THE PINCHAS SAPIR CENTER
FOR DEVELOPMENT
Tel Aviv University

המרכז לפיתוח ע"ש פנחס ספיר
ליד אוניברסיטת תל אביב
"עמותה רשומה"
580010221

THE PINHAS SAPIR CENTER FOR DEVELOPMENT
TEL AVIV UNIVERSITY

Challenging Encounters and Within-Physician Practice Variability

Daniel Fershtman and Alessandro Pavan

Discussion Paper No. 2-2023

Matching Auctions*

Daniel Fershtman[†] **Alessandro Pavan**[‡]

Abstract

We study platform markets in which agents arrive gradually, experience changes to their preferences over time, and are frequently re-matched. We introduce simple auctions specifically designed for such markets. Upon joining, agents select a status that determines the weight assigned to their future bids. Each match is then assigned a score that depends on the agents' reciprocal bids and status. The matches maximizing the sum of the bilateral scores are implemented. Under certain conditions, such auctions maximize profits, welfare, or a combination of the two. We use the results to shed light on the distortions due to platforms' market power.

JEL Classification Numbers: D82, C73, L1.

Keywords: platforms, matching, asymmetric information, learning

*For helpful comments and suggestions, we are grateful to Dirk Bergemann, Yeon-Koo Che, Eddie Dekel, Drew Fudenberg, Benjamin Hermalin, Terence Johnson, Fuhito Kojima, Benny Moldovanu, Larry Samuelson, Rakesh Vohra, Jean Tirole, Asher Wolinsky, Leeat Yariv, and seminar participants at various conferences and workshops where the paper was presented. Matteo Camboni provided excellent research assistance. Pavan also thanks the National Science Foundation for financial support under the grant SES 1530798 and Bocconi University for its hospitality during the 2017-2018 academic year. Fershtman gratefully acknowledges funding from the Pinhas Sapir Center for Development. The usual disclaimer applies.

[†]Eitan Berglas School of Economics, Tel-Aviv University (danielfer@tauex.tau.ac.il)

[‡]Northwestern University (alepavan@northwestern.edu)

1 Introduction

In recent years, matching markets have been growing at an unprecedented rate, reflecting the sharing economy's role in the organization of modern business activities. In e-commerce, for example, a sizable fraction of trade is mediated by business-to-business (B2B) platforms matching vendors with procurers in search of business opportunities. Likewise, a sizable fraction of online advertising is mediated by ad exchanges, search engines, media outlets, online malls, and videogame consoles, matching advertisers with either consumers or content providers.[1] Platform-mediated matching also plays an important role in the growing market for scientific outsourcing, where intermediaries such as Science Exchange match labs with idle equipment with research units seeking to conduct experiments off-site;[2] project finance, where consulting firms match startups with lenders; lobbying, where commercial firms match interest groups with policy makers;[3] the market for private medical-tourism services, where intermediaries such as MEDIGO match patients outside of the US with US physicians providing specialized treatments;[4] and the market for organized events, where platforms such as meetings.com match clients in search of hospitality services with providers of such services.

These markets are highly dynamic, with agents experiencing frequent changes to their preferences for interacting with other agents, and with new agents arriving at the market gradually over time. As a result, agents on both sides of the market are frequently re-matched. For example, Science Exchange frequently re-matches the research units with the facilities renting out their lab equipment. Such re-matching reflects the arrival of new labs and research units over time but also changes in the research units' needs and preferences and in the

[1]See, e.g., "Marketing in the digital age: A brand new game", *The Economist*, August, 2015.

[2]See, e.g., "Uber for experiments," *The Economist*, December 6, 2014.

[3]See Allard (2008) and Kang and You (2016) for how lobbying firms provide tailored (many-to-many) matching services and dynamically price-discriminate each side of the lobbying market. See Dekel, Jackson and Wolinsky (2008) for a detailed account of how intermediaries help to buy and sell votes.

[4]See http://www.nytimes.com/2013/08/07/us/the-growing-popularity-of-having-surgery-overseas.html.

facilities' equipment and staff. Frequent re-matching is also a prominent feature of markets in which platforms mediate the interactions between advertisers and content providers.

Leveraging on recent technological advances, many platforms are also evolving from a non-discriminatory business model in which each agent on board is granted access to any other agent from the opposite side, to one where matching is *customized* (e.g., each advertiser is matched only to a subset of the content providers in the ad exchange's network). The same technological advances have also expanded the scope for using auctions to select the appropriate matches.[5] Because the match values are typically the agents' private information, and because platforms can entertain only a limited number of matches in each period due to individual (i.e., agent-specific) and aggregate (i.e., platform-specific) capacity constraints, more and more platforms are contemplating using auctions to mediate the interactions between the different sides of the market (see, e.g., Martens (2016) where it is argued that search rankings and price auctions will soon become the main tools to facilitate online matching). However, standard auction formats used for physical goods or services (e.g., first-price, second-price, English, or clock auctions) need not be appropriate for markets in which agents play the double role of buyers and inputs and where the match values are expected to evolve frequently over time. Nor are those auctions designed for the sale of multiple, but homogeneous, goods, such as the various versions of the double auction used in financial markets to sell securities, or by the Government to issue Treasuries.

In this paper, we introduce and study auctions specifically designed for dynamic matching markets in which agents arrive over time, experience frequent changes to their match values, and are repeatedly re-matched in response to variations in market conditions. The goal of the analysis is primarily normative: to identify practical mechanisms (auctions) that can be used in such markets and show that certain versions of such auctions maximize the platform's

[5] See, e.g., Pinker, Seidmann, and Vakrat (2003).

profits, welfare, or a combination of the two.

An important feature of the model is that agents have a demand for repeated interactions with agents from the opposite side. The (flow) payoff that each agent derives from each possible match is governed by two components: (i) a time-invariant vertical characteristic capturing the overall importance the agent assigns to interacting with partners from the opposite side of the market (the "vertical" type); and (ii) a vector of time-varying match-specific values capturing the evolution of the agent's information and preferences for specific partners (the "horizontal" types). These values evolve (stochastically) over time and may turn negative, reflecting the idea that certain agents may dislike certain interactions. Both the vertical and the horizontal types are the agents' private information. Agents learn their vertical types prior to joining the platform and learn their horizontal types over time after discovering who is on board. The model allows for limits on both the number of matches that each agent may participate in within a period (individual capacity constraints), as well as on the total number of matches the platform can accommodate within a period (aggregate capacity constraint). Such limits may reflect time, resource, or facility constraints, but also capture certain non-separabilities and decreasing returns to scale in the agents' preferences.

The matching auctions we propose have the following features. When joining the platform, each agent is asked to select a membership status. At any subsequent period, each agent on board is then asked to submit a vector of bids, one for each possible partner from the opposite side of the market (counterpart). Each bilateral match then receives a "score" that depends on the two agents' reciprocal bids, their membership status, and the platform's cost of implementing the match.[6] Different scoring rules reflect different objectives of the platform, ranging from profit to welfare maximization. In each period, the platform imple-

[6] We allow the platform's costs to take on negative values, reflecting the possibility that the platform may benefit from certain interactions (e.g., they may help promote the platform's matching capabilities).

ments those matches that maximize the sum of the bilateral scores, taking individual and aggregate capacity constraints into account. As in the Vickrey-Clark-Groves (VCG) and Generalized Second Price (GSP) auctions, the payments the platform asks of each agent reflect the externalities the agent imposes on others due to the individual and aggregate capacity constraints. However, contrary to these auctions, such externalities may also account for the effects of matching on the agents' informational rents. In particular, the platform may find it optimal to subsidize certain interactions and favor matches that generate lower surplus when this permits the platform to raise more revenue. In addition to charging the agents for the matches they receive over time, the platform also charges each agent a participation fee at the time of joining that depends on the selected membership status. The pricing of status reflects how the latter influences the intertemporal surplus expected by an agent at the time of joining.

To the best of our knowledge, the specific auctions we propose are not used in any real-world market. However, many of their features seem to play a role that has been recognized as important in concrete markets. For example, the use of the agents' status to distort matching and pricing to favor more prominent agents, both within and across sides, is something that resonates well with practices followed by many real-world platforms. In this respect, the weighted scores in our auctions play a role similar to that played by the "compatibility scores" used by various ad exchanges to select matches between advertises and content providers (see, e.g., Chapter 6 in Moghaddam and Shimon, 2016). In fact, online ad exchanges such as Google's DoubleClick and Microsoft's Exchange already use auctions to match advertisers with content providers. In such auctions, as in ours, advertisers bid repeatedly over time to place their ads on the website of multiple content providers and, over time, content providers ask different ad-specific prices to display the ads.[7] Dynamic compatibility scores where the

[7] See, e.g., Mansour, Muthukrishnan, and Nisan (2012).

bids of agents from different sides of the market are weighted so as to respond to changes in the fit and efficacy of the advertisement, without damaging broadcaster reputation, play an important role in such auctions.[8] Furthermore, participating agents from both sides of the market are often re-matched and are charged fees to join the platform.[9]

In the auctions we propose, at all histories (including those off-path), equilibrium bidding is truthful (Theorem 1). Bidding one's myopic values for all matches is optimal because the matches under truthful bidding maximize a weighted sum of all agents' payoffs, net of the platform's matching costs and, in case of profit-maximization, net of the agents' information rents. This property, together with the fact that the auction's payments reflect the imposed externalities, guarantees that each agent finds it optimal to bid truthfully at all periods, irrespective of the agent's beliefs about other agents' current and past types, and independently of past matches. As a result, the proposed auctions can be made fully transparent: At the end of each period, all membership statuses, bids, matches, and payments are disclosed to all agents on board.

That, once on board, agents find it optimal to bid truthfully at each period follows from arguments similar to those in the literature on VCG mechanisms (but adjusted for the fact that the induced matches need not be efficient). That agents find it optimal to join immediately upon arrival and select the "right" membership status is not obvious and follows from the interaction of three novel monotonicities. First, the intertemporal match quality that each agent expects from joining the auction declines over time. This property guarantees that agents do not benefit from postponing their joining. Second, when all agents bid truthfully, the net present value of the quality of the matches that each agent expects from selecting

[8]See, e.g., https://support.google.com/adxseller/answer/2913506?hl=en&ref topic=3376095.

[9]Auctions are also used in the market for personalized display ads by search engines such as Yahoo! and Google. These intermediaries have originally used variations of the second-price auction, the so-called GSP auction. In 2012, however, Google switched from the GSP auction to a VCG auction on the grounds that dynamic bid re-optimization is easier under the VCG protocol. See, e.g., Edelman, Ostrovsky, and Schwarz (2007), Gomes and Sweeney (2014), Harris and Varian (2014), and Arnosti, Beck and Milgrom (2016).

the membership status designed for his true vertical type is nondecreasing in the agent's true vertical type. This property guarantees that if low types find it optimal to participate, so do higher ones. Third, when all agents bid truthfully in all periods, the intertemporal match quality that each agent expects at the time of joining is nondecreasing in the selected status. This last property guarantees that agents find it optimal to select the status designed for their vertical type.

In a matching environment, the above monotonicities should not be taken for granted. In fact, contrary to standard screening problems, each agent plays the double role of a buyer and of an input provider for the matches the platform sells to the other side. Furthermore, the private information that each agent receives after joining the platform is multidimensional, and although certain dimensions may contribute to higher match quality, others may contribute to a lower one.[10] Lastly, agents may dislike interacting with certain other agents (i.e., experience a payoff below their outside option) and, for such interactions, a larger vertical type or a higher status may imply a larger loss.[11]

We show that, under certain conditions, the matching auctions we propose include auctions that maximize the platform's profits, as well as auctions that maximize welfare, or a convex combination between the two, *over all possible mechanisms* (Theorems 2 and 3). We then use the results to shed light on distortions due to market power (Theorem 4). The last few years have witnessed growing concerns about possible inefficiencies in matching markets dominated by a few platforms with strong market power, spurring an active debate on how to regulate such platforms.[12] We show that, in markets where all agents assign a nonnegative value to all interactions and none of the capacity constraints binds, in each period, each match

[10] In other words, a higher vertical type, or a higher status, may bring higher match quality when paired with some dimensions but not with others.

[11] Technically, the agents' payoffs at the time of joining need not satisfy the familiar increasing-difference property.

[12] See, e.g., "Online Platforms: Nostrums for Rostrums" and "Regulating Technology Companies: Taming the Beasts," (*The Economist*, May 28, 2016).

induced by a monopolistic profit-maximizing platform is welfare maximizing (however, many welfare-maximizing matches need not be implemented under profit maximization). In the presence of capacity constraints, instead, some of the matches induced by a profit-maximizing platform may be socially inefficient. However, when the only binding capacity constraint is the aggregate one (i.e., the platform's), the total number of matches induced in each period by a monopolistic profit-maximizing platform is always inefficiently low. Interestingly, these conclusions do not extend to markets in which individual capacity constraints may be binding. In this case, a platform enjoying strong market power may induce an inefficiently large number of matches, for any number of periods. The same is true in markets in which certain agents dislike certain interactions (that is, derive a payoff lower than their outside option from certain matches). In such markets, simple policy interventions based on tax reductions or match subsidies may prove counterproductive and yield lower welfare.

Outline. The rest of the paper is organized as follows. The remainder of the introduction briefly discusses the most pertinent literature. Section 2 describes the environment. Section 3 introduces the dynamic matching auctions and derives their equilibrium properties. Section 4 identifies auctions maximizing the platform's profits, total welfare, or a convex combination of the two, and uses the results to characterize the inefficiencies due to market power. Section 5 concludes with a brief discussion of the robustness of the insights and of lines of future research. All proofs are in the Appendix at the end of the document.

1.1 Related Literature

Markets where agents purchase access to other agents are the focus of a vast literature on two-sided markets pioneered by Caillaud and Jullien (2003), Rochet and Tirole (2003, 2006) and Armstrong (2006)—see Belleflamme and Peitz (2017) for a recent overview. This literature assumes that all agents on board interact with all other agents from the opposite side (i.e., it restricts attention to a single network, or to mutually exclusive networks). Most importantly, it focuses primarily on static environments (see Cabral (2011) for a dynamic model with complete information, and Jullien and Pavan (2019) for a dynamic extension with asymmetric information). In these works, match values are constant over time and there is no customization.

Damiano and Li (2007) and Johnson (2013) consider the mechanism-design problem of a profit-maximizing platform facing agents with private information on their vertical types. Hoppe, Moldovanu and Ozdenoren (2011) quantify the benefit of a coarse matching scheme in terms of matching surplus, revenue, and welfare. In these papers, in equilibrium, matching is one-to-one. In contrast, Board (2009) considers the problem of a profit-maximizing platform allocating agents to mutually exclusive groups (e.g., teams), whereas Gomes and Pavan (2016, 2019) study a problem similar to the one in Board (2009) but where agents differ both in their preferences and in their attractiveness, and where matching is non-partitional. Hoppe, Moldovanu and Sela (2009) show that assortative matching can arise in a Bayesian equilibrium of a bilateral (costly) signaling game. Dizdar and Moldovanu (2016) study a model where agents are characterized by private, multi-dimensional, attributes which jointly determine the surplus from a match, and give a possible explanation for the prevalence of rules that divide surplus in a fixed proportion. Matching in all of these papers is static. Fershtman and Pavan (2017) study a simple model of dynamic matching in which valuations change only once, after the first interaction, a single match is formed in each period, and the new private information agents receive in each period is uni-dimensional.

Contrary to these works, in the present paper, agents arrive stochastically over time, experience multiple shocks to their match values, which are privately observed, and are repeatedly re-matched.

The paper is also related to the literature on scoring auctions. In this literature, scores are used by the procurers to aggregate the various dimensions of the suppliers' offers (price, product design, delivery time). See, for example, Che (1993) and Asker and Cantillon (2008). Our matching auctions share with this literature the idea that the desired allocations can be induced through an appropriate design of the bilateral scores. Contrary to this literature, however, the scores in our auctions aggregate the preferences of different agents from different sides of the market, instead of the various dimensions of each seller's own offer.

Another related literature is the one studying auctions for sponsored links by search engines and for contextualized ads by ad exchanges. For example, Varian (2007), Edelman, Ostrovsky and Schwarz (2007), and Gomes and Sweeney (2014) study the properties of the GSP auction used by online search engines to allocate ads, whereas Mansour, Muthukrishnan, and Nisan (2012), Harris and Varian (2014), and Arnosti, Beck and Milgrom (2016) study auctions used by ad exchanges to match advertisers with content providers.[13] Our matching auctions are relevant also for these markets, modulo the fact that, in online search, searchers typically do not pay for the matches (or, more precisely, the "currency" used for the services they receive is the release of their privacy). In ad exchange auctions, instead, it is becoming customary for content providers to specify dynamic reservation prices for different types of ads, which is a form of dynamic bidding. In these auctions, both sides thus repeatedly bid for all matches, bids are aggregated into bilateral compatibility scores, and participating bidders are charged upfront fees at the time of joining, as in our model. The key contribution of our paper vis-a-vis this literature is the characterization of simple auctions maximizing the platform's profits when agents arrive over time and experience frequent changes to their preferences, as well as the discussion of the inefficiencies due to market power.

Most of the recent literature on centralized dynamic matching focuses on markets without transfers, in which matching is irreversible and dynamics originate entirely in the arrival and

[13] See also Athey and Ellison (2011), Börgers, Cox, Pesendorfer and Petricek (2013), and Gomes (2014).

departure of agents at and from the market. In the context of kidney exchange, for example, Ünver (2010) studies optimal mechanisms for exchanges minimizing total waiting costs.[14] Optimal dynamic matching is also the focus of Anderson, Ashlagi, Gamarnik, and Kanoria (2017), Akbarpour, Li, and Oveis Gharan (2020), and Baccara, Lee, and Yariv (2020). A central trade-off in these papers is between avoiding waiting costs and waiting for the market to thicken. The key differences with respect to the present paper are that, in this literature, agents are not re-matched (matching is irreversible), match values are constant over time, and payments are absent.

From a methodological standpoint, we draw from recent developments in the dynamic mechanism design literature. In particular, the conditions for incentive compatibility in the present paper build on results in Pavan, Segal, and Toikka (2014), adapted to our matching environment.[15] The key difference is that, in this paper, agents arrive stochastically over time, strategically choose the time at which they join the platform, and receive multi-dimensional private information in each period, after learning the identities of the other agents on board.

2 The Environment

Arrivals, Match Values, and Payoffs. A platform matches agents from two sides of a market, A and B. Agents arrive at the market stochastically over time. As usual, an agent's arrival can be interpreted as the agent becoming aware of the platform, or developing an interest for the platform's services. To simplify the exposition, we assume that, once an agent arrives, he stays forever. This assumption, however, is inconsequential for the results. The

[14] See Damiano and Lam (2005), Kurino (2009), and Doval (2015) for appropriate stability concepts for such environments.

[15] See also Baron and Besanko (1984), Besanko (1985), Courty and Li (2000), Board (2007), Eso and Szentes (2007), and Bergemann and Valimaki (2010) for some of the earlier contributions, and Gershkov and Moldovanu (2014), Börgers (2015), Bergemann and Pavan (2015), Pavan (2017) and Bergemann and Valimaki (2019) for overviews of this literature.

case where some of the agents depart from the market is formally equivalent to one where all agents stay forever but some agents, after a certain (possibly stochastic) number of periods, experience sufficiently negative payoffs from that moment onwards for interacting with all agents from the opposite side, making it optimal for the platform to exclude them from all future interactions. This is a special case of the environment we consider.

The time at which the agents arrive at the market is their own private information. Hence agents can strategically choose the time at which they join the platform. The market is ex-ante anonymous, that is, individual names bear no information to the platform about match values.[16] Each agent learns about the identities of the agents who are on-board already, from both sides of the market, only after joining the platform.

Time is discrete and indexed by $t = 0, ..., \infty$. For each t, and each side $k = A, B$, $n_t^k \in \mathbb{N}$ denotes the number of side-k agents who joined the platform prior to period t (i.e., in periods $0, ..., t-1$). We use natural numbers to denote the positions occupied by the agents in the platform, and assume that positions are occupied in the order of arrival. The set of side-k positions occupied at the beginning of period t is denoted by $N_t^k \equiv \{1, ..., n_t^k\}$, with $N_0^k = \emptyset$. The set of new positions occupied in period t is therefore given by $N_{t+1}^k \setminus N_t^k = \left\{n_t^k + 1, ..., n_{t+1}^k\right\}$. If multiple agents join the platform at the same time, the new positions are assigned randomly to the newly arrived agents according to a uniform distribution. Denote by $N_t^{AB} \equiv N_t^A \times N_t^B$ the collection of possible matches that can be formed in period t, given the positions that have been filled. As we explain below, not all such matches are, however, feasible, as both the platform and the agents may face capacity constraints—more below.

Agents arrive at the market according to the process $\mathcal{P}$. For example, the arrival of each agent could be governed by a Poisson process, with draws independent across agents. The

[16]This assumption is not essential to the results but simplifies the exposition by permitting us to describe the platform's mechanisms in an anonymous way (that is, by dropping the conditioning of the allocations and payments on individual names).

details of $\mathcal{P}$ do not play a role for our results and hence no specific structure on $\mathcal{P}$ is imposed.

To ease the exposition, hereafter we often refer to the generic agent occupying the i-th position on side k as the side-k i-th agent. Below we describe various features of the environment focusing on the i-th agent from side A, with the understanding that a similar description applies to any of the side-B agents.

The period-t flow payoff v_{ijt}^A that the side-A i-th agent derives from interacting with the side-B j-th agent depends on t, and on the identities of the two agents, but not on the time at which the two agents arrived at the market, and is given by

$$v_{ijt}^A = \theta_i^A \varepsilon_{ijt}^A. \tag{1}$$

The term θ_i^A denotes the average (unconditional) payoff that the side-A i-th agent derives from interacting with any of the side-B agents, before doing any profiling, i.e., before learning the identity of the side-B j-th agent. Agent i learns θ_i^A prior to joining the platform. In contrast, the term ε_{ijt}^A denotes the (time-varying) attractiveness of the side-B j-th agent in the eyes of the side-A i-th agent. Naturally, ε_{ijt}^A depends on the identities of the two individuals and, as such, is learned by the side-A i-th agent only after joining the platform. Each ε_{ijt}^A evolves over time reflecting the accumulation of information, or simply variations in the environment affecting the profitability of the match.

The representation in (1) permits us to interpret θ_i^A as the agent's "*vertical type*," that is, the overall importance that the side-A agent assigns to interacting with the side-B agents, and $\varepsilon_{it}^A \equiv (\varepsilon_{ijt}^A)_{j \in N_t^B}$ as the agent's period-t "*horizontal types*," that is, the agent's preferences over the specific agents occupying the various positions on side B in period t. Hereafter, we refer to $v_{it}^A \equiv (v_{ijt}^A)_{j \in N_t^B}$ as the agent's period-t "*match values*" which combine the vertical and the horizontal types.

For any $t \geq 1$, and any pair of positions $(i,j) \in \mathbb{N}^2$, let $X_{ijt} \equiv \{0,1\}$, with $x_{ijt} = 1$ in case the pair (i,j) is matched in period t, and with $x_{ijt} = 0$ otherwise. The ex-post payoff that the side-A i-th agent derives from joining the platform in period t is given by

$$U_{it}^A = \sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{j \in N_s^B} x_{ijs} v_{ijs}^A - \sum_{s=t+1}^{\infty} \delta^{s-t} p_{is}^A - f_{it}^A, \tag{2}$$

where $\delta \in (0,1]$ is the common discount factor, f_{it}^A is the fee paid by the agent to the platform at the time of joining, and $(p_{is}^A)_{s=t+1}^{\infty}$ are the payments made to the platform in all subsequent periods. Note that, although the agent's first match does not occur prior to period $t+1$, the agent's first payment to the platform, f_{it}^A, occurs at the time the agent joins the platform. Payments in each period can be negative, reflecting the possibility that (a) agents may dislike interacting with other agents and ask to be compensated, or (b) the platform's cross-subsidization of the two sides.

The platform's ex-post payoff (its profit) is given by

$$U_0 = \sum_{t=0}^{\infty} \delta^t \left(\sum_{i \in N_t^A} p_{it}^A + \sum_{j \in N_t^B} p_{jt}^B + \sum_{i \in N_{t+1}^A \setminus N_t^A} f_{it}^A + \sum_{j \in N_{t+1}^B \setminus N_t^B} f_{jt}^B \right) - \sum_{t=1}^{\infty} \delta^t \left(\sum_{(i,j) \in N_t^{AB}} c_{ijt} x_{ijt} \right).$$

The platform's payoff is thus equal to the discounted sum of the payments collected from the two sides of the market, net of possible costs of implementing the matches, where $c_{ijt} \in \mathbb{R}$ is the period-t "cost" of matching the pair $(i,j) \in N_t^{AB}$. We allow these costs to take on negative values so as to capture the possibility that the platform may derive positive benefit from certain matches. These costs may also incorporate auxiliary services the platform provides to the agents, over and above matching the agents from the two sides.

All agents, as well as the platform, are expected-utility maximizers.

Evolution of Match Values. Each θ_i^A is drawn from an absolutely continuous cumulative distribution function G^A with density g^A such that $g^A(\theta_i^A) > 0$ if and only if $\theta_i^A \in \Theta^A = [\underline{\theta}^A, \overline{\theta}^A]$, with $\underline{\theta}^A > 0$. Vertical types are drawn independently across agents, independently of the time at which the agents arrive at the market, and independently of the horizontal types. The horizontal types, $\varepsilon \equiv (\varepsilon_{ijt}^k)_{i,j,t \in \mathbb{N}}^{k=A,B}$, instead, are possibly correlated across agents and over time. Importantly, although we restrict the vertical types to be nonnegative, we allow the horizontal types to take on negative values, reflecting the possibility that certain agents may dislike certain interactions; that is, certain agents may derive a payoff lower than their outside option, which is assumed to be equal to zero, from interacting with certain agents.[17] Although not essential to the results, we find it convenient to think of each ε_{ijt}^k as being drawn from a set $\mathcal{E}_{ijt}^k \subseteq \mathbb{R}$ which either coincides with the entire real line, or with a compact and connected subset of it. To guarantee that each agent's expected payoff is well defined at all histories, we assume that, for each, $i \in \mathbb{N}$, and $t \in \mathbb{N}$, and any feasible matching rule χ specifying who is matched with whom at each period (formally described below), $\mathbb{E}\left[\sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{j \in N_s^B} \chi_{ijs} |\varepsilon_{ijt}^A|\right] \leq E_i^A$, for some constant $E_i^A > 0$. A similar condition applies to side B.

Both the vertical and the horizontal types are the agents' private information.

Capacity Constraints. At each period t, each side-k agent who has joined prior to period t can be matched to at most m_t^k of the agents from the opposite side, among those who have also joined prior to period t. Such constraints admit as special cases the case of *one-to-one* matching ($m_t^k = 1$ for all t, $k = A, B$) and the case of *many-to-many* matching with no binding individual constraints ($m_t^k = \infty$ for all t, $k = A, B$).

In addition to such individual capacity constraints, in each period t, the platform may face an aggregate capacity constraint imposing that no more than M_t matches be formed. In

[17] Under the specification in (1), allowing the vertical types to also take on negative values would introduce confusion, given that the horizontal types ε_{ijt}^k are already allowed to take on negative values.

each period, the platform can discontinue any of the previously formed matches and create new ones. The total number of existing matches, however, cannot exceed M_t in each period. For example, the number of ads and articles jointly displayed on a given outlet may be naturally limited by the outlet's physical capacity, which may evolve over time. More generally, these limits may reflect space, time, and resource constraints, but also capture certain non-separabilities, i.e., decreasing returns, in payoffs (in the case of individual constraints). For simplicity, we assume the individual capacity limits are the same for all agents on a given side. All key results, however, extend to settings in which such limits are individual-specific and are the agents' private information.[18]

A period-t *matching allocation* $x_t \in \prod_{(i,j)\in\mathbb{N}^2} X_{ijt}$ is *feasible* if (a) $x_{ijt} = 1$ only if $(i,j) \in N_t^{AB}$, and (b) none of the capacity constraints, individual or aggregate, are violated.[19] We denote by X_t the set of feasible period-t allocations, and by $X \equiv \prod_{t=1}^{\infty} X_t$ the set of sequences of feasible allocations.

Matching Mechanisms. We are interested in mechanisms that take the form of simple auctions in which agents, after paying a membership fee, repeatedly bid for the interaction with other agents from the opposite side. We describe such auctions in the next section. In Section 4, we then show that a certain version of such auctions maximizes the platform's profits, whereas another maximizes total welfare, across all possible mechanisms. To establish these results, below we define arbitrary matching mechanisms.

An arbitrary matching mechanism comprises a payment, a matching, and a disclosure rule, with the latter specifying the information revealed to the agents over time (for example, such a rule may disclose the identities of the agents involved in the matches implemented in the

[18] When the capacity limits are the agents' private information, and possibly evolve over time, the agents must be asked to report them at the beginning of each period. In this case, all the qualitative predictions continue to hold, provided that agents learn their capacities after joining the platform and selecting their status.

[19] The constraint $x_{ijt} = 1$ only if $(i,j) \in N_t^{AB}$ simply means that the match between the i-th position from side A and the j-th position from side B is formed only if the two positions have been filled.

past, the payments made by the agents, or some coarser statistics of such information). Under the assumptions discussed below, the optimal matching mechanisms will turn out to be fully transparent.

Formally, a *matching mechanism* $\Gamma \equiv (\mathcal{M}, \mathcal{S}, \chi, \psi, \rho)$ consists of: (i) a collection of message sets $\mathcal{M} \equiv (\mathcal{M}_t)_{t=0}^{\infty}$, where, for each t, $\mathcal{M}_t \equiv \mathcal{M}_t^A \times \mathcal{M}_t^B$, with $\mathcal{M}_t^k = \prod_{i \in \mathbb{N}} \mathcal{M}_{it}^k$, and with each $\mathcal{M}_{it}^k$ denoting the set of messages for the side-k i-th position; (ii) a collection of sets of signals $\mathcal{S} \equiv (\mathcal{S}_t)_{t=0}^{\infty}$ that the platform may disclose to the agents, where, for each t, $\mathcal{S}_t \equiv \mathcal{S}_t^A \times \mathcal{S}_t^B$, with $\mathcal{S}_t^k = \prod_{i \in \mathbb{N}} \mathcal{S}_{it}^k$, and with each $\mathcal{S}_{it}^k$ denoting the set of signals the platform may disclose to the agent occupying the side-k i-th position; (iii) a *matching rule* $\chi \equiv (\chi_t)_{t=1}^{\infty}$ describing, for each $t \geq 1$, the matches $x_t \in X_t$ implemented given the history of received messages, with $\mathcal{M}^t = \prod_{s=0}^{t} \mathcal{M}_s$ and $\chi_t : \mathcal{M}^t \to X_t$; (iv) a *payment rule* $\psi \equiv (\psi_t)_{t=0}^{\infty}$ describing, for each $t \geq 0$, the payments (positive or negative) asked to the agents who joined the mechanism prior to, or in, period t; and (v) a *disclosure policy* $\rho \equiv (\rho_t)_{t=0}^{\infty}$ specifying the information disclosed to the agents over time, with $\rho_t : \mathcal{M}^t \to \Delta(\mathcal{S}_t)$. Each ρ_t must reveal to each agent his own matches and payments. It may also reveal additional information, but it cannot conceal the matches the individual is involved in, or the payments from/to the individual. We assume that the sets $\mathcal{M}_{it}^k$ and $\mathcal{S}_{it}^k$ contain elements $\bar{m}_{it}^k \in \mathcal{M}_{it}^k$ and $\bar{s}_{it}^k \in \mathcal{S}_{it}^k$ such that, when the i-th position is not occupied in period t, the default message sent by the position is $\bar{m}_{it}^k$ and the default information disclosed to the position is $\bar{s}_{it}^k$.

Note that, because the matches implemented in each period, the payments collected from the agents, and the information disclosed to the agents, depend on the entire history of messages sent in current and past periods, and because the history reveals when agents arrived, the allocations implemented under such mechanisms may naturally condition on the timing of the agents' arrivals.

Once an agent arrives at the market, he chooses when to join the mechanism, after observing his vertical type but before observing the number and identities of the agents who joined in previous periods. Upon joining the mechanism, at each subsequent period, after learning his own position, the identities of the other agents on-board, and his horizontal types, the agent sends a message from the set $\mathcal{M}_{it}^{k}$. In the auctions we introduce in the next section, such messages correspond to the selection of a membership status (at the period at which the agent joins) and a collection of bids, one for each occupied position on the opposite side (in the subsequent periods). A matching mechanism Γ is *feasible* if, for any sequence of messages $m \in \mathcal{M}$, the implemented allocations are feasible.[20]

Solution concept. We use perfect Bayesian equilibrium (PBE) as our solution concept.

2.1 Discussion of the model.

Private information before joining. The assumption that agents possess some private information prior to joining the platform, and that such information is independent among the agents, guarantees that the matching dynamics under profit maximization differ from their counterparts under welfare maximization. In environments in which the platform maximizes welfare, the information the agents possess at the time they join the platform plays no role. When, instead, the platform maximizes profits, the private information the agents possess at the time they join plays a fundamental role. If the agents possess no private information, the platform can extract the entire surplus. If the agents' private information is multidimensional, the optimal mechanisms can be more complicated to characterize (for example, it may be impossible to determine which participation constraints bind). In this respect, the assumption that the agents do not know the identities of those agents already on board is made only to guarantee that their private information at the moment of joining is unidimensional, and can

[20]Note that the matching rule χ is deterministic. This is because, in this environment, the platform never gains from inducing random matches.

easily be dispensed with if one assumes that the match-specific values are learned only after joining the platform.

Non-separability in payoffs. The assumption that the value each agent derives from interacting with any other agent is invariant to the composition of the two agents' matching set favors a certain simplicity in the description of the scoring rules. Note, however, that the individual capacity constraints already capture non-separabilities in payoffs. The analysis can be extended to accommodate for richer non-separabilities, albeit at the cost of an increase in the complexity of the scores.

Assortative matching. The assumption that each agent's utility is invariant to other agents' vertical types guarantees that the payoffs feature private values. Allowing for interdependent values complicates the structure of the auctions; as is well known, with interdependent values, payments linked to the (virtual) externalities agents impose on others are not guaranteed to induce truthful bidding. The model can nevertheless accommodate for assortative matching. Because ε are allowed to be correlated across agents, heterogeneity along a vertical dimension can be captured not only through θ but also through ε (in particular, an attractive agent j from side B can always be captured as someone delivering a high ε_{ij}^{A} to most if not all agents i from side A).

Multiplicative structure of match values. The multiplicative structure of the match values v permits us to accommodate for a rich multi-dimensionality of the horizontal types ε (an alternative convenient structure is one in which the match values are additively separable in θ and ε).[21] If each agent's horizontal type in each period is uni-dimensional (that is,

[21]In particular, a multiplicatively-separable structure (alternatively, an additively-separable one) implies that the flow "virtual surplus" in each period (which controls for the cost to the platform of leaving rents to the agents to induce them to reveal their initial private information) is multiplicatively-separable (alternatively, additively-separable) in the agents' initial private information, θ. Such a separability in turn permits us to use VCG-type of payments to induce truthful bidding not only in the case of welfare maximization but also in the case of profit maximization. Without such a separability, the optimal mechanisms may be more complicated, due to the multi-dimensionality of the private information ε the agents receive in each period, after joining the platform.

$\varepsilon_{ijt}^{A} = \varepsilon_{it}^{A}$ for all j), then one can accommodate for a general function $v_{ijt}^{A}\left(\theta_i^A, \left(\varepsilon_i^A\right)^t\right)$ of the vertical and horizontal types. In this case, however, to capture heterogeneity in the match values, one would need the match value functions to be identity-specific (that is, each v_{ijt}^{A} to depend on the identity j of agent i's period-t partner), thus losing the convenience of the anonymity of the matching mechanisms assumed throughout.[22]

Endogenous capacity constraints. The capacity constraints capture sharp non-linearities in the costs of the implemented matches. We assume for simplicity that such constraints are exogenous. In practice, the platform could control the number of slots it can accommodate in each period. This would amount to a smoother cost of the implemented matches. The analysis below already accommodates for match-specific costs in addition to the exogenous capacity constraints, but assumes such costs are linear in the implemented matches. One could envision richer cost specifications. However, such costs could interfere with the possibility of writing the platform's flow payoff in terms of the sum of the bilateral scores of the implemented matches, thus making the structure of the optimal mechanisms more complicated.

We see the combination of the above assumptions as a convenient way to retain tractability, while permitting us to capture key trade-offs in the design of matching auctions.

3 Matching Auctions

We now introduce a class of matching mechanisms in which (a) upon joining the platform, agents are invited to select a membership status; (b) in each subsequent period, after learning the identities of the agents occupying the various positions, each agent is asked to submit a vector of bids, one for each position on the opposite side; (c) bids are aggregated into a collection of *bilateral scores*, one for each possible match, with each score depending only on

[22] In this case, the flow payments implementing the desired allocations would have a structure similar to the one in Pavan, Segal, and Toikka (2014).

the pair of agents' reciprocal current bids and their membership status; (d) in each period, the matches maximizing the sum of the bilateral scores are implemented, subject to individual and aggregate capacity constraints. Formally, matching auctions are defined as follows.

Definition 1 (Matching Auctions). In a *matching auction*:

- Each agent, upon joining the platform, is asked to select a membership status. This status determines the weight the platform assigns to the agent's bids in the subsequent auctions, with a higher status corresponding to a higher weight. Because different statuses are meant for agents with different vertical types (recall that agents with a higher vertical type are agents who value interacting with agents from the opposite side more), we find it convenient to label the statuses directly with the vertical types they are meant for. We thus assume there exist functions $\beta \equiv \left(\beta^k(\cdot)\right)_{k=A,B}$, with $\beta^k : \Theta^k \to \mathbb{R}_{++}$, $k = A, B$, such that each agent from side k selecting the status meant for type θ^k is assigned the weight $\beta^k(\theta^k)$.

- Each agent from side k joining the platform in period t and selecting the status meant for type θ^k is asked to make an upfront payment equal to $f_t^k = \bar{\psi}_t^k(\theta^k)$, where $\bar{\psi}_t^k \equiv \left(\bar{\psi}_t^k(\cdot)\right)_{k=A,B}$ are non-decreasing functions. The payment f_t^k should be interpreted as an "entry fee" which naturally depends on the selected status.[23]

- In each period following the one at which the agent joined the platform, after learning the identities of those other agents on board and the positions they occupy, the agent is asked to submit a vector of bids, one for each filled position on the opposite side of the

[23]The reader may wonder why the fee f_t^k does not depend on the information revealed by those agents (from each side) who already joined the platform. As we show below, conditioning on such information does not help the platform attain higher profits, or higher welfare.

market. Each pair of filled positions $(i,j) \in N_t^{AB}$ is then assigned a "*score*"

$$S_{ijt} \equiv \beta_i^A b_{ijt}^A + \beta_j^B b_{ijt}^B - c_{ijt}, \tag{3}$$

where $\beta_i^A = \beta^A(\theta_i^A)$ and $\beta_j^B = \beta^B(\theta_j^B)$ represent the statuses of the side-A i-th agent and of the side-B j-th agent, respectively. The auction then implements the matches maximizing the sum of the bilateral scores subject to individual and aggregate capacity constraints, with ties broken arbitrarily.[24]

- All agents who joined prior to period t and who are unmatched in period t pay nothing in period t. Agents who joined prior to period t and are matched to some of the agents from the opposite side in period t make (or receive) payments in period t according to the rule ψ^β described as follows. Let

$$S_t \equiv \sum_{(i,j)\in N_t^{AB}} S_{ijt} x_{ijt} \tag{4}$$

denote the sum of the implemented scores. Similarly, let $S_t^{-i,k}$ denote the sum of the implemented scores under the same matching rule but in a fictitious environment in which all bilateral scores involving agent i from side k are identically equal to zero. Each agent $i \in N_t^A$ from side A who, in period t, is matched to some of the side-B agents is asked to make a payment equal to

$$\psi_{it}^A = \sum_{j\in N_t^B} b_{ijt}^A x_{ijt} - \frac{S_t - S_t^{-i,A}}{\beta_i^A}. \tag{5}$$

The payments for the side-B agents are defined in a similar manner.

[24] The specific tie-breaking rule plays no role in the analysis. For concreteness, assume ties are broken at random, from a uniform distribution, independently over time.

- In each period, the platform discloses to each agent on board (i.e., who joined prior to that period) all information collected in previous periods (i.e., the status selected, the bids submitted, and the payments made by any of the agents who joined in previous periods).

An agent's status thus determines the importance the platform assigns to the agent's bids, relative to those of others. Suppose that, in period t, the agent occupying the side-A i-th position submits a positive bid for the side-B agent occupying the j-th position, whereas the latter agent submits a negative bid for the same match, thus asking to be compensated. For given bids (b^A_{ijt}, b^B_{ijt}), a higher status of the side-A i-th agent implies a higher score for the match (i, j), tilting the allocation in favor of the side-A i-th agent. Symmetrically, a higher status for the side-B j-th agent reduces the score for the match (i, j), tilting the allocation in favor of the side-B j-th agent. A higher status thus grants an agent preferential treatment, both with respect to the competition the agent faces from other agents from his own side (when the capacity constraints bind) and with respect to the competition the agent faces from agents from the opposite side, shifting the matching to his benefit.

As far as the payments are concerned, the agent occupying the side-A i-th position is asked to make a payment equal to the total value the agent derives from all matches implemented in period t in which the agent is involved (this value is naturally computed using the agent's own bids), net of a discount proportional to the agent's contribution to the total score, with a coefficient of proportionality inversely related to the agent's membership status. As in standard VCG auctions, these payments reflect the flow externalities the agents impose on other agents from both sides of the market, as well as the costs (or benefits) they impose on the platform. Such externalities may be positive or negative, and therefore the payments may be positive or negative. For example, if an agent is valued highly by the agents he is matched to, he may receive a positive transfer from the platform, reflecting cross-subsidization. Contrary to standard VCG auctions, however, such externalities are not computed with respect to a standard welfare aggregator; they are calculated with respect to the aggregate score they induce. In particular, due to the different statuses the agents may hold, in contrast to a standard VCG mechanism, the implemented matches need not be efficient, that is, need not maximize the sum of the agents' and the platform's payoffs.[25] Furthermore, multiple agents (from both sides of the market) may be charged for the same externality they impose on others. The following example illustrates:

Example 1 (**Many-to-Many Matching with Aggregate Capacity Constraints**)**.** Suppose that, in period t, the only relevant capacity constraint is the aggregate one, i.e., the one pertaining to the platform. Consider an agent from side A in the i-th position (a similar description applies to the side-B agents). The agent's period-t payment to the platform is equal to

$$\psi_{it}^{A} = \frac{1}{\beta_i^A}\left(\mathbf{B}_{it}^{A}(K) + \sum_{j \in N_t^B} (c_{ijt} - \beta_j^B b_{ijt}^B) x_{ijt}\right),$$

where $K \leq n_t^B$ is the number of agents from side B with whom agent i is matched, and $\mathbf{B}_{it}^{A}(K)$ is the sum of the K highest nonnegative scores among the pairs that are unmatched in period t and that do not include agent i. In the special case in which the platform's costs are identically equal to zero for all matches and $\beta_j^k = 1$ for all agents, the above payments reduce to

$$\sum_{l \in N_t^A \setminus \{i\}} \sum_{j \in N_t^B} \left(b_{ljt}^A + b_{ljt}^B\right) x_{ljt}^{-i,A} - \sum_{j \in N_t^B} b_{ijt}^B x_{ljt},$$

[25] As shown in Section 4, β functions different from the identity one may permit the platform to implement allocations that maximize objectives other than welfare maximization (such as profit maximization or a combination of profit and welfare).

where

$$x_{ljt}^{-i,A} \in \arg\max_{\tilde{x}_t \in X_t} \left\{ \sum_{(l,j)\in N_t^{AB}} (b_{ljt}^A + b_{ljt}^B)\tilde{x}_{lj} : \sum_{(l,j)\in N_t^{AB}} \tilde{x}_{lj} \leq M_t,\ x_{ij} = 0 \text{ all } j \in N_t^B \right\}$$

are the matches that maximize the sum of bilateral scores in the absence of agent i (in this case, such a rule also maximizes total surplus, in the absence of agent i). □

3.1 Immediate Participation and Truthful Bidding

Definition 2 (**Straight-forward equilibrium**). A strategy profile σ for the above matching auctions is *straight-forward* if (a) each agent joins the platform as soon as he arrives at the market and never leaves thereafter, (b) upon joining, the status that each agent selects is the one designed for his vertical type, and (c) at each period following the one at which the agent joined, the agent submits bids that coincide with his true match values for all agents from the other side who are on board (i.e., for each $(i,j) \in N_t^{AB}$, $b_{ijt}^k = v_{ijt}^k$, $k = A, B$). A *straight-forward equilibrium* is a PBE in which the agents' strategies are straight-forward.

For any $\beta \equiv \left(\beta^k(\cdot)\right)_{k=A,B}$, let χ^β be the matching rule that obtains under a straight-forward equilibrium of an auction with weights β. Then, for any t and any $k = A, B$, let[26]

$$D_t^k(\theta^k;\beta) \equiv \begin{cases} \mathbb{E}\left[\sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{j\in N_s^B} \varepsilon_{ijs}^A \chi_{ijs}^\beta \mid \theta^k\right] & \text{if } k = A \\ \mathbb{E}\left[\sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{i\in N_s^A} \varepsilon_{ijs}^B \chi_{ijs}^\beta \mid \theta^k\right] & \text{if } k = B \end{cases} \tag{6}$$

denote the "quality" of the matches expected under such a straight-forward equilibrium by any side-k agent with vertical type θ^k when joining the platform in period t, before learning the identities of the agents already on board. Note that the expectation is with respect to the

[26]The reason why we distinguish between the case in which $k = A$ and the one in which $k = B$ is that the order in the subscripts of the allocations χ_{ijs}^β, as well as the order in the subscripts in the horizontal types ε_{ijs}^k, is not permutable. The first index always refers to side A, whereas the second to side B.

agent's own position in the platform, the identities and arrival times (and hence the positions) of all other agents on board, the arrival of new agents in future periods, all other agents' vertical types, and the evolution of all agents' horizontal types. Finally, let $\psi^{\beta} \equiv (\psi_{ls}^{k,\beta}(\cdot))_{l,s\in\mathbb{N}}^{k=A,B}$ be the payment rule defined in (5) when the weights are given by β.

Theorem 1. *A matching auction with weights $\beta \equiv \left(\beta^k(\cdot)\right)_{k=A,B}$ admits a straight-forward equilibrium if (a) the functions $\beta^k(\cdot)$ are non-decreasing, $k = A, B$, and (b) there exist non-negative scalars $(Q_t^k)_{t=0,\ldots,\infty}^{k=A,B}$ such that, for any $t \geq 0$, $\theta^k \in \Theta^k$, and $k = A, B$, the price the agent pays for each status θ^k is given by*

$$\bar{\psi}_t^{k,\beta}(\theta^k) = \theta^k D_t^k(\theta^k;\beta) - \int_{\underline{\theta}^k}^{\theta^k} D_t^k(y;\beta)dy - \mathbb{E}\left[\sum_{s=t+1}^{\infty} \delta^{s-t}\psi_{ls}^{k,\beta} \mid \theta^k\right] - Q_t^k \tag{7}$$

and the dynamics of the match quality satisfies

$$\int_{\underline{\theta}^k}^{\theta^k} D_t^k(y;\beta)dy + Q_t^k \geq \delta\left[\int_{\underline{\theta}^k}^{\theta^k} D_{t+1}^k(y;\beta)dy + Q_{t+1}^k\right] \geq 0. \tag{8}$$

The formal proof is in the Appendix. Here we illustrate the key ideas heuristically. The payments in (5) are such that, at each period $t \geq 1$, any agent on board who expects all other agents on board to bid truthfully has incentives to do the same. To see this, observe that, under such payments, regardless of the membership status selected by any of the agents (both those already on board and those joining in subsequent periods), the bids that maximize each agent's continuation payoff (net of the payments) are invariant to the bids the same agent submitted in previous periods and those he will submit in subsequent periods. Furthermore, under such payments, the agent's period-t flow net payoff is proportional to his flow marginal contribution to the *total score* (with the latter defined as in (4), and with the coefficient of proportionality given by $1/\beta_i^k$). Because, given the received bids, the matches implemented maximize the flow total score at all histories (including those off the equilibrium path), agents have incentives to remain on board and bid truthfully at all histories, irrespective of the beliefs they may have about past and current types of other agents, the status selected by other agents, and the bids submitted by the other agents in the current period and in the past. In other words, remaining on board and bidding truthfully is periodic ex-post optimal for each agent *after he joins the platform*, at any history.[27,28]

The key difficulty is in showing that all agents find it optimal to join the platform immediately upon arrival and select the membership status designed for their true vertical type. In the Appendix, we show that this follows from the following properties. First, we show that, when all other agents follow straight-forward strategies, the match quality

$$\sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{j \in N_s^A} \varepsilon_{ijs}^A x_{ijs}$$

expected by each agent from side A of vertical type θ^A who arrives at the market in period t, joins the platform immediately upon arrival, selects the membership status designed for type $\hat{\theta}^A$, and then bids truthfully at all periods, is nondecreasing in the selected status $\hat{\theta}^A$ (a similar property applies to the side-B agents). Together with the fact that bidding truthfully at all subsequent periods is optimal for the agent irrespective of the selected status (as explained above), such monotonicity in turn implies that, when the price of status is given by (7), the agent prefers to select the status designed for his true type to any other status. This is true

[27] Periodic ex-post means optimal for an agent who expects other agents to follow straight-forward strategies, regardless of the agent's beliefs over the past and current types of any of the other agents on board. Straight-forward strategies need not be (weakly) dominant though because bidding truthfully need not be optimal when other agents' strategies condition on the (information that other agents may have about) agent i's own behavior and/or the allocations induced by agent i's behavior.

[28] By making the price each agent pays for the selected status depend on the status selected by the other agents (equivalently, by making f_i^k depend on θ_{-i}^k) the platform can guarantee that the mechanism is periodic ex-post IC and IR also in period $t = 0$ (meaning that each agent finds it optimal to participate in the mechanism and choose the status designed for his true vertical type, regardless of his beliefs over the other agents' vertical types).

irrespectively of the time at which the agent joins.

Next, we show that, when all other agents follow a straight-forward strategy, the payoff that a side-A agent of vertical type equal to θ^A obtains by joining the platform in period t, selecting the status θ^A designed for his true vertical type, and bidding truthfully in each subsequent period is equal to $\int_{\underline{\theta}^A}^{\theta^A} D_t^A(y;\beta)dy + Q_t^A$, and that the latter is non-decreasing in θ^A. Because the agent receives no information while waiting outside the platform, the inequality in (8) then implies that the agent prefers following a straight-forward strategy to waiting an arbitrary number of periods before joining the platform, selecting the status designed for his true type upon joining, and then bidding truthfully. Jointly, the above properties imply that, under the conditions in the theorem, the agent maximizes his payoff by following a straight-forward strategy.[29] □

The following example illustrates the workings of the matching auctions in a simple setting.

Example 2. Matching occurs only in period 1. To ease the notation, we therefore drop the time index from all the relevant functions. There are two agents on side A and one agent on side B, and this is common knowledge (that is, $N^A = \{1, 2\}$ and $N^B = \{1\}$). The platform

[29] In standard dynamic mechanism design, incentive compatibility is established by selecting the payments so that the payoffs in all periods satisfy an appropriate envelope formula, and by showing that the allocations satisfy an "integral monotonicity" condition in each period (cf. Pavan, Segal, and Toikka, 2014). As anticipated above, such a characterization holds for settings where the private information in all periods is unidimensional, and need not extend to settings where the agents' private information is multi-dimensional. Incentive compatibility is established here by using VCG-type of transfers in all periods (other than the one at which the agent joins) that link the payments to the externalities the agents impose on one another, weighted by their membership statuses. Incentive compatibility at the time of joining cannot be established with VCG-type of payments because the initial membership choices determine the relative weights the platform assigns to the agents' utility in the subsequent periods. In other words, the endogeneity of the weights introduces interdependencies in the agents' payoffs, which are known to conflict with the VCG arguments. Incentive compatibility at the time of joining is established by showing that the derivative of each agent's discounted expected payoff with respect to the selected status is non-decreasing, which (under the assumed multiplicative structure of the agents' flow payoffs) coincides with the requirement that an agent's expected match quality be nondecreasing in the agent's status. That the latter property holds in turn follows from the fact that the matches implemented in each period maximize the sum of the scores. Lastly, because arrivals are the agents' private information, incentive compatibility also requires that each agent's expected payoff from joining the platform be non-increasing in time, accounting for discounting. Whether such a property binds or not depends on the process governing the agents' arrivals.

incurs no costs from matching any of the agents (that is, $c_{i1} = 0$, $i = 1, 2$) and matching is one-to-one.[30] The vertical types of the two side-A agents are drawn from $\Theta^A = [1.5, 2.5]$, with G^A uniform over Θ^A. For the side-B agent, $\Theta^B = \{1\}$, meaning that the side-B agent's vertical type is commonly known. Likewise, the horizontal types of each of the two side-A agents are known and equal to $\varepsilon^A_{i1} = 1$, $i = 1, 2$. For the side-B agent, instead, ε^B_{11} is drawn uniformly from $\{2, 3\}$, and $\varepsilon^B_{21} = 5 - \varepsilon^B_{11}$.

The weights are given by $\beta^k(\theta^k) = 1 - \left[1 - G^k(\theta^k)\right] / g^k(\theta^k)\theta^k$, $k = A, B$, implying that $\beta^A(\theta^A) = \left(2\theta^A - \frac{5}{2}\right) / \theta^A$ and $\beta^B(\theta^B) = 1$. Using (5), we then have that each side-A agent's payment at the end of period 1 is equal to zero if the agent is not matched, and otherwise is equal to $\psi^A_1 = \left[-\beta^B_1 b^B_{11} + \beta^A_2 b^A_{21} + \beta^B_1 b^B_{21}\right] / \beta^A_1$. In a straightforward equilibrium, agent 1 from side A is matched only if the side-B agent is of horizontal type $\varepsilon^B_{11} = 3$, implying that, in case he is matched, his payment at the end of period 1 is uniquely determined by his own vertical type and the vertical type of the other side-A agent, and is equal to $\psi^A_1 = \theta^A_1 \left[2\theta^A_2 - \frac{7}{2}\right] / \left[2\theta^A_1 - \frac{5}{2}\right]$. Note that for $\theta^A_2 > 1.75$, ψ^A_1 is positive and decreasing in the agent's own status, θ^A_1, but increasing in the other side-A agent's status, θ^A_2, reflecting the fact that a higher θ^A_2 implies a higher externality imposed by agent 1 on agent 2. When $\theta^A_2 < 1.75$, ψ^A_1 is negative, i.e., the agent is compensated for the interaction. A similar payment applies to agent 2 from side A. Turning to the side-B agent, if she is not matched her payment at the end of period 1 is equal to zero. If, instead, she is matched, her payment is equal to $\psi^B_1 = b^B_{11} - \left(\beta^B_1 b^B_{11} + \beta^A_1 b^A_{11}\right) / \beta^B_1$. Using the fact that, under truthful bidding, the side-B agent is matched with one of the two side-A agents only if her horizontal type for that agent is 3, ψ^B_1 reduces to $\psi^B_1 = -\left[2\theta^A - \frac{5}{2}\right]$. As a result, each of the two side-A agents is matched with probability 1/2 and the match qualities are equal to $D^A_0(\theta^A; \beta) = \frac{1}{2}$ and $D^B_0(\theta^B; \beta) = 3$.

[30]In this example, that matching is one-to-one may either reflect an aggregate capacity constraint of $M = 1$, or individual capacity constraints.

Using (7), with $Q_t^k = 0$, $k = A, B$, we then have that the price for status for the side-A agents is equal to

$$\overline{\psi}_0^A(\theta^A) = \theta^A \frac{1}{2} - \int_{1.5}^{\theta^A} \frac{1}{2} dy - \mathbb{E}\left[\psi_1^A \mid \theta^A\right] = \frac{1}{2}\left(5\theta^A - 7.5\right) / \left(4\theta^A - 5\right),$$

whereas the price for status for the side-B agent is equal to $\overline{\psi}_0^B(\theta^B) = \overline{\psi}_0^B = 4.5$. Note that $\overline{\psi}_0^A(\theta^A)$ is increasing in θ^A.

Hence, each side-A agent's interim expected payoff from participating in the auctions is equal to

$$\frac{1}{2}\left[\theta^A - \left(6\theta^A - 7.5\right) / \left(4\theta^A - 5\right)\right],$$

which is positive and (linearly) increasing in θ^A, whereas, for the side-B agent, her interim expected payoff is equal to zero, because she possesses no private information prior to joining the platform. As we show in Section 4, the auction described above not only admits a straight-forward equilibrium, but is in fact profit-maximizing among all possible matching mechanisms. □

Comments. At this point, the reader may wonder whether the matching auctions described above are just a complicated description of standard direct-revelation mechanisms. The answer is no. The matches sustained under the straight-forward equilibria of matching auctions are selected by comparing bilateral scores. As a result, many matching allocations cannot be sustained under the proposed auctions, despite being implementable with other direct-revelation mechanisms. In the next section, however, we show that, under certain conditions, a specific version of the proposed auctions maximizes the platform's profits whereas another version maximizes welfare, over all possible mechanisms.

Due to the assumption that the processes governing the match values are exogenous, the payments in the straight-forward equilibria are history-independent. In the online Supplementary Material, we show how the results extend to certain environments in which the match values evolve endogenously over time, as a function of past matches (e.g., reflecting experimentation or habit formation).[31] In such environments, the payments in the corresponding matching auctions are history-dependent.

4 Profit- and Welfare-Maximizing Auctions

Hereafter, when we say that a feasible mechanism, paired with a specific equilibrium it induces, is *profit maximizing* we mean that the platform's profits under the equilibrium of the proposed mechanism are as high as under any equilibrium of any other feasible mechanism. A *welfare-maximizing* mechanism is defined in a similar way, with welfare replacing profits in the platform's objective.

Let $\beta^P \equiv (\beta^{k,P}(\cdot))_{k=A,B}$ be the weights given by

$$\beta^{k,P}(\theta^k) = 1 - \frac{1 - G^k(\theta^k)}{g^k(\theta^k)\theta^k}, \tag{9}$$

$k = A, B$. We then have the following result:

Theorem 2 (Profit-Maximizing Auctions). *Consider the matching auctions in which the weights are given by β^P and the access fees are given by (7) with $Q_t^k = 0$ all t, $k = A, B$. Suppose that the functions $\beta^{k,P}$ are strictly increasing and satisfy $\beta^{k,P}(\underline{\theta}^k) > 0$, $k = A, B$, and that, for any t, any $k = A, B$, (a) $D_t^k(\underline{\theta}^k; \beta^P) \geq 0$ and (b) for any $\theta^k \in \Theta^k$,*

$$\int_{\underline{\theta}^k}^{\theta^k} D_t^k(y; \beta^P)dy \geq \delta \left[\int_{\underline{\theta}^k}^{\theta^k} D_{t+1}^k(y; \beta^P)dy \right].$$

[31] Endogenously evolving match value can also capture a form of satiation in the agents' preferences. See the online Supplementary Material for the details.

The matching auctions in which the weights are given by β^P and the access fees are given by (7), with $Q_t^k = 0$ all t, $k = A, B$, along with the straight-forward equilibria they induce, are profit maximizing.[32]

The monotonicity of the weights β^P guarantees that the agents' "virtual" match values $\beta^{k,P}(\theta^k)v_{ijt}^k = \left(\theta^k - \left(1 - G^k(\theta^k)\right)/g^k(\theta^k)\right)\varepsilon_{ijt}^k$ respect the same rankings as the true match values, $v_{ijt}^k = \theta_i^k \varepsilon_{ijt}^k$. Conditions (a) and (b) in the theorem, instead, are joint conditions on the capacity constraints and the process governing the arrivals of the agents and the evolution of their match values over time. Note that Condition (a) is vacuously satisfied if horizontal types ε are nonnegative, that is, if no agent dislikes interacting with any other agents from the opposite side. Condition (b) requires that each agent's expected discounted match quality at the time of joining be non-increasing with time. The condition is satisfied, for example, when the disutility that an agent derives from interacting with certain agents from the opposite side does not decline too fast due to the arrival of other agents from either side of the market. It is also satisfied when the positive utility the agent expects from interacting with some of the agents already on board declines with time due to the competition the agent expects from other agents from his own side joining in subsequent periods.

The proof in the Appendix is in three steps. First, we show that, given any matching mechanism $\Gamma \equiv (\mathcal{M}, \mathcal{S}, \chi, \psi, \rho)$ and any *Bayes Nash equilibrium* (and hence any PBE) σ of the game induced by Γ, the interim expected payoff $\tilde{U}_t^k(\theta^k; \mathcal{I}_t^k)$ of any agent arriving in period t with vertical type θ^k and receiving information $\mathcal{I}_t^k$ about the number, the identity, the arrival date, and the past messages of all other agents arriving prior to, or at, period t, is given by

$$\tilde{U}_t^k(\theta^k; \mathcal{I}_t^k) = \tilde{U}_t^k(\underline{\theta}^k; \mathcal{I}_t) + \int_{\underline{\theta}^k}^{\theta^k} \tilde{D}_t^k(y; \mathcal{I}_t^k)dy. \tag{10}$$

[32] Note that the proposed auctions may admit multiple straight-forward equilibria, which differ in the agents' out-of-equilibrium beliefs. However, the platform's profits are the same under any such equilibrium.

Here

$$\tilde{D}_t^k(\theta^k;\mathcal{I}_t^k) \equiv \begin{cases} \mathbb{E}\left[\sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{h\in N_s^B} \varepsilon_{lhs}^A \tilde{\chi}_{lhs} \mid \theta^A, \mathcal{I}_t^A\right] & \text{if } k=A \\ \mathbb{E}\left[\sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{h\in N_s^A} \varepsilon_{hls}^k \tilde{\chi}_{lhs} \mid \theta^B, \mathcal{I}_t^B\right] & \text{if } k=B \end{cases}$$

denotes the match quality expected by an agent arriving in period t with a vertical type equal to θ^k who receives information $\mathcal{I}_t^k$, with $\tilde{\chi}$ denoting the matches induced in period t under σ.

Next, we use the above representation of the agents' equilibrium interim expected payoffs to show that, given any mechanism Γ and any BNE σ of Γ, the platform's profits are given by

$$\mathbb{E}\Bigg[\sum_{t=1}^{\infty}\delta^t \sum_{(i,j)\in N_t^{AB}} \left(\beta^{A,P}(\theta_i^A)\theta_i^A\varepsilon_{ijt}^A + \beta^{B,P}(\theta_j^B)\theta_j^B\varepsilon_{ijt}^B - c_{ijt}\right)\tilde{\chi}_{ijt} - \sum_{t=0}^{\infty}\sum_{k=A,B}\sum_{l\in N_{t+1}^k\setminus N_t^k} \delta^t \tilde{U}_t^k(\underline{\theta}^k;\mathcal{I}_t^k)\Bigg], \tag{11}$$

where the weights β^P are as in (9). Finally, we show that, under the straight-forward equilibria of the proposed auctions, (a) the induced state-contingent matches maximize the first component of the function in (11), and (b) for each t, $\tilde{U}_t^k(\underline{\theta}^k;\mathcal{I}_t^k) = 0$. Together, the above properties imply that, provided the agents find it optimal to join the auctions upon arrival, select the status designed for their true vertical type, and bid truthfully, the proposed auctions are profit maximizing. Conditions (a) and (b) in the theorem, along with the results in Theorem 1, guarantee that this is indeed the case.

We now turn to welfare maximization. Let $\beta^W \equiv (\beta^{k,W}(\cdot))_{k=A,B}$ be the weights given by $\beta^{k,W}(\theta^k) = 1$, all θ^k, $k = A, B$.

Theorem 3 (**Welfare-Maximizing Auctions**). *(i) Any matching auctions in which the weights are given by β^W and the participation fees satisfy the properties of Theorem 1, along with the straight-forward equilibria they induce, are welfare maximizing. (ii) Suppose that (a)*

$D_t^k(\underline{\theta}^k;\beta^W) \geq 0$ *for all* t*, all* $k = A, B$*, and (b) for any* $t \geq 0$*,* $\theta^k \in \Theta^k$*, and* $k = A, B$*,*

$$\int_{\underline{\theta}^k}^{\theta^k} D_t^k(y;\beta^W)dy \geq \delta \left[\int_{\underline{\theta}^k}^{\theta^k} D_{t+1}^k(y;\beta^W)dy \right].$$

The matching auctions in which the weights are given by β^W *and the participation fees are given by (7) with* $Q_t^k = 0$ *all* t*,* $k = A, B$*, admit straight-forward equilibria. Furthermore, the platform's profits under the straight-forward equilibria of these auctions are as high as under any equilibrium of any mechanism implementing welfare-maximizing matches at all periods and inducing the agents to join as soon as they arrive at the market.*

The result in part (i) follows from the fact that, when the weights are given by β^W, the matches sustained under the straight-forward equilibria of the auctions under consideration maximize welfare after each history. The conditions in part (ii) guarantee that, when the payments are as in (5), with weights β^W, and the participation fees are given by (7) with $Q_t^k = 0$ all t, $k = A, B$, each agent finds it optimal to join upon arrival. This follows from the fact that the payoff expected by each agent with vertical type θ^k when joining the platform in period t is given by

$$U_t^k(\theta^k;\beta^W) = \int_{\underline{\theta}^k}^{\theta^k} D_t^k(y;\beta^W)dy,$$

along with the fact that $D_t^k(\cdot;\beta^W)$ is non-decreasing. That the straight-forward equilibria of the proposed auctions maximize the platform's profits over all equilibria of all feasible mechanisms implementing welfare-maximizing matches and inducing the agents to join upon arrival then follows from arguments similar to those establishing the optimality of the matching auctions in Theorem 2 above. To see this, note that the platform's profits under any BNE (and hence under any PBE) of any feasible mechanism implementing the welfare-maximizing matches in each period and inducing the agents to join upon arrival are given by (11), with

$\tilde{\chi}$ representing the efficient assignment rule, and with $\tilde{U}_t^k(\underline{\theta}^k; \mathcal{I}_t^k) \geq 0$, for all t. Under the straight-forward equilibria of the proposed auctions, for each t, the payoff $\tilde{U}_t^k(\underline{\theta}^k; \mathcal{I}_t^k)$ expected by an agent arriving in period t with the lowest vertical type is exactly equal to zero. Hence, the platform's profits under the straight-forward equilibria of the proposed matching auctions are at least as high as under any equilibrium of any feasible mechanism inducing the agents to join upon arrival and implementing the efficient matches.

Remark 1. The above results are for the two polar cases of profit and welfare maximization. It should be clear that similar results apply to the case of a platform maximizing a combination of these two objectives. Formally, suppose that the platform maximizes a convex combination between profits and welfare, with the weight on profit equal to α and the one on welfare equal to $1 - \alpha$. Then let $\beta^\alpha \equiv (\beta^{k,\alpha}(\cdot))_{k=A,B}$ be such that $\beta^{k,\alpha}(\theta^k) = \alpha\beta^{k,P}(\theta^k) + 1 - \alpha$, for all θ^k, $k = A, B$. Suppose that $D_t^k(\underline{\theta}^k; \beta^\alpha) \geq 0$ for all t, all $k = A, B$, and for any $t \geq 0$, $\theta^k \in \Theta^k$, and $k = A, B$,

$$\int_{\underline{\theta}^k}^{\theta^k} D_t^k(y; \beta^\alpha) dy \geq \delta \left[\int_{\underline{\theta}^k}^{\theta^k} D_{t+1}^k(y; \beta^\alpha) dy \right].$$

The matching auctions in which the weights used to compute the scores are given by β^α and the participation fees are given by (7), with $Q_t^k = 0$ all t, $k = A, B$, admit straight-forward equilibria. Furthermore, under such equilibria, the combination of the platform's profits and welfare is as high as under any equilibrium of any other mechanism.

We conclude by discussing the distortions brought in by market power under profit maximization. Suppose the conditions in Theorems 2 and 3 hold and let $\chi^P = \left(\chi_t^P(\theta, \varepsilon)\right)_{t=1}^\infty$ and $\chi^W = \left(\chi_t^W(\theta, \varepsilon)\right)_{t=1}^\infty$ denote the profit- and the welfare-maximizing matches implemented under the straight-forward equilibria of the auctions of Theorems 2 and 3, respectively, in state (θ, ε), where (θ, ε) is a complete description of the vertical and horizontal types of all agents arriving to the platform over time. Because (θ, ε) is exogenous, we drop it from the arguments of χ^P and χ^W below to ease the notation. We then have the following result[33]:

Theorem 4 (Dynamic inefficiencies). *Suppose that all agents derive a nonnegative utility from interacting with all other agents from the opposite side.*

1. *If none of the capacity constraints bind, then at any period $t \geq 1$ and for any pair of agents $(i, j) \in N_t^{AB}$, $\chi_{ijt}^P = 1$ implies $\chi_{ijt}^W = 1$.*

2. *If only the platform's aggregate capacity constraint is potentially binding, then at each period, certain matches active under profit-maximization need not be active under welfare maximization. However, for all $t \geq 1$, $\sum_{(i,j)\in N_t^{AB}} \chi_{ijt}^W \geq \sum_{(i,j)\in N_t^{AB}} \chi_{ijt}^P$.*

3. *If some of the individual capacity constraints may bind then it need not be the case that the total number of matches active under welfare maximization is at least as large as under profit maximization. However, for all $t \geq 1$, $\sum_{(i,j)\in N_t^{AB}} \chi_{ijt}^P > 0$ implies $\sum_{(i,j)\in N_t^{AB}} \chi_{ijt}^W > 0$.*

As in other screening problems, distortions are introduced under profit maximization to reduce the agents' information rents (that is, the surplus the platform must leave to the agents to induce them to reveal their private information). When all agents value positively interacting with all other agents from the opposite side and none of the capacity constraints binds, in each period, a profit-maximizing platform induces fewer interactions than a welfare-maximizing one. In particular, any match that is active under profit maximization is also active under welfare maximization. This is because, when none of the capacity constraints binds, both under profit and under welfare maximization, the platform implements in each period all matches for which the score is non-negative. The result then follows from the fact

[33]To facilitate the comparison between profit and welfare maximization, we assume that, in each period, both the profit- and the welfare-maximizing auctions match all pairs for which the bilateral score is zero when the capacity constraints are not binding.

that, at each period, and for each match, the score under profit maximization is smaller than under welfare maximization due to the handicaps $[1 - G^k(\theta^k)]/g^k(\theta^k)\theta^k$ the platform applies to the scores under profit maximization to account for the cost of leaving rents to the agents.

When, instead, some of the capacity constraints potentially bind, certain matches active under profit maximization need not be active under welfare maximization. Yet, when the only potentially binding constraints are the platform's, at any point in time, the total number of matches under welfare maximization is at least as high as under profit maximization. This property follows again from the fact that each score under welfare maximization is at least as large as the corresponding one under profit maximization, along with the fact that the total number of matches active under both profit and welfare maximization is the minimum between the number of matches for which the score is non-negative and the platform's aggregate capacity. That some of the matches active under profit maximization need not be active under welfare maximization follows from the fact the ranking of the scores when the weights are given by β^P need not coincide with the ranking of the scores given the weights β^W.

Interestingly, in markets in which some of the individual capacity constraints potentially bind, it need not be the case that the total number of matches active under welfare maximization is at least as large as under profit maximization. To see this, suppose, in period t, there are 2 agents on each side of the market, that matching is one-to-one, and that $M_t > 4$, so that the only relevant capacity constraints in period t are the individual ones. Let S^P_{ijt} and S^W_{ijt} be the period-t scores under the weights β^P and β^W, respectively. Further suppose that $S^W_{22t} < 0$, and that $S^P_{ijt} > 0$ for any $(i,j) \neq (2,2)$. Lastly, assume that $S^W_{11t} > S^W_{12t} + S^W_{21t}$, whereas $S^P_{11t} < S^P_{12t} + S^P_{21t}$. Then, despite all scores being higher under welfare maximization than under profit maximization, the welfare-maximizing auction implements a single match in period t, the one corresponding to $(1,1)$, whereas the profit-maximizing auction implements two matches in the same period, $(1,2)$ and $(2,1)$. The reason is that, once the match

$(1,1)$ is formed, the matches $(1,2)$ and $(2,1)$ become infeasible, due to the individual capacity constraints. In this case, what remains true though is that, if matching is not completely shut down under profit maximization (i.e., at least one match is active), matching is also not completely shut down under welfare maximization (part 3 in Theorem 4).

The above conclusions need not extend to settings in which certain agents dislike certain interactions (formally, the horizontal types may take on negative values for certain pairs). In such settings, a profit-maximizing platform may induce an inefficiently *high* number of matches within each period and, over time, each pair may experience a larger number of matches under profit maximization than under welfare maximization, irrespective of the capacity constraints. Formally, when certain interactions may generate negative match values, a pair's score under profit maximization may be greater than its counterpart under welfare maximization, which implies that *the total number of interactions under profit maximization may exceed the efficient level.* The reason why a profit-maximizing platform may induce an inefficiently high number of interactions is that this may discourage the agents from purchasing a lower membership status. By locking agents selecting a low status into unpleasant interactions, the platform makes it costly for those agents with a high vertical type to pretend to have a low type. In turn, this permits the platform to extract more surplus from the high-type agents. The following example illustrates.

Example 3 (**Upward Distortions under Negative Values**)**.** Consider the following environment where $N_t^A = N_t^B = \{1\}$, $M = 1$, and $c_{11t} = 0$, all $t \geq 1$. The vertical types are given by $\Theta^B = \{1\}$ and $\Theta^A = [1+\varsigma, 2+\varsigma]$, $\varsigma > 0$, with G^A uniform over Θ^A. At each period $t \geq 1$, regardless of past realizations, $\varepsilon_{11t}^B = 1$, whereas ε_{11t}^A is drawn uniformly from $\{-3, +3\}$. Suppose the realized vertical type of agent 1 from side A is equal to $1+\varsigma$, in which case the weights used under profit maximization to scale the two agents' bids are given by $\beta_1^{A,P}(\theta_1^A) = \frac{\varsigma}{1+\varsigma}$ and $\beta_1^{B,P}(\theta_1^B) = 1$. Furthermore, consider a realized sequence $(\varepsilon_{11t}^A)_{t=1}^{\infty}$ of horizontal types for agent 1 from side A such that $\varepsilon^A_{11t} = -3$, all $t \geq 1$. Then, for sufficiently small ς and δ, the pair is matched in each period under profit maximization, despite matching being inefficient.

Finally, note that the familiar result of *"no distortion at the top"* from standard screening problems does not apply to a matching environment. A profit-maximizing platform may distort the matches of all agents, including those "at the top" of the distribution, for whom the vertical type is the highest. The reason is that, contrary to standard screening problems in which the cost of procuring inputs is exogenous, in a matching market, the cost of "procuring" agents-inputs from the opposite side of the market is endogenous and is higher than under welfare maximization, due to the informational rents the platform must provide to such agents-inputs to induce them to reveal their private information.

The above results bear certain implications for government intervention in matching markets, a topic that is receiving increasing attention in recent years. In markets in which capacity constraints are unlikely to be binding and agents are unlikely to suffer losses from interacting with others, platforms should be encouraged to implement more matches. Importantly, though, even if platforms could be induced to run welfare-maximizing auctions (more generally, to implement the welfare-maximizing matches), such auctions are not guaranteed to yield positive profits to the platforms (despite the fact that they minimize the platforms' losses over all mechanisms implementing the welfare-maximizing matches, as established in part (ii) of Theorem 3). As a result, subsidizing such markets may be necessary. On the other hand, in markets where certain agents may incur a cost for interacting with certain agents from the opposite side and/or where binding capacity constraints are prominent features, profit-maximizing platforms may implement too many matches in each period, and those implemented need not be the most valuable from a welfare standpoint. Policy makers may thus need to discourage certain interactions but inducing the platforms to implement the right allocations with simple taxes and subsidies may be impossible. It may be easier for the government to take over the private sector and provide directly the matching services.

5 Conclusions

This paper studies dynamic matching markets in which agents arrive stochastically over time, experience shocks to their match values, and are repeatedly re-matched with the help of a platform. We introduce a class of auctions that are specifically designed for such markets and that are fairly simple to operate. Upon joining the platform, agents are asked to select a membership status which determines the weight assigned to their bids in the subsequent auctions. They then bid repeatedly for potential partners from the opposite side of the market. In each period, the platform computes bilateral scores, one for each match, and implements those matches that maximize the scores, subject to individual and aggregate capacity constraints. We show that, under certain conditions, such auctions admit as special cases auctions maximizing profits, welfare, or a convex combination of the two, over all possible mechanisms. The analysis also sheds light on the inefficiencies that arise when platforms enjoy strong market power, which can be useful to guide government interventions in platform markets.

In the online Supplementary Material, we show how similar auctions can be used in markets in which the evolution of the match values is endogenous, be it the result of experimentation (whereby agents learn the attractiveness of their partners by interacting with them), a preference for variety (whereby agents gradually lose interest in those partners they already interacted with, reflecting various forms of satiation), or habit formation (whereby match values increase with the number of past interactions). In these markets, the bilateral scores take the form of forward-looking "indices" that account for (a) the benefit of generating new information (in the case of experimentation), (b) the opportunity cost of reducing future match values (in the case of a preference for variety), or (c) the value of enhancing future match values (in the case of habit formation). With endogenous processes, the conditions under which such matching auctions are fully optimal (i.e., profit- or welfare-maximizing) are more restrictive. In particular, they require the match values to evolve independently across matches, each match value to remain frozen when the match is not active, and the capacity constraints to satisfy a certain separability condition, which always holds when none of the capacity constraints binds or when a single match is formed in each period, but is restrictive in more general environments. Such conditions are typical in dynamic problems with endogenous processes. Notwithstanding these limitations, the insights are similar to those for the case of exogenous processes. Importantly, even when such conditions are not satisfied, the proposed auctions (in which the bilateral scores take the form of indexes) continue to admit straight-forward equilibria. Furthermore, in many problems of interest, such auctions yield the platform a reasonable fraction of the maximal profits, even if they are not fully optimal.[34]

In future work, it would be interesting to study how the auctions must be adapted if one side cannot bid for the matches, as is currently the case in sponsored search auctions. It would also be interesting to consider markets in which the platform must incur a cost only when it matches a pair that was not matched in the preceding period. Such costs introduce additional non-separabilities in the matching allocations that translate into a certain form of inertia and give rise to a trade-off between improving the quality of the existing matches and economizing on future re-matching costs, which is absent in the present analysis.

It would also be interesting to endogenize the process governing the arrival of the agents, for example by allowing the platform to invest in marketing activities that promote its services as

[34]When the aforementioned conditions are violated, it is typically impossible to identify either the profit-maximizing or the welfare-maximizing mechanisms. Nonetheless, using approximation results, one can show that the match allocations under the straight-forward equilibria of the proposed auctions yield a decent fraction of the maximal profits (alternatively, of maximal welfare).

a function of the evolution of the match values of those agents already on board. It would also be interesting to investigate whether distortions are more pronounced for "young" platforms with a small base or for "more mature" ones with many agents on board.

Another interesting direction would be to consider dynamic competition between platforms. Finally, it would be interesting to compare the matching dynamics in centralized markets such as those investigated in the present paper to their counterparts in decentralized markets, where agents match with other agents without the help of a platform. To the best of our knowledge, there is no tractable model of decentralized matching where agents perform on-the-job search and re-match over time. Developing such a model is challenging but is an important next step for this literature.

Appendix

Proof of Theorem 1. The proof is in two steps. Step 1 shows that, for any period $t \geq 1$ and any period-$(t-1)$ history, any agent who is on board at period t and who expects all other agents to follow a straight-forward strategy, finds it optimal to remain on board and bid truthfully. Step 2 first shows that the matches implemented under straight-forward strategies satisfy two key monotonicity conditions (defined below). It then shows that such monotonicities imply that joining the platform immediately upon arrival and selecting the membership status corresponding to an agent's true vertical type are optimal for any individual who expects all other individuals to follow straight-forward strategies.

Step 1. We start by establishing the following result:

Lemma 1 (Optimality for $t \geq 1$). *Consider an agent who joined the platform at any period $t' < t$ and who expects all other agents to follow straight-forward strategies. It is* periodically ex-post *optimal for the agent to bid truthfully in period t.*[35]

Proof of Lemma 1. Consider an agent from side A occupying the l-th position (in short, agent l from side A)—the problem for any side-B agent is analogous and hence not considered). Suppose that the profile of vertical types of all agents who arrived prior to period t is $\theta_t \equiv (\theta_i^k)_{i \in N_t^k}^{k=A,B}$, and that the history of horizontal types for such agents is $\varepsilon_t \equiv (\varepsilon_{ijs}^k)_{(i,j) \in N_s^{AB}, s=1,\ldots,t}^{k=A,B}$. Observe that, in the proposed auctions, the matches and the payments at each period $s > t$ are invariant to the period-t bids. Likewise, the period-t matches and payments are invariant to the bids submitted and the matches implemented in previous periods. Furthermore, when all other agents follow straight-forward strategies, their behavior in the continuation game starting with period t does not depend on their past behavior or on agent l's bids. To prove the result, it thus suffices to show that the flow payoff the agent obtains in period t by bidding truthfully is higher than the flow payoff he obtains by submitting any other vector of bids. Below we show that this is the case, no matter the status of the agents who are on board at period t and of the period-t bids submitted by the other agents.

Fix the status of all agents on board and let $\left(\beta^k(\hat{\theta}_l^k)\right)_{l \in N_t^k}^{k=A,B}$ denote the vector of corresponding weights. Because these weights are kept constant, throughout, to ease the notation, we drop the statuses from the arguments of all relevant functions and denote all the simplified functions (which thus depend only on bids) with hats.

Let $\hat{\chi}_t(b_t)$ denote the period-t matches that, given the period-t bids $b_t \equiv (b_{ijt}^k)_{(i,j) \in N_t^{AB}}^{k=A,B}$, maximize the *aggregate period-t score*

$$\sum_{(i,j) \in N_t^{AB}} \hat{S}_{ijt}(b_t) x_{ijt} \tag{12}$$

[35] As explained in the main text, "periodic ex-post optimal" means for any beliefs the agent may have about the history at the beginning of period-t, i.e., the time at which the other agents present at the beginning of period t joined, the history of the vertical and horizontal types of all such agents, and the behavior of such agents in previous periods.

subject to aggregate and individual capacity constraints, where, for any match $(i, j) \in N_t^{AB}$, the period-t bilateral score for the match between position i from side A and position j from side B is equal to

$$\hat{S}_{ijt}(b_t) \equiv \hat{\beta}_i^A b_{ijt}^A + \hat{\beta}_j^B b_{ijt}^B - c_{ijt}. \tag{13}$$

Similarly, let $\hat{\chi}_t^{-l,k}(b_t)$ denote the period-t matches that maximize the aggregate period-t score (as defined in (12)) subject to individual and aggregate capacity constraints, when the period-t bilateral score of any match involving position l from side k instead of taking the form in (13) is identically equal to zero (in this case, $\hat{\chi}^{-l,k}$ can be assumed to implement no match involving position l from side k). Let

$$\hat{S}_t\left(b_t\right) \equiv \sum_{(i,j)\in N_t^{AB}} \hat{S}_{ijt}(b_t)\hat{\chi}_{ijt}\left(b_t\right)$$

denote the aggregate period-t score induced by $\hat{\chi}_t(b_t)$ and

$$\hat{S}_t^{-l,k}\left(b_t\right) \equiv \sum_{(i,j)\in N_t^{AB}} \hat{S}_{ijt}(b_t)\hat{\chi}_{ijt}^{-l,k}\left(b_t\right)$$

the corresponding aggregate period-t score under $\hat{\chi}_t^{-l,k}(b_t)$. Finally, denote by $\hat{\psi}_{lt}^A\left(b_t\right)$ the period-t payment from the agent occupying the lth position on side A given b_t, as specified by the auction's payment rule (5).

Given the arguments above, it suffices to show that, for all possible match values $v_{lt}^A \equiv (\theta_l^A \varepsilon_{ljt}^A)_{j\in N_t^B}$ and period-t bids $b_{lt}^A \equiv (b_{ljt}^A)_{j\in N_t^B}$ by agent l from side A, any profile of bids $b_t^{-l,A} \equiv \left((b_{ijt}^A)_{(i,j)\in N_t^A\setminus\{l\}\times N_t^B}, (b_{ijt}^B)_{(i,j)\in N_t^A\times N_t^B}\right)$ by all agents other than agent l from side A,

$$\sum_{j\in N_t^B} v_{ljt}^A \hat{\chi}_{ljt}\left(v_{lt}^A, b_t^{-l,A}\right) - \hat{\psi}_{lt}^A\left(v_{lt}^A, b_t^{-l,A}\right) \geq \sum_{j\in N_t^B} v_{ljt}^A \hat{\chi}_{ljt}\left(b_{lt}^A, b_t^{-l,A}\right) - \hat{\psi}_{lt}^A\left(b_{lt}^A, b_t^{-l,A}\right), \tag{14}$$

When the payments are given by the rule in (5), the left hand side of the above inequality is equal to

$$\frac{1}{\hat{\beta}_l^A}\left[\hat{S}_t\left(v_{lt}^A, b_t^{-l,A}\right) - \hat{S}_t^{-l,A}\left(v_{lt}^A, b_t^{-l,A}\right)\right]. \tag{15}$$

Furthermore, the period-t payment in the right-hand side of (14) is equal to

$$\begin{aligned}
&\hat{\psi}_{lt}^A\left(b_{lt}^A, b_t^{-l,A}\right)\\
&= \sum_{j\in N_t^B} b_{ljt}^A \hat{\chi}_{ljt}\left(b_{lt}^A, b_t^{-l,A}\right) - \frac{1}{\hat{\beta}_l^A}\left[\hat{S}_t\left(b_{lt}^A, b_t^{-l,A}\right) - \hat{S}_t^{-l,A}\left(b_{lt}^A, b_t^{-l,A}\right)\right]\\
&= -\frac{1}{\hat{\beta}_l^A}\sum_{i\in N_t^A\setminus\{l\}}\sum_{j\in N_t^B}\hat{S}_{ijt}(b_{lt}^A, b_t^{-l,A})\hat{\chi}_{ijt}\left(b_{lt}^A, b_t^{-l,A}\right)\\
&\quad -\frac{1}{\hat{\beta}_l^A}\sum_{j\in N_t^B}\left(\hat{\beta}_j^B b_{ljt}^B - c_{ljt}\right)\hat{\chi}_{ljt}\left(b_{lt}^A, b_t^{-l,A}\right) + \frac{1}{\hat{\beta}_l^A}\hat{S}_t^{-l,A}\left(b_{lt}^A, b_t^{-l,A}\right).
\end{aligned}$$

This means that the right hand side of the inequality in (14) is equal to

$$\frac{1}{\hat{\beta}_l^A}\sum_{(i,j)\in N_t^{AB}}\hat{S}_{ijt}(v_{lt}^A, b_t^{-l,A})\hat{\chi}_{ijt}\left(b_{lt}^A, b_t^{-l,A}\right) - \frac{1}{\hat{\beta}_l^A}\hat{S}_t^{-l,A}\left(b_{lt}^A, b_t^{-l,A}\right).$$

Next, observe that $\hat{S}_t^{-l,A}$ is invariant to the bids by agent l from side A, meaning that

$$\hat{S}_t^{-l,A}\left(v_{lt}^A, b_t^{-l,A}\right) = \hat{S}_t^{-l,A}\left(b_{lt}^A, b_t^{-l,A}\right).$$

It follows that the inequality in (14) holds if and only if

$$\hat{S}_t\left(v_{lt}^A, b_t^{-l,A}\right) \geq \sum_{(i,j)\in N_t^{AB}}\hat{S}_{ijt}(v_{lt}^A, b_t^{-l,A})\hat{\chi}_{ijt}\left(b_{lt}^A, b_t^{-l,A}\right). \tag{16}$$

The inequality in (16) follows from the definition of the matching rule $\hat{\chi}_t(\cdot)$. This means that bidding truthfully yields the agent a payoff at least as high as under any other profile of bids.

Because the inequality holds no matter the period-t bids submitted by the other agents, the period-$(t-1)$ history, and the period-t profile of vertical and horizontal types of all agents on board, bidding truthfully is periodically ex-post optimal for agent l, as claimed. $\square$

Next, we establish the following result:

Lemma 2. *Consider an agent who joined the platform at any period $t' < t$ and who expects all other agents to follow straight-forward strategies. It is* periodically ex-post *optimal for the agent to remain in the auction in period t.*

Proof of Lemma 2. To see that it is periodically ex-post optimal for the agent to remain on board, note that, no matter the period-$(t-1)$ history and the period-t profile of true vertical and horizontal types, when all agents, including agent l from side A, follow straight-forward strategies, in the continuation game starting with period t, agent l's flow period-s payoff at any period $s \geq t$, and any period-s history, is proportional to his expected contribution to the aggregate score, which is always nonnegative given that $\hat{\beta}_l^A > 0$, and that at each period $\tau \geq t$, $\hat{S}_\tau(\cdot) - \hat{S}_\tau^{-l,A}(\cdot) \geq 0$. $\square$

Step 2. We now show that, when the membership fees are as in (7), joining immediately upon arrival and selecting the membership status designed for the agent's true vertical type are optimal for any agent who expects all other agents to follow straight-forward strategies.

Consider an agent from side k with vertical type θ^k. Let

$$\hat{D}_t^k(\hat{\theta}^k, \theta^k; \beta) \equiv \begin{cases} \mathbb{E}\left[\sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{j \in N_s^B} \varepsilon_{ijs}^A \chi_{ijs}^\beta(\hat{\theta}_i^A, \theta_s^{-i,A}, b_s) \mid \theta^k\right] & \text{if } k = A \\ \mathbb{E}\left[\sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{i \in N_s^A} \varepsilon_{ijs}^B \chi_{ijs}^\beta(\hat{\theta}_j^B, \theta_s^{-j,B}, b_s) \mid \theta^k\right] & \text{if } k = B \end{cases},$$

denote the "match quality" the agent expects by joining in period t and selecting the status designed for type $\hat{\theta}^k$, when the agent plans to bid truthfully at all subsequent periods, and expects all other agents to follow straight-forward strategies. Note that the expectation is over the position occupying by the individual, the identities and vertical types of the other agents on board, all future arrivals, the evolution of the agents' match values, and all bids. Note that when $\hat{\theta}^k = \theta^k$,

$$\hat{D}_t^k(\hat{\theta}^k, \theta^k; \beta) = D_t^k(\theta^k; \beta), \tag{17}$$

with the functions D_t^k as defined in (6).

The next lemma, which is the key step in the proof, shows that, no matter the true type θ^k, the expected match quality $\hat{D}_t^k(\hat{\theta}^k, \theta^k; \beta)$ is nondecreasing in the selected status $\hat{\theta}^k$, and that the match quality expected when selecting the status corresponding to the true vertical type, $D_t^k(\theta^k; \beta)$, is non-decreasing in the true vertical type, θ^k. The first property guarantees the optimality of selecting the status corresponding to the true vertical type. The second monotonicity guarantees the monotonicity of the agents' equilibrium payoffs in their true vertical type, which in turn guarantees that when low types find it optimal to participate, so do higher vertical types.

Lemma 3 (**Key monotonicities**)**.** *For any $t \geq 0$, and any $k = A, B$, the following monotonicities hold:*

(i) For any θ^k, $\hat{D}_t^k(\hat{\theta}^k, \theta^k; \beta)$ is non-decreasing in $\hat{\theta}^k$;

(ii) $D_t^k(\theta^k; \beta)$ is non-decreasing in θ^k.

Proof of Lemma 3. Consider an arbitrary agent from side A who arrives at period t and, upon joining the platform, occupies the i-th position (the arguments for the side-B agents are analogous). Fix the profile of vertical types $\theta_t^{-i,A}$ for the other agents on board.

We establish part (ii) first.

Part (*ii*). Take any agent from side A (the arguments for the side-B agents are analogous and hence omitted) and take any pair $\theta^A, \hat{\theta}^A \in \Theta^A$, with $\theta^A < \hat{\theta}^A$. Denote by i the generic position that the agent will occupy. Recall that the agent learns i only after joining and selecting his status.

That, under straight-forward strategies, the matches implemented in each period maximize the aggregate score, i.e., the sum of the individual bilateral scores (with the weights determined by the agents' statuses), subject to the aggregate and the individual capacity constraints, implies that

$$
\begin{aligned}
&\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{r\in N_s^A\setminus\{i\}}\sum_{j\in N_s^B}\left(\beta_r^A b_{rjs}^A+\beta_j^B b_{rjs}^B-c_{rjs}\right)\chi_{rjs}^{\beta}\left(\left(\hat{\theta}^A,\theta_s^{-i,A}\right),\left(\left(\hat{\theta}^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\hat{\theta}^A\right]\\
&+\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{j\in N_s^B}\left(\beta^A(\hat{\theta}^A)\left(\hat{\theta}^A\varepsilon_{ijs}^A\right)+\beta_j^B b_{ijs}^B-c_{ijs}\right)\chi_{ijs}^{\beta}\left(\left(\hat{\theta}^A,\theta_s^{-i,A}\right),\left(\left(\hat{\theta}^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\hat{\theta}^A\right]\\
&\geq\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{r\in N_s^A\setminus\{i\}}\sum_{j\in N_s^B}\left(\beta_r^A b_{rjs}^A+\beta_j^B b_{rjs}^B-c_{rjs}\right)\chi_{rjs}^{\beta}\left(\left(\theta^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\hat{\theta}^A\right]\\
&+\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{j\in N_s^B}\left(\beta^A(\hat{\theta}^A)\left(\hat{\theta}^A\varepsilon_{ijs}^A\right)+\beta_j^B b_{ijs}^B-c_{ijs}\right)\chi_{ijs}^{\beta}\left(\left(\theta^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\hat{\theta}^A\right]\\
&=\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{r\in N_s^A\setminus\{i\}}\sum_{j\in N_s^B}\left(\beta_r^A b_{rjs}^A+\beta_j^B b_{rjs}^B-c_{rjs}\right)\chi_{rjs}^{\beta}\left(\left(\theta^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right]\\
&+\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{j\in N_s^B}\left(\beta^A(\hat{\theta}^A)\left(\hat{\theta}^A\varepsilon_{ijs}^A\right)+\beta_j^B b_{ijs}^B-c_{ijs}\right)\chi_{ijs}^{\beta}\left(\left(\theta^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right].
\end{aligned}
$$

The left-hand side of the above inequality is the discounted sum of the aggregate score the agent expects when his true vertical type is $\hat{\theta}^A$ and all agents follow straight-forward strategies. Note that, no matter the position i the agent will occupy, the vertical type (and status) corresponding to his position is always equal to $\hat{\theta}^A$. The vector $\theta_s^{-i,A}$ denotes the vertical types (and hence the statuses) of all positions occupied in period s, excluding position i from side A. A similar notation applies to the bids $b_s^{-i,A}$.

The right-hand side, instead, is the discounted sum of the aggregate score expected by the same agent (of true vertical type $\hat{\theta}^A$) when the following two properties hold:

1. The agent replicates the behavior of an agent from side A of vertical type equal to θ^A in all periods (that is, he selects the membership status θ^A and then, at each subsequent

period, after learning his position i and his horizontal types $\varepsilon_{is}^A \equiv (\varepsilon_{ijs}^A)_{j=1,\ldots,n_s^B}$, he submits bids equal to $b_{ijs}^A = \theta^A \varepsilon_{ijs}^A$, as if his true vertical type was θ^A);

2. The weight the individual assigns to his own match values in the calculation of the aggregate score in each period is $\beta^A(\hat{\theta}^A)$.

The inequality follows from the fact that, when the individual (whose type is $\hat{\theta}^A$) assigns a weight $\beta^A(\hat{\theta}^A)$ to his own match values in the computation of the aggregate score in each period, the matches implemented by

$$\chi_s^\beta \left((\hat{\theta}^A, \theta_s^{-i,A}), \left((\hat{\theta}^A \varepsilon_{ijs}^A)_{j \in N_s^B}, b_s^{-i,A} \right) \right)$$

maximize the aggregate score when the membership status selected by the individual coincides with $\beta^A(\hat{\theta}^A)$ and the bids the individual submits correspond to his true match values.

The equality in the above expression in turn follows from the independence of the horizontal types from the vertical ones.

The above inequality is not to be confused with agent i's incentive-compatibility constraints. As explained in the main text, the monotonicity of match quality in the lemma does not follow from standard arguments in (both static and dynamic) screening models where the monotonicity of the allocations follows from the combination of incentive compatibility with the supermodularity of the agents' payoffs (between vertical types and allocations).

Now, inverting the role of $\hat{\theta}^A$ and θ^A in the above inequality, we have that

$$\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{r\in N_s^A\setminus\{i\}}\sum_{j\in N_s^B}\left(\beta_r^A b_{rjs}^A+\beta_j^B b_{rjs}^B-c_{rjs}\right)\chi_{rjs}^{\beta}\left(\left(\theta^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right]$$
$$+\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{j\in N_s^B}\left(\beta^A(\theta^A)\left(\theta^A\varepsilon_{ijs}^A\right)+\beta_j^B b_{ijs}^B-c_{ijs}\right)\chi_{ijs}^{\beta}\left(\left(\theta^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right]$$
$$\geq\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{r\in N_s^A\setminus\{i\}}\sum_{j\in N_s^B}\left(\beta_r^A b_{rjs}^A+\beta_j^B b_{rjs}^B-c_{rjs}\right)\chi_{rjs}^{\beta}\left(\left(\hat{\theta}^A,\theta_s^{-i,A}\right),\left(\left(\hat{\theta}^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\hat{\theta}^A\right]$$
$$+\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{j\in N_s^B}\left(\beta^A(\theta^A)\left(\theta^A\varepsilon_{ijs}^A\right)+\beta_j^B b_{ijs}^B-c_{ijs}\right)\chi_{ijs}^{\beta}\left(\left(\hat{\theta}^A,\theta_s^{-i,A}\right),\left(\left(\hat{\theta}^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\hat{\theta}^A\right].$$

Combining the above two inequalities, and using the definition of the $D_t^k(\theta^k;\beta)$ functions, we have that

$$\left(\beta^A(\hat{\theta}^A)\hat{\theta}^A-\beta^A(\theta^A)\theta^A\right)\cdot\left(D_t^A(\hat{\theta}^A;\beta)-D_t^A(\theta^A;\beta)\right)\geq 0.$$

Because $\beta^A(\cdot)$ is strictly positive and non-decreasing and $\hat{\theta}^A>\theta^A$, it must be that

$$D_t^A(\hat{\theta}^A;\beta)\geq D_t^A(\theta^A;\beta).$$

Part (i). Because in each period the matches implemented under straight-forward strategies maximize the sum of the scores subject to the aggregate and the individual capacity constraints, we have that, for any $\hat{\theta}^A,\theta^A\in\Theta^A$, with $\theta^A<\hat{\theta}^A$,

$$\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{r\in N_s^A\setminus\{i\}}\sum_{j\in N_s^B}\left(\beta_r^A b_{rjs}^A+\beta_j^B b_{rjs}^B-c_{rjs}\right)\chi_{rjs}^{\beta}\left(\left(\hat{\theta}^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right]$$
$$+\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{j\in N_s^B}\left(\beta^A(\hat{\theta}^A)\left(\theta^A\varepsilon_{ijs}^A\right)+\beta_j^B b_{ijs}^B-c_{ijs}\right)\chi_{ijs}^{\beta}\left(\left(\hat{\theta}^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right]$$
$$\geq\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{r\in N_s^A\setminus\{i\}}\sum_{j\in N_s^B}\left(\beta_r^A b_{rjs}^A+\beta_j^B b_{rjs}^B-c_{rjs}\right)\chi_{rjs}^{\beta}\left(\left(\theta^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right]$$
$$+\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{j\in N_s^B}\left(\beta^A(\hat{\theta}^A)\left(\theta^A\varepsilon_{ijs}^A\right)+\beta_j^B b_{ijs}^B-c_{ijs}\right)\chi_{ijs}^{\beta}\left(\left(\theta^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right].$$

The left-hand side of the above inequality is the discounted sum of the aggregate score expected by an individual of true vertical type θ^A, when the weight the individual assigns to his own match values in the computation of the aggregate score is $\beta^A(\hat{\theta}^A)$ in each period, all other agents follow straight-forward strategies, and the individual selects the membership status $\hat{\theta}^A$ and then bids truthfully at all periods. The right-hand side, instead, is the discounted sum of the aggregate score expected by the same individual, when the weight the individual assigns to his own match values to compute the aggregate score continues to be $\beta^A(\hat{\theta}^A)$ in each period, but this time the individual selects the status designed for type θ^A and then bids truthfully in all periods. Note that, in this latter case, the selection of the actual matches occurs by the auction scaling agent i's bids by $\beta^A(\theta^A)$ given that the agent's selected status is θ^A.

The inequality follows from the fact that, when the agent computes the desirability of each match that involves him by weighting his match values by $\beta^A(\hat{\theta}^A)$, the aggregate score in each period is higher when, for given bids, the matches are selected by a rule that assigns the same weight $\beta^A(\hat{\theta}^A)$ to the agent's bids as the one used by the individual himself than when the weight used by the rule is $\beta^A(\theta^A)$ whereas the one used by the individual is $\beta^A(\hat{\theta}^A)$.

By the same arguments, inverting the role of $\hat{\theta}^A$ and θ^A we have that,

$$\begin{aligned}
&\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{r\in N_s^A\setminus\{i\}}\sum_{j\in N_s^B}\left(\beta_r^A b_{rjs}^A+\beta_j^B b_{rjs}^B-c_{rjs}\right)\chi_{rjs}^{\beta}\left(\left(\theta^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right]\\
&+\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{j\in N_s^B}\left(\beta^A(\theta^A)\left(\theta^A\varepsilon_{ijs}^A\right)+\beta_j^B b_{ijs}^B-c_{ijs}\right)\chi_{ijs}^{\beta}\left(\left(\theta^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right]\\
&\geq\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{r\in N_s^A\setminus\{i\}}\sum_{j\in N_s^B}\left(\beta_r^A b_{rjs}^A+\beta_j^B b_{rjs}^B-c_{rjs}\right)\chi_{rjs}^{\beta}\left(\left(\hat{\theta}^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right]\\
&+\mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{j\in N_s^B}\left(\beta^A(\theta^A)\left(\theta^A\varepsilon_{ijs}^A\right)+\beta_j^B b_{ijs}^B-c_{ijs}\right)\chi_{ijs}^{\beta}\left(\left(\hat{\theta}^A,\theta_s^{-i,A}\right),\left(\left(\theta^A\varepsilon_{ijs}^A\right)_{j\in N_s^B},b_s^{-i,A}\right)\right)|\theta^A\right].
\end{aligned}$$

Combining the above two inequalities, we obtain that

$$\left(\beta^A(\hat{\theta}^A) - \beta^A(\theta^A)\right) \cdot \theta^A \cdot \left(\hat{D}_t^A(\hat{\theta}^A, \theta^A; \beta) - \hat{D}_t^A(\theta^A, \theta^A; \beta)\right) \geq 0.$$

Because $\theta^A > 0$ and because $\beta^A(\cdot)$ is strictly positive and non-decreasing, we conclude that $\hat{D}_t^A(\cdot, \theta^A; \beta)$ is non-decreasing.[36] □

We now show that when status is priced according to the formula in (7), the monotonicities in the previous lemma imply that each agent finds it optimal to select the membership status designed for his true vertical type. Without loss of generality, take an arbitrary agent from side A (the arguments for the side-B agents are analogous), and let t denote again the period in which the agent arrives. Let

$$\begin{aligned} \hat{U}_t^A(\hat{\theta}^A, \theta^A; \beta) \equiv & \mathbb{E}\left[\sum_{s=t+1}^{\infty} \delta^{s-t} \sum_{j \in N_s^B} \theta^A \varepsilon_{ijs}^A \chi_{ijs}^{\beta}\left((\hat{\theta}^A, \theta_s^{-i,A}), b_s\right) \mid \theta^A\right] \\ & - \mathbb{E}\left[\sum_{s=t}^{\infty} \delta^{s-t} \psi_{is}^{A,\beta}\left((\hat{\theta}^A, \theta_s^{-i,A}), b_s\right) \mid \theta^A\right] \end{aligned}$$

denote the intertemporal payoff that the agent expects when his true vertical type is θ^A, he chooses the membership status designed for type $\hat{\theta}^A$, he then bids truthfully at all periods, and all other agents follow straight-forward strategies. Then let

$$U_t^A(\theta^A; \beta) \equiv \hat{U}_t^A(\theta^A, \theta^A; \beta) \tag{18}$$

be the payoff expected by the same agent under straight-forward strategies by all agents,

[36] Note that, because the only influence of the membership statuses on match quality is through their impact on the weights, if $\beta^A(\hat{\theta}^A) = \beta^A(\theta^A)$, then $\hat{D}_t^A(\hat{\theta}^A, \theta^A; \beta) = \hat{D}_t^A(\theta^A, \theta^A; \beta)$.

including himself. Using the definition of the payments, it is easy to verify that

$$U_t^A(\theta^A;\beta) = \int_{\underline{\theta}^A}^{\theta^A} D_t^k(y;\beta)dy + Q_t^k. \tag{19}$$

Recall from Step 1 that, irrespective of the selection of the membership status, once an agent joins, remaining in the auctions and bidding truthfully is periodic ex-post optimal for the agent in the continuation game starting from any period-s history, $s > t$ (including those reached off path, by previous deviations). Now consider a fictitious environment where the agent is constrained to select the status $\hat{\theta}^A$ in period t but is otherwise free to choose any strategy of his choice for the continuation game that starts with period $t+1$. The above results imply that $\hat{U}_t^A(\hat{\theta}^A, \theta^A; \beta)$ is a value function for the problem the agent faces in the fictitious environment. Standard envelope arguments (see, e.g., Milgrom and Segal (2002)) then imply that $\hat{U}_t^A(\hat{\theta}^A, \theta^A; \beta)$ is Lipschitz continuous in the agent's true vertical type and that it admits the following representation

$$\hat{U}_t^A(\hat{\theta}^A, \theta^A;\beta) = \hat{U}_t^A(\hat{\theta}^A, \hat{\theta}^A;\beta) + \int_{\hat{\theta}^A}^{\theta^A} \hat{D}_t^A(\hat{\theta}^A, y;\beta)dy. \tag{20}$$

Combining (19) with (20), we then have that

$$\begin{aligned}
\hat{U}_t^A(\hat{\theta}^A, \theta^A;\beta) &= \hat{U}_t^A(\hat{\theta}^A, \hat{\theta}^A;\beta) + \int_{\hat{\theta}^A}^{\theta^A} \hat{D}_t^A(\hat{\theta}^A, y;\beta)dy \\
&\leq \hat{U}_t^A(\hat{\theta}^A, \hat{\theta}^A;\beta) + \int_{\hat{\theta}^A}^{\theta^A} \hat{D}_t^A(y, y;\beta)dy \\
&= U_t^k(\hat{\theta}^A;\beta) + \int_{\hat{\theta}^A}^{\theta^A} D_t^A(y;\beta)dy \\
&= U_t^k(\theta^A;\beta),
\end{aligned}$$

where the first equality follows from (20), the inequality follows from part (i) in Lemma 3, the second equality follows from (17) and (18), and the last equality from (19).

Hence the results above imply that, for any period t, an agent arriving at period t prefers following a straight-forward strategy to joining the platform in period t and then following any other strategy. Provided that Q_t^A is large enough, the same property also implies that the agent prefers following a straight-forward strategy rather than never joining the platform. The existence of Q_t^A satisfying the above requirement in turn follows from the fact that D_t^A is uniformly bounded by E_t^A, which guarantees that any $Q_t^A \geq (\bar{\theta}^A - \underline{\theta}^A)E_t^A$ is large enough.

Finally, observe that the maximal payoff that the agent can obtain by joining the platform at any period $s > t$ is attained by the agent selecting the status designed for his true vertical type upon joining and then bidding truthfully in all subsequent periods. This means that the maximal expected payoff he obtains by postponing his joining to any period $s > t$ is equal to

$$\delta^{s-t}\left[\int_{\underline{\theta}^k}^{\theta^k} D_s^k(y;\beta)dy + Q_s^k\right] \geq 0.$$

Condition (8), along with the other properties above, therefore implies that following a straight-forward strategy is also better than joining at any subsequent period and then following any arbitrary strategy upon joining. ■

Proof of Theorem 2. Consider any feasible mechanism Γ and any BNE σ of the game induced by Γ. Denote by $\tilde{\chi}$ and $\tilde{\psi}$ the functions describing the matches and payments induced by σ in Γ, as a function of the period-t exogenous state variables $(\theta_{t+1}, \varepsilon_t)$, where $\theta_{t+1} \equiv (\theta_i^k)_{i\in N_{t+1}^k}^{k=A,B}$ is the collection of the vertical types of all agents who arrived prior to or in period t (recall that N_{t+1}^k is the set of agents present at the beginning of period $t+1$, i.e., who arrived in period $s = 0, ..., t$) and where $\varepsilon_t \equiv (\varepsilon_{ijs}^k)_{(i,j)\in N_s^{AB}, s=1,...,t}^{k=A,B}$ is the history of horizontal types for those agents who arrived before period t (recall that each agent arriving at period t learns his horizontal types for the other agents already on board at the beginning of period t or who join in period t only at the beginning of period $t+1$). Also recall that the period-t matches can involve only individuals already on board at the beginning of period t, whereas the period-t payments can specify payments also for those agents joining in period t. Finally, note that we are allowing here for *any* feasible mechanism; that is, the message and signal spaces may be different than those in the matching auctions.

The platform's profits under (Γ, σ) are equal to

$$\mathbb{E}\left[\sum_{k=A,B}\sum_{t=0}^{\infty}\delta^t \sum_{l\in N_{t+1}^k}\tilde{\psi}_{lt}^k - \sum_{t=1}^{\infty}\delta^t \sum_{(i,j)\in N_t^{AB}} c_{ijt}\tilde{\chi}_{ijt}\right]. \tag{21}$$

Alternatively, (21) can be rewritten as follows:

$$\mathbb{E}\left[\sum_{s=1}^{\infty}\sum_{(i,j)\in N_s^{AB}}\delta^s\left(\left(\theta_i^A\varepsilon_{ijs}^A+\theta_j^B\varepsilon_{ijs}^B-c_{ijs}\right)\tilde{\chi}_{ijs}\right)-\sum_{s=0}^{\infty}\sum_{k=A,B}\sum_{l\in N_{s+1}^k\setminus N_s^k}\delta^s\tilde{U}_s^k(\theta_l^k;\mathcal{I}_s^k)\right], \tag{22}$$

where $\tilde{U}_s^k(\theta^k;\mathcal{I}_s^k)$ denotes the equilibrium payoff expected by any agent from side k with vertical type θ^k joining at period s.

Following steps similar to those in Pavan, Segal and Toikka (2014, Theorem 1), we can then show that the equilibrium payoffs must satisfy the envelope condition

$$\tilde{U}_t^k(\theta^k;\mathcal{I}_t^k) = \tilde{U}_t^k(\underline{\theta}^k;\mathcal{I}_t^k) + \int_{\underline{\theta}^k}^{\theta^k}\tilde{D}_t^k(y;\mathcal{I}_t^k)dy. \tag{23}$$

where

$$\tilde{D}_t^k(\theta^k;\mathcal{I}_t^k) \equiv \begin{cases} \mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{h\in N_s^B}\varepsilon_{lhs}^A\tilde{\chi}_{lhs} \mid \theta^A,\mathcal{I}_t^A\right] & \text{if } k=A \\ \mathbb{E}\left[\sum_{s=t+1}^{\infty}\delta^{s-t}\sum_{h\in N_s^A}\varepsilon_{hls}^k\tilde{\chi}_{lhs} \mid \theta^B,\mathcal{I}_t^B\right] & \text{if } k=B \end{cases}$$

denotes the match quality expected by any individual from side k of vertical type θ^k when joining at period t. The above envelope condition, together with integration by parts, yields the following representation of the platform's profits,

$$\mathbb{E}\left[\sum_{s=1}^{\infty}\delta^{s}\sum_{(i,j)\in N_s^{AB}}\left(\left(1-\frac{1-G^A(\theta_i^A)}{g^A(\theta_i^A)\theta_i^A}\right)\theta_i^A\varepsilon_{ijs}^A+\left(1-\frac{1-G^B(\theta_j^B)}{g^B(\theta_j^B)\theta_j^B}\right)\theta_j^B\varepsilon_{ijs}^B\right.\right.$$
$$\left.\left.-c_{ijs}\right)\tilde{\chi}_{ijs}\right]-\sum_{s=0}^{\infty}\sum_{k=A,B}\sum_{l\in N_{s+1}^k\setminus N_s^k}\tilde{U}_s^k(\underline{\theta}^k;\mathcal{I}_s^k). \tag{24}$$

Clearly, because such a representation applies to any BNE of any mechanism, it also applies to the straight-forward equilibria of the matching auctions.

Now observe that, when the weights are given by β^P, the matches implemented under straight-forward strategies in the auctions with weights β^P maximize

$$\sum_{(i,j)\in N_s^{AB}}\left(\left(1-\frac{1-G^A(\theta_i^A)}{g^A(\theta_i^A)\theta_i^A}\right)\theta_i^A\varepsilon_{ijs}^A+\left(1-\frac{1-G^B(\theta_j^B)}{g^B(\theta_j^B)\theta_j^B}\right)\theta_j^B\varepsilon_{ijs}^B-c_{ijs}\right)x_{ijs}$$

in each period and each state.

Next, observe that, when the weights are given by β^P and the access fees are given by (7), with $Q_t^k=0$ all t, $k=A,B$, the payoff expected under straight-forward strategies by any agent joining the platform at each period t with the lowest vertical type is equal to zero. This means that, under the straight-forward strategies of the matching auctions with weights β^P and with access fees given by (7), with $Q_t^k=0$ all t, $k=A,B$, the platform's profits are as high as under any BNE of any mechanism Γ. Provided that such auctions admit straight-forward equilibria, we thus have that such auctions, with the associated straight-forward equilibria, are profit-maximizing. Finally, as established in Theorem 1, the proposed auctions admit straight-forward equilibria if and only if conditions (a) and (b) in the theorem hold. ■

Proof of Theorem 3. As explained in the main text, part (i) follows directly from Theorem 1. Part (ii) follows from arguments similar to those establishing the optimality of the auctions of Theorem 2. In particular, it follows from the fact that, (a) under any BNE of any mechanism Γ implementing the welfare-maximizing matches in all periods, the platform's expected profits satisfy the representation in (24), with $\tilde{U}_t^k(\underline{\theta}^k; \mathcal{I}_t^k) \geq 0$, and (b) in the auctions under consideration, when all agents follow straight-forward strategies, the payoff expected by each agent with the lowest vertical type upon joining is equal to zero in all periods. The above properties imply that, when the auctions under consideration admit straight-forward equilibria, the platform's profits under such equilibria are as high as under any BNE of any mechanism implementing the efficient matches in all periods. As established in Theorem 1, the proposed auctions admit straight-forward equilibria if and only if conditions (a) and (b) in the theorem hold. ■

Proof of Theorem 4. Let χ_t^P and χ_t^W denote the matches implemented in period t under the straight-forward equilibria of the profit-maximizing and the welfare-maximizing auctions of Theorems 2 and 3, respectively.

Similarly, let S_{ijt}^P and S_{ijt}^W denote the period-t scores under the straight-forward equilibria of the profit-maximizing and the welfare-maximizing auctions of Theorem 2 and 3, respectively.

First, observe that because $\beta^{k,P}(\theta^k) \leq 1 = \beta^{k,W}(\theta^k)$ for all $\theta^k \in \Theta^k$, $k = A, B$, when all horizontal types are nonnegative, for any $t \geq 1$ and any pair of agents $(i, j) \in N_t^{AB}$, $S_{ijt}^P \leq S_{ijt}^W$.

Part 1. When none of the capacity constraints are binding, in each period $t \geq 1$, the matches implemented under the straight-forward equilibria of the profit-maximizing auctions (alternatively, the welfare-maximizing auctions) are all those for which the scores $S_{ijt}^P \geq 0$ (alternatively, $S_{ijt}^W \geq 0$). This property, along with the fact that for any $t \geq 1$ and any $(i, j) \in N_t^{AB}$, $S_{ijt}^W \geq S_{ijt}^P$, immediately yields the result.

Part 2. Next, consider the case in which only the aggregate capacity constraint is potentially binding. The result then follows from the following two properties: (a) in each period

$t \geq 1$, the set of matches for which $S^P_{ijt} \geq 0$ is a subset of the set of matches for which $S^W_{ijt} \geq 0$, (b) the cardinality of the set of matches implemented in each period in a profit-maximizing auction (alternatively, in a welfare-maximizing auction) is the minimum between M and the cardinality of the set of matches for which $S^P_{ijt} \geq 0$ (alternatively, $S^W_{ijt} \geq 0$).

Part 3. Finally, consider the case in which some of the individual capacity constraints are potentially binding, i.e., $m^k_t < n^{-k}_t$, for some t, $k = A, B$. The fact that $S^W_{ijt} \geq S^P_{ijt}$ for all $t \geq 1$ and all $(i,j) \in N^{AB}_t$ implies that, at any period $t \geq 1$ at which $|\{(i,j) \in N^{AB}_t : \chi^P_{ijt} = 1\}| > 0$, it must be the case that $|\{(i,j) \in N^{AB}_t : \chi^W_{ijt} = 1\}| > 0$. This is because $\chi^P_{ijt} = 1$ implies that $S^P_{ijt} \geq 0$, and hence $S^W_{ijt} \geq 0$. By matching the pair (i,j), the platform then weakly increases welfare relative to the case in which no pair is matched in period t. ■

References

[1] **Akbarpour, M., Li, S., and S. Oveis Gharan** (2020), "Thickness and Information in Dynamic Matching Markets," *Journal of Political Economy, 128*(3), 783-815.

[2] **Allard, N.,** (2008), "Lobbying Is An Honorable Profession: The Right To Petition and The Competition To Be Right," *Stanford Law Review*, Vol. 19(1), 23-68.

[3] **Anderson, R., Ashlagi, I., Gamarnik, D., and Y. Kanoria** (2017), "Efficient Dynamic Barter Exchange," *Operations Research, 65*(6), 1446-1459.

[4] **Armstrong, M.** (2006), "Competition in Two-sided Markets," *The RAND Journal of Economics, 37*(3), 668-691.

[5] **Arnosti, N., Beck, M., and P. Milgrom** (2016), "Adverse Selection and Auction Design for Internet Display Advertising," *American Economic Review*, 106(10), 2852-66.

[6] **Asker, J., and E. Cantillon** (2008), "Properties of Scoring Auctions," *The RAND Journal of Economics,* 39, 69-85.

[7] **Athey, S. and G. Ellison** (2011), "Position Auctions with Consumer Search," *Quarterly Journal of Economics,* Vol. 126, 1213–1270.

[8] **Baccara, M., S. Lee, and L. Yariv** (2020), "Optimal Dynamic Matching," *Theoretical Economics, 15*(3), 1221-1278.

[9] **Baron, D. and D. Besanko** (1984), "Regulation and Information in a Continuing Relationship," *Information Economics and Policy*, 1, 447–470.

[10] **Belleflamme, P., and M. Peitz,** (2017), "Platforms and Network Effects," In Corchon, L. and M. Marini (Eds). *Handbook of Game Theory and Industrial Organization.* Edward Elgar. Chelthenam, UK, Northampton, MA, USA.

[11] **Bergemann, D. and A. Pavan** (2015), "Introduction to Symposium on Dynamic Contracts and Mechanism Design," *Journal of Economic Theory* 159, 679-701.

[12] **Bergemann, D. and J. Valimaki** (2010), "The Dynamic Pivot Mechanism," *Econometrica*, 78(2), 771–789.

[13] **Bergemann, D. and J. Valimaki** (2019), "Dynamic Mechanism Design: An Introduction," *Journal of Economic Literature, 57*(2), 235-74.

[14] **Besanko, D.** (1985), "Multiperiod Contracts Between Principal and Agent with Adverse Selection," *Economics Letters*, 17, 33–37.

[15] **Board, S.** (2007), "Selling Options," *Journal of Economic Theory* 136, 324–340.

[16] **Board, S.** (2009), "Monopoly Group Design with Peer Effects," *Theoretical Economics*, Vol. 4, 89-125.

[17] **Börgers, T.** (2015), "An Introduction to the Theory of Mechanism Design," Oxford University Press, Oxford.

[18] **Börgers, T., Cox, I., Pesendorfer, M., and V. Petricek** (2013), "Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence," *American Economic Journal: Microeconomics,* Vol. 5, 163-187.

[19] **Cabral L.** (2011): "Dynamic Price Competition with Network Effects," *Review of Economic Studies* 78 (1): 83-111.

[20] **Caillaud, B. and B. Jullien**, (2003): "Chicken and Egg: Competition among Intermediation Service Providers," *The RAND Journal of Economics*, 4(2), 309-328.

[21] **Che, Y. K.** (1993): "Design Competition through Multi-Dimensional Auctions," *RAND Journal of Economics* Vol. 24, pp. 668-680.

[22] **Cremer, J. and R. P. McLean** (1988), "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56(6), 1247–1257.

[23] **Courty, P. and H. Li** (2000), "Sequential Screening," *Review of Economic Studies*, 67, 697–717.

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

[25] **Damiano, E. and R. Lam,** (2005), "Stability in dynamic matching markets," *Games and Economic Behavior,* 52(1):34–53.

[26] **Dekel, E., Jackson, M.O. and A. Wolinsky,** (2008) "Vote Buying: General Elections," *Journal of Political Economy,* 116(2), 351–380.

[27] **Dizdar, D., and B. Moldovanu** (2016). "On the Importance of Uniform Sharing Rules for Efficient Matching," *Journal of Economic Theory,* 165, 106-123.

[28] **Doval, L.** (2016), "A Theory of Stability in Dynamic Matching Markets," *Theoretical Economics*, forthcoming..

[29] **Edelman, B., Ostrovsky, M., and M. Schwarz** (2007), "Internet Advertising and the Generalized Second-Price Auction," *American Economic Review*, Vol. 97, 242-259.

[30] **Eső, P. and B. Szentes (2007),** "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 74(3), 705–731.

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

[32] **Gershkov, A. and B. Moldovanu** (2014), "Dynamic Allocation and Pricing: A Mechanism Design Approach," MIT Press.

[33] **Gomes, R.** (2014), "Optimal auction design in two-sided markets," *The RAND Journal of Economics,* 45: 248–272.

[34] **Gomes, R. and A. Pavan** (2016), "Many-to-Many Matching and Price Discrimination," *Theoretical Economics*, 11(3), 1005–1052.

[35] **Gomes, R. and A. Pavan** (2019), "Price Customization and Targeting in Platform Markets," mimeo Northwestern University and Toulouse School of Economics.

[36] **Gomes, R. and A. Sweeney** (2014), "Bayes–Nash Equilibria of the Generalized Second-price Auction," *Games and Economic Behavior*, 86, 421–437.

[37] **Hoppe, H. C., Moldovanu, B., and A. Sela** (2009), "The Theory of Assortative Matching Based on Costly Signals," *Review of Economic Studies*, 76, 253–281.

[38] **Hoppe, H. C., Moldovanu, B., and E. Ozdenoren** (2011). “Coarse Matching with Incomplete Information,” *Economic Theory*, 47(1), 75-104.

[39] **Johnson, T.,** (2013), “Matching Through Position Auctions,” *Journal of Economic Theory*, 148, 1700-1713.

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

[41] **Kang, K. and H. You** (2016), “Lobbyists as Matchmakers in the Market for Access,” mimeo Carnegie Mellon University.

[42] **Kurino, M.** (2009) “Credibility, efficiency and stability: a theory of dynamic matching markets,” *Jena economic research papers, JENA.*

[43] **Mansour, Y,. S. Muthukrishnan, and N. Nisan** (2012) “Doubleclick Ad Exchange Auction,” mimeo, Tel Aviv University.

[44] **Martens, B.** (2016). “An Economic Policy Perspective on Online Platforms,” Institute for Prospective Technological Studies Digital Economy Working Paper, 5.

[45] **Moghaddam, M and S. Y. Nof** (2016), “Best Matching: Theory and Applications,” Springer.

[46] **Pavan, A.** (2017), “Dynamic Mechanism Design: Robustness and Endogenous Types,” *Advances in Economics and Econometrics*, Proceedings of the 11th World Congress of Econometric Society, Volume 1, 1-62.

[47] **Pavan, A., I. Segal, and J. Toikka** (2014), “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82, 601–653.

[48] **Pinker, E. J., Seidmann, A., and Y. Vakrat** (2003). "Managing Online Auctions: Current Business and Research Issues," *Management Science, 49*(11), 1457-1484.

[49] **Rochet, J. C., and J. Tirole** (2003), "Platform Competition in Two-sided Markets," *Journal of the European Economic Association, 1*(4), 990-1029.

[50] **Rochet, J. C., and J. Tirole** (2006), "Two-sided Markets: A Progress Report," *The RAND Journal of Economics, 37*(3), 645-667.

[51] Ü**nver, U.** (2010), "Dynamic Kidney Exchange," *Review of Economic Studies,* 77, 372-414.

[52] **Varian, H.** (2007), "Position Auctions," *International Journal of Industrial Organization,* Vol. 25, pp. 1163- 1178.

[53] **Varian, H. and C. Harris** (2014). "The VCG Auction in Theory and Practice," *American Economic Review*, 104(5), 442-45.