Matching Auctions: Experimentation and Cross-Subsidization*

Daniel Fershtman
Northwestern University

Alessandro Pavan
Northwestern University

July 2016

Abstract

We study mediated many-to-many matching in dynamic two-sided markets in which agents' valuations for potential partners evolve over time, as the result of shocks, learning through experimentation, or a preference for variety, by which valuations for current partners drop with the number of past interactions. The matching dynamics that maximize either the platform's profits or welfare can be sustained through a sequence of matching auctions in which agents bid for all possible partners and where the matches with the highest bilateral score are implemented. Each score is either myopic or takes the form of an index, accounting for the endogenous changes in the partners' expected match values. In equilibrium, bidding is straight-forward and myopic. The analysis also sheds light on the merits of regulating, taxing, or subsidizing such markets. When all agents benefit from interacting with all other agents from the opposite side, profit maximization involves fewer and shorter interactions than what is efficient, for all agents, including those at the top of the distribution. This conclusion, however, need not extend to environments where certain agents dislike certain interactions.

JEL Classification Numbers: D82, C73, L1.
Keywords: matching, experimentation, platforms, bandit problems, asymmetric information, learning, dynamic auctions.

*A previous version circulated under the title "Re-Matching, Experimentation, and Cross-Subsidization." For helpful comments and suggestions, we thank Yeon-Koo Che, Eddie Dekel, Drew Fudenberg, Benjamin Hermalin, Terrence Johnson, Fuhito Kojima, Larry Samuelson, Rakesh Vohra, Jean Tirole, Asher Wolinsky, Leeat Yariv, and seminar participants at the University of California Berkeley, the University of Chicago, Columbia University, the University of North Carolina, Duke University, the Toulouse School of Economics, the University of Western Ontario, the University of Southern California, the University of British Columbia, Stanford University, the University of Illinois at Urbana Champaign, IESE, the 2015 Paris Workshop on Information Economics and Dynamics, the 2016 Warwick Economic Theory Conference, the 2016 North American Summer Meetings of the Econometric Society, and the 2016 World Congress of the Game Theory Society. Pavan thanks the National Science Foundation for financial support under the grant SES 1530798. The usual disclaimer applies.

1 Introduction

Matching markets have been growing at an unprecedented rate in the last few years, reflecting the role that the "shared economy" is taking in the organization of modern business activities. In electronic commerce, for example, a sizeable fraction of trade is mediated by business-to-business (B2B) platforms matching vendors with procurers in search of business opportunities.[1] Likewise, a sizeable fraction of online advertising is mediated by media outlets, content providers, online malls, videogame consoles, and search engines, matching consumers with advertisers (see, e.g., "Marketing in the digital age: A brand new game", *The Economist*, August 29, 2015).

Mediated matching plays an important role also in the growing market for scientific outsourcing. Arrangements in the form of "contract-experiments," in which researchers or startup firms contract with laboratories that carry out research on their behalf, have existed for a long time.[2] Recently, however, new intermediary firms, such as Science Exchange, have revolutionized the market by introducing the ideas of the shared economy. These platforms match research labs with idle equipment with research units that wish to conduct experiments off-site (see, e.g., "Uber for Experiments," *The Economist*, December 6, 2014). They introduce parties that otherwise would be unlikely to come in contact, and also provide crucial services that remove most of the complexities involved in such ad-hoc relationships. For example, Science Exchange creates an agreement governing the nature of the interaction between the parties, a service without which negotiations, especially for startups, would be costly and lengthy. In addition, it offers a concierge service, which includes the services of a dedicated scientist with a master's degree or a PhD, who follows up the entire matching process. Finally, it provides access to a software that allows projects to be tracked through a webpage where parties can monitor their joint activity and exchange data, and that includes an integrated payment system.

Most matching markets are intrinsically dynamic, due to the gradual resolution of uncertainty about match values, shocks that alter the desirability of the existing matching allocations, or a preference for variety by which agents' valuations for existing partners drop with the number of past interactions. Furthermore, matching in these markets is typically many-to-many — the same agent is matched to multiple agents from the other side of the market.

In the market for scientific outsourcing, for example, the participating research labs, which are highly differentiated in size, equipment, and specialization, typically allow numerous firms to conduct experiments simultaneously within their facilities. Likewise, firms typically seek to run a battery of experiments across different labs within the same period.[3] Over time, labs learn about the reliability of the participating research firms and firms learn about the characteristics of the participating labs.

[1] According to the U.S Census Bureau, electronic commerce accounted for 47 percent of manufacturing total value in the US in 2010 (http://www.census.gov/econ/estats/2010/2010reportfinal.pdf).

[2] For example, facilities that own wind tunnels perform tests for the aerospace industry, for architectural design, for auto and ship design, and for bicycle design.

[3] The list of laboratories that work with Science Exchange includes 75 of the top 100 recipients of NIH grant funding. Examples include Johns Hopkins University, the Mayo Clinic, and Harvard Medical School.

Natural capacity constraints preclude the possibility of all firms using all labs simultaneously. As a result, experimentation and re-matching plays an important role in such exchanges. Over time, firms change labs as a function of the evolution of their needs and the information they gather about the participating labs. The prices asked by the labs are also dynamic, reflecting the information the labs obtain about the participating research units, as well as the competition among research firms.

Other matching markets sharing similar features include project finance, where consulting firms match start-ups seeking project-finance with banks;[4] lobbying, where commercial firms mediate the interactions between policy makers and interest groups;[5] the market for private medical-tourism services, where intermediaries match patients from abroad seeking specialized treatments with local physicians providing such treatments;[6] the market for organized events, where online platforms such as meetings.com match clients in search of conference venues, meeting spaces, corporate hotel travel plans, or other hospitality services, with hotels and venues.[7]

In all the above examples, matching is many-to-many and mediated by a profit-maximizing platform collecting payments from the various sides of the market. Matching is also dynamic, with agents changing partners over time using the same platform, and with prices also varying over time in response to the evolution of the agents' preferences and information. Furthermore, in most cases of interest, the services provided by the matching intermediary extend well beyond simply introducing agents from the various sides of the market, as the example of Science Exchange discussed above illustrates. The platforms' costs of such services are also naturally dynamic and match-specific. Finally, platforms often face capacity constraints that restrict the number of interactions they may accommodate.[8]

[4] A similar role is played by peer-to-peer lending platforms such as Prosper and LendingClub. Such platforms match borrowers requesting loans with individual or institutional investors. In addition to connecting borrowers with lenders, the platform verifies the borrowers' identities and personal data and manages all stages of funded loans. Prosper operated as an online auction marketplace between 2006 and 2009. It recently switched to a system of pre-set rates determined by an algorithm evaluating borrowers' credit risks.

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

[6] For example, MEDIGO matches providers of medical services abroad with potential patients. Similarly to Science Exchange, it also offers a "Custom A-to-Z Concierge Package" including a variety of costly services. See the New York Times article http://www.nytimes.com/2013/08/07/us/the-growing-popularity-of-having-surgery-overseas.html for details about the growing popularity of overseas surgeries.

[7] As in the case of Science Exchange, meetings.com does not simply introduce the parties, it also offers tailored services (through its local affiliates), such as site selection, contract negotiations, and on-site event management. Clients of meetings.com include corporations, online travel agencies, sport organizations, colleges, and government institutions. Matching is many-to-many, as hotels and venues hold multiple events within the same period, and clients seek a variety of locations for different types of events. Matching is also dynamic, as clients seek repeated services and both their preferences as well as those of the venues typically change over time as the result of previous experiences and variations in needs and availability.

[8] For instance, for medical-tourism intermediaries, the capacity constraint is determined by the medical facilities they contract with. For lobbying firms, the capacity constraint comes from the limited time they can devote to follow up on individual relationships. For online content providers, the capacity constraint comes from the impossibility to place too many adds on the same page.

A question that is receiving growing interest in the design of these markets is how prices and matches should respond to the dynamics of market conditions, variations in preferences, and the arrival of information. In large markets such as those for shared transportation services, it is natural for intermediaries such as Uber and Lift to price discriminate as a function of variations in demand and supply. In markets with a smaller number of participating agents such as those for scientific outsourcing, analysts have predicted that intermediaries will eventually resort to auctions to dynamically allocate their limited capacity. In fact, auctions are already used in the market for personalized display ads by online search engines such as Google and Yahoo!. These intermediaries have been using variations of the second-price auction, the so-called GSP auction (Generalized Second Price) to match ads with viewers — see, e.g., Edelman, Ostrovsky, and Schwarz 2007, and Gomes and Sweeney (2014). Such auctions, however, have been criticized for being static, and in particular for not taking into account the value that both the platform and the advertisers assign to learning the click-through rates through experimentation (see, e.g., Li, Mahdian, and McAfee (2010) and the discussion therein).

In this paper we propose a new class of matching auctions that account explicitly for the fact that, in many environments of interest, match values evolve over time, either exogenously (as the result of changes in preferences and needs) or endogenously (as the result of learning and experimentation). We investigate the properties of matching auctions that maximize either profits or welfare.

The key ingredients of our model are the following. The payoff that each agent derives from being matched to any other agent from the opposite side is governed by two components: a time-invariant vertical characteristic that is responsible for the overall importance that the agent assigns to interacting with agents from the opposite side of the market; and a vector of time-varying relation-specific values capturing the evolution of the agent's information and preferences for interacting with specific partners. The latter values evolve stochastically over time and may turn negative, reflecting the idea that agents may dislike certain interactions. We consider both the case in which the relation-specific values evolve exogenously, as well as the case in which they evolve endogenously as a function of previous interactions. The latter case may capture either the possibility that agents gradually learn the attractiveness of their partners via individual interactions, or a preference for variety, by which agents' values change (possibly stochastically) with the number of past interactions. Both the vertical and the horizontal components are the agents' private information. The model also accounts for the possibility that the level of activity the platform can accommodate within each period may be limited, reflecting either time, resource, or facility constraints.

The analysis first introduces and then studies the properties of a class of matching auctions that operate as follows. Upon joining the platform, agents select a membership status whose level determines the weight their bids receive in the subsequent auctions. At any period, agents are then asked to bid for each possible partner from the opposite side of the market. Each bilateral match then receives a "score" that depends on the involved agents' reciprocal bids, on their membership status, and on the number of past interactions between the pair of agents. The matches with the highest nonnegative score are then implemented, up to capacity. As in the case of the GSP auction used by search engines for display advertisement, the payments the platform asks to each agent reflect the externalities the agent imposes on others due to the capacity constraint. Contrary to the GSP auction, however, such externalities are computed taking into account the value of experimentation and the dynamics of the agents' informational rents. In particular, the platform may find it optimal to cross-subsidize certain interactions to generate information that can be used in future periods.

When match values evolve exogenously over time, the proposed scoring rules are *myopic*: The score the platform assigns to each pair of agents coincides with the pair's joint flow surplus, adjusted by (a) "handicaps" controlling for the informational rents the platform must leave to the agents, and (b) the platforms' own cost of implementing the matches, which reflect the relation-specific services provided by the platform, as discussed above.

When, instead, the match values evolve endogenously, as in the case of experimentation, or when agents have a preference for variety, the scores take the form of an *index*. As in other experimentation environments, such indexes summarize both the current and future expected profitability of each match, taking into account the dynamics of (i) the agents' joint values, (ii) informational rents, and (iii) the platforms' cost of implementing the matches.

In both cases, at all histories, including those off-path, in equilibrium, agents bid truthfully their myopic values for all partners. The reason why bidding the myopic values is optimal under the proposed auctions is that the matching allocations under truthful bidding maximize continuation weighted surplus, that is, the sum of all agents' current and future payoffs, net of the platform's matching costs, and net of the agents' information rents. This property, together with the fact that the payments make each agent's continuation payoff proportional to the agent's contribution to continuation weighted surplus, then guarantees that each agent finds it optimal to stay in the mechanism and bid truthfully at all periods, irrespective of the agent's beliefs about other agents' current and past types, as well as of the implemented past matches. As as result, the proposed auctions can be made fully transparent. At the end of each period, agents see the membership status selected by any other agent as well as all the bids received by the platform over time.

One of the advantages of the proposed auctions is that they explicitly account for the value of experimentation, as advocated by many analysts and market designers, without, however, expecting the agents to adjust their bids to incorporate such value. The complexity of computing indexes to optimally control for the trade-offs between "experimentation" and "exploitation" is entirely on the platform's side. Once the agents understand the proposed scoring rules and their associated payment schemes (which can be facilitated by the platform explaining such role to the bidders), it is in their interest to bid straightforwardly their myopic values in all periods.

The proposed matching auctions allow for a high sensitivity of matches and prices to the evolution of the agents' preferences and information. As mentioned above, such sensitivity is desirable in markets with a small number of participating agents, in which interactions are far from being anonymous. Certain results, however, have implications also for larger markets in which matching is more anonymous but where pricing is dynamic as it responds to the evolution of market conditions.

The results also shed light on the value of regulating such markets. The last few years have witnessed great interest (both from policy makers and academics) on how to regulate matching markets dominated by large platforms (see, e.g., the articles "Online Platforms: Nostrums for Rostrums" and "Regulating Technology Companies: Taming the Beasts," recently published in *The Economist*, May 28, 2016). We contrast matching dynamics under public (welfare-maximizing) and private (profit-maximizing) provision of matching services. When agents assign a nonnegative value to all interactions at every period, profit maximization results in fewer and shorter interactions. Specifically, when the capacity constraint is not binding, a profit-maximizing intermediary matches each pair of agents for an inefficiently short period of time. When, instead, the capacity constraint is binding, certain interactions may last longer under profit maximization than under welfare maximization. However, the aggregate number of interactions in each period under profit maximization is always lower than under welfare maximization. Interestingly, the above conclusions need not extend to markets in which certain agents derive a negative payoff (equivalently, a payoff lower than their outside option) from interacting with certain other agents. In this case, profit maximization may result in an inefficiently large volume of matches, for any number of periods. The above conclusions have implications for how governments should subsidize certain platforms while taxing others, as well as for the costs and benefits of leaving certain matching markets unregulated.

The rest of the paper is organized as follows. We wrap up the introduction with a brief discussion of the most pertinent literature. Section 2 describes the model. Section 3 introduces the matching auctions. Section 4 derives equilibrium properties of the proposed mechanisms. Section 5 identifies a subclass of matching auctions that are profit-maximizing. Section 6, instead, identifies a class of auctions that are welfare-maximizing and contrasts matching dynamics under profit maximization with their counterparts under welfare maximization. Section 7 concludes. All proofs are in the Appendix at the end of the document.

1.1 Related Literature

Centralized dynamic matching. Much of the recent literature on centralized dynamic matching focuses on markets without transfers, in which agents are matched once, with dynamics stemming from the arrival and possibly departure of agents to and from the market. In the context of kidney exchange, Ünver (2010) studies optimal mechanisms for two-way and multi-way exchanges, minimizing total waiting costs, in a market with stochastic arrivals of donors and recipients. Optimal dynamic matching in markets in which agents gradually arrive over time and may be matched at most once is also the focus of Anderson, Ashlagi, Gamarnik, and Kanoria (2015), Baccara, Lee, and Yariv (2015), Akbarpour, Li, and Oveis Gharan (2016), and Herbst and Schickner (2016). A key trade-off in such environments is between avoiding waiting costs and waiting for the market to thicken. A related strand of papers study the assignment of objects through waiting lists (see Leshno (2015), Thakral (2015), Bloch and Cantala (2016), and Schummer (2016) for recent developments).[9]

The key difference with respect to this literature is that, in the present paper, agents change partners over time in response to changes in their valuations for one another. In particular, dynamics stem from changes in actual or perceived match values rather than the arrival and departure of agents to and from the market.

Matching with transfers. The paper is also related to the literature on profit-maximization in matching markets with private information and transfers. Damiano and Li (2007) and Johnson (2013) consider a *one-to-one* matching environment where the intermediary faces asymmetric information about agents' vertical characteristics responsible for match values. Board (2009) studies the problem of a profit-maximizing platform (e.g., a school) that can induce agents to self-select into mutually exclusive groups (e.g., classes). These papers derive conditions on primitives for a profit-maximizing intermediary to induce positive assortative matching. Gomes and Pavan (2015) study many-to-many matching in a setting where agents differ both in their consumer value (willingness-to-pay) and input value (salience) and where matching is non-partitional.

Matching in all of these papers is static. In contrast, the present paper considers dynamic matching in an environment in which match values evolve, either exogenously, or endogenously, over time.

Position and scoring auctions. In procurement auctions, scoring rules are often used to aggregate the various dimensions of the sellers' offers (price, product design, delivery time, etc.). 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 scoring rules governing the auctions. However, while the above literature focuses on static settings, our scoring rules are for dynamic environments in which preferences evolve over time. Another difference is that our scores aggregate the preferences of different agents from different sides of the market, as opposed to the various dimensions of each seller's own offer.

Another related literature studies auctions for sponsored links. 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.[10] Our matching auctions could be used also for these markets, modulo the fact that, in online search, searchers typically do not pay for the matches. Notwithstanding this qualification, one of the advantages of our auctions is that they explicitly account for the value of generating information that can be used in future periods.

Two-sided markets. Markets where agents purchase access to other agents are the focus of the literature on two-sided markets (see Rysman (2009) for a survey, and Weyl (2010), Bedre-Defolie and Calvano (2013), Lee (2013), and Jullien and Pavan (2016) for some of the recent developments). This literature, however, restricts attention to a single network or to mutually exclusive networks. In contrast, the present paper allows for general matching rules and for more flexible payoff structures.

[9] See also Damiano and Lam (2005), Kurino (2009), and Doval (2015) for appropriate stability notions in dynamic matching environments.

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

In particular, it does not restrict agents' willingness to pay to coincide with their attractiveness. Most importantly, it focuses on a dynamic environment in which match values change over time. Cabral (2011) also considers a dynamic model with network effects but in which values are constant over time.

Screening under experimentation. The paper is also related to the literature on screening agents who are experimenting about the value of a good. In particular, the mechanisms proposed here can be seen as the matching analogs of the bandit auctions of Pavan et al. (2014) and Kakade et al. (2013) for the sale of an indivisible object. See also Bergemann and Valimaki (2006) for a survey on bandit problems in economics.

Mechanism design. From a methodological standpoint, we draw from recent results in the dynamic mechanism design literature. In particular, the conditions for incentive compatibility in the present paper adapt to the environment under examination results in Pavan, Segal, and Toikka (2014). See also Baron and Besanko (1984), Besanko (1985), Courty and Li (2000), Board (2007), and Eso and Szentes (2007), for some of the earlier contributions, Borgers (2015) and Bergemann and Pavan (2015) for a general overview of this literature, and Bergemann and Strack (2015) for recent developments in continuous time. In most of these models, the agents' private information is exogenous. In contrast, the agents' private information is endogenous in the present paper, in the bandit single-unit auction of Pavan et al. (2014), in the dynamic virtual pivot mechanism of Kakade et al. (2013), and in the taxation model of Makris and Pavan (2015).

Particularly related is also the strand of the dynamic mechanism design literature that investigates how to implement dynamically efficient allocations in settings in which the agents' types change over time, thus extending the Vickrey–Clarke–Groves (VCG) and d'Aspremont–Gérard-Varet (AGV) results from static to dynamic settings (see, for example, Bergemann and Valimaki, 2010, Athey and Segal, 2013, and the references therein). The payments in the auctions we propose here are determined by a pricing formula similar to those in Bergemann and Valimaki (2010), and in Kakade et al. (2013), adapted to the fact that the private information the agents receive in each period is multi-dimensional.

Another stream of the dynamic mechanism design literature considers both efficient and profit-maximizing mechanisms in settings where the agents' private information is static, but where the dynamics originate in the arrival and/or departure of objects and agents to and from the market (for an overview of this literature, see Bergemann and Said (2011), and Gershkov and Moldovanu (2014)).

2 Model

Agents, matches, and preferences

A platform mediates the interactions among agents from two sides of a market, A and B. There are $n_A \in \mathbb{N}$ agents on side A and $n_B \in \mathbb{N}$ agents on side B, with $N_A \equiv \{1, ..., n_A\}$ and $N_B \equiv \{1, ..., n_B\}$ denoting the corresponding sets of agents on the two sides. Time is discrete, indexed by $t = 0, 1, ..., \infty$. Agents live for infinitely many periods and can change partners infinitely many times.

Below, we describe various features of the environment from the perspective of a generic agent from side A. A similar description applies to side B.

The flow period-t utility that agent $i \in N_A$ derives from being matched to agent $j \in N_B$ is given by

$$u_{ijt}^A(\theta_i^A, \varepsilon_{ijt}^A) = \theta_i^A \cdot \varepsilon_{ijt}^A. \tag{1}$$

The parameter θ_i^A is time- and match-invariant and controls for the overall importance that agent i assigns to interacting with agents from the opposite side of the market. That is, θ_i^A parametrizes the value that agent i assigns to an interaction with a generic agent from the opposite side, prior to conditioning on the specific profile of the latter agent. The parameter ε_{ijt}^A, instead, is match-specific and time-variant and controls for the attractiveness of agent j from side B in the eyes of agent i. These match-specific values evolve over time, reflecting the change in the agents' true (or perceived) attractiveness. They can either represent the evolution of the agents' beliefs about fixed, but unknown, match qualities, or variations in attractiveness triggered by stochastic changes in the environment. Hereafter we refer to θ_i^A as the agent's "*vertical type*" and to $\varepsilon_{it}^A \equiv (\varepsilon_{ijt}^A)_{j \in N_B}$ as the profile of the agent's period-t "*horizontal types*". We refer to $u_{it}^A \equiv (u_{ijt}^A)_{j \in N_B}$ as the agent's period-t "*match values*".

Both the agent's vertical and horizontal types are his own private information. Agent i's vertical type θ_i^A is drawn from an absolutely continuous cumulative distribution function F_i^A with density f_i^A strictly positive over $\Theta_i^A = [\underline{\theta}_i^A, \overline{\theta}_i^A]$, with $\underline{\theta}_i^A > 0$. Vertical types are drawn independently across agents and from the horizontal types $\varepsilon \equiv (\varepsilon_{ijt}^k)_{t=1,\ldots,\infty}^{(i,j) \in N_A \times N_B, k=A,B} \in \mathcal{E}$. Importantly, while we restrict the agents' vertical types to be nonnegative, we allow the horizontal types to be negative, reflecting the possibility that an agent may derive a negative utility from interacting with certain agents from the opposite side.[11]

For any $t \geq 1$, and any pair of agents $(i, j) \in N_A \times N_B$, let $X_{ijt} \equiv \{0, 1\}$, with $x_{ijt} = 1$ denoting the decision to match the pair (i, j) in period t, and with $x_{ijt} = 0$ denoting the decision to leave the two agents unmatched.[12] Matching is many-to-many, meaning that the same agent may be matched to multiple agents from the opposite side.

All agents are expected-utility maximizers and maximize the expected discounted sum of their flow payoffs using the common discount factor $\delta \in (0, 1]$. Let $p_t \equiv (p_{lt}^k)_{l \in N_k, k=A,B}$ denote the payments collected by the platform from the two sides of the market in period t, and $p \equiv (p_t)_{t=0}^{\infty}$ an entire sequence of payments. Note that, while the matching starts in period one, we allow the platform to start collecting payments from the agents in period zero, after the agents have observed their vertical types, but before they observe their horizontal partner-specific types. This assumption is motivated by the idea that, in many markets of interest, at the time the agents join the platform, they do not

[11] Under the assumed multiplicative structure $u_{ijt}^A = \theta_i^A \cdot \varepsilon_{ijt}^A$, allowing the vertical types to also take on negative values is not desirable as it facilitates confusion given that the horizontal types ε_{ijt}^A are already allowed to take on negative values.

[12] Implicit in this formalization is a *reciprocity condition* imposing that whenever agent $i \in N_A$ is matched to agent $j \in N_B$, then agent j is matched to agent i (see also Gomes and Pavan (2015) for the role played by a similar condition in a different matching environment).

know yet the specific profile of the agents who join from the opposite side. In other words, agents learn about agents from the opposite side only after "getting on board". Also note that payments are allowed to be negative, reflecting the possibility that (a) certain agents may dislike certain interactions and ask to be compensated, or (b) even if all agents like interacting with all other agents, the platform may want to cross-subsidize certain matches.

The (Bernoulli) payoff function for each agent $i \in N_A$ is given by

$$U_i^A = \sum_{t=1}^{\infty} \delta^t \sum_{j \in N_B} x_{ijt} u_{ijt}^A - \sum_{t=0}^{\infty} \delta^t p_{it}^A. \tag{2}$$

The platform's (Bernoulli) payoff function is the discounted sum of the payments collected from the two sides of the market, net of possible costs of implementing the matches:

$$U_0 = \sum_{t=0}^{\infty} \delta^t \left(\sum_{i \in N_A} p_{it}^A + \sum_{j \in N_B} p_{jt}^B \right) - \sum_{t=1}^{\infty} \delta^t \left(\sum_{i \in N_A} \sum_{j \in N_B} c_{ijt}(x_{ij}^{t-1}) \cdot x_{ijt} \right),$$

where $c_{ijt} \geq 0$ is the period-t cost of matching the pair (i, j), with $x_{ij}^{t-1} \equiv (x_{ijs})_{s=1}^{t-1} \in X_{ij}^{t-1} \equiv \prod_{s=1}^{t-1} X_{ijs}$ denoting the history of past interactions between the pair (i, j). These costs incorporate all the auxiliary services that the platform provides to the agents, over and above putting the agents in contact one with the other.

Evolution of match values

The match values u_{ijt}^A that agent i from side A derives from interacting with agent j from side B are correlated over time, both through the fully persistent vertical component θ_i^A and through the partially persistent horizontal components ε_{ijt}^A. In particular, we assume that the latter evolve over time according to the following process. For each pair $(i, j) \in N_A \times N_B$, and each period $t \geq 1$, ε_{ijt}^A may take any value in $\mathcal{E}_{ijt}^A \subseteq \mathbb{R}$. While not essential to the results, we find it convenient to think of $\mathcal{E}_{ijt}^A$ either as the entire real line, or as a compact and connected subset of it.

The evolution of agent i's horizontal types is governed by a collection of match-specific kernels $G_{ij}^A \equiv (G_{ijt}^A : \mathcal{E}_{ijt}^A \times \mathcal{E}_{ijt-1}^A \times X_{ij}^{t-1} \to [0, 1])_{t=1}^{\infty}$, with $X_{ij}^0 \equiv \varnothing$. We then let $G \equiv (G_{ij}^k)_{(i,j) \in N_A \times N_B}^{k=A,B}$ denote the complete collection of kernels for all matches. The interpretation of these kernels is the following. Each period-1 horizontal type ε_{ij1}^A is drawn from the cdf G_{ij1}^A. In each subsequent period $t > 1$, the period-t type ε_{ijt}^A is drawn from the cdf $G_{ijt}^A(\varepsilon_{ijt}^A \mid \varepsilon_{ijt-1}^A, x^{t-1})$, where ε_{ijt-1}^A is the agent's horizontal match-specific type from the preceding period, and where $x_{ij}^{t-1} = (x_s)_{s=1}^{t-1}$ is the history of past interactions between the pair (i, j). Importantly, while the support of each kernel G_{ijt}^A is a subset of $\mathcal{E}_{ijt}^A$, we allow for the possibility that, for certain histories $(\varepsilon_{ijt-1}^A, x^{t-1})$, it is a strict subset of $\mathcal{E}_{ijt}^A$.

To guarantee that the expected payoff of each agent is well defined at all histories, and satisfies a certain envelope formula (more below), we assume that, for all $i \in N_A$, there exists a constant $E_i^A > 0$ such that, for any sequence of matches x, $\mathbb{E}\left[\sum_{t=1}^{\infty} \delta^t \sum_{j \in N_B} |\varepsilon_{ijt}^A| \cdot x_{ijt}\right] \leq E_i^A$, where the expectation is taken with respect the distribution over $\mathcal{E}$ generated by the kernels G, under the matches x.

We will focus on two environments that capture different features of various dynamic matching markets with private information. The first environment is one in which the evolution of the agents' match values and of the platform's costs is exogenous. The second environment is one in which the match values and the platform's costs depend on past interactions, with properties reflecting either private experimentation or a preference for variety.

• **Exogenous processes.** For all $(i,j) \in N_A \times N_B$, $t > 1$, $\varepsilon^A_{ijt-1} \in \mathcal{E}^A_{ijt-1}$, $G^A_{ijt}(\cdot \mid \varepsilon^A_{ijt-1}, x^{t-1})$ is invariant in x^{t-1}. Furthermore, c_{ijt} does not depend on x^{t-1}.

• **Endogenous processes.** For any $(i,j) \in N_A \times N_B$, the following properties hold: (i) whenever $x_{ijt-1} = 1$, the dependence of $G^A_{ijt}(\varepsilon^A_{ijt} \mid \varepsilon^A_{ijt-1}, x^{t-1})$ on x^{t-1} is only through $\sum_{s=1}^{t-1} x_{ijs}$; (ii) whenever $x_{ijt-1} = 0$, $G^A_{ijt}(\cdot \mid \varepsilon^A_{ijt-1}, x^{t-1})$ is a Dirac measure at $\varepsilon^A_{ijt} = \varepsilon^A_{ijt-1}$, i.e., $G^A_{ijt}(\varepsilon^A_{ijt} \mid \varepsilon^A_{ijt-1}, x^{t-1}) = \mathbf{1}_{\{\varepsilon^A_{ijt} \geq \varepsilon^A_{ijt-1}\}}$; (iii) there exists a sequence $(\omega^A_{ijs})_{s=1}^{\infty}$ drawn from an exogenous distribution, such that, for any number R_{ij} of past interactions between agent $i \in N_A$ and agent $j \in N_B$, the period-t match value ε^A_{ijt} is given by a deterministic function of $(\omega^A_{ijs})_{s=1}^{R_{ij}}$, uniformly over t; (iv) the cost $c_{ijt}(x^{t-1}_{ij})$ depends on x^{t-1}_{ij} only through $\sum_{s=1}^{t-1} x_{ijs}$.

In the case of endogenous processes, the above assumptions imply the following properties: (1) agents' match-specific values change only upon interacting with partners; (2) The processes governing the agents' match values are Markov time-homogeneous — if agents $(i,j) \in N_A \times N_B$ are matched in period $t-1$, the distribution of ε^A_{ijt} depends only on ε^A_{ijt-1} and the number of past interactions between the pair (i,j). These assumptions capture important features of environments in which agents either (a) learn gradually about unknown but constant values for interacting with agents from the other side of the market, or (b) have a preference for variety. That, under endogenous processes, the platform's cost of linking a pair of agents depends on the number of times the pair interacted in the past reflects the idea that, in most cases of interest, costs decrease with the number of past interactions.[13]

Example 1 (Gaussian learning) Suppose that every agent $i \in N_A$ derives a constant value v^A_{ij} from interacting with each agent $j \in N_B$, and that this value is unknown to the platform and to all agents. Agent i starts with a prior belief that $v^A_{ij} \sim N(\varepsilon^A_{ij1}, \tau^A_{ij})$, where the variance τ^A_{ij} is common knowledge but where the initial prior mean ε^A_{ij1} is the agent's private information. The agent's prior mean ε^A_{ij1} is drawn from a distribution G^A_{ij1}. Each time agent i is matched to agent j, agent i receives a conditionally i.i.d. private signal $\xi^A_{ij} \sim N(v^A_{ij}, \vartheta^A_{ij})$ about v^A_{ij} and updates his expectation of v^A_{ij} using standard projection formulae.[14] Let ε^A_{ijt} be agent i's posterior mean about v^A_{ij} in period t. Such environment satisfies the assumptions of the above endogenous-processes model. In particular, when $x_{ijt-1} = 1$, agent i's posterior mean ε^A_{ijt} about v^A_{ij} is drawn from a (Gaussian) distribution $G^A_{ijt}(\cdot \mid \varepsilon^A_{ijt-1}, x^{t-1})$ that depends only on the agent's mean ε^A_{ijt-1} in the previous period and the number of past interactions $\sum_{s=1}^{t-1} x_{ijs}$. □

[13] A case of special interest is when the cost vanishes after the first interaction.

[14] Specifically, each signal ξ^A_{ij} can be written as $\xi^A_{ij} = v^A_{ij} + \varsigma^A_{ij}$ with the innovations ς^A_{ij} drawn from a normal distribution with mean 0 and variance ϑ^A_{ij}, independently from all other random variables.

Example 2 (preference for variety) Suppose that the value each agent $i \in N_A$ derives from interacting with each agent $j \in N_B$ decreases (possibly stochastically) with the number of past interactions with agent j. Precisely, for all $t \geq 1$, all $j \in N_B$, all $(\varepsilon^A_{ijt-1}, \varepsilon^A_{ijt}) \in \mathcal{E}^A_{ijt-1} \times \mathcal{E}^A_{ijt}$, $G^A_{ijt}(\varepsilon^A_{ijt} \mid \varepsilon^A_{ijt-1}, x^{t-1})$ is non-decreasing in $\sum_{s=1}^{t-1} x_{ijs}$. This assumption captures the idea that agents gradually lose interest in partners with whom they interacted already. Alternatively, agents may have a fixed demand for interactions with each partner, or each agent may provide only up to a fixed number of services to each of his partners. Such environments can be captured as special cases of the endogenous-processes model. For example, the case in which each agent has a fixed demand for interactions with each partner can be captured by assuming that, for each agent $i \in N_A$, each partner $j \in N_B$, there exists $\alpha^A_{ij} \in \mathbb{N}$ such that, whenever $\sum_{s=1}^{t-1} x_{ijs} > \alpha^A_{ij}$, $G^A_{ijt}(\varepsilon^A_{ijt} \mid \varepsilon^A_{ijt-1}, x^{t-1}) = 1$ for all $\varepsilon^A_{ijt} \geq 0$. $\square$

Remarks. Given the matches x, agent i's horizontal types $\varepsilon^A_{ij} \equiv (\varepsilon^A_{ijt})_{t=1,..,\infty}$ for interacting with agent j from side B are drawn independently from the values $\varepsilon^A_{ij'} = (\varepsilon^A_{ij't})_{t=1,..,\infty}$ the agent derives from interacting with any other agent $j' \in N_B$. In the case of endogenous processes, such an assumption facilitates the characterization of the optimal mechanism by favoring an index representation of the optimal policy (more below). This assumption, however, can be dispensed with in case of exogenous processes, where types can be allowed to be correlated. That the values agent i derives from interacting with agent j are independent of the values that other agents derive from interacting with the same agent in turn avoids the possibility that the platform trivially extracts the entire surplus from each agent using payments similar to those in Cremer and McLean (1988).

Finally, that the value agent i derives from interacting with agent j is invariant in the composition of agent i's matching set (that is, it does not depend on who else the individual interacts with) is also meant to facilitate the description of the scoring rules in the matching auctions below. This assumption can be dispensed with in the case of exogenous processes (albeit at the cost of an increase in the complexity of the scores), but is more difficult to dispense with in the case of endogenous processes.[15] Importantly, we see both assumptions (type independence and preference separability) as a reasonable description of markets in which the platform lacks detailed information about the agents' precise preference structure.

Capacity constraints

In each period $t \geq 1$, the platform can match at most M agents from opposite sides, independently of past matching allocations. That is, M is a constraint on the stock of existing matches. In each period, the platform can delete some of the previously formed matches and create new ones. We only impose that the total number of existing matches be no greater than M in all periods. As mentioned in the Introduction, such limited capacity may capture various constraints on the platform's side, such as limited facilities, time, or services availability.

[15] The case where each agent values interacting with at most one agent from the opposite side (one-to-one matching) corresponds to a particular relaxation of this assumption.

We will consider both the case in which M is sufficiently large that this constraint never binds (i.e., $M \geq n_A \cdot n_B$), as well as the case $M < n_A \cdot n_B$ in which this constraint may be binding. The set of feasible period-t matches is therefore given by

$$X_t \equiv \left\{ x_t \in \prod_{(i,j)\in N_A \times N_B} X_{ijt} : \sum_{i\in N_A} \sum_{j\in N_B} x_{ijt} \leq M \right\},$$

and is independent of past allocations. The set of sequences of feasible matching allocations is denoted by $X \equiv \prod_{t=1}^{\infty} X_t$.

Matching mechanisms, efficiency, and profit maximization

The interactions between the two sides can be thought of as governed by a *matching mechanism* $\Gamma \equiv (\mathcal{M}, \mathcal{S}, \chi, \psi, \rho)$. The latter consists of: (i) a collection of matching sets $\mathcal{M} \equiv (\mathcal{M}_t)_{t=0}^{\infty}$, with each $\mathcal{M}_t \equiv \prod_{l\in N_k, k=A,B} \mathcal{M}_{lt}^k$ denoting the set of messages the agents can send in period t; (ii) signals $\mathcal{S} \equiv (\mathcal{S}_t)_{t=0}^{\infty}$ that the platform may disclose to the agents, with $\mathcal{S}_t \equiv \prod_{l\in N_k, k=A,B} \mathcal{S}_{lt}^k$; (iii) a *matching rule* $\chi \equiv (\chi_t)_{t=1}^{\infty}$ describing, for each $t \geq 1$, the matches $\chi_t : \mathcal{M}^t \to X_t$ implemented given the history of received messages; (iv) a *payment rule* $\psi \equiv (\psi_t)_{t=0}^{\infty}$ describing, for each $t \geq 0$, the payments $\psi_t : \mathcal{M}^t \to \mathbb{R}^{N_A+N_B}$ asked from, or made to, the agents; and (v) a *disclosure policy* $\rho \equiv (\rho_t)_{t=0}^{\infty}$ specifying the information $\rho_t : \mathcal{M}^t \to \mathcal{S}_t$ disclosed to the agents over time. Each ρ_t must reveal to each agents his own matches. It may also reveal additional information, but it cannot conceal the matches in which the individual is involved.

A matching mechanism Γ is *feasible* if, for any sequence of messages $m \in \mathcal{M}$, the implemented allocations are feasible, that is, $\chi(m) \in X$.[16]

Each agent must choose whether or not to join the platform after observing his vertical type θ_i^A but before observing his horizontal types. As explained above, this assumption reflects the idea that, in most markets of interest, agents choose whether or not to get "on-board" before learning the profile of potential partners. In environments in which the platform maximizes welfare (a possibility we consider below), the information the agents possess at the time they join the platform plays no role (they can be assumed to know only their vertical types θ, or also the period-1 horizontal types ε_1, or to have no private information at all). If, instead, the platform maximizes profits, then the private information the agents possess at the time they join plays a fundamental role. If the agents possess no private information at all, the platform can extract the entire surplus. If, in addition to their vertical types, they also possess private information about their period-1 types ε_1, the optimal mechanism is significantly more complicated, because of the multi-dimensionality of the agents' initial private information. To maintain tractability of the analysis, we thus assume that the agents' private information at the time of joining is limited to the vertical types θ.

[16] Note that the rule χ is deterministic. This is because, in this environment, the platform never gains from inducing random matches, irrespective of whether its objective is profit maximization, or welfare maximization — see the discussion in the proof of Theorem 2 below.

Upon joining the platform, at each period $t \geq 0$, each agent $l \in N_k$, from each side $k = A, B$, is asked to send a message m_{lt}^k from the set $\mathcal{M}_{lt}^k$. In the auctions we introduce in the next section, such messages correspond to the selection of a membership status along with a collection of bids, one for each partner. Note that, while the game starts in period 0, the actual matching begins in period 1.

We use perfect Bayesian equilibrium (PBE) as our solution concept. A matching mechanism will be referred to as *profit-maximizing* if it is feasible and admits a PBE such that, under such equilibrium, the platform's profits are at least as high as under any other PBE of any other feasible mechanism. A *welfare maximizing* mechanism is defined in a similar way, but with welfare replacing profits in the platform's objective.

Remark. The equilibria in our dynamic matching auctions satisfy properties stronger than those required by PBE. In particular, the equilibrium strategies remain optimal no matter the information each agent may possess about other agents' past messages and match allocations, and no matter the beliefs the agent may have about other agents' past and current types. Consistently with the rest of the dynamic mechanism design literature, we will refer to such equilibria as *periodic ex-post* (see, e.g., Bergemann and Valimaki (2010), Athey and Segal (2013), and Pavan, Segal, and Toikka (2014)). As a result, both the profit-maximizing and the welfare-maximizing mechanisms we propose can be made *fully transparent*, in the sense that all bids, matches, and payments are made public after each period. Lastly, while we only require that the mechanisms induce participation in period zero, in the specific mechanisms we propose, agents find it optimal to remain in the mechanism after each history.

3 Matching auctions

We now introduce a class of mechanisms, which we refer to as "matching auctions," in which agents bid repeatedly to be matched with potential partners. The structure is the following.

Definition 1 (matching auctions) *A* matching auction *is a mechanism with the following properties:*

- *At period $t = 0$, upon joining the platform, each agent $l \in N_k$, from each side $k = A, B$, is asked to select a membership status $\theta_{l0}^k \in \Theta_l^k$. Each status is conveniently indexed by the vertical type it is designed for. A higher status translates into a more favorable treatment, on average, in the subsequent auctions. Accordingly, a higher status comes with a higher fee $p_{l0}^k \in \mathbb{R}$.*

- *In each subsequent period $t \geq 1$, each agent $l \in N_k$, from each side $k = A, B$, is first offered the possibility to revise his membership status by selecting $\theta_{lt}^k \in \Theta_l^k$. Each agent is then invited to submit bids $b_{lt}^k \equiv (b_{ljt}^k)_{j \in N_{-k}}$, one for each potential partner from the opposite side. Each pair of agents $(i, j) \in N_A \times N_B$ is then assigned a "score" $S_{ijt} \in \mathbb{R}$ that depends on the pair's period-0 and period-t status, the pair's reciprocal bids (b_{ijt}^A, b_{ijt}^B) and, possibly, the number of times the pair interacted in the past. All pairs of agents with the highest nonnegative score are matched, up to capacity. All agents who are unmatched in period t, pay nothing. Matched agents make (or receive) payments $p_{lt}^k \in \mathbb{R}$ that may depend also on other agents' bids and membership status.*

- *All past bids, payments, membership choices, and matches are public.*

The details of the scoring rules and of the corresponding payment functions naturally depend on the environment under examination and, in particular, the platform's capacity constraint and the nature of the processes governing the evolution of the agents' match values and the platform's costs.

3.1 Scoring rules

A *scoring rule* S consists of a sequence of mappings $(S_{ijt})_{t=1,\ldots,\infty}^{(i,j)\in N_A \times N_b}$ assigning to each pair of agents from opposite sides $(i,j) \in N_A \times N_B$, each $t \geq 1$, a score $S_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) \in \mathbb{R}$ that depends on the pair's period-0 and period-t membership status, current bids, and number of past interactions. As mentioned above, each score S_{ijt} depends only on information pertaining to the pair (i,j). The conditioning of S_{ijt} on the entire profile $(\theta_0, \theta_t, b_t, x^{t-1})$ is only to facilitate the notation.

Formally speaking, the matching auctions defined above correspond to a matching mechanism in which the message spaces are given by $\mathcal{M}_{i0}^A = \Theta_i^A$ for $t = 0$, and, for any $t \geq 1$, by $\mathcal{M}_{it}^A = \Theta_i^A \times \mathbb{R}^{N_B}$, and where the matching rule χ matches in each period $t \geq 1$ all pairs with the highest nonnegative score $S_{ijt}(\theta_0, \theta_t, b_t, x^{t-1})$, up to capacity, with ties broken arbitrarily. Obviously, the history of past matches $x^{t-1} = \chi^{t-1}(\theta^{t-1}, b^{t-1})$ is itself determined by past bids and membership choices. With an abuse of notation, whenever there is no risk of confusion, hereafter we drop the dependence of x^{t-1} on past bids and membership choices to ease the notation.

Now let

$$Q_t^S(\theta_0, \theta_t, b_t, x^{t-1}) \equiv \left\{ \begin{array}{l} (i,j) \in N_A \times N_B \text{ s.t. (i) } S_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) \geq 0 \text{ and} \\ \text{(ii) } \# \left\{ (l,m) \in N_A \times N_B : S_{lmt}(\theta_0, \theta_t, b_t, x^{t-1}) > S_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) \right\} < M \end{array} \right\}$$

denote the set of period-t matches $(i,j) \in N_A \times N_B$ for which the period-t score is nonnegative and for which there are at most $M-1$ matches with a strictly higher score. The matching rule χ corresponding to the scoring rule S then satisfies the following properties: (a) $(i,j) \in Q_t^S(\theta_0, \theta_t, b_t, x^{t-1})$ whenever $\chi_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) = 1$; (b) if $S_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) > 0$ and

$$\# \left\{ (l,m) \in N_A \times N_B : S_{lmt}(\theta_0, \theta_t, b_t, x^{t-1}) \geq S_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) \right\} \leq M,$$

then $\chi_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) = 1$; and (c) if $(i,j) \in Q_t^S(\theta_0, \theta_t, b_t, x^{t-1})$, $S_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) > 0$, and $\chi_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) = 0$, then $\# \left\{ (l,m) \in N_A \times N_B : \chi_{lmt}(\theta_0, \theta_t, b_t, x^{t-1}) = 1 \right\} = M$. In words, if the pair (i,j) is matched in period t, it must be that its period-t score is nonnegative and that there are at most $M-1$ matches with a period-t score strictly higher than the one for (i,j). Furthermore, the pair (i,j) is *necessarily* matched in period t if its period-t score is strictly positive, and there are at most M matches with a period-t score weakly higher than that for (i,j). Lastly, condition (c) says that if there are no more than $M-1$ matches with a score strictly higher than that of (i,j), (i,j)'s score is strictly positive, and the match (i,j) is not formed, then it must be that capacity is fully utilized and that the platform is matching other pairs with a score identical to (i,j)'s.

Importantly, agents' bids are allowed to be negative. An agent submitting a negative bid for a potential partner may nonetheless be matched with that partner if the platform finds it optimal to cross-subsidize the interaction. In considering the profitability of such cross-subsidization, the platform may also take into account the value of generating information about the profitability of such a match, which can be used in future periods.

We now introduce two scoring rules that play a central role in the analysis.

Definition 2 *A myopic scoring rule $S^{m;\beta}$ (with weights β) is one in which, for each $t \geq 1$, each pair $(i, j) \in N_A \times N_B$, the score is given by*

$$S^{m;\beta}_{ijt}\left(\theta_0, \theta_t, b_t, x^{t-1}\right) \equiv \beta^A_i(\theta^A_{i0}) \cdot b^A_{ijt} + \beta^B_j(\theta^B_{j0}) \cdot b^B_{ijt} - c_{ijt}(x^{t-1}), \tag{3}$$

where $\beta \equiv (\beta^k_l)_{l \in N_k, k=A,B}$ are time-invariant, non-decreasing, strictly positive, and bounded functions of the period-0 membership statuses. A myopic matching rule $\chi^{m;\beta}$ is a matching rule in which the scoring rule is given by $S^{m;\beta}$.

A myopic score is thus a weighted average of the bids b^A_{ijt} and b^B_{ijt} submitted by the pair $(i, j) \in N_A \times N_B$ in period t, $t \geq 1$, where the weights are given by strictly positive and non-decreasing functions of the agents' period-0 membership statuses, net of the platform's period-t cost of matching the pair. Note that myopic scores may depend on past bids only through the induced past allocations x^{t-1}, and are invariant in all membership choices except the period-0 ones. In the case of exogenous processes, because costs do not depend on past allocations, the myopic scores are also invariant in past bids.

The next scoring rule is a forward-looking one, where the scores take into account not only the flow surpluses generated by the current interactions, but also the value of generating information that can be used in future periods. Let $\lambda_{ij}|\theta_t, b_t, x^{t-1}$ denote the stochastic process over the pair (i, j)'s current and future match values (u^A_{ijs}, u^B_{ijs}) and costs c_{ijs} that one obtains under any matching rule that matches the pair (i, j) at all periods $s \geq t$, when the pair's vertical types (θ^A_i, θ^B_j) are the ones specified in the period-t vector θ_t, and when the horizontal types are given by

$$\varepsilon^k_{ijt} = \begin{cases} b^k_{ijt}/\theta^k_{it} & \text{if } b^k_{ijt}/\theta^k_{it} \in \mathcal{E}^k_{ijt} \\ \arg\min\limits_{\hat{\varepsilon}^k_{ijt} \in \mathcal{E}^k_{ijt}} \left\{ |b^k_{ijt}/\theta^k_{it} - \hat{\varepsilon}^k_{ijt}| \right\} & \text{otherwise,} \end{cases} \tag{4}$$

$k = A, B$.

Definition 3 *An index scoring rule $S^{I;\beta}$ (with weights β) is one in which, for each $t \geq 1$, each pair $(i, j) \in N_A \times N_B$,*

$$S^{I;\beta}_{ijt}\left(\theta_0, \theta_t, b_t, x^{t-1}\right) \equiv \max_{\tau} \left\{ \frac{\mathbb{E}^{\lambda_{ij}|\theta_t, b_t, x^{t-1}} \left[\sum_{s=t}^{\tau} \delta^{s-t} \left(\beta^A_i(\theta^A_{i0}) u^A_{ijs} + \beta^B_j(\theta^B_{j0}) u^B_{ijs} - c_{ijs} \right) \right]}{\mathbb{E}^{\lambda_{ij}|\theta_t, b_t, x^{t-1}} \left[\sum_{s=t}^{\tau} \delta^{s-t} \right]} \right\}, \tag{5}$$

where τ denotes a stopping time, and where $\beta \equiv (\beta^k_l)_{l \in N_k, k=A,B}$ are time-invariant, non-decreasing, strictly positive, and bounded functions of the period-0 membership statuses. An index matching rule $\chi^{I;\beta}$ is a matching rule in which the scoring rule is given by $S^{I;\beta}$.

The score $S^{I;\beta}_{ijt}$ thus corresponds to a Gittins index for a process for which the "rewards" are given by the weighted total surplus of the match (i,j), with the weights given by the functions β of the agents' period-0 membership statuses. Note that, under our assumptions for endogenous processes, each score $S^{I;\beta}_{ijt}$ depends on x^{t-1} only through the total number of past interactions $\sum_{s=1}^{t-1} x_{ijs}$ between the pair (i,j). Also note that, contrary to the myopic scores, the indexes S^{I}_{ijt} depend on the entire history of past and current membership choices. However, while the dependence on the period-0 and on the period-t choice is direct, the dependence on other periods' choices is only via the number of past interactions.

3.2 Role of the membership status

Intuitively, each agent's membership status determines the importance the platform assigns to the agent's bids, relative to those of others. Consider, for example, an environment with exogenous processes, and suppose that matches are determined according to the myopic rule $\chi^{m;\beta}$. Suppose, in period $t \geq 1$, agent i from side A submits a positive bid for agent j from side B, whereas agent j submits a negative bid. For given bids (b^A_{ijt}, b^B_{ijt}), a higher period-0 status of agent i implies a higher score for the match (i,j), thus tilting the matching allocation in favor of agent i. Symmetrically, a higher status for agent j reduces the score for the match (i,j), thus tilting the allocation in favor of agent j. A higher period-0 status thus grants an agent preferential treatment in all subsequent auctions, both with respect to the competition the agent faces with other agents from his own side (when the capacity constraint is binding) and with respect to the competition the agent faces with agents from the opposite side, shifting the length of each interaction to his benefit.

The role of the agents' status in subsequent periods (that is, θ^A_{it}, with $t \geq 1$), is somewhat different. It permits the platform to decouple the agents' horizontal types ε^A_{ijt} from the agents' bids, which is useful to form expectations about the agents' future match values. In fact, as we show below, while the agents' bids convey information about the agents' total match values u^A_{ijt} for each possible interaction, they are not sufficient statistics with respect to the agents' overall private information, when it comes to predicting future values. For instance, a high bid b^A_{ijt} by agent i for agent j may either reflect a persistent high value for interacting with all agents from the opposite side (i.e., a high vertical type θ^A_i), or a high temporary appreciation for interacting with agent j (i.e., a high horizontal type, ε^A_{ijt}). Accordingly, when the evolution of the agents' match values is endogenous, the platform uses the combination of the submitted bids with the membership statuses to trade off the value of current matches with the value of generating information useful in future periods.

Finally, note that the reason why the platform allows the agents to revise their status over time, despite the fact that, under the specification in (1), the vertical types are perfectly persistent, is that this favors equilibria in which the agents bid myopically over time, irrespective of past bids and membership selections, in the following sense:

Definition 4 *A strategy profile $\sigma = (\sigma^k_l)^{k=A,B}_{l\in N_k}$ for the above matching auctions is* truthful *if for all $t \geq 0$, all private histories, the selection of the membership status is given by $\theta^k_{lt} = \theta^k_l$, and the* *submitted bids are given by* $b^k_{ijt} = u^k_{ijt} = \theta^k_l \cdot \varepsilon^k_{ijt}$*, all* $(i,j) \in N_A \times N_B$*,* $k = A, B$*. A truthful equilibrium is an equilibrium in which the strategy profile is truthful.*

That is, under truthful strategies, irrespective of the history, agents select a membership status equal to their true vertical type, and then submit bids equal to the myopic match value they assign to each interaction.

3.3 Payments

We now complete the description of the matching auctions by describing the associated payment rules. Let χ be the matching rule corresponding to the scoring rule S. For any $t \geq 1$, any $(\theta_0, \theta_t, b_t, x^{t-1})$, any weights $\beta \equiv (\beta^k_l)_{l \in N_k, k=A,B}$, let

$$W_t(\theta_0, \theta_t, b_t, x^{t-1}; \beta) \equiv \mathbb{E}^{\lambda[\chi]|\theta_t, b_t, x^{t-1}} \left[\sum_{s=t}^{\infty} \delta^{s-t} \sum_{i \in N_A} \sum_{j \in N_B} S^{m;\beta}_{ijs}(\theta_0, \theta_s, b_s, x^{s-1}) \cdot \chi_{ijs}(\theta_0, \theta_s, b_s, x^{s-1}) \right] \tag{6}$$

denote the continuation weighted surplus. Here $\lambda[\chi]|\theta_t, b_t, x^{t-1}$ denotes the stochastic process over (θ_s, b_s, x^{s-1}), $s \geq t$, under the rule χ, when the selected period-t membership statuses are θ_t, the period-t bids are b_t, all agents follow truthful strategies from period $s > t$ onwards, the true vertical types are the ones corresponding to the selected period-t statuses (i.e., $\theta^k_l = \theta^k_{lt}$, $l \in N_k$, $k = A, B$), and the true period-t horizontal types are given by (4).[17]

Similarly, let $W^{-l,k}_t(\theta_0, \theta_t, b_t, x^{t-1}; \beta)$ denote the continuation weighted surplus, as defined in (6), when the myopic scores involving agent l from side k are identically equal to zero at all periods $s \geq t$.

Next, let

$$R^k_{lt}(\theta_0, \theta_t, b_t, x^{t-1}; \beta) \equiv W_t(\theta_0, \theta_t, b_t, x^{t-1}; \beta) - W^{-l,k}_t(\theta_0, \theta_t, b_t, x^{t-1}; \beta),$$

denote the *contribution* of agent $l \in N_k$ to the continuation weighted surplus and

$$r^k_{lt}(\theta_0, \theta_t, b_t, x^{t-1}; \beta) \equiv R^k_{lt}(\theta_0, \theta_t, b_t, x^{t-1}; \beta) - \delta \mathbb{E}^{\lambda[\chi]|\theta_t, b_t, x^{t-1}} \left[R^k_{lt+1}(\theta_0, \theta_{t+1}, b_{t+1}, x^t; \beta) \right] \tag{7}$$

the corresponding *flow marginal contribution.*

In each period $t \geq 1$, the payment asked to each agent $i \in N_A$ (an analogous construction holds for each agent from side B) is given by

$$\psi^A_{it}(\theta_0, \theta_t, b_t, x^{t-1}; \beta) = \sum_{j \in N_B} b^A_{ijt} \cdot \chi_{ijt}(\theta_0, \theta_t, b_t, x^{t-1}) - \frac{1}{\beta^A_i(\theta^A_{i0})} r^A_{it}(\theta_0, \theta_t, b_t, x^{t-1}; \beta). \tag{8}$$

In words, each agent i is asked to make a payment that equals the total flow value the agent derives from all the matches implemented in period t in which the agent is involved, net of a discount that is proportional to the agent's flow marginal contribution to weighted surplus, with a coefficient of proportionality given by $1/\beta^A_i(\theta^A_{i0})$.

[17] Note that we are dropping the argument χ from the W_t functions, to ease the exposition.

The description of the payments is then completed by the specification of the fees charged to the agents in period zero, as a function of the profile of selected membership statuses. The latter are given by

$$\psi_{l0}^k(\theta_0;\beta) = \theta_{l0}^k D_l^k(\theta_0) - \int_{\underline{\theta}_l^k}^{\theta_{l0}^k} D_l^k(\theta_{-l}^k, y)dy - \mathbb{E}^{\lambda[\chi]|\theta_0}\left[\sum_{t=1}^{\infty} \delta^t \psi_{lt}^k(\theta_0, \theta_t, b_t, x^{t-1};\beta)\right] - L_l^k, \tag{9}$$

where L_l^k is a scalar whose role is to guarantee that the agent's participation constraint is satisfied, $\theta_{-l}^k \equiv (\theta_j^m)_{j\in N_m,(m,j)\neq(k,l)}^{m\in\{A,B\}} \in \Theta_{-l}^k \equiv \prod_{j\in N_m,(m,j)\neq(k,l),}^{m\in\{A,B\}} \Theta_j^m$ is a profile of vertical types excluding agent l from side k, and

$$D_l^k(\theta_0) \equiv \begin{cases} \mathbb{E}^{\lambda[\chi]|\theta_0}\left[\sum_{t=1}^{\infty} \delta^t \sum_{h\in N_{-k}} \varepsilon_{lht}^k \chi_{lht}(\theta_0, \theta_t, b_t, x^{t-1})\right] & \text{if } k = A \\ \mathbb{E}^{\lambda[\chi]|\theta_0}\left[\sum_{t=1}^{\infty} \delta^t \sum_{h\in N_{-k}} \varepsilon_{hlt}^k \chi_{hlt}(\theta_0, \theta_t, b_t, x^{t-1})\right] & \text{if } k = B \end{cases} \tag{10}$$

is the "quality" of the matches agent l from side k expects from all subsequent auctions, when the profile of selected period-0 membership statuses is θ_0 (here $\lambda[\chi]|\theta_0$ denotes the distribution over future membership statuses, bids, and matches, under truthful strategies). Note that the reason 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_{ijt}, as well as the order in the subscripts in the horizontal types ε_{ijt}^k is not permutable: the first index always refers to side A, while the second to side B. Also note that, according to (9), the fee the platform charges to each agent depends on the entire profile of membership statuses selected by all agents. In other words, the price for status depends on the aggregate demand of status from each side of the market. Such dependence plays a role analogous to that of "insulating tariffs" in two-sided markets (see, e.g., Weyl, 2010): it guarantees that all agents find it optimal to join the platform in period zero, irrespective of their beliefs about the types of other agents. Such insulating tariffs can always be replaced by (perhaps more familiar) tariffs in which membership fees depend only on each agent's own status, by taking expectations over the types of other agents.

The period-t payments, $t \geq 1$, on the other hand, reflect the externalities the agents impose both in the current and in future periods on all other agents (from both sides of the market), as well as the costs they impose on the platform. Such externalities are calculated with respect to the "weighted surplus" associated with each of the matches, with the weights determined by the period-0 membership statuses. Note that 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 – cross subsidization.

As for the membership fees, because higher status implies higher average match quality in all future auctions (a property we establish formally below), higher status naturally comes with a higher fee. In particular the period-0 membership fees are equal to the agent's expected match quality, net of the payments the agent expects to make in the subsequent auctions, and net of a discount $\int_{\underline{\theta}_l^k}^{\theta_{l0}^k} D_l^k(\theta_{-l}^k, y)dy$ whose role is to incentivize the agent to select the status corresponding to his true vertical type.

Example 3 (payments - exogenous processes) To illustrate, suppose the processes are exogenous. If $i \in N_A$ is not matched to any agent $j \in N_B$ in period t, he does not make any payments in that period. Otherwise, i's payment depends on the total number of matches he secures in the period-t auction. Specifically, if i is matched to $K \geq 1$ agents from side B, his payment is equal to

$$\psi_{it}^A(\theta_0, \theta_t, b_t, x^{t-1}; \beta) = \frac{1}{\beta_i^A(\theta_{i0}^A)} \left(\mathbf{B}_{it}^A(K) + \sum_{j \in N_B} (c_{ijt} - \beta_j^B(\theta_{j0}^B) b_{ijt}^B) \chi_{ijt}^m \right),$$

where $\mathbf{B}_{it}^A(K)$ is the sum of the K highest nonnegative myopic scores among the pairs that are unmatched in period t and that do not include agent i. Similar payments apply to each side-B agent.

Note that, in this environment, the period-t payments depend only on the period-0 membership choices, the period-t bids, and the platform's costs. Interestingly, contrary to the payments in standard Vickrey-type auctions such as the GSP auctions, multiple agents (from both sides of the market) may be charged for the same externality they impose on others. $\square$

The payments for the case of endogenous processes differ based on the capacity constraint in place. When the capacity constraint is binding, payments take into account also the expected future externalities under the processes induced by the agent's period-t matches. Notwithstanding this difference, the period-t payments under endogenous processes are also independent of all past bids and membership choices, except for the period-0 choices θ_0 and for the effect of past bids and membership choices on the number of past interactions.[18]

4 Equilibrium

We now turn to the equilibria of the matching auctions introduced above. We start by describing a special class of environments that plays a role in case of endogenous processes. Let $S^{I;\beta}$ denote the index scoring rule defined by the weights β. Then let

$$\underline{S}_{ijt}^{\beta}(\theta^t, u^t, x^{t-1}) \equiv \inf_{s \leq t} \left\{ S_{ijs}^{I;\beta}(\theta_0, \theta_s, u_s, x^{s-1}) \right\} \tag{11}$$

denote the historical infimum of the indexes for the pair (i, j), when the period-t history of bids is given by $b^t = u^t \equiv (u_s)_{s=1}^t$.[19]

Definition 5 (separability) *The environment is separable under the rule $S^{I;\beta}$ if, for any $t \geq 1$, any two pairs of agents, $(i, j), (i', j') \in N_A \times N_B$, any (θ^t, u^t, x^{t-1}),*

$$\underline{S}_{ijt}^{\beta}(\theta^t, u^t, x^{t-1}) > \underline{S}_{i'j't}^{\beta}(\theta^t, u^t, x^{t-1}) > 0 \Rightarrow \underline{S}_{ijt}^{\beta}(\theta^t, u^t, x^{t-1}) \cdot (1 - \delta) \geq \underline{S}_{i'j't}^{\beta}(\theta^t, u^t, x^{t-1}). \tag{12}$$

Note that the definition imposes a restriction on the initial heterogeneity of the positive indexes as well as on the way such indexes decrease over time. It imposes no restriction on the way the indexes respond to positive shocks, or on the way they evolve once they turn negative.

[18] One might expect the payments under endogenous processes to be analogous to those under exogenous processes, but with the indexes $S_{ijt}^{I;\beta}$ replacing the myopic scores $S_{ijt}^{m;\beta}$. However, this is not correct. In fact, such payments are greater than those defined by the functions ψ above and would lead agents to shade their bids.

[19] Such auxiliary processes are also referred to as *lower envelope processes* (see, e.g., Mandelbaum 1986).

Example 4 (bad news) Consider an environment in which the period-1 indexes $S^{I;\beta}_{ij1}$ are sufficiently heterogeneous, in the sense of (12). Further assume that, for each pair of agents (i, j), in each period $t \geq 2$, either $S^{I;\beta}_{ijt} \geq S^{I;\beta}_{ijt-1}$ or $S^{I;\beta}_{ijt} < 0$. The environment is then separable. The above structure emerges, for example, in the familiar "no-news-is-good-news" class of environments considered in the experimentation literature (e.g., Bonatti and Horner (2015), Keller and Rady (2015)).[20] In such environments, in each period, agents either receive no news or negative news revealing the unprofitability of the match. In the absence of bad news, agents become increasingly more optimistic about the profitability of their relationships.[21] □

We then have the following result:

Theorem 1 (equilibrium) (*i*) *Suppose the processes are exogenous. Any matching auction in which (1) the scoring rule is myopic with weights β, and (2) the payments are given by (8) and (9), with the same weights β as in the scoring rule, and with L^k_l large enough, all $l \in N_k$, $k = A, B$, admits an equilibrium in which all agents participate in each period and follow truthful strategies.*

(*ii*) *Suppose processes are endogenous and assume that either (a) $M = 1$, or (b) $M \geq n_A \cdot n_B$, or (c) $1 < M < n_A \cdot n_B$ and, in this latter case, the environment is separable under the rule $S^{I;\beta}$. The matching auction in which (1) the scoring rule is the index rule with weights β, and (2) the payments are given by (8) and (9) with the same weights β as in the scoring rule, and with L^k_l large enough, all $l \in N_k$, $k = A, B$, admits an equilibrium in which all agents participate in each period and follow truthful strategies.*

(*iii*) *The equilibria of the auctions in parts* (*i*) *and* (*ii*) *above are* periodic ex-post*; that is, the agents' strategies are sequentially rational, regardless of the agents' beliefs about other agents' past and current types.*

The formal proof is in the Appendix. Here we illustrate the key ideas in an heuristic way. We organize the arguments in three steps. The first step shows that, in the continuation game that starts with period $t \geq 1$, independently of past behavior, when the agents follow truthful strategies, the matches implemented under the myopic and the index rules maximize continuation weighted surplus, as defined in (6), in the respective environments of Theorem 1. Step 2 in turn establishes that, when the payments are the ones in (8) above, in the continuation game that starts with period $t \geq 1$, irrespective of past behavior, all agents have incentives to bid truthfully their myopic values u^k_{ijt}. Step 3 completes the proof by showing that, when the period-0 membership fees are the ones in (9), all

[20] See Keller, Rady and Cripps (2005) for the opposite case of fully revealing "breakthroughs" and Horner and Skrzypacz (2016) for a recent survey of the literature on experimentation in strategic environments.

[21] A similar structure obtains in environments in which, in each period and for each match, with some positive probability (possibly history dependent), the match is revealed to be unprofitable, in which case the index turns negative, while with the complementary probability the index either increases or decreases by a small amount that does not violate (12).

agents find it optimal to join the platform in period zero, and select the period-0 membership status corresponding to their true vertical types, again, regardless of their beliefs about other agents' types.

Step 1. When the processes are exogenous, that the matches implemented under truthful strategies maximize continuation weighted surplus follows directly from the fact that a myopic rule (with weights β) selects in each period the matches with the highest nonnegative myopic score.

The result for endogenous processes is more delicate. In this case, the problem of maximizing continuation weighted surplus is a *multiarmed bandit problem.* Such a problem admits an index-policy solution when either the capacity constraint is not binding ($M \geq n_A \cdot n_B$), or when the platform can match at most one pair of agents in each period ($M = 1$). In both such cases, continuation weighted surplus is maximized by selecting in each period the matches with the highest nonnegative index, up to capacity.

For intermediate capacity constraints ($1 < M < n_A \cdot n_B$), however, without further restrictions on the environment, there is no guarantee that an index policy maximizes continuation weighted surplus. This is because, *in general*, multiarmed bandit problems in which multiple arms can be activated simultaneously fail to admit a simple index solution. The following example, in which the separability condition in Definition 5 is violated, illustrates the problems in the context of our matching environment.

Example 5 (suboptimality of $\chi^{I;\beta}$ for $1 < M < n_A \cdot n_B$) Suppose $N_A = \{1, 2, 3\}$, $N_B = \{1\}$, and $M = 2$. Assume that the myopic score $S^{m;\beta}_{ijt}$ of each match (i, j) evolves according to the following process, where n denotes here the number of past interactions:

	$n=0$	$n=1$	$n=2$	$n \geq 3$
$S^{m;\beta}_{11}$	3	3	3	-1
$S^{m;\beta}_{21}$	4	2	2	-1
$S^{m;\beta}_{31}$	5	1	1	-1

Since each $S^{m;\beta}_{ijt}(n)$ decreases with n, the associated index rule (as defined in Definition 3) is myopic, and hence implements the matches $(2, 1)$ and $(3, 1)$ in period 1, the matches $(1, 1)$ and $(2, 1)$ in period 2 and again in period 3, the matches $(1, 1)$ and $(3, 1)$ in period 4, the match $(3, 1)$ in period 5, and no match from period 6 onwards.[22] Consider now an alternative rule, implementing the matches $(1, 1)$ and $(3, 1)$ in period 1, the matches $(1, 1)$ and $(2, 1)$ in period 2 and again in period 3, the matches $(2, 1)$ and $(3, 1)$ in period 4, the match $(3, 1)$ in period 5, and no match thereafter. The difference in the corresponding profits is $1 - 2\delta + \delta^3$, which is negative for $\delta > \frac{\sqrt{5}-1}{2}$. Thus, for all $\delta > \frac{\sqrt{5}-1}{2}$,

[22] More generally, suppose processes are endogenous and match quality deteriorates over time, in the sense that, for all $t \geq 1$, $(i, j) \in N_A \times N_B$, $k = A, B$, $x^{t-1} \in X^{t-1}$, whenever $x_{ijt} = 1$, then $\varepsilon^k_{ijt+1} \leq \varepsilon^k_{ijt}$ a.s. and $c_{ijt+1} \geq c_{ijt}$. In such environment, irrespective of the weights β, for any pair $(i, j) \in N_A \times N_B$, any $t \geq 1$, $S^{I;\beta}_{ijt} = S^{m;\beta}_{ijt}$. This can be seen by noting that the optimal stopping time in (5) satisfies the property $\tau_{ijt} = \inf\{s > t \mid S^{I;\beta}_{ijs} \leq S^{I;\beta}_{ijt}\}$. That is, it is the first time at which the process of $S^{m;\beta}_{ij}$ reaches a state in which $S^{I;\beta}_{ij}$ drops (weakly) below its period-t value.

an index rule does not maximize weighted surplus. Finally, note that, in this example, separability is violated for $\delta > \frac{\sqrt{5}-1}{2}$. $\square$

In the case of intermediate capacity constraints, the result in Theorem 1 above thus assumes the environment is separable under the rule $S^{I;\beta}$, in the sense of Definition 5. As we show in the Appendix, the role of the separability condition is to guarantee that the optimal rule is myopic in a fictitious environment in which the reward processes are the auxiliary ones, $\underline{S}^{\beta}_{ijt}$, as defined in (11). The proof then proceeds by showing that, given any matching rule χ, continuation weighted surplus is weakly higher when the rewards are given by the auxiliary processes (i.e., $\underline{S}^{m;\beta}_{ijt}$) than when the rewards are the primitive ones (i.e., $S^{m;\beta}_{ijt}$). Furthermore, in the special case in which the scoring rule is the index rule (for the primitive environment), continuation weighted surplus is the same in the two environments.

We then show that, when the environment is separable, a myopic rule maximizes continuation weighted surplus when the rewards are given by the auxiliary processes. Recall that such processes drift downwards. The maximal continuation surplus that can be expected from each match is thus equal to the "annuity" $\underline{S}^{\beta}_{ijt}/(1-\delta)$ of its current auxiliary value. Separability then implies that it is always optimal to select the matches for which the current auxiliary reward is the highest.

That the matches implemented by the index rule maximize continuation weighted surplus in the primitive environment then follows from the above properties along with the fact that the matches implemented by such rule in the primitive environment coincide with those implemented by the myopic rule in the fictitious environment.

Step 2. The second step of the proof shows that, when the payments are the ones specified in (8) above, starting from each period $t \geq 1$, and irrespective of past behavior, all agents have incentives to bid truthfully their myopic values $b^k_{ijt} = u^k_{ijt}$ in the continuation game that starts with period t. This step combines arguments from Bergemann and Valimaki (2010) (see also Kakade et al. (2013)) with arguments from Pavan, Segal, and Toikka (2014). A notorious difficulty in dynamic incentives problems is the need to control for multi-period contingent deviations. In matching auctions, agents who deviated from a truthful strategy in the past may wish to do so again at present and/or in future periods. An additional difficulty in the matching environment under consideration here is the multi-dimensionality of the new private information each agent receives in each period.

Given the profile of period-0 membership choices, the payments in (8) are designed to make each agent's continuation payoff (net of the payments) proportional to his flow marginal contribution to *weighted surplus* (with the latter defined as in (7), and with the coefficient of proportionality given by $1/\beta^k_l(\theta^k_{l0})$). Because the matches under the proposed scoring rules maximize the continuation weighted surplus after all histories (including those off the equilibrium path), agents have incentives to go back to the truthful strategies (and remain in the mechanism) after all histories, and irrespective of the beliefs they may have about past and current types of all other agents. In other words, participating and then following truthful strategies constitutes a periodic ex-post continuation equilibrium, after any history. Importantly, when processes are endogenous, it is important that agents have the possibility to adjust their membership status at the beginning of any period. This helps establishing the sequential optimality of truthful strategies by guaranteeing that, irrespective of past behavior, the matches implemented in the continuation game under truthful strategies maximize the continuation weighed surplus, at all histories.[23]

Step 3. The final step completes the proof by showing that, when the membership fees are the ones in (9), all agents find it optimal to join the platform in period zero, and select the membership status designed for their true vertical type, regardless of the agents' beliefs about other agents' types.

Specifically, the period-0 payments are added to the payments in the subsequent periods so that the payoff that each agent obtains in equilibrium satisfies a certain envelope condition, which relates the agents' period-0 interim expected payoffs to a measure of their expected discounted future "match quality," as defined in (10). The proof then shows that the myopic and index rules satisfy an "average monotonicity" property, according to which match quality increases both with each agent's period-0 membership status (for fixed true vertical type) and with the agent's true vertical type (under truthful strategies). Along with the fact that the equilibrium payoffs satisfy the aforementioned envelope formula, such "average monotonicity" property guarantees that truthful strategies are optimal also in period zero.

Importantly, the quality of an agent's interactions need not increase with his period-0 membership status, or with his true vertical type, in each state of the world (if fact, it does not, under the index rule[24]). It suffices that it increases, on average, where the averaging is across time and states.[25] ■

In environments in which agents have a fixed demand for interactions with each potential partner (as in Example 2 above), the separability condition in part (ii) of the theorem can be relaxed. More generally, separability is sufficient for the index rule to maximize continuation weighted surplus, and hence for the agents to find it optimal to follow truthful strategies in the proposed matching auctions, but it is by no means necessary.[26]

Also, note that the periodic ex-post nature of the equilibria of Theorem 1 is what justifies making the auctions fully transparent.

[23] A similar approach of enlarging the message space so as to give agents the possibility to reveal their true types after possible deviations in previous periods is used in Doepke and Townsend (2006) and Kakade et al. (2013).

[24] This is because an increase, for example, in agent i's period-0 membership status may lead to reversals in the ordering of the agent's indexes. Because these indexes are forward looking, such reversals, while optimal based on the period-t information, might not be optimal ex-post. As a result, the monotonicity of an agent's discounted sum of current and future horizontal types under an index rule need not hold ex-post. That is, along certain paths, higher membership status may result in lower match quality.

[25] The above condition can be viewed as the analog of Pavan, Segal and Toikka (2014)'s average monotonicity condition in a setting, the present one, in which agents receive multi-dimensional new private information in each period. In fact, in the present environment, the horizontal types $\varepsilon_{it}^A \equiv (\varepsilon_{ijt}^A)_{j\in N_B}$ correspond to the "impulse responses" of the period-t match values $u_{it}^A \equiv (u_{ijt}^A)_{j\in N_B}$ to the vertical type θ_i^A, as defined in Pavan, Segal, and Toikka (2014).

[26] To the best of our knowledge, a condition that is jointly necessary and sufficient for the optimality of index policies in multi-armed bandit problems in which an arbitrary number of arms is activated in each period remains elusive, except in special environments.

As the theorem illustrates, one of the appeals of the proposed auctions is that they admit simple equilibria in which agents bid their myopic match values in each period. In other words, the agents do not need to know how to solve complex dynamic-programming problems, or be able to compute the indexes (although, nowadays, there is software that does so). Once the scoring and payments rules are understood, it is in the agents' interest to bid "straight-forwardly" in all periods, irrespective of the beliefs they may have about other agents' past and current types, and irrespective of their own, as well as other agents', past behavior. As in other market design settings, it is, however, important that the market designer educates the bidders by carefully explaining them the structure of the scoring and payment rules and why the proposed auctions admit such simple equilibria. Also note that the proposed auctions are, in spirit, the analogs of the familiar second-price Vickrey auctions, but adapted to control for the dynamic nature of the externalities the agents impose on one another, and for the fact that the matches implemented in equilibrium need not maximize total surplus. In fact, as we show in the next section, by properly selecting the weights β, the designer can guarantee that the auctions maximize the platform's profits as opposed to welfare.

5 Profit maximization

We now show that the class of matching auctions introduced above includes a subclass that maximizes the platform's profits across all possible mechanisms. We make the following additional assumption on the distribution of the agents' vertical types.

Assumption 1 *For all $l \in N_k$, $k = A, B$, the Mills ratio $\left[1 - F_l^k(\theta_l^k)\right] / f_l^k(\theta_l^k)$ is non-increasing. Furthermore, $\underline{\theta}_l^k \cdot f_l^k(\underline{\theta}_l^k) > 1$.*

The first part of Assumption 1 is standard in mechanism design and guarantees that the agents' "virtual" vertical types $\theta_l^k - \left[1 - F_l^k(\theta_l^k)\right] / f_l^k(\theta_l^k)$ are non-decreasing in the true types. The second part is added to guarantee that the virtual vertical types are strictly positive. Given the multiplicative structure of the match values in (1), this assumption guarantees that the virtual match values $\left[\theta_l^k - \left[1 - F_l^k(\theta_l^k)\right] / f_l^k(\theta_l^k)\right] \varepsilon_{ijt}^k$ respect the same ranking as the true ones, $u_{ijt}^k = \theta_l^k \varepsilon_{ijt}^k$, thus avoiding confusion in the interpretation.

Because the myopic and the index rules are completely characterized by the weights β, hereafter, we denote by $D_l^k(\theta_0; m, \beta)$ and by $D_l^k(\theta_0; I, \beta)$, respectively, the expected match quality under a myopic and under an index rule, with weights β. We then have the following result:

Theorem 2 (profit maximization) *Let $\hat{\beta}$ be the weights of the period-0 membership statuses given by*[27]

$$\hat{\beta}_l^k(\theta_{l0}^k) \equiv 1 - \frac{1 - F_l^k(\theta_{l0}^k)}{f_l^k(\theta_{l0}^k)\theta_{l0}^k}, \text{ all } l \in N_k, \ k = A, B.$$

[27] Note that Assumption 1 guarantees that these weights are non-decreasing, strictly positive, and bounded, as required by the scoring rules of Theorem 1.

(i) Suppose processes are exogenous and $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; m, \hat{\beta}) \geq 0$*, all* $l \in N_k$*,* $k = A, B$*,* $\theta_{-l}^k \in \Theta_{-l}^k$*. The matching auctions in which (1) the scoring rule is the myopic rule with weights* $\hat{\beta}$*, and (2) the payments are given by (8) and (9) with weights* $\hat{\beta}$ *and* $L_l^k = 0$*, all* $l \in N_k$*,* $k = A, B$*, are profit-maximizing.*

(ii) Suppose processes are endogenous and $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; I, \hat{\beta}) \geq 0$*, all* $l \in N_k$*,* $k = A, B$*,* $\theta_{-l}^k \in \Theta_{-l}^k$*. In each of the three environments of part (ii) of Theorem 1, the matching auctions in which (1) the scoring rule is the index rule with weights* $\hat{\beta}$*, and (2) the payment are given by (8) and (9) with weights* $\hat{\beta}$ *and* $L_l^k = 0$*, all* $l \in N_k$*,* $k = A, B$*, are profit-maximizing.*[28]

(iii) The equilibria of the matching auctions in parts (i) and (ii) above that maximize the platform's profits are in truthful strategies, and are periodic ex-post.

The proof of Theorem 2, in the Appendix, is in three steps. First, using an approach similar to the one in Pavan, Segal and Toikka (2014), we show that, given any matching mechanism Γ and any *Bayes Nash equilibrium* (BNE) σ of the game induced by Γ, the period-0 interim expected payoff of each agent $l \in N_k$, $k = A, B$ must satisfy the envelope condition[29]

$$U_l^k(\theta_l^k) = U_l^k(\underline{\theta}_l^k) + \int_{\underline{\theta}_l^k}^{\theta_l^k} \mathbb{E}\left[D_l^k(\theta; \hat{\chi})|y\right] dy, \tag{13}$$

where the rule $\hat{\chi} = \left(\hat{\chi}_t(\theta, \varepsilon^t)\right)_{t=1}^{\infty}$ describes the state-contingent matches induced by the strategy profile σ in Γ, and where the expectation in (13) is with respect to the entire profile of vertical types θ, given the agent's own vertical type.[30]

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 the following weighted surplus function:

$$\begin{aligned} &\mathbb{E}^{\lambda[\hat{\chi}]}\left[\sum_{t=1}^{\infty} \delta^t \sum_{i \in N_A} \sum_{j \in N_B} \left(\hat{\beta}_i^A(\theta_i^A)\theta_i^A \varepsilon_{ijt}^A + \hat{\beta}_j^B(\theta_j^B)\theta_j^B \varepsilon_{ijt}^B - c_{ijt}(\hat{\chi}^{t-1}(\theta, \varepsilon^{t-1})\right) \hat{\chi}_{ijt}(\theta, \varepsilon^t)\right] \\ &\quad - \sum_{k=A,B} \sum_{l \in N_k} U_l^k(\underline{\theta}_l^k), \end{aligned} \tag{14}$$

for which the first term is only a function of the matching rule.

The optimality of the auctions in Theorem 2 is then established by showing that, under the truthful equilibria of the proposed auctions, (a) the induced state-contingent matches maximize the first component of the weighted surplus function (14), and (b) the participation constraints for the

[28] If $1 < M < n_A \cdot n_B$, the optimality of the proposed matching auctions requires the environment to be separable with respect to the rule $S^{I;\hat{\beta}}$ corresponding to the weights $\hat{\beta}$.

[29] By period-0 interim expected payoff, we mean the payoff the agent expects, under the equilibrium startegy profile σ, in the game induced by Γ, when his period-0 vertical type is equal to θ_l^k.

[30] In case of endogenous processes, these functions are defined only for those histories that are consistent with equilibrium play in previous periods. To ease the exposition, hereafter we omit the formal specification of the domains of such functions.

lowest vertical types are binding. That the lowest vertical type of each agent expects a nonnegative match quality (formally, that $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; r, \hat{\beta}) \geq 0$, $r = m, I$), together with the fact that match quality $D_l^k(\theta_{-l}^k, \cdot; r, \hat{\beta})$, $r = m, I$, is non-decreasing in the agent's vertical type θ_l^k under the equilibria of Theorem 2 and that the interim expected payoffs satisfy the envelope conditions in (13), in turn guarantees that participation in period zero is optimal for all agents. In fact, we show in the Appendix that payoffs satisfy a stronger envelope condition, which guarantees the optimality of period zero participation irrespective of the agents' beliefs about other agents' types.

Note that the conditions in Theorem 2 pertaining to expected match quality D_l^k can be relaxed for standard PBE solution concepts, and are vacuously satisfied, for example, if horizontal types are nonnegative, that is, if no agent ever dislikes interacting with any other agent. Furthermore, note that the proof of Theorem 2 shows that the truthful equilibria of the proposed matching auctions (with weights $\hat{\beta}$) maximize profits over all *BNE* (not just PBE) of all indirect mechanisms.

6 Welfare vs profit maximization

We now turn to the distortions in the dynamics of the matches due to the fact that the platform maximizes profits instead of welfare. Theorem 3 below identifies conditions under which the matching auctions introduced in Section 3 above admit a subclass for which the matches sustained in equilibrium are welfare maximizing.

Theorem 3 (welfare maximization) *Let β^W be the weights given by $\beta_l^{k,W}(\theta_l^k) = 1$, all $\theta_l^k \in \Theta_l^k$, $l \in N_k$, $k = A, B$.*

(i) *Suppose processes are exogenous. The matching auctions in which (1) the scoring rule is the myopic rule with weights β^W, and (2) the payments are given by (8) and (9) with weights β^W and with L_l^k large enough, all $l \in N_k$, $k = A, B$, are welfare-maximizing.*

(ii) *Suppose processes are endogenous. In each of the three environments of part (ii) of Theorem 1, the matching auctions in which (1) the scoring rule is the index rule with weights β^W, and (2) the payments are given by (8) and (9) with weights β^W and with L_l^k large enough, all $l \in N_k$, $k = A, B$, are welfare-maximizing.*[31]

(iii) *Suppose (a) processes are exogenous and $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; m, \beta^W) \geq 0$, all $l \in N_k$, $k = A, B$, $\theta_{-l}^k \in \Theta_{-l}^k$, or (b) processes are endogenous and $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; I, \beta^W) \geq 0$, all $l \in N_k$, $k = A, B$, $\theta_{-l}^k \in \Theta_{-l}^k$. In their respective environments of parts* (i) *and* (ii) *above, the matching auctions with payments given by (8) and (9), with $L_l^k = 0$, all $l \in N_k$, $k = A, B$, admit ex-post periodic equilibria in which agents participate and follow truthful strategies at all histories. Furthermore, such auctions maximize the platform's profits over all mechanisms implementing welfare-maximizing matches and inducing the agents to join the platform in period zero.*

[31] If $1 < M < n_A \cdot n_B$, the separability condition must hold with respect to the rule $S^{I;\beta^W}$ with weights β^W.

The results in parts (i) and (ii) follow directly from Theorem 1 by noting that, when the weights are given by β^W, the matching allocations sustained under truthful strategies maximize welfare after each history. The conditions in part (iii) of Theorem 3 in turn guarantee that, when the payments are as in (8) and (9) with weights β^W and with $L_l^k = 0$, all $l \in N_k$, $k = A, B$, agents find it optimal to participate and follow truthful strategies in each period, irrespective of their beliefs about other agents' current and past types, and past matches. That the proposed auctions maximize the platform's profits among all mechanisms implementing welfare-maximizing matches and inducing the agents to participate in period zero follows from arguments similar to those establishing the optimality of the matching auctions in Theorem 2 above. Note that the conditions in part (iii), though, do not guarantee that the platform's profits be positive. In general, the government may thus need to subsidize the platform to induce it to implement welfare-maximizing matching allocations, while also recouping its costs.

We now turn to the distortions brought in by profit maximization. Let $\chi^P \equiv \left(\chi_t^P(\theta,\omega)\right)_{t=1}^{\infty}$ and $\chi^W \equiv \left(\chi_t^W(\theta,\omega)\right)_{t=1}^{\infty}$ denote the state-contingent profit- and welfare-maximizing matching allocations, under the equilibria of the auctions of Theorems 2 and 3, respectively. Because the period-t history of horizontal types ε^t may differ in the two auctions (under endogenous processes), these allocations are expressed as a function of the sequence $\omega \equiv (\omega_{ijs}^k)_{(i,j)\in N_A\times N_B, k\in\{A,B\}}^{s=1,\ldots,\infty}$ of exogenous random variables that, together with the history of past matches (in case of endogenous processes), generate the sequence of horizontal types ε.[32] Because (θ,ω) are exogenous, we then drop them from the arguments of χ^P and χ^W to ease the notation. We then have the following result:[33]

Theorem 4 (distortions) *Suppose all agents derive a nonnegative utility from interacting with all other agents from the opposite side (formally, $\varepsilon_{ijt}^k \geq 0$, all $(i,j) \in N_A \times N_B$, $k = A, B$, $t \geq 1$).*

(1) If $M \geq n_A \cdot n_B$, then under both exogenous and endogenous processes, for all $(i,j) \in N_A \times N_B$, all $t \geq 1$,

$$\chi_{ijt}^P = 1 \;\Rightarrow\; \chi_{ijt}^W = 1.$$

(2) Suppose that either (a) the processes are exogenous and $M < n_A \cdot n_B$, or (b) the processes are endogenous and $M = 1$. Then, for all $t \geq 1$:

$$\sum_{(i,j)\in N_A\times N_B} \chi_{ijt}^W \geq \sum_{(i,j)\in N_A\times N_B} \chi_{ijt}^P.$$

(3) Suppose processes are endogenous, the environment is separable under both $S^{I;\hat{\beta}}$ and $S^{I;\beta^W}$,

[32] In case of exogenous processes, one can think of the horizontal types ε as coinciding with the random variables ω. With endogenous processes, instead, the horizontal types ε are generated from the exogenous random variables ω and the past decisions using the construction explained in Section 2.

[33] For this result, to facilitate the comparison between profit and welfare maximization, we assume that, in each period, both the profit- and the welfare-maximizing rules match pairs for which the scores are zero. That is, both rules use all the available M slots, unless the number of matches for which the score is nonnegative is strictly less than M.

and $1 < M < n_A \cdot n_B$. If, under profit-maximization, matching stops in finite time, then:

$$\sum_{t=1}^{\infty} \sum_{(i,j) \in N_A \times N_B} \chi_{ijt}^{W} \geq \sum_{t=1}^{\infty} \sum_{(i,j) \in N_A \times N_B} \chi_{ijt}^{P}.$$

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, a profit-maximizing platform induces fewer interactions than what is efficient. The precise nature of the inefficiency, however, depends on whether the processes governing the evolution of the agents' match values are endogenous or exogenous, and on the platform's capacity constraint.

When the capacity constraint is not binding, irrespective of the nature of the processes, in each period $t \geq 1$, the set of matches accommodated by a profit-maximizing platform is a subset of those maximizing welfare. In case of endogenous processes, when $M \geq n_A \cdot n_B$, once a match is severed, it is never reactivated again. Therefore, matches are gradually broken over time both under profit and under welfare maximization. In this case, all relationships are severed too early (weakly) under profit maximization.

This last property, however, does not extend to environments with binding capacity constraints. Suppose, for example, that $M = 1$ and that processes are endogenous. For any pair of agents $(i, j) \in N_A \times N_B$ and any history of past interactions between the pair, the index assigned to the pair in a profit-maximizing auction is always smaller than the index in the corresponding welfare-maximizing auction. Formally, this can be seen by noticing that the weights $\hat{\beta}$ used to compute the indexes under profit maximization are strictly smaller than the corresponding weights in a welfare-maximizing auction. However, the ranking of the indexes across pairs of agents under profit maximization need not coincide with the ranking under welfare maximization. As a result, certain interactions may last longer under profit maximization than under welfare maximization. What remains true, though, is that, if at a given point in time matching shuts down under welfare maximization, then, under profit maximization, either matching shut down already in previous periods, or it shuts down in the present period.

Similarly, when processes are exogenous, for any M, the aggregate level of interactions at any period under profit maximization is always (weakly) lower than under welfare maximization, although some individual interactions may last longer under profit maximization.

Interestingly, under endogenous processes, for intermediate capacity levels (that is, for $1 < M < n_A \cdot n_B$), the above conclusions about the aggregate level of activity being too low under profit maximization need not hold. Welfare maximization does not necessarily induce more aggregate activity in each period than profit maximization. Nor is it necessarily the case that, intertemporally, more matches are implemented under welfare maximization than under profit maximization. This latter property is guaranteed only if matching terminates in finite time under profit maximization. Example

6 illustrates.

Example 6 (endogenous processes with $1 < M < n_A \cdot n_B$) Suppose $N_A = \{1, 2, 3\}$, $N_B = \{1\}$ and $M = 2$. Consider the following deterministic processes governing the evolution of the myopic scores, respectively under welfare maximization (left panel), and under profit maximization (right panel):[34]

	$n = 0$	$n = 1$	$n = 2$	$n \geq 3$		$n = 0$	$n = 1$	$n = 2$	$n \geq 3$
$S^{m;\beta^W}_{1,1}$	7	-1	-1	-1	$S^{m;\hat{\beta}}_{1,1}$	3	-1	-1	-1
$S^{m;\beta^W}_{2,1}$	4	2	2	2	$S^{m;\hat{\beta}}_{2,1}$	4	2	2	2
$S^{m;\beta^W}_{3,1}$	5	1	-1	-1	$S^{m;\hat{\beta}}_{3,1}$	5	1	-1	-1

As in Example 5 above, since the myopic scores decline over time, the indexes coincide with the myopic scores, and hence both the profit-maximizing and the welfare-maximizing auctions simply match the two pairs of agents with the highest myopic score.[35] The welfare-maximizing auction matches in period 1 the pairs $(1, 1)$ and $(3, 1)$, in period 2 the pairs $(2, 1)$ and $(3, 1)$, and in any subsequent period only the pair $(2, 1)$. The profit-maximizing auction matches in period 1 the pairs $(2, 1)$ and $(3, 1)$, in period 2 the pairs $(1, 1)$ and $(2, 1)$, in period 3 the pairs $(2, 1)$ and $(3, 1)$, and from period 4 onwards only the pair $(2, 1)$. Thus, at any point in time, the total number of matches under profit maximization is weakly higher than under welfare maximization (strictly higher in period 3).$\square$

Interestingly, when certain agents may dislike certain interactions (formally, when horizontal types may be negative for certain pairs), the conclusions in Theorem 4 need not hold, and profit maximization may distort matching in the opposite direction. That is, profit-maximizing auctions may induce an inefficiently *high* volume of matches within each period, and interactions may last longer under profit maximization than under welfare maximization. Formally, when match values are negative, a pair's myopic score under profit maximization may be greater than its counterpart under welfare maximization. As a result, the conclusions in the above theorem can be overturned. Intuitively, the reason why a profit-maximizing platform may induce an inefficiently high volume of interactions is

[34] In turn, these processes can be generated by the truthful-strategies equilibria of Theorems 2 and 3, in the following environment. The vertical types are given by $\Theta^A_2 = \Theta^A_3 = \Theta^B_1 = \{1\}$, whereas $\Theta^A_1 = [2, 3]$, with F^A_1 uniform over $[2, 3]$. The platform's costs are such that $c_{11t} = c_{31t} = 1$, $c_{21t} = 0$. Lastly, the horizontal types are such that $\varepsilon^B_{i1t} = 0$, all $i = 1, 2, 3$ all $t \geq 1$. For the side-A agents, instead, the horizontal types evolve deterministically over time according to the table below (with n indicating the number of previous interactions with agent 1 from side B). The processes in the example then correspond to those for the realized vertical type $\theta^A_i = 2$.

	$n = 0$	$n = 1$	$n = 2$	$n \geq 3$
ε^A_{11}	4	0	0	0
ε^A_{21}	4	2	2	2
ε^A_{31}	6	2	0	0

[35] Also note that, for sufficiently low δ, the environment in this example is separable both under $S^{I;\hat{\beta}}$ and $S^{I;\beta^W}$.

that this may permit it to economize on the informational rents that it must leave to the agents. By locking those agents selecting a low membership status (equivalently, claiming to have a low vertical type) 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 those high-type agents. The following example illustrates:

Example 7 (upward distortions under negative values) Consider the following environment where processes are exogenous, $N_A = N_B = \{1\}$, $M = 1$, and $c_{11t} = 0$, all $t \geq 1$. The vertical types are given by $\Theta_1^B = \{1\}$ and $\Theta_1^A = [1+\varsigma, 2+\varsigma]$, $\varsigma > 0$, with F_1^A uniform over Θ_1^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 $\hat{\beta}_1^A(\theta_1^A) = \varsigma$ and $\hat{\beta}_1^B(\theta_1^B) = 1$. Furthermore, consider a realized sequence $(\varepsilon_{11t}^A)_{t=1}^{\infty}$ of horizontal types for agent 1 from side A such that $\varepsilon_{11t}^A = -3$, 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. $\square$

Finally, note that the familiar result of "no distortion at the top" from standard mechanism design 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 that the platform must provide to such agents-inputs to induce them to reveal their private information.

The above results have implications for the design of government intervention in matching markets. The results in Theorem 3 indicate that, in many markets of interest, the government could simply impose that the scoring rules be tilted by adopting the welfare-maximizing weights β^W in lieu of the profit-maximizing ones $\hat{\beta}$. However, because the welfare-maximizing auctions need not yield positive profits (even under the profit-maximizing payments of part (iii) in Theorem 3), the government may need to subsidize the intermediary.

The results in Theorem 4, in turn, suggest that simple subsidies aimed at inducing platforms to increase the volume of matches they accommodate in each period can be welfare-increasing in markets with non-binding capacity constraints, but can be counterproductive when such constraints bind. In this latter case, regulators may need to target the subsidizes to specific matches, which may be feasible in markets with a small number of agents, but seems difficult in large anonymous markets.

Lastly, the above results indicate that, when certain agents dislike certain interactions, the government may want to tax certain interactions while subsidizing others. Again, whether this is at all possible is likely to depend on how anonymous the interactions in such markets are.

7 Conclusions

This paper introduces and then studies a class of auctions that platforms can use to match agents from different sides of the market in environments in which agents' preferences for potential partners evolve over time, either exogenously, or as a function of previous interactions.

The proposed auctions are fairly simple and can be used in a variety of markets including those for scientific outsourcing, peer-to-peer lending, and organized events. Upon joining the platform, agents 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. The platform then computes bilateral scores for each match and implements the matches with the highest nonnegative score, up to capacity.

In case the processes governing the evolution of the match values are exogenous, the scores are myopic and reflect the flow values the agents assign to the matches, net of the platform's costs of providing all the auxiliary services necessary to implement the matches, and net of handicaps controlling for the agents' information rents. When, instead, the match values depend on past interactions, the scores are forward-looking and take the form of indexes similar to those in the operation research literature (e.g., Gittins, 1979), but adjusted to account for the costs of information rents.

The framework is flexible enough to admit as special cases such environments in which learning occurs immediately upon matching, as well as such environments in which agents enjoy interacting at most finitely many times with each partner. The results are then used to shed light on the inefficiencies associated with the private provision of matching services, which in turn can be used to evaluate the merits of certain government interventions.

Many extensions appear interesting. Certain results can be adapted to accommodate for the possibility that agents' match values depend on the entire history of past interactions (e.g., agents may care about their partners' previous partners). Extending the analysis to allow for more general forms of correlation in the agents' preferences is challenging but also worth exploring. It also appears interesting to study how the results specialize in markets in which payments cannot be collected from one side of the market (as is the case with sponsored search).

In would also be interesting to consider a broader class of preferences and processes governing the evolution of the match values. The ones considered in the present paper have the advantage of maintaining tractability. As noticed above, no general solutions are known in the experimentation literature for multi-armed bandit problems in which multiple arms can be activated simultaneously. However, there are environments in which asymptotic results can be established for large markets. For example, Bergemann and Valimaki (2001) show that an index policy is optimal for stationary multi-armed bandit problems, in which there are countable infinitely many ex-ante identical arms (and approximately optimal in the limit as the number of arms goes to infinity). For the more general case of *restless bandits* (which involves arms that evolve and yield rewards even when they are not activated), under the infinite-horizon average reward criterion, Weber and Weiss (1990) provide conditions guaranteeing that an index policy that selects in each period the arms with the highest Whittle index is asymptotically optimal (see also Whittle (1988)).[36] Such asymptotic results could also be useful in designing scoring rules for large matching auctions.

We conclude by noting that, while matching dynamics in the present paper originate in changes in agents' preferences for potential partners, another line of recent research explores matching dynamics driven by the (stochastic) arrival or departure of agents to and from the market (see, for example, Anderson et al. 2015, Baccara et al. 2015, and Akbarpour et al. 2016). Combining the two literatures, for example by incorporating stochastic arrivals into the scores of the matching auctions introduced in the present paper is another direction of future research that is expected to generate interesting insights about matching dynamics under profit- and welfare-maximizing mechanisms.

Appendix

Proof of Theorem 1. The proof is in three steps. Step 1 shows that, when all agents follow truthful strategies, in each of the respective environments of Theorem 1, the matches under the rules $\chi^{m;\beta}$ and $\chi^{I;\beta}$ maximize the continuation weighted surplus, as defined in (6), starting from any period-t history, any $t \geq 1$. Step 2 shows that when, in addition, the period-t payments are as in (8), any $t \geq 1$, then participating in the auctions and following truthful strategies constitutes a periodic ex-post continuation equilibrium, after any period-t history, any $t \geq 1$. Finally, Step 3 completes the proof by showing that the matching rules $\chi^{m;\beta}$ and $\chi^{I;\beta}$ satisfy a certain dynamic monotonicity condition (defined below), which guarantees that when, in addition, the period-0 membership fees are as in (9), then participating and following truthful strategies is a periodic ex-post equilibrium starting from period zero.

Step 1. We first establish that, in each of the respective environments of Theorem 1, the matching rules $\chi^{m;\beta}$ and $\chi^{I;\beta}$ maximize continuation weighted surplus (6). That is, irrespective of the particular history that led to the selection of the period-0 membership statuses θ_0 and of the past matches x^{t-1}, in the continuation game that starts with period t, $t \geq 1$, when the true vertical type profile is θ_t, the true profile of horizontal types is ε_t (with ε_t obtained from θ_t and b_t using (4)), and when agents follow truthful strategies from period t onwards, the matches under $\chi^{m;\beta}$ and $\chi^{I;\beta}$ maximize $W_t(\theta_0, \theta_t, b_t, x^{t-1}; \beta)$, over the entire set $\mathcal{X}$ of all feasible matching rules, with message spaces defined as in the matching auctions (i.e., $\mathcal{M}^k_{l0} = \Theta^k_l$, and $\mathcal{M}^k_{lt} = \Theta^k_l \times \mathbb{R}^{N_{-k}}$, $l \in N_k$, $k = A, B$, $t \geq 1$).

Lemma 1 below considers the case of exogenous processes and the case of endogenous processes with extreme capacity constraints ($M = 1$ and $M \geq n_A \cdot n_B$). Proposition 1 below considers the case of endogenous processes with intermediate capacity constraints, under the separability assumption.

Lemma 1 *(i) Suppose the processes are exogenous. Then $\chi^{m;\beta}$ maximizes continuation weighted surplus at all histories. (ii) Suppose processes are endogenous and either $M \geq n_A \cdot n_B$, or $M = 1$. Then $\chi^{I;\beta}$ maximizes continuation weighted surplus at all histories.*

[36] In the special case in which passive arms are static and yield no reward, this index reduces to the Gittins index.

Proof of Lemma 1. Part (*i*) follows directly from the fact that $\chi^{m;\beta}$ maximizes (6) history by history. For part (*ii*), note that the problem of maximizing (6) can be viewed as a multiarmed bandit problem, with each arm corresponding to a potential match, and with the flow period-t reward of activating each arm (i,j) given by the myopic score $S_{ijt}^{m;\beta}$. When $M=1$, that $\chi^{I;\beta}$ maximizes (6) at all histories is then immediate, as the score $S_{ijt}^{I;\beta}$ corresponds to the arm's Gittins Index (see, for example, Whittle (1982)). Similarly, when $M \geq n_A \cdot n_B$, since the capacity constraint never binds, the platform's problem can be viewed as a collection of $n_A \cdot n_B$ separate two-armed bandit problems, one for each potential pair of agents, with the reward from matching the pair (i,j) given by $S_{ijt}^{m;\beta}$ and the reward from activating the "safe arm" identically equal to zero. In both cases, $\chi^{I;\beta}$ maximizes continuation weighted surplus. ■

Next, consider the case of endogenous processes and arbitrary capacity constraints, under the separability assumption.

Proposition 1 *Suppose processes are endogenous and the environment is separable under the rule $S^{I;\beta}$. For any $M \in \mathbb{N}$, the matching rule $\chi^{I;\beta}$ maximizes continuation weighted surplus (6) starting from any period-t history, any $t \geq 1$.*

Proof of Proposition 1. The proof follows from the four lemmas below. Consider a fictitious environment in which the rewards are given by the auxiliary processes defined in (11). That is, for any $t \geq 1$, $(i,j) \in N_A \times N_B$, (θ^t, b^t, x^{t-1}), suppose the period-t reward from the match (i,j) is given by $\underline{S}_{ijt}^{\beta}$. Denote by $\underline{\chi}^{m;\beta}$ the myopic matching rule that, in each period, matches the pairs with the highest nonnegative auxiliary rewards $\underline{S}_{ijt}^{\beta}$, subject to the platform's capacity constraint (ties broken arbitrarily). Formally, let

$$\underline{Q}_t^{\beta}(\theta^t, b^t, x^{t-1}) \equiv \left\{ \begin{array}{l} (i,j) \in N_A \times N_B \text{ s.t. (i) } \underline{S}_{ijt}^{\beta}(\theta^t, b^t, x^{t-1}) \geq 0 \text{ and} \\ \text{(ii) } \#\left\{(l,m) \in N_A \times N_B : \underline{S}_{lmt}^{\beta}(\theta^t, b^t, x^{t-1}) > \underline{S}_{ijt}^{\beta}(\theta^t, b^t, x^{t-1})\right\} < M \end{array} \right\}.$$

Then (a) $\underline{\chi}_{ijt}^{m;\beta}(\theta^t, b^t, x^{t-1}) = 1$ only if $(i,j) \in \underline{Q}_t^{\beta}(\theta^t, b^t, x^{t-1})$. Furthermore, (b) if $\underline{S}_{ijt}^{\beta}(\theta^t, b^t, x^{t-1}) > 0$ and $\#\left\{(l,m) \in N_A \times N_B : \underline{S}_{lmt}^{\beta}(\theta^t, b^t, x^{t-1}) \geq \underline{S}_{ijt}^{\beta}(\theta^t, b^t, x^{t-1})\right\} \leq M$, then $\underline{\chi}_{ijt}^{m;\beta}(\theta^t, b^t, x^{t-1}) = 1$. Finally, (c), if $(i,j) \in \underline{Q}_t^{\beta}(\theta^t, b^t, x^{t-1})$, $\underline{S}_{ijt}^{\beta}(\theta^t, b^t, x^{t-1}) > 0$, and $\underline{\chi}_{ijt}^{m;\beta}(\theta^t, b^t, x^{t-1}) = 0$, then $\#\left\{(l,m) \in N_A \times N_B : \underline{\chi}_{lmt}^{m;\beta}(\theta^t, b^t, x^{t-1}) = 1\right\} = M$.

Lemma 2 *Suppose processes are endogenous and the environment is separable under the rule $S^{I;\beta}$. Then in the corresponding fictitious environment in which the flow rewards are given by $\underline{S}^{\beta}$, the rule $\underline{\chi}^{m;\beta}$ maximizes the expected discounted sum of the auxiliary rewards starting from any period-t history, any $t \geq 1$.*[37]

[37] The formula for the expected discounted sum of the auxiliary rewards is the same as the one in (6), but with $\underline{S}_{ijs}$ replacing S_{ijs}, and each χ_{ijs} defined over (θ^s, b^s, x^{s-1}) as opposed to $(\theta_0, \theta_s, b_s, x^{s-1})$, all $(i,j) \in N_A \times N_B$, $s \geq t$. As in the case of the formula in (6), the expectation over future values and types is under truthful strategies.

Proof of Lemma 2. Suppose, towards a contradiction, that the claim is not true. This means that there exists a period $t \geq 1$, a history (θ^t, b^t, x^{t-1}), and a feasible rule $\underline{\chi} \neq \underline{\chi}^{m;\beta}$ such that the expected discounted sum of auxiliary rewards from period t onwards is higher under $\underline{\chi}$ than under $\underline{\chi}^{m;\beta}$. For this to be the case, there must exist a period $s \geq t$, and a set of histories (θ^s, b^s, x^{s-1}) of strictly positive probability under $\lambda[\underline{\chi}]|\theta^t, b^t, x^{t-1}$, for which the matches under the two rules differ, meaning that one of the following two properties (or both) must hold, for some $(i,j) \in N_A \times N_B$: (a) either $\underline{\chi}_{ijs}(\theta^s, b^s, x^{s-1}) = 1$ and $(i,j) \notin \underline{Q}_s^{\beta}(\theta^s, b^s, x^{s-1})$, in which case $\underline{\chi}_{ijs}^{m}(\theta^s, b^s, x^{s-1}) = 0$; (b) or $\underline{\chi}_{ijs}(\theta^s, b^s, x^{s-1}) = 0$, $\underline{\chi}_{ijs}^{m;\beta}(\theta^s, b^s, x^{s-1}) = 1$, $\underline{S}_{ijs}^{\beta}(\theta^s, b^s, x^{s-1}) > 0$, and

$$\#\left\{(l,m) \in N_A \times N_B : \underline{S}_{lms}^{\beta}(\theta^s, b^s, x^{s-1}) \geq \underline{S}_{ijs}^{\beta}(\theta^s, b^s, x^{s-1}),\ \underline{\chi}_{lms}(\theta^s, b^s, x^{s-1}) = 1\right\} < M.$$

Note that any other case in which the two rules $\underline{\chi}^{m}$ and $\underline{\chi}$ implement different allocations for the pair (i,j) and neither (a) nor (b) are satisfied is inconsequential for the difference in the expected continuation weighted surplus under the two rules (in fact, the difference in the allocations in any such other case simply reflects the way ties are broken).

First, consider case (a). Because $(i,j) \notin \underline{Q}_s^{\beta}(\theta^s, b^s, x^{s-1})$, either (a1) $\underline{S}_{ijs}^{\beta}(\theta^s, b^s, x^{s-1}) < 0$ or (a2)

$$\#\left\{(l,m) \in N_A \times N_B : \underline{S}_{lms}^{\beta}(\theta^s, b^s, x^{s-1}) > \underline{S}_{ijs}^{\beta}(\theta^s, b^s, x^{s-1})\right\} \geq M.$$

In the first case (case a1), given that the auxiliary rewards are non-increasing, it is immediate that $\underline{\chi}^{m;\beta}$, by leaving the pair (i,j) unmatched in period s, improves upon $\underline{\chi}$. Thus consider the second possibility, case (a2). There must exist another pair $(i',j') \in N_A \times N_B$ such that $\underline{S}_{i'j's}^{\beta}(\theta^s, b^s, x^{s-1}) > \underline{S}_{ijs}^{\beta}(\theta^s, b^s, x^{s-1}) \geq 0$, $\underline{\chi}_{i'j's}(\theta^s, b^s, x^{s-1}) = 0$ and $\underline{\chi}_{i'j's}^{m}(\theta^s, b^s, x^{s-1}) = 1$. Because the environment is separable under the rule $S^{I;\beta}$, this implies that $\underline{S}_{i'j's}^{\beta}(\theta^s, b^s, x^{s-1}) \geq \underline{S}_{ijs}^{\beta}(\theta^s, b^s, x^{s-1})/(1-\delta)$. By definition, the auxiliary processes are non-increasing. This means that the expected discounted sum of the stream of auxiliary rewards that can be obtained by matching (i,j) in period s (and possibly in some of the subsequent periods) is no greater than $\underline{S}_{ijs}^{\beta}(\theta^s, b^s, x^{s-1})/(1-\delta)$, which is smaller than the flow reward $\underline{S}_{i'j's}^{\beta}(\theta^s, b^s, x^{s-1})$ obtained by matching (i',j') in period s. Hence, by favoring (i',j') over (i,j) in period t, $\underline{\chi}^{m;\beta}$ again improves upon the rule $\underline{\chi}$.

Next, consider case (b). The myopic rule $\underline{\chi}^{m;\beta}$ improves upon $\underline{\chi}$ by either adding the match (i,j) to the set of matches implemented under $\underline{\chi}$ (in case $\underline{\chi}$ matches fewer than M pairs) or by matching the pair (i,j) instead of another pair (i',j') that is matched under the rule $\underline{\chi}$ but not under the rule $\underline{\chi}^{m;\beta}$ and for which the flow auxiliary reward is strictly smaller than (i,j)'s (in which case the separability assumption again guarantees that such a change improves upon $\underline{\chi}$).

In either cases (a) and (b) above, $\underline{\chi}^{m;\beta}$ thus improves upon $\underline{\chi}$. Applying the arguments above to all histories and all pairs of agents yields a contradiction to the fact that $\underline{\chi}$ strictly dominates $\underline{\chi}^{m;\beta}$.
■

The next two lemmas fix the matching rule and relate the expected discounted sum of the rewards in the primitive environment (where the flow rewards are given by $S^{m;\beta}$) to the expected discounted sum of the rewards in the fictitious environment (where the rewards are the auxiliary ones, $\underline{S}^{\beta}$). The proof of the next two lemmas follows from Ishikida and Varaiya (1994) (see also Mandelbaum, 1986, and Pandelis and Teneketzis, 1999), adapted to the matching environment under examination.

Lemma 3 *Suppose the agents follow truthful strategies. The (period-*0*) expected discounted sum of the auxiliary rewards* $\underline{S}^{\beta}$ *under any rule* χ *is weakly higher than the (period-*0*) expected discounted sum of the primitive rewards* $S^{m;\beta}$ *under the same rule* χ*.*

To understand the result, note first that the auxiliary rewards are themselves Gittins indexes, and are therefore defined selecting the stopping times that maximize the expected average discounted payoff per unit of expected discounted time. Also note that the optimal stopping times coincide with the first time at which the index drops weakly below its initial value. Therefore, under any arbitrary matching rule where the stopping times are possibly different from the optimal ones, the expected discounted sum of the true rewards $S^{m;\beta}$ can never exceed the expected discounted sum of the auxiliary rewards $\underline{S}^{\beta}$.

In the special case where the matching rule is the myopic rule for the auxiliary rewards, $\underline{\chi}^{m;\beta}$, the stopping times implemented under such rule coincide with those under the Gittins policy for the primitive rewards. As a result, the expected discounted sum of the primitive rewards is the same as that for the auxiliary rewards.

Lemma 4 *Suppose the agents follow truthful strategies. The (period-*0*) expected discounted sum of the auxiliary rewards* $\underline{S}^{\beta}$ *under the myopic rule* $\underline{\chi}^{m;\beta}$ *for the auxiliary processes is the same as the (period-*0*) expected discounted sum of the primitive rewards* $S^{m;\beta}$ *under the same rule* $\underline{\chi}^{m;\beta}$*.*

Finally, the following property is a direct consequence of the definition of the index rule $\chi^{I;\beta}$ for the primitive environment.

Lemma 5 *Suppose the agents follow truthful strategies. The dynamics of the matches under* $\underline{\chi}^{m;\beta}$ *are the same as under* $\chi^{I;\beta}$*.*

The proof of Proposition 1 then follows from combining the results in the above four lemmas. ■

Step 2. Fix the weights β and denote by $\tilde{\chi}$ a matching rule that maximizes the continuation weighted surplus at all histories, and by $\tilde{\chi}^{k}_{-l}$ a matching rule that does so in the absence of agent l from side $k \in \{A, B\}$ (equivalently, that maximizes continuation weighted surplus when the myopic score of any match that involves agent l from side k is identically equal to zero, in which case $\tilde{\chi}^{k}_{-l}$ can be assumed to never implement any match involving agent l). The existence of such rules, in the corresponding environments of Theorem 1, has been shown in Step 1 above.[38] Denote by $\tilde{\psi}_{t\geq 1} \equiv (\tilde{\psi}_s)_{s\geq 1}$ the collection of payment functions, given by (8), starting from period one, defined with respect to the matching rule $\tilde{\chi}$. Henceforth, the weighted surpluses W_t and $W_t^{-l,k}$, as well as the resulting marginal

[38] These rules correspond to the myopic and index rules for exogenous and endogenous processes, respectively.

and flow contributions to weighted surplus, as defined in the main text, unless otherwise specified, are with respect to the matching rules $\tilde{\chi}$ and $\tilde{\chi}^k_{-l}$, respectively. Note that, because the weights β are held fixed, to ease the notation we drop them from all functions below, when there is no risk of confusion.

Lemma 6 *Consider an auction in which the matching rule is* $\tilde{\chi}$ *and the payments from period* 1 *onwards are given by* $\tilde{\psi}_{t\geq 1}$. *In such an auction, participating and following truthful strategies constitutes a periodic ex-post continuation equilibrium, after any period-t history,* $t \geq 1$.

Proof of Lemma 6. We show that, in the continuation game that starts in period $t \geq 1$, irrespective of the history of past play, of the true vertical type profile θ, and of the history of past and current horizontal types $\varepsilon^t \equiv (\varepsilon_s)^t_{s=1}$, any agent $l \in N_k$, $k = A, B$, who expects all other agents to participate and follow truthful strategies from period t (included) onwards, finds it optimal to do the same.

Consider agent l from side A (the problem for any agent from side B is similar). Suppose that the true profile of vertical types is θ, the true profile of period-t match values is u_t, the profile of period-0 membership choices is θ_0, and the history of past matches is x^{t-1}. Denote by

$$\mathbb{E}^{\tilde{\lambda}[\tilde{\chi}]|\theta,u_t,x^{t-1};(\hat{\theta}^A_{lt},\hat{b}^A_{lt})}\left[R^A_{lt+1}(\theta_0,\theta_{t+1},b_{t+1},x^t)\right]$$

the expected contribution of agent l to continuation weighted surplus from period $t+1$ onwards, when, in period t, the agent selects the period-t membership status $\hat{\theta}^A_{lt}$, submits the period-t bids $\hat{b}^A_{lt}$, follows a truthful strategy from period $t+1$ onwards, and expects all other agents to follow truthful strategies at all periods $s \geq t$.[39] Note that, when the agent follows the truthful strategy also in period t (i.e., when $\hat{\theta}^A_{lt} = \theta^A_l$ and $\hat{b}^k_{lt} = u^A_{lt}$), then

$$\tilde{\lambda}[\tilde{\chi}]|\theta, u_t, x^{t-1}; (\theta^A_l, u^A_{lt}) = \lambda[\tilde{\chi}]|\theta, u_t, x^{t-1},$$

where the process $\lambda[\tilde{\chi}]|\theta, u_t, x^{t-1}$ is as defined in the main text.

Since the agent can adjust his membership status in any of the subsequent periods, any deviation from the truthful strategy in period t can be corrected in period $t+1$. This means that, to prove the result, it suffices to show that the agent prefers to follow the truthful strategy from period t onwards than deviating in period t and then reverting to the truthful strategy from period $t+1$ onwards.

Under the proposed auction rules, when the agent follows the truthful strategy from period $t+1$ onwards, his continuation payoff from period $t+1$ onwards is given by

$$\frac{1}{\beta^A_l(\theta^A_{l0})} R^A_{lt+1}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)$$

[39] The stochastic process $\tilde{\lambda}[\tilde{\chi}]|\theta, \varepsilon_t, x^{t-1}; (\hat{\theta}^A_{lt}, \hat{b}^A_{lt})$ is here over future matches, bids and membership choices, under the rule $\tilde{\chi}$, when agent l's period-t choices are $(\hat{\theta}^A_{lt}, \hat{b}^A_{lt})$, the profile of vertical types is θ, the true profile of period-t horizontal types is ε_t, the history of past matches is x^{t-1}, the agent plans to follow a truthful strategy from $t+1$ onwards and all other agents follow truthful strategies from period t onwards.

Therefore, it is enough to show that, for any period-t selection $(\hat{\theta}_{lt}^{A}, \hat{b}_{lt}^{A})$,

$$\begin{aligned}
&\sum_{j\in N_B} u_{ljt}^{A}\tilde{\chi}_{ljt}\left(\theta_0,(\theta_l^A,\theta_{-l}^A),(u_{lt}^A,u_{-lt}^A),x^{t-1}\right)-\tilde{\psi}_{lt}^{A}\left(\theta_0,(\theta_l^A,\theta_{-l}^A),(u_{lt}^A,u_{-lt}^A),x^{t-1}\right)\\
&\quad+\frac{\delta}{\beta_l^A(\theta_{l0}^A)}\mathbb{E}^{\lambda[\tilde{\chi}]|\theta_l^A,\theta_{-l}^A,u_{lt}^A,u_{-lt}^A,x^{t-1}}\left[R_{lt+1}^{A}(\theta_0,\theta_{t+1},b_{t+1},x^t)\right]\\
&\quad\geq\sum_{j\in N_B} u_{ljt}^{A}\tilde{\chi}_{ljt}\left(\theta_0,(\hat{\theta}_{lt}^A,\theta_{-l}^A),(\hat{b}_{lt}^A,u_{-lt}^A),x^{t-1}\right)-\tilde{\psi}_{lt}^{A}\left(\theta_0,(\hat{\theta}_{lt}^A,\theta_{-l}^A),(\hat{b}_{lt}^A,u_{-lt}^A),x^{t-1}\right)\\
&\qquad+\frac{\delta}{\beta_l^A(\theta_{l0}^A)}\mathbb{E}^{\tilde{\lambda}[\tilde{\chi}]|\theta_l^A,\theta_{-l}^A,u_{lt}^A,u_{-lt}^A,x^{t-1};(\hat{\theta}_{lt}^A,\hat{b}_{lt}^A)}\left[R_{lt+1}^{A}(\theta_0,\theta_{t+1},b_{t+1},x^t)\right].
\end{aligned}\tag{15}$$

The left hand side of the above inequality can be rewritten in terms of the functions W_t and $W_t^{-l,A}$ as follows:

$$\frac{1}{\beta_l^A(\theta_{l0}^A)}\left[W_t\left(\theta_0,(\theta_l^A,\theta_{-l}^A),(u_{lt}^A,u_{-lt}^A),x^{t-1}\right)-W_t^{-l,A}\left(\theta_0,(\theta_l^A,\theta_{-l}^A),(u_{lt}^A,u_{-lt}^A),x^{t-1}\right)\right].\tag{16}$$

That is, agent l's expected continuation payoff when he follows the truthful strategy from period t onward is equal to his expected contribution to the maximal continuation weighted surplus, scaled by the weight $\beta_l^A(\theta_{l0}^A)$. It then suffices to show that (16) is weakly greater than the right hand side of (15).

Next, note that the flow contribution r_{lt}^k can be rewritten as

$$\begin{aligned}
r_{lt}^{k}(\theta_0,\theta_t,b_t,x^{t-1})=&\sum_{i\in N_A}\sum_{j\in N_B} S_{ijt}^{m;\beta}\left(\theta_0,\theta_t,b_t,x^{t-1}\right)\cdot\tilde{\chi}_{ijt}(\theta_0,\theta_t,b_t,x^{t-1})\\
&-\sum_{i\in N_A}\sum_{j\in N_B} S_{ijt}^{m;\beta}\left(\theta_0,\theta_t,b_t,x^{t-1}\right)\tilde{\chi}_{ijt}^{-l,k}(\theta_0,\theta_t,b_t,x^{t-1})\\
&+\delta\mathbb{E}^{\lambda[\tilde{\chi}]|\theta_t,b_t,x^{t-1}}\left[W_{t+1}^{-l,k}(\theta_0,\theta_{t+1},b_{t+1},x^t)\right]\\
&-\delta\mathbb{E}^{\lambda[\tilde{\chi}^{-l,k}]|\theta_t,b_t,x^{t-1}}\left[W_{t+1}^{-l,k}(\theta_0,\theta_{t+1},b_{t+1},x^t)\right].
\end{aligned}\tag{17}$$

Using (17), we can rewrite the period-t payment in the right-hand side of (15) as follows:

$$
\begin{aligned}
&\tilde{\psi}_{lt}^{A}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right) \\
&= \sum_{j \in N_B} \hat{b}_{ljt}^{A} \tilde{\chi}_{ljt}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right) - \frac{1}{\beta_l^A(\theta_{l0}^A)} r_{lt}^{A}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right) \\
&\quad = \sum_{j \in N_B} \hat{b}_{ljt}^{A} \tilde{\chi}_{ljt}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right) \\
&\quad - \frac{1}{\beta_l^A(\theta_{l0}^A)} \sum_{i \in N_A} \sum_{j \in N_B} S_{ijt}^{m;\beta}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right) \cdot \tilde{\chi}_{ijt}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right) \\
&\quad - \frac{\delta}{\beta_l^A(\theta_{l0}^A)} \mathbb{E}^{\tilde{\lambda}[\tilde{\chi}] | \hat{\theta}_{lt}^{A}, \theta_{-l}^{A}, \hat{b}_{lt}^{A}, u_{-lt}^{A}, x^{t-1}} \left[W_{t+1}^{-l,A}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)\right] \\
&\quad + \frac{1}{\beta_l^A(\theta_{l0}^A)} W_t^{-l,A}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right) \\
&\quad = -\frac{1}{\beta_l^A(\theta_{l0}^A)} \sum_{i \in N_A \setminus \{l\}} \sum_{j \in N_B} S_{ijt}^{m;\beta}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right) \cdot \tilde{\chi}_{ijt}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right) \\
&\quad - \frac{1}{\beta_l^A(\theta_{l0}^A)} \sum_{j \in N_B} \left(\beta_j^B(\theta_{j0}^B) u_{ljt}^B - c_{ljt}(x^{t-1})\right) \cdot \tilde{\chi}_{ljt}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right) \\
&\quad - \frac{\delta}{\beta_l^A(\theta_{l0}^A)} \mathbb{E}^{\tilde{\lambda}[\tilde{\chi}] | \hat{\theta}_{lt}^{A}, \theta_{-l}^{A}, \hat{b}_{lt}^{A}, u_{-lt}^{A}, x^{t-1}} \left[W_{t+1}^{-l,A}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)\right] \\
&\quad + \frac{1}{\beta_l^A(\theta_{l0}^A)} W_t^{-l,A}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right).
\end{aligned}
$$

Furthermore, note that

$$
\begin{aligned}
&\mathbb{E}^{\tilde{\lambda}[\tilde{\chi}] | \theta_l^A, \theta_{-l}^A, u_{lt}^A, u_{-lt}^A, x^{t-1}; (\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)} \left[R_{lt+1}^{A}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)\right] \\
&\quad = \mathbb{E}^{\tilde{\lambda}[\tilde{\chi}] | \theta_l^A, \theta_{-l}^A, u_{lt}^A, u_{-lt}^A, x^{t-1}; (\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)} \left[W_{t+1}(\theta_0, \theta_{t+1}, b_{t+1}, x^t) - W_{t+1}^{-l,A}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)\right] \\
&\quad = \mathbb{E}^{\tilde{\lambda}[\tilde{\chi}] | \theta_l^A, \theta_{-l}^A, u_{lt}^A, u_{-lt}^A, x^{t-1}; (\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)} \left[W_{t+1}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)\right] \\
&\qquad - \mathbb{E}^{\lambda[\tilde{\chi}] | \hat{\theta}_{lt}^{A}, \theta_{-l}^{A}, \hat{b}_{lt}^{A}, u_{-lt}^{A}, x^{t-1}} \left[W_{t+1}^{-l,A}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)\right],
\end{aligned}
$$

where the last equality uses the fact that, given x^t, $W_{t+1}^{-l,A}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)$ is invariant in agent l's period-t bids and that the period-t decisions x_t are invariant in the agent's *true* types. Therefore, the right hand side of the inequality (15) is equal to

$$
\begin{aligned}
&\frac{1}{\beta_l^A(\theta_{l0}^A)} \sum_{i \in N_A} \sum_{j \in N_B} S_{ijt}^{m;\beta}\left(\theta_0, \theta, u_t, x^{t-1}\right) \cdot \tilde{\chi}_{ijt}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right) \\
&\quad + \frac{\delta}{\beta_l^A(\theta_{l0}^A)} \mathbb{E}^{\tilde{\lambda}[\tilde{\chi}] | \theta_l^A, \theta_{-l}^A, u_{lt}^A, u_{-lt}^A, x^{t-1}; (\hat{\theta}_{lt}^A, \hat{b}_{lt}^A)} \left[W_{t+1}(\theta_0, \theta_{t+1}, b_{t+1}, x^t)\right] \\
&\quad - \frac{1}{\beta_l^A(\theta_{l0}^A)} W_t^{-l,A}\left(\theta_0, (\hat{\theta}_{lt}^{A}, \theta_{-l}^{A}), (\hat{b}_{lt}^{A}, u_{-lt}^{A}), x^{t-1}\right).
\end{aligned}
$$

Since, again, $W_t^{-l,A}$ is independent of agent l's period-t bids, to establish that the inequality in

(15) holds, it suffices to show that

$$W_t(\theta_0,\theta,u_t,x^{t-1}) \geq \sum_{i\in N_A}\sum_{j\in N_B} S^{m;\beta}_{ijt}\left(\theta_0,\theta,u_t,x^{t-1}\right)\cdot \tilde{\chi}_{ijt}\left(\theta_0,(\hat{\theta}^A_{lt},\theta^A_{-l}),(\hat{b}^A_{lt},u^A_{-lt}),x^{t-1}\right) \\ + \mathbb{E}^{\tilde{\lambda}[\tilde{\chi}]|\theta^A_l,\theta^A_{-l},u^A_{lt},u^A_{-lt},x^{t-1};(\hat{\theta}^A_{lt},\hat{b}^A_{lt})}\left[W_{t+1}(\theta_0,\theta_{t+1},b_{t+1},x^t)\right]. \tag{18}$$

The inequality in (18) follows from the definition of the matching rule $\tilde{\chi}$.

That it is (periodic ex-post) optimal for each agent to participate at all periods $t \geq 1$, and after all histories, follows from the fact that each agent's continuation payoff under truthful strategies coincides with his expected contribution to continuation weighted surplus, which is always nonnegative, scaled by a strictly positive weight.

The arguments above therefore establish that, when the matching rule is $\tilde{\chi}$ and the payments from period 1 onwards are $\tilde{\psi}_{\geq 1}$, participating and following truthful strategies constitutes a periodic ex-post continuation equilibrium, after any period-t history, $t \geq 1$. ■

Step 3. We now show that, when the period-0 membership fees are as in (9), participating in period zero and then following a truthful strategy at all periods (including period zero) is a periodic ex-post equilibrium.

Let

$$\tilde{D}^k_l(\theta^k_{l0},\theta;\chi) \equiv \begin{cases} \mathbb{E}^{\lambda[\chi]|\theta}\left[\sum_{t=1}^{\infty}\delta^t\sum_{h\in N_{-k}}\varepsilon^k_{lht}\chi_{lht}(\theta_0,\theta_t,b_t,x^{t-1})\right] & \text{if } k=A \\ \mathbb{E}^{\lambda[\chi]|\theta}\left[\sum_{t=1}^{\infty}\delta^t\sum_{h\in N_{-k}}\varepsilon^k_{hlt}\chi_{hlt}(\theta_0,\theta_t,b_t,x^{t-1})\right] & \text{if } k=B \end{cases},$$

denote the match quality that agent $l \in N_k$ from side $k = A, B$ expects under the rule χ when the true profile of vertical types is $\theta \in \Theta$, the agent selects the membership status θ^k_{l0} in period zero and then conforms to the truthful strategy from period $t = 1$ onwards, and all agents other than l (from side k) follow truthful strategies at each period. Note that $\lambda[\chi]|\theta$ denotes the stochastic process over matches, bids, and membership choices, when the true vertical types are θ, and all agents follow truthful strategies from period $t = 1$ onwards. Also note that $\tilde{D}^k_l(\theta^k_l,(\theta^k_{-l},\theta^k_l);\chi) = D^k_l(\theta;\chi)$, with the function $D^k_l(\theta;\chi)$ as defined in (10) — hereafter we highlight the dependence of the function $D^k_l(\theta;\chi)$ on the matching rule χ to avoid possible confusion.

For any agent $i \in N_A$ (the arguments for the side-B agents are analogous), let

$$\hat{U}^A_i(\theta) \equiv \mathbb{E}^{\lambda[\chi]|\theta}\left[\sum_{t=1}^{\infty}\delta^t\sum_{j\in N_B}\theta^A_i\varepsilon^A_{ijt}\chi_{ijt}(\theta,\theta_t,b_t,x^{t-1})\right] - \mathbb{E}^{\lambda[\chi]|\theta}\left[\sum_{t=0}^{\infty}\delta^t\psi^A_{it}(\theta,\theta_t,b_t,x^{t-1};\beta)\right] \tag{19}$$

denote the payoff that the agent expects in the matching auctions defined by the rules (χ,ψ) when the true vertical type profile is θ and all agents follow truthful strategies at all periods.[40]

Let $\tilde{\chi}$ be a matching rule that maximizes continuation weighted surplus and $\tilde{\psi} = (\tilde{\psi}_0,\tilde{\psi}_{\geq 1})$ the associated payment rule, as defined in the main text. Note that, for each agent $l \in N_k$, each profile θ

[40] The dependence of the payoff on the mechanism (χ,ψ) is omitted for convenience.

of true vertical types,

$$\mathbb{E}^{\lambda[\tilde{\chi}]|\theta}\left[\tilde{\psi}^k_{l0}(\theta;\beta)+\sum_{t=1}^{\infty}\delta^t\tilde{\psi}^k_{lt}(\theta,\theta_t,b_t,x^{t-1};\beta)\right]=\theta^k_l D^k_l(\theta;\tilde{\chi})-\int_{\underline{\theta}^k_l}^{\theta^k_l}D^k_l(\theta^k_{-l},y;\tilde{\chi})dy-L^k_l,$$

which guarantees that, when all agents follow truthful strategies in each period, including period zero, the period-zero expected payoffs, $\hat{U}^k_l(\theta)$, are given by

$$\hat{U}^k_l(\theta)=\int_{\underline{\theta}^k_l}^{\theta^k_l}D^k_l(\theta^k_{-l},y;\tilde{\chi})dy+L^k_l. \tag{20}$$

The next lemma shows that any matching rule $\tilde{\chi}$ that maximizes continuation weighted surplus (6) at all histories satisfies a certain monotonicity condition which plays a central role in establishing the optimality of the truthful strategies at period zero.

Lemma 7 *Suppose $\tilde{\chi}$ maximizes continuation weighted surplus (6) at all histories. For all $l \in N_k$, $k = A, B$, the following monotonicities hold:*
(i) $D^k_l\left((\theta^k_{-l},\theta^k_l);\tilde{\chi}\right)$ is non-decreasing in θ^k_l, all $\theta^k_{-l}\in\Theta^k_{-l}$;
(ii) $\tilde{D}^k_l\left(\theta^k_{l0},\theta;\tilde{\chi}\right)$ is non-decreasing in θ^k_{l0}, all $\theta\in\Theta$.

Proof of Lemma 7. Consider an arbitrary agent $i \in N_A$ from side A (the arguments for the side-B agents are analogous) and fix the profile of types θ^A_{-i} for the other agents.

Part (i). Take any pair of types $\theta^A_i, \hat{\theta}^A_i \in \Theta^A_i$, with $\theta^A_i < \hat{\theta}^A_i$. That $\tilde{\chi}$ maximizes continuation weighted surplus implies that[41]

$$\begin{aligned}
&\mathbb{E}^{\lambda[\tilde{\chi}]|\hat{\theta}^A_i,\theta^A_{-i}}\left[\sum_{t=1}^{\infty}\delta^t\sum_{r\in N_A\setminus\{i\}}\sum_{j\in N_B}\left(\beta^A_r(\theta^A_r)b^A_{rjt}+\beta^B_j(\theta^B_j)b^B_{rjt}-c_{rjt}(x^{t-1})\right)\tilde{\chi}_{rjt}\left((\theta^A_{-i},\hat{\theta}^A_i),(\theta^A_{-i},\hat{\theta}^A_i),b_t,x^{t-1}\right)\right]\\
&+\mathbb{E}^{\lambda[\tilde{\chi}]|\hat{\theta}^A_i,\theta^A_{-i}}\left[\sum_{t=1}^{\infty}\delta^t\sum_{j\in N_B}\left(\beta^A_i(\hat{\theta}^A_i)b^A_{ijt}+\beta^B_j(\theta^B_j)b^B_{ijt}-c_{ijt}(x^{t-1})\right)\tilde{\chi}_{ijt}\left((\theta^A_{-i},\hat{\theta}^A_i),(\theta^A_{-i},\hat{\theta}^A_i),b_t,x^{t-1}\right)\right]\\
&\geq\mathbb{E}^{\lambda[\tilde{\chi}]|\theta^A_i,\theta^A_{-i}}\left[\sum_{t=1}^{\infty}\delta^t\sum_{r\in N_A\setminus\{i\}}\sum_{j\in N_B}\left(\beta^A_r(\theta^A_r)b^A_{rjt}+\beta^B_r(\theta^B_r)b^B_{rjt}-c_{rjt}(x^{t-1})\right)\tilde{\chi}_{rjt}\left((\theta^A_{-i},\theta^A_i),(\theta^A_{-i},\theta^A_i),b_t,x^{t-1}\right)\right]\\
&+\mathbb{E}^{\lambda[\tilde{\chi}]|\theta^A_i,\theta^A_{-i}}\left[\sum_{t=1}^{\infty}\delta^t\sum_{j\in N_B}\left(\beta^A_i(\hat{\theta}^A_i)b^A_{ijt}\frac{\hat{\theta}^A_i}{\theta^A_i}+\beta^B_j(\theta^B_j)b^B_{ijt}-c_{ijt}(x^{t-1})\right)\tilde{\chi}_{ijt}\left((\theta^A_{-i},\theta^A_i),(\theta^A_{-i},\theta^A_i),b_t,x^{t-1}\right)\right].
\end{aligned}$$

The left-hand side of the previous inequality is the expected weighted surplus when all agents follow truthful strategies from period $t = 0$ onward and the true profile of vertical types is $(\hat{\theta}^A_i, \theta^A_{-i})$. The right-hand side is the expected weighted surplus when, under the same profile of true vertical types $(\hat{\theta}^A_i, \theta^A_{-i})$, all agents other than agent i from side A follow truthful strategies in all periods whereas agent i follows the strategy of type θ^A_i in all periods (that is, at each period he selects the membership status $\theta^A_{it} = \theta^A_i$ and then submits bids equal to $b^A_{ijt} = \theta^A_i\varepsilon^A_{ijt}$, where ε^A_{ijt} are the true horizontal types).

[41] Recall that $\lambda[\tilde{\chi}]|\theta^A_i,\theta^A_{-i}$ denotes the process over matches, bids, and membership choices from period $t = 1$ onwards, when all agents follow truthful strategies from period zero onwards, and the true vertical types are $(\theta^A_i, \theta^A_{-i})$.

Similarly, inverting the role of $\hat{\theta}_i^A$ and θ_i^A, we have that

$$\mathbb{E}^{\lambda[\tilde{\chi}]|\theta_i^A,\theta_{-i}^A}\left[\sum_{t=1}^{\infty}\delta^t \sum_{r\in N_A\setminus\{i\}}\sum_{j\in N_B}\left(\beta_r^A(\theta_r^A)b_{rjt}^A+\beta_j^B(\theta_j^B)b_{rjt}^B-c_{rjt}(x^{t-1})\right)\tilde{\chi}_{rjt}\left((\theta_{-i}^A,\theta_i^A),(\theta_{-i}^A,\theta_i^A),b_t,x^{t-1}\right)\right]$$

$$+\,\mathbb{E}^{\lambda[\tilde{\chi}]|\theta_i^A,\theta_{-i}^A}\left[\sum_{t=1}^{\infty}\delta^t \sum_{j\in N_B}\left(\beta_i^A(\theta_i^A)b_{ijt}^A+\beta_j^B(\theta_j^B)b_{ijt}^B-c_{ijt}(x^{t-1})\right)\tilde{\chi}_{ijt}\left((\theta_{-i}^A,\theta_i^A),(\theta_{-i}^A,\theta_i^A),b_t,x^{t-1}\right)\right]$$

$$\geq \mathbb{E}^{\lambda[\tilde{\chi}]|\hat{\theta}_i^A,\theta_{-i}^A}\left[\sum_{t=1}^{\infty}\delta^t \sum_{r\in N_A\setminus\{i\}}\sum_{j\in N_B}\left(\beta_r^A(\theta_r^A)b_{rjt}^A+\beta_j^B(\theta_j^B)b_{rjt}^B-c_{rjt}(x^{t-1})\right)\tilde{\chi}_{rjt}\left((\theta_{-i}^A,\hat{\theta}_i^A),(\theta_{-i}^A,\hat{\theta}_i^A),b_t,x^{t-1}\right)\right]$$

$$+\,\mathbb{E}^{\lambda[\tilde{\chi}]|\hat{\theta}_i^A,\theta_{-i}^A}\left[\sum_{t=1}^{\infty}\delta^t \sum_{j\in N_B}\left(\beta_i^A(\theta_i^A)b_{ijt}^A\frac{\theta_i^A}{\hat{\theta}_i^A}+\beta_j^B(\theta_j^B)b_{ijt}^B-c_{ijt}(x^{t-1})\right)\tilde{\chi}_{ijt}\left((\theta_{-i}^A,\hat{\theta}_i^A),(\theta_{-i}^A,\hat{\theta}_i^A),b_t,x^{t-1}\right)\right].$$

Combining the last two inequalities, and using the fact that horizontal and vertical types are independent, and that vertical types are drawn independently across agents, we have that

$$\left(\beta_i^A(\hat{\theta}_i^A)\hat{\theta}_i^A-\beta_i^A(\theta_i^A)\theta_i^A\right)\cdot\left(D_i^A((\theta_{-i}^A,\hat{\theta}_i^A);\tilde{\chi})-D_i^A((\theta_{-i}^A,\theta_i^A);\tilde{\chi})\right)\geq 0.$$

Because $\beta_i^A(\cdot)$ is strictly positive and non-decreasing, it must be that

$$D_i^A((\theta_{-i}^A,\hat{\theta}_i^A);\tilde{\chi})\geq D_i^A((\theta_{-i}^A,\theta_i^A);\tilde{\chi}).$$

Part (ii). Since $\tilde{\chi}$ maximizes continuation weighted surplus after any history, arguments similar to those used to establish part (i) above imply that, for any $\hat{\theta}_{i0}^A,\theta_{i0}^A\in\Theta_i^A$, $\theta\in\Theta$,

$$\left(\beta_i^A(\hat{\theta}_{i0}^A)-\beta_i^A(\theta_{i0}^A)\right)\cdot\theta_i^A\cdot\left(\tilde{D}_i^A(\hat{\theta}_{i0}^A,\theta;\tilde{\chi})-\tilde{D}_i^A(\theta_{i0}^A,\theta;\tilde{\chi})\right)\geq 0,$$

Because $\theta_i^A>0$ and $\beta_i^A(\cdot)$ is strictly positive and non-decreasing, it must be that[42]

$$\tilde{D}_i^A(\hat{\theta}_{i0}^A,\theta;\tilde{\chi})\geq\tilde{D}_i^A(\theta_{i0}^A,\theta;\tilde{\chi}).$$

■

Next, for any agent $i\in N_A$ (the arguments for the side-B agents are analogous), let

$$\tilde{U}_i^A(\theta_{i0}^A;\theta)\equiv\mathbb{E}^{\lambda[\chi]|\theta}\left[\sum_{t=1}^{\infty}\delta^t\sum_{j\in N_B}\theta_i^A\varepsilon_{ijt}^A\chi_{ijt}((\theta_{-i}^A,\theta_{i0}^A),\theta_t,b_t,x^{t-1})\right]-\mathbb{E}^{\lambda[\chi]|\theta}\left[\sum_{t=0}^{\infty}\delta^t\psi_{it}^A((\theta_{-i}^A,\theta_{i0}^A),\theta_t,b_t,x^{t-1};\beta)\right]$$

denote the payoff that the agent expects in the auctions defined by the rules (χ,ψ), when the true vertical type profile is θ, the agent chooses the membership status θ_{i0}^A in period zero, he follows a truthful strategy from period $t=1$ onwards, and all other agents follow truthful strategies from period $t=0$ onwards.[43] Note that $\tilde{U}_i^A(\theta_i^A;(\theta_{-i}^A,\theta_i^A))=\hat{U}_i^A(\theta_{-i}^A,\theta_i^A)$, where $\hat{U}$ is as in (20).

[42] Note that if $\beta_i^A(\hat{\theta}_{i0}^A)=\beta_i^A(\theta_{i0}^A)$, then $\tilde{D}_i^A(\hat{\theta}_{i0}^A,\theta;\tilde{\chi})=\tilde{D}_i^A(\theta_{i0}^A,\theta;\tilde{\chi})$.

[43] The dependence of the payoff on the mechanism (χ,ψ) is omitted for convenience.

From Step 2, under the rules $(\tilde{\chi}, \tilde{\psi})$, participating and following truthful strategies is a periodic ex-post continuation equilibrium starting from any period-1 history (including those reached off path, by deviations in period zero).[44] Standard arguments can then be used to show that the following envelope condition must be satisfied for all $l \in N_k$, $k = A, B$, all $\theta_{l0}^k \in \Theta_l^k$, all $\theta \in \Theta$,

$$\tilde{U}_l^k(\theta_{l0}^k; \theta) = \tilde{U}_l^k(\theta_{l0}^k; (\theta_{-l}^k, \underline{\theta}_l^k)) + \int_{\underline{\theta}_l^k}^{\theta_l^k} \tilde{D}_l^k(\theta_{l0}^k, (\theta_{-l}^k, y); \tilde{\chi}) dy. \tag{21}$$

The payoff that agent $l \in N_k$ from side $k = A, B$ obtains by selecting the membership θ_{l0}^k when the true vertical type profile is θ is thus given by

$$\tilde{U}_l^k(\theta_{l0}^k; \theta) = \tilde{U}_l^k(\theta_{l0}^k; (\theta_{-l}^k, \theta_{l0}^k)) + \int_{\theta_{l0}^k}^{\theta_l^k} \tilde{D}_l^k(\theta_{l0}^k, (\theta_{-l}^k, y); \tilde{\chi}) dy \leq \tilde{U}_l^k(\theta_{l0}^k; (\theta_{-l}^k, \theta_{l0}^k)) + \int_{\theta_{l0}^k}^{\theta_l^k} \tilde{D}_l^k(y, (\theta_{-l}^k, y); \tilde{\chi}) dy$$

$$= \hat{U}_l^k(\theta_{-l}^k, \theta_{l0}^k) + \int_{\theta_{l0}^k}^{\theta_l^k} D_l^k((\theta_{-l}^k, y); \tilde{\chi}) dy = \hat{U}_l^k(\theta),$$

where the first equality follows from (21), the inequality follows from part (ii) in Lemma 7, and the other equalities follow from (20) and the definition of the interim expected payoffs. Hence, given θ, the agent is better off following a truthful strategy from period zero onwards than deviating in period zero and then following a truthful strategy from period one onwards.

Finally, note that, since for any $l \in N_k$, $k = A, B$, any $\theta \in \Theta$, $\hat{U}_l^k(\theta)$ is bounded, participation constraints can always be satisfied by choosing the constants L_l^k appropriately.

Combining the results in Step 2 with those in Step 1, we thus have that, when the scoring rules and the associated payment functions satisfy the conditions in the theorem, participating and following a truthful strategy is a periodic ex-post equilibrium in the entire game. Q.E.D.

Proof of Theorem 2. Consider any feasible mechanism Γ and any BNE σ of the game induced by Γ. Denote by $\hat{\chi} = (\hat{\chi}_t(\theta, \varepsilon^t))_{t=1}^{\infty}$ and $\hat{\psi} = (\hat{\psi}_t(\theta, \varepsilon^t))_{t=1}^{\infty}$ the matching and payment rules, as a function of the true state, induced by σ in Γ. As mentioned in the main text, in case of endogenous processes, these functions are defined only for histories that are consistent with equilibrium play in previous periods. Also 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

$$\mathbb{E}^{\lambda[\hat{\chi}]} \left[\sum_{k=A,B} \sum_{l \in N_k} \sum_{t=0}^{\infty} \delta^t \hat{\psi}_{lt}^k(\theta, \varepsilon^t) - \sum_{t=1}^{\infty} \delta^t \sum_{i \in N_A} \sum_{j \in N_B} c_{ijt}(\hat{\chi}^{t-1}(\theta, \varepsilon^{t-1})) \hat{\chi}_{ijt}(\theta, \varepsilon^t) \right], \tag{22}$$

where $\lambda[\hat{\chi}]$ denotes the process over vertical and horizontal types under the matching rule $\hat{\chi}$ induced by the strategies σ in Γ. Alternatively, (22) can be rewritten as follows:

$$\mathbb{E}^{\lambda[\hat{\chi}]} \left[\sum_{t=1}^{\infty} \sum_{i \in N_A} \sum_{j \in N_B} \delta^t \left(\left(\theta_i^A \varepsilon_{ijt}^A + \theta_j^B \varepsilon_{ijt}^B - c_{ijt}(\hat{\chi}^{t-1}(\theta, \varepsilon^{t-1})) \right) \hat{\chi}_{ijt}(\theta, \varepsilon^t) \right) - \sum_{k=A,B} \sum_{l \in N_k} U_l^k(\theta_l^k) \right], \tag{23}$$

[44] Recall that $\tilde{\chi}$ maximizes continuation weighted surplus and that the associated payment rule $\tilde{\psi}$ satisfies the conditions in (8) and (9) in the main text. From Step 1, $\tilde{\chi}$ coincides with either $\chi^{m;\beta}$, or with $\chi^{I;\beta}$, depending on whether the environment is the one in part (i) or part (ii) in the theorem.

where $U_l^k(\theta_l^k)$ denotes the period-0 interim expected payoff of agent $l \in N_k$ from side $k = A, B$ when his vertical type is θ_l^k, under the equilibrium σ in the mechanism Γ. Note that we denote the interim payoffs by U_l^k do differentiate them from the interim payoff functions $\hat{U}_l^k$ in the auction (defined in (19)), where the domain of these latter functions is the entire profile of vertical types.

The period-0 participation constraints are satisfied if for all $l \in N_k$, $k = A, B$, $\theta_l^k \in \Theta_l^k$,

$$U_l^k(\theta_l^k) \geq 0.$$

Following an approach similar to the one in Pavan, Segal and Toikka (2014, Theorem 1), we can show that the period-0 (interim) expected payoff of each agent $l \in N_k$, $k = A, B$ must satisfy the following envelope condition:

$$U_l^k(\theta_l^k) = U_l^k(\underline{\theta}_l^k) + \int_{\underline{\theta}_l^k}^{\theta_l^k} \mathbb{E}\left[D_l^k(\theta; \hat{\chi})|y\right] dy, \tag{24}$$

where the expectation is taken over the entire profile of vertical types θ given agent l's own vertical type. This envelope condition, together with integration by parts, yields the following representation of the platform's profits,

$$\mathbb{E}^{\lambda[\hat{\chi}]}\left[\sum_{t=1}^{\infty}\delta^t \sum_{i\in N_A}\sum_{j\in N_B}\left(\left(\theta_i^A - \frac{1-F_i^A(\theta_i^A)}{f_i^A(\theta_i^A)}\right)\varepsilon_{ijt}^A + \left(\theta_j^B - \frac{1-F_j^B(\theta_j^B)}{f_j^B(\theta_j^B)}\right)\varepsilon_{ijt}^B \right.\right.$$
$$\left.\left. - c_{ijt}(\hat{\chi}^{t-1}(\theta, \varepsilon^{t-1}))\right)\hat{\chi}_{ijt}(\theta, \varepsilon^t)\right] - \sum_{k=A,B}\sum_{l\in N_k} U_l^k(\underline{\theta}_l^k). \tag{25}$$

The first term, which is only a function of the matching rule $\hat{\chi}$, is the *expected dynamic virtual surplus* (DVS) generated by the matching rule $\hat{\chi}$. Clearly, because such a representation applies to the matching rule generated by any BNE of any mechanism, the above representation applies also to the state-contingent matching rules generated by the truthful strategies in the matching auctions.

Now observe that, when the weights are given by $\hat{\beta} \equiv (\hat{\beta}_l^k(\cdot))_{l\in N_k, k=A,B}$, the state-contingent matches implemented under the truthful equilibria of the auctions of Theorem 1 maximize DVS over all feasible state-contingent rules $\hat{\chi}$. More precisely, when the processes are exogenous, the matches implemented under a myopic scoring rule maximize DVS. Similarly, when the processes are endogenous and either $M = 1$, or $M \geq n_A \cdot n_B$, or the environment is separable, the matches implemented under the index rule maximize DVS. This is because (i) these matches have been shown to maximize the continuation weighted surplus (6) for any strictly positive and non-decreasing weights $\beta = (\beta_l^k(\cdot))_{l\in N_k, k=A,B}$ (step 1 in the proof of Theorem 1), (ii) the ex-ante weighted surplus when the weights are given by $\hat{\beta}$ coincides with DVS, and (iii) the weights $\hat{\beta}$ are strictly positive and non-decreasing by Assumption 1.

Next, suppose the processes are exogenous. Let $\psi^{m;\hat{\beta}}$ be the payment scheme associated with the myopic scoring rule $\chi^{m;\hat{\beta}}$ with weights given by $\hat{\beta}$ (as defined by (8) and (9) in the main text) and in which $L_l^k = 0$, all $l \in N_k$, $k = A, B$. From (8) and (9), it is easy to see that, under the equilibria in truthful strategies of the matching auctions where (a) the scoring rule is the myopic one with weights given by $\hat{\beta}$, and (b) the payments are $\psi^{m;\hat{\beta}}$, the payoff expected by the lowest vertical type of each agent is exactly equal to zero (that is, $U_l^k(\underline{\theta}_l^k) = 0$, all $l \in N_k, k = A, B$). This means that the truthful equilibria of the above matching auctions maximize both terms of (25). Provided all the period-0 participation constraints are satisfied (something we verify below), we then have that the platform's profits are maximized under the truthful equilibria of the proposed matching auctions.

Also note that, while we have restricted attention to deterministic mechanisms, the platform cannot increase its profits by using a randomized mechanism. This is because any randomized mechanism is equivalent to a deterministic one that conditions on the type realizations of a fictitious agent. The platform's profits under any equilibrium of such a mechanism thus continue to be given by the expression in (25), but with the matching rule conditioning on the type realizations of such fictitious agent. This means that the platform's profits are equal to the weighted average of the platform's profits under the deterministic matching rules obtained by conditioning on the various types of the fictitious agent. Because (25) is maximized over all possible deterministic rules under the equilibria in truthful strategies of the proposed matching auctions, we thus have that stochastic mechanisms can never improve upon the equilibria of the proposed auctions when it comes to the platform's profits.

We now complete the proof by establishing that all period-0 participation constraints are satisfied under the equilibria in truthful strategies of the proposed auctions. To see this, it suffices to observe that, for all $\theta_{-l}^k \in \Theta_{-l}^k$, $l \in N_k$, $k = A, B$, $D_l^k((\theta_{-l}^k, \theta_l^k); m, \hat{\beta})$ is non-decreasing in θ_l^k, as established in Lemma 7 above. The assumption in the Theorem that $D_l^k((\theta_{-l}^k, \underline{\theta}_l^k); m, \hat{\beta}) \geq 0$, all $l \in N_k$, $k = A, B$, all $\theta_{-l}^k \in \Theta_{-l}^k$ then guarantees that $D_l^k(\theta; m, \hat{\beta}) \geq 0$, all $\theta \in \Theta$, $l \in N_k$, $k = A, B$. Because the period-0 interim payoffs satisfy the envelope condition (20), we then have that $\hat{U}_l^k(\theta) \geq 0$, all $\theta \in \Theta$, $l \in N_k$, $k = A, B$, which means that all the period-0 participation constraints are satisfied (in a periodic ex-post sense, i.e., for any θ, and not just in expectation over θ_{-l}^k given θ_l^k).[45]

The optimality of the matching auctions with index rules in part (ii) of the theorem follow from similar arguments. Q.E.D.

Proof of Theorem 3. As explained in the main text, parts (i) and (ii) follow from Theorem 1. Part (iii) follows from arguments similar to those establishing the optimality of the auctions of Theorem 2. In particular, it follow from the fact that (a) the platform's expected profits under any BNE of any mechanism Γ implementing the welfare-maximizing matches satisfy the representation in (25); (b) each agent's period-0 expected payoff satisfies Condition (24); (c) in the proposed auctions, $U_l^k(\underline{\theta}_l^k) = 0$ if, and only if, the payments in (8) and (9) (for $\beta = \beta^W$) are such that $L_l^k = 0$, all $l \in N_k$, $k = A, B$, and (d) when the payments are given by (8) and (9) with $\beta = \beta^W$ and $L_l^k = 0$, all $l \in N_k$, $k = A, B$, all agents' period-0 participation constraints are satisfied, regardless of their beliefs over other agents' types, if, and only if, $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; m, \beta^W) \geq 0$, or $D_l^k(\theta_{-l}^k, \underline{\theta}_l^k; I, \beta^W) \geq 0$ (under exogenous or endogenous processes, respectively). The latter property is a result of the fact that the functions $D_l^k(\theta_{-l}^k, \cdot; m, \beta^W)$ and $D_l^k(\theta_{-l}^k, \cdot; I, \beta^W)$ are non-decreasing, which follows from Lemma 7. Q.E.D.

[45]Note that, under $\psi^{m;\hat{\beta}}$, $\hat{U}_l^k(\theta_{-l}^k, \underline{\theta}_l^k) = 0$ all θ_{-l}^k.

Proof of Theorem 4. Let $\chi_t^P(\theta,\omega)$ and $\chi_t^W(\theta,\omega)$ denote the state-contingent matches implemented in period $t \geq 1$, under the truthful equilibria of, respectively, the profit-maximizing and the welfare-maximizing auctions of Theorems 2 and 3. Note that the arguments of these functions are the exogenous vertical types θ, and the sequences of exogenous innovations $\omega \equiv (\omega_{ijs}^k)_{(i,j)\in N_A\times N_B, k\in\{A,B\}}^{s=1,\ldots,\infty}$ that, along with the matches implemented in previous periods (in case of endogenous processes), generate the horizontal types ε. As explained in the main text, such a representation favors the comparison of the matches sustained under the two auctions by making the "state" exogenous, thus eliminating the confusion that may originate from the fact that the histories of horizontal types need not coincide under the two auctions.

Similarly, let $S_{ijt}^{m;P}(\theta,\omega)$ and $S_{ijt}^{I;P}(\theta,\omega)$ denote, respectively, the period-t state-contingent myopic and index scores under the truthful equilibria of the profit-maximizing auctions of Theorem 2. Likewise, let $S_{ijt}^{m;W}(\theta,\omega)$ and $S_{ijt}^{I;W}(\theta,\omega)$ be the counterparts of $S_{ijt}^{m;P}(\theta,\omega)$ and $S_{ijt}^{I;P}(\theta,\omega)$ under the truthful equilibria of the welfare-maximizing auctions of Theorem 3. Because (θ,ω) are exogenous and time-invariant, they are dropped from all the functions χ^P, χ^W, $S^{m;P}$, $S^{m;W}$, $S^{I;P}$, and $S^{I;W}$ below.

First, observe that, because $\hat{\beta}_l^k(\theta_l^k) \leq 1 = \beta_l^{k,W}(\theta_l^k)$, all $\theta_l^k \in \Theta_l^k$, $l \in N_k, k = A, B$, and because the horizontal types are nonnegative, when processes are exogenous, for any $(i,j) \in N_A \times N_B$, $S_{ijt}^{m;P} \leq S_{ijt}^{m;W}$. Likewise, when processes are endogenous, for any $(i,j) \in N_A \times N_B$, any $t,\tau \geq 1$,

$$\sum\nolimits_{s=1}^{t-1} \chi_{ijs}^W = \sum\nolimits_{s=1}^{\tau-1} \chi_{ijs}^P \;\Rightarrow\; S_{ijt}^{I;W} \geq S_{ij\tau}^{I;P}. \tag{26}$$

Part 1. First, consider the case of exogenous processes. Because the capacity constraint is not binding, in each period $t \geq 1$, the matches implemented under the equilibria of the profit-maximizing auctions (alternatively, of the welfare-maximizing auctions) are all those for which the myopic scores $S_{ijt}^{m;P} \geq 0$ (alternatively, $S_{ijt}^{m;W} \geq 0$). The above property, along with the fact that, for any $(i,j) \in N_A \times N_B$, $t \geq 1$, $S_{ijt}^{m;W} \geq S_{ijt}^{m;P}$, then yields the result.

Next, consider the case of endogenous processes. Again, because the capacity constraint is not binding, in each period $t \geq 1$, the matches implemented under the equilibria of the profit-maximizing auctions (alternatively, the welfare-maximizing auctions) are all those for which the index $S_{ijt}^{I;P} \geq 0$ (alternatively, $S_{ijt}^{I;W} \geq 0$). The result then follows from the fact that, for any $(i,j) \in N_A \times N_B$, $t \geq 1$, $S_{ijt}^{m;W} \geq S_{ijt}^{m;P}$. The last property, in turn, follows by induction. First observe that the property is necessarily true at $t = 1$, given (26) and the fact that, in period $t = 1$, the number of past interactions is necessarily the same under profit and welfare maximization. Now suppose the result holds for all $1 \leq s < t$. Note that any match for which $S_{ijt}^{I;P} \geq 0$ has been active at each preceding period $s < t$, both under profit maximization and under welfare maximization. The result then follows again from (26), which implies that $S_{ijt}^{I;W} \geq S_{ijt}^{I;P}$.

Part 2. Consider first the case of exogenous processes and $M < n_A \cdot n_B$. The result follows directly from the following two properties: (a) in each period $t \geq 1$, the set of matches for which $S_{ijt}^{m;P} \geq 0$ is a subset of the set of matches for which $S_{ijt}^{m;W} \geq 0$, (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 $S_{ijt}^{m;P} \geq 0$ (alternatively, $S_{ijt}^{mW} \geq 0$).

Next, consider the case of endogenous processes with $M = 1$. First observe that, under the equilibria of the profit-maximizing auction, if at some period $t \geq 1$, $\chi_{ijt}^{P} = 0$, all $(i,j) \in N_A \times N_B$, then $\chi_{ijs}^{P} = 0$, all $s > t$, all $(i,j) \in N_A \times N_B$. The same property holds for χ^W. Next, observe that, if matching stops at period t under profit maximization (alternatively, welfare maximization), then $S_{ijt}^{I;P} < 0$ all $(i,j) \in N_A \times N_B$ (alternatively, $S_{ijt}^{I;W} < 0$ all $(i,j) \in N_A \times N_B$). Now suppose that, under profit maximization, matching is still active in period t (meaning, there exists $(i,j) \in N_A \times N_B$ such that $\chi_{ijt}^{P} = 1$). Then there are two cases. (1) Either $\sum_{s=1}^{t-1} \chi_{ijs}^{W} = \sum_{s=1}^{t-1} \chi_{ijs}^{P}$, all $(i,j) \in N_A \times N_B$, in which case (26) implies that $S_{ijt}^{I;W} \geq S_{ijt}^{I;P}$ for all $(i,j) \in N_A \times N_B$, which implies the result. Or, (2) there exists $(i,j) \in N_A \times N_B$ such that $\sum_{s=1}^{t-1} \chi_{ijs}^{W} < \sum_{s=1}^{t-1} \chi_{ijs}^{P}$. In this case, there must exist $\tau < t$ such that $\sum_{s=1}^{\tau-1} \chi_{ijs}^{P} = \sum_{s=1}^{t-1} \chi_{ijs}^{W}$, and $\chi_{ij\tau}^{P} = 1$. That $\chi_{ij\tau}^{P} = 1$ in turn implies $S_{ij\tau}^{I;P} \geq 0$. That $\sum_{s=1}^{\tau-1} \chi_{ijs}^{P} = \sum_{s=1}^{t-1} \chi_{ijs}^{W}$ in turn implies that

$$S_{ijt}^{I;W} \geq S_{ij\tau}^{I;P}, \tag{27}$$

where the result follows again from (26). From the discussion above, (27) in turn implies that matching must be active in period t also under welfare maximization.

Part 3. Since, under profit maximization, matching terminates after a finite number of periods $T \in \mathbb{N}$, at period $T+1$, the index $S_{ijT+1}^{I;P} < 0$ all $(i,j) \in N_A \times N_B$. Take any pair (i,j) and denote the number of times the match (i,j) has been active under profit maximization prior to period $T+1$ by R_{ij}^{P}. For any $n \leq R_{ij}^{P}$, the index $S_{ij}^{I;P}$ at the n-th time the pair was matched must have been nonnegative. The same must then be true for the index $S_{ij}^{I;W}$ (this follows again from (26)). Therefore, matching will not stop under welfare maximization until the pair (i,j) is matched at least R_{ij}^{P} times. Since this holds for each pair $(i,j) \in N_A \times N_B$, the result in Part (3) follows. Q.E.D.

References

[1] Akbarpour, M., Li, S., and S. Oveis Gharan (2016), "Thickness and Information in Dynamic Matching Markets," mimeo Stanford University.

[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), pp. 23-68.

[3] Anderson, R., Ashlagi, I., Gamarnik, D., and Y. Kanoria (2015), "Efficient Dynamic Barter Exchange," mimeo, MIT.

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

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

[6] Athey, S. and I. Segal (2013), "An Efficient Dynamic Mechanism," *Econometrica*, 81(6), 2463–2485.

[7] Baccara, M., S. Lee, and L. Yariv (2015): "Optimal Dynamic Matching," mimeo Caltech.

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

[9] Bedre-Defolie, O. and E. Calvano, (2013),"Pricing Payment Cards," *American Economic Journal: Microeconomics*, Vol. 5(3), pp. 206-231.

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

[11] Bergemann, D. and M. Said (2011), "Dynamic Auctions," *Wiley Encyclopedia of Operations Research and Management Science.*

[12] Bergemann, D. and P. Strack (2015), "Dynamic Revenue Maximization: A Continuous Time Approach," *Journal of Economic Theory* 159, 819-853.

[13] Bergemann, D. and J. Välimäki (2001), "Stationary Multi-Choice Bandit Problems," *Journal of Economic dynamics and Control,* 25(10), 1585-1594.

[14] Bergemann, D. and J. Välimäki (2008), "Bandit Problems," *New Palgrave Dictionary of Economics*, Durlauf and Blume eds.

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

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

[17] Bloch, F. and D. Cantala (2016), "Dynamic Allocation of Objects to Queuing Agents," *American Economic Journal: Microeconomics*, forthcoming.

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

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

[20] Bonatti, A. and J. Horner (2015), "Learning to Disagree in a Game of Experimentation," mimeo Yale University.

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

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

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

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

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

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

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

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

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

[30] Doepke, M and R.M. Townsend, (2006) "Dynamic Mechanism Design with Hidden Income and Hidden Actions," *Journal of Economic Theory* 126, 235 – 285.

[31] Doval, L. (2016), "A Theory of Stability in Dynamic Matching Markets," mimeo Northwestern University.

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

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

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

[35] Gittins, J. C. (1979), "Bandit processes and dynamic allocation indices," *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2):148–177.

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

[37] Gomes, R. and A. Pavan (2015), "Many-to-Many Matching and Price Discrimination," *Theoretical Economics*, forthcoming.

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

[39] Herbst, H., and B. Schickner (2016), "Dynamic Formation of Teams: When Does Waiting for Good Matches Pay Off?,"mimeo University of Bonn.

[40] Horner, J., and A. Skrzypacz (2016), "Learning, Experimentation and Information Design,"mimeo Yale University.

[41] Ishikida, T., and P. Varaiya (1994), "Multi-Armed Bandit Problem Revisited," *Journal of Optimization Theory and Applications,* 83(1), 113-154.

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

[43] Jullien, B. and A. Pavan (2016), "Platform Pricing under Dispersed Information," mimeo Northwestern University and Toulouse School of Economics.

[44] Kakade, S., Lobel, I., and H. Nazerzadeh (2011), "Optimal Dynamic Mechanism Design and the Virtual Pivot Mechanism," *Operations Research*, 61, 837–854.

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

[46] Kapicka, M. (2013), "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach," *Review of Economic Studies*, 80(3), 1027-1054.

[47] Keller, G. and S. Rady (2015), "Breakdowns," *Theoretical Economics,* 10(1), 175-202.

[48] Keller, G., Rady, S. and M. Cripps (2005), "Strategic experimentation with exponential bandits," *Econometrica,* 73(1), 39-68.

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

[50] Lee, R. (2013), "Vertical Integration and Exclusivity in Platform and Two-Sided Markets," *American Economic Review*, Vol. 103(7), pp. 2960-3000.

[51] Leshno, J. (2015). "Dynamic Matching in Overloaded Waiting Lists," mimeo Columbia University.

[52] Li, S., M. Mahdian, and P. McAfee (2010) "Value of Learning in Sponsored Search Auctions," *Proceedings of the 6th international Workshop on Internet and Network Economics* (WINE).

[53] Makris, M., and A. Pavan (2015), "Incentives for Endogeneous Types: Taxation under Learning by Doing," mimeo Northwestern University.

[54] Mandelbaum, A. (1986), "Discrete Multi-Armed Bandits and Multi-Parameter Processes," *Probability Theory and Related Fields,* 71(1), 129-147.

[55] Pandelis, D. G., and D. Teneketzis (1999), "On the Optimality of the Gittins Index Rule for Multi-Armed Bandits with Multiple Plays," *Mathematical Methods of Operations Research,* 50(3), 449-461.

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

[57] Rysman, M., (2009),"The Economics of Two-Sided Markets," *Journal of Economic Perspectives*, Vol. 23(3), pp. 125-143.

[58] Schummer, J., (2016), "Influencing Waiting Lists," mimeo Northwestern University.

[59] Thakral, N., (2015), "Matching with Stochastic Arrival: Theory and an Application to Public Housing," mimeo Harvard University.

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

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

[62] Weber, R. R., and G. Weiss (1990), "On an Index Policy for Restless Bandits," *Journal of Applied Probability,* 637-648.

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

[64] Whittle, P. (1982), "Optimization Over Time," John Wiley & Sons, Inc.

[65] Whittle, P. (1988), "Restless Bandits: Activity Allocation in a Changing World," *Journal of applied probability,* 287-298.