ELSEVIER

Incentives in decentralized random matching markets

Joana Pais

ISEG/Technical University of Lisbon and UECE (Research Unit on Complexity in Economics), Rua Miguel Lupi, 20, 1249-078 Lisbon, Portugal

Abstract

Decentralized markets are modeled by means of a sequential game where, starting from any matching situation, firms are randomly given the opportunity to make job offers. In this random context, we prove the existence of ordinal subgame perfect equilibria where firms act according to a list of preferences. Moreover, every such equilibrium preserves stability for a particular profile of preferences. In particular, when firms best reply by acting truthfully, every equilibrium outcome is stable for the true preferences. Conversely, when the initial matching is the empty matching, every stable matching can be reached as the outcome of an ordinal equilibrium play of the game.
© 2008 Elsevier Inc. All rights reserved.

1. Introduction

The study of centralized markets has been privileged in the two-sided matching literature. The introduction of centralized matching procedures in markets that experienced certain kinds of failures is partially responsible for such dedication. In fact, a number of markets—for physicians, lawyers, dentists, and osteopaths, among others—have adopted central clearinghouses after periods of uncontrolled unraveling of appointment dates and chaotic recontracting.[1] These markets now work by having each agent of the two sides of the market submit a rank ordered preference list of acceptable matches to the central clearinghouse, which then produces a matching by

---

E-mail address: jpais@iseg.utl.pt.

[1] See Roth and Xing (1994) and Niederle and Roth (2003).

processing all the preference lists according to an algorithm. Roth (1984a, 1991) showed that the algorithms used in most of the successful clearinghouses roughly follow the lines of Gale and Shapley's deferred acceptance algorithm (Gale and Shapley, 1962). This procedure generates a matching of workers to positions that is stable in terms of the submitted preferences in the sense that no worker and firm that are not matched to each other would prefer to be so matched.[2]

In contrast, decentralized markets have received relatively little attention.[3] The exact set of rules that governs a centralized market, making it particularly amenable to analysis, is no longer present when matching is organized in a decentralized way. Namely, decentralized decision making in complex environments may introduce randomness in the order in which job offers are made. The speed of the mail, the telephone network, or the internal structure of firms making some react faster than others determine the success in establishing communication with the desired workers. Moreover, decentralized markets involve different strategic issues from those of centralized markets. In fact, when a clearinghouse exists, agents must simply decide what preference lists to submit to the matchmaker, after which the match is created. However, in a decentralized market agents do not submit lists; instead, they can decide, after each interview or telephone call what to do next. The size of the strategy space is thus extremely large and, together with the underlying uncertainty, has precluded analysis by means of standard matching tools.

The purpose of this paper is to apply the extremely simple marriage model to the study of decentralized labor markets. In this model, agents' preferences are ordinal in nature. Hence, we merely take for granted their ability to order the potential working partners, without assuming that they are capable of giving cardinal content to their preferences.[4] The starting point of the analysis is any matching situation, providing a framework to the study of both entry-level and senior level markets. The matching process is then modeled as an extensive form game, where firms sequentially offer their positions. The inherently uncertain features of the market are modeled as chance moves that determine the order of play: at each moment in time, any firm—even if already matched—is randomly selected and given the opportunity to offer its position to a worker. This worker compares it with any offer he may be holding and rejects one, while (temporarily) holding the other, pending the possible arrival of even better offers. We assume that, once rejected, the firm is not willing to propose to the same worker again, but it may obviously offer its position to a different worker when given the opportunity to act.

There are different options when it comes to modeling the functioning of a decentralized labor market. A fundamental feature is that an agent's strategy does not necessarily comply with a list of preferences. In this model, for instance, a worker may entertain the strategy of accepting the first proposal received and rejecting all subsequent proposals, which is not consistent with a list of preferences for different orders of proposals. This, together with the fact that the initial matching is not, necessarily, the empty matching, where every agent is unmatched, assures that the process we describe cannot be reduced to a mere implementation of Gale and Shapley's algorithm. Moreover, chance moves introduce uncertainty in which matchings are achieved: it may happen that starting with the same initial matching, different plays of the game yield different outcomes for the same strategy profile.[5] This contrasts with two other papers aimed at modeling decentralized labor markets: Blum et al. (1997), which studies the vacancy chain problem and where chance

---

[2] See Roth and Sotomayor (1990) for a comprehensive study of two-sided matching markets.

[3] There are notable exceptions, namely Blum et al. (1997), Haeringer and Wooders (2004), Roth and Vate (1991), Roth and Xing (1997), among others.

[4] It follows that monetary transfers are embodied in agents' preferences, i.e., there is more to a job than just a salary.

[5] See Example 1 below for illustration.

moves do not add any randomness to the outcome, and Haeringer and Wooders (2004), where there is no uncertainty whatsoever, suggesting that markets are small worlds. Another dissimilar feature concerns agents' commitment: in our setting agents are greedy and always willing to look for a better partner, while in Blum et al. (1997) only firms with vacancies can actually make proposals and in Haeringer and Wooders (2004), once a proposal is accepted, the agents involved may simply leave the market. Finally, the results in this paper are proven for the low information scenario in Blum et al. (1997)—where each firm knows which of her offers have been rejected and by whom, and each worker recalls he proposals he obtained—but we briefly discuss an extension to the case in which agents know every proposal made, accepted, and rejected, which is close to the perfect information scenario in Haeringer and Wooders (2004). Hence, this paper does not compete with the cited literature, but rather complements it by providing yet another perspective on the functioning of decentralized markets.

The study of equilibria requires us to consider agents' preferences over lotteries. Since preferences are ordinal—and no natural utility representation of these preferences exists—we cannot compute expected utilities to compare different probability distributions over matchings. We thus adopt the strong solution concept of ordinal Nash equilibrium, which guarantees that each agent is an expected utility maximizer for *every* utility representation of his preferences.[6] We go beyond this concept to account for the dynamic nature of the game and characterize subgame perfect ordinal Nash equilibria. Despite the strength of this concept, we prove the existence of subgame perfect ordinal Nash equilibria and, in particular, equilibria where firms use preference strategies (i.e., strategies that can, up to some point, be identified with a list of preferences). On the other hand, every such equilibrium delivers matchings that are stable with respect to a particular profile of preferences. This has two appealing implications. First, for any equilibrium where firms adhere to preference lists, all outcomes are such that the set of unmatched agents is the same, so that no agent ending up unmatched can claim it was due to an instance of misfortune. Second, in the particular case that firms act according to their true preferences, stability with respect to the true preferences is guaranteed in a subgame perfect ordinal Nash equilibrium. The (partially) converse statement holds when we start from a situation where all agents are unmatched: every stable matching for the true preferences can be reached as the outcome of an equilibrium play where firms act straightforwardly according to their true preferences.

These results may be viewed as an extension of some known features of the game induced by Gale and Shapley's centralized mechanism (Roth, 1984b), where the underlying strategy space is confined to the set of preference strategies and the initial matching is the empty matching, and of those in Blum et al. (1997) when the initial matching is firm-quasi-stable, i.e., a matching whose stability has been disrupted by the creation of a new position or the retirement of a worker.[7] Moreover, it provides a direction into which the reason behind the success of some decentralized markets may be sought: equilibrium entails stability of the final matching—and no individual agent or pair of agents (consisting of a firm and a worker) has the incentive to circumvent it—as long as firms best reply by acting straightforwardly. Finally, even though focusing on equilibria where firms use lists of preferences as strategies or even reveal their true preferences can be considered a shortcoming of the analysis, such forms of behavior may be justified in some settings.

---

[6] This concept was introduced in d'Aspremont and Peleg (1988); it has been used in the context of voting theory in Majumdar and Sen (2004) and in matching markets in Ehlers and Massó (2007), Majumdar (2003), and Pais (2006, 2008).

[7] As for Haeringer and Wooders (2004), its deterministic approach produces a strikingly sharp result: every subgame perfect equilibrium delivers a matching that is the optimal stable matching for workers.

The decisions of a firm do not usually reflect the opinion of a single individual; instead, such actions embody a complex process of assembling the opinions of several individuals. We may conjecture that establishing a list of candidates and using it as guidance is—despite the apparent myopia—a more plausible form of behavior than deciding, at each moment in time, whom to propose to. In addition, in some settings firms obey objective criteria to admit workers, so that strategic behavior on the firms' side looses its meaning.

The paper is organized as follows. In Section 2 we introduce the matching model and review some results on matching markets. We formally present the model in Section 3. In Section 4 we turn our attention to questions related to individual decision making and characterize equilibria. The robustness of the results is discussed in Section 5. We conclude in Section 6. Appendix A collects the formal statement of some results.

2. The marriage model

Consider two finite and disjoint sets $F = \{f_1, \ldots, f_n\}$ and $W = \{w_1, \ldots, w_p\}$, where F is the set of firms and W is the set of workers. We let $V = W \cup F$ and sometimes refer to a generic agent by v, while w and f represent a generic worker and firm, respectively. Each agent has a strict, complete, and transitive preference relation over the agents on the other side of the market and the perspective of being unmatched. The preferences of a firm f, for example, can be represented by $P_f = w_3, w_1, f, w_2, \ldots, w_4$, indicating that f's first choice is to be matched to w_3, its second choice is w_1 and it prefers remaining unmatched to being assigned to any other worker. Sometimes it is sufficient to describe only f's ranking of workers it prefers to remaining unmatched, so that the above preferences can be abbreviated as $P_f = w_3, w_1$. Let $P = (P_{f_1}, \ldots, P_{f_n}, P_{w_1}, \ldots, P_{w_p})$ denote the profile of all agents' preferences; we sometimes write it as $P = (P_v, P_{-v})$ where P_{-v} is the set of preferences of all agents other than v. Further, we may use P_S, where $S \subseteq V$, to denote the profile of preferences $(P_v)_{v \in S}$. We write $v' P_v v''$ when v' is preferred to v'' under preferences P_v and we say that v *prefers* v' to v''. We write $v' R_v v''$, when v likes v' *at least as well as* v'' (it may be the case that v' and v'' are the same agent). A worker is *acceptable* if the firm prefers to employ him rather than having its position unfilled; similarly, a firm is acceptable to a worker if he prefers occupying its position, rather than being unemployed.

Formally, a *marriage market* is a triple (F, W, P). An outcome for a marriage market, a *matching*, is a function $\mu : V \longrightarrow V$ satisfying the following: (i) for each f in F and for each w in W, $\mu(f) = w$ if and only if $\mu(w) = f$; (ii) if $\mu(f) \neq f$ then $\mu(f) \in W$; (iii) if $\mu(w) \neq w$ then $\mu(w) \in F$. If $\mu(v) = v$, then v is *unmatched* under μ, while if $\mu(w) = f$, we say that f and w are *matched* to one another. A description of a matching is given by $\mu = \{(f_1, w_2), (f_2, w_3)\}$, indicating that f_1 is matched to w_2, f_2 is matched to w_3 and the remaining agents in the market are unmatched. A matching μ is *individually rational* if each agent is acceptable to its partner, i.e., $\mu(v) R_v v$, for all $v \in V$. We denote the set of all individually rational matchings by $IR(P)$. Two agents f and w form a *blocking pair* for μ if they prefer each other to the agents they are actually assigned to under μ, i.e., $f P_w \mu(w)$ and $w P_f \mu(f)$. A matching μ is *stable* if it is individually rational and it is not blocked by any pair of agents. We denote the set of all stable matchings by $S(P)$.

3. The decentralized job matching game

3.1. Description of the game

In this section, we define the Decentralized Game. The game is given by a market (F, W, P) and an initial matching μ^I. In general, we consider μ^I to be individually rational under the true preferences. The rules of the game are as follows.

The game begins with a node at which nature chooses a sequence of firms at random. Each sequence corresponds to an order at which firms are given the opportunity to make proposals. Following nature's move, the first firm in the selected sequence has the chance to make a proposal. If unmatched under μ^I, the firm may propose to any worker or pass its turn. If matched under μ^I, it may simply fire its initial partner, propose to a different worker, or pass its turn and keep the initial partner.

In the case that a proposal is actually made, the game continues by having the proposed worker deciding whether to accept or to reject the offer. If he accepts, a new matching is formed where this worker and the proposing firm are together and their previous partners, if any, are unmatched. If he rejects, μ^I goes on unchanged. In the case that the firm simply chose to fire its initial worker, a new matching is formed where the firm and its former partner are unmatched, whereas if the firm chose "pass," the initial matching is preserved.

The second firm then moves and the game continues by giving firms the opportunity to make offers, in accordance with the order of the sequence. Each time a firm is called to play, the available moves depend on whether its position is vacant or not. If vacant, the firm may propose to any worker to whom it has not proposed before or simply pass its turn. Otherwise, it may fire the worker it holds, propose to a worker different from its current match and from any worker it has already proposed to, or pass its turn. When a worker receives a proposal, he may accept the offer or reject it and keep his former partner.

The game continues as long as there is at least one firm wishing to make a new offer or to fire the incumbent worker. As soon as every firm in the market sequentially passes its turn, the game ends.

Two remarks are in order. First, the fact that firms are reluctant to repeating offers guarantees that every play of the game ends in a finite number of steps. Still, relaxing this simplifying assumption and allowing for a finite number of repeated offers would not compromise the results. Second, a matched worker is only allowed to reject his current position if he obtains and accepts an alternative offer. This is a one-side-proposing game: workers do not have the spirit to make offers or even to quit their current job. In Section 5, we discuss the consequences of having both firms and workers playing an active role.

To complete the description of the game, we still have to specify the information that each agent possesses throughout the game. It is sensible to assume that in labor markets where myriads of firms and workers interact, each agent only becomes aware of events as they directly impinge on him. In the particular case of a firm, this means that it learns only if the proposal it made was accepted or rejected, or if its position became vacant. Hence, whenever called upon to move, a firm knows the identity of its initial partner and the ordered list of workers to whom it proposed, along with their reactions. Similarly, a worker is only aware of events that directly affect him. Whenever called upon to make a decision, a worker knows his initial position, as well as the

ordered list of proposals he received, his own responses, and firings. The initial chance move is never observable.[8]

3.2. Chance move

Let us now focus on nature's move. At each moment in time, a randomly selected firm is given the chance to play. This random selection should not be interpreted merely as every firm having equal probability of proposing at each step. It may reflect some institutional—and perhaps inherently uncertain—features of the market which are not modeled. In fact, in decentralized markets matching is performed over the telephone network, using the mail, or through the Internet. In such environments, randomness determines the order in which agents communicate: it may depend on which telephone call goes through, on the speed of the mail, or on how fast firms react to eventual proposals. Or it may even be the case that there exists a natural order in which firms are expected to propose—firms that have potentially more to gain will certainly devote more resources into finding the right worker for their position and are, therefore, more likely to make offers.

To be precise, the game starts with a lottery prescribing a sequence of firms that defines the subsequent moves. A sequence corresponds to one of the innumerable possible orders in which firms are allowed to act. We assume that every sequence is infinite and that, in each sequence, every single firm appears infinitely many times. We also assume that every sequence has positive probability of occurring. The sample space over which this probability distribution is defined is denoted by O and o is an arbitrary sample point, a sequence of firms.

Note that, even though we consider infinite sequences of firms, every play of the game ends in finite time. In fact, as firms are only allowed to propose to each worker once and, obviously, firing is possible only if matched, the moment comes when every firm chooses to pass its turn if called to play, either because keeping the incumbent worker is part of its strategy, or because passing is the only available action. Moreover, the fact that a firm appears infinitely many times in each sequence guarantees that this moment comes and the firm is actually called to play. The end is then reached for every play of the game.

3.3. The strategy space

In this section, we describe agents' strategies and introduce some notation. A player's strategy in the Decentralized Game complies with the usual definition of behavioral strategy in an extensive form game, i.e., a plan of action for each information set where he is called to act. However, in the context of a matching market there is a class of strategies worth emphasizing, strategies that resemble those used in a centralized market. Following Blum et al. (1997) we will call these strategies "preference strategies." Such strategies obey a consistency criterion in which agents decide how to move at any information set basing on a list of preferences, including those information sets that would not be reached had that list actually been used. Hence, deviations are regarded as temporary mistakes and further moves fit in the original list. When using a preference strategy, a firm selects an ordered list of potential matches and, whenever called to propose, makes the offer to the best worker on its list to whom it has not proposed before;

---

[8] Such low information environment may be enriched. It may be the case that agents learn of the actions of the others, even though they are not immediately affected by them. The validity of the results that follow will be discussed for broader information structures.

likewise, a worker decides whether to accept or to reject a new proposal by comparing it with his current position on his list.

In what follows, we provide the formal definition of a preference strategy. We let σ_v denote the strategy of agent v; given one of firm f's decision nodes x, we let W^x denote the set of workers f has not proposed to when node x is reached and we let v^x denote f's partner when node x is reached; similarly, y is a decision node of worker w and we let v^y denote w's matching partner when node y is reached.

Definition 1. Let x be one of firm f's decision nodes and let Q_f be a preference ordering for f. We say that the strategy σ_f is *consistent with Q_f at x* if, when using σ_f, the following conditions hold:

1. f proposes to w at x if and only if $w Q_f f$ and there exists no worker $w' \in W^x \cup \{v^x\}$ such that $w' Q_f w$;
2. f passes its turn at x if and only if there exists no worker $w' \in W^x$ such that $w' Q_f v^x$ and either (i) $v^x Q_f f$ or (ii) f is unmatched when x is reached;
3. f fires v^x at x if and only if $f Q_f v^x$ and there exists no worker $w' \in W^x$ such that $w' Q_f f$.

Definition 2. Let y be one of worker w's decision nodes and let Q_w be a preference ordering for w. We say that the strategy σ_w is *consistent with Q_w at y* when w rejects firm f's proposal if and only if $v^y Q_w f$.

Definition 3. The strategy σ_v of agent v is a *preference strategy for v* if there exists a preference ordering Q_v such that σ_v is consistent with Q_v at every decision node of v.

Even though the lists of preferences that serve as guidance do not have to faithfully reveal agents' true preferences, the set of preference strategies represents merely a small part of the set of feasible strategies.[9] For example, a worker w's strategy of accepting only the first proposal he gets and rejecting all the others is not consistent with any list of preferences. In fact, different plays of the game induce different orders of proposals; thus, depending on the play of the game, w's first proposal may be from, say, f and f'. It follows that f may be revealed preferred to f' by w or vice versa, which clearly cannot be consistent with a preference list.

A final remark on notation. We let σ denote a strategy profile; we sometimes write $\sigma = (\sigma_v, \sigma_{-v})$, where σ_v denotes an arbitrary strategy for agent v and σ_{-v} denotes the strategy profile of the other agents. We abuse language slightly and sometimes refer to a preference strategy for agent v by means of the corresponding preference ordering, say Q_v; the preference ordering P_v always denotes v's true preferences. A sequence of firms o and a strategy profile σ determine a play of the game, denoted by π.

3.4. Random matching and ordinal Nash equilibria

In the Decentralized Game, different plays of the game with the same strategy profile may yield different output matchings, depending on the order of proposals. This applies even in the case that agents use preference strategies, as the following example illustrates.

---

[9] We refer to Pais (2006) for the analysis of a job matching game where the strategy space is confined to the set of preference strategies.

Example 1. The outcome depends on the selection of the order by which firms propose.

Let (F, W, P) be a marriage market with P such that

$$P_{w_1} = f_2, f_1 \qquad P_{f_1} = w_1, w_2$$
$$P_{w_2} = f_1, f_3, f_2 \qquad P_{f_2} = w_2, w_1$$
$$P_{f_3} = w_2.$$

Note that the unique stable matching for this market is $\mu = \{(f_1, w_2), (f_2, w_1)\}$. Now consider the Decentralized Game with $\mu^I = \{(f_1, w_2)\}$ when agents play according to their true preferences P.

Start by considering the case in which f_3 is the first to make an offer. Given that f_3 is using P_{f_3}, it proposes to the only acceptable worker, w_2, and w_2 rejects this proposal, as he is initially matched to f_1, the best firm on his list. Then, it may be the case that either f_1's or f_2's opportunity comes. Let us say f_1 makes an offer; it proposes to w_1, the first worker in P_{f_1}, who is currently unmatched and thus accepts the proposal. Once this proposal is accepted, w_2 is left unmatched. Hence, when f_2 is given the chance to propose, w_2 accepts its offer. In the following moves every firm passes its turn, so that the game ends with the final non-stable matching $\hat{\mu} = \{(f_1, w_1), (f_2, w_2)\}$.

Nevertheless, if the first randomly chosen firm is f_2, its proposal to w_2 is refused, as this worker is still matched to f_1 and f_1 is preferred to f_2 in P_{w_2}. The next firm to propose can either be f_1, f_2, or f_3. Assume f_2 is the first to propose. It proposes to w_1, the second worker on its list, and w_1 accepts. Next, if f_1's turn comes, it proposes to w_1, who rejects this offer, since he is matched to his top choice f_2. So imagine f_1 is called to propose once more, tendering an offer to w_2, who accepts it. When finally f_3 proposes to w_2, he rejects the offer, given that he is already holding the highest ranked firm in his preference list. This play of the game terminates when the three firms are given the chance to pass their turns and the matching $\mu = \{(f_1, w_2), (f_2, w_1)\}$ is reached as the outcome of the game.

Given an initial matching and a strategy profile, all the uncertainty on the order of play as described above is fully translated into a probability distribution over the set of matchings. Hence, fix a probability distribution on O and take an initial matching μ^I, a preference profile P, and an arbitrary worker w (what follows also holds for a representative firm, with obvious modifications). We will let $\widetilde{G}^{\mu^I}[\sigma]$ denote the probability distribution over the set of matchings induced by the Decentralized Game starting from μ^I when the strategy profile σ is used and $\widetilde{G}^{\mu^I}_w[\sigma]$ is the distribution that $\widetilde{G}^{\mu^I}[\sigma]$ induces over $F \cup \{w\}$. The expression $\widetilde{G}^{\mu^I}[\sigma](\{\mu\})$ represents the probability that μ is the final matching of the Decentralized Game with the strategy profile σ. Moreover, defining $U(v, P_w)$ as the set of agents w considers at least as good as v according to his true preferences P_w, $\widetilde{G}^{\mu^I}_w[\sigma](U(v, P_w))$ is the probability that, in the Decentralized Game, w obtains a partner at least as good as v when σ is adopted. Observe that these probabilities rest on the probability distribution on O, but all the results that follow hold regardless of this lottery.

To address strategic questions we need to develop ideas about what constitutes a "best decision" to be taken by an agent. With this purpose in mind, let $\sigma = (\sigma_w, \sigma_{-w})$ be a strategy profile and consider an alternative strategy σ'_w for worker w. If there existed a utility function representing w's true preferences, the strategy σ_w would be a better decision for w than σ'_w against σ_{-w} if, when using σ_w, worker w's expected utility would be higher than when using σ'_w. Since in this setting agents' preferences are ordinal and have no natural utility representation, we consider σ_w to be a better decision than σ'_w against σ_{-w} if it induces a

higher expected utility level for *every* utility representation of w's preferences. This is equivalent to the following requirement in terms of stochastic dominance: for all $v \in F \cup \{w\}$, $\widetilde{G}_w^{\mu^I}[\sigma_w, \sigma_{-w}](U(v, P_w)) \geqslant \widetilde{G}_w^{\mu^I}[\sigma'_w, \sigma_{-w}](U(v, P_w))$.[10] In this case, we say that σ_w *stochastically P_w-dominates* σ'_w given σ_{-w}. Note that this criterion only entails a partial order of strategies and it may be the case that for some σ_{-w} there exists no σ_w that stochastically P_w-dominates every alternative strategy. Still, it provides the basis for the solution concepts we will use throughout the paper, namely the concept of ordinal Nash equilibrium, whose existence is not compromised, as shown in the next section.

Definition 4. Let (F, W, P) be a matching market and let μ^I be the initial matching. The profile of strategies σ is an *ordinal Nash equilibrium* (*ON equilibrium*) in the Decentralized Game if, for each player v in V, σ_v stochastically P_v-dominates every alternative strategy σ'_v given σ_{-v}.

Thus, by using a strategy other than σ_v, each agent v will not be able to strictly increase the probability of obtaining any v' (an agent with whom it may end up matched) and all agents ranked higher than v' in its true preference list P_v. Alternatively, the profile of strategies σ is an ordinal Nash equilibrium if, once adopted by the agents, no one can profit by unilaterally deviating for all possible utility representations of the agents' preferences.[11]

The notion of ordinal Nash equilibrium can be refined to account for the dynamic nature of the Decentralized Game.

Definition 5. Let (F, W, P) be a matching market and let μ^I be the initial matching. The profile of strategies σ is an *ordinal subgame perfect Nash equilibrium* (*OSPN equilibrium*) in the Decentralized Game if it induces an ordinal Nash equilibrium in every subgame of the Decentralized Game.

Finally, in the particular case that every strategy used in an ordinal (subgame perfect) Nash equilibrium is a preference strategy, we can speak of ordinal (subgame perfect) Nash equilibria in preference strategies.

Definition 6. Let (F, W, P) be a matching market and let μ^I be the initial matching. The profile of strategies σ is an *ordinal* (*subgame perfect*) *Nash equilibrium in preference strategies* in the Decentralized Game if σ is an ordinal (subgame perfect) Nash equilibrium and σ_v is a preference strategy for every agent v in V.

4. Equilibrium analysis

We begin this section by exploring the relationship between ordinal Nash and subgame perfect ordinal Nash equilibria.

Proposition 1. *Let $|F| \geqslant 2$. Then, no information set is a singleton.*

---

[10] We refer to d'Aspremont and Peleg (1988) for a proof of this equivalence.

[11] For illustration, in Example 1, truthtelling is not an ordinal Nash equilibrium: the preference strategy $Q_{w_2} = f_1$ is a profitable deviation for w_2, ensuring that he gets his preferred match, firm f_1, with probability one.

Lemma 1. *Let* $|F| \geqslant 2$. *Let* x *and* x' *be the two last decision nodes of the play of the game* π, *such that* x' *precedes* x. *Then,* x *and* x' *belong to two different firms and both firms choose the action "pass" at these nodes.*

Proof. First, notice that the game ends when every firm has sequentially chosen "pass." Given that x and x' precede the terminal node reached with π and that $|F| \geqslant 2$, it follows that the action taken at these nodes must be "pass." Now suppose, by contradiction, that both x and x' are firm f's decision nodes. Since, when π is considered, the game ends after f chooses "pass" at x, every firm other than f must have chosen "pass" in the nodes that precede x. Hence, every firm other than f has passed its turn in the nodes that precede x'. The rules of the game thus imply that the game ends immediately after f chooses "pass" at x' and we reach a contradiction: x is not a decision node. $\square$

Lemma 2. *Let* $|F| \geqslant 2$. *Let* π *be a play of the game and let* x *be a node of* f *reached along* π, *such that the game does not end after* f's *choice at* x *along* π. *Then, there exists a firm* f' *that still has a chance to act in* π.

Proof. Immediate from Lemma 1. $\square$

Proof of Proposition 1. Let x be a node that belongs to f in π when nature draws the sequence o. Let f's move at x correspond to the kth element of o. We will prove that there exists a sequence o' and a node x' reached when nature draws o', such that x and x' belong to the same information set.

First, assume that the game does not end after f's choice at x along π. By Lemma 2, there exists a firm f' that still has the chance to act along π. Now let o' be a sequence whose k first elements are the same as those in o, but that differs from o in that f is inserted in position $k+1$ and all the remaining elements are identical. Consider any play of the game where nature draws o' and every agent chooses exactly the same actions as along π up to the point where o'_{k+1} is called to play. Let x' be the node corresponding to f's move in position k of the sequence o'. It is clear that x' belongs to the same information set as x, since every action, except for the unobservable nature's move, is the same along π and π'.

Now let x be a node of f, reached along π, such that f's action at x is the last action in π. By Lemma 1, there exists a firm $f' \neq f$ that has had the chance to move immediately before f moves at x, i.e., in position o_{k-1} of the sequence, and both have chosen "pass." Now let o' be a sequence whose first $k-1$ elements coincide with those of o, but where f' occupies the position o_k and f occupies the position o_{k+1}. Consider the play of the game π' where nature draws o', every agent up to the element $k-1$ in the sequence chooses exactly the same action as in π, and f' chooses "pass" when called to play at the kth position of the sequence. Let x' be the node reached in π' where f acts in position $k+1$. Since f cannot observe nature's moves nor f''s action, it holds exactly the same information in both x and x'. Hence, x' belongs to the same information set as x.

Now consider π where nature draws o and along which some worker w may accept or reject a proposal made by firm f. Let x be the node where w acts and let f's proposal correspond to the kth element of o. Lemma 1 ensures that the game does not end after w's move at x. Hence, let o' be any sequence whose k first elements are the same as those in o, but such that the elements in position $k+1$ are different. Define π' as a play of the game in which nature draws o' and every other player chooses the same actions as along π up to the point where w reacts to f's proposal.

Let x' be the node where w takes such decision. Since nature's draws are not observable, w's information is exactly the same in x and in x'. It follows that the information set containing x is not a singleton. $\quad\square$

An immediate implication of this result is that the set of ordinal Nash and subgame perfect ordinal Nash equilibria coincide. In fact, given that all information sets are non-singletons, the Decentralized Game has no proper subgames. It may be conjectured that this is due to the low information environment we have assumed. And there are labor markets in which agents may become aware of events that do not affect them directly—acquaintances and social networks in general may play an important role. However, considering an enriched information environment where agents perceive all the offers that are made, as well as the proposed workers' reactions, the arguments in the above proof remain valid, as long as nature's move remains unobservable. Roughly speaking, for every decision node x along some play of the game that includes a draw of nature o, it is always possible to find a decision node x' belonging to the same information set of x by building a different play of the game in the following way: add a single firm to o, let it choose "pass" in its new decision node, and let agents choose exactly the same actions as in the original play in every other node. The conclusion follows since every proposal, acceptance, and rejection is made respecting the original order. Hence, even in this extreme case, ordinal subgame perfect Nash coincide with ordinal Nash equilibria. In what follows, we will refer to these concepts indistinctly as ordinal equilibria.

The following results show the existence of ordinal equilibria when the initial matching is individually rational.

Definition 7. Let μ^I be an arbitrary matching. We say that μ is *individually rational with respect to μ^I* if $\mu \in IR(P)$ and if, for every $f \in F$ such that $w' = \mu^I(f) P_f \mu(f)$, we have $\mu(w') \neq w'$.

We will denote by $IR^{\mu^I}(P)$ the set of all individually rational matchings with respect to μ^I. For illustration, in the particular case that μ^I is the empty matching, the set of all individually rational matchings with respect to this initial matching coincides with the set of individually rational matchings (i.e., $IR^{\emptyset}(P) = IR(P)$).[12]

Proposition 2. *Let μ^I be an individually rational matching for (F, W, P). Then, $S(P)$ is a subset of $IR^{\mu^I}(P)$.*

Proof. Consider $\mu \in S(P)$. We will prove that $\mu \in IR^{\mu^I}(P)$ using a contradiction argument. Assume that $\mu \in S(P)$. By definition of stability, this implies $\mu \in IR(P)$, but assume that there exists a firm f such that $w' = \mu^I(f) P_f \mu(f)$ and $\mu(w') = w'$. Stability of μ implies that $w' P_{w'} f$ and we get a contradiction: μ^I is not individually rational. Therefore, every stable matching is an element of $IR^{\mu^I}(P)$. $\quad\square$

Since a stable matching exists for every marriage market (Gale and Shapley, 1962), $IR^{\mu^I}(P)$ is not empty for every individually rational matching μ^I.

---

[12] This holds since, when $\mu^I = \emptyset$, the second condition in Definition 7 is trivially satisfied: $f P_f \mu(f)$ implies $\mu(f) \neq f$ by definition of P_f.

Proposition 3. *Let μ^I be an individually rational matching for (F, W, P) and let $\mu \in IR^{\mu^I}(P)$. Consider the profile of preference strategies Q such that $Q_v = \mu(v)$, for every agent v. Then, Q is an ordinal equilibrium in preference strategies in the Decentralized Game and leads to μ with probability one.*

Proof. Given $Q_v = \mu(v)$, it is clear that every play of the game with the profile σ will lead to the output matching μ.

Let us show that for every firm f, Q_f is a best reply to Q_{-f}. First, as long as $\mu(f) \neq \mu^I(f)$, f can never hold its initial match under μ. Indeed, it is clear that if $\mu^I(f) P_f \mu(f)$, then $\mu^I(f)$ receives and accepts another firm's proposal (and in the case that $\mu(f) P_f \mu^I(f)$, $\mu^I(f)$ is not a temptation). Hence, when $\mu(f) \in W$, given that the only worker willing to accept f's proposal is $\mu(f)$, the only choice f can actually make is between being assigned to this worker or staying alone. From individual rationality we have $\mu(f) P_f f$ which implies that f will not be able to profit from deviating from Q_f. Obviously, for f such that $\mu(f) = f$, no worker accepts f's proposal and it can do no better than staying alone.

Finally, for any w, Q_w is a best reply to Q_{-w}. In fact, given firms' strategies, w gets at most one proposal and, considering μ is individually rational, the best he can do is to accept it. This completes the proof. □

Note that, in Proposition 3, we use strategies that are somewhat degenerate, by considering that a unique matching partner is considered acceptable. It serves the purpose of exposing the multiplicity of ordinal equilibrium outcomes: virtually every individually rational matching can be reached in such an ordinal equilibrium. Still, the set of equilibrium strategies is wider, as the following example shows.

Example 2. Ordinal equilibria with non-degenerate strategies.

Consider the matching market in Example 1 and again consider the Decentralized Game with $\mu^I = \{(f_1, w_2)\}$ when every agent plays according to its true preferences P, except for w_2 who follows the preference strategy $Q_{w_2} = f_1$. This profile of strategies is an ordinal equilibrium of the Decentralized Game leading to $\mu = \{(f_1, w_2), (f_2, w_1)\}$ with probability one. In fact, no agent has profitable deviations: workers obtain their preferred matches for every order of proposals and cannot improve their positions; given Q_{w_2}, w_2 never accepts f_3 nor f_2, so that neither firm can end up matched to w_2; and w_1 ends up receiving an offer from f_2 for every order of proposals, undermining every chance of f_1 hiring w_1.

Even though the set of possible equilibrium outcomes is vast and agents may misrepresent their preferences in an ordinal equilibrium, there are some necessary conditions that every ordinal equilibrium outcome has to fulfill. First, it is obvious that every outcome obtained in an equilibrium play of the game must be individually rational. Second, in Theorem 1, we state that under every ordinal equilibrium play of the Decentralized Game where firms use preference strategies, *some* form of stability is preserved. To be more precise, every matching that can be obtained under such a play is stable for the same profile of preferences. The following remark is used in the proof of the theorem.

Remark 1. When using a preference strategy, a firm will not fire a worker it proposed to nor exchange him for another worker along any play of the Decentralized Game. In fact, when a proposal is made, the firm reveals that this particular worker is the best among all who have not

rejected it. If the worker accepts, the only occasion under which the firm makes a proposal again is when the worker it holds resigns from his position.

Theorem 1. *Let μ^I be an individually rational input matching for $(F, W, (Q_F, P_W))$. Assume that the strategy profile $\sigma = (Q_F, \sigma_W)$ is an ordinal equilibrium in the Decentralized Game. Then, the probability distribution obtained over the set of matchings is such that every element in its support is a member of $S(Q_F, P_W)$.*

Proof. Suppose that $\{\mu_1, \ldots, \mu_k\}$ is the support of the distribution induced over the set of matchings when agents use σ. Assume that for some $i \in \{1, \ldots, k\}$, $\mu_i \notin S(Q_F, \sigma_W)$. We will prove that σ is not an ordinal equilibrium.

Let π be a play of the game that results in μ_i. To start, notice that for every firm f it must be the case that its assignment, $\mu_i(f)$, is acceptable with respect to Q_f. In fact, once using Q_f, f never proposes, under any play of the game, to a worker that, according to Q_f, is considered worse than being unmatched. On the other hand, every worker must consider his partner acceptable with respect to P. Assume that this is not the case and that there exists a worker, say w, such that $w P_w \mu_i(w)$. Individual rationality of the matching μ^I implies $\mu_i(w) \neq \mu^I(w)$. Hence, σ_w must include, at some point along π, accepting $\mu_i(w)$'s proposal. Now take an alternative strategy $\hat{\sigma}_w$ according to which no offer is accepted by w. By using $\hat{\sigma}_w$, w may end up unmatched or, if initially matched, keep his original partner $\mu^I(w)$, but he is never assigned to a firm considered unacceptable under P_w. Thus, the following holds:

$$1 = \widetilde{G}_w^{\mu^I}[\hat{\sigma}_w, \sigma_{-w}]\big(U(w, P_w)\big) > \widetilde{G}_w^{\mu^I}[\sigma]\big(U(w, P_w)\big)$$

and σ_w is not a best reply to σ_{-w}.

We have proved that μ_i is individually rational. Thus, there must exist a blocking pair for μ_i when the preference profile (Q_F, P_W) is considered. Let us say (f, w) blocks μ_i, i.e., $f P_w \mu_i(w)$ and $w Q_f \mu_i(f)$. This implies that f proposed to w and, by Remark 1, was rejected by w in the course of π. Hence, σ_w includes rejecting f in at least one of w's information sets. Now, define $\hat{\sigma}_w$ as the strategy according to which w chooses the same actions as under σ_w at every information set, except for those that lead to rejecting f. When using $\hat{\sigma}_w$, if f proposes to w, w accepts this proposal and rejects every subsequent offer.

First, let us prove that the probability of being assigned to f is strictly higher under $\hat{\sigma}_w$ than under σ_w. Recall that π is a play of the game leading to μ_i and let o be nature's move in π. We know that firm f must have proposed to w along π. If, instead of using σ_w, w deviates and acts according to $\hat{\sigma}_w$, by Remark 1, w will end up matched to f when nature draws o and σ_{-w} is used. Now let π' be the play of the game in which nature draws o', with $o' \neq o$, and players use $(\hat{\sigma}_w, \sigma_{-w})$. If f does not propose to w along π', w acts exactly as if using σ_w and ends up matched to the same partner as when nature draws o and players use σ. Otherwise, f and w are matched in the final matching. It follows that the probability of having f and w matched is strictly increased when w uses $\hat{\sigma}_w$.

In order to prove σ_w is not a best reply to σ_{-w}, assume, without loss of generality, that $P_w = f_1, f_2, \ldots, f_{m-1}, f, f_{m+1}, \ldots, w, \ldots, f_n$. Consider a firm f_j, with $j = 1, \ldots, m - 1$, and consider all the plays of the game where σ is used and where w and f_j end up together in the final matching. Some of these plays may not give f_j assigned to w when he deviates and acts according to $\hat{\sigma}_w$. However, the only occasion under which this happens is when w obtains a pro-

posal from f and ends up matched to f. Hence, the probability of having w matched to f or to a firm he considers better than f is strictly increased when w uses $\hat{\sigma}_w$. We have

$$\widetilde{G}_w^{\mu^I}[\hat{\sigma}_w, \sigma_{-w}]\big(U(f, P_w)\big) > \widetilde{G}_w^{\mu^I}[\sigma]\big(U(f, P_w)\big),$$

contradicting that σ is an ordinal equilibrium. $\quad\square$

The importance of this result lies in two of its implications. Since the set of unmatched agents is the same for every matching that is stable in a matching market (McVitie and Wilson, 1970, and Roth, 1982), the same agents remain unmatched in every possible outcome of an ordinal equilibrium where firms use lists of workers to guide their decisions. To fix a list of workers and myopically take all decisions basing on that list may be reasonable in firms where hiring a worker embodies a process of aggregating the opinions of different individuals that compose the recruiting committee, whereas having to decide what to do next at each moment in time may be a less plausible form of behavior. The second implication of Theorem 1 is that, when we focus on equilibria where firms act according to their true preferences, stability with respect to the true preferences is guaranteed.[13] Furthermore, such straightforward form of behavior can be justified. In some settings, firms obey objective criteria when selecting whom to hire (e.g., universities select students according to their grades, some firms choose their workers basing on scores given by a recruiting agency, student placement mechanisms assign students to public schools according to the area of residence, etc.). Even when firms are not constrained to follow such rules, reverting to the true preferences is always an easy resort, given the multiplicity of available strategies and the complexity of the environment.

In Proposition 4, ordinal equilibria where firms act straightforwardly according to their true preferences are shown to exist in the particular case that the starting point of the Decentralized Game is the empty matching. In this case, the game is a stylized model of an entry-level labor market without commitment, where cohorts of vacant positions and cohorts of candidates become simultaneously available, and decisions are taken in a decentralized way.

Proposition 4. *Let μ^I be the empty matching in (F, W, P) and let $\mu \in S(P)$. Let $Q = (P_F, Q_W)$ be such that $Q_w = \mu(w)$ for every worker w. Then, Q is an ordinal equilibrium in preference strategies in the Decentralized Game and leads to μ with probability one.*

Proof. We start by showing that the profile of strategies Q always leads to the matching μ, i.e., $\widetilde{G}^{\mu^I}[Q](\{\mu\}) = 1$. If this is not the case, then there exists a play of the game leading to $\hat{\mu} \neq \mu$. But this is equivalent to having a firm, say f, that ends up matched to a partner, $\hat{\mu}(f)$, different from $\mu(f)$ for some instance of the game. Given that f is unmatched in the initial matching and that the only worker willing to accept f is $\mu(f)$, we must have $\hat{\mu}(f) = f$ (as long as $f \neq \mu(f)$; otherwise it must be the case that $\hat{\mu}(f) = \mu(f)$ and we have a contradiction). So assume that $\hat{\mu}(f) = f$. Since $\mu(f)$ would accept f's proposal and f is acting according to its true preferences, it must be the case that $f P_f \mu(f)$. Hence, $\mu(f)$ is not acceptable and the stability of μ is contradicted.

Let us now prove that, for every firm f, P_f stochastically P_f-dominates every other strategy σ_f. We will start by considering the case in which $\mu(f) \neq f$. Given that the only worker who is willing to accept f is $\mu(f)$, by choosing its strategy appropriately, f can either be alone

---

[13] Going back to the setup of Example 1, this result rules out truth telling as a candidate to an ordinal equilibrium.

or hold $\mu(f)$ under the output matching. By stability of μ, $\mu(f) P_f f$; since truth telling guarantees that $\mu(f)$ is assigned to f with probability one, f cannot improve by switching its strategy. In the case that $\mu(f) = f$, no worker accepts its proposal, and the best it can achieve is staying unmatched. It follows that f cannot do better than being assigned to $\mu(f)$ and P_f stochastically P_f-dominates every other strategy σ_f.

Now take the case of an arbitrary worker, w. Suppose, by contradiction, that Q_w does not stochastically P_w-dominate strategy σ_w. This implies that $\widetilde{G}_w^{\mu^I}[P_F, \sigma_w, \sigma_{-w}](U(\mu(w), P_w)) = 1$

and that there exists a firm, say f, such that the following holds: $\widetilde{G}_w^{\mu^I}[P_F, \sigma_w, \sigma_{-w}](\{f\}) > 0$ with $f P_w \mu(w)$. But this means that, for some draw of nature, f approaches w before making an offer to $\mu(f)$. In fact, it cannot be the case that f proposes to $\mu(f)$ first and he does not accept it, as $\mu(f)$ is acting according to his original strategy, $Q_{\mu(f)}$, defined above. Thus, f must prefer w to $\mu(f)$. However, in this case (f, w) forms a blocking pair for μ, contradicting the fact that μ is stable. $\square$

Hence, in a decentralized entry-level labor market every stable matching can be reached as the outcome of an ordinal equilibrium play of the game where firms stick to their true rankings. Still, a remark is due: the above result requires workers to use strategies where they only report as acceptable the firm to which they are matched under a particular stable matching. This obviously requires a lot of coordination, which may seem unnatural in a decentralized market. The (partially) converse statement is given by Theorem 1, ensuring that every such ordinal equilibrium guarantees stability.

5. Discussion

In this section we discuss the robustness of the results to some changes in the rules of the game. First, as already mentioned, the simplifying assumption that firms do not repeat offers ensures that the game is finite. Still, admitting that firms are persistent does not compromise the above results, as long as we allow for a finite number of repetitions.

Second, throughout the game matchings are formed and dissolved as agents act in what they perceive to be their own best interest. We may think of this as a mere negotiation process, where no contracts are signed and where these temporary matchings would be the ones prevailing should negotiations suddenly stop. Alternatively, considering that provisional matchings are indeed consummated amounts to assuming that agents are free to recontract without any restrictions whatsoever. In the other extreme, we can consider that it is too costly to fire a worker: only firms with vacancies will actually make proposals and the Decentralized Game falls in the realms of Blum et al.'s analysis. Hence, Theorem 1, which allows for having any individually rational as an initial matching, no longer holds. The validity of Proposition 3 is also compromised: if we start from an initial matching where every firm is matched, no firm will be allowed to hire a new worker and the initial matching situation will be preserved, independently of the strategies used. In this setting, the majority of the results on equilibria depend on having a firm-quasi-stable matching as a starting point.

A different issue concerns providing workers with the initiative to quit and to propose. In some real labor markets, not only firms, but also workers may defy their preferred firms and we may account for this in the Decentralized Game. Hence, suppose that at each moment in time, an agent, either a firm or a worker, is randomly selected and makes an offer to someone in the other side of the market to whom it has never proposed to nor received a proposal from. The

agent that receives the offer can only accept, or reject and keep his former partner, if a former partner existed. The game ends when every agent in the market passes its turn. It turns out that, starting from an arbitrary matching, every individually rational matching can be obtained in an ordinal equilibrium play of the game, so that the scope of Proposition 3 is enhanced. In what Theorem 1 and its implications are concerned, stability is robust to sophisticated behavior by one side of the market, provided that the other side acts in accordance with the true preferences. To be precise, every matching sustained at an ordinal equilibrium is stable with respect to the true preferences whenever each firm (respectively worker) faithfully transmits its preferences and each worker (respectively firm) behaves strategically by using a strategy that is consistent with a list of preferences along each play of the game, but that may reveal different orderings of the other side of the market in different executions of the algorithm. Finally, in the particular case that the initial matching is the empty matching, every stable matching can be reached with probability one in an equilibrium where one side of the market truthfully reveals its preferences.[14]

6. Concluding remarks

The present paper attempts to extend the two-sided matching theory by constructing a game that mimics the behavior of some decentralized labor markets. Equilibrium analysis in a random context is performed at the expense of using an ordinal equilibrium concept that allows for obtaining some interesting results. Namely, equilibria where firms use preference strategies always exist and lead to matchings that preserve stability for a particular profile of preferences. Furthermore, when we consider an ordinal equilibrium where firms act truthfully, stability for the true preferences is achieved in every outcome matching. A case of particular interest has the empty matching as the starting point of the game. Here we give a fairly complete characterization of ordinal equilibria.

It is natural to ask to what extent the stylized model constructed here can serve as a description of a real decentralized labor market. The marriage model is perhaps too simple. Aside from the assumption that each firm has a unique position to fill, the important unrealistic feature lies in considering that the salary associated with each position is a fixed part of the job description, rather than something to be negotiated between each firm and prospective worker. Furthermore, the concept of ordinal equilibrium—justified by the ordinal nature of agents' preferences—is quite demanding. It thus remains important to explore models where these assumptions are relaxed and a milder solution concept is used, even though we believe that the present analysis provides a good starting point to understand the functioning of some decentralized labor markets.

---

[14] In Appendix A we define "a strategy consistent with a list of preferences along a play of the game" and provide the formal statement of these results.

J. Pais / Games and Economic Behavior 64 (2008) 632–649

Appendix A

In this section we provide the definition of a strategy consistent with a list of preferences along a play of the game and the formal statement of the results obtained when both sides of the market are able to tender offers.

Definition 8. Let π be a play of the game and let Q_v be a preference ordering for agent v. The strategy σ_v *is consistent with* Q_v *along the play of the game* π if and only if σ_v is consistent with Q_v at every decision node of v reached along π.

Proposition 5. *Let* $|V| \geqslant 2$. *Then, no information set is a singleton.*

It follows that subgame perfect ordinal Nash equilibria coincide with ordinal Nash equilibria. The remaining results are as follows.

Proposition 6. *Let* μ^I *be an arbitrary matching in* (F, W, P) *and let* $\mu \in IR(P)$. *Consider the profile of preference strategies* Q *such that* $Q_v = \mu(v)$, *for every agent* v. *Then,* Q *is an ordinal equilibrium in preference strategies in the Decentralized Game and leads to* μ *with probability one.*

Proposition 7. *Let* μ^I *be an arbitrary matching in* (F, W, P). *Assume that the strategy profile* $\sigma = (Q_F, \sigma_W)$, *where* σ_w *is consistent with a list of preferences along each play of the game for every worker* w, *is an ordinal equilibrium in the Decentralized Game. Then, the probability distribution obtained over the set of matchings is such that every element in its support is a member of* $S(Q_F, P_W)$.

Proposition 8. *Let* μ^I *be the empty matching in* (F, W, P) *and let* $\mu \in S(P)$. *Then, there exists an ordinal equilibrium in the Decentralized Game where one side of the market reveals its true preferences that yields* μ *with probability one.*

References

Blum, Y., Roth, A., Rothblum, U., 1997. Vacancy chains and equilibration in senior-level labor markets. J. Econ. Theory 76, 362–411.

d'Aspremont, C., Peleg, B., 1988. Ordinal Bayesian incentive compatible representations of committees. Soc. Choice Welfare 5, 261–279.

Ehlers, L., Massó, J., 2007. Incomplete information and singleton cores in matching markets. J. Econ. Theory 136, 587–600.

Gale, D., Shapley, L., 1962. College admissions and the stability of marriage. Amer. Math. Monthly 69, 9–15.

Haeringer, G., Wooders, M., 2004. Decentralized job matching. Mimeo, Universitat Autònoma de Barcelona.

Majumdar, D., 2003. Ordinally Bayesian incentive compatible stable matchings. Mimeo, Concordia University.

Majumdar, D., Sen, A., 2004. Ordinally Bayesian incentive compatible voting rules. Econometrica 72, 523–540.

McVitie, D., Wilson, B., 1970. Stable marriage assignment for unequal sets. BIT 10, 295–309.

Niederle, M., Roth, A., 2003. Unraveling reduces mobility in a labor market: Gastroentrology with and without a centralized match. J. Polit. Econ. 111, 1342–1352.

Pais, J., 2006. On random matching markets: Properties and equilibria. UECE WP 11/2006/DE/UECE.

Pais, J., 2008. Random matching in the college admissions problem. Econ. Theory 35, 99–116.

Roth, A., 1982. The economics of matching: Stability and incentives. Math. Operations Res. 7, 617–628.

Roth, A., 1984a. The evolution of the labor market for medical interns and residents: A case study in game theory. J. Polit. Econ. 92, 991–1016.

Roth, A., 1984b. Misrepresentation and stability in the marriage problem. J. Econ. Theory 34, 383–387.

Roth, A., 1991. A natural experiment in the organization of entry level labor markets: Regional markets for new physicians and surgeons in the UK. Amer. Econ. Rev. 81, 415–440.

Roth, A., Sotomayor, M., 1990. Two-Sided Matching: A Study in Game-Theoretical Modeling and Analysis. Econ. Soc. Monogr. Ser., vol. 18. Cambridge Univ. Press, New York.

Roth, A., Vate, J., 1991. Incentives in two-sided matching with random stable mechanisms. Econ. Theory 1, 31–44.

Roth, A., Xing, X., 1994. Jumping the gun: Imperfections and institutions related to the timing of market transactions. Amer. Econ. Rev. 84, 992–1044.

Roth, A., Xing, X., 1997. Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists. J. Polit. Econ. 105, 284–329.