part 1 of Proposition 5 hold and, for any dp_a/dq and x_b , the convexity $CD_a \left(\frac{dp_a}{dq} | x_b; x_a \right)$ of the demands by the x_a -agents for the q -th unit of the x_b -agents declines with the distance $|x_a - x_b|$. Then welfare of the side- a agents is higher in a decentralized market than in a centralized one.*

2. *Suppose the assumptions in part 2 of Proposition 5 hold and, for any dp_a/dq and x_b , the convexity $CD_a \left(\frac{dp_a}{dq} | x_b; x_a \right)$ of the demands by the x_a -agents for the q -th unit of the x_b -agents increases with the distance $|x_a - x_b|$. Then, again, welfare of the side- a agents is higher in a decentralized market than in a centralized one.*

The proposition thus identifies primitive conditions under which the transition from a centralized market where matching is mediated to a decentralized one where sellers post prices and matching is un-mediated increases consumer surplus (i.e., the welfare of the side- a agents). The result may help policymakers identify markets where the transition to a decentralized structure should be promoted, for example through fiscal incentives and/or direct subsidies.

5 Concluding Remarks

This paper studies many-to-many matching in markets in which agents' preferences are both vertically and horizontally differentiated. The analysis delivers the following results. First, it identifies primitive conditions under which profit-maximizing platforms engage in price customization. That is, they offer agents the possibility to customize their matching set by including partners of different profiles based on their horizontal preferences, with the price for such customization varying with the configuration of the baseline plan. We show that the optimal tariffs induce negative assortative matching at the margin. As the matching sets expand, the marginal agents from each location included in the set are always those with the lowest value for matching. The composition of the pool of marginal agents, however, naturally respects horizontal differences in preferences, with most of the marginal agents coming from "locations" close to the subscriber's bliss-point. We then provide a formula relating the optimal prices to location-specific elasticities of the demands on both sides of the market that can be used in empirical work for testing and structural estimation and that permits us to study how the distortions in the provision of matching services vary with the horizontal dimension of the agents' preferences.

Second, the paper studies the effects on prices, the composition of matching sets, and welfare,

of uniform-pricing obligations aimed at dissuading platforms from offering menus of matching plans where the prices for different matching sets depend directly, or indirectly, on the subscribers' horizontal dimensions.

Finally, the paper contributes to the policy debate about the desirability of mediation in matching markets by offering a new angle relating the welfare effects of decentralization to the targeting and price-customization implications of different market structures.

We believe the results have useful implications for various markets. Consider, for example, ad-exchanges (see, among others, Bonatti and Bergemann (2011) for an overview of such markets). As mentioned in the introduction, platforms such as AppNexus, AOL's Marketplace, Microsoft Ad Exchange, OpenX, Rubicon Project Exchange, and Smaato, use sophisticated pricing algorithms where prices depend not only on volumes but on advertisers' and publishers' profiles. Such algorithms thus permit ad exchanges to engage in price-customization practices that appear similar, at least in spirit, to those studied in the present paper. While such algorithms have initially been praised for the customization possibilities they offer, more recently they have been associated with targeting and price-discriminatory practices believed to be detrimental to welfare. The policy debate about the merits of such algorithms and, related, about the desirability of regulations imposing uniform pricing, lacks a formal model shedding light on how platforms will respond to the introduction of such regulations. Our paper contributes a stylized but flexible framework that one can use to study both the distortions associated with price customization as well as platforms' response to the imposition of uniform pricing.

Next, consider the market for cable TV. Most providers price discriminate on the viewer side by offering viewers packages of channels whereby the baseline configuration can be customized by adding channels at a cost that depends on the baseline configuration originally selected (see, among others, Crawford (2000), and Crawford and Yurukoglu (2012)). For example, in the US, Direct TV offers various vertically differentiated (i.e., nested) packages (both in English and in Spanish). It then allows viewers to add to these packages various (horizontally differentiated) premium packages, which bundle together channels specialized in movies, sports, news, and games. In addition, viewers can further customize the packages by adding individual sports, news, and movie channels. Similar combinations of packages with different degrees of horizontal and vertical differentiation are offered by other providers. This form of price customization appears highly consistent with what predicted by our model. Such industry is evolving fast and many analysts predict a transition to a market structure whereby viewers purchase broadcasting directly from the channels, thus bypassing the intermediation of the current providers. Our analysis sheds some light on how prices set by individual channels compare to their counterparts in markets where the interactions between the channels and the viewers are mediated by cable companies and identifies conditions under which the transition to a decentralized structure is advantageous to the viewers.

We conclude by discussing a few limitations of our analysis and venues for future research. First, our analysis abstracts from platform competition. Second, and related, it assumes platform have

the power to set prices on both sides of the market. While these assumptions are a natural starting point, there are many markets where multiple platforms compete on multiple sides and their ability to set prices is hindered by their lack of bargaining power. For example, the market for cable TV is populated by multiple providers. Furthermore, as indicated in Crawford and Yurukoglu (2012), large channel conglomerates enjoy nontrivial bargaining power vis-a-vis cable TV providers, which suggests that prices are likely to be negotiated on the channel side instead of being set directly by the platforms. Extending the analysis to accommodate for platform competition and limited bargaining power on one, or multiple, sides of the market will provide further insights on the bundling and pricing strategies of many platforms.

Furthermore, certain platforms, most notably B2B platforms, have recently expanded their services to include e-billing and supply-management. These additional services open the door to more sophisticated price-discriminatory practices that use instruments other than the composition of the matching sets. Extending the analysis to accommodate for richer instruments is another interesting direction for future research (see, e.g., Jeon, Kim, and Menicucci (2016)).

Lastly, in future work, it seems desirable to extend the analysis to accommodate for "within-side" network effects (e.g., congestion and limited attention) and dynamics.

6 Appendix

Proof of Proposition 1. We establish the result using mechanism design techniques. Let

$$(\mathbf{s}_k(\theta_k), p_k(\theta_k))_{\theta_k \in \Theta_k}^{k=a,b}$$

denote a direct revelation mechanism, where agents are asked to report their types and where $(\mathbf{s}_k(\theta_k), p_k(\theta_k))$ denotes the allocation (matching set and total transfer) specified by the mechanism for each side- k agent reporting θ_k .

By familiar envelope arguments, a necessary condition for each type $\theta_k = (v_k, x_k) \in \Theta_k$, $k = a, b$, to prefer reporting truthfully to lying with respect to the vertical dimension v_k while reporting truthfully the horizontal dimension x_k is that transfers satisfy the envelope conditions

$$p_k(\theta_k) = \int_{\mathbf{s}_k(\theta_k)} u_k(v_k, |x_k - x_l|) dF_l(\theta_l) - \int_{\underline{v}_k}^{v_k} \int_{\mathbf{s}_k(y, x_k)} \frac{\partial u_k}{\partial v}(y, |x_k - x_l|) dF_l(\theta_l) dy, \quad (17)$$

$$- U_k(\underline{v}_k, x_k),$$

where $U_k(\underline{v}_k, x_k)$ is the payoff of a side- k agent with type (\underline{v}_k, x_k) .

Using (17), the platform's profits under any incentive-compatible mechanism can then be written as

$$\sum_{k=a,b} \int_{\Theta_k} \left\{ \int_{\mathbf{s}_k(\theta_k)} \left[u_k(v_k, |x_k - x_l|) - \frac{1 - F_k^v(v_k | x_k)}{f_k^v(v_k | x_k)} \cdot \frac{\partial u_k}{\partial v}(v_k, |x_k - x_l|) \right] dF_l(\theta_l) \right\} dF_k(\theta_k).$$

Using the definition of the virtual-value functions $\varphi_k(\theta_k, \theta_l)$ in the main text, we then have that the platform's profits are maximal when $U_k(\underline{v}_k, x_k) = 0$ for all $x_k \in X_k$, $k = a, b$, and when the matching sets are chosen so as to maximize

$$\sum_{k=a,b} \int_{\Theta_k} \left\{ \int_{\mathbf{s}_k(\theta_k)} \varphi_k(\theta_k, \theta_l) dF_l(\theta_l) \right\} dF_k(\theta_k) \quad (18)$$

subject to the reciprocity condition

$$\theta_l \in \mathbf{s}_k(\theta_k) \iff \theta_k \in \mathbf{s}_l(\theta_l), \quad l, k \in \{a, b\}, k \neq l. \quad (19)$$

Hereafter, we first describe the matching sets that maximize (18) subject to the above reciprocity condition and then show that, under the assumptions in the proposition, the platform can implement the allocations $(\mathbf{s}_k(\theta_k), p_k(\theta_k))_{\theta_k \in \Theta_k}^{k=a,b}$, where the functions $s_k(\cdot)$ are those that maximize (18) subject to (19), and where the functions $p_k(\cdot)$ are as in (17), with $U_k(\underline{v}_k, x_k) = 0$, all $x_k \in X_k$, $k = a, b$.

Define the indicator function $m_k(\theta_k, \theta_l) \in \{0, 1\}$ taking value one if and only if $\theta_l \in \mathbf{s}_k(\theta_k)$, that is, if and only if the two types θ_k and θ_l are matched. Then define the following measure on the Borel sigma-algebra over $\Theta_k \times \Theta_l$:

$$\nu_k(E) \equiv \int_E m_k(\theta_k, \theta_l) dF_k(\theta_k) dF_l(\theta_l). \quad (20)$$

Reciprocity implies that $m_k(\theta_k, \theta_l) = m_l(\theta_l, \theta_k)$. As a consequence, the measures ν_k and ν_l satisfy $d\nu_k(\theta_k, \theta_l) = d\nu_l(\theta_l, \theta_k)$. Equipped with this notation, the expression in (18) can be rewritten as

$$\begin{aligned} & \sum_{k,l=a,b, l \neq k} \int_{\Theta_k \times \Theta_l} \varphi_k(\theta_k, \theta_l) d\nu_k(\theta_k, \theta_l) \\ &= \int_{\Theta_k \times \Theta_l} \Delta_k(\theta_k, \theta_l) m_k(\theta_k, \theta_l) dF_k(\theta_k) dF_l(\theta_l), \end{aligned} \quad (21)$$

where, for $k, l = a, b$, $l \neq k$,

$$\Delta_k(\theta_k, \theta_l) \equiv \varphi_k(\theta_k, \theta_l) + \varphi_l(\theta_l, \theta_k).$$

Note that the functions $\Delta_a(\theta_a, \theta_b) = \Delta_b(\theta_b, \theta_a)$ represent the marginal effects on the platform's profits of matching types θ_a and θ_b . It is then immediate that the rule $(m_k(\cdot))_{k=a,b}$ that maximizes the expression in (21) is such that, for any $(\theta_k, \theta_l) \in \Theta_k \times \Theta_l$, $k, l = a, b$, $l \neq k$, $m_k(\theta_k, \theta_l) = 1$ if and only if

$$\Delta_k(\theta_k, \theta_l) \geq 0.$$

Next, observe that, under Condition R, the function

$$\varphi_k(\theta_k, \theta_l) \equiv u_k(v_k, |x_k - x_l|) - \frac{1 - F_k^v(v_k | x_k)}{f_k^v(v_k | x_k)} \cdot \frac{\partial u_k}{\partial v}(v_k, |x_k - x_l|)$$

is strictly increasing in v_k , $k, l = a, b$, $l \neq k$. We conclude that the matching rule that maximizes (18) subject to the reciprocity condition (19) can be described by means of a collection of threshold

functions $t_k^* : \Theta_k \times X_l \rightarrow V_l$, $k, l = a, b$, $l \neq k$, such that, for any $\theta_k = (v_k, x_k)$, any $\theta_l = (v_l, x_l)$, $\theta_l \in \mathbf{s}_k(\theta_k)$ if, and only if, $v_l \geq t_k^*(\theta_k, x_l)$. The threshold functions $t_k^*(\cdot)$ are such that, for any $\theta_k \in \Theta_k$, any $x_l \in [0, 1]$, $t_k^*(\theta_k, x_l) = \underline{v}_l$ if $\Delta_k(\theta_k, (\underline{v}_l, x_l)) > 0$, $t_k^*(\theta_k, x_l) = \bar{v}_l$ if $\Delta_k(\theta_k, (\bar{v}_l, x_l)) < 0$, and $t_k^*(\theta_k, x_l)$ is the unique solution to $\Delta_k(\theta_k, (t_k^*(\theta_k, x_l), x_l)) = 0$ if

$$\Delta_k(\theta_k, (\underline{v}_l, x_l)) \leq 0 \leq \Delta_k(\theta_k, (\bar{v}_l, x_l)).$$

Condition R also implies that, for any $x_k, x_l \in [0, 1]^2$, the threshold $t_k^*(\theta_k, x_l)$ is decreasing in v_k . Finally, if, in addition to Condition R, Condition VD also holds (in which case, the the virtual values φ_k are non-increasing in $|x_k - x_l|$, $k = a, b$), for any $\theta_k = (v_k, x_k)$, the threshold functions $t_k^*(\theta_k, x_l)$ and non-decreasing in $|x_l - x_k|$.

Equipped with the above result, we now show that, in each of the environments described by the conditions in the proposition, the platform can implement the allocations $(\mathbf{s}_k(\theta_k), p_k(\theta_k))_{\theta_k \in \Theta_k}^{k=a,b}$, where $\mathbf{s}_k(\theta_k)$ are the matching sets described by the above threshold rule, and where the payment functions $p_k(\theta_k)$ are the ones in (17), with $U_k(\underline{v}_k, x_k) = 0$, all $x_k \in X_k$, $k = a, b$.

First observe that the payoff that each type θ_k obtains in the above direct revelation mechanism when reporting truthfully is equal to

$$U_k(\theta_k) = \int_{\underline{v}_k}^{v_k} \int_{\mathbf{s}_k(y, x_k)} \frac{\partial u_k}{\partial v}(y, |x_k - x_l|) dF_l(\theta_l) dy.$$

That $U_k(\theta_k) \geq 0$ follows directly from the fact that u_k is non-decreasing in v_k . This means that the mechanism is individually rational (meaning that each type θ_k prefers participating in the mechanism and receiving the allocation $(\mathbf{s}_k(\theta_k), p_k(\theta_k))$ to refusing to participate and receiving the allocation $(\emptyset, 0)$ yielding a payoff equal to zero).

Below we show that either the above direct mechanism is also incentive-compatible (meaning that each type θ_k prefers the allocation $(\mathbf{s}_k(\theta_k), p_k(\theta_k))$ designed for him to the allocation $(\mathbf{s}_k(\theta'_k), p_k(\theta'_k))$ designed for any other type θ'_k), or it can be turned, at no cost to the platform, into a mechanism implementing the same allocations as the above ones which is both incentive compatible and individually rational.

Definition 6. (nested matching) A matching rule $\mathbf{s}_k(\theta_k)$ is nested if, for any pair $\theta_k = (v_k, x_k)$ and $\hat{\theta}_k = (\hat{v}_k, \hat{x}_k)$ such that $x_k = \hat{x}_k$, either $\mathbf{s}_k(\theta_k) \subseteq \mathbf{s}_k(\hat{\theta}_k)$, or $\mathbf{s}_k(\theta_k) \supseteq \mathbf{s}_k(\hat{\theta}_k)$. A direct revelation mechanism is nested if its matching rule is nested.

Clearly, the direct mechanism defined above where the matching rule is described by the threshold function $t_k^*(\theta_k, x_l)$ is nested. Now let $\Pi_k(\theta_k; \hat{\theta}_k)$ denote the payoff that type θ_k obtains in a direct revelation mechanism $(\mathbf{s}_k(\theta_k), p_k(\theta_k))_{\theta_k \in \Theta_k}^{k=a,b}$ by mimicking type $\hat{\theta}_k$.

Definition 7. (ICV) A direct revelation mechanism $(\mathbf{s}_k(\theta_k), p_k(\theta_k))_{\theta_k \in \Theta_k}^{k=a,b}$ satisfies incentive compatibility along the v dimension (ICV) if, for any $\theta_k = (v_k, x_k)$ and $\hat{\theta}_k = (\hat{v}_k, \hat{x}_k)$ with $x_k = \hat{x}_k$, $U_k(\theta_k) \geq \Pi_k(\theta_k; \hat{\theta}_k)$.

The following results are then true (the proof is standard and hence omitted):

Lemma 1. *A nested direct revelation mechanism $(\mathbf{s}_k(\theta_k), p_k(\theta_k))_{\theta_k \in \Theta_k}^{k=a,b}$ satisfies ICV if, and only if, the following conditions jointly hold:*

1. *for any $\theta_k = (v_k, x_k)$ and $\hat{\theta}_k = (\hat{v}_k, \hat{x}_k)$ such that $x_k = \hat{x}_k$, $v_k > \hat{v}_k$ implies that $\mathbf{s}_k(\theta_k) \supseteq \mathbf{s}_k(\hat{\theta}_k)$;*
2. *the payment functions $p_k(\theta_k)$ satisfy the envelope formula (17).*

Clearly, the direct revelation mechanism where the matching rule is the one corresponding to the threshold functions $t_k^*(\cdot)$ described above and where the payment functions $p_k(\theta_k)$ are the ones in (17), with $U_k(\underline{v}_k, x_k) = 0$, all $x_k \in X_k$, $k = a, b$, is not only nested but satisfies the two conditions in the lemma. It follows that such a mechanism satisfies ICV.

Equipped with the above results, we now show that, in each of the environments corresponding to the combination of conditions described in the proposition, the above direct revelation mechanism is either incentive-compatible, or it can be augmented to implement the same allocations prescribed by $(\mathbf{s}_k(\theta_k), p_k(\theta_k))_{\theta_k \in \Theta_k}^{k=a,b}$ at no extra cost to the platform.

Consider first Scenario (i). Recall that, in this case, locations are public on both sides. That the mechanism is ICV implies that any deviation along the vertical dimension is unprofitable. Furthermore, because locations are public on both sides, any deviation along the horizontal dimension is detectable. It is then immediate that the platform can augment the above direct revelation mechanism by adding to it punishments (in the form of large fines) for those agents lying along the horizontal dimension. The augmented mechanism is both individually rational and incentive compatible and implements the same allocations as the original mechanism $(\mathbf{s}_k(\theta_k), p_k(\theta_k))_{\theta_k \in \Theta_k}^{k=a,b}$ at no extra cost to the platform.

Next suppose the environment satisfies the properties of Scenario (ii) and, in addition, Conditions I_a , S_b , and VD hold. Again, because locations are public on side b , incentive compatibility on side b can be guaranteed by augmenting the mechanism as described above for Scenario (i). Thus consider incentive compatibility on side a . The latter requires that

$$U_a(v_a, x_a) \geq \Pi_a((v_a, x_a); (\hat{v}_a, \hat{x}_a)),$$

for all $(x_a, \hat{x}_a, v_a, \hat{v}_a) \in X_a^2 \times V_a^2$. The above inequality is equivalent to

$$\begin{aligned} \int_{\underline{v}_a}^{v_a} \int_{\mathbf{s}_a(y, x_a)} \frac{\partial u_a}{\partial v}(y, |x_a - x_b|) dF_b(\theta_b) dy &\geq \int_{\underline{v}_a}^{\hat{v}_a} \int_{\mathbf{s}_a(y, \hat{x}_a)} \frac{\partial u_a}{\partial v}(y, |\hat{x}_a - x_b|) dF_b(\theta_b) dy \\ &+ \int_{\mathbf{s}_a(\hat{v}_a, \hat{x}_a)} [u_a(v_a, |x_a - x_b|) - u_k(\hat{v}_a, |\hat{x}_a - x_b|)] dF_b(b). \end{aligned} \quad (22)$$

It is easy to see that, for any $\theta_a = (v_a, x_a) \in \Theta_a$,

$$\int_{\mathbf{s}_a(v_a, x_a)} \frac{\partial u_a}{\partial v}(v_a, |x_a - x_b|) dF_b(\theta_b) = \int_{d \in [0, 1/2]} \frac{\partial u_a(v_a, d)}{\partial v} dW(d; \theta_a), \quad (23)$$

where $W(d; \theta_a)$ is the measure of agents whose distance from x_a is at most d included in the matching set $\mathbf{s}_a(v_a, x_a)$ of type θ_a under the proposed mechanism. It is also easy to see that, under Conditions

I_a and S_b , the expression in (23) is invariant in x_a . That is, $W(d; \theta_a) = W(d; \theta'_a)$ for any $d \in [0, 1/2]$, any $\theta_a, \theta'_a \in \Theta_a$ with $v_a = v'_a$.¹⁷ This means that

$$\int_{\underline{v}_a}^{\hat{v}_a} \int_{\mathbf{s}_a(y, \hat{x}_a)} \frac{\partial u_a}{\partial v}(y, |\hat{x}_a - x_b|) dF_b(\theta_b) dy = \int_{\underline{v}_a}^{\hat{v}_a} \int_{\mathbf{s}_a(y, x_a)} \frac{\partial u_a}{\partial v}(y, |x_a - x_b|) dF_b(\theta_b) dy.$$

By the same arguments,

$$\int_{\mathbf{s}_a(\hat{v}_a, \hat{x}_a)} u_a(\hat{v}_a, |\hat{x}_a - x_b|) dF_b(\theta_b) = \int_{\mathbf{s}_a(\hat{v}_a, x_a)} u_a(\hat{v}_a, |x_a - x_b|) dF_b(\theta_b).$$

Furthermore, under condition VD, the threshold functions $t_k^*(\theta_k, x_l)$ and non-decreasing in the distance $|x_l - x_k|$. In turn, this implies that

$$\int_{\mathbf{s}_a(\hat{v}_a, \hat{x}_a)} u_a(v_a, |x_a - x_b|) dF_b(\theta_b) \leq \int_{\mathbf{s}_a(\hat{v}_a, x_b)} u_a(v_a, |x_b - x_a|) dF_b(\theta_b).$$

It follows that the right hand side of (22) is smaller than

$$\begin{aligned} & \int_{\underline{v}_a}^{\hat{v}_a} \int_{\mathbf{s}_a(y, x_a)} \frac{\partial u_k}{\partial v}(y, |x_a - x_b|) dF_b(\theta_b) dy \\ & + \int_{\mathbf{s}_a(\hat{v}_a, x_a)} [u_a(v_a, |x_a - x_b|) - u_a(\hat{v}_a, |x_a - x_b|)] dF_b(\theta_b), \end{aligned}$$

which is the payoff that type $\theta_a = (v_a, x_a)$ obtains by announcing (\hat{v}_a, x_a) (that is, by lying about the vertical dimension but reporting truthfully the horizontal one). That the inequality in (22) holds then follows from the fact that the direct revelation mechanism $(\mathbf{s}_k(\theta_k), p_k(\theta_k))_{\theta_k \in \Theta_k}^{k=a,b}$ satisfies ICV.

The arguments for an environment satisfying the properties of Scenario (iii) along with Conditions VD, I_b and S_a are symmetric to those for an environment satisfying the properties of Scenario (ii) along with Conditions VD, I_a and S_b , and hence the proof is omitted.

Finally, consider an environment satisfying the properties of Scenario (iv) along with Conditions VD, S_a and S_b . That the proposed mechanism is incentive compatible follows from the same arguments as for Scenario (ii) above, now applied to both sides of the market.

We conclude that, in each of the environments considered in the proposition, the allocations $(\mathbf{s}_k(\theta_k), p_k(\theta_k))_{\theta_k \in \Theta_k}^{k=a,b}$, where the matching sets $\mathbf{s}_k(\theta_k)$ are the ones specified by the threshold functions $t_k^*(\cdot)$ described above, and where the payments are the ones in (17) with $U_k(\underline{v}_k, x_k) = 0$, all $x_k \in X_k$, $k = a, b$ can be sustained in a mechanism that is both individually rational and incentive compatible. The result in the proposition then follows from the fact that (a) such allocations are profit-maximizing among those consistent with the rationality of the agents (i.e., satisfying the IC and IR constraints), and (b) can be induced by offering customized tariffs

$$\{(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathbf{S}_k(x_k)) : x_k \in [0, 1]\}$$

¹⁷Conditions I_k , $k = a, b$, suffice to guarantee that the function $\Delta_k(\theta_k, \theta_l)$ depends only on v_k, v_l , and $|x_l - x_k|$. The strengthening of Condition I_b to S_b is, however, necessary to guarantee that the mass of agents of a given distance d included in the matching sets of any pair of types $\theta_a, \theta'_a \in \Theta_a$ with $v_a = v'_a$ is the same.

satisfying the properties described below. For each plan $x_k \in [0, 1]$, the baseline configuration is given by

$$\underline{\mathbf{s}}_k(x_k) = \mathbf{s}_k(\underline{v}_k, x_k),$$

the baseline price is given by

$$\underline{T}_k(x_k) = p_k(\underline{v}_k, x_k) = \int_{\underline{\mathbf{s}}_k(\underline{v}_k, x_k)} u_k(\underline{v}_k, |x_k - x_l|) dF_l(\theta_l),$$

the set of possible customizations is given by

$$\mathbf{S}_k(x_k) = \{\mathbf{s}_k(v_k, x_k) : v_k \in V_k\},$$

and the price schedules $\rho_k(q|x_l; x_k)$ are such that, for $q = q_{x_l}(\mathbf{s}_k(\underline{v}_k, x_k))$, $\rho_k(q|x_l; x_k) = 0$, while for $q \in (q_{x_l}(\mathbf{s}_k(\underline{v}_k, x_k)), q_{x_l}(\mathbf{s}_k(\bar{v}_k, x_k))]$,

$$\rho_k(q|x_l; x_k) = q u_k(v_k(q; x_k, x_l), |x_k - x_l|) - \int_{\underline{v}_k}^{v_k(q; x_k, x_l)} q_{x_l}(\mathbf{s}_k(y, x_k)) \frac{\partial u_k}{\partial v}(y, |x_k - x_l|) dy - \underline{T}_k(x_k) \quad (24)$$

where

$$v_k(q; x_k, x_l) = \inf \{v_k \in V_k : q_{x_l}(\mathbf{s}_k(v_k, x_k)) = q\}.$$

Any agent selecting the plan $(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathbf{S}_k(x_k))$ and then choosing a matching set $\mathbf{s}_k \notin \mathbf{S}_k(x_k)$ is charged a fine large enough to make the utility of such a set, net of the payment, negative for all types. Likewise, when locations are public on side k , any side- k agent selecting a plan other than $(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathbf{S}_k(x_k))$ is charged a large enough fine to make the choice unprofitable for any type. Note that the existence of such fines is guaranteed by the assumption that u_k is bounded, $k = a, b$.

That the above customized tariff implements the same allocations as the direct mechanism $(\mathbf{s}_k(\theta_k), p_k(\theta_k))_{\theta_k \in \Theta_k}^{k=a,b}$ then follows from the following considerations. Each type $\theta_k = (v_k, x_k)$, by selecting the plan $(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathbf{S}_k(x_k))$ designed for agents with the same location as type θ_k and then choosing the customization $\mathbf{s}_k(v_k, x_k)$ specified by the direct mechanism for type θ_k is charged a total payment equal to

$$\begin{aligned} & \underline{T}_k(x_k) + \int_0^1 \left[q_{x_l}(\mathbf{s}_k(v_k, x_k)) u_k(v_k, |x_k - x_l|) - \int_{\underline{v}_k}^{v_k} q_{x_l}(\mathbf{s}_k(y, x_k)) \frac{\partial u_k}{\partial v}(y, |x_k - x_l|) dy \right] dx_l - \underline{T}_k(x_k) \\ &= \int_{\mathbf{s}_k(\theta_k)} u_k(v_k, |x_k - x_l|) dF_l(\theta_l) - \int_{\underline{v}_k}^{v_k} \int_{\mathbf{s}_k(y, x_k)} \frac{\partial u_k}{\partial v}(y, |x_k - x_l|) dF_l(\theta_l) dy \\ &= p_k(\theta_k), \end{aligned}$$

exactly as in the direct mechanism. That each type θ_k maximizes his payoff by selecting the plan $(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathbf{S}_k(x_k))$ and then choosing the customization $\mathbf{s}_k(v_k, x_k)$ specified for him by the direct mechanism then follows from the fact that (a) the direct mechanism is incentive compatible, (b) the payment associated with any other plan $(\underline{\mathbf{s}}_k(\hat{x}_k), \underline{T}_k(\hat{x}_k), \rho_k(\cdot|\cdot; \hat{x}_k), \mathbf{S}_k(\hat{x}_k))$ followed

by the selection of a set \mathbf{s}_k is either equal to the payment specified by the direct mechanism for some report (\hat{v}_k, \hat{x}_k) , or is so large to make the net payoff of such selection negative.

Finally, to see that, when locations are public on side k , without loss of optimality, the side- k customized tariff does not need to restrict the agents' ability to customize their matching sets (that is, $\mathbf{S}_k(x_k) = \Sigma(\Theta_l)$, all x_k) recall that, in this case, each side- k agent located at x_k can be induced to select the matching plan $(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathbf{S}_k(x_k))$ designed for agents located at x_k by setting the fee associated with the selection of any other plan sufficiently high. The separability of the agents' preferences then implies that, once the plan $\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathbf{S}_k(x_k)$ is selected, even if $\mathbf{S}_k(x_k) = \Sigma(\Theta_l)$, because the price schedules $\rho_k(\cdot|\cdot; x_k)$ satisfy (24), type θ_k prefers to select $q_{x_l}(\mathbf{s}_k(v_k, x_k))$ agents from each location x_l to any other mass of agents from the same location x_l , irrespective of the mass of agents from other locations type θ_k includes in his matching set. Q.E.D.

Proof of Proposition 2. Fix a pair of locations $x_a, x_b \in [0, 1]$. From Proposition 1, the profit-maximizing tariffs are customized and induce agents to select matching sets satisfying the threshold property of Proposition 1. Furthermore, from the proof of Proposition 1, for any $\theta_k = (v_k, x_k)$, any $x_l \in [0, 1]$, the threshold t_k^* is such that $t_k^*(\theta_k, x_l) = \underline{v}_l$ if $\Delta_k(\theta_k, (\underline{v}_l, x_l)) > 0$, $t_k^*(\theta_k, x_l) = \bar{v}_l$ if $\Delta_k(\theta_k, (\bar{v}_l, x_l)) < 0$, and $t_k^*(\theta_k, x_l)$ is the unique solution to $\Delta_k(\theta_k, (t_k^*(\theta_k, x_l), x_l)) = 0$ if

$$\Delta_k(\theta_k, (\underline{v}_l, x_l)) \leq 0 \leq \Delta_k(\theta_k, (\bar{v}_l, x_l)).$$

This means that, for any $q_k \in (0, f_l^x(x_l))$, either there exists no $v_k \in V_k$ such that $q_{x_l}(\mathbf{s}_k(v_k, x_k)) = q_k$, or there exists a unique $v_k \in V_k$ such that $q_{x_l}(\mathbf{s}_k(v_k, x_k)) = q_k$. Now take any $q_k \in (0, f_l^x(x_l))$ for which there exists $v_k \in V_k$ such that $q_{x_l}(\mathbf{s}_k(v_k, x_k)) = q_k$. As explained in the main text, for any such q_k , the unique value of v_k such that $q_{x_l}(\mathbf{s}_k(v_k, x_k)) = q_k$ is also the unique value of v_k that solves

$$u_k(v_k, |x_k - x_l|) = \frac{d\rho_k}{dq}(q_k|x_l; x_k). \quad (25)$$

Now let $\hat{v}_{x_l} \left(\frac{d\rho_k}{dq}|x_k \right)$ be the unique solution to (25) and $v'_l(q_k; x_l)$ be the unique solution to

$$\left[1 - F_l^{v|x} (v'_l(q_k; x_l)|x_l) \right] f_l^x(x_l) = q_k.$$

That the demands under the profit-maximizing tariffs satisfy the threshold structure of Proposition 1 implies that

$$t_k^* \left(\left(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq}|x_k \right), x_k \right), x_l \right) = v'_l(q_k; x_l)$$

and that

$$\varphi_k \left(\left(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq}|x_k \right), x_k \right), (v'_l(q_k; x_l), x_l) \right) + \varphi_l \left(v'_l(q_k; x_l), x_l, \left(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq}|x_k \right), x_k \right) \right) = 0. \quad (26)$$

Lastly, observe that, for any such q_k ,

$$\varepsilon_k \left(\frac{d\rho_k}{dq}|x_l; x_k \right) = \frac{f_k^{v|x}(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq}|x_k \right)|x_k)}{1 - F_k^{v|x}(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq}|x_k \right)|x_k)} \left[\frac{\partial u_k}{\partial v} \left(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq}|x_k \right), |x_k - x_l| \right) \right]^{-1} \frac{d\rho_k}{dq}(q_k|x_l; x_k). \quad (27)$$

Using the definition of φ_k from the main text together with (25) and (27), we then have that, for any such q_k ,

$$\varphi_k \left(\left(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq} | x_k \right), x_k \right), (v'_l(q_k; x_l), x_l) \right) = \frac{d\rho_k}{dq} (q_k | x_l; x_k) \left[1 - \frac{1}{\varepsilon_k \left(\frac{d\rho_k}{dq} | x_l; x_k \right)} \right]. \quad (28)$$

Likewise, when $q_l = \left[1 - F_k^{v|x} \left(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq} | x_k \right) | x_k \right) \right] f_k^x(x_k)$,

$$\varphi_l \left((v'_l(q_k; x_l), x_l), \left(\hat{v}_{x_l} \left(\frac{d\rho_k}{dq} | x_k \right), x_k \right) \right) = \frac{d\rho_l}{dq} (q_l | x_k; x_l) \left[1 - \frac{1}{\varepsilon_l \left(\frac{d\rho_l}{dq} | x_k; x_l \right)} \right]. \quad (29)$$

Combining (28) and (29) with (26), we obtain the result in the proposition. Q.E.D.

Proof of Example 3. Recall that, for any (θ_a, x_b) such that $t_a^*(\theta_a, x_b) \in \text{Int}[V_b]$, $t_a^*(\theta_a, x_b)$ is defined by the unique solution to

$$\varphi_a(\theta_a, (t_a^*(\theta_a, x_b), x_b)) + \varphi_b((t_a^*(\theta_a, x_b), x_b), \theta_a) = 0.$$

When preferences are as in Example 2, with $\delta = 1$, and vertical and horizontal dimensions are independently distributed,

$$u_a(v_a, |x_a - x_b|) = \alpha \cdot v_a + (1 - \alpha) \cdot \phi(|x_a - x_b|),$$

$u_b(v_b, |x_b - x_a|) = v_b$, and, for any $x_k \in [0, 1]$, $k = a, b$,

$$\frac{1 - F_k^{v|x}(v_k | x_k)}{f_k^{v|x}(v_k | x_k)} = \frac{1 - F_k^v(v_k)}{f_k^v(v_k)} \equiv [\lambda_k^v(v_k)]^{-1}.$$

Hence, in this case,

$$\varphi_a(\theta_a, \theta_b) \equiv \alpha \cdot v_a + (1 - \alpha) \cdot \phi(|x_a - x_b|) - \frac{\alpha}{\lambda_a^v(v_a)}$$

and

$$\varphi_b(\theta_b, \theta_a) = v_b - \frac{1}{\lambda_b^v(v_b)}.$$

This means that, for any (θ_a, x_b) such that $t_a^*(\theta_a, x_b) \in \text{Int}[V_b]$, $t_a^*(\theta_a, x_b)$ is given by the unique solution to

$$\alpha \cdot v_a + (1 - \alpha) \cdot \phi(|x_a - x_b|) - \frac{\alpha}{\lambda_a^v(v_a)} + t_a^*(\theta_a, x_b) - \frac{1}{\lambda_b^v(t_a^*(\theta_a, x_b))} = 0.$$

On the other hand, the threshold functions $t_a^e(\theta_a, x_b)$ defining the matching sets under welfare maximizations are such that, for any (θ_a, x_b) such that $t_a^e(\theta_a, x_b) \in \text{Int}[V_b]$, $t_a^e(\theta_a, x_b)$ is given by the unique solutions to

$$u_a(v_a, |x_a - x_b|) + u_b(t_a^e(\theta_a, x_b), |x_b - x_a|) = 0$$

which means that $t_a^e(\theta_a, x_b) = -[\alpha v_a + (1 - \alpha)\phi(|x_a - x_b|)]$.

We conclude that, when $|K|$ is sufficiently large that, for any $(\theta_a, x_b) \in \Theta_a \times [0, 1]$, $t_a^*(\theta_a, x_b), t_a^e(\theta_a, x_b) \in \text{Int}[V_b]$,

$$t_a^*(\theta_a, x_b) - t_a^e(\theta_a, x_b) = \frac{1}{\lambda_b^v(t_a^*(\theta_a, x_b))} + \frac{\alpha}{\lambda_a^v(v_a)}.$$

Using the fact that, when vertical types are drawn from an exponential distribution on side b ,

$$\frac{1}{\lambda_b^v(t_a^*(\theta_a, x_b))} = \frac{1}{\tilde{\lambda}_b}$$

we then obtain the result in the example. Q.E.D.

Proof of Proposition 3. Take any $(\theta_k, x_l) \in \Theta_k \times [0, 1]$ for which $t_k^*(\theta_k, x_l), t_k^e(\theta_k, x_l) \in \text{Int}[V_l]$. Recall that, in this case, $t_k^*(\theta_k, x_l)$ is given by the unique solution to

$$\varphi_k(\theta_k, (t_k^*(\theta_k, x_l), x_l)) + \varphi_l((t_k^*(\theta_k, x_l), x_l), \theta_k) = 0$$

whereas $t_k^e(\theta_k, x_l)$ is given by the unique solution to

$$u_k(v_k, |x_k - x_l|) + u_l(t_k^e(\theta_k, x_l), |x_l - x_k|) = 0.$$

This means that

$$\begin{aligned} & u_l(t_k^*(\theta_k, x_l), |x_l - x_k|) - u_l(t_k^e(\theta_k, x_l), |x_l - x_k|) \\ &= \frac{1 - F_k^v(v_k)}{f_k^v(v_k)} \frac{\partial u_k}{\partial v_k}(v_k, |x_k - x_l|) + \frac{1 - F_l^v(t_k^*(\theta_k, x_l))}{f_l^v(t_k^*(\theta_k, x_l))} \frac{\partial u_l}{\partial v_l}(t_k^*(\theta_k, x_l), |x_k - x_l|). \end{aligned} \tag{30}$$

Now fix $\theta_k = (v_k, x_k)$, let $\mu(x_k, x_l) = |x_k - x_l|$ and, without loss of generality, assume $x_l > x_k$ (the results for the case $x_l < x_k$ are analogous to those for the case $x_l > x_k$, after the obvious change in sign due to the fact that, in this case, increasing distance means decreasing x_l). Differentiating the expression in the right-hand-side of (30) with respect to x_l at some x_l for which $\mu(x_k, x_l) \in (0, \frac{1}{2})$, we obtain that

$$\begin{aligned} \frac{\partial}{\partial x_l} [u_l(t_k^*(\theta_k, x_l), |x_l - x_k|) - u_l(t_k^e(\theta_k, x_l), |x_l - x_k|)] &= \frac{1 - F_k^v(v_k)}{f_k^v(v_k)} \cdot \frac{\partial^2 u_k(v_k, |x_k - x_l|)}{\partial v \partial \mu} \\ &\quad - \frac{(\lambda_l^v)'(t_k^*(\theta_k, x_l))}{\lambda_l^v(t_k^*(\theta_k, x_l))^2} \frac{\partial t_k^*}{\partial x_l}(\theta_k, x_l) \frac{\partial u_l(t_k^*(\theta_k, x_l), |x_k - x_l|)}{\partial v_l} \\ &\quad + \frac{1}{\lambda_l^v(t_k^*(\theta_k, x_l))} \left(\frac{\partial^2 u_l(t_k^*(\theta_k, x_l), |x_k - x_l|)}{\partial v_l^2} \frac{\partial t_k^*(\theta_k, x_l)}{\partial x_l} + \frac{\partial^2 u_l(t_k^*(\theta_k, x_l), |x_k - x_l|)}{\partial v_l \partial \mu} \right). \end{aligned}$$

By Proposition 1, $\frac{\partial t_k^*}{\partial x_l}(\theta_k, x_l) \geq 0$. Therefore, under the assumptions of part 1, the expression above is negative. Conversely, under the assumptions of part 2, the expression above is positive. The result in the proposition then follows from the above properties. Q.E.D.

Proof of Proposition 4. The platform's problem consists in choosing a collection of side- a uniform price schedules $p_a(\cdot|x_b)$, one for each side- b location $x_b \in [0, 1]$, along with a collection of side- b price schedules $\rho_b(\cdot|x_a; x_b)$, one for each pair $(x_a, x_b) \in [0, 1]^2$, that jointly maximize its profits, which can be conveniently expressed as

$$\int_0^1 \int_0^{f_b^x(x_b)} \bar{D}_a \left(\frac{dp_a}{dq}(q|x_b)|x_b \right) \frac{dp_a}{dq}(q|x_b) dq dx_b \\ + \int_0^1 \int_0^1 \int_0^{f_a^x(x_a)} D_b \left(\frac{d\rho_b}{dq}(q|x_a; x_b)|x_a; x_b \right) \frac{d\rho_b}{dq}(q|x_a; x_b) dq dx_a dx_b,$$

subject to the feasibility constraint (3).

For any x_b , $q \leq f_b^x(x_b)$, and $\frac{dp_a}{dq}(q|x_b)$, let

$$\hat{v}_{x_b} \left(\frac{dp_a}{dq} | x_a \right) = \begin{cases} v_a \text{ s.t. } u_a(v_a, |x_a - x_b|) = \frac{dp_a}{dq} & \text{if } \frac{dp_a}{dq} \in [u_a(\underline{v}_a, |x_a - x_b|), u_a(\bar{v}_a, |x_a - x_b|)] \\ \underline{v}_a & \text{if } \frac{dp_a}{dq} < u_a(\underline{v}_a, |x_a - x_b|) \\ \bar{v}_a & \text{if } \frac{dp_a}{dq} > u_a(\bar{v}_a, |x_a - x_b|). \end{cases} \quad (31)$$

Given the above definition, we have that the demand by the x_a -agents for the q -th unit of the x_b -agents at the marginal price $\frac{dp_a}{dq}(q|x_b)$ is equal to

$$D_a \left(\frac{dp_a}{dq}(q|x_b)|x_b; x_a \right) = \left[1 - F_a^{v|x} \left(\hat{v}_{x_b} \left(\frac{dp_a}{dq} | x_a \right) | x_a \right) \right] f_a^x(x_a).$$

Also, for any $q \leq f_b^x(x_b)$, recall that $v'_b(q; x_b)$ is the unique solution to $\left[1 - F_b^{v|x} (v'_b(q; x_b)|x_b) \right] f_b^x(x_b) = q$. Reciprocity, along with optimality, implies that the most profitable way to deliver q units of x_b -agents to each x_a -agent demanding to be matched to q units of x_b -agents is to match the x_a -agent to every x_b -agent whose vertical type exceeds $v'_b(q; x_b)$. In other words, the optimal tariffs induce matching demands with a threshold structure, as in the case where tariffs are customized on both sides of the market (cfr Proposition 1). Now for each $x_a, x_b \in [0, 1]$, each $q \leq f_b^x(x_b)$, let

$$\hat{q}_b(q; x_a; x_b) \equiv D_a \left(\frac{dp_a}{dq}(q|x_b)|x_b; x_a \right).$$

Given $\frac{dp_a}{dq}(q|x_b)$, the platform thus optimally selects customized prices for the x_b -agents for each quantity $\hat{q}_b(q; x_a; x_b)$ of the x_a -agents equal to

$$\frac{d\rho_b}{dq}(\hat{q}_b(q; x_a; x_b)|x_a; x_b) = u_b(v'_b(q; x_b), |x_b - x_a|). \quad (32)$$

Such prices guarantee that, for each $x_a \in [0, 1]$, $D_b \left(\frac{d\rho_b}{dq}(\hat{q}_b(q; x_a; x_b)|x_a; x_b)|x_a; x_b \right) = q$, thus clearing the market.

The function $\frac{dp_a}{dq}(q|x_b) : \mathbb{R} \times [0, 1] \rightarrow \mathbb{R}$ thus uniquely defines the matching sets on both sides of the market. Now, from the arguments in the proof of Proposition 1, we know that the maximal revenue the platform receives from the side- b agents when each x_b -agent with vertical type v_b is assigned a matching set equal to $\mathbf{s}_b(v_b, x_b)$ is given by

$$\int_{\Theta_b} \left\{ \int_0^1 \left\{ u_b(v_b, |x_b - x_a|) - \frac{1 - F_b^{v|x}(v_b|x_b)}{f_b^{v|x}(v_b|x_b)} \cdot \frac{\partial u_b}{\partial v}(v_b, |x_b - x_a|) \right\} q_{x_a}(\mathbf{s}_b(v_b, x_b)) dx_a \right\} dF_b(\theta_b).$$

In turn, this means that the platform's problem can be re-casted as choosing a function $\frac{dp_a}{dq}(q|x_b) : \mathbb{R} \times [0, 1] \rightarrow \mathbb{R}$ that maximizes

$$\int_0^1 \int_0^{f_b^x(x_b)} \left\{ \bar{D}_a \left(\frac{dp_a}{dq}(q|x_b)|x_b \right) \frac{dp_a}{dq}(q|x_b) - \mathcal{C} \left[\frac{dp_a}{dq}(q|x_b) \right] \right\} dq dx_b$$

where, for each $x_b \in [0, 1]$, each $q \leq f_b^x(x_b)$, the function

$$\begin{aligned} \mathcal{C} \left[\frac{dp_a}{dq}(q|x_b) \right] \equiv \\ - \int_0^1 \left\{ u_b(v'_b(q; x_b), |x_b - x_a|) - \frac{1 - F_b^{v|x}(v'_b(q; x_b)|x_b)}{f_b^{v|x}(v'_b(q; x_b)|x_b)} \cdot \frac{\partial u_b}{\partial v}(v'_b(q; x_b), |x_b - x_a|) \right\} D_a \left(\frac{dp_a}{dq}(q|x_b)|x_b; x_a \right) dx_a \end{aligned}$$

captures the ‘‘procurement costs’’ of clearing the matching demands of all side- a agents that demand at least q matches with the x_b -agents. This problem can be solved by point-wise maximization of the above objective function, i.e., by selecting for each $x_b \in [0, 1]$, $q \leq f_b^x(x_b)$ (equivalently, for each $(x_b, v_b) \in [0, 1] \times V_b$), $\frac{dp_a}{dq}(q|x_b)$ so as to maximize

$$\bar{D}_a \left(\frac{dp_a}{dq}(q|x_b)|x_b \right) \frac{dp_a}{dq}(q|x_b) - \mathcal{C} \left[\frac{dp_a}{dq}(q|x_b) \right].$$

The first-order conditions for such a problem are given by

$$\frac{dp_a}{dq}(q|x_b) \frac{\partial \bar{D}_a \left(\frac{dp_a}{dq}(q|x_b)|x_b \right)}{\partial \left(\frac{dp_a}{dq} \right)} \left[1 - \frac{1}{\bar{\varepsilon}_a \left(\frac{dp_a}{dq}(q|x_b)|x_b \right)} \right] - \mathcal{C}' \left[\frac{dp_a}{dq}(q|x_b) \right] = 0,$$

where

$$\begin{aligned} \mathcal{C}' \left[\frac{dp_a}{dq}(q|x_b) \right] \\ = - \int_0^1 \left\{ u_b(v'_b(q; x_b), |x_b - x_a|) - \frac{1 - F_b^{v|x}(v'_b(q; x_b)|x_b)}{f_b^{v|x}(v'_b(q; x_b)|x_b)} \cdot \frac{\partial u_b}{\partial v}(v'_b(q; x_b), |x_b - x_a|) \right\} \frac{\partial D_a \left(\frac{dp_a}{dq}(q|x_b)|x_b; x_a \right)}{\partial \left(\frac{dp_a}{dq} \right)} dx_a. \end{aligned}$$

Now observe that (32) implies that

$$\begin{aligned} u_b(v'_b(q; x_b), |x_b - x_a|) - \frac{1 - F_b^{v|x}(v'_b(q; x_b)|x_b)}{f_b^{v|x}(v'_b(q; x_b)|x_b)} \cdot \frac{\partial u_b}{\partial v}(v'_b(q; x_b), |x_b - x_a|) \\ = \frac{d\rho_b}{dq}(\hat{q}_b(q; x_a; x_b)|x_a; x_b) \left(1 - \frac{1}{\varepsilon_b \left(\frac{d\rho_b}{dq}(\hat{q}_b(q; x_a; x_b)|x_a; x_b)|x_a; x_b \right)} \right). \end{aligned}$$

This means that the above first-order conditions can be rewritten as

$$\begin{aligned} & \frac{dp_a}{dq}(q|x_b) \left[1 - \frac{1}{\bar{\varepsilon}_a\left(\frac{dp_a}{dq}(q|x_b)|x_b\right)} \right] \\ & + \mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a}{dq}(q|x_b))} \left[\frac{d\rho_b}{dq}(\hat{q}_b(q; \tilde{x}_a; x_b)|\tilde{x}_a; x_b) \left(1 - \frac{1}{\varepsilon_b\left(\frac{d\rho_b}{dq}(\hat{q}_b(q; \tilde{x}_a; x_b)|\tilde{x}_a; x_b)|\tilde{x}_a; x_b\right)} \right) \right] = 0, \end{aligned}$$

where $H(x_a|x_b, q)$ is the distribution over $X_a = [0, 1]$ whose density is given by

$$h_a\left(x_a|x_b, \frac{dp_a}{dq}(q|x_b)\right) \equiv \frac{\frac{\partial D_a\left(\frac{dp_a}{dq}(q|x_b)|x_b; x_a\right)}{\partial\left(\frac{dp_a}{dq}\right)}}{\frac{\partial \bar{D}_a\left(\frac{dp_a}{dq}(q|x_b)|x_b\right)}{\partial\left(\frac{dp_a}{dq}\right)}}.$$

The above properties imply the result in the proposition. Q.E.D.

Proof of Proposition 5. Fix $\theta_b = (v_b, x_b)$ and let $q = f_b^x(x_b) [1 - F_b^v(v_b)]$. The result in Proposition 4 implies that, under uniform pricing on side a and customized pricing on side b , for any $x_a \in X_a$ such that $t_b^u(\theta_b, x_a) \in \text{Int}[V_a]$, $t_b^u(\theta_b, x_a)$ is such that

$$\begin{aligned} & u_a(t_b^u(\theta_b, x_a), |x_b - x_a|) - \mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a^u}{dq})} \left[\frac{1 - F_a^v\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right)\right)}{f_a^v\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right)\right)} \cdot \frac{\partial u_a}{\partial v}\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right), |\tilde{x}_a - x_b|\right) \right] \\ & + \mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a^u}{dq})} \left[\varphi_b\left(\theta_b, \left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right), \tilde{x}_a\right)\right) \right] = 0, \end{aligned} \quad (33)$$

where $H\left(x_a|x_b, \frac{dp_a^u}{dq}\right)$ is the distribution over $X_a = [0, 1]$ whose density is given by

$$h\left(x_a|x_b, \frac{dp_a^u}{dq}\right) \equiv \frac{\frac{\partial D_a\left(\frac{dp_a^u}{dq}|x_b; x_a\right)}{\partial\left(\frac{dp_a^u}{dq}\right)}}{\frac{\partial \bar{D}_a\left(\frac{dp_a^u}{dq}|x_b\right)}{\partial\left(\frac{dp_a^u}{dq}\right)}},$$

and where $\frac{dp_a^u}{dq}$ is a shortcut for $\frac{dp_a^u}{dq}(q|x_b)$ with the latter equal to $\frac{dp_a^u}{dq}(q|x_b) = u_a(t_b^u(\theta_b, x_a), |x_a - x_b|)$. Note that, to arrive at (33), we used the result in Proposition 4 along with the property in (15) and the fact that, for any x_a such that $\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|x_a\right) \notin \text{Int}[V_a]$, $h\left(x_a|x_b, \frac{dp_a^u}{dq}\right) = 0$, whereas for any x_a such that $\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|x_a\right) \in \text{Int}[V_a]$,

$$\frac{\frac{dp_a^u}{dq}}{\varepsilon_a\left(\frac{dp_a^u}{dq}|x_b; x_a\right)} = \frac{1 - F_a^v\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|x_a\right)\right)}{f_a^v\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|x_a\right)\right)} \cdot \frac{\partial u_a}{\partial v}\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|x_a\right), |x_a - x_b|\right).$$

We also used the fact that, for any x_a such that $h\left(x_a|x_b, \frac{dp_a^u}{dq}\right) > 0$ (equivalently, $\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|x_a\right) \in \text{Int}[V_a]$),

$$\begin{aligned} & \frac{d\rho_b}{dq}(\hat{q}_b(q; x_a; x_b)|x_a; x_b) \left(1 - \frac{1}{\varepsilon_b\left(\frac{d\rho_b}{dq}(\hat{q}_b(q; x_a; x_b)|x_a; x_b)|x_a; x_b\right)}\right) \\ &= \varphi_b\left(\theta_b, \left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|x_a\right), x_a\right)\right), \end{aligned}$$

as shown in the proof of Proposition 4.

On the other hand, under customized pricing on both sides, for any such $\theta_b = (v_b, x_b)$, any $x_a \in X_a$ such that $t_b^*(\theta_b, x_a) \in \text{Int}[V_a]$, the threshold $t_b^*(\theta_b, x_a)$ is such that

$$u_a(t_b^*(\theta_b, x_a), |x_b - x_a|) - \frac{1 - F_a^v(t_b^*(\theta_b, x_a))}{f_a^v(t_b^*(\theta_b, x_a))} \cdot \frac{\partial u_a}{\partial v}(t_b^*(\theta_b, x_a), |x_b - x_a|) + \varphi_b(\theta_b, (t_b^*(\theta_b, x_a), x_a)) = 0.$$

It is then immediate that, for any x_a such that

$$\begin{aligned} & -\mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a^u}{dq})} \left[\frac{1 - F_a^v\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right)\right)}{f_a^v\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right)\right)} \cdot \frac{\partial u_a}{\partial v}\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right), |x_b - \tilde{x}_a|\right) \right] \\ & + \mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a^u}{dq})} \left[\varphi_b\left(\theta_b, \left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right), \tilde{x}_a\right)\right) \right] \\ & \leq -\frac{1 - F_a^v(t_b^*(\theta_b, x_a))}{f_a^v(t_b^*(\theta_b, x_a))} \cdot \frac{\partial u_a}{\partial v}(t_b^*(\theta_b, x_a), |x_b - x_a|) + \varphi_b(\theta_b, (t_b^*(\theta_b, x_a), x_a)) \end{aligned}$$

we have that $t_b^u(\theta_b, x_a) \geq t_b^*(\theta_b, x_a)$, whereas, for any x_a such that

$$\begin{aligned} & -\mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a^u}{dq})} \left[\frac{1 - F_a^v\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right)\right)}{f_a^v\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right)\right)} \cdot \frac{\partial u_a}{\partial v}\left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right), |x_b - \tilde{x}_a|\right) \right] \\ & + \mathbb{E}_{H(\tilde{x}_a|x_b, \frac{dp_a^u}{dq})} \left[\varphi_b\left(\theta_b, \left(\hat{v}_{x_b}\left(\frac{dp_a^u}{dq}|\tilde{x}_a\right), \tilde{x}_a\right)\right) \right] \\ & \geq -\frac{1 - F_a^v(t_b^*(\theta_b, x_a))}{f_a^v(t_b^*(\theta_b, x_a))} \cdot \frac{\partial u_a}{\partial v}(t_b^*(\theta_b, x_a), |x_b - x_a|) + \varphi_b(\theta_b, (t_b^*(\theta_b, x_a), x_a)) \end{aligned}$$

we have that $t_b^u(\theta_b, x_a) \leq t_b^*(\theta_b, x_a)$.

The results in the proposition then follow from the monotonicity of the function

$$\varphi_b(\theta_b, (t_b^*(\theta_b, x_a), x_a)) - \frac{1 - F_a^v(t_b^*(\theta_b, x_a))}{f_a^v(t_b^*(\theta_b, x_a))} \cdot \frac{\partial u_a}{\partial v}(t_b^*(\theta_b, x_a), |x_b - x_a|)$$

in the distance $|x_a - x_b|$ (holding θ_b fixed), along with the fact that, by virtue of reciprocity, $t_b^u(\theta_b, x_a) \leq t_b^*(\theta_b, x_a)$ if and only if

$$t_a^u((t_b^*(\theta_b, x_a), x_a), x_b) \leq t_b^*((t_b^*(\theta_b, x_a), x_a), x_b)$$

and, likewise, $t_b^u(\theta_b, x_a) \geq t_b^*(\theta_b, x_a)$ if and only if

$$t_a^u((t_b^*(\theta_b, x_a), x_a), x_b) \geq t_b^*((t_b^*(\theta_b, x_a), x_a), x_b).$$

Q.E.D.

Proof of Proposition 6. The proof follows from the combination of the results in Proposition 5 with the results in Proposition 1 in Aguirre et al (2010). When the environment satisfies the conditions in part 1 of Proposition 5, starting from uniform pricing on side a , the introduction of customized pricing on side a leads to an increase in prices for nearby locations and a reduction in prices for distant locations. Proposition 1 in Aguirre et al (2010), along with the fact that the environment satisfies Condition IR and that, for any x_b and dp_a/dq , the convexity $CD_a\left(\frac{dp_a}{dq}|x_b; x_a\right)$ of the demands by the x_a -agents for the q -th unit of the x_b -agents declines with the distance $|x_a - x_b|$, then implies that welfare of the side- a agents is higher under uniform pricing. Likewise, under the conditions in part 2 of Proposition 5, that welfare of the side- a agents is higher under uniform pricing follows from the fact that, starting from uniform pricing on side a , the introduction of customized pricing on side a leads to an increase in prices for distant locations and a reduction in prices for nearby locations. The welfare implications of such price adjustments then follow again from Proposition 1 in Aguirre et al (2010), along with the fact that Condition IR holds and that, for any x_b and dp_a/dq , the convexity $CD_a\left(\frac{dp_a}{dq}|x_b; x_a\right)$ of the demands by the x_a -agents for the q -th unit of the x_b -agents increases with the distance $|x_b - x_a|$. Q.E.D.

Proof of Proposition 7. The results in Proposition 1) imply that, in a centralized market, the platform engages in customized pricing on both sides of the market. Furthermore, the profit-maximizing tariffs induce matching sets with a threshold structure. In particular, the proof of Proposition 1 implies that, for any $\theta_b = (v_b, x_b)$, any $x_a \in X_a$ such that

$$\varphi_a((\bar{v}_a, x_a), \theta_b) + \varphi_b(\theta_b, (\bar{v}_a, x_a)) < 0,$$

the threshold $t_b^*(\theta_b, x_a)$ is equal to $t_b^*(\theta_b, x_a) = \bar{v}_a$. For any $x_a \in X_a$ such that

$$\varphi_a((\underline{v}_a, x_a), \theta_b) + \varphi_b(\theta_b, (\underline{v}_a, x_a)) > 0$$

the threshold $t_b^*(\theta_b, x_a)$ is equal to $t_b^*(\theta_b, x_a) = \underline{v}_a$. Finally, for any $x_a \in X_a$ such that

$$\varphi_a((\underline{v}_a, x_a), \theta_b) + \varphi_b(\theta_b, (\underline{v}_a, x_a)) \leq 0 \leq \varphi_a((\bar{v}_a, x_a), \theta_b) + \varphi_b(\theta_b, (\bar{v}_a, x_a))$$

the threshold $t_b^*(\theta_b, x_a)$ is given by the unique solution to

$$\varphi_a((t_b^*(\theta_b, x_a), x_a), \theta_b) + \varphi_b(\theta_b, (t_b^*(\theta_b, x_a), x_a)) = 0,$$

or, equivalently,¹⁸

$$u_a(t_b^*(\theta_b, x_a), |x_a - x_b|) - \frac{1 - F_a^v(t_b^*(\theta_b, x_a))}{f_a^v(t_b^*(\theta_b, x_a))} \cdot \frac{\partial u_a}{\partial v}(t_b^*(\theta_b, x_a), |x_a - x_b|) + \varphi_b(\theta_b, (t_b^*(\theta_b, x_a), x_a)) = 0. \quad (34)$$

¹⁸Note that, in writing (34), we used the fact that Condition I_a holds.

Next, consider the prices optimally set by the sellers in a decentralized market. Because each seller can sell at most a single unit to each buyer, no seller can engage in price discrimination. Furthermore because buyers have separable demands, there is no competition among the sellers. This means that the price $p_b^d(\theta_b)$ optimally selected by each seller with type $\theta_b = (v_b, x_b)$ maximizes

$$p\bar{D}_a(p|x_b) + \int_0^1 u_b(v_b, |x_b - x_a) D_a(p|x_b; x_a) dx_a$$

where $\bar{D}_a(p|x_b) \equiv \int_0^1 D_a(p|x_b; x_a) dx_a$, with $D_a(p|x_b; x_a) = f_a^x(x_a) [1 - F_a^v(\hat{v}_{x_b}(p|x_a))]$ and with the function $\hat{v}_{x_b}(p|x_a)$ as in (31). The optimal price $p_b^d(\theta_b)$ thus solves optimality conditions analogous to those for the marginal price $dp_a^u(q|x_b)/dq$ set by the platform under uniform pricing on side a , but with the procurement cost adjusted by removing the monopsony markup

$$-\frac{p}{\varepsilon_b(p|x_a; x_b)}.$$

In particular, for any θ_b such that $p_b^d(\theta_b) \in (u_a(\underline{v}_a, \frac{1}{2}), u_a(\bar{v}_a, 0))$, the price $p_b^d(\theta_b)$ must solve the optimality condition

$$p_b^d(\theta_b) \left[1 - \frac{1}{\bar{\varepsilon}_a(p_b^d(\theta_b)|x_b)} \right] + \mathbb{E}_{H(\tilde{x}_a|x_b, p_b^d(\theta_b))} [u_b(v_b, |x_b - \tilde{x}_a|)] = 0,$$

where $H(x_a|x_b, \frac{dp_a^u}{dq})$ is the distribution over $X_a = [0, 1]$ whose density is given by

$$h(x_a|x_b, p) \equiv \frac{\frac{\partial D_a(p|x_b; x_a)}{\partial p}}{\frac{\partial \bar{D}_a(p|x_b)}{\partial p}}.$$

Note that $u_a(\underline{v}_a, \frac{1}{2})$ is the lowest willingness to pay, among all side- a agents, for each x_b -product, whereas $u_a(\bar{v}_a, 0)$ is the highest.

The price $p_b^d(\theta_b)$ optimally set by each seller of type $\theta_b = (v_b, x_b)$ thus induces location-specific demands on side a with a threshold structure where the threshold $t_b^d(\theta_b, x_a) = \hat{v}_{x_b}(p_b^d(\theta_b)|x_a)$. Hence, for any x_a such that $t_b^d(\theta_b, x_a) \in \text{Int}[V_a]$ (equivalently, for which $u_a(\underline{v}_a, |x_b - x_a|) \leq p_b^d(\theta_b) \leq u_a(\bar{v}_a, |x_b - x_a|)$), the threshold $t_b^d(\theta_b, x_a)$ is such that

$$u_a(t_b^d(\theta_b, x_a), |x_a - x_b|) - \mathbb{E}_{H(\tilde{x}_a|x_b, p_b^d(\theta_b))} \left[\frac{1 - F_a^v(\hat{v}_{x_b}(p_b^d(\theta_b)|\tilde{x}_a))}{f_a^v(\hat{v}_{x_b}(p_b^d(\theta_b)|\tilde{x}_a))} \frac{\partial u_a}{\partial v}(\hat{v}_{x_b}(p_b^d(\theta_b)|\tilde{x}_a), |\tilde{x}_a - x_b|) \right] \\ + \mathbb{E}_{H(\tilde{x}_a|x_b, p_b^d(\theta_b))} [u_b(v_b, |x_b - \tilde{x}_a|)] = 0.$$

The results in the proposition then follow from arguments that parallel those establishing Propositions 5 and 6. First, one can show that the assumptions in part 1 (alternatively, part 2) imply that targeting is higher in a decentralized market (alternatively, in a centralized market). That, in the two scenarios covered by parts 1 and 2 in the proposition, welfare of the side- a agents is higher in a decentralized market than in a centralized one then follows from the above properties along with the fact that the environment also satisfies condition IR and the additional convexity/concavity properties of the $CD_a\left(\frac{dp_a}{dq}|x_b; x_a\right)$ function in the proposition (again, this follows from arguments identical to those establishing Proposition 1 in Aguirre et al (2010)). Q.E.D.

References

- [1] Aguirre I., S. Cowan, and J. Vickers, 2010, “Monopoly Price Discrimination and Demand Curvature,” *American Economic Review*, Vol. 100, no. 4, pp. 1601-15.
- [2] Balestrieri, F. and S. Izmalkov, 2015, “Informed Seller in a Hotelling Market,” Working Paper, New Economic School.
- [3] Belleflamme, P. and M. Peitz, 2017, “Platforms and Network Effects,” In Corchon, L. and M. Marini (Eds). *Handbook of Game Theory and Industrial Organization*. Edward Elgar, Cheltenham, UK, Northampton, MA, USA.
- [4] Bergemann, D., and A. Bonatti, 2011, “Targeting in Advertising Markets: Implications for Offline versus Online Media,” *Rand Journal of Economics*, Volume 42, Issue 3, pp. 417–443.
- [5] Bergemann, D., and A. Bonatti, 2015, “Selling Cookies,” *American Economic Journal: Microeconomics*, 7(3), pp. 259–294.
- [6] Bergemann, D., B. Brooks, and S. Morris, 2015, “The Limits of Price Discrimination,” *American Economic Review*, 105(3), pp. 921-57.
- [7] Bulow, J. and D. Roberts, 1989, “The Simple Economics of Optimal Auctions.” *Journal of Political Economy*, Vol. 97, pp. 1060-1090.
- [8] Cox, H., B. Straathof, and G. Zwart, 2017, “Targeted Advertising, Platform Competition, and Privacy,” *Journal of Economics and Management Strategy*, 26, 557-570.
- [9] Crawford, G., 2000, “The Impact of the 1992 Cable Act on Household Demand and Welfare,” *Rand Journal of Economics*, Vol. 31(3), pp 422-450.
- [10] Crawford, G. and A. Yurukoglu, 2012, “The Welfare Effects of Bundling in Multichannel Television Markets,” *American Economic Review*, 102(2), pp. 643-85.
- [11] Damiano, E. and H. Li, 2007, “Price Discrimination and Efficient Matching,” *Economic Theory*, Vol. 30, pp. 243-263.
- [12] Fershtman D. and A. Pavan, 2017, “Matching Auctions,” Working Paper, Northwestern University.
- [13] Federal Communications Commission (FCC), 2004, “Report on the Packaging and Sale of Video Programming to the Public.” <http://www.fcc.gov/mb/csrptpg.html>.
- [14] Federal Communications Commission (FCC), 2006, “Further Report on the Packaging and Sale of Video Programming to the Public.” <http://www.fcc.gov/mb/csrptpg.html>.

- [15] Goettler R. and R. Shachar, 2001, “Spatial Competition in the Network Television Industry,” *RAND Journal of Economics* Vol. 32, No. 4, pp. 624–656.
- [16] Gomes, R. and A. Pavan, 2016, “Many-to-Many Matching and Price Discrimination,” *Theoretical Economics*, 11, 1005–1052.
- [17] Hart, S. and P. Reny, 2015, “Maximal Revenue with Multiple Goods: Nonmonotonicity and Other Observations,” *Theoretical Economics*, 10, 893–922.
- [18] Jeon, D., B. Kim, and D. Menicucci, 2016, “Second-degree Price Discrimination by a Two-sided Monopoly Platform,” Working Paper, Toulouse School of Economics.
- [19] Jonhson, T., 2013, “Matching Through Position Auctions,” *Journal of Economic Theory*, Vol. 148, pp. 1700-1713.
- [20] Jullien B. and A. Pavan (2017), “Information Management and Pricing in Platform Markets,” Working paper, Toulouse School of Economics and Northwestern University.
- [21] Kang, Karam and Hye Young You (2016), “Lobbyists as Matchmakers in the Market for Access,” Working Paper, Carnegie Mellon University and Vanderbilt University.
- [22] Maskin, E., and J. Riley, 1984, “Monopoly with Incomplete Information,” *Rand Journal of Economics*, Vol. 15, pp. 171-196.
- [23] Mirrokni, V. and H. Nazerzadeh, 2017, “Deals or No Deals: Contract Design for Online Advertising,” *World Wide Web Conference (WWW2017)*.
- [24] Myerson, R., 1981, “Optimal Auction Design,” *Mathematics of Operations Research*, Vol. 6(1), pp. 58-73.
- [25] Mussa, M. and S. Rosen, 1978, “Monopoly and Product Quality,” *Journal of Economic Theory*, Vol. 18, pp. 301-317.
- [26] Office of Fair Trading, 2010, “Online Targeting of Advertising and Prices: A Market Study.”
- [27] Tan, G. and J. Zhou (2017), “Price Competition in Multi-sided Markets,” mimeo University of Southern California and National University of Singapore.
- [28] Wilson, R., 1993, “Non-Linear Pricing,” *Oxford University Press*.