Price Customization and Targeting in Matching Markets*

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

October 2022

Abstract

We introduce a model of (platform-mediated) many-to-many matching in which agents' preferences are both vertically and horizontally differentiated. We first show how the model can be used to derive the profit-maximizing matching plans under customized pricing. We then investigate the implications for targeting and welfare of uniform pricing (be it explicitly mandated or induced by privacy regulation), preventing the platform from conditioning prices on agents' profiles. The model can be applied to study ad exchanges, online retailing, and media markets.

JEL classification: D82

Keywords: Many-to-many matching, asymmetric information, platforms, uniform pricing, second- and third-degree price discrimination, targeting

---

*For useful comments and suggestions, we thank the Editor Gary Biglaiser, anonymous referees, 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, INSPER, FGV-EPGE, and various other workshops where the article was presented. We are grateful to Stephan Lauermann and Markus Reisinger for discussing the article. Gomes gratefully acknowledges financial support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 759733 - PLATFORM). Gomes also acknowledges funding from the ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir programme). Pavan gratefully acknowledges financial support from the NSF (Research Grant SES 1530798) and from Bocconi University where part of this research was conducted.

[†]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.

In advertising exchanges, for example, 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 typically comes from data brokers, who collect and sell personal information (in the form of demographics, geolocation, and browsing history).[4]

In the markets mentioned above, price customization is easy to enforce, as the agents' profiles (i.e., their "horizontal" characteristics that are relevant for price customization) are observable. For instance, in ad exchanges, the advertisers' profile is often revealed by the ads' content, or can be learnt from third parties, whereas in online retailing, information about consumers can often be obtained from data brokers or affiliated websites. In other markets, instead, the agents' profiles 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) where sophisticated pricing strategies are used to 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). See Montes et al. (2018) for a discussion of the value of privacy in online markets and Bounie et al. (2021) for an analysis of how data brokers may optimally partition information for sale to competing firms.

[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

Although having a long history in the policy debate, price customization has attracted renewed attention in the last decade due to the two-sided nature of matching intermediaries and the sheer amount of information now available for pricing.[7] One 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 hamper price customization based on data from third parties.[8,9] Another concern is that consumers often perceive customized prices as being inherently unfair or exploitative.[10] The requirement that prices be uniform may potentially address this issue while still permitting that consumers benefit from a high level of targeting (as their data would determine the matching assignment but not the price).

It is however challenging to assess the impact of customized pricing or, alternatively, of policies constraining it. Part of the difficulty lies in having an analytically amenable model of matching design that is rich enough to accommodate for both *horizontal differentiation* across consumers (capturing disagreements over the most desirable matches) and *vertical differentiation* (i.e., allowing for elastic demands). The main contribution of the present paper is to introduce a tractable model featuring these two dimensions of differentiation. We use the model to show how price customization shapes the matching opportunities offered by a profit-maximizing platform, and to study the impact on targeting and consumer welfare of uniform-price obligations on one of the two sides of the market (whereby payments to the platforms do not depend on the agents' profiles).

Specifically, we capture vertical and horizontal differentiation by letting the agents' types be located on a cylinder, where the height represents the vertical dimension, whereas the radial position determines the horizontal dimension, i.e., the agent's profile, with the latter summarizing a combination of relevant characteristics that, depending on the application, may include demographics, education, zip-code, but also the description of a product or a publisher's website, as in the case

configuration originally selected (see, among others, Crawford (2000), and Crawford and Yurukoglu (2012)). For example, in the US, Direct TV offers various vertically differentiated (i.e., nested) packages (both in English and in Spanish). It then allows viewers to add to these packages various (horizontally differentiated) premium packages, which bundle together channels specialized in movies, sports, news, and games. In addition, viewers can further customize the packages by adding individual sports, news, and movie channels.

[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 data brokers 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]Although certainly relevant, privacy and fairness concerns are out of the scope of this paper.

of ad exchanges - see Figure 1.[11] 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 "ideal match" on the opposite side. Accordingly, each agent's utility for interacting with any other agent from the opposite side decreases with the circular distance between the agent's ideal match (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 used in the empirical literature on media and advertising markets (see, for example, Goettler and Shachar 2001).[12]

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 the matches with the agents from 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 pricing for TV channels (for a discussion, see Crawford and Yurukoglu (2012)). Stringent privacy policies that limit the use of consumers' browsing histories also induce uniform pricing, by limiting online retailers' and market places' ability to condition their offers on the characteristics/profiles of potential buyers.

[FIGURE 1 HERE]

Our first main result shows that, absent any regulation, platforms offer *customized* tariffs on both sides, which discriminate according to the agents' horizontal characteristics (third-degree price discrimination). Crucially, the marginal prices for matches with agents on the opposite side vary both with an agent's own location (his profile) *and* with his partner's location (his partner's profile). As marginal prices are not constant across the volume of the matches, customized tariffs involve profile-specific quantity premia (second-degree price discrimination). In online advertising markets, for instance, this corresponds to advertisers being charged differential marginal prices for access to consumers of a given profile. The complex pricing algorithms used by ad exchanges, combining publisher- and advertiser-specific scores with nonlinear prices, are similar in spirit to the customized tariffs predicted by our model.

[11]The reason for considering a cylinder model instead of a rectangular one (i.e., a Hoteling line augmented by a vertical dimension) is that the former model favors symmetry which in turn simplifies some of the analytical expressions with no major effect on the qualitative results.

[12]In the Appendix, we consider a more general payoff structure whereby match values may also depend on the partners' vertical dimensions, as in the case of firms with different productivities matching with workers with different abilities.

An alternative way of achieving this customization consists in 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 nonlinear cost to the subscriber of customizing the plan by adding extra matches. In the market for cable TV, for instance, 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)).

On technical grounds, the tractability of the model favors a convenient representation of the profit-maximizing tariffs linking location- and volume-specific prices to the various local elasticities of the demands on the two sides of the market. The representation constitutes the analog in a matching market of the familiar Lerner-Wilson formula of optimal non-linear pricing (see, for instance, Wilson (1993)) for standard goods.

Our second set of results provides a characterization of the effects on prices, targeting, and welfare of uniform-price obligations on one of the two sides of the market (be them explicitly mandated or induced by privacy regulation). We consider uniform pricing on a single side, rather than on both sides, for the following reason: When imposed on both sides, uniform tariffs are not sufficiently flexibile to accomodate horizontal preferences while clearing the market across all pairs of locations. Accordingly, uniform pricing on both sides is typically infeasible, in that matching demands on both sides fail to be mutually compatible (see also the discussion in Section 4).

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 with each profile of agents from the opposite side. Relative to the case of customized pricing, this new pricing formula identifies the relevant *aggregate elasticities* in environments where customized (i.e., profile-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 on the side where customization is not feasible, but also how the procurement costs of the matches are affected by the horizontal component of the agents' preferences.

From a 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 quantities, which are typically easier to analyze using standard techniques).

We then use such a characterization to study 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, fails 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 profiles. To give empirical content to this finding, we relate elasticities to match values and type distributions. For an illustration, consider the case of an ad exchange. Under natural conditions on the payoffs of the advertisers and publishers, we show that price customization leads to less targeting than uniform pricing if the distribution of profits per sale of advertisers has thin tails (in the sense of an increasing hazard rate). Accordingly, anonymous pricing for advertising slots (e.g., as a result of regulation banning the use of scores) results in the advertisers being more often matched (relative to laissez-faire) to those publishers whose profile is closer to their ideal audience. That is, uniform pricing leads to more targeting in this case.

We conclude by looking 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 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 surplus. []For instance, in the ad exchange example, advertisers' profits are higher under uniform pricing on the advertisers' side if the distribution of profits per sale of advertisers satisfies a (testable) convexity property.

We believe the model could be used more broadly to study the design of regulatory interventions in markets in which platforms enjoy significant power and price customization is a concern.

Outline of the paper. The rest of the paper is organized as follows. Section 2 presents the model. Section 3 identifies properties of profit-maximizing tariffs and of the induced matching demands, under customized pricing. Section 4 studies the effects of uniform-price obligations. Section 5 discusses the role of various assumptions and the robustness of the analysis to richer specifications. Section 6 briefly reviews the pertinent literature. Section 7 concludes. All proofs are in the Appendix at the end of the document.

2 The cylinder 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 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 $v_k \in V_k \equiv [\underline{v}_k, \overline{v}_k] \subseteq \mathbb{R} \cup \{+\infty\}$ is a shifter that captures heterogeneity in preferences along a *vertical* dimension. It controls for the overall utility the agent derives from interacting with a generic agent from the opposite side, before doing any profiling. The location parameter $x_k \in X_k \equiv [0, 1]$, instead, describes the agent's profile and captures heterogeneity in preferences along a *horizontal* dimension. The latter parameter captures personal traits (such as demographics, gender, education, residence, income, etc) that jointly determine one's relative preferences over any two agents from the opposite side, as well as one's attractiveness in the eyes of agents from the opposite side. Figure 1 depicts the above structure. 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 height of the cylinder measuring the vertical type and the position on the circle measuring the agent's profile (i.e., the horizontal type).

Agents derive higher utility from being matched to agents who are "closer" to them. Their utility also increases, over all profiles, with their vertical type. We assume the utility that an agent from side k with type $\theta_k = (v_k, x_k)$ derives from being matched to each 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' profiles. 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.

Each agent's type $\theta_k = (v_k, x_k)$ is an independent draw from the absolutely continuous distribution function F_k with support Θ_k. The total payoff that type $\theta_k = (v_k, x_k)$ obtains from being matched, at a price p, to a measurable set of types $\mathbf{s}_k \subseteq \Theta_l$ from side $l \neq k$ is given by

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

Accordingly, matches are non-rival, in that agents always benefit from having "access" to more agents from the other side of the market.[13] The payoff that the same agent obtains outside of the platform is equal to zero.[14]

We assume that the vertical dimensions v_k are the agents' private information whereas the profiles, x_k, are publicly observable.[15] We also assume that the agents' profiles are uniformly distributed over the circle and that the vertical types are drawn independently from the profiles according to

[13] The utilities u_k and u_l should be interpreted as ex-ante expected payoffs. Ex-post, the agents may learn that the match is unattractive (to one or both agents) and refrain from interacting.

[14] 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 agents' types, as opposed to their identities, reflects the property that, under both the welfare- and the profit-maximizing tariffs, each agent from each side k 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. In a previous version, we considered a more general setting where such assumptions are relaxed. We opted here for the representation in (1) because it permits us to simplify the exposition and favors sharper conclusions. See also Section 5 for further discussion.

[15] See also Akbarpour et al. (2021) for an alternative model in which certain characteristics of the agents (referred to as "labels" in that paper) are publicly observable whereas others such as the agents' social welfare weights and their willingness to pay for quality are the agents' private information. In that paper, though, the designer assigns physical goods to the agents, whereas in the present paper the designer matches the agents with other agents who are also privately informed.

an absolutely-continuous distribution F_k^v with density f_k^v strictly positive over V_k and hazard rate $\lambda_k^v \equiv f_k^v/[1-F_k^v]$.[16] In Section 5, we discuss how the results extend to settings in which the vertical and horizontal types are correlated (with the marginal distribution of the horizontal types possibly different from the uniform one), as well as to settings with richer payoff specifications that allow the utility that each type θ_k derives from being matched to each type θ_l to depend also on type θ_l's vertical dimension v_l. Finally, we discuss how the results extend to certain settings in which the agents' profiles are private information. The primary reason for abstracting from these enrichments in the baseline model is that this permits us to simplify a lot the exposition.

The next two examples illustrate the type of markets the analysis can be applied to.

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

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

where v_a is the advertiser's profit per sale and where the strictly decreasing function $\phi : [0, \frac{1}{2}] \to [0, 1]$ 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 profile, x_a. By contrast, publishers can be viewed (to a first approximation) as indifferent with respect to the kind of advertisement displayed at their websites. The matching (dis)utility of a publisher reflects the opportunity cost of not using the advertisement space to sell its own products, or from not selling the ad slot outside of the platform. Accordingly, the profit that a publisher of type $\theta_b = (v_b, x_b)$ derives from displaying the ad of an advertiser of type $\theta_a = (v_a, x_a)$ is given by $u_b(v_b, |x_a - x_b|) = v_b \leq 0$, all $x_a, x_b \in [0, 1]$. $\diamondsuit$

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

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

where $\alpha \in [0, 1]$ captures the relative importance of the viewer's vertical preferences v_a, describing her willingness to consume media content, and her horizontal preferences x_a, describing her ideal type of content (e.g., sports, news, movies, etc). In turn, the strictly decreasing function $\phi : [0, \frac{1}{2}] \to \mathbb{R}_+$ describes how the viewer's utility declines with the distance between the viewer's profile, x_a, and the provider's profile, x_b. Finally, $\delta \in \mathbb{R}/\{0\}$ measures the elasticity of substitution between profile compatibility and willingness to consume. The matching (dis)utility $u_b(v_b, |x_b - x_a|))$ of the content

[16] Hence, for any $\theta_k = (v_k, x_k)$, $F_k(\theta_k) = F_k^v(v_k)x_k$.

provider may reflect the extra revenue from advertisers (which may depend on the profile of the 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).[17] $\diamondsuit$

Another example that shares the preference structure of Example 1 is that of business-to-business platforms matching firms engaging in procurement (from side a) with firms supplying services (from side b). Each supplier's profile, x_b, identifies the type of service offered, whereas each procurer's profile, x_a, identifies the type of service demanded. The function $\phi(|x_a - x_b|)$ describes the probability that the match results in trade, which obviously decreases with the discrepancy between x_a and x_b. The vertical parameters $v_a > 0$ and $v_b < 0$ capture the willingness to pay and marginal costs for the procuring and supplying firms, respectively.

For future reference, we 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"

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

satisfy the same monotonicity properties of the true values. Namely, the functions $\varphi_k\left(\theta_k, \theta_l\right)$ are strictly increasing in v_k and weakly decreasing in $|x_k - x_l|$, which is the analog of Myerson regularity condition (e.g., Myerson (1981)) in a matching environment.

Tariffs and matching demands

The platform offers matching tariffs on each side of the market. A *matching tariff* T_k specifies the (possibly negative) total payment $T_k\left(\mathbf{s}_k | x_k\right)$ that each agent with profile $x_k \in X_k$ is asked to pay to be matched to each set of types $\mathbf{s}_k$ from the opposite side of the market.[18] To guarantee participation by all agents, we require that, for all x_k, $T_k\left(\mathbf{s}_k | x_k\right) = 0$ if $\mathbf{s}_k = \emptyset$.

Given the tariff T_k, we say that the function $\mathbf{s}_k : \Theta_k \to \Sigma(\Theta_l)$ is a *matching demand* consistent with the tariff T_k if, for any $\theta_k = (v_k, x_k) \in \Theta_k$,

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

Definition 1. The tariff profile $(T_k)_{k=a,b}$ is *feasible* if there exists a pair of matching demands $(\mathbf{s}_k)_{k=a,b}$ consistent with $(T_k)_{k=a,b}$ satisfying the following *reciprocity condition*, for all $(\theta_k, \theta_l) \in \Theta_k \times \Theta_l$, $k, l \in \{a, b\}$, $l \neq k$:

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

[17] 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).

[18] In their most general form, matching tariffs might condition such payments on the agents' own profiles (which are observable by the platform), but not on the agents' vertical dimensions, which are the agents' private information.

That is, if an agent from side k with type θ_k demands 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 demand to be matched to all agents from side k with type θ_k.

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

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

A pair of tariffs $(T_k^*)_{k=a,b}$ is *profit-maximizing* if there exist matching demands $(\mathbf{s}_k^*)_{k=a,b}$ consistent with $(T_k^*)_{k=a,b}$ such that the platform's profits under $(T_k^*, \mathbf{s}_k^*)_{k=a,b}$ are as high as under any other quadruple $(T_k, \mathbf{s}_k)_{k=a,b}$, where $(T_k)_{k=a,b}$ is a pair of feasible tariffs and $(\mathbf{s}_k)_{k=a,b}$ are demands consistent with $(T_k)_{k=a,b}$. Hereafter we denote by $(T_k^*)_{k=a,b}$ a pair of profit-maximizing tariffs, and by $(\mathbf{s}_k^*)_{k=a,b}$ the matching demands that, together with $(T_k^*)_{k=a,b}$, maximize the platform's profits.

3 Customized pricing

We now introduce a class of tariffs that plays an important role in the analysis below. Under such tariffs, which we call *customized*, the platform offers to each side-k agent a baseline matching set at a baseline price, along with a collection of personalized prices that the agent can use to customize his matching set. The total price of the customization is *separable* across agents' profiles, but may vary non-linearly with the amount of agents from the opposite side included in the matching set (a form of second-degree price discrimination). Importantly, the personalized prices the agents pay for the customizations depend on the agents' *own profiles* (a form of third degree price discrimination). Customized tariffs capture important features of the matching plans offered by platforms such as cable TV providers, ad exchanges, and online retailers. Before proceeding to the definition, we need to introduce the following piece of notation: Given any matching set $\mathbf{s}_k$, and any profile x_l, we let $q_{x_l}(\mathbf{s}_k)$ denote the "mass" of side-l agents with profile x_l included in the matching set $\mathbf{s}_k$.[19]

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

$$\left\{(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k)) : x_k \in X_k\right\},$$

one for each profile $x_k \in X_k$, such that each side-k agent with profile x_k choosing the matching set $\mathbf{s}_k \in \Sigma(\Theta_l)$ is asked to make a total payment equal to

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

with $\rho_k\left(q_{x_l}\left(\underline{\mathbf{s}}_k(x_k)\right)|x_l; x_k\right) = 0$ for all $x_l \in X_l$.

[19]We abuse terminology by referring to the density of agents of a certain type as the "mass" of agents of that type.

A customized tariff can thus be thought of as a collection of *matching plans*, one for each profile x_k. Each plan comes with a baseline configuration, given by the default set of types $\underline{\mathbf{s}}_k(x_k)$ from side $l \neq k$ included in the package, and a baseline price $\underline{T}_k(x_k)$. Each agent selecting the plan $(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k))$ can then customize his matching set by adding extra matches. The cost of the customization is separable in the type of matches added to the baseline configuration, with each schedule $\rho_k(q|x_l; x_k)$ describing the non-linear fee for adjusting the amount of x_l-agents from the default level $q_{x_l}(\underline{\mathbf{s}}_k(x_k))$ to q.[20] Importantly, for each quantity q of x_l-agents, the price $\rho_k(q|x_l; x_k)$ depends on the baseline plan, which is conveniently indexed by the profile x_k targeted by the plan. 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*. In particular, note that a customized tariff combines elements of *second-degree* price discrimination (each price function $\rho_k(q|x_l; x_k)$ is possibly non-linear in q) with elements of *third-degree* price discrimination (each non-linear price function $\rho_k(q|x_l; x_k)$ depends on the plan, and hence on the agent's own profile). The baseline configurations are designed for those agents with the lowest vertical type, $\underline{v}_k$, whereas the customizations are designed for those agents with higher vertical types.

Clearly, not all tariffs are customized, in the sense of Definition 2. For instance, tariffs that condition the price for the x_l-matches on the demand for the x_l'-matches, with $x_l' \neq x_l$, are not customized. The following result shows that this extra degree of freedom is inconsequential for profits, in that the platform's optimum is achieved by offering a pair of customized tariffs.

Lemma 1. ***(properties of the optimum)*** *The following are true:*

1. *there exists a pair of customized tariffs $(T_k^*)_{k=a,b}$ that are profit-maximizing;*

2. *the matching demands $(\mathbf{s}_k^*)_{k=a,b}$ consistent with the profit-maximizing customized tariffs $(T_k^*)_{k=a,b}$ are described by threshold functions $t_k^* : \Theta_k \times X_l \to V_l$ such that*

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

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

Under the profit-maximizing tariffs, for any given profile x_k, the matching sets demanded by those agents with higher vertical types are supersets of those demanded by agents with lower vertical types. Moreover, side-l agents of profile 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 agents' vertical types v_k are large enough. Finally, the range of vertical types $[t_k(\theta_k, x_l), \bar{v}_l]$ of side-l agents with profile x_l that each side-k agent

[20]The schedules $\rho_k(q_{x_l}|x_l; x_k)$ may also specify the price for reducing the amount of x_l-agents below the default level. However, as we show in the Appendix, in equilibrium, the induced demands are such that $\mathbf{s}_k(\theta_k) \supset \underline{\mathbf{s}}_k(x_k)$ for all $\theta_k = (v_k, x_k)$, $k = a, b$, meaning that no agent asks to reduce the number of matches below the level specified in the baseline configuration. Accordingly, $\underline{\mathbf{s}}_k(x_k)$ and $\underline{T}_k(x_k)$ correspond to the matching sets and payments designed for the x_k-agents with the lowest vertical type, $\underline{v}_k$.

of type $\theta_k = (v_k, x_k)$ is matched with is smaller the larger is the discrepancy $|x_k - x_l|$ between the profiles. Figure 2 illustrates these properties.

[FIGURE 2 HERE]

To gain intuition, note that the marginal profits the platform obtains by matching type $\theta_l = (v_l, x_l)$ from side l with 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 strictly increase with the agents' vertical types implies existence of a threshold $t_k^*(\theta_k, x_l)$ such that Condition (6) is satisfied if, and only if, fixing (θ_k, x_l), $v_l \geq t_k^*(\theta_k, x_l)$, 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 side-l agents (of all profiles) with lower vertical types. These properties are the analogues of those in Gomes and Pavan (2016) in a setting with horizontally-differentiated preferences.[21]

Lerner-Wilson formula for matching schedules

We now derive further properties of the customized tariffs that maximize the platform's profits. Consider first the problem of a side-k agent of type $\theta_k = (v_k, x_k)$ under the plan x_k (recall that this is the plan designed for all side-k agents with profile x_k). The mass of x_l-agents demanded by type θ_k is given by

$$\hat{q}_{x_l}(\theta_k) \in \arg\max_{q \in [0,1]} \left\{ u_k(v_k, |x_k - x_l|) \cdot q - \rho_k(q|x_l; x_k) \right\}.$$

Assuming the price schedule $\rho_k(\cdot|x_l; x_k)$ is convex and differentiable in q (to be confirmed below), with derivative $\rho_k'(\cdot|x_l; x_k)$, we have that $\hat{q}_{x_l}(\theta_k)$ is a solution to the following first-order condition

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

whenever $\hat{q}_{x_l}(\theta_k)$ is interior, i.e., whenever $\hat{q}_{x_l}(\theta_k) \in (0, 1)$.

Next, for any pair of profiles $x_k, x_l \in [0, 1]$, and any "interior" marginal price,[22] let $\hat{v}_{x_l}(\rho_k'|x_k)$ denote the value of v_k that makes each x_k-agent indifferent between adding an extra unit of x_l-matches or not, given the marginal price ρ_k'. Note that $\hat{v}_{x_l}(\rho_k'|x_k)$ is implicitly defined by:[23]

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

[21] In a setting with purely vertically-differentiated preferences, Gomes and Pavan (2016) identify conditions under which the profit-maximizing tariffs induce demands with a threshold structure.

[22] A marginal price is interior if $\rho_k' \in [u_k(\underline{v}_k, |x_k - x_l|), u_k(\bar{v}_k, |x_k - x_l|)]$.

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

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 demanded quantity q. Furthermore, because agents with higher vertical types purchase larger matching sets, the *demand* for the q-th unit of x_l-agents by x_k-agents, at the *marginal price* ρ_k', is given by:[24]

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

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 x_k-agents 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 x_k-agents for the q-th unit of x_l-agents with respect to its marginal price ρ_k' by

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

where, once again, the arguments of the marginal price ρ_k' are dropped to lighten the notation. The next proposition characterizes the price schedules associated with the profit-maximizing customized tariffs of Lemma 1 in terms of the profile-specific elasticities of the demands on both sides of the market. To ease the exposition, the dependence of the marginal prices, $\rho_k^{*\prime}$, of the demands D_k, and of the elasticities, ε_k, on the profiles (x_a, x_b) 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.*[25] *Moreover, for all pair of profiles (x_a, x_b), and all interior pairs of 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:*

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

The Lerner-Wilson formulas (11) jointly determine the price schedules on both sides of the market.[26] Intuitively, these formulas require that the marginal revenue from adding to the matching sets of x_k-agents the q_k-th unit of x_l-agents compensate for the marginal cost from adding to the matching sets of x_l-agents the q_l-th unit of x_k-agents, where q_k and q_l are jointly related by the reciprocity condition in the proposition.

[24] 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.

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

[26] The reason why we refer to the formulas in the proposition as Lerner-Wilson is that, in his book "Nonlinear Pricing," Robert Wilson was among the first to illustrate how Lerner's elasticity formulas for textbook monopolistic pricing problems extend to a wide range of non-linear pricing with rich quantity discounts.

As for the standard Lerner-Wilson formula for monopoly/monopsony pricing, on each side, the marginal revenue/cost has two terms: the term $\rho_k^{*\prime}$ captures the impact on the platform's profits of the matches involving the marginal agents, whereas the semi-inverse-elasticity term $\rho_k^{*\prime}\left[\varepsilon_k\left(\rho_k^{*\prime}\right)\right]^{-1}$ captures the impact of the adjustment in the price for the infra-marginal matches.

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), which is one of the key features or matching. This is manifested in the fact 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. This property is not specific to our matching model. Quantity premia are a common feature of all screening models with multiplicative payoffs.

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. This is illustrated in the next example.

Example 3. (***ad exchange* - continued**) Consider the ad exchange market of Example 1, and assume that the cdf's F_a^v and F_b^v are uniform with support on $[0,1]$ and $[-1,0]$, respectively. Proposition 1 then implies that the marginal price schedules $\rho_a^{*\prime}$ and $\rho_b^{*\prime}$ are such that

$$\rho_a^{*\prime}(q) = q + \frac{\phi\left(|x_a - x_b|\right)}{2} \qquad \text{and} \qquad \rho_b^{*\prime}(q) = \left(q - \frac{1}{2}\right)\phi\left(|x_a - x_b|\right),$$

for any $q \in [0,1]$. Accordingly, for any q, the marginal price that the advertisers pay for the q-th impression on each publisher's website decreases with the discrepancy between the publishers' and the advertisers' profiles. Likewise, the marginal subsidy to the publishers decreases with the discrepancy between the publishers' and the advertisers' profiles. $\diamondsuit$

In Example 3, the marginal schedule $\rho_b^{*\prime}$ faced by the side-b publishers depends on the conversion probabilities ϕ, despite the fact that the side-b publishers do not care about such conversions per se (or about the profiles of the advertisers whose ads they display). Such dependence, however, is a natural consequence of the need for the platform to clear the market respecting the reciprocity conditions required by matching. Note that the use of scoring auctions by ad exchanges, which make the procurement prices paid to the publishers explicitly depend on the discrepancy in profiles is broadly consistent with the results in Proposition 1.

More generally, the price schedules offered to any two profiles x_k and x_l are a function of the *profile-specific demand elasticities* $\varepsilon_k\left(\cdot|x_l;x_k\right)$. 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 both sides. As we show below, the ability to tailor price schedules to profiles (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 relate the cut-off types on both sides of the market.[27] The spirit of these formulas is the same as in the reduced-form approach pioneered by Saez (2001) in the context of optimal taxation. Under the *assumption* that a platform prices matches optimally, Proposition 1 can be used by the econometrician to estimate the distribution of the agents' preferences from data on price schedules and match volumes.

Alternatively, Proposition 1 can be employed to assess the optimality of mechanisms *currently* used. In online advertising markets, for instance, a complex system of prices is employed by ad exchanges to match publishers with advertisers. These prices employ user- and advertiser-specific scores, and are non-linear in the number of impressions, which is consistent with what predicted by Proposition 1. The Lerner-Wilson formulas (11) can serve as a test for the optimality of these "indirect" mechanisms.

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.[28] In this section, we study platforms' behavior when subject to uniform-price obligations on one of the two sides of the market.

Uniform pricing and aggregate demand elasticities

Definition 3. The side-k tariff T_k is consistent with *uniform pricing* if there exists a collection of (possibly non-linear) *price schedules* $p_k(q|x_l)$, one for each side-l profile $x_l \in X_l$, such that the total payment asked by the platform to the side-k agents for each matching set $\mathbf{s}_k \in \Sigma(\Theta_l)$ is given by[29]

$$T_k(\mathbf{s}_k) = \int_0^1 p_k(q_{x_l}(\mathbf{s}_k)|x_l)dx_l. \tag{12}$$

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 profile of the side-l agents.

[27] To see this, fix (x_a, x_b) and drop it to ease the notation. For any q_a, the result in Lemma 1 implies that the most economical way of providing the x_a-agents with access to q_a agents with profile x_b is to match the former agents to all x_b-agents whose vertical type is above $\tilde{v}_b$, with $\tilde{v}_b$ defined by $1 - F_b^v(\tilde{v}_b) = q_a$. For any q_b, 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 $u_a(v_a, |x_a - x_b|) = \rho_a^{*\prime}(q_a)$.

[28] As indicated above, customized pricing is a special form of bundling, whereby the marginal prices $\rho_k(q|x_l; x_k)$ for the q-th matches with the x_l-agents depends on the x_k package and hence, indirectly, on the quantity of x_l'-agents included in the baseline configuration of the x_k-plan.

[29] Recall that $q_{x_l}(\mathbf{s}_k)$ is the quantity of x_l-agents included in the set $\mathbf{s}_k$.

Namely, each schedule $p_k(q|x_l)$ specifies the total price each side-k agent has to pay to be matched to q x_l-agents. Importantly, contrary to the case of price customization, the price $p_k(q|x_l)$ is independent of the side-k agent's own profile, x_k.

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 profile $x_b \in [0,1]$, and each quantity q, this schedule specifies the price that side-a agents have to pay to be matched to q x_b-agents. The side-a aggregate demand (over x_a) for the q-th unit of x_b-agents at the marginal price $p'_a(q|x_b)$ is then equal to

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

where $D_a\left(p'_a|x_b;x_a\right)$ is the mass of agents located at x_a that demand q units or more of x_b-agents, and $\hat{v}_{x_b}\left(p'_a|x_a\right)$ is the value of v_a for which $u_a(v_a, |x_b - x_a|) = p'_a$. As in the previous section, the arguments $(q|x_b)$ of the marginal prices $p'_a(q|x_b)$ have been dropped to ease the exposition.

The elasticity of the side-a *aggregate demand* for the q-th unit of 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 local elasticity 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

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

The elasticity $\bar{\varepsilon}_a(p'_a|x_b)$ measures the percentage variation in the mass of agents from side a that demand at least q matches with the x_b-agents in response to a percentage change in the marginal price for the q-th unit of such agents. It is also equal to the average elasticity (over side-a profiles) of the x_a-demands for the q-th unit of the x_b-agents with respect to the marginal price p'_a. This average is taken under a distribution that assigns to each profile x_a a weight proportional to the mass of agents $D_a\left(p'_a|x_b,x_a\right)$ with profile x_a demanding q units, or more, of x_b-agents.

The next proposition derives properties of the profit-maximizing tariffs T_a^u and T_b^u offered by a platform that is constrained to price uniformly on side a. To ease the exposition, the dependence of the marginal price $p_a^{u\prime}$ and of the aggregate elasticity $\bar{\varepsilon}_a$ on x_b, as well as the dependence of the marginal price $\hat{\rho}_b^{u\prime}$ and of the local elasticity $\hat{\varepsilon}_b$ on (x_a, x_b), are dropped from all the formulas in the proposition.

Proposition 2. ***(uniform pricing)*** *Suppose the platform is constrained to price uniformly on side a, but is free to offer any tariff on side b. The profit-maximizing tariffs $(T_k^u)_{k=a,b}$ are such that T_b^u is customized. The price schedules p_a^u and ρ_b^u associated with the profit-maximizing tariffs $(T_k^u)_{k=a,b}$ are differentiable and convex over the equilibrium ranges.*[30] *Moreover, for all profiles $x_b \in X_b$, and*

[30] 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 $q_b \in [q_{x_a}(\mathbf{s}_b(\underline{v}_b, x_b)), q_{x_a}(\mathbf{s}_a(\overline{v}_b, x_b))]$.

all interior quantity pairs $(q_a, q_b(\tilde{x}_a))$, $\tilde{x}_a \in X_a$, *such that*

$$q_a = D_b\left(\rho_b^{u\prime}\left(q_b|x_b;x_a\right)|x_b;x_a\right) \quad \textit{and} \quad q_b(x_a) \equiv D_a\left(p_a^{u\prime}(q_a)|x_b;x_a\right),$$

the marginal prices schedules $p_a^{u\prime}$ *and* $\rho_b^{u\prime}$ *jointly satisfy the following optimality condition:*

$$\underbrace{p_a^{u\prime}(q_a)\left(1-\frac{1}{\bar{\varepsilon}_a\left(p_a^{u\prime}(q_a)\right)}\right)}_{\textit{net effect on side-a profits}}+\mathbb{E}_{H(\tilde{x}_a|x_b,p_a^{u\prime}(q_a))}\left[\underbrace{\rho_b^{u\prime}(q_b(\tilde{x}_a))\left(1-\frac{1}{\varepsilon_b\left(\rho_b^{u\prime}(q_b(\tilde{x}_a))\right)}\right)}_{\textit{net effect on side-b profits}}\right]=0, \tag{13}$$

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

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

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 x_b-agents cannot condition on the profile of side-a agents. As a result, the markup on the q-th unit of x_b-matches is constant across all side-a profiles x_a. The relevant elasticity for determining this markup is then the aggregate elasticity $\bar{\varepsilon}_a(\cdot|x_b)$, rather than the profile-specific elasticities $\varepsilon_a\left(\cdot|x_b;x_a\right)$ in the Lerner-Wilson formula (11) of Proposition 1. Interestingly, the cost of procuring x_b-agents is also an average; namely, that of the (mark-up augmented) prices

$$\mathbb{E}_{H(\tilde{x}_a|x_b,p_a^{u\prime}(q_a))}\left[\rho_b^{u\prime}(q_b(\tilde{x}_a))\left(1-\frac{1}{\varepsilon_b\left(\rho_b^{u\prime}(q_b(\tilde{x}_a))\right)}\right)\right]$$

charged to the x_b-agents for their interactions with the side-a agents demanding q_a or more x_b-matches. The next example illustrates how the uniformity requirement affects the optimal marginal price schedules in the advertising market.

Example 4. (***ad exchange* - continued**) Consider the ad exchange market under the assumptions of Example 3, and let $\phi(|x_a - x_b|) = (1+\mu|x_a - x_b|)^{-1}$, where $\mu > 0$. Proposition 2 then implies that the marginal price schedules are given by

$$p_a^{u\prime}(q) = q + \frac{1}{2+\frac{\mu}{2}} \qquad \text{and} \qquad \rho_b^{u\prime}(q) = \frac{q-1}{1+\mu|x_a-x_b|} + \frac{1}{2+\frac{\mu}{2}},$$

for any $q \in [0,1]$. As a consequence of uniformity, marginal prices on side a are invariant to the discrepancy between the publishers and the advertisers' profiles. $\diamondsuit$

The analysis above investigates the platform's optimum when uniform pricing is imposed in *only one* side of the market (side a). In this case, the ability to price discriminate on side b guarantees that the platform has enough flexibility to procure the side-b matches demanded by the side-a agents, while respecting reciprocity. This property is not guaranteed to hold when uniform pricing is imposed on both sides. Intuitively, in this case, the price schedule that clears the market for the matches between the x_a-agents and the x_b-agents need not clear the market for the matches between the x_a-agents and the x'_b-agents, when $x'_b \neq x_b$. This is formally discussed in the next remark.

Remark 1. ***(uniform pricing on both sides)*** Consider matching demands $(\mathbf{s}_k(\theta_k))_{k=a,b}$ described by threshold functions $t_k : \Theta_k \times X_l \to V_l$ for $k, l = a, b$, $l \neq k$ (in the sense of Lemma 1). Such matching demands can be implemented by a pair of tariffs T_k and T_l, each consistent with uniform pricing, if and only if, for each $\theta_k \in \Theta_k$,

$$u_l(t_k(\theta_k, 0), |x_k|) = u_l(t_k(\theta_k, x_l), |x_k - x_l|) \tag{14}$$

for all $x_l \in X_l$. The set of threshold rules $(t_k)_{k=a,b}$ satisfying (14) for each $\theta_k \in \Theta_k$ and $x_l \in X_l$ is non-empty only in knife-edge cases. For instance, if u_l is invariant in $|x_k - x_l|$, as in the ad exchange application of Example 1, then any threshold function $t_k(\theta_k, x_l)$ satisfying (14) must be invariant in x_l. In this case, its associated matching demands can be induced by tariffs satisfying uniform pricing on both sides only if u_k is also invariant in $|x_k - x_l|$.[31,32]

Relatedly, requiring that the platform prices *linearly* on both sides of the market (that is, imposing that the marginal prices $\rho'_k(q|x_l; x_k)$ be invariant in the volume q of the matches demanded) may preclude the possibility of clearing the market when some matches are excluded. This is so *even* if such prices are allowed to depend on the profiles of the involved agents, on both sides of the market.[33] Accordingly, banning second-degree price discrimination is typically feasible only if done on one side only.

Targeting

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 matches enjoyed by the agents obviously depend on the pricing practices followed by the platform. Price customization

[31] Remark 1 focuses on demands with a threshold structure for simplicity. The difficulty to satisfy reciprocity under uniform pricing on both sides of the market extends to more general demands.

[32] Tariffs that require each participating agent to pay a fixed access fee for interacting with all participating agents from the opposite side of the market (and offer no other matching alternatives) are not uniform in the sense of Definition 3, for they are not separable. Importantly, such tariffs do not allow for any form of discrimination within sides, and abstract from matching design and targeting. These are the tariffs considered in the two-sided market literature - see, for instance, Rochet and Tirole (2006) and Jullien et al (2021), for a recent overview.

[33] To see this, note that, given the marginal price ρ'_k charged to x_k-agents for interacting with x_l-agents, each x_k-agent whose vertical type exceeds $\hat{v}_{x_l}(\rho'_k|x_k)$, where $\hat{v}_{x_l}(\rho'_k|x_k)$ is implicitly defined by $u_k(v_k, |x_k - x_l|) = \rho'_k$, demands to interact with *all* x_l-agents. To procure such x_l-agents, the price ρ'_l the platform must charge x_l-agents equals $\rho'_l = u_l(\underline{v}_l, |x_k - x_l|)$. Given such a price, however, each x_l-agent demands to interact with *all* x_k-agents, implying that the only linear prices that clear the (x_k, x_l)-market must induce full participation on either side.

enables a platform to charge agents prices that depend on their horizontal characteristics (either directly, when the latter are observable, as assumed here, or indirectly, through bundling, as discussed in the Appendix). 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-price obligations. 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 broadly consistent with this line of reasoning. This intuition, however, is incomplete, as it ignores the (endogenous) changes in prices that platforms undertake in response to uniform-price obligations. The analysis below provides some guidelines as to the effects of uniform-price obligations on one of the two sides of the market on targeting.

Broadly, targeting captures the idea that agents are induced to demand relatively more of those matches close to their ideal profile. That is, the composition of the matching sets is geared towards the most-preferred matches, even if the total amount of matches may differ across alternative tariffs.

Definition 4. (targeting) Consider tariffs $(T_k)_{k=a,b}$ and $(T'_k)_{k=a,b}$ inducing the matching demands $(\mathbf{s}_k)_{k=a,b}$ and $(\mathbf{s}'_k)_{k=a,b}$, respectively. The tariffs $(T_k)_{k=a,b}$ lead to more targeting than the tariffs $(T'_k)_{k=a,b}$ if, for each $\theta_k = (v_k, x_k)$, there exists a critical distance $\delta(\theta_k) \in (0, \frac{1}{2})$ such that

$$q_{x_l}(\mathbf{s}_k(\theta_k)) - q_{x_l}(\mathbf{s}'_k(\theta_k)) \begin{cases} > 0 & \text{if} \quad |x_k - x_l| < \delta(\theta_k) \\ < 0 & \text{if} \quad |x_k - x_l| > \delta(\theta_k), \end{cases}$$

where, recall, $q_{x_l}(\mathbf{s}_k)$ is the "mass" of side-l agents with profile x_l included in the matching set $\mathbf{s}_k$.

If matching demands exhibit a threshold structure, as it occurs at the optimum, Definition 4 boils down to

$$t_k(\theta_k, x_l) - t'_k(\theta_k, x_l) \begin{cases} < 0 & \text{if} \quad |x_k - x_l| < \delta(\theta_k) \\ > 0 & \text{if} \quad |x_k - x_l| > \delta(\theta_k), \end{cases}$$

where $(t_k)_{k=a,b}$ and $(t'_k)_{k=a,b}$ are the thresholds describing the demands induced by $(T_k)_{k=a,b}$ and $(T'_k)_{k=a,b}$, respectively.

Clearly, Definition 4 describes a partial order, in that one can construct tariffs that are not comparable in terms of targeting (for instance, those whose matching demands are described by threshold functions with multiple crossings). Yet, it is always possible to rank the optimal tariffs under, respectively, customized pricing on both sides, and uniform pricing on side a but customized on side b. In fact, the results in Propositions 1 and 2 imply that the threshold functions $t^*_a(\theta_a, x_b)$ and $t^u_a(\theta_a, x_b)$ cross once and only once. This follows from the dependence of the thresholds on local and average elasticities on the two sides of the market (more on this below). Figure 3 provides an illustration.

[FIGURE 3 HERE]

Note that, because matching is reciprocal, we only need to compare thresholds on one of the two sides. 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 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 example compares targeting under price customization on both sides to targeting under uniform pricing on side a and customized pricing on side b in the context of the ad exchange application:

Example 5. (***ad exchange* - continued**) Consider the ad exchange market under the assumptions of Example 4. The side-b marginal price schedules under customized pricing (on both sides) and under uniform pricing (on side a and customized pricing on side b) are such that, for any pair of profiles (x_a, x_b),

$$\rho_b^{*\prime}(q|x_a; x_b) > \rho_b^{u\prime}(q|x_a; x_b) \ \ \forall q \in [0,1] \qquad \Longleftrightarrow \qquad |x_a - x_b| < \frac{1}{4}.$$

Therefore, $t_b^*(\theta_b, x_a) > t_b^u(\theta_b, x_a)$ if and only if $|x_a - x_b| < \frac{1}{4}$. Accordingly, price customization leads to less targeting than uniform pricing. $\diamondsuit$

More generally, as revealed by the pricing formulas (11) and (13), the effects on targeting of uniform pricing on side a hinge on the comparison between (i) the local elasticities on the two sides of the market and (ii) the aggregate inverse elasticity on side a and the average inverse elasticity on side b.

The next proposition leverages on this observation to identify conditions under which uniform pricing on side a (for short, uniform pricing) leads to more (alternatively, less) targeting than customized pricing on both sides (for short, customized pricing). For simplicity, the result in Proposition 3 is for a market in which preferences on side b are profile-invariant, as in the ad-exchange application of Example 1 (we discuss more general conditions in the proof of Proposition 3 in the Appendix).

Proposition 3. ***(comparison: targeting)*** *Suppose preferences on side b are profile-invariant. Uniform pricing on side a leads to more (alternatively, less) targeting than customized pricing on both sides when the side-a semi-elasticities are increasing (alternatively, decreasing) in both distance and price.*

Fix the side-b profile 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-matches is invariant to the distance $|x_b - x_a|$, as the marginal price for the q-th unit of the x_b-matches is the same for all x_a-profiles. As a consequence, when the semi-elasticities of the side-a demands increase (alternatively, decrease) with distance, the marginal price for the q-th unit of the x_b-matches charged to the x_a-agents under customized pricing is lower than the corresponding price under uniform pricing when profiles are far apart, whereas the opposite is true for nearby profiles. Accordingly, there is more targeting under uniform pricing than under customized pricing.

The reader may find it surprising that targeting can be lower under customized pricing (on both sides) than under customized pricing on side b but uniform pricing on side a. However, it is worth recalling that the platform has two margins to raise its revenues, price and quantity. Fixing the profile x_b of the side-b agents, we have that, when the elasticity of the demands by the side-a agents for the x_b-agents increases with distance (meaning that those side-a agents who like the x_b-agents the most are also the least price-elastic ones), then, under customized pricing, the platform maximizes its revenues by asking a high price to those side-a agents who value the x_b-agents the most, optimally sacrificing high volumes of trade for larger markups.

That targeting is higher under one pair of tariffs than another, however, does not mean that the agents are better off under the tariffs inducing more targeting. This is because the volume of the matches may be lower under the tariffs inducing more targeting. Furthermore, the total payments need not be the same across the two regimes (see the discussion surrounding Proposition 4 below).

One can also use the characterization of the matching demands in the previous section to translate the result in Proposition 3 in terms of conditions on match values and type distributions. For example, one can show that the side-a semi-elasticities are increasing in both distance and price when the hazard rate for F_a^v is increasing in v_a, and u_a is submodular and concave in v_a. Alternatively, they are decreasing in both distance and price when the hazard rate for F_a^v is decreasing in v_a, and u_a is supermodular and convex in v_a.

Accordingly, we are able to generalize the conclusion of Example 5 to a broad class of conversion probability functions $\phi(\cdot)$ and distributions of profits per sale.

Example 6. (***ad exchange* - continued**) Consider the ad exchange market of Example 1. Because the match function u_a is submodular and linear in v_a, price customization leads to less targeting than uniform pricing when F_a^v has an increasing hazard rate (e.g., a uniform or exponential cdf). $\diamondsuit$

We next apply Proposition 3 to assess the impact of price customization (or lack thereof) in the media market application of Example 2.

Example 7. (***media platform* - continued**) Consider the media market of Example 2. Price customization leads to more targeting than uniform pricing when the viewers' elasticity of substitution is high (namely, when $\delta > 1$) and F_a^v has a decreasing hazard rate (e.g., a Pareto cdf). By contrast, price customization leads to less targeting than uniform pricing when the viewers' elasticity of substitution is low (namely, when $\delta < 1$) and F_a^v has an increasing hazard rate (e.g., a uniform cdf).[34] The effects on targeting of regulation requiring stand-alone pricing for media content therefore *jointly* depend on the viewers' elasticity of substitution and the distribution of their vertical preferences. $\diamondsuit$

[34] When the hazard rate is constant (that is, F_a^v is a shifted exponential distribution), the effect of uniform (as opposed to customized) pricing on targeting solely depends on the viewers' elasticity of substitution.

Accordingly, anonymous pricing for advertising slots (e.g., as a result of regulation banning the use of scores) makes advertisers be more often matched (relative to laissez-faire) to those publishers whose profile is closer to their ideal audience if the distribution F_a^v has thin tails. This condition is testable. Analogous testable conditions can be derived for other applications of our model.

Welfare

The result in Proposition 3 can also be used to study the welfare implications of uniform-price obligations. To see this, suppose that targeting is higher under uniform pricing (on side a) than under customized pricing (on both sides). Then, under uniform pricing, 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 profile is far from their ideal one.

This observation permits us to adapt results from the third-degree price discrimination literature to the matching environment under consideration here to identify conditions under which the 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 the q-th unit of the x_b-agents at the marginal price p_a' is given by

$$D_a\left(p_a'|x_b;x_a\right) = 1 - F_a^v\left(\hat{v}_{x_b}\left(p_a'|x_a\right)\right)$$

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

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

denote the convexity of the demand by the x_a-agents for the q-th unit of the x_b-agents.[35] Before proceeding, we impose the following additional regularity condition.

Condition 1. [NDR] Nondecreasing Ratio: For any $(x_a, x_b) \in X_a \times X_b$, any q, 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.

Condition 1 is typically satisfied in the main applications of our model.[36] For instance, in the context of Example 1, this condition holds when F_a^v is uniform, or exponential. We then have the following result:

Proposition 4. ***(comparison: welfare side** a**)*** *Suppose Condition NDR holds and either one of the following alternatives is satisfied:*

[35]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'.

[36]More broadly, Condition 1 guarantees the quasi-concavity of an auxiliary problem in the proof of Proposition 4. We see it mostly as a technical requirement.

1. targeting is higher under uniform pricing than under customized pricing and, for any p_a' and x_b, $CD_a\left(p_a'|x_b;x_a\right)$ is decreasing in $|x_a - x_b|$;

2. targeting is higher under customized pricing than under uniform pricing and, for any p_a' and x_b, $CD_a\left(p_a'|x_b;x_a\right)$ is increasing in $|x_a - x_b|$.

Then welfare of the side-a agents is higher under uniform pricing on side a (and customized pricing on side b) than under customized pricing on both sides.

The next example applies Proposition 4 to the ad-exchange application. We assume that Condition NDR holds.

Example 8. (***ad exchange* - continued**) Consider the ad exchange market of Example 1. Advertisers' profits are higher under uniform pricing, and so is the level of targeting, when F_a^v has an increasing hazard rate and its convexity function

$$CF_a^v(v_a) \equiv -\frac{d^2F_a^v(v_a)}{dv_a^2}\left(\frac{dF_a^v}{dv_a}\right)^{-1}v_a = -\frac{f_a^{v\prime}\left(v_a\right)}{f_a^v\left(v_a\right)}v_a$$

is weakly decreasing. The latter condition is satisfied by the uniform distribution of Examples 3, 4, and 5. More broadly, a sufficient condition for the above monotonicity property is that the density $f_a(v_a)$ is log-convex and is either increasing or does not decrease "too fast". $\diamondsuit$

Condition NDR, as well as the convexity properties of the demand functions in Proposition 4, 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 3 is key to the result in Proposition 4. It permits us to identify "stronger markets," in the sense of Aguirre et al. (2010), with those for matches involving agents with closer profiles (alternative 1) or more distant profiles (alternative 2). Once the connection between targeting and price customization is at hand, the welfare implications of customized pricing naturally parallel those in the third-degree price discrimination literature.

We now investigate the welfare effect on side b that results from the imposition of uniform pricing on side a. For this purpose, we introduce the following condition, which implies Condition NDR.

Condition 2. [M] **Monotonicity**: For any $(x_a, x_b) \in X_a \times X_b$, any q, the convexity function $CD_a\left(p_a'|x_b;x_a\right)$ is weakly positive and nondecreasing in the marginal price p_a' for the q-th unit of the x_b-agents.

Similarly to Condition 1, the regularity Condition 2 is commonly satisfied. For instance, in the context of the the ad exchange market, this condition holds when F_a^v is uniform (as assumed in Examples 3, 4, and 5), or exponential.

Proposition 5. ***(comparison: welfare side** b)* *Suppose Condition M holds and preferences on side b are profile-invariant. Then, under either one of the alternatives of Proposition 4, welfare of* *the side-b agents is higher under uniform pricing on side a (and customized pricing on side b) than under customized pricing on both sides.*

Proposition 5 adapts to the matching environment under consideration a result on how third-degree price discrimination affects the quantity traded. When demand is more convex in strong than weak markets (as implied by either alternative from Proposition 4 and Condition M), the quantity of side-a agents consuming every q-th unit of each x_b-agents is higher under uniform pricing. By reciprocity, this implies that every side-b agent is assigned a larger quantity of matches. Incentive compatibility then implies that the indirect utility of every side-b agent increases.

For simplicity, Proposition 5 requires that preferences on side b are profile-invariant. In the absence of this assumption, the result of Proposition 5 remains true provided uniform pricing on side a increases targeting (alternative 1 of Proposition 5). If, however, uniform pricing on side a decreases targeting relative to customized pricing on both sides (alternative 2), then the effect of uniform pricing on the welfare of the side-b agents hinges on the comparison between the loss from less targeting and the gain in information rents from the increase in the quantity of matches each side-b agent is granted access to.

More broadly, Propositions 4 and 5 illustrate the usefulness of Proposition 3 in connecting our matching environment to the literature on third-degree price discrimination. We expect this parallelism to be useful in studying other forms of price regulation in matching.

5 Discussion

Consistently with what assumed in most models of horizontal differentiation, the analysis above assumes that agents' profiles are uniformly distributed over the circle. It also assumes that vertical dimensions are drawn independently from the horizontal ones. Both assumptions can be relaxed (allowing for correlation between both dimensions and non-uniformity of the horizontal dimension), while retaining all the qualitative results of our analysis. The proofs in the Appendix consider this more general environment.

The analysis in the previous sections also assumes that the agents' profiles are observable on *both* sides of the market. All the results extend to settings in which the profiles are private information on both sides, provided that one continues to assume that the profiles are uniformly distributed over the circle and that the vertical dimensions are drawn independently from the horizontal ones. As we show in the Appendix, this is because the equilibrium matching demands when profiles are observable are "horizontal translations" across the various profiles. In this case, when the profiles are the agents' private information, all agents prefer to select the matching plans designed for their profile (paying the same price as when the profiles are observable) than those designed for alternative profiles. This property, however, may require that the platform restricts the set of possible customizations that each agent can choose from. Namely, the available customizations must coincide with the collection

$(\mathbf{s}_a^*(v_a, x_a))_{v_a \in V_a}$ of matching sets designed for x_a-agents of various vertical types v_a, when the profiles are public on both sides.[37] This is carefully shown in the Appendix.

Similar conclusions extend to settings in which profiles are observable on one side but private information on the other. Suppose, for example, that profiles are private on side a and public on side b. By offering the same profit-maximizing tariffs as when the profiles are public on both sides, the platform induces all agents to select the matching plans designed for them when (i) the vertical dimensions are drawn independently from the horizontal ones on both sides, and (ii) the horizontal dimensions are distributed uniformly on the side where they are public (side b) but are distributed arbitrarily on the side where they are private (side a). In fact, under such distributional assumptions, the platform's pricing problem on side a is symmetric across any two profiles. This is because of two reasons. First, the profile of any agent from side a provides no information about the agent's vertical preferences (by virtue of the vertical dimensions being drawn independently from the horizontal ones on side a). Second, when the profiles are public on both sides, the gross utility that each type (v_a, x_a) obtains from the matching set $\mathbf{s}_a^*(v_a, x_a)$ coincides with the gross utility obtained by each type $(v_a, x_a + \delta)$ from the matching set $\mathbf{s}_a^*(v_a, x_a + \delta)$, and this is true for all $\delta \in (0, 1/2]$. In fact, the matching set $\mathbf{s}_a^*(v_a, x_a + \delta)$ is a *parallel translation* of the matching set $\mathbf{s}_a^*(v_a, x_a)$ by δ units of distance along the horizontal dimension.[38] As a result, when profiles are public only on the b-side and the platform offers the same tariffs as when they are public on both sides, all agents find it optimal to choose the same matching sets as when profiles are public on both sides.

In the absence of distributional restrictions, the platform need not be able to implement the same allocations as when profiles are public on both sides. This should not surprise given the multi-dimensionality of the agents' private information. Identifying the profit-maximizing tariffs in the general case is notoriously complex. The restrictions above permit one to capture many of the key trade-offs that platforms face in the design of the matching tariffs, while retaining tractability.

Lastly, the analysis in the previous sections assumes that match values depend on profiles and on own vertical types, but not on the partners' vertical types. This assumption may be appropriate for settings in which the vertical types capture the partners' willingness to pay, their physical or opportunity cost of trading, and, more generally, traits that are unlikely to directly determine one's attractiveness. For instance, in the ad-exchange application of Example 1, it seems unlikely that an advertiser's value of placing an ad to a publisher's website depends on the publisher's opportunity cost of space. Likewise, in the media application of Example 2, it seems unlikely that viewers' utility from accessing a content provider's material depends on the royalties paid by the provider to the content suppliers (e.g., the channels, in case of cable TV).

[37] In the absence of such restrictions, an agent of type $\theta_a = (v_a, x_a)$ may find it optimal to choose the matching plan designed for agents with profile $x_a' \neq x_a$ and then select a matching set that no x_a'-agent would have selected when profiles are public on both sides. In other words, without the aforementioned restrictions, the matching allocations induced when profiles are public on both sides need not be implementable (at the same cost) when profiles are private.

[38] in turn, this follows from the assumption that vertical dimensions are drawn independently from the horizontal ones on both sides, together with the assumption that the horizontal dimensions are uniformly distributed on the b-side.

In certain settings, though, an agent's vertical type may be a proxy for various vertical traits that, together with the agent's horizontal profile, determine the agent's attractiveness. In employment relationships, for example, a worker's vertical type often coincides with the worker's productivity, in which case the value that firms derive from being matched with the worker clearly depends on the latter's vertical type. In the Appendix, we show that all the results in the previous sections extend to such richer settings (in fact, the proofs for all the results are established under this richer payoff specification). This extra generality makes some of the formulas a little heavier but does not change the nature of the results, provided that both the true and the virtual payoffs are (weakly) increasing in the partners' vertical types. In the absence of such a monotonicity condition, the optimal tariffs cannot be guaranteed to induce demands with a threshold structure and the characterization of the corresponding matching sets is significantly more complex.

Finally, all the results in the article are for settings in which payoffs are (i) additively separable in the matches, and (ii) invariant to who else from the same side is matched with those agents from the opposite side included in one's matching set. The first assumption (that payoffs are additively separable) simplifies the characterization of the optimal thresholds but is not key to the results (see, e.g., Gomes and Pavan (2016) for a setting with non-separable payoffs, but in which payoffs are heterogenous only over a vertical dimension). Relaxing the second assumption (that payoffs are invariant to who else from the same side is matched with the same partner) introduces additional complexity, but is worth examining in future work. In a recent article, Valenzuela-Stookey (2021) shows that optimal tariffs induce demands with a threshold structure in markets with congestion effects but where the horizontal dimensions are not known to the agents at the time of matching and are only imperfectly correlated with the vertical ones. Extending the analysis in that article to a setting in which the horizontal dimensions are known at the time the agents choose the matching sets is challenging but important for this literature.

In the same vein, the analysis assumes that the only relevant feasibility constraints are the reciprocity ones. In certain settings, platforms may face additional feasibility constraints. For instance, in the ad-exchange application of Example 1, each publisher may be able to display only a certain number of ads within the relevant time frame. Likewise, in labor relationships, a worker who is employed by one firm may not be able to work for others. In addition, the platform itself may face constraints on the number of matches it can induce in each period (think of a mall with limited space, or an hospital with limited surgery rooms). Such capacity constraints introduce additional complexities in the design of the optimal matching plans. However, we expect the key trade-offs identified in this article to continue to play an important role also in these richer settings.[39]

[39]See also Fershtman and Pavan (2021) for a discussion of the role of such constraints in the design of matching auctions.

6 Related literature

This article studies many-to-many matching (with monetary transfers) in markets in which the agents' preferences are both vertically and horizontally differentiated. Related are Jeon et al. (2022) and Gomes and Pavan (2016). The first article studies the provision of quality by a platform in a setting where quality provision enhances match values. The second article studies the inefficiencies of the matching allocations under profit maximization. Both articles assume that preferences are only vertically differentiated, thus ignoring the issues of targeting and price customization that emerge when preferences are also horizontally differentiated and that are the focus of the present article. Importantly, neither of the above works studies the implications of uniform-price obligations, which is one of the key contributions of the present article.[40]

Fershtman and Pavan (2021) also studies many-to-many matching in a model in which preferences are both vertically and horizontally differentiated. The focus of that article, however, is bidding in dynamic matching markets in which agents arrive over time, experience shocks to their match values, and are repeatedly re-matched. The present article, instead, abstracts from dynamics and focuses on how uniform-price obligations impact targeting and welfare.

Related are also Jullien and Pavan (2019), and Tan and Zhou (2021). The former article studies platform competition in markets where agents' preferences for the products of different platforms are heterogenous but where all agents have the same preferences for interacting with agents from the opposite side of the market. The latter article 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 and exhibit both within-side and across-sides network effects. These articles, however, abstract from (second and third-degree) price discrimination, which is the focus of the present article. Price discrimination in matching markets is examined in Damiano and Li (2007) and Johnson (2013). Contrary to this article, these works consider markets where matching is one-to-one and where agents' preferences are differentiated only along a vertical dimension.[41]

The present article 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 preference structures are examined in the matching literature surveyed in Roth (2018) (see also Kojima (2017) and Pathak (2017) for a detailed discussion of some of the recent contributions). This literature is methodologically distinct from the current article, in that it focuses on solution concepts

[40] Both the present article and Gomes and Pavan (2016) use mechanism design to characterize the properties of the optimal matching sets and to establish that the latter have a threshold structure (with the thresholds profile-specific in the present article). However, to derive the optimal tariffs under uniform pricing, the present article tackles a new problem in which the designer faces a constraint directly on the shape of the implementing tariffs. Such constraint is novel to the mechanism design literature (where transfers are typically obtained as a residual and are pinned down by the familiar envelope formula). Tackling this additional constraint is the major technical contribution of the present article.

[41] See also Belleflamme and Peitz (2020) for a recent study of price discrimination in platform market with network effects.

such as stability and typically does not allow for transfers.[42]

More broadly, markets where agents purchase access to other agents are the focus of the literature on two-sided markets (see Belleflamme and Peitz (2017, 2021) and Jullien et al. (2021) for some recent overviews). Most of this literature, however, restricts attention to a single network, or to mutually exclusive networks. Ambrus et al. (2016) relax this structure by proposing a model of competing media platforms with overlapping viewerships (i.e., multi-homing). By contrast, the present article assumes a monopolistic market, but introduces a richer preferences structure (allowing for horizontal tastes for matches), which enables us to study targeting and price customization in such markets.[43]

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 (1993) study the provision of quality/quantity in markets where agents possess private information about a vertical dimension of their preferences. 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.[44]

The article also contributes to the literature on third-degree price discrimination. In addition to the article by Aguirre et al. (2010) mentioned above, see Bergemann, Brooks, and Morris (2015) for an overview of this literature and for recent developments. The latter article 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). The present article differs from the above two articles in its preferences structure and in the two-sidedness of the platform's problem.

Related is also the literature on bundling (see, among others, Armstrong (2013), Hart and Reny (2015), and the references therein). The present article differs from that literature in two important aspects. First, in our setting, preferences can be decomposed into a vertical and a horizontal dimension. The bundling literature, instead, assumes a more general preferences 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

[42] See the book "Market Design" by Haeringer (2018) and the forthcoming book "Online and Matching-Based Market Design" for a connection between the two literatures.

[43] Most of the literature on two-sided markets assumes that platforms price discriminate across sides but not within side. For models of within-side price discrimination, see also Halaburda and Yehezkel (2013), Reisinger (2014), Gomes and Pavan (2016), Belleflamme and Peitz (2020), and Jeon et al. (2022).

[44] Related is also Balestrieri and Izmalkov (2015). That article 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 article is information disclosure, whereas the focus of the present article is matching, targeting, and price customization.

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 article 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 focusing on the matching outcomes that emerge in platform markets where the matching between the advertisers and the publishers (or content providers) is mediated. Contrary to some of the articles 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 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 preferences signals provided by the consumers. Relatedly, Eliaz and Spiegler (2020) consider the problem of a profit-maximizing advertising platform eliciting the advertisers' profiles so as to match them to consumers with preferences for diversity. These articles do not investigate the effects of uniform pricing, but rather focus on the incentives of firms to truthfully reveal their "ideal audiences."

7 Concluding remarks

The article proposes a new model of (platform-mediated) many-to-many matching for markets in which agents' preferences are both vertically and horizontally differentiated. The model can be used to study the effects on prices, the composition of the matching sets, and welfare, of uniform-price obligations that hinder platforms' ability to condition prices on agents' "profiles," as in the case of privacy regulations preventing online retailers from conditioning prices on buyers' age, gender, physical location, and various other demographic traits.

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, depending on testable characteristics of consumer demand. Related conditions can also be used to evaluate whether or not the imposition of uniform-price 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 a market). Ad exchanges such as AppNexus, AOL's Marketplace, Microsoft Ad Exchange, OpenX, Rubicon Project Exchange, and Smaato, use sophisticated pricing algorithms whereby prices depend not only on volumes but also 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 article. Although 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 lacks a formal model shedding light on how matching demands and welfare are affected by such changes. Our article contributes to such a debate by offering a stylized, yet quite flexible, framework to analyze market outcomes under uniform pricing.

We conclude by discussing a few limitations of our analysis and avenues for future research. First, our analysis abstracts from platform competition. Second, and related, it assumes platforms have the power to set prices on both sides of the market. Although 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 et al. (2022) for related issues).

8 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 (a) the vertical and horizontal dimensions need not be independent, (b) the agents need not be uniformly distributed over the circle of the cylinder, (c) the profiles x_k need not be observable on either side of the market, and (d) the payoff that each side-k agent of type $\theta_k = (v_k, x_k)$ derives from being matched to each side-l agent of type $\theta_l = (v_l, x_l)$ may depend also on the latter agent's vertical type v_k (in the baseline model, it depends on θ_l only through the horizontal type x_l). To accommodate for these enrichments, continue to denote by F_k^v the marginal distribution of F_k with respect to the vertical dimension v_k, but now let F_k^x denote the marginal distribution of F_k with respect to the horizontal dimension x_k (with density f_k^x) — recall that, in the baseline model, $F_k^x(x_k) = x_k$ and $f_k^x(x_k) = 1$, for all $x_k \in X_k = [0,1]$). Then let $F_k^{v|x}$ denote the distribution of v_k conditional on x_k and assume that the latter is absolute continuous with density $f_k^{v|x}$ and hazard rate $\lambda_k^{v|x} \equiv f_k^{v|x}/[1 - F_k^{v|x}]$. Next, assume that the payoff that each side-k agent of type $\theta_k = (v_k, x_k)$ derives from being matched to each side-l agent of type $\theta_l = (v_l, x_l)$ is given by $u_k(v_k, v_l, |x_k - x_l|)$, with $u_k(v_k, v_l, |x_k - x_l|)$ strictly increasing in v_k, weakly increasing in v_l and weakly decreasing in the circular distance $|x_k - x_l|$ (strictly, on one of the two sides). Continue to denote by $\varphi_k(\theta_k, \theta_l)$ the "virtual" value that each side-k agent of type $\theta_k = (v_k, x_k)$ derives from being batched to each side-l agent of type $\theta_l = (v_l, x_l)$, but with the function now taking the more general form

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

Assume that $\varphi_k(\theta_k, \theta_l)$ satisfies the same properties as $u_k(v_k, v_l, |x_k - x_l|)$. By this we mean the following. First, φ_k is strictly increasing in v_k, and weakly increasing in v_l. Second, fixing (v_k, v_l, x_l), φ_k is increasing (alternatively, decreasing) in x_k if and only if u_k is increasing (alternatively, decreasing) in x_k. Third, fixing (v_k, x_k, v_l), φ_k is increasing (alternatively, decreasing) in x_l if and only if u_k is increasing (alternatively, decreasing) in x_l.

To discuss the possibility that agents may possess private information also about their profiles, i.e., about their horizontal dimensions x_k, consider the following four scenarios:

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

The baseline model corresponds to Scenario (i). Under this scenario, all the results in the main body extend to the more-general structure described above (as one can verify from the various proofs below). Under Scenarios (ii)-(iv), instead, all the results in the main body continue to hold provided that a certain combination of the following two conditions holds (see the proof of Lemma 1 below):

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

[Sy$_k$] 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 I$_k$ requires that the vertical dimensions v_k be drawn independently from the profiles x_k. Condition Sy$_k$ strengthens the independence condition by further requiring that profiles be uniformly distributed over $X_k = [0, 1]$, as assumed in the baseline model.[45]

To accommodate for the possibility that profiles are private information, we need to generalize the notion of customized tariffs, as follows:

Definition 5. The tariff T_k is *customized* if there exists a collection of quadruples

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

one for each profile $x_k \in X_k$, with $\mathcal{S}_k(x_k) \subseteq \Sigma(\Theta_l)$ denoting a set of *permissible customizations*, such that each side-k agent selecting the plan indexed by x_k and then choosing the customization $\mathbf{s}_k \in \mathcal{S}_k(x_k)$ from the set of permissible customizations $\mathcal{S}_k(x_k)$, is asked to make a total payment equal to[46]

$$T_k(\mathbf{s}_k|x_k) = \underline{T}_k(x_k) + \int_0^1 \rho_k(q_{x_l}(\mathbf{s}_k)|x_l; x_k)dx_l, \tag{15}$$

with $\rho_k\left(q|x_l; x_k\right) = 0$ if $q = q_{x_l}\left(\underline{\mathbf{s}}_k(x_k)\right)$, i.e., if the quantity of x_l-agents included in the matching set coincides with the level specified in the baseline configuration, $x_l \in [0, 1]$.

Relative to Definition 2 in the main text, Definition 5 adds the requirement that a customization must be permissible, that is, it has to belong to the collection of possible customizations $\mathcal{S}_k(x_k) \subseteq \Sigma(\Theta_l)$. As we show below, when profiles are public on side k, without loss of optimality, the platform can set $\mathcal{S}_k(x_k) \equiv \Sigma(\Theta_l)$, in which case Definitions 2 and 5 coincide.

Proof of Lemma 1. The proof below establishes the following result, for which the claim in the main text is a special case. *Consider the more general environment described above and suppose that the environment satisfies the properties of one of the following four cases: Scenario (i); Scenario*

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

[46] The payment specified by the tariff for any non-permissible customization $\mathbf{s}_k \notin \{\cup \mathcal{S}_k(x_k) : x_k \in X_k\}$ can be taken to be arbitrarily large to guarantee that no type finds it optimal to select any such customization. The existence of such payments is guaranteed by the assumption that u_k is bounded, $k = a, b$. Furthermore, in case profiles 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, profiles are public, the collection of matching plans $\{(\underline{\mathbf{s}}_k(x_k), \underline{T}_k(x_k), \rho_k(\cdot|\cdot; x_k), \mathcal{S}_k(x_k)) : x_k \in X_k\}$ 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 profile, de facto each agent faces a tariff specifying a single price for each set.

(ii) along with Conditions I_a and Sy_b; Scenario (iii) along with Conditions Sy_a and I_b; Scenario (iv) along with Conditions Sy_a and Sy_b. Then, the following are true:

1. *there exists a pair of customized tariffs $(T_k^*)_{k=a,b}$ that is profit-maximizing;*

2. *the matching demands $(\mathbf{s}_k^*)_{k=a,b}$ consistent with the profit-maximizing customized tariffs $(T_k^*)_{k=a,b}$ are described by threshold functions $t_k^* : \Theta_k \times X_l \to V_l$ such that*

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

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

3. *When profiles are public on side $k \in \{a, b\}$, $\mathcal{S}_k(x_k) = \Sigma(\Theta_l)$,* for all $x_k \in X_k$.

As anticipated in Section 5, appropriate combinations of the above two conditions guarantee that the platform can price discriminate along the agents' profiles, without leaving the agents extra rents for the private information they may possess regarding the horizontal dimension of their preferences. To gain some intuition, consider Scenario (ii) and assume that Conditions I_a and Sy_b hold. Suppose that the platform offers the same tariffs as under Scenario (i), i.e., those that it offers when profiles are observable on both sides of the market. Then, note that each type $\theta_a = (v_a, x_a)$ prefers the matching set $\mathbf{s}_a^*(v_a, x_a)$ designed for him to the set $\mathbf{s}_a^*(v_a, x_a + \delta)$ designed for an agent with the same vertical type but a different profile $x_a + \delta$, and this is true for all $\delta \in [0, 1/2]$. This is because, under Conditions I_a and Sy_b, $\mathbf{s}_a^*(v_a, x_a + \delta)$ is a *parallel translation* of the matching set $\mathbf{s}_a^*(v_a, x_a)$ by δ units of distance, along the horizontal dimension. The price of $\mathbf{s}_a^*(v_a, x_a)$ is the same as that of $\mathbf{s}_a^*(v_a, x_a+\delta)$ and the two sets contain the same total measure of agents. However, $\mathbf{s}_a^*(v_a, x_a)$ contains more types closer to θ_a's ideal profile. Hence, θ_a must prefer $\mathbf{s}_a^*(v_a, x_a)$ to $\mathbf{s}_a^*(v_a, x_a+\delta)$. As explained in Section 5, the platform must however restrict the set of possible customizations $\mathcal{S}_a(x_a)$ that each agent selecting the x_a-plan may choose from. The set $\mathcal{S}_a(x_a)$ must coincide with the collection $(\mathbf{s}_a^*(v_a, x_a))_{v_a \in V_a}$ of matching sets designed for the various x_a-agents of different vertical types.[47] When the sets $\mathcal{S}_a(x_a)$ are so restricted, each type θ_a finds it optimal to select the plan designed for the profile x_a and then select the set $\mathbf{s}_a^*(v_a, x_a)$. By offering the same tariffs as in Scenario (i) and restricting the customizations to satisfy the properties above, the platform thus induces all agents to choose the same matching set as when profiles are observable on both sides.

A symmetric situation applies to Scenario (iii). Finally consider Scenario (iv). Arguments similar to those for Scenarios (ii) and (iii) imply that, when Conditions Sy_a and Sy_b jointly hold, the platform can implement the same allocations as in Scenario (i) by letting the agents reveal their profiles.

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

[47] In the absence of such restrictions, an agent of type $\theta_k = (v_k, x_k)$ choosing the plan $x'_k \neq x_k$ may find it optimal to select a matching set that no x'_k-agent would have selected under Scenario (i).

$(\mathbf{s}_k(\theta_k), p_k(\theta_k))$ denotes the allocation (matching set and total transfer) specified by the mechanism for each side-k agent reporting θ_k.

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

$$p_k(\theta_k) = \int_{\mathbf{s}_k(\theta_k)} u_k\left(v_k, v_l, |x_k - x_l|\right) dF_l(\theta_l) - \int_{\underline{v}_k}^{v_k} \int_{\mathbf{s}_k(y,x_k)} \frac{\partial u_k}{\partial v_k}\left(y, v_l, |x_k - x_l|\right) dF_l(\theta_l) dy, \qquad (16)$$
$$- U_k(\underline{v}_k, x_k),$$

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

Using (16), the platform's profits under any incentive-compatible mechanism can then be written as

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

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

$$\sum_{k=a,b} \int_{\Theta_k} \left\{ \int_{\mathbf{s}_k(\theta_k)} \varphi_k\left(\theta_k, \theta_l\right) dF_l(\theta_l) \right\} dF_k(\theta_k) \qquad (17)$$

subject to the reciprocity condition

$$\theta_l \in \mathbf{s}_k(\theta_k) \Longleftrightarrow \theta_k \in \mathbf{s}_l(\theta_l), \quad l, k \in \{a, b\},\ k \neq l. \qquad (18)$$

Hereafter, we first describe the matching sets that maximize (17) 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 (17) subject to (18), and where the functions $p_k(\cdot)$ are as in (16), 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 $\Theta_k \times \Theta_l$:

$$\nu_k(E) \equiv \int_E m_k(\theta_k, \theta_l) dF_k(\theta_k) dF_l(\theta_l). \qquad (19)$$

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 (17) can be rewritten as

$$\sum_{k,l=a,b,\, l \neq k} \int_{\Theta_k \times \Theta_l} \varphi_k\left(\theta_k, \theta_l\right) d\nu_k(\theta_k, \theta_l) = \int_{\Theta_k \times \Theta_l} \triangle_k(\theta_k, \theta_l) m_k(\theta_k, \theta_l) dF_k(\theta_k) dF_l(\theta_l), \qquad (20)$$

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 (20) 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, and weakly increasing in v_l, $k, l = a, b,\ l \neq k$, then implies that the matching rule that maximizes (17) subject to the reciprocity condition (18) 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 $\varphi_k(\theta_k, \theta_l)$ satisfy the same monotonicity properties as the true values[48] also implies that, for ay $x_k, x_l \in [0, 1]^2$, $t_k^*(\theta_k, x_l)$ is non-increasing in v_k, and, 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 (16), 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

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

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.

[48] Recall that this means that φ_k (1) is strictly increasing in v_k and weakly increasing in v_l, (2) is increasing (alternatively, decreasing) in x_k if and only if u_k is increasing (alternatively, decreasing) in x_k, and (3) is increasing (alternatively, decreasing) in x_l if and only if u_k is increasing (alternatively, decreasing) in x_l.

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

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

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

The following 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 (16).*

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 (16), 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 of Property 1. 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, the above direct revelation mechanism is either incentive-compatible, or it can be augmented to implement the same allocations specified 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, profiles are public on both sides. That the mechanism is ICV implies that any deviation along the vertical dimension is unprofitable. Furthermore, because profiles 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 Sy_b hold. Again, because profiles 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

$$\int_{\underline{v}_a}^{v_a} \int_{\mathbf{s}_a(y,x_a)} \frac{\partial u_a}{\partial v_a} (y, v_b, |x_a - x_b|) \, dF_b(\theta_b) dy \geq \int_{\underline{v}_a}^{\hat{v}_a} \int_{\mathbf{s}_a(y,\hat{x}_a)} \frac{\partial u_a}{\partial v_a} (y, v_b, |\hat{x}_a - x_b|) \, dF_b(\theta_b) dy \tag{21}$$
$$+ \int_{\mathbf{s}_a(\hat{v}_a,\hat{x}_a)} [u_a (v_a, v_b, |x_a - x_b|) - u_k (\hat{v}_a, v_b, |\hat{x}_a - x_b|)] dF_b(\theta_b).$$

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

$$\int_{\mathbf{s}_a(v_a,x_a)} \frac{\partial u_a}{\partial v_a} (v_a, v_b, |x_a - x_b|) \, dF_b(\theta_b) = \int_{v_b \in V_b} \int_{\delta \in [0,1/2]} \frac{\partial u_a (v_a, v_b, \delta)}{\partial v_a} W(\mathrm{d}(v_b, \delta); \theta_a), \tag{22}$$

where $W(\cdot; \theta_a)$ is the measure over $V_b \times [0, 1/2]$ defined by the types included in the matching set $\mathbf{s}_a(v_a, x_a)$, under the proposed mechanism.[49] It is also easy to see that, when Conditions I_a and Sy_b hold, under the proposed mechanism, the expression in (22) is invariant in x_a. That is, for any (v_b, δ), $W((v_b, \delta); \theta_a) = W((v_b, \delta); \theta'_a)$ for any $\theta_a, \theta'_a \in \Theta_a$ with $v_a = v'_a$.[50] This means that

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

By the same arguments,

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

Furthermore, because the threshold functions $t_k^* (\theta_k, x_l)$ are non-decreasing in the distance $|x_l - x_k|$, we have that

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

It follows that the right hand side of (21) is smaller than

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

$$+ \textstyle\int_{\mathbf{s}_a(\hat{v}_a,x_a)} [u_a (v_a, v_b, |x_a - x_b|) - u_a (\hat{v}_a, v_b, |x_a - x_b|)] \, dF_b(\theta_b),$$

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 (21) 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.

[49] That is, for any $(v_b, \delta) \in V_b \times [0, 1/2]$, $W((v_b, \delta); \theta_a\}$ is the measure of agents from side b whose vertical type is less than v_b and whose distance from x_a is less than δ included in the matching set $\mathbf{s}_a(v_a, x_a)$ of type θ_a under the proposed mechanism.

[50] 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 Sy_b is, however, necessary to guarantee, for any $(v_b, \delta) \in V_b \times [0, 1/2]$, any $v_a \in V_a$, and any $x_a, x'_a \in X_a$, the mass of side-b agents with vertical type no greater than v_b and with a profile x_b whose distance from x_a is no greater than δ included in the matching set $\mathbf{s}_a(v_a, x_a)$ is the same as the mass of side-b agents with vertical type no greater than v_b and with a profile x_b whose distance from x'_a is no greater than δ included in the matching set $\mathbf{s}_a(v_a, x'_a)$.

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

Finally, consider an environment satisfying the properties of Scenario (iv) along with Conditions Sy_a and Sy_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 (16) 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

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

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 $q \in \left(q_{x_l}\left(\mathbf{s}_k(\underline{v}_k, x_k)\right), q_{x_l}\left(\mathbf{s}_k(\bar{v}_k, x_k)\right)\right]$,

$$\begin{aligned}\rho_k(q|x_l; x_k) = &\textstyle\int_{v_l'(q;x_l)}^{\bar{v}_l} u_k\left(v_k(q; x_k, x_l), v_l, |x_k - x_l|\right) dF_l^{v|x}(v_l|x_l) f_l^x(x_l) \\ &- \textstyle\int_{\underline{v}_k}^{v_k(q;x_k,x_l)} \int_{v_l'(q_{x_l}(\mathbf{s}_k(y,x_k));x_l)}^{\bar{v}_l} \frac{\partial u_k}{\partial v_k}\left(y, v_l, |x_k - x_l|\right) dF_l^{v|x}(v_l|x_l) f_l^x(x_l) dy - \underline{T}_k(x_k)\end{aligned} \tag{23}$$

where

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

and[51]

$$v_l'(q; x_l) \equiv \left(F_l^{v|x}\right)^{-1}\left(1 - \frac{q}{f_l^x(x_l)}|x_l\right).$$

Note that in (23) we used the fact that, under the proposed mechanism, the x_l-agents included in the matching set of any x_k-agent interacting with q x_l-agents are those whose vertical type exceeds $v_l'(q; x_l)$.

[51] In other words, $v_l'(q; x_l)$ is given by the unique solution to $\left[1 - F_l^{v|x}\left(v_l'(q; x_l)|x_l\right)\right] f_l^x(x_l) = q$.

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 profiles 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 profile 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

$$\underline{T}_k(x_k) + \int_0^1 \left[\int_{v_l'(q_{x_l}(\mathbf{s}_k(v_k,x_k));x_l)}^{\bar{v}_l} u_k\left(v_k, v_l, |x_k - x_l|\right) dF_l^{v|x}(v_l|x_l) f_l^x(x_l) \right.$$

$$\left. - \int_{\underline{v}_k}^{v_k} \int_{v_l'(q_{x_l}(\mathbf{s}_k(y,x_k));x_l)}^{\bar{v}_l} \frac{\partial u_k}{\partial v}\left(y, v_l, |x_k - x_l|\right) dF_l^{v|x}(v_l|x_l) f_l^x(x_l) dy \right] dx_l - \underline{T}_k(x_k)$$

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

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 a selection negative.

Finally, to see that, when profiles 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 with profile 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 with profile x_k by setting the fee associated with the selection of any other plan sufficiently high. The separability of the agents' payoffs in the matches 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 (23), type θ_k prefers to interact with $q_{x_l}(\mathbf{s}_k(v_k, x_k))$ agents with profile x_l to any other mass of agents with the same profile x_l, irrespective of the masses of agents with profiles other than x_l that type θ_k includes in his matching set. Q.E.D.

Proof of Proposition 1. Fix a pair of profiles $x_a, x_b \in [0, 1]$. 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

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

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 a quantity, 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

$$u_k\left(v_k, v_l'(q_k; x_l), |x_k - x_l|\right) = \rho_k'\left(q_k | x_l; x_k\right), \tag{24}$$

where recall that $v_l'(q_k; x_l) \equiv \left(F_l^{v|x}\right)^{-1}\left(1 - \frac{q_k}{f_l^x(x_l)} | x_l\right)$. This is because the marginal x_l-agent that is brought to the matching set when expanding the demand for the x_l-agents starting from q_k has a vertical type $v_l = v_l'(q_k; x_l)$.

Now let ρ_k' be a shortcut for $\rho_k'(q_k | x_l; x_k)$ — because q_k and the profiles x_l and x_k are held fixed, there is no risk of confusion. Then, let $\hat{v}_{x_l}(\rho_k' | x_k)$ be the unique solution to (24). 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

$$\varphi_k\left(\left(\hat{v}_{x_l}\left(\rho_k' | x_k\right), x_k\right), \left(v_l'(q_k; x_l), x_l\right)\right) + \varphi_l\left(v_l'(q_k; x_l), x_l), \left(\hat{v}_{x_l}\left(\rho_k' | x_k\right), x_k\right)\right) = 0. \tag{25}$$

Lastly, observe that, for any such q_k,

$$\frac{\rho_k'}{\varepsilon_k\left(\rho_k' | x_l; x_k\right)} = \frac{1 - F_k^{v|x}(\hat{v}_{x_l}\left(\rho_k' | x_k\right) | x_k)}{f_k^{v|x}(\hat{v}_{x_l}\left(\rho_k' | x_k\right) | x_k)} \frac{\partial u_k}{\partial v_k}\left(\hat{v}_{x_l}\left(\rho_k' | x_k\right), v_l'(q_k; x_l), |x_k - x_l|\right). \tag{26}$$

Using the definition of φ_k from the main text together with (24) and (26), we then have that, for any such a q_k,

$$\varphi_k\left(\left(\hat{v}_{x_l}\left(\rho_k' | x_k\right), x_k\right), \left(v_l'(q_k; x_l), x_l\right)\right) = \rho_k'(q_k | x_l; x_k)\left[1 - \frac{1}{\varepsilon_k\left(\rho_k' | x_l; x_k\right)}\right]. \tag{27}$$

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

$$\varphi_l\left(\left(v_l'(q_k; x_l), x_l\right), \left(\hat{v}_{x_l}\left(\rho_k' | x_k\right), x_k\right)\right) = \rho_l'\left(q_l | x_k; x_l\right)\left[1 - \frac{1}{\varepsilon_l\left(\rho_l' | x_k; x_l\right)}\right]. \tag{28}$$

Combining (27) and (28) with (25), we obtain the result in the proposition. Q.E.D.

Proof of Proposition 2. 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 profile $x_b \in [0, 1]$, 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

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

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

subject to the feasibility constraint (3), where $D_b\left(\rho_b'(q|x_a;x_b)|x_a;x_b\right)$ is the total measure of x_b-agents that demand q or more matches with the x_a-agent at marginal price $\rho_b'(q|x_a;x_b)$, and where

$$\bar{D}_a\left(p_a'(q|x_b)|x_b\right)\equiv\int_0^1 D_a\left(p_a'(q|x_b)|x_b;x_a\right)dx_a$$

is the total measure of side-a agents that demand q or more interactions with the x_b-agents at marginal price $p_a'(q|x_b)$.

Now fix x_b and, for any $q\le f_b^x(x_b)$, recall that $v_b'(q;x_b)\equiv\left(F_b^{(v|x)}\right)^{-1}\left(1-\frac{q}{f_b^x(x_b)}|x_b\right)$. 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 any x_b and $q\le f_b^x(x_b)$, let p_a' be a short cut for $p_a'(q|x_b)$. For any p_a', then let

$$\hat{v}_{x_b}\left(p_a'|x_a\right)=\begin{cases}v_a\text{ s.t. }u_a(v_a,v_b'(q;x_b),|x_a-x_b|)=p_a' & \text{if } p_a'\in\begin{bmatrix}u_a(\underline{v}_a,v_b'(q;x_b),|x_a-x_b|),\\ u_a(\bar{v}_a,v_b'(q;x_b),|x_a-x_b|)\end{bmatrix}\\ \underline{v}_a & \text{if } p_a'<u_a(\underline{v}_a,v_b'(q;x_b),|x_a-x_b|)\\ \bar{v}_a & \text{if } p_a'>u_a(\bar{v}_a,v_b'(q;x_b),|x_a-x_b|),\end{cases}\tag{29}$$

Given the above definitions, we have that the demand by the x_a-agents for the q-th unit of the x_b-agents at the marginal price $p_a'=p_a'(q|x_b)$ is equal to

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

Now for each $x_a,x_b\in[0,1]$, each $q\le f_b^x(x_b)$, let

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

where we reintroduced the arguments of the p_a' function for clarity. 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

$$\rho_b'(\hat{q}_b(q;x_a;x_b)|x_a;x_b)=u_b(v_b'(q;x_b),\hat{v}_{x_b}\left(p_a'|x_a\right),|x_b-x_a|).\tag{30}$$

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 pins down 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 type $\theta_b = (v_b, x_b)$ is assigned a matching set equal to $\mathbf{s}_b(v_b, x_b)$ is given by

$$\int_{\Theta_b} \left\{ \int_0^1 \left[\int_{v_a'(q_{x_a}(\mathbf{s}_b(v_b,x_b));x_a)}^{\bar{v}_a} \left(u_b\left(v_b, v_a, |x_b - x_a|\right) - \frac{1-F_b^{v|x}(v_b|x_b)}{f_b^{v|x}(v_b|x_b)} \frac{\partial u_b}{\partial v_b}\left(v_b, v_a, |x_b - x_a|\right) \right) \times \right.\right.$$

$$\left.\left. \times\, dF_a^{v|x}(v_a|x_a) \right] f_a^x(x_a) dx_a \right\} dF_b(\theta_b),$$

where, for any $q \leq f_a^x(x_a)$, $v_a'(q;x_a) = F_a^{(v|x)-1}\left(1 - \frac{q}{f_a^x(x_a)}|x_a\right)$. In turn, this means that the platform's problem can be re-casted as choosing a function $p_a'(q|x_b)$ 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 any $x_b \in [0,1]$, any $q \leq f_b^x(x_b)$, the function

$$\mathcal{C}\left[p_a'(q|x_b)\right] \equiv -\int_0^1 \left\{ \int_{v_a'(D_a(p_a'(q|x_b)|x_b;x_a);x_a)}^{\bar{v}_a} \left[u_b\left(v_b'(q;x_b), v_a, |x_b - x_a|\right) - \frac{1-F_b^{v|x}(v_b'(q;x_b)|x_b)}{f_b^{v|x}(v_b'(q;x_b)|x_b)} \frac{\partial u_b}{\partial v_b}\left(v_b'(q;x_b), v_a, |x_b - x_a|\right) \right] \times \right.$$

$$\left. \times dF_a^{v|x}(v_a|x_a) \right\} f_a^x(x_a) dx_a$$

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

$$\bar{\varepsilon}_a\left(p_a'|x_b\right) \equiv -\frac{\partial \bar{D}_a\left(p_a'|x_b\right)}{\partial\left(p_a'\right)} \frac{p_a'}{\bar{D}_a\left(p_a'|x_b\right)}$$

and

$$\mathcal{C}'\left[p_a'(q|x_b)\right] = -\int_0^1 \left\{ u_b\left(v_b'(q;x_b), v_a'(D_a\left(p_a'(q|x_b)|x_b;x_a\right);x_a), |x_b - x_a|\right) \right.$$

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

$$\times f_a^{v|x}(v_a'(D_a\left(p_a'(q|x_b)|x_b;x_a\right);x_a)|x_a) f_a^x(x_a) \frac{\partial v_a'(D_a(p_a'(q|x_b)|x_b;x_a);x_a)}{\partial q} \frac{\partial D_a(p_a'(q|x_b)|x_b;x_a)}{\partial(p_a')} dx_a.$$

Now observe that

$$\frac{\partial v_a'(D_a\left(p_a'(q|x_b)|x_b;x_a\right);x_a)}{\partial q}=-\frac{1}{f_a^{(v|x)}\left(v_a'(D_a\left(p_a'(q|x_b)|x_b;x_a\right);x_a)|x_a\right)f_a^x(x_a)},$$

implying that

$$\mathcal{C}'\left[p_a'(q|x_b)\right]=-\int_0^1\left\{u_b\left(v_b'(q;x_b),v_a'(D_a\left(p_a'(q|x_b)|x_b;x_a\right);x_a),|x_b-x_a|\right)\right.$$

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

$$\times\frac{\partial D_a(p_a'(q|x_b)|x_b;x_a)}{\partial(p_a')}dx_a.$$

Also note that(30) implies that

$$u_b\left(v_b'(q;x_b),v_a'(D_a\left(p_a'(q|x_b)|x_b;x_a\right);x_a),|x_b-x_a|\right)$$

$$-\frac{1-F_b^{v|x}(v_b'(q;x_b)|x_b)}{f_b^{v|x}(v_b'(q;x_b)|x_b)}\frac{\partial u_b}{\partial v_b}\left(v_b'(q;x_b),v_a'(D_a\left(p_a'(q|x_b)|x_b;x_a\right);x_a),|x_b-x_a|\right)$$
$$=\rho_b'(\hat{q}_b(q;x_a;x_b)|x_a;x_b)\left(1-\frac{1}{\varepsilon_b\left(\rho_b'(\hat{q}_b(q;x_a;x_b)|x_a;x_b)|x_a;x_b\right)}\right).$$

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 3. The proof below is for the more general case in which the side-b preferences may depend on the profiles.

Fix $\theta_b=(v_b,x_b)$ and let $q=f_b^x(x_b)\left[1-F_b^{v|x}(v_b|x_b)\right]$. The result in Proposition 2 implies that, under uniform pricing on side a and customized pricing on side b, for any $x_a\in X_a$ such that $t_b^u(\theta_b,x_a)\in\text{Int}[V_a]$, $t_b^u(\theta_b,x_a)$ is such that[52]

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

$$+\mathbb{E}_{H(\tilde{x}_a|x_b,p_a^{u\prime})}\left[\varphi_b\left(\theta_b,\left(\hat{v}_{x_b}\left(p_a^{u\prime}|\tilde{x}_a\right),\tilde{x}_a\right)\right)\right]=0, \tag{31}$$

[52]Note that, for any $\tilde{x}_a$, $u_a(\hat{v}_{x_b}\left(p_a^{u\prime}|\tilde{x}_a\right),v_b,|x_b-\tilde{x}_a|)=u_a(t_b^u(\theta_b,x_a),v_b,|x_b-x_a|)=p_a^{u\prime}$, where, as usual, $p_a^{u\prime}$ is a shortcut for $p_a^{u\prime}(q|x_b)$.

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 $p_a^{u\prime}(q|x_b) = u_a(t_b^u(\theta_b, x_a), v_b, |x_a - x_b|)$. Note that, to arrive at (31), we used the result in Proposition 2 and the fact that, for any x_a such that $\hat{v}_{x_b}\left(p_a^{u\prime}|x_a\right) \notin Int[V_a]$, $h\left(x_a|x_b, p_a^{u\prime}\right) = 0$, whereas for any x_a such that $\hat{v}_{x_b}\left(p_a^{u\prime}|x_a\right) \in Int[V_a]$,

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

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

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

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

as shown in the proof of Proposition 2.

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

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

It is then easy to see that, for any x_a such that

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

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

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

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

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

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

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

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 $\theta_b = (v_b, x_b)$, the function

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

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

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

Fixing $\theta_b = (v_b, x_b)$, the function $L(x_a|\theta_b)$ is non-decreasing 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. Q.E.D.

Proof of Proposition 4. The proof follows from the combination of the results in Proposition 3 with the results in Proposition 1 in Aguirre et al (2010). When the environment satisfies the conditions in Part 1 of Proposition 3, starting from uniform pricing on side a, the introduction of customized pricing on side a leads to an increase in prices for nearby profiles and a reduction in prices for distant profiles. Proposition 1 in Aguirre et al (2010), along with the fact that the environment satisfies Condition NDR 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 3, 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 profiles and a reduction in prices for nearby profiles. The welfare implications of such price adjustments then follow again from Proposition 1 in Aguirre et al (2010), along with the fact that Condition NDR 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.

Proof of Proposition 5. The proof follows from the combination of the results in Proposition 3 with the results in Proposition 4 in Aguirre et al (2010). When the environment satisfies the conditions in Part 1 of Proposition 3, starting from uniform pricing on side a, the introduction of customized pricing on side a leads to an increase in prices for nearby profiles and a reduction in prices for distant profiles. Proposition 4(ii) in Aguirre et al (2010), along with the fact that the environment satisfies Condition M 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 the quantity of the side-a agents obtaining the q-th unit of the x_b-agents increases under uniform pricing. By reciprocity, this implies that every side-b agent is assigned a larger quantity of matches. Incentive compatibility then implies that the indirect utility of every type v_b increases. Likewise, when the environment satisfies the conditions in Part 2 of Proposition 3, the quantity of the side-a agents obtaining the q-th unit of the x_b-agents increases under uniform pricing. This 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 profiles and a reduction in prices for nearby profiles. Again, the result follows from Proposition 4(ii) in Aguirre et al (2010), along with the fact that Condition M 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 $|x_b - x_a|$. By reciprocity, every side-b agent is assigned a larger quantity of matches. Incentive compatibility then implies that the indirect utility of every type v_b increases. Q.E.D.

References

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

[2] Akbarpour, M., Dworczak, P. and Duke Kominers, S. “Redistributive Allocation Mechanisms.” Working Paper, Northwestern University, 2021.

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

[4] Armstrong, M. “A More General Theory of Commodity Bundling.” *Journal of Economic Theory*, Vol. 148 (2013), pp. 448-472.

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

[6] Belleflamme, P., and Peitz, M. “Platforms and Network Effects.” In L. Corchon and M. Marini, eds., *The Handbook of Game Theory and Industrial Organization.* Chelthenam, UK, Northampton, MA, USA: Edward Elgar, 2017.

[7] Belleflamme, P., and Peitz, M. "Network Goods, Price Discrimination, and Two-sided Platforms." Working Paper, University of Mannheim, 2020.

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

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

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

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

[12] Bounie, D., Dubus, A., and Waelbroeck, P. “Selling Strategic Information in Digital Competitive Markets.” *RAND Journal of Economics*, Vol. 52 (2021), pp. 283-313.

[13] Bulow, J., and Roberts, D. “The Simple Economics of Optimal Auctions.” *Journal of Political Economy*, Vol. 97 (1989), pp. 1060-1090.

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

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

[16] Damiano, E., and Li, H. “Price Discrimination and Efficient Matching*.” Economic Theory, Vol.* 30 (2007), pp. 243-263.

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

[18] Eliaz, K., and Spiegler, R. “Incentive-Compatible Advertising on Non-Retail Platforms.” *RAND Journal of Economics*, Vol. 51 (2020), pp. 323-345.

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

[20] Federal Communications Commission (FCC). “Further Report on the Packaging and Sale of Video Programming to the Public.” http://www.fcc.gov/mb/csrptpg.html, 2006.

[21] Fershtman, D., and Pavan, A. “Matching Auctions.” *RAND Journal of Economics*, Vol. 53 (2021), pp. 32-62.

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

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

[24] Haeringer, G. *Market Design: Auctions and Matching*. Boston: MIT Press, 2018.

[25] Halaburda, H. and Yehezkel, Y. "Platform competition under asymmetric information." *American Economic Journal: Microeconomics*, Vol. 5 (2013), pp. 22–68.

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

[27] Jeon, D.-S., Kim, B.-C. , and Menicucci, D. "Second-degree Price Discrimination by a Two-sided Monopoly Platform." *American Economic Journal: Microeconomics*, Vol. 14 (2022), pp. 322–369.

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

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

[30] Jullien, B., Pavan, A. , and Rysman, M. "Two-sided Markets, Pricing, and Network Effects," In A. Hortacsu, K. Ho, and A. Lizzeri, ed., *Handbook of Industrial Organization*, New York: North-Holland, 2021.

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

[32] Kojima, F. "Recent Developments in Matching Theory and their Practical Applications." *Advances in Economics and Econometrics*, Vol. 1 (2017), pp. 138-176.

[33] Kox, H., Straathof, B., and Zwart, G. "Targeted Advertising, Platform Competition, and Privacy." *Journal of Economics and Management Strategy*, Vol. 26 (2017), pp. 557-570.

[34] Maskin, E., and Riley, J. "Monopoly with Incomplete Information." *RAND Journal of Economics*, Vol. 15 (1984), pp. 171-196.

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

[36] Montes, R., Sand-Zantman, W. , and Valletti, T. "The value of personal information in online markets with endogenous privacy." *Management Science*, Vol. 65 (2019), pp. 1342-1362.

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

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

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

[40] Pathak, P. A. "What Really Matters in Designing School Choice Mechanisms." *Advances in Economics and Econometrics*, Vol. 1 (2017), pp. 176–214.

[41] Reisinger, M. "Two-part Tariff Competition between Two-sided Platforms." *European Economic Review*, Vol. 68 (2014), pp. 168–180.

[42] Rochet, J.-C., and Tirole, J. "Two-Sided Markets: A Progress Report." *RAND Journal of Economics*, Vol. 37 (2006), pp. 645-667.

[43] Roth, A. E. "Marketplaces, Markets, and Market Design." *American Economic Review*, Vol. 108 (2018), pp. 1609–1658.

[44] Saez, E. "Using Elasticities to Derive Optimal Tax Rates." *Review of Economic Studies*, Vol. 68 (2001), pp. 205-229.

[45] Tan, G., and Zhou, J. "The Effects of Competition and Entry in Multi-sided Markets" *Review of Economic Studies*, Vol. 88 (2021), pp. 1002–1030.

[46] Valenzuela-Stookey, Q. "Platform-Mediated Competition." Working Paper, Northwestern University, 2020.

[47] Wilson, R. *Non-Linear Pricing.* Oxford: Oxford University Press, 1993.

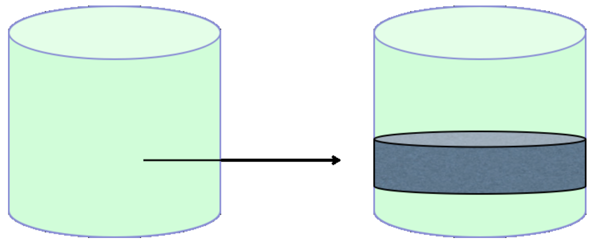

Figure 1: The cylinder model

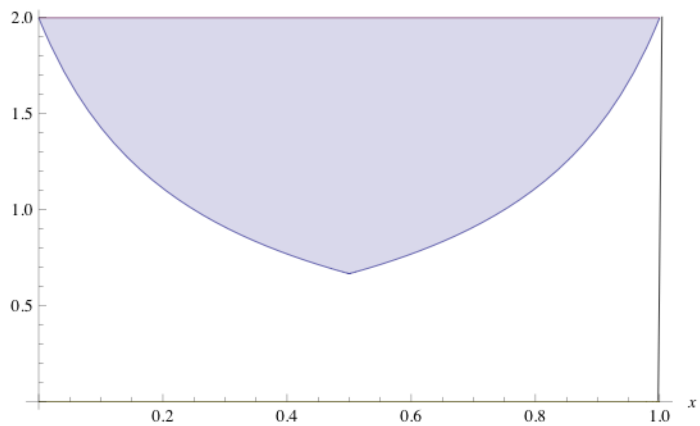

Figure 2: Matching sets under profit-maximizing tariffs. The shaded area in the figure describes the matching set of a representative agent from side a with profile $x_a = 1/2$ (the agent's vertical type is omitted for simplicity).

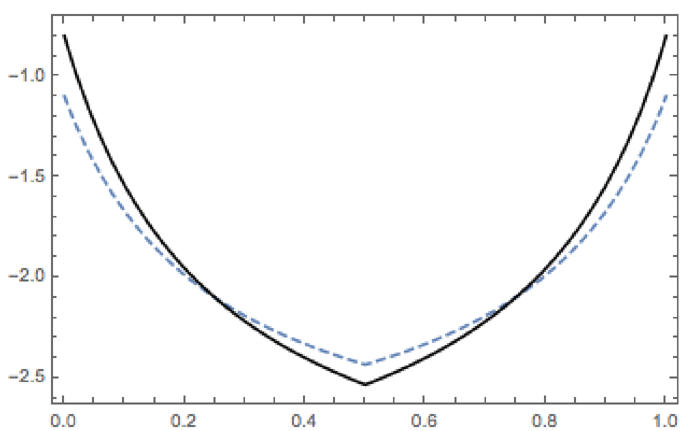

Figure 3: The threshold function $t_a^*(\theta_a, x_b)$ under customized pricing (solid black curve) and uniform pricing $t_a^u(\theta_a, x_b)$ (dashed blue curve) when customized pricing (on both sides) leads to more targeting than uniform pricing (on side a).