Price Customization and Targeting in Matching Markets*

Renato Gomes[†] **Alessandro Pavan**[‡]

January 2019

Abstract

Recent technologies enable matching intermediaries to engage in unprecedented levels of *targeting*, whereby matches finely depend on the agents' characteristics, but also favor *customized* (i.e., match-specific) pricing. Yet, recent regulations prohibiting the transfer of personal data and imposing restrictions on bundling are expected to hinder price customization and favor *uniform* pricing (whereby the price of a match charged to agents on a given side of a market is invariant to the agents' characteristics). To assess the impact of these developments, we build a model in which agents' preferences are both vertically and horizontally differentiated. We show how, absent uniform-price obligations, platforms maximize profits through price customization and link the latter to structural elasticities. Perhaps surprisingly, we show that uniform-pricing obligations may either increase or decrease targeting, depending on testable properties of the demand. We then identify conditions under which uniform pricing is beneficial to consumer welfare. Finally, we show how the same analysis can shed light on whether the transition from a centralized to a decentralized structure (whereby one side posts prices and the other determines the composition of the matching sets) is welfare enhancing. The results have implications for online shopping, ad-exchanges, and media platforms.

JEL classification: D82

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

*For useful comments and suggestions, we thank Igal Hendel, Wilfried Sand-Zantmann, Steve Tadelis, Jean Tirole, and seminar participants at the 2017 Barcelona Conference on the Economics of Platforms, the 2018 EIEF Conference on Platforms, the 2018 CEPR-Gerzensee Summer School in Economic Theory, and various other workshops where the paper was presented. A special thank is to our discussants Stephan Lauermann and Markus Reisinger. Pavan also thanks the NSF (Research Grant SES 1530798) for financial support, and Bocconi University for its hospitality during the 2017-2018 academic year.

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

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

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 ad exchanges, matching publishers with advertisers, business-to-business platforms, matching firms with mutually beneficial commercial interests, and 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 depends on characteristics of the matching partners.

For instance, in advertising exchanges, the assignment of, and payments from, advertisers depend on scores that summarize the compatibility of the ads with each publisher's content.[1] A similar trend can be found in other markets, not traditionally analyzed through the lens of matching. In online shopping, for example, it is common practice among retailers to use customers' personal data to set personalized prices. In one of the most publicized cases, Orbitz, an online travel agency, reportedly used information about customers' demographics to charge targeted customers higher hotel fees.[2] Similarly, Safeway, an online grocery chain, often proposes individualized price offers and quantity discounts to customers with certain profiles.[3] The retailers' knowledge about consumers' characteristics often comes from data brokers, who collect and sell personal information (in the form of demographics, geolocation, browsing history, etc).[4]

In the markets mentioned above, price customization is easy to enforce, as the agents' "horizontal" characteristics are observable (for instance, in ad exchanges, the advertisers' profile is often revealed by the ads' content) or can be learnt from third parties (for example, in online retailing, information about consumers can often be obtained from data brokers or affiliated websites). In other markets, instead, the agents' horizontal characteristics that are relevant for price customization have to be indirectly elicited, and this may require bundling.[5] A case in point is that of media markets (for instance, satellite/cable TV providers), that use sophisticated pricing strategies that condition payments on the entire bundle of channels selected by the subscribers.[6]

[1] See, for example, https://support.google.com/adxseller/answer/2913506?hl=en&ref_topic=3376095. Moreover, ad exchanges use advertiser-specific reservation prices which are easily automated using proxy-bidding tools. Ad exchanges also price discriminate on the publisher side, by making the payments to the publishers depend on the publishers' profile and on the volume of impressions.

[2] See the article "On Orbitz, Mac Users Steered to Pricier Hotels," the *Wall Street Journal*, August 23, 2012.

[3] See https://www.bloomberg.com/news/articles/2013-11-14/2014-outlook-supermarkets-offer-personalized-pricing.

[4] According to *The New York Times*, the data broker industry's revenue reached $156 billion in 2013 (see the article "The Dark Market for Personal Data," August 16, 2014).

[5] For instance, Ad exchanges have recently developed new contractual arrangements that allow them to bundle different ads as a way of screening the publishers' unobservable preferences (see, Mirrokni and Nazerzadeh (2017)).

[6] Most satellite/cable TV 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

While having a long history in the policy debate, price customization has attracted renewed attention due to the two-sided nature of matching intermediaries and the amount of information now available for target pricing.[7] The concern is that, by leveraging the platforms' market power, price customization hinders the efficiency gains permitted by better targeting technologies. Recent regulations speak directly to these issues. In the European Union, for example, the General Data Protection Regulation (GDPR) and the ePrivacy Regulation (ePR) mandate that businesses ask for consumers' consent prior to collecting and transmitting personal data. Such regulations hinder price customization based on data from third parties.[8,9]

Market decentralization poses yet another challenge to price customization. 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, while the other side (typically, the "buyer" side) then constructs the matching sets. For example, in the market for media content, several analysts believe that the increase in the speed of fiber-optic and broadband internet connection will favor a gradual transition to a market structure whereby viewers will pay directly the content providers, bypassing the intermediation of media platforms (such as satellite or cable TV providers).[10] In decentralized markets, price customization is often hindered by the absence of data about consumers' tastes (for instance, due to privacy regulation), along with the difficulty of tracking consumers' purchases with other vendors. This is in contrast to centralized matching markets where price customization can often be implemented through bundling (see footnote 6).

The aim of this paper is to understand how targeting and price customization shape the matching opportunities offered by profit-maximizing platforms, and study the impact on targeting and consumer welfare of uniform pricing (whereby payments to the platforms do not depend on the "horizontal characteristics" of the agents' profiles), be it a result of regulation or market decentralization.

To examine these issues, we develop a model where agents' preferences exhibit elements of both

---

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.

[7] In the case of media markets, see, for example, the Federal Communications Commission 2004 and 2006 reports on the potential harm of price customization through bundling. In the case of online retailing, see the UK Office of Fair Trading 2010 eponymous report on online targeting in advertising and pricing.

[8] See Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the processing of personal data and on the free movement of such data.

[9] In the US, the Federal Trade Commission (FTC) recommended in 2014 legislation increasing the transparency of the data broker industry and giving consumers greater control over their personal information. See https://www.ftc.gov/news-events/press-releases/2014/05/ftc-recommends-congress-require-data-broker-industry-be-more.

[10] A similar trend is also taking place in online advertising, where many publishers now prefer employing direct-sales channels so as to avoid the commissions charged by the platforms. See, for instance, https://digiday.com/media/advertisers-are-often-left-flying-blind-in-ad-exchanges/.

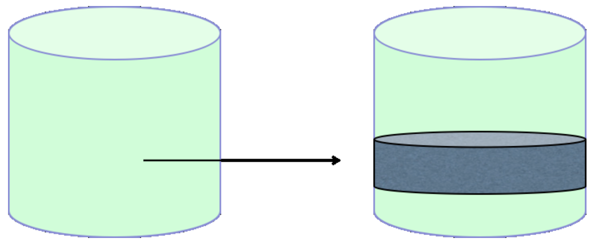

Figure 1: Market geometry

vertical and horizontal differentiation. Certain agents value being matched with agents from the 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 (see Figure 1).

Each agent's utility from interacting with any other agent from the opposite side increases 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 circular 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 sets demanded by the agents. A tariff exhibits *uniform pricing* if all agents from a given side face the same price schedule for different quantities of the matches with agents at a given location on 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 offer *customized* tariffs on both sides, whereby the customization is obtained by offering agents menus of *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," such tariffs thus also display a form of third-degree price discrimination. When neither the vertical nor the horizontal dimensions defining the agents' preferences are observable by the platform, information about such dimensions is elicited via bundling, through a careful design of the baseline plans and the associated price-customizing schedules.

Our first main result derives properties of the demands under profit-maximizing tariffs. We then offer a convenient representation of the profit-maximizing tariffs in terms of elasticities on both sides of the market. 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 non-linear 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 bundling and self-selection).

Our characterization of the profit-maximizing matching tariffs reveals interesting patterns of cross-subsidization, unique to matching environments. Namely, at any given location, agents with a low value for matching (a low 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 (both in terms of elasticities and in terms of the agents' match values and type distributions) 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 standard models of price discrimination, which are not amenable to study the effects of market power on targeting levels.

Our second set of results investigates the effects on prices, the composition of the matching demands, and consumer welfare of uniform pricing on a given side of the market. 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 uniformity of prices, but also how the procurement costs of matches are affected by the horizontal component of the agents' 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 the familiar 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 affects targeting and welfare. Intuition might suggest that uniform pricing should increase targeting by preventing platforms 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 can either decrease or increase the equilibrium level of targeting, depending on how match-demand elasticities vary with location. We also relate elasticities to match values and type distributions and identify conditions under which targeting is higher (alternatively, lower) under uniform pricing (alternatively, customized pricing).

We then use such characterization to look into the welfare effects of uniform pricing. 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. The results identify sufficient conditions (both in terms of elasticities and in terms of match utilities and type distributions) for uniform pricing to increase consumer surplus (in the side where prices are uniform). 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 that the same analysis that permits us to uncover the welfare effects of uniform pricing can be adapted to shed light on the welfare effects of a transition from a centralized to a decentralized market structure (where sellers independently post stand-alone prices), thus contributing a novel angle to the policy debate over whether or not such transition should be encouraged.

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 properties of profit-maximizing tariffs and of the induced matching demands. Section 4 studies the effects of uniform-pricing obligations, whereas Section 5 discusses the effects of the transition from a centralized to a decentralized market structure. Section 6 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 the partners and experience shocks to their preferences over time.

Related are also Jullien and Pavan (2017), and Tan and Zhou (2017). The former 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 latter paper studies price competition in a 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. While considering rich preference structures, both papers abstract from price discrimination. The latter is examined in Damiano and Li (2007) and Johnson (2013). Contrary to this paper, these works consider markets where matching is one-to-one and agents' preferences are differentiated along a vertical dimension only.

This paper considers a many-to-many matching market where agents might disagree on the relative attractiveness of any two agents from the other side (horizontal differentiation). Similar preferences structures are examined in the matching literature surveyed in Roth and Sotomayor (1990) (for a more recent treatment, see Hatfield and Milgrom (2005)), and in the literature on the school assignment problem (see, for example, Abdulkadiroglu, Pathak and Roth (2005a, 2005b) and Abdulkadiroglu and Sonmez (2003)). These literatures are methodologically distinct from the current paper, in that they focus on solution concepts such as stability and do not allow for transfers.

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). Most of this literature, however, restricts attention to a single network, or to mutually exclusive networks. Ambrus, Calvano and Reisinger (2016) relax this structure by proposing a model of competing media platforms with overlapping viewerships (i.e., multi-homing). By contrast, we stick to a monopolistic market, but introduce a richer preferences structure (displaying horizontal tastes for matches), which enables us to study targeting and price customization in such markets.

Our discussion of market decentralization is related to Loertscher and Niedermayer (2017), who compare the agency and wholesale business models by a platform that faces competition from an independent bilateral exchange. To focus on the substitutability between market places, that paper, however, abstracts from targeting and price customization.

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 analysis differs from this literature in two important dimensions. First, the platform's feasibility constraint (namely, the reciprocity of the matches) has no equivalent in standard 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 novel cross-subsidization patterns.[11]

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 for recent developments. The latter 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 the applications under consideration.

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 certain special cases. 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 Kox 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 advertises and the publishers (or content providers) is mediated. Contrary to some of the papers in this literature, however, we abstain from platform competition. Importantly, we also assume that agents can perfectly communicate their preferences and face no informational frictions regarding the desirability of the matches. Eliaz and

[11] Related 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.

Spiegler (2016) relax these assumptions and consider the mechanism design problem of a platform that wants to allocate firms into search pools created in response to noisy preference signals provided by the consumers. Relatedly, Eliaz and Spiegler (2017) consider the problem of a profit-maximizing advertising platform who wants to elicit the advertisers' profiles so as to match them to consumers with preferences for diversity. These papers do not investigate the effects of uniform pricing, but rather focus on the incentives of firms to truthfully reveal their "ideal audiences."

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 attractiveness.

The parameter x_k, which describes the agent's location, captures *horizontal* differentiation in preferences. We assume that agents are located on a circle of perimeter one, in which case $X_k = [0, 1]$, $k = a, b$. The parameter $v_k \in V_k \equiv [\underline{v}_k, \overline{v}_k] \subseteq \mathbb{R} \cup \{+\infty\}$ captures heterogeneity in preferences along a *vertical* dimension. It controls for the intensity of the 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). See Figure 1; the cylinder on each side represents the population on that side of the market. Each individual is located on the external surface of the cylinder with the hight of the cylinder describing the individual's vertical type and his position on the circle describing his horizontal type.

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 $\theta_k = (v_k, x_k)$ derives 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 circular (minimal) 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 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.

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 $\mathbf{s}_k \subseteq \Theta_l$ from side $l \neq k$ is given by

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

The payoff that the same agent obtains outside of the platform is equal to zero.

We assume that the vertical dimensions v_k are the agents' private information. As for the horizontal dimensions, i.e., the locations x_k, to ease the exposition, in the main text, we assume they are publicly observable. In the Appendix, however, we explain how the analysis can accommodate for the possibility that locations are private information, under additional assumptions on the distributions F_k (see the proof of Lemma 1).

Let F_k^v (alternatively, F_k^x) denote the marginal distribution of F_k with respect to v_k (alternatively, x_k), and $F_k^{v|x}$ the distribution of v_k conditional on x_k. Then let f_k^v be the density of F_k^v and $\lambda_k^v \equiv f_k^v/[1-F_k^v]$ its hazard rate. An analogous notation applies to the densities and hazard rates of the conditional distributions $F_k^{v|x}$. Finally, let $\mathrm{Int}[V_k]$ denote the interior of the set V_k and $\Sigma(\Theta_l)$ the collection of all F_l-measurable subsets of Θ_l. Hereafter, we assume that the "virtual values"

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

respect the same rankings as the true values, which is the natural analog of standard regularity conditions (e.g., Myerson (1981)) in matching environments.[12]

Remark 1. The representation in (1) assumes the agent is matched to all agents from side $l \neq k$ whose type is in $\mathbf{s}_k$. That matching sets are described by the types from the opposite side an agent has access to, as opposed to the identities of the agents in the matching set, reflects the property that, under both the welfare- and the profit-maximizing tariffs, each agent from each side $k = a, b$ who decides to include in his matching set some agent from side $l \neq k$ whose type is θ_l, optimally chooses to include in his matching set *all* agents from side l whose type is θ_l. The specification in (1) also implies that the utility that agent i from side k derives from being matched to agent j from side $l \neq k$ is invariant to who else the agent is matched with, as well as who else from the agent's own side is matched to agent j. It also implies that such utility is invariant to agent j's vertical type, v_k. In a previous version, we considered a more general setting where such assumptions are relaxed. In particular, we assumed that (a) each agent's attractiveness depends on both the agent's location and other traits that are correlated with the agent's vertical dimension, and (b) that the utility that each agent derives from each matching set $\mathbf{s}_k$ is given by a concave (alternatively, convex) transformation of $\int_{\mathbf{s}_k} u_k(\cdot) dF_l(\theta_l)$, where the concavity/convexity of the transformation captures decreasing/increasing returns to match quality. Because both the exposition and the analysis are more tedious under such richer specifications, we opted here for the convenience of the representation in (1).

The following examples illustrate the type of markets the analysis is meant for.

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

[12] By this we mean the following. Take any pair $\left(\theta_k, \theta_l\right), \left(\theta'_k, \theta'_l\right) \in \Theta_k \times \Theta_l$. Then $\varphi_k\left(\theta_k, \theta_l\right) \geq \varphi_k\left(\theta'_k, \theta'_l\right)$ if and only if $u_k\left(v_k, |x_k - x_l|\right) \geq u_k\left(v'_k, |x'_k - x'_l|\right)$.

$v_a\phi\left(|x_a - x_b|\right)$, where v_a is the advertiser's profit per sale and where the strictly decreasing function $\phi : [0, \frac{1}{2}] \rightarrow [0, 1]$ 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. The matching (dis)utility $u_b(x_b, |x_a - x_b|)$ that a publisher of type $\theta_b = (v_b, x_b)$ obtains from displaying the ad of advertiser $\theta_a = (v_a, x_a)$ may reflect both 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, as well as the (positive or negative) impact that the ad may have on the publisher's viewership. In this example, both the advertisers' ideal type of audience, x_a, and each content provider's profile, x_b, are likely to be observable by the platform.[13] $\diamondsuit$

(media platform) The platform is a media outlet matching viewers from side a with content providers from side b. The utility $u_a(v_a, |x_a - x_b|)$ that a viewer of type $\theta_a = (v_a, x_a)$ derives from the content of a provider of type $\theta_b = (v_b, x_b)$ is increasing in the overall importance that the viewer assigns to having access to content, captured by the parameter $v_a \in V_a \subset \mathbb{R}_+$, and decreasing in the distance between the viewer's ideal type of content, x_a, and the provider's content, x_b. The matching (dis)utility $u_b(v_b, |x_b - x_a|))$ of the content provider may reflect the extra revenue from advertisers (which may depend on the type of viewers reached, as advertisers typically pay more to content providers with a higher exposure to viewers of certain characteristics), or the expenses from broadcasting rights paid to third parties (which are typically invariant to the type of audience reached). While each content provider's profile (the type of content provided) is observable, each viewer's ideal type of content is likely to be the viewer's private information (this situation amounts to Scenario (ii) in the proof of Lemma 1 in the Appendix).[14] $\diamondsuit$

(online shopping) The platform is an online intermediary matching consumers from side a with sellers from side b. Each seller's location, x_b, identifies the seller's product variety, whereas each buyer's location, x_a, identifies the buyer's most preferred variety. The utility that a buyer of type $\theta_a = (v_a, x_a)$ derives from purchasing a product of variety x_b is decreasing in the distance between the buyer's ideal type of product and the one purchased. In this example, the sellers' preferences are invariant in the distance $|x_b - x_a|$ and the vertical parameter $v_b < 0$ captures the sellers' constant marginal costs. As with the previous example, each seller's product variety, x_b, is likely to be publicly observable, whereas each buyer's most preferred variety (her bliss point) may be known to the platform (for example, when the latter has access to data about consumers' characteristics), or be the buyer's private information (for example, as the result of privacy-protecting regulations). $\diamondsuit$

[13] Another example that shares the preference structure of Example 1 is that of a lobbying firm (platform) matching interest groups from side a with public officials from side b (see Kang and You (2016) for a detailed description of this market).

[14] The structure of this example follows closely the one typically assumed in the empirical literature on media markets (see, e.g., Goettler and Shachar 2001).

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 side k must pay to the platform for being matched to the set of types $\mathbf{s}_k$ 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

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

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 they admit demands satisfying the following *reciprocity condition*, for all $(\theta_k, \theta_l) \in \Theta_k \times \Theta_l$, $k, l \in \{a, b\}$, $l \neq k$,

$$\theta_l \in \hat{\mathbf{s}}_k(\theta_k) \Longleftrightarrow \theta_k \in \hat{\mathbf{s}}_l(\theta_l). \tag{3}$$

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 $\mathbf{s}_k$, any location x_l, we denote by $q_{x_l}(\mathbf{s}_k)$ the "mass" of side-l agents located at x_l included in the matching set $\mathbf{s}_k$.[15]

The platform's problem consists of choosing a pair of feasible matching tariffs T_k, $k = a, b$, along with a pair of demands consistent with the selected tariffs, that jointly maximize the platform's profits, which are given by

$$\sum_{k=a,b} \int_{\Theta_k} T_k(\hat{\mathbf{s}}_k(\theta_k)) dF_k(\theta_k). \tag{4}$$

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

3 Customized Tariffs

We start by studying the platform's problem when no restrictions are imposed on the matching tariffs it can offer. The next definition describes a type of tariffs that plays an important role in the analysis below.

Definition 2. The tariff T_k is *customized* if there exists a collection of matching plans

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

[15]Hereafter, we abuse of terminology by referring to the density of agents of a certain type as the "mass" of agents of that type.

one for each side-k location $x_k \in [0,1]$, such that each side-k agent selecting the plan x_k and then the customization $\mathbf{s}_k \in \mathcal{S}_k(x_k) \subseteq \Sigma(\Theta_l)$ is asked to make a total payment equal to[16]

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

with $\rho_k\left(q_{x_l}\left(\underline{\mathbf{s}}_k(x_k)\right)|x_l; x_k\right) = 0$ all $x_l \in [0,1]$, where $\mathcal{S}_k(x_k)$ is the collection of possible customizations under the plan x_k.

Under customized tariffs, the platform thus offers 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 $\mathcal{S}_k(x_k) \subseteq \Sigma(\Theta_l)$, 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. 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. The price depends on the plan selected by the side-k agent, which is conveniently indexed by the side-k location, x_k. The dependence of the price $\rho_k(q|x_l; x_k)$ on the plan x_k is a manifestation of a particular form of *bundling* practiced by the platform. In particular, note that a menu of customized tariffs combines elements of *second-degree* price discrimination (each price function $\rho_k\left(q|x_l; x_k\right)$ is possibly non-linear in q) with elements of *third-degree* price discrimination (each non-linear price function $\rho_k\left(q|x_l; x_k\right)$ depends on the plan, and hence the location, of the side-k agents). That $\rho_k\left(q_{x_l}\left(\underline{\mathbf{s}}_k(x_k)\right)|x_l; x_k\right) = 0$, 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 online retailers.

We then have the following result:

Lemma 1. ***(properties of the optimum)*** *The profit-maximizing tariffs along with the demands they induce satisfy the following properties, for any $k \in \{a, b\}$,*

1. *the matching tariff T_k^* is customized;*

2. *there exist functions $t_k^* : \Theta_k \times [0,1] \to V_l$ such that the matching demands under T_k^*, $k = a, b$, are given by*

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

[16] The payment specified by the tariff for any matching set $\mathbf{s}_k \notin \{\cup\mathcal{S}_k(x_k) : x_k \in [0,1]\}$ 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 $\mathbf{s}_k \in \mathcal{S}_k(x_k) \cap \mathcal{S}_k(x_k')$, 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 $\{(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathcal{S}_k(x_k)) : x_k \in [0,1]\}$ may entail multiple prices for the same matching set $\mathbf{s}_k$. However, because, in this case, each agent can be constrained to choosing the plan designed for his location, de facto each agent faces a tariff specifying a single price for each set.

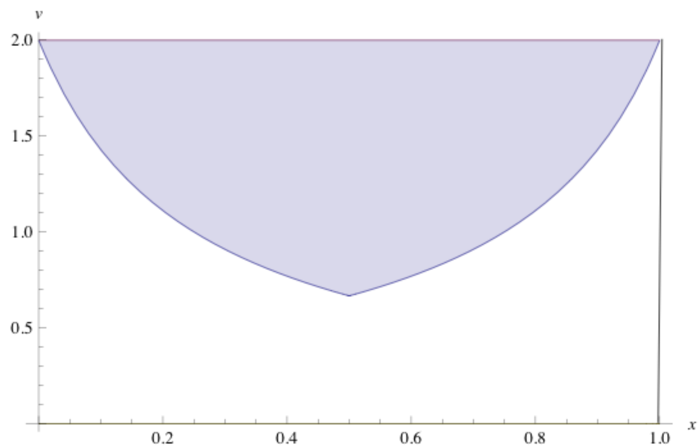

Figure 2: Matching sets under profit-maximizing tariffs. The shaded area in the figure describes the matching set for an agent from side a located at $x_a = 1/2$.

with the threshold function t_k^* *non-increasing in* v_k *and non-decreasing in* $|x_k - x_l|$.

Figure 2 illustrates the matching set of a representative agent from side a located at $x_a = .5$ under a pair of profit-maximizing tariffs.

Under the profit-maximizing tariffs, 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 gain intuition, note that 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,

$$\varphi_k(\theta_k, \theta_l) + \varphi_l(\theta_l, \theta_k) \geq 0. \tag{6}$$

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). That virtual values respect the same rankings of the true values implies existence of a threshold $t_k^*(\theta_k, x_l)$ such that Condition (6) is satisfied if, and only if, $v_l \geq t_k^*(\theta_k, x_l)$, with the threshold $t_k^*(\theta_k, x_l)$ non-increasing in v_k and non-decreasing in $|x_k - x_l|$. Jointly, these properties imply that, as v_k increases, the matching set of type θ_k expands to include new agents with lower vertical types. Moreover, 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.

Remark 2. In the Appendix, we show that the result in Lemma 1 extends to markets in which locations are private information, provided that the distributions from which the types are drawn guarantee that the platform can price discriminate along the agents' locations, without leaving the agents extra rents for the private information they may possess regarding their locations. These conditions are met for example when the distributions are symmetric across locations, in which case the matching demands of any two types $\theta_k = (v_k, x_k)$ and $\theta_k' = (v_k, x_k')$ with the same vertical dimension are location translations of one another. The same conditions also rule out the optimality of bunching and lotteries over matching sets, thus permitting us to highlight the role of targeting and price customization in the simplest possible manner.

3.1 Lerner-Wilson formula for matching schedules

We now derive further properties of the customized tariffs that maximize the platform's profits. As explained in the proof of Lemma 1, in markets in which locations are public on both sides, the platform does not need to impose restrictions on the set of possible customizations, i.e., without loss of optimality, the platform can set $\mathcal{S}_k(x_k) = \Sigma(\Theta_l)$, for all $x_k \in [0,1]$. To ease the notation, hereafter, we thus drop the sets $\mathcal{S}_k(x_k)$ from the specification of the matching plans.

Now consider the problem of a side-k agent of type $\theta_k = (v_k, x_k)$ under the plan x_k. The mass of agents located at x_l demanded by type θ_k is given by

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

Assuming the price schedule $\rho_k(\cdot|x_l; x_k)$ is convex and differentiable in q, with derivative $\rho_k'(\cdot|x_l; x_k)$, it follows that $\hat{q}_{x_l}(\theta_k)$ is a solution to the following first-order condition[17]

$$u_k(v_k, |x_k - x_l|) = \rho_k'\left(\hat{q}_{x_l}(\theta_k)|x_l; x_k\right) \tag{7}$$

whenever $\hat{q}_{x_l}(\theta_k)$ is interior, i.e., whenever $\hat{q}_{x_l}(\theta_k) \in (0, f_l^x(x_l))$.

Next, for any pair of locations $x_k, x_l \in [0,1]$, and any "interior" marginal price

$$\rho_k' \in \left[u_k(\underline{v}_k, |x_k - x_l|), u_k(\bar{v}_k, |x_k - x_l|)\right],$$

let $\hat{v}_{x_l}\left(\rho_k'|x_k\right)$ denote the value of v_k that makes each agent from side k located at x_k indifferent between adding the extra q-th unit of x_l-agents or not, given the marginal price ρ_k'. Note that, irrespective of q, $\hat{v}_{x_l}\left(\rho_k'|x_k\right)$ is implicitly defined by

$$u_k(v_k, |x_k - x_l|) = \rho_k'. \tag{8}$$

If, instead, $\rho_k' \notin [u_k(\underline{v}_k, |x_k - x_l|), u_k(\bar{v}_k, |x_k - x_l|)]$, let $\hat{v}_{x_l}\left(\rho_k'|x_k\right) = \underline{v}_k$ for all $\rho_k' < u_k(\underline{v}_k, |x_k - x_l|)$, and $\hat{v}_{x_l}\left(\rho_k'|x_k\right) = \bar{v}_k$ for all $\rho_k' > u_k(\bar{v}_k, |x_k - x_l|)$.

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 ρ_k' uniquely identifies the quantity q. Furthermore,

[17] The 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).

because agents with higher vertical types demand larger matching sets, the *demand* for the q-th unit of the x_l-agents by the x_k-agents, at the *marginal price* ρ_k', is given by[18]

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

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

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 ρ_k' by (once again, the arguments of the marginal price ρ_k' are dropped to ease the notation)

$$\varepsilon_k\left(\rho_k'|x_l;x_k\right) \equiv -\frac{\partial D_k\left(\rho_k'|x_l;x_k\right)}{\partial\left(\rho_k'\right)} \cdot \frac{\rho_k'}{D_k\left(\rho_k'|x_l;x_k\right)}. \tag{10}$$

The next proposition characterizes the price schedules associated with the profit-maximizing customized tariffs of Lemma 1 in terms of the location-specific elasticities of the demands on both sides of the market [To ease he exposition, the dependence of the marginal prices, $\rho_k^{*\prime}$, the demands. D_k, and the elasticities, ε_k, on the locations $(x_l;x_k)$ is dropped from all the formulas in the proposition].

Proposition 1. ***(Lerner-Wilson price schedules)*** *The price schedules ρ_k^* associated with the profit-maximizing customized tariffs T_k^* are differentiable and convex over the equilibrium range.*[19] *Moreover, for all pair of locations (x_a, x_b), and all pairs of reciprocal demands (q_a, q_b) such that $q_a = D_b\left(\rho_b^{*\prime}(q_b)\right)$ and $q_b = D_a\left(\rho_a^{*\prime}(q_a)\right)$, the marginal prices $\rho_a^{*\prime}(q_a)$ and $\rho_b^{*\prime}(q_b)$ jointly satisfy the following Lerner-Wilson formulas*

$$\underbrace{\rho_a^{*\prime}(q_a)\left(1 - \frac{1}{\varepsilon_a\left(\rho_a^{*\prime}(q_a)\right)}\right)}_{\textit{net effect on side-a profits}} + \underbrace{\rho_b^{*\prime}(q_b)\left(1 - \frac{1}{\varepsilon_b\left(\rho_b^{*\prime}(q_b)\right)}\right)}_{\textit{net effect on side-b profits}} = 0. \tag{11}$$

The Lerner-Wilson formulas (11) 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. 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 $\rho_k^{*\prime}$ captures the marginal benefit from adding the extra agents, whereas the semi-inverse-elasticity term $\rho_k^{*\prime}\left[\varepsilon_k\left(\rho_k^{*\prime}\right)\right]^{-1}$ capture its associated infra-marginal losses.

[18]By 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.

[19]Namely, at any $q_l \in [q_{x_l}(\mathbf{s}_k(\underline{v}_k, x_l + .5)), q_{x_l}(\mathbf{s}_k(\overline{v}_k, x_l))]$, $k, l = a, b$, $l \neq k$.

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 Lemma 1 to establish that the mass of x_k-agents that, at the marginal price $\rho_k^{*\prime}(q_k)$, demand q_k agents or more of type x_l coincide with the mass $D_k\left(\rho_k^{*\prime}\left(q_k\right)\right)$ of x_k-agents with vertical type above $\hat{v}_{x_l}\left(\rho_k^{*\prime}|x_k\right)$, where recall that $\hat{v}_{x_l}\left(\rho_k^{*\prime}|x_k\right)$ is the threshold type for whom the utility of interacting with the x_l-agents equals the marginal price ρ_k^*, as defined in (8). Together with reciprocity, Lemma 1 then also implies that the mass q_k of x_l-agents that, at the marginal price $\rho_l^{*\prime}(q_l)$, demand $q_l = D_k\left(\rho_k^{*\prime}\left(q_k\right)\right)$ or more of the x_k-agents coincides with the mass of x_l-agents with vertical type above $\hat{v}_{x_k}\left(\rho_l^{*\prime}|x_l\right)$.

Finally, that the price schedules $\rho_k^*(q_k)$ 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 $\rho_k^{*\prime}(q_k)$ for the q_k-unit of the x_l-agents has to increase with q_k.

The formulas in (11) 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* $\varepsilon_k\left(\cdot|x_l;x_k\right)$ 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 formulas in (11) define a system of structural equations that the econometrician can use to identify the cut-off types. [20] The formulas in (11) can also be used by the econometrician to estimate the distribution of the agents' preferences from data on 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.

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 matches involving partners farther away from an agent's bliss point. The analysis in this section, which, to the best of our knowledge, has

[20] To see this, fix (x_a, x_b) and drop it to ease the notation. For any $q_a \in [0, f_b^x(x_b)]$, the result in Lemma 1 implies that the most economical way of giving the x_a-agents access to q_a agents located at x_b is to match them to all x_b-agents whose vertical type is above $\tilde{v}_b$, with $\tilde{v}_b$ defined by $[1 - F_b^{v|x}(\tilde{v}_b|x_b)]f_b^x(x_b) = q_a$. For any $q_b \in [0, f_a^x(x_a)]$, the marginal price $\rho_b^{*\prime}(q_b)$ is then equal to $u_b(\tilde{v}_b, |x_a - x_b|)$. Given q_a and $\rho_b^{*\prime}(q_b)$, the marginal price $\rho_a^{*\prime}(q_a)$ is then given by equation (11). Once $\rho_a^{*\prime}(q_a)$ is identified, the threshold $v_a = t_b^*((\tilde{v}_b, x_b), x_a)$ is given by the unique solution to $u_a(v_a, |x_a - x_b|) = \rho_a^{*\prime}(q_a)$.

no parallel in the price discrimination literature, has important implications for how policy makers may consider regulating certain platform markets.

Let

$$\sum_{k=a,b} \int_{\Theta_k} \int_{\hat{\mathbf{s}}_k(\theta_k)} u_k\left(v_k, |x_k - x_l|\right) dF_j(\theta_j) dF_k(\theta_k), \tag{12}$$

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 matching demands they induce satisfy the following property: For any two agents with types $\theta_k = (v_k, x_k)$ and $\theta_l = (v_l, x_l)$,

$$\theta_l \in \hat{\mathbf{s}}_k(\theta_k) \quad \Longleftrightarrow \quad u_k\left(v_k, |x_k - x_l|\right) + u_l\left(v_l, |x_k - x_l|\right) \geq 0.$$

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

$$\rho_a^{e\prime}(q_a) + \rho_b^{e\prime}(q_b) = 0,$$

for ant pair of locations (x_a, x_b) and any pair q_a and q_b such that $q_a = D_b\left(\rho_b^{e\prime}(q_b)\right)$ and $q_b = D_a\left(\rho_a^{e\prime}(q_a)\right)$.

Definition 3. (distortions and distance) Distortions on side $k \in \{a, b\}$ decrease (alternatively, increase) with distance if, and only if, for any $\theta_k \in \Theta_k$,

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

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

Hence, fixing the type $\theta_k = (v_k, x_k)$ of a side-k agent, distortions decrease with distance when the difference between the minimal utility asked by a profit-maximizing platform and a welfare-maximizing platform to an x_l-agent to be matched with each side-k agent of type θ_k decreases with the distance between the two agents' locations. Note that the difference in utilities reduces to the difference in the thresholds when the side-l's preferences are invariant to locations. Figure 3 illustrates the case of a market in which distortions decrease with distance and the side-b preferences are invariant to locations.

Definition 4. The side-k demands' inverse-semi-elasticities are increasing (alternatively, decreasing) in distance if, and only if, given any $x_k \in [0, 1]$, $q \in [0, 1]$, and ρ_k', $\rho_k' \left[\varepsilon_k\left(\rho_k' | x_l; x_k\right)\right]^{-1}$ are increasing (alternatively, decreasing) in $|x_l - x_k|$. They are increasing (alternatively, decreasing) in price if, and only if, given any $x_l, x_k \in [0, 1]$ and $q \in [0, f_l^x(x_l)]$, $\rho_k' \left[\varepsilon_k\left(\rho_k' | x_l; x_k\right)\right]^{-1}$ are increasing (alternatively, decreasing) in ρ_k', where ρ_k' is the marginal price asked by the platform to the x_k-agents for the q-th unit of the x_l-matches.

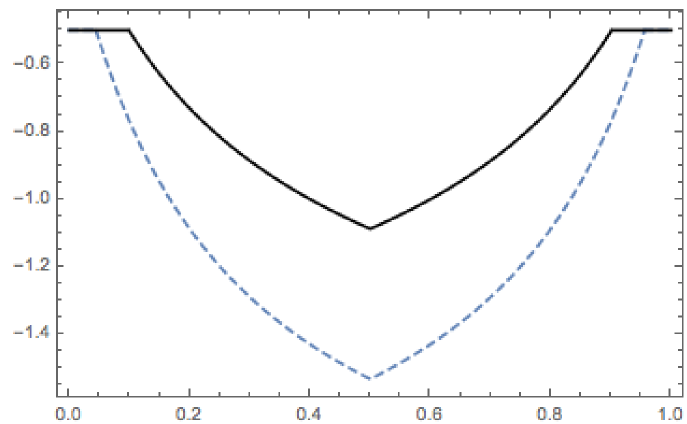

Figure 3: The welfare-maximizing demand threshold $t_a^e(\theta_a, x_b)$ (dashed blue curve) and the profit-maximizing demand threshold $t_a^*(\theta_a, x_b)$ (solid black curve) for agents on side a located at $x_a = .5$ in a market in which side-b preferences are location-invariant and side-a distortions decrease with distance.

Proposition 2. ***(distortions under customized pricing)*** *Either of the following two sets of conditions suffice for distortions on side k to decrease with distance:*[21]

(1.a) the side-k inverse-semi-elasticities are decreasing in distance and increasing in price, whereas the side-l inverse-semi-elasticities are decreasing in both distance and price;

(1.b) u_k submodular, x_l and v_l independent, the hazard rate for F_l^v increasing in v_l, u_l submodular and concave in v_l.

Either of the following two sets of conditions suffice for distortions on side k to increase with distance:

(2.a) the side-k inverse-semi-elasticities are increasing in distance and decreasing in price, whereas the side-l inverse-semi-elasticities are increasing in both distance and price.

(2.b) u_k supermodular, x_l and v_l independent, the hazard rate for F_l^v decreasing in v_l, u_l supermodular and convex in v_l.

Recall that, under the profit-maximizing tariffs, a pair of agents with types θ_k and θ_l is matched if, and only if $\varphi_k(\theta_k, \theta_l) + \varphi_l(\theta_l, \theta_k) \geq 0$, whereas, under the welfare-maximizing tariffs, the same pair is matched if, and only if $u_k(v_k, |x_k - x_l|) + u_l(v_l, |x_k - x_l|) \geq 0$. Also note that the virtual values $\varphi_k(\theta_k, \theta_l)$ differ from the true values $u_k(v_k, |x_k - x_l|)$ by the "markup handicaps"

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

the platform applies to the prices to maximize profits and that, when prices are interior,[22] the latter

[21] Strictly so, if at least one of the properties is strict.

[22] That is, when $\rho_k' \in [u_k(\underline{v}_k, |x_k - x_l|), u_k(\bar{v}_k, |x_k - x_l|)]$.

are related to the inverse-semi-elasticities of the demands by the formula

$$\frac{1-F_k^{v|x}(v_k|x_k)}{f_k^{v|x}(v_k|x_k)}\frac{\partial u_k}{\partial v_k}\left(v_k,|x_k-x_l|\right)=\left.\frac{\rho_k'}{\varepsilon_k\left(\rho_k'|x_l;x_k\right)}\right|_{\rho_k'=u_k(v_k,|x_k-x_l|)}$$

It follows that

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

$$\begin{aligned}&=\tfrac{1-F_k^{v|x}(v_k|x_k)}{f_k^{v|x}(v_k|x_k)}\tfrac{\partial u_k}{\partial v_k}\left(v_k,|x_k-x_l|\right)+\tfrac{1-F_l^{v|x}\left(t_k^*(\theta_k,x_l)|x_l\right)}{f_l^{v|x}\left(t_k^*(\theta_k,x_l)|x_l\right)}\tfrac{\partial u_l}{\partial v_l}\left(t_k^*\left(\theta_k,x_l\right),|x_k-x_l|\right)\\&=\left.\tfrac{\rho_k'}{\varepsilon_k\left(\rho_k'|x_l;x_k\right)}\right|_{\rho_k'=u_k(v_k,|x_k-x_l|)}+\left.\tfrac{\rho_l'}{\varepsilon_l\left(\rho_l'|x_k;x_l\right)}\right|_{\rho_l'=u_l\left(t_k^*(\theta_k,x_l),|x_k-x_l|\right)}\end{aligned}\tag{13}$$

The results in parts (1.a) and (2.a) in the proposition then follow directly from the monotonicities of the handicaps (equivalently of the inverse-semi-elasticities), along with the fact that, fixing $\theta_k=(v_k,x_k)$, as the distance $|x_k-x_l|$ between θ_k and θ_l increases, the maximal price $\rho_k'=u_k\left(v_k,|x_k-x_l|\right)$ that type θ_k is willing to pay to interact with the x_l-agents declines, whereas the price $\rho_l'=u_l\left(t_k^*\left(\theta_k,x_l\right),|x_k-x_l|\right)$ that the platform asks to the x_l-agents to interact with type θ_k increases.[23] The results in parts (1.b) and (2.b), instead, use the specific relation between the demand elasticities and the primitives in (13) to identify *joint conditions* on match values and type distributions guaranteeing that *sum of the markups* on the two sides is monotone in distance, accounting for the joint variations in the marginal prices. Importantly, both the conditions involving the demand elasticities as well as those involving the match values and the type distributions can be empirically estimated using techniques similar to those in the empirical auction literature. Whether the econometrician may prefer working with the elasticities or with the primitive conditions is likely to depend on the application under consideration and, in particular, on the type of information available.

As we show in the next two sections, conditions similar to those in Proposition 2 are also responsible for whether targeting and welfare are higher under customized or uniform pricing, as well as for whether welfare is higher in a decentralized or in a centralized market.

4 Uniform Pricing

Stringent regulations on the transfer of personal data together with restrictions on bundling imposed on certain platforms are expected to hinder the customization of prices and favor instead uniform pricing. In this section, we study platforms' behavior when subject to uniform-pricing obligations.

[23] These properties follow from the fact that the virtual values respect the same rankings as the true values.

Uniform Pricing and Aggregate Demand Elasticities

Definition 5. The tariff T_k is consistent with *uniform pricing* if there exists a collection of *price schedules* $p_k : [0,1]^2 \to \mathbb{R}$, 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

$$T_k(\mathbf{s}_k) = \int_0^1 p_k(q_{x_l}(\mathbf{s}_k)|x_l)dx_l. \tag{14}$$

Hence, under uniform pricing, the tariff offered by the platform to the side-k agents consists of a collection of non-linear price schedules, $(p_k(\cdot|x_l))_{x_l\in[0,1]}$, one for each type of side-l agents, 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, contrary to the case of price customization, the price $p_k(q|x_l)$ is independent of the agent's own characteristics, θ_k, and of the mass of agents from locations other than x_l included in the matching set $\mathbf{s}_k$ (i.e., no bundling).

Suppose the platform is forced to adopt a uniform price schedule $p_a(\cdot|x_b)$ on side a (with marginal schedule $p'_a(\cdot|x_b)$). Recall that, for each location $x_b \in [0,1]$, and each quantity $q \in [0, f_b^x(x_b)]$, 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 $p'_a(q|x_b)$ is equal to

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

where, as in the previous section, $D_a\left(p'_a|x_b;x_a\right)$ 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 $p'_a\left(q|x_b\right)$ have been dropped, to ease the exposition.

The elasticity of the side-a aggregate demand for the q-th unit of the x_b-agents with respect to the marginal price p'_a is then equal to

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

where $\varepsilon_a\left(p'_a|x_b;x_a\right)$ is the elasticity of the demand by the x_a-agents, as defined in (10), and where the expectation is over $X_a = [0,1]$, under the distribution $\bar{H}\left(\cdot|x_b,p'_a\right)$ whose density is equal to

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

Hereafter, we refer to $\bar{\varepsilon}_a\left(\cdot|x_b\right)$ as the *aggregate elasticity* of the side-a demand for the q-th unit of 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 p'_a, where the average is under a distribution that assigns to each location x_a a weight proportional to the mass of agents $D_a\left(p'_a|x_b, x_a\right)$ located at x_a demanding q units, or more, of the x_b-agents.

The next proposition derives the profit-maximizing tariffs T^u_a and T^u_b offered by a platform that is constrained to price uniformly on side a [To ease the exposition, the arguments q and x_b have been dropped from the formulas for the prices and the elasticities].

Proposition 3. ***(uniform pricing)*** *Suppose 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^u_a and ρ^u_b are differentiable and convex over the equilibrium ranges,*[24] *and jointly satisfy the following optimality conditions for all $x_b \in [0,1]$, all $q \leq f^x_b(x_b)$,*

$$\underbrace{p^{u\prime}_a\left(1-\frac{1}{\bar{\varepsilon}_a\left(p^{u\prime}_a\right)}\right)}_{\text{net effect on side-}a\text{ profits}} + \mathbb{E}_{H(\tilde{x}_a|p^{u\prime}_a)}\left[\underbrace{\hat{\rho}^{u\prime}_b(\tilde{x}_a)\left(1-\frac{1}{\hat{\varepsilon}_b\left(\hat{\rho}^{u\prime}_b(\tilde{x}_a)\right)}\right)}_{\text{net effect on side-}b\text{ profits}}\right] = 0, \tag{15}$$

where $H\left(\tilde{x}_a|p^{u\prime}_a\right)$ is the distribution over $X_a = [0,1]$ whose density is given by

$$h\left(x_a|p^{u\prime}_a\right) \equiv \frac{\partial D_a\left(p^{u\prime}_a|x_b;x_a\right)}{\partial\left(p^{u\prime}_a\right)}\left(\frac{\partial \bar{D}_a\left(p^{u\prime}_a|x_b\right)}{\partial\left(p^{u\prime}_a\right)}\right)^{-1},$$

$\hat{q}_b(x_a) \equiv D_a\left(p^{u\prime}_a|x_b;x_a\right)$ 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 $p^{u\prime}_a$, $\hat{\rho}^{u\prime}_b(\tilde{x}_a) \equiv \rho^{u\prime}_b(\hat{q}_b(x_a)|x_a;x_b)$ is the marginal price for the $\hat{q}_b(x_a)$-unit of the x_a-agents charged to the x_b-agents, and $\hat{\varepsilon}_b\left(\hat{\rho}^{u\prime}_b(\tilde{x}_a)\right) \equiv \varepsilon_b\left(\hat{\rho}^{u\prime}_b(\tilde{x}_a)|x_a;x_b\right)$ is the elasticity of the demand of the x_b-agents for the x_a-agents when the marginal price for the $\hat{q}_b(x_a)$-unit of the x_a-agent is equal to $\hat{\rho}^{u\prime}_b(\tilde{x}_a)$.

The result in the proposition provides structural equations similar to those in Proposition 1, 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 x_b-agents 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 $\varepsilon_a\left(\cdot|x_b;x_a\right)$ in the Lerner-Wilson formula (11). 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 x_b-agents is the average (mark-up augmented) price

$$\mathbb{E}_{H(\tilde{x}_a|p^{u\prime}_a)}\left[\underbrace{\hat{\rho}^{u\prime}_b(\tilde{x}_a)\left(1-\frac{1}{\hat{\varepsilon}_b\left(\hat{\rho}^{u\prime}_b(\tilde{x}_a)\right)}\right)}_{\text{net effect on side-}b\text{ profits}}\right]$$

[24] Namely, at any $q_a \in [q_{x_b}(\mathbf{s}_a(\underline{v}_a, x_b+.5)), q_{x_b}(\mathbf{s}_a(\overline{v}_a, x_b))]$ and $q_b \in [q_{x_a}(\mathbf{s}_b(\underline{v}_b, x_b)), q_{x_a}(\mathbf{s}_a(\overline{v}_b, x_b))]$.

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 (11) and (15), 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 $1/\bar{\varepsilon}_a(\cdot|\cdot)$ and the location-specific inverse-elasticities $1/\varepsilon_a(\cdot|\cdot; x_a)$ on side a, as well as the comparison between the average inverse-elasticity $\mathbb{E}_{H(\tilde{x}_a|p_a')}\left[1/\varepsilon_b(\cdot|\tilde{x}_a;\cdot)\right]$ of the x_b-demands for the various x_a-matches and the inverse-elasticities $1/\varepsilon_b(\cdot|\tilde{x}_a;\cdot)$ 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, given the marginal price $p_a^{u\prime}$ for the q-th unit of the x_b agents,

$$\frac{1}{\bar{\varepsilon}_a\left(p_a^{u\prime}\right)} = \mathbb{E}_{H(\tilde{x}_a|p_a^{u\prime})}\left[\frac{1}{\hat{\varepsilon}_a\left(p_a^{u\prime}|\tilde{x}_a\right)}\right]. \tag{16}$$

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 $H\left(x_a|p_a'\right)$ introduced in the proposition.

Next, let $\check{v}_b$ be implicitly defined by

$$\left(1 - F_b^{v|x}(\check{v}_b|x_b)\right) f_b^x(x_b) = q.$$

Note that $\check{v}_b$ denotes 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 $\check{v}_b$ is equal to q. Using the characterization in Lemma 1, we then have that the optimality condition (15) can be re-written as

$$\underbrace{\mathbb{E}_{H(\tilde{x}_a|p_a^{u\prime})}\left[\varphi_a\left(\left(\hat{v}_{x_b}\left(p_a^{u\prime}|\tilde{x}_a\right), \tilde{x}_a\right), (\check{v}_b, x_b)\right)\right] + \underbrace{\mathbb{E}_{H(\tilde{x}_a|p_a^{u\prime})}\left[\varphi_b\left((\check{v}_b, x_b)\right), \left(\hat{v}_{x_b}\left(p_a^{u\prime}|\tilde{x}_a\right), \tilde{x}_a\right)\right]}_{\text{net effect on side-}b\text{ profits}}}_{\text{net effect on side-}a\text{ profits}} = 0. \tag{17}$$

Hereafter, we assume that the left-hand side of (17) is monotone in the marginal price p_a', 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 3. The above property implies that the necessary condition in (17) is also sufficient for optimality.

Now recall that, under price customization (on both sides), the platform matches a 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, when $p'_a = u_a(v_a, |x_b - x_a|)$,

$$\mathbb{E}_{H(\tilde{x}_a|p'_a)}\left[\varphi_a\left(\left(\hat{v}_{x_b}\left(p'_a|\tilde{x}_a\right), \tilde{x}_a\right), \theta_b\right)\right] + \mathbb{E}_{H(\tilde{x}_a|p'_a)}\left[\varphi_b\left(\theta_b, \left(\hat{v}_{x_b}\left(p'_a|\tilde{x}_a\right), \tilde{x}_a\right)\right)\right] \geq 0.$$

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 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 stringent protection of consumer privacy (de facto banning price customization), stand-alone pricing for media content (thus banning bundling), 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 platforms undertake in response to uniform-pricing obligations. The analysis below provides some guidelines when it comes to the effects of uniform and customized pricing on targeting.

Definition 6. (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

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

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

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

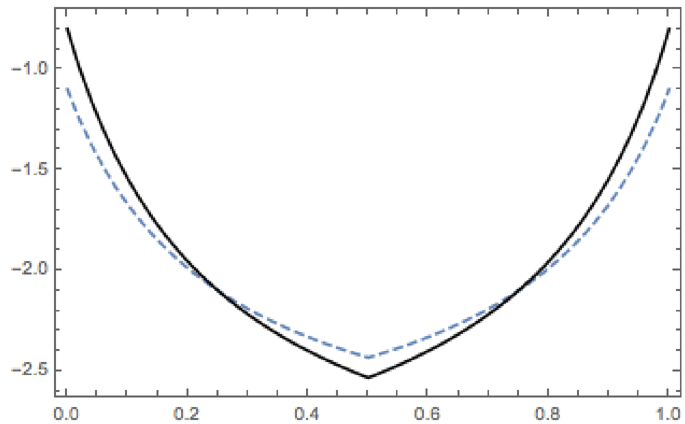

Figure 4: The threshold function $t_a^*(\theta_a, x_b)$ under customized pricing (solid black curve) and uniform pricing $t_a^u(\theta_a, x_b)$ (dashed blue curve) when customized pricing (on both sides) leads to more targeting than uniform pricing (on side a).

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.

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 conditions under which uniform pricing on side a (for short, uniform pricing) leads to more targeting that customized pricing on both sides (for short, customized pricing). For simplicity, the result is for a market in which preferences on side b are location-invariant (the conditions for the more general case are discussed after the proposition).

Proposition 4. ***(comparison: targeting)*** *Suppose preferences on side b are location-invariant.*

1. Uniform pricing (on side a) leads to more (alternatively, less) targeting than customized pricing (on both sides) when the side-a inverse-semi-elasticities are decreasing (alternatively, increasing) in both distance and price.

2. The side-a inverse-semi-elasticities are decreasing (alternatively, increasing) in both distance and price when x_a and v_a are independent, the hazard rate for F_a^v is increasing in v_a, and u_a is submodular and concave in v_a (alternatively, x_a and v_a are independent, the hazard rate for F_a^v is *decreasing in v_a, and u_a is supermodular and convex in v_a).*

Consider Part 1 and fix the side-b location x_b. Under uniform pricing, the elasticity of the aggregate demand by the side-a agents for the q-th unit of the x_b-agents is invariant to the distance $|x_b - x_a|$, as the marginal price for the q-th unit of x_b-matches is the same for all x_a-locations. As a consequence, when the inverse-semi-elasticities of the side-a demands decrease (alternatively, increase) with distance, the marginal price for the q-th unit of x_b-matches charged to the x_a-agents under customized pricing is lower than the corresponding price under uniform pricing 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. The result in Part 2 then uses the characterization of the matching demands in the previous section to translate the result in Part 1 in terms of conditions on match values and type distributions.

Remark 3. To simplify the exposition, Proposition 4 assumes that, on side b, preferences are location-invariant. A similar result, however, applies more generally. 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 $L : [0,1] \times \Theta_b \to \mathbb{R}$ be the function defined by

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

$$= \rho_b'\left(1 - \frac{1}{\varepsilon_b\left(\rho_b'|x_a;x_b\right)}\right)\Bigg|_{\rho_b'=u_b(v_b,|x_a-x_b|)} - \frac{\rho_a'}{\varepsilon_a\left(\rho_a'|x_b;x_a\right)}\Bigg|_{\rho_a'=u_a(t_b^*(\theta_b,x_a),|x_a-x_b|)}$$

As we show in the Appendix, uniform pricing (on side a) leads to more (alternatively, less) targeting than customized pricing, on both sides, if, for any θ_b, the function $L(\cdot|\theta_b)$ is non-decreasing (alternatively, non-increasing) in $|x_b - x_a|$. When preferences on side b are location-invariant, the above property reduces to the monotonicity conditions in the proposition. More generally, uniform pricing (on side a) leads to more targeting than customized pricing (on both sides) when the inverse-semi-elasticities on both sides of the market decrease sufficiently sharply with distance relative to the utility that the side-b agents derive from interacting with the side-a agents when the distance increases.

4.2 Welfare under Uniform and Customized Pricing

The result in Proposition 4 can also be used to study the welfare implications of uniform-pricing obligations. 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 $p_a'(q|x_b)$ 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 at the marginal price p'_a is given by

$$D_a\left(p'_a|x_b;x_a\right)=\left[1-F_a^{v|x}\left(\hat{v}_{x_b}\left(p'_a|x_a\right)|x_a\right)\right]f_a^x(x_a)$$

where, to ease the notation, we dropped $(q|x_b)$ from the arguments of the marginal price $p'_a(q|x_b)$. Now let

$$CD_a\left(p'_a|x_b;x_a\right)=-\frac{\partial^2 D_a\left(p'_a|x_b;x_a\right)}{\partial\left(p'_a\right)^2}\left(\frac{\partial D_a\left(p'_a|x_b;x_a\right)}{\partial\left(p'_a\right)}\right)^{-1}p'_a$$

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 p'_a.[25] Before proceeding, we have to impose the following regularity condition.

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

$$z_a\left(p'_a|x_b;x_a\right)\equiv\frac{p'_a}{2-CD_a\left(p'_a|x_b;x_a\right)}$$

is nondecreasing in the marginal price p'_a for the q-th unit of the x_b-agents.

We then have the following result:

Proposition 5. ***(comparison: welfare)*** *Suppose Condition IR holds and either one of the following alternatives is satisfied:*

1. targeting is higher under uniform pricing than under customized pricing and, for any p'_a and x_b, $CD_a\left(p'_a|x_b;x_a\right)$ declines with $|x_a-x_b|$.

2. targeting is higher under customized pricing than under uniform pricing and, for any p'_a and x_b, $CD_a\left(p'_a|x_b;x_a\right)$ increases with $|x_a-x_b|$.

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 5, 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 4 is key to the result in Proposition 5. 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.

[25] Note that $CD_a(p'_a|x_b;x_a)$ is also the elasticity of the marginal demand $\partial D_a(p'_a|x_b;x_a)/\partial p'_a$ with respect to the marginal price p'_a.

Also note that the result in Proposition 5 is just an illustration of the type of welfare results that Proposition 4 permits. Paralleling the analysis in Proposition 2 in Aguirre et al. (2010), for example, one can also identify primitive conditions under which welfare is higher under price customization than under uniform pricing, as well as conditions under which price customization impacts negatively one side of the market and positively the other.

5 Centralized vs Decentralized Markets

In the market for media content, several analysts believe that 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, thus bypassing current market intermediaries such as cable TV providers. This corresponds to a decentralized market (where matching is governed by agents on one side of the market posting prices for their matches with the other side), as opposed to the centralized markets considered so far in this paper (where matching and pricing are controlled by a platform). As it turns out, the results above also permit us to assess the effects of market decentralization on matching allocations and welfare.

To illustrate, consider a matching market where side a is populated by "buyers" and side b is populated by "sellers". Unlike in the baseline model, suppose sellers post prices independently and sell the "match" directly to the buyer (i.e., the market is decentralized). Absent data about consumers' tastes (for instance, due to privacy regulation), we can assume that buyers' locations are private information, whereas sellers' locations (their product variety) is observable by the buyers. In this environment, provided sellers are not able to monitor the purchases that each buyer makes with the other sellers (which amounts to no bundling), customized pricing on side a is infeasible. This is to be contrasted to a centralized matching market, where, as explained in the proof of Lemma 1 in the Appendix, location-dependent prices can be implemented by the platform (even when locations are private information) through bundling.

Therefore, a move from a centralized to a decentralized structure is akin to the imposition of uniform-pricing obligations in a centralized market, in that sellers have to set (for lack of information and pricing instruments) the same prices for all buyers.[26] But there is more: 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 in a centralized market. Combining this extra effect with the ones identified in Propositions 4 and 5 permits us to conclude that, whenever uniform pricing increases the welfare of the side-a agents, so does market decentralization. The conditions in Proposition 5 are therefore sufficient to conclude that consumer surplus (i.e., the welfare of the side-a agents) is higher in a

[26]To facilitate the comparison with the centralized case, we assume that, in a decentralized market, each seller of variety x_b sell a single unit to each buyer (unit demands). The marginal price schedule $p'_a(q|x_b)$ should then be interpreted as the inverse *aggregate* market supply curve resulting from the aggregation of the behavior of different sellers posting different stand-alone prices as a function of their different marginal costs (captured by the term v_b).

decentralized market than in a centralized one. This observation may help policymakers identify markets where the transition to a decentralized structure should be promoted, for example through fiscal incentives and/or direct subsidies.

6 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 characterizes properties of the matching demands when 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 plans. 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 agent'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 that hinder platforms' possibility to condition prices on agents' "locations", as in the case of privacy regulations preventing online retailers from conditioning prices on buyers' age, gender, of physical location. 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, online shopping. As mentioned in the introduction, recent regulations requiring consumers' consent for the diffusion of personal information are expected to hinder price customization when third-party data are needed. Perhaps surprisingly, our analysis shows that this may either increase or decrease targeting levels, depending on testable characteristics of consumer demand. Related conditions can also be used to evaluate whether or not the imposition of uniform-pricing obligations increase consumer welfare.

Another natural application of our framework is the market for online advertising (see, among others, Bergemann and Bonatti (2011) for an overview of such markets). Ad exchanges 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 enable 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 often seen with suspicion by consumers and regulators. The policy debate about the desirability of regulations imposing uniform pricing, or about the pros and cons of market decentralization, lacks a formal model shedding light on how matching demands and welfare are affected by such changes. Our paper contributes to such debate by offering a stylized yet flexible framework that one can use to study both the distortions associated with price customization, as well as the market outcomes under 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)). Such industry is evolving fast and many analysts predict a transition to a market structure whereby viewers will purchase content 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 plausible 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 is an important step for future research.

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 such richer instruments is another interesting direction for future work (see Jeon, Kim, and Menicucci (2016) for related ideas).

Lastly, in future work, it would be desirable to extend the analysis to accommodate for "within-side" network effects (e.g., congestion and limited attention) and dynamics (whereby agents gradually learn the attractiveness of the partners and platforms indirectly control the speed of such learning with their pricing strategies). All the above extensions are challenging but worth exploring.

7 Appendix

In this Appendix, we provide the proofs for all the results in the main text. The proofs are established for more general environments in which, in addition to possessing private information about the vertical dimensions, v_k, the agents may possess private information also about their horizontal dimensions, i.e., about their locations x_k. In particular, we consider the following four scenarios:

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

All the results in the main body are for Scenario (i). Below, we discuss how the results extend to Scenarios (ii)-(iv) provided a certain combination of the following two conditions holds.

Condition 2. [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 that the vertical dimensions v_k be drawn independently from the locations x_k. This condition implies that knowing an agent's "bliss point" carries no information about the overall importance the agent assigns to matching with agents from the opposite side.

Condition 3. [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.[27]

Proof of Lemma 1. The proof below establishes the following result, for which the claim in the main text is a special case. *Suppose the environment satisfies the properties of one of the following four cases: Scenario (i); Scenario (ii) along with Conditions I_a and S_b; Scenario (iii) along with Conditions S_a and I_b; Scenario (iv) along with Conditions 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 $t_k^* : \Theta_k \times [0, 1] \to V_l$ such that*

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

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

with the threshold function t_k^* non-increasing in v_k and 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 $x_k \in [0, 1]$, $\mathcal{S}_k(x_k) = \Sigma(\Theta_l)$.

Conditions I_k and S_k 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 obtains the same profits as when locations are public on both sides, as in Scenario (i). To gain some intuition, 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 $\mathbf{s}_a^*(v_a, x_a + d)$, $d \in [0, 1/2]$. This occurs because 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 matching plans and 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. Note that, to guarantee that agents reveal their locations when the latter are the agents' private information, the platform may need to restrict the set of possible customizations at each location, $\mathcal{S}_k(x_k)$, to coincide with the matching sets demanded under Scenario (i).[28] The above properties, together with a judicious restriction on the set of possible customizations $\mathcal{S}_k(x_k)$, imply that the matching demands and payments induced in Scenario (i) are implementable also under Scenario (ii). A symmetric situation applies to Scenario (iii). Arguments similar to the ones above for Scenarios (ii) and (iii) imply that, in Scenario (iv), where locations are private on both sides, when the platform offers the same matching plans and tariffs as in Scenario (i) to each side, agents continue to choose the same matching sets as in Scenario (i), provided that Condition S_k holds on both sides of the market.

We establish the above results 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.

[28] In the absence of such restrictions, an agent of type $\theta_k = (v_k, x_k)$ misrepresenting his location to be $x_k' \neq x_k$ may find it optimal to select a matching set that no agent located at x_k' would have demanded under Scenario (i).

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

$$\begin{aligned} p_k(\theta_k) = &\int_{\mathbf{s}_k(\theta_k)} u_k\left(v_k, |x_k - x_l|\right) dF_l(\theta_l) - \int_{\underline{v}_k}^{v_k} \int_{\mathbf{s}_k(y,x_k)} \frac{\partial u_k}{\partial v}\left(y, |x_k - x_l|\right) dF_l(\theta_l) dy, \\ &- U_k(\underline{v}_k, x_k), \end{aligned} \tag{18}$$

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

Using (18), the platform's profits under any incentive-compatible mechanism can then be written as

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

Using the definition of the virtual-value functions $\varphi_k\left(\theta_k, \theta_l\right)$ 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

$$\sum_{k=a,b} \int_{\Theta_k} \left\{ \int_{\mathbf{s}_k(\theta_k)} \varphi_k\left(\theta_k, \theta_l\right) dF_l(\theta_l) \right\} dF_k(\theta_k) \tag{19}$$

subject to the reciprocity condition

$$\theta_l \in \mathbf{s}_k(\theta_k) \Longleftrightarrow \theta_k \in \mathbf{s}_l(\theta_l), \quad l, k \in \{a, b\},\ k \neq l. \tag{20}$$

Hereafter, we first describe the matching sets that maximize (19) subject to the above reciprocity condition and then show that, under the assumptions in the lemma, 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 (19) subject to (20), and where the functions $p_k(\cdot)$ are as in (18), 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 $\Theta_k \times \Theta_l$:

$$\nu_k(E) \equiv \int_E m_k(\theta_k, \theta_l) dF_k(\theta_k) dF_l(\theta_l). \tag{21}$$

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 (19) can be rewritten as

$$\begin{aligned} &\sum_{k,l=a,b,\, l \neq k} \int_{\Theta_k \times \Theta_l} \varphi_k\left(\theta_k, \theta_l\right) d\nu_k(\theta_k, \theta_l) \\ &= \int_{\Theta_k \times \Theta_l} \triangle_k(\theta_k, \theta_l) m_k(\theta_k, \theta_l) dF_k(\theta_k) dF_l(\theta_l), \end{aligned} \tag{22}$$

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

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

Note that the functions $\triangle_a(\theta_a, \theta_b) = \triangle_b(\theta_b, \theta_a)$ 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 (22) 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

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

That the virtual values $\varphi_k(\theta_k, \theta_l)$ are strictly increasing in v_k, $k, l = a, b$, $l \neq k$, then implies that the matching rule that maximizes (19) subject to the reciprocity condition (20) can be described by means of a collection of threshold functions $t_k^* : \Theta_k \times X_l \to V_l$, $k, l = a, b$, $l \neq k$, 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 $\triangle_k(\theta_k, (\underline{v}_l, x_l)) > 0$, $t_k^*(\theta_k, x_l) = \bar{v}_l$ if $\triangle_k(\theta_k, (\bar{v}_l, x_l)) < 0$, and $t_k^*(\theta_k, x_l)$ is the unique solution to $\triangle_k(\theta_k, (t_k^*(\theta_k, x_l), x_l)) = 0$ if

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

That the virtual values $\varphi_k(\theta_k, \theta_l)$ are increasing in v_k and decreasing in $|x_k - x_l|$ also implies that, for ay $x_k, x_l \in [0,1]^2$, the threshold $t_k^*(\theta_k, x_l)$ is decreasing in v_k, and that, for any v_k, $t_k^*(\theta_k, x_l)$ is non-decreasing in $|x_l - x_k|$.

Equipped with the above result, we now show that, in each of the environments stated in the generalized version of the lemma reported above, 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 (18), 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 7. (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 8. (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 property is then true (the proof is standard and hence omitted):

Property 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 (18).*

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 (18), 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 general version of the lemma presented above, 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 and S_b 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 $(x_a, \hat{x}_a, v_a, v_a) \in X_a^2 \times V_a^2$. The above inequality is equivalent to

$$\int_{\underline{v}_a}^{v_a} \int_{\mathbf{s}_a(y,x_a)} \frac{\partial u_a}{\partial v}\left(y, |x_a - x_b|\right) dF_b(\theta_b) dy \geq \int_{\underline{v}_a}^{\hat{v}_a} \int_{\mathbf{s}_a(y,\hat{x}_a)} \frac{\partial u_a}{\partial v}\left(y, |\hat{x}_a - x_b|\right) dF_b(\theta_b) dy \tag{23}$$
$$+ \int_{\mathbf{s}_a(\hat{v}_a,\hat{x}_a)} [u_a\left(v_a, |x_a - x_b|\right) - u_k\left(\hat{v}_a, |\hat{x}_a - x_b|\right)] dF_b(b).$$

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

$$\int_{\mathbf{s}_a(v_a,x_a)} \frac{\partial u_a}{\partial v}\left(v_a, |x_a - x_b|\right) dF_b(\theta_b) = \int_{d \in [0,1/2]} \frac{\partial u_a\left(v_a, d\right)}{\partial v} dW(d; \theta_a), \tag{24}$$

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 (24) 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 $v_a = v_a'$.[29] This means that

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

By the same arguments,

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

Furthermore, because virtual values respect the same ranking as the true values, the threshold functions $t_k^*(\theta_k, x_l)$ are non-decreasing in the distance $|x_l - x_k|$. In turn, this implies that

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

It follows that the right hand side of (23) is smaller than

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

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

[29]Conditions 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.

Finally, consider an environment satisfying the properties of Scenario (iv) along with Conditions 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 general version of the lemma reported above, 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 (18) 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 we wanted to establish 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

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

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

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

the set of possible customizations is given by

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

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

$$\rho_k(q|x_l; x_k) = q u_k\left(v_k(q; x_k, x_l), |x_k - x_l|\right) - \int_{\underline{v}_k}^{v_k(q; x_k, x_l)} q_{x_l}\left(\mathbf{s}_k(y, x_k)\right) \frac{\partial u_k}{\partial v}\left(y, |x_k - x_l|\right) dy - \underline{T}_k(x_k) \tag{25}$$

where

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

Any agent selecting the plan $(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathcal{S}_k(x_k))$ and then choosing a matching set $\mathbf{s}_k \notin \mathcal{S}_k(x_k)$ 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 $(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathcal{S}_k(x_k))$ 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 $(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot;x_k), \mathcal{S}_k(x_k))$ 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

$$\underline{T}_k(x_k) + \int_0^1 \left[q_{x_l}(\mathbf{s}_k(v_k,x_k))u_k\left(v_k,|x_k-x_l|\right) - \int_{\underline{v}_k}^{v_k} q_{x_l}(\mathbf{s}_k(y,x_k))\frac{\partial u_k}{\partial v}\left(y,|x_k-x_l|\right)dy\right]dx_l - \underline{T}_k(x_k)$$

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

$$= p_k(\theta_k),$$

exactly as in the direct mechanism. That each type θ_k maximizes his payoff by selecting the plan $(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot;x_k), \mathcal{S}_k(x_k))$ 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 $(\underline{\mathbf{s}}_k(\hat{x}_k), \underline{T}_k(\hat{x}_k), \rho_k(\cdot|\cdot;\hat{x}_k), \mathcal{S}_k(\hat{x}_k))$ 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, $\mathcal{S}_k(x_k) = \Sigma(\Theta_l)$, all x_k) recall that, in this case, each side-k agent located at x_k can be induced to select the matching plan $(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot;x_k), \mathcal{S}_k(x_k))$ 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 $\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot;x_k), \mathcal{S}_k(x_k))$ is selected, even if $\mathcal{S}_k(x_k) = \Sigma(\Theta_l)$, because the price schedules $\rho_k(\cdot|\cdot;x_k)$ satisfy (25), 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 1. Fix a pair of locations $x_a, x_b \in [0,1]$. From Lemma 1, the profit-maximizing tariffs are customized and induce agents to select matching sets satisfying the threshold property of Lemma 1. Furthermore, from the proof of Lemma 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 $\triangle_k(\theta_k, (\underline{v}_l, x_l)) > 0$, $t_k^*(\theta_k, x_l) = \bar{v}_l$ if $\triangle_k(\theta_k, (\bar{v}_l, x_l)) < 0$, and $t_k^*(\theta_k, x_l)$ is the unique solution to $\triangle_k(\theta_k, (t_k^*(\theta_k, x_l), x_l)) = 0$ if

$$\triangle_k(\theta_k, (\underline{v}_l, x_l)) \le 0 \le \triangle_k(\theta_k, (\bar{v}_l, x_l)).$$

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

$$u_k\left(v_k, |x_k - x_l|\right) = \rho_k'\left(q_k|x_l;x_k\right). \tag{26}$$

Now let $\hat{v}_{x_l}\left(\rho'_k|x_k\right)$ be the unique solution to (26) and $v'_l(q_k;x_l)$ be the unique solution to

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

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

$$t_k^*\left(\left(\hat{v}_{x_l}\left(\rho'_k|x_k\right),x_k\right),x_l\right)=v'_l(q_k;x_l)$$

and that

$$\varphi_k\left(\left(\hat{v}_{x_l}\left(\rho'_k|x_k\right),x_k\right),\left(v'_l(q_k;x_l),x_l\right)\right)+\varphi_l\left(v'_l(q_k;x_l),x_l),\left(\hat{v}_{x_l}\left(\rho'_k|x_k\right),x_k\right)\right)=0. \tag{27}$$

Lastly, observe that, for any such q_k,

$$\frac{\rho'_k\left(q_k|x_l;x_k\right)}{\varepsilon_k\left(\rho'_k|x_l;x_k\right)}=\frac{1-F_k^{v|x}(\hat{v}_{x_l}\left(\rho'_k|x_k\right)|x_k)}{f_k^{v|x}(\hat{v}_{x_l}\left(\rho'_k|x_k\right)|x_k)}\frac{\partial u_k}{\partial v}\left(\hat{v}_{x_l}\left(\rho'_k|x_k\right),|x_k-x_l|\right). \tag{28}$$

Using the definition of φ_k from the main text together with (26) and (28), we then have that, for any such q_k,

$$\varphi_k\left(\left(\hat{v}_{x_l}\left(\rho'_k|x_k\right),x_k\right),\left(v'_l(q_k;x_l),x_l\right)\right)=\rho'_k\left(q_k|x_l;x_k\right)\left[1-\frac{1}{\varepsilon_k\left(\rho'_k|x_l;x_k\right)}\right]. \tag{29}$$

Likewise, when $q_l=\left[1-F_k^{v|x}\left(\hat{v}_{x_l}\left(\rho'_k|x_k\right)|x_k\right)\right]f_k^x(x_k)$,

$$\varphi_l\left(\left(v'_l(q_k;x_l),x_l\right),\left(\hat{v}_{x_l}\left(\rho'_k|x_k\right),x_k\right)\right)=\rho'_l\left(q_l|x_k;x_l\right)\left[1-\frac{1}{\varepsilon_l\left(\rho'_l|x_k;x_l\right)}\right]. \tag{30}$$

Combining (29) and (30) with (27), we obtain the result in the proposition. Q.E.D.

Proof of Proposition 2. Take any $(\theta_k,x_l)\in\Theta_k\times[0,1]$ for which $t_k^*\left(\theta_k,x_l\right),t_k^e\left(\theta_k,x_l\right)\in Int[V_l]$. Recall that, in this case, $t_k^*\left(\theta_k,x_l\right)$ is given by the unique solution to

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

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

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

This means that $u_l\left(t_k^*\left(\theta_k,x_l\right),|x_l-x_k|\right)-u_l\left(t_k^e\left(\theta_k,x_l\right),|x_l-x_k|\right)$ is given by the expression in (13) in the main text. That the conditions in (1.a) and (2.a) suffice for the distortions on side k to decrease (alternatively, increase) with distance then follows from the triangular property, along with the fact that, fixing $\theta_k=(v_k,x_k)$, $\rho'_k=u_k\left(v_k,|x_k-x_l|\right)$ is decreasing in $\mu(x_k,x_l)\equiv|x_k-x_l|$, whereas $\rho'_l=u_l\left(t_k^*\left(\theta_k,x_l\right),|x_k-x_l|\right)$ is increasing in $\mu(x_k,x_l)$. In turn, these properties follow from the fact that the virtual values respect the same rankings as the true values.

Then consider parts (1.b) and (2.b). 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 the second equality in (13) with respect to x_l at some x_l for which $\mu(x_k, x_l) \in \left(0, \frac{1}{2}\right)$, we obtain that

$$\frac{\partial}{\partial x_l}\left[u_l\left(t_k^*\left(\theta_k, x_l\right), |x_l - x_k|\right) - u_l\left(t_k^e\left(\theta_k, x_l\right), |x_l - x_k|\right)\right] = \frac{1-F_k^v(v_k)}{f_k^v(v_k)} \cdot \frac{\partial^2 u_k(v_k, |x_k - x_l|)}{\partial v \partial \mu}$$

$$-\frac{\left(\lambda_l^v\right)'\left(t_k^*(\theta_k, x_l)\right)}{\lambda_l^v\left(t_k^*(\theta_k, x_l)\right)^2}\frac{\partial t_k^*}{\partial x_l}\left(\theta_k, x_l\right)\frac{\partial u_l\left(t_k^*(\theta_k, x_l), |x_k - x_l|\right)}{\partial v_l}$$

$$+\frac{1}{\lambda_l^v\left(t_k^*(\theta_k, x_l)\right)}\left(\frac{\partial^2 u_l\left(t_k^*(\theta_k, x_l), |x_k - x_l|\right)}{\partial v_l^2}\frac{\partial t_k^*(\theta_k, x_l)}{\partial x_l} + \frac{\partial^2 u_l\left(t_k^*(\theta_k, x_l), |x_k - x_l|\right)}{\partial v_l \partial \mu}\right),$$

where $\lambda_l^v(v_l) \equiv f_l^v(v_l)/[1 - F_l^v(v_l)]$ is a shortcut for the hazard rate of the marginal distribution of the vertical dimension v_l, when v_l and x_l are independent on side $l \neq k$. By Lemma 1, $\frac{\partial t_k^*}{\partial x_l}\left(\theta_k, x_l\right) \geq 0$. Therefore, under the conditions of Part (1.b), the expression above is negative, whereas, under the conditions of Part (2.b), the above expression is positive. That distortions decrease (alternatively, increase) with distance under the conditions of Part (1.b) (alternatively, (2.b)) then follows from the above properties. Q.E.D.

Proof of Proposition 3. 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

$$\int_0^1 \int_0^{f_b^x(x_b)} \bar{D}_a\left(p_a'(q|x_b)|x_b\right) p_a'(q|x_b) dq dx_b$$

$$+\int_0^1 \int_0^1 \int_0^{f_a^x(x_a)} D_b\left(\rho_b'(q|x_a; x_b)|x_a; x_b\right) \rho_b'(q|x_a; x_b) dq dx_a dx_b,$$

subject to the feasibility constraint (3).

For any x_b, $q \leq f_b^x(x_b)$, and $p_a'(q|x_b)$, let

$$\hat{v}_{x_b}\left(p_a'|x_a\right) = \begin{cases} v_a \; s.t. \; u_a(v_a, |x_a - x_b|) = p_a' & \text{if } p_a' \in [u_a(\underline{v}_a, |x_a - x_b|), u_a(\bar{v}_a, |x_a - x_b|)] \\ \underline{v}_a & \text{if } p_a' < u_a(\underline{v}_a, |x_a - x_b|) \\ \bar{v}_a & \text{if } p_a' > u_a(\bar{v}_a, |x_a - x_b|). \end{cases} \tag{31}$$

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 $p_a'(q|x_b)$ is equal to

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

Also, for any $q \leq f_b^x(x_b)$, recall that $v_b'(q;x_b)$ is the unique solution to $\left[1 - F_b^{v|x}\left(v_b'(q;x_b)|x_b\right)\right] f_b^x(x_b) = q$. 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 Lemma 1). Now for each $x_a, x_b \in [0,1]$, each $q \leq f_b^x(x_b)$, let

$$\hat{q}_b(q;x_a;x_b) \equiv D_a\left(p_a'(q|x_b)|x_b;x_a\right).$$

Given $p_a'(q|x_b)$, 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

$$\rho_b'(\hat{q}_b(q;x_a;x_b)|x_a;x_b) = u_b(v_b'(q;x_b), |x_b - x_a|). \tag{32}$$

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

The function $p_a'(q|x_b) : \mathbb{R} \times [0,1] \to \mathbb{R}$ thus uniquely defines the matching sets on both sides of the market. Now, from the arguments in the proof of Lemma 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

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

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

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

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

$$\mathcal{C}\left[p_a'(q|x_b)\right] \equiv$$

$$-\int_0^1\left\{u_b\left(v_b'(q;x_b), |x_b - x_a|\right) - \frac{1 - F_b^{v|x}(v_b'(q;x_b)|x_b)}{f_b^{v|x}(v_b'(q;x_b)|x_b)} \cdot \frac{\partial u_b}{\partial v}\left(v_b'(q;x_b), |x_b - x_a|\right)\right\} D_a\left(p_a'(q|x_b)|x_b;x_a\right) dx_a$$

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 $(x_b, v_b) \in [0,1] \times V_b$), $p_a'(q|x_b)$ so as to maximize

$$\bar{D}_a\left(p_a'(q|x_b)|x_b\right) p_a'(q|x_l) - \mathcal{C}\left[p_a'(q|x_b)\right].$$

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

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

where

$$\mathcal{C}'\left[p'_a(q_a|x_b)\right]$$

$$=-\int_0^1\left\{u_b\left(v'_b(q;x_b),|x_b-x_a|\right)-\frac{1-F_b^{v|x}(v'_b(q;x_b)|x_b)}{f_b^{v|x}(v'_b(q;x_b)|x_b)}\cdot\frac{\partial u_b}{\partial v}\left(v'_b(q;x_b),|x_b-x_a|\right)\right\}\frac{\partial D_a(p'_a(q|x_b)|x_b;x_a)}{\partial(p'_a)}dx_a.$$

Now observe that (32) implies that

$$u_b\left(v'_b(q;x_b),|x_b-x_a|\right)-\frac{1-F_b^{v|x}(v'_b(q;x_b)|x_b)}{f_b^{v|x}(v'_b(q;x_b)|x_b)}\cdot\frac{\partial u_b}{\partial v}\left(v'_b(q;x_b),|x_b-x_a|\right)$$

$$=\rho'_b(\hat{q}_b(q;x_a;x_b)|x_a;x_b)\left(1-\frac{1}{\varepsilon_b\left(\rho'_b(\hat{q}_b(q;x_a;x_b)|x_a;x_b)|x_a;x_b\right)}\right).$$

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

$$p'_a(q|x_b)\left[1-\frac{1}{\bar{\varepsilon}_a(p'_a(q|x_b)|x_b)}\right]$$

$$+\mathbb{E}_{H(\tilde{x}_a|x_b,p'_a(q|x_b))}\left[\rho'_b(\hat{q}_b(q;\tilde{x}_a;x_b)|\tilde{x}_a;x_b)\left(1-\frac{1}{\varepsilon_b\left(\rho'_b(\hat{q}_b(q;\tilde{x}_a;x_b)|\tilde{x}_a;x_b)|\tilde{x}_a;x_b\right)}\right)\right]=0,$$

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

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

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

Proof of Proposition 4. As explained in the main text, the proof below is for the more general case in which the side-b preferences may depend on the locations.

Fix $\theta_b=(v_b,x_b)$ and let $q=f_b^x(x_b)\left[1-F_b^v(v_b)\right]$. The result in Proposition 3 implies that, under uniform pricing on side a and customized pricing on side b, for any $x_a\in X_a$ such that $t_b^u(\theta_b,x_a)\in$ Int$[V_a]$, $t_b^u(\theta_b,x_a)$ is such that

$$u_a(t_b^u(\theta_b,x_a),|x_b-x_a|)-\mathbb{E}_{H(\tilde{x}_a|x_b,p_a^{u\prime})}\left[\frac{1-F_a^{v|x}\left(\hat{v}_{x_b}(p_a^{u\prime}|\tilde{x}_a)\right)}{f_a^{v|x}\left(\hat{v}_{x_b}(p_a^{u\prime}|\tilde{x}_a)\right)}\cdot\frac{\partial u_a}{\partial v}\left(\hat{v}_{x_b}\left(p_a^{u\prime}|\tilde{x}_a\right),|\tilde{x}_a-x_b|\right)\right]$$

$$+\mathbb{E}_{H(\tilde{x}_a|x_b,p_a^{u\prime})}\left[\varphi_b\left(\theta_b,\left(\hat{v}_{x_b}\left(p_a^{u\prime}|\tilde{x}_a\right),\tilde{x}_a\right)\right)\right]=0, \tag{33}$$

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

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

and where $p_a^{u\prime}$ is a shortcut for $p_a^{u\prime}(q|x_b)$ with the latter equal to $p_a^{u\prime}(q|x_b) = u_a(t_b^u(\theta_b, x_a), |x_a - x_b|)$. Note that, to arrive at (33), we used the result in Proposition 3 along with the property in (16) and the fact that, for any x_a such that $\hat{v}_{x_b}(p_a^{u\prime}|x_a) \notin Int[V_a]$, $h(x_a|x_b, p_a^{u\prime}) = 0$, whereas for any x_a such that $\hat{v}_{x_b}(p_a^{u\prime}|x_a) \in Int[V_a]$,

$$\frac{p_a^{u\prime}}{\varepsilon_a(p_a^{u\prime}|x_b; x_a)} = \frac{1 - F_a^{v|x}(\hat{v}_{x_b}(p_a^{u\prime}|x_a))}{f_a^{v|x}(\hat{v}_{x_b}(p_a^{u\prime}|x_a))} \cdot \frac{\partial u_a}{\partial v}(\hat{v}_{x_b}(p_a^{u\prime}|x_a), |x_a - x_b|).$$

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

$$\rho_b'(\hat{q}_b(q; x_a; x_b)|x_a; x_b)\left(1 - \frac{1}{\varepsilon_b\left(\rho_b'(\hat{q}_b(q; x_a; x_b)|x_a; x_b)|x_a; x_b\right)}\right)$$

$$= \varphi_b(\theta_b, (\hat{v}_{x_b}(p_a^{u\prime}|x_a), x_a)),$$

as shown in the proof of Proposition 3.

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 $t_b^*(\theta_b, x_a) \in Int[V_a]$, the threshold $t_b^*(\theta_b, x_a)$ is such that

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

It is then immediate that, for any x_a such that

$$-\mathbb{E}_{H(\tilde{x}_a|x_b, p_a^{u\prime})}\left[\frac{1 - F_a^{v|x}(\hat{v}_{x_b}(p_a^{u\prime}|\tilde{x}_a))}{f_a^{v|x}(\hat{v}_{x_b}(p_a^{u\prime}|\tilde{x}_a))} \cdot \frac{\partial u_a}{\partial v}(\hat{v}_{x_b}(p_a^{u\prime}|\tilde{x}_a), |x_b - \tilde{x}_a|)\right]$$

$$+\mathbb{E}_{H(\tilde{x}_a|x_b, p_a^{u\prime})}\left[\varphi_b(\theta_b, (\hat{v}_{x_b}(p_a^{u\prime}|\tilde{x}_a), \tilde{x}_a))\right]$$

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

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

$$-\mathbb{E}_{H(\tilde{x}_a|x_b, p_a^{u\prime})}\left[\frac{1 - F_a^{v|x}(\hat{v}_{x_b}(p_a^{u\prime}|\tilde{x}_a))}{f_a^{v|x}(\hat{v}_{x_b}(p_a^{u\prime}|\tilde{x}_a))} \cdot \frac{\partial u_a}{\partial v}(\hat{v}_{x_b}(p_a^{u\prime}|\tilde{x}_a), |x_b - \tilde{x}_a|)\right]$$

$$+\mathbb{E}_{H(\tilde{x}_a|x_b, p_a^{u\prime})}\left[\varphi_b(\theta_b, (\hat{v}_{x_b}(p_a^{u\prime}|\tilde{x}_a), \tilde{x}_a))\right]$$

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

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

Also note that, by virtue of reciprocity, $t_b^u(\theta_b, x_a) \leq t_b^*(\theta_b, x_a)$ if and only if

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

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

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

The above properties imply that uniform pricing (on side a) leads to more (alternatively, less) targeting than customized pricing (on both sides), if, for any θ_b, the function

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

$$= \rho_b'\left(1 - \frac{1}{\varepsilon_b\left(\rho_b'|x_a;x_b\right)}\right)\Bigg|_{\rho_b'=u_b(v_b,|x_a-x_b|)} - \frac{\rho_a'}{\varepsilon_a\left(\rho_a'|x_b;x_a\right)}\Bigg|_{\rho_a'=u_a(t_b^*(\theta_b,x_a),|x_a-x_b|)}$$

is non-decreasing (alternatively, non-increasing) in the distance $|x_a - x_b|$.

Fixing θ_b, the function $L(x_a|\theta_b)$ is nondecreasing in $|x_a - x_b|$ when the side-a inverse-semi-elasticities are decreasing in distance and in price and the side-b preferences are invariant to distance. It is non-increasing in $|x_a - x_b|$ when the side-a inverse-semi-elasticities are increasing in distance and in price and the side-b preferences are invariant to distance, These properties establish the result in Part 1 in the proposition. The result in Part 2 then follows from the result in Part 1 along with the fact that the side-a inverse-semi-elasticities are decreasing (alternatively, increasing) in both distance and price when x_a and v_a are independent, the hazard rate for F_a^v is increasing in v_a, and u_a is submodular and concave in v_a (alternatively, x_a and v_a are independent, the hazard rate for F_a^v is decreasing in v_a, and u_a is supermodular and convex in v_a). Q.E.D.

Proof of Proposition 5. The proof follows from the combination of the results in Proposition 4 with the results in Proposition 1 in Aguirre et al (2010). When the environment satisfies the conditions in Part 1 of Proposition 4, 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 p_a', the convexity $CD_a\left(p_a'|x_b; x_a\right)$ 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 4, 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 p_a', the convexity $CD_a\left(p_a'|x_b; x_a\right)$ 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.

References

[1] Abdulkadiroglu, A., P. Pathak, and A. Roth, 2005a, "Boston Public School Matching," *American Economic Review*, 95(2), pp. 368-71.

[2] Abdulkadiroglu, A., P. Pathak, and A. Roth, 2005b, "The New York City High School Match," *American Economic Review*, 95(2), pp. 364-367.

[3] Abdulkadiroglu, A., and T. Sonmez, 2003, "School Choice: A Mechanism Design Approach," *American Economic Review*, 93(3), pp. 729-47.

[4] Aguirre, I., S. Cowan, and J. Vickers, 2010, "Monopoly Price Discrimination and Demand Curvature," *American Economic Review*, 100(4), pp. 1601-15.

[5] Ambrus, A., E. Calvano, and M. Reisinger, 2016, "Either or Both Competition: A "Two-Sided" Theory of Advertising with Overlapping Viewerships," *American Economic Journal: Microeconomics*, 8, pp. 189-222.

[6] Armstrong, M., 2013, "A More General Theory of Commodity Bundling," *Journal of Economic Theory*, 148, pp. 448-472.

[7] Balestrieri, F., and S. Izmalkov, 2015, "Informed Seller in a Hotelling Market," Working Paper, New Economic School.

[8] 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. Chelthenam, UK, Northampton, MA, USA.

[9] Bergemann, D., and A. Bonatti, 2011, "Targeting in Advertising Markets: Implications for Offline versus Online Media," *Rand Journal of Economics*, 42(3), pp. 417-443.

[10] Bergemann, D., and A. Bonatti, 2015, "Selling Cookies," *American Economic Journal: Microeconomics*, 7(3), pp. 259-294.

[11] Bergemann, D., B. Brooks, and S. Morris, 2015, "The Limits of Price Discrimination," *American Economic Review*, 105(3), pp. 921-57.

[12] Bulow, J., and D. Roberts, 1989, "The Simple Economics of Optimal Auctions." *Journal of Political Economy*, 97, pp. 1060-1090.

[13] Cox, H., B. Straathof, and G. Zwart, 2017, "Targeted Advertising, Platform Competition, and Privacy," *Journal of Economics and Management Strategy*, 26, 557-570.

[14] Crawford, G., 2000, "The Impact of the 1992 Cable Act on Household Demand and Welfare," *Rand Journal of Economics*, 31(3), pp. 422-450.

[15] Crawford, G., and A. Yurukoglu, 2012, "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, 102(2), pp. 643-85.

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

[17] Eliaz, K., and R. Spiegler, 2016, "Search Design and Broad Matching," *American Economic Review*, 105, pp. 563-586.

[18] Eliaz, K., and R. Spiegler, 2017, "Incentive-Compatible Advertising on Non-Retail Platforms," Working Paper, Tel Aviv University.

[19] Fershtman, D., and A. Pavan, 2017, "Matching Auctions," Working Paper, Northwestern University.

[20] Federal Communications Commission (FCC), 2004, "Report on the Packaging and Sale of Video Programming to the Public." http://www.fcc.gov/mb/csrptpg.html.

[21] 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.

[22] Goettler, R., and R. Shachar, 2001, "Spatial Competition in the Network Television Industry," *RAND Journal of Economics*, 32(4), pp. 624-656.

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

[24] Hatfield, J., and P. Milgrom, 2005, Matching with Contracts, *American Economic Review*, Vol. 95(4), pp. 913-935.

[25] Hart, S., and P. Reny, 2015, "Maximal Revenue with Multiple Goods: Nonmonotonicity and Other Observations," *Theoretical Economics*, 10, pp. 893-922.

[26] Jeon, D., B. Kim, and D. Menicucci, 2016, "Second-degree Price Discrimination by a Two-sided Monopoly Platform," Working Paper, Toulouse School of Economics.

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

[28] Jullien, B., and A. Pavan, 2017, "Information Management and Pricing in Platform Markets," Working paper, Toulouse School of Economics and Northwestern University.

[29] Kang, K., and H. Y. You, 2016, "Lobbyists as Matchmakers in the Market for Access," Working Paper, Carnegie Mellon University and Vanderbilt University.

[30] Loertscher, S., and A. Niedermayer, 2017, "Entry-Deterring Agency," Working Paper, University of Melbourne and Université Paris-Dauphine.

[31] Maskin, E., and J. Riley, 1984, "Monopoly with Incomplete Information," *Rand Journal of Economics*, 15, pp. 171-196.

[32] Mirrokni, V., and H. Nazerzadeh, 2017, "Deals or No Deals: Contract Design for Online Advertising," *World Wide Web Conference (WWW2017).*

[33] Myerson, R., 1981, "Optimal Auction Design," *Mathematics of Operations Research*, 6(1), pp. 58-73.

[34] Mussa, M., and S. Rosen, 1978, "Monopoly and Product Quality," *Journal of Economic Theory,* 18, pp. 301-317.

[35] Office of Fair Trading, 2010, "Online Targeting of Advertising and Prices: A Market Study."

[36] Roth, A., and M. Sotomayor, 1990, Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis, *Econometric Society Monograph Series*, Cambridge University Press, Cambridge.

[37] Tan, G., and J. Zhou, 2017, "Price Competition in Multi-sided Markets," mimeo University of Southern California and National University of Singapore.

[38] Wilson, R., 1993, "Non-Linear Pricing," *Oxford University Press.*