ISSN 1471-0498

DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

THE EVOLUTION OF CORE STABILITY IN DECENTRALIZED MATCHING MARKETS

Heinrich H. Nax, Bary S.R. Pradelski and H. Peyton Young

Number 607
June 2012

Manor Road Building, Oxford OX1 3UQ

The Evolution of Core Stability in Decentralized Matching Markets

Heinrich H. Nax, Bary S. R. Pradelski & H. Peyton Young[1]

Abstract

Decentralized matching platforms on the internet allow large numbers of agents to interact anonymously at virtually no cost. Very little information is available to market participants and trade takes place at many different prices simultaneously. We propose a decentralized learning process in such environments that leads to stable and efficient outcomes. Agents on each side of the market make demands of potential partners and are matched if their demands are mutually compatible. Matched agents occasionally experiment with higher demands, while unmatched agents lower their demands in the hope of attracting partners. This learning process implements core allocations even though agents have no knowledge of other agents' strategies, payoffs, or the structure of the game, and there is no central authority with such knowledge either.

JEL classifications: C71, C73, C78, D83

Keywords: assignment games, cooperative games, core, evolutionary game theory, learning, matching markets, stochastic stability

[1]Acknowledgements. This research was supported by the United States Air Force Office of Scientific Research Grant FA9550-09-1-0538.

1. Introduction

Electronic technology has created new forms of markets that involve large numbers of agents who interact in real time at virtually no cost. Interactions are driven by repeated online participation over extended periods of time without public announcements of bids, offers, or realized prices. Even after many encounters, agents may learn little or nothing about the preferences and past actions of other market participants. Our goal is to construct a dynamic model that incorporates these features, and to explore its convergence and welfare properties. We see this as a first step toward developing a better understanding of how such markets operate, and how they might be more effectively designed.

We shall be particularly interested in bilateral markets where agents on each side of the market submit "demands" and are matched provided that their demands are mutually compatible. Examples include online platforms for matching buyers and sellers of goods, for matching workers and firms, and for matching hotels with hotel clients. These matching markets have traditionally been analyzed using game-theoretic methods (Gale & Shapley [1962], Shapley & Shubik [1972], Roth & Sotomayor [1990]). In much of this literature, however, it is assumed that agents submit preference menus to a central authority, which then employs a suitably designed algorithm to match them. The model we propose here is different in character: agents submit bids that are conditional on the characteristics of those with whom they are matched, and the only role of the central authority is to create a compatible (not necessarily optimal) set of matches at each point in time. There is no presumption that agents or the central authority know anything about the preferences of the others, or that they can deduce such information from prior rounds. Instead agents employ a trial-and-error learning model: matched agents occasionally ratchet up their demands to see if they can "get away with it", while unmatched agents ratchet down their demands in the hope of attracting partners.

We show that a class of learning rules with simple adjustment dynamics of this type implements the core. The main contribution of the paper is to show how little information may be necessary to implement efficient outcomes, in particular, that such outcomes are achieved even though agents have no knowledge of other agents' strategies or the structure of the game, and there is no central authority with such knowledge either.

The paper is structured as follows. The next section discusses the related literature on matching and core implementation. Section 3 formally introduces assignment games and the concepts of bilateral stability and the core. Section 4 describes the process of adjustment and search by individual agents. In section 5 we prove that this process converges to the core. Section 6 concludes.

2. Related literature

To the best of our knowledge there is no previous work on dynamic learning models applied to decentralized matching markets. However, there is a sizeable literature on matching algorithms that grows out of the seminal paper by Gale & Shapley [1962]. In this approach agents submit preferences for being matched with agents on the other side of the market, and a central clearing algorithm matches them in a way that yields a core outcome (provided that the reports are truthful).[2] These algorithms have been successfully applied in situations where agents engage in a formal application process, such as students applying for admission to universities, or doctors applying for hospital residencies.[3]

In the present paper, by contrast, we consider situations where the market is fluid and decentralized. Agents are matched and rematched over time, and the information they submit takes the form of prices rather than preferences. We shall show that even when agents have minimal amounts of information and use very simple price adjustment rules, the market evolves towards core outcomes. During the necessary adjustment process, there is a positive probability that the relevant resulting adjustment chains mirror key features of the Hungarian (primal-dual) algorithm (Kuhn [1955]). In particular, sequences occur with positive probability where matches, breakups, and rematches occur in such a way that shares of the surplus move from one side of the market to the other in a random fashion. We prove that this leads in expectation to higher degrees of stability and fewer single agents in the market. Crucially, however, unlike the Hungarian method or related algorithms, neither a central authority nor the agents know that they are implementing any paths with global significance.

This result fits into a growing literature showing how cooperative game solutions can be understood as outcomes of a dynamic learning process (Agastya [1997], [1999]; Arnold & Schwalbe [2002], Rozen [2010a], [2010b]; Newton [2010], [2011]; Sawa [2011]). To illustrate the differences between these approaches and ours, we shall briefly outline Newton's model here; the others are similar in spirit.[4] In each period a player is selected at random and demands a share of the surplus from some targeted coalition of players. He chooses a demand that amounts to a best reply to the expected demands of the others in the coalition, where his expectations are based on a random sample of the other players' past demands. In fact he chooses a best reply with probability close to one, but with small probability he may make some other demand. This noisy best response process leads to a Markov chain whose ergodic distribution can be characterized using the theory of large deviations. Newton shows that, subject to various regularity conditions, this process converges to a core allocation in games that possess a nonempty interior core.[5]

The approach we take here requires considerably less information on the part of the agents. Unlike in Newton's model, players know nothing about current or past behavior

[2]See Crawford & Knoer [1981], Kelso & Crawford [1982], Demange & Gale [1985], Demange, Gale & Sotomayor [1986] for examples. See Shimer [2005], Elliott [2010], [2011] for models with costly search.

[3]See, for example, Roth [1984] for a discussion of the medical resident market in the US and the National Residency Matching Program.

[4]Newton's [2011] model nests the models of Agastya [1997], [1999] and Rozen [2010a], [2010b] as special cases. Unlike the other papers, Arnold & Schwalbe [2002] assume further random perturbations outside the core.

[5]The core of an assignment game typically has an empty interior, hence the model's results cannot be translated directly in the present structure.

of other market participants. No information is available about other players' choices, intentions, or payoffs. Thus, they have no basis on which to best respond to the other players' strategies; they simply experiment to see whether they might be able to do better. Rules of this type have a long history in the psychology literature (Thorndike [1898], Hoppe [1931], Estes [1950], Bush & Mosteller [1955], Herrnstein [1961]). Furthermore it has recently been shown that there are families of such rules that lead to equilibrium behavior in generic noncooperative games (Karandikar, Mookherjee, Ray & Vega-Redondo [1998], Foster & Young [2006], Germano & Lugosi [2007], Marden, Young, Arslan & Shamma [2009], Young [2009], Pradelski & Young [2010]).

To the best of our knowledge this framework has not previously been used to study learning dynamics in cooperative games.[6] It seems especially well-suited to modeling behavior in large decentralized markets, where agents have little information about the overall game and the identity of the other market participants. Here we shall restrict our attention to the analysis of learning dynamics in matching (assignment) games, which constitutes a particularly important class in practice.

3. Assignment games

In this section we shall introduce the conceptual framework for analyzing matching markets. The next section will introduce the learning process itself.

A population $N = \{F \cup W\}$ consisting of firms $F = \{f_1, ..., f_m\}$ and workers $W = \{w_1, ..., w_n\}$ repeatedly interacts in a two-sided market submitting demands to a central market authority ("the Center"); one-to-one partnerships form between firms $i \in F$ and workers $j \in W$ if they are compatible.[7]

Willingness to pay. Each firm i has a willingness to pay, p^+_{ij}, for being matched to worker j.

Willingness to accept. Each worker j has a willingness to accept, q^-_{ij}, for being matched with firm i.

We assume that these numbers are specific to the agents and are not known to the other market participants or to the Center.

Value of trade. Assume that utility is linear and separable in money. For each partnership $(i, j) \in F \times W$, the value of trade is the potential surplus of a match which we assume to be positive, that is,

$$\alpha_{ij} = p^+_{ij} - q^-_{ij} \geq 0.$$

Bids and payoffs. The way in which the value of a trade is distributed among the partners depends on their bids. Each agent submits conditional bids to the Center. Agent $i \in F$ submits a vector of numbers $b^t_i = (p^t_{i1}, ..., p^t_{in})$ where p^t_{ij} is the amount i

[6] Sandholm [2008] reviews many of the previous applications.

[7] The two sides may also represent buyers and sellers, or men and women.

would pay if matched with $j \in W$. Similarly, agent $j \in W$ submits a vector of bids $b_j^t = (q_{1j}^t, ..., q_{mj}^t)$ where q_{ij}^t is the amount j would accept if matched with $i \in F$.

Only these bids are known to the Center who can deduce which matches are *compatible*, that is, for whom $p_{ij}^t \geq q_{ij}^t$, a subset of which is implemented.

Central information. At any moment in time, the Center observes

- current bids p_{ij}, q_{ij} for all i, j,
- who is matched with whom,
- compatibility of any possible match i, j.

If i is matched with j, given their bids $p_{ij} \geq q_{ij}$, we assume the resulting "price" is the larger of the two bids. Thus, in any matched pair, the entire surplus is allocated by definition. The resulting payoffs to i and j are $\phi_i = p_{ij}^+ - p_{ij}$ and $\phi_j = p_{ij} - q_{ij}^-$. We can assume that $\phi_i, \phi_j \geq 0$, that is, no one bids more than his willingness to pay or less than his willingness to accept: for all (i, j) and for all t, $p_{ij}^t \leq p_{ij}^+$ and $q_{ij}^t \geq q_{ij}^-$. Furthermore, we shall assume that each bid by a given agent would yield the same surplus if realized. In other words, an agent factors the expected benefits from each possible match into his bid. This means that each agent's bid vector has a very simple structure: there exist numbers ϕ_i, ϕ_j such that

$$\text{for every } i, \; p_{ij} = p_{ij}^+ - \phi_i \;\text{ for all } j, \tag{1}$$

$$\text{for every } j, \; q_{ij} = \phi_j - q_{ij}^- \;\text{ for all } i. \tag{2}$$

We can think of each agent's **strategy** as a demand for a certain level of surplus (ϕ_i or ϕ_j). Two bids are compatible if $p_{ij}^t \geq q_{ij}^t$, which the Center can observe. This is equivalent to saying that their demands d_i, d_j satisfy the inequality $d_i + d_j \leq \alpha_{ij}$ where the latter is not known to the Center (or even to the players). Equations 1 and 2 imply that when a pair is matched, we have

$$d_i + d_j = \alpha_{ij}. \tag{3}$$

In what follows, however, it will be convenient to formulate the adjustment algorithm in terms of the surplus demanded rather than in terms of the bids themselves.

Assignment market. The assignment market is now described by $[F, W, \boldsymbol{\alpha}, \mathbf{A}]$ in the standard way:

- $F = \{f_1, ..., f_m\}$ is a set of m firms (or men or sellers).
- $W = \{w_1, ..., w_n\}$ is a set of n workers (or women or buyers).
- $\boldsymbol{\alpha} = \begin{pmatrix} \alpha_{11} & \dots & \alpha_{1n} \\ \vdots & \alpha_{ij} & \vdots \\ \alpha_{m1} & \dots & \alpha_{mn} \end{pmatrix}$ is the matrix of partnership values.

$\boldsymbol{\alpha}$ specifies the total surplus generated by any possible match in a given assignment.

- $\mathbf{A} = \begin{pmatrix} a_{11} & \dots & a_{1n} \\ \vdots & a_{ij} & \vdots \\ a_{m1} & \dots & a_{mn} \end{pmatrix}$ is a particular assignment matrix with

$$a_{ij} = 0 \text{ or } 1 \quad \sum_{j \in W} a_{ij} \leq 1 \text{ and } \sum_{i \in F} a_{ij} \leq 1 \text{ for all } (i,j) \in F \times W.$$

  The set of all possible assignments is denoted by $\mathcal{A}$.

Cooperative assignment game. *Given $[F, W, \boldsymbol{\alpha}, \mathcal{A}]$, the cooperative assignment game $G(v, N)$ is defined as follows. Let $N = F \cup W$ and define $v : S \subseteq N \to \mathbb{R}$ such that*

- $v(i) = v(\emptyset) = 0$ for all singletons $i \in N$,
- $v(S) = \alpha_{ij}$ for all $S = (i,j)$ such that $i \in F$ and $j \in W$,
- $v(S) = \max\{v(i_1, j_1) + ... + v(i_k, j_k)\}$ for every $S \subseteq N$,

where the maximum is taken over all sets $\{(i_1, j_1), ..., (i_k, j_k)\}$ consisting of disjoint pairs that can be formed by matching firms and workers in S. The number $v(N)$ specifies the value it takes of an optimal assignment.

Optimality.
An assignment $\mathbf{A}$ is optimal if $\sum_{(i,j) \in F \times W} \alpha_{ij} a_{ij} = v(N)$.

To simplify the notation, we shall add dummies to the smaller side of the market so that there is an equal number of firms and workers in the market, that is, $|F| = |W| = n$, and the value of any match with a dummy is zero.

Now we turn to feasibility of an *outcome*, by which we mean an assignment $\mathbf{A} \in \mathcal{A}$ together with a payoff vector $\boldsymbol{\phi} \in \mathbf{R}^{2n}_{+}$.

Feasibility.
A partnership (i,j) is feasible if $d_i + d_j \leq \alpha_{ij}$ (equivalently, their observable demands are compatible). An outcome $[\mathbf{A}, \boldsymbol{\phi}]$ is feasible if $\phi_i + \phi_j \leq \alpha_{ij}$ for all pairs (i,j) with $a_{ij} = 1$.

Henceforth we shall restrict our attention to feasible outcomes. However, a feasible outcome may be unstable if alternative partnerships can form that improve the deviating partners' positions.

Pairwise stability.
$[\mathbf{A}, \boldsymbol{\phi}]$ is pairwise stable if for all (i,j) with $a_{ij} = 1$, $\phi_i + \phi_j = \alpha_{ij}$ and $\phi_{i'} + \phi_j \geq \alpha_{i'j}$ for every alternative firm i' and $\phi_i + \phi_{j'} \geq \alpha_{ij'}$ for every alternative worker j'.

If $[\mathbf{A}, \boldsymbol{\phi}]$ is not pairwise stable, two agents exist who have a common incentive to deviate and form a new partnership.

The Core.
The core of an assignment game, $G(v, N)$, consists of the set $\mathbf{C}$ of all outcomes, $[\mathbf{A}, \boldsymbol{\phi}]$, such that $\mathbf{A}$ is an optimal assignment and $\boldsymbol{\phi}$ is pairwise stable.

Shapley & Shubik [1972] show that the core of any assignment game is always nonempty and coincides with the set of pairwise stable allocations that are supported by optimal assignments. In particular $[\mathbf{A}, \boldsymbol{\phi}]$ is in the core if all $a_{ij} = 0$ or 1, all $\phi_i \geq 0$ and the following conditions hold [8]:

(*i*) $\forall i, \sum_{j \in N} a_{ij} \leq 1,$

(*ii*) $\forall i, j, \; \phi_i + \phi_j \geq \alpha_{ij},$

(*iii*) $\forall i, \sum_{j \in N} a_{ij} < 1 \; \Rightarrow \; \phi_i = 0.$

(*iv*) $\forall i, j, \; a_{ij} = 1 \; \Rightarrow \; \phi_i + \phi_j = \alpha_{ij}.$

Subsequent literature has investigated the structure of the assignment game core, which turns out to be very rich.[9]

4. Evolving play

A fixed population of agents, $N = F \cup W$, repeatedly plays the assignment game $G(v, N)$ by submitting nonnegative real number demands to an electronic market maker and adjusting them dynamically as the game evolves.

States. At any period of time t, the *state* Z^t consists of a triple $[\mathbf{A}^t, d^t, \boldsymbol{\phi}^t]$:

- the assignment is $\mathbf{A}^t$,
- the demand vector is $d^t = \{d_1^t, ..., d_{2n}^t\}$,
- the payoff vector is $\boldsymbol{\phi}^t = \{\phi_1^t, ..., \phi_{2n}^t\}$ where,
  for any $i \in F$,
  $$\phi_i^t = \phi_i^t(A^t, d^t) = \begin{cases} d_i^t & \text{if } a_{ij}^t = 1 \text{ for some } j \\ 0 & \text{if } a_{ij}^t = 0 \text{ for all } j \end{cases}$$
  and, for any $j \in W$,
  $$\phi_j^t = \phi_j^t(\mathbf{A}^t, d^t) = \begin{cases} \alpha_{ij} - d_i^t & \text{if } a_{ij}^t = 1 \text{ for some } i \\ 0 & \text{if } a_{ij}^t = 0 \text{ for all } i \end{cases}$$

These payoffs follow from the assumption that one side's (the firms') bids are accepted.[10]

The set of all states will be denoted by Ω. To simplify the analysis, we shall assume that all demands, d_i, and all partnership values, α_{ij}, are multiples of an incremental rate of demand adjustment, $\delta \in \mathbf{R}^+$.

Note that the demand vectors constitute the invisible dual variables of the process and are such that, for all matched pairs i, j in any period t, $d_i^t + d_j^t = \alpha_{ij}$.

[8] These are the feasibility and complementary slackness conditions for the associated linear program and its dual.

[9] See, for example, Roth & Sotomayor [1992], Balinski & Gale [1987], Sotomayor [2003].

[10] Another assumption could be made here.

4.1. Behavioral dynamics

The essential features of the learning process are as follows:

1. Matched agents occasionally try to ratchet up their demands.
2. Unmatched agents ratchet down their demands in the hope of being rematched.
3. Existing matches are preserved unless made infeasible due to changes in demands or rematches with an active agent.

We shall now describe these transitions more precisely.

Given state Z^t at the beginning of period $t+1$, an agent $i \in N$ is drawn uniformly at random and becomes **activated**. All other agents' demands are held fixed at their period-t aspiration levels (at the previous demand if single, or the previous payoff if matched): $d_j^{t+1} = \max\{d_j^t, \phi_j^t\}$. The agent i who is selected makes new demands with probabilities to be specified below. The updating rule differs according to whether i is currently *assigned* ($a_{ij}^t = 1$ for some j) or is *unassigned*. Note that unassigned agents currently receive a payoff of zero, while assigned agents receive what they demand.

(*i.*) *Assigned agent*

If selected at the start of period $t+1$, an agent starts "looking around" with positive probability. An agent i (currently matched with j) probes for just a little bit more, i.e., he temporarily raises his demand by δ; otherwise he sticks with his previous demand throughout the period. (See Figure 1 for illustration.)

If there exists at least one $j' \in N$ such that $(d_i^t+\delta)+d_{j'}^t \leq \alpha_{ij'}$, such a j' is drawn uniformly at random (by the Center) and accepts the proposed match with i with probability $p \in (0,1)$. As a result, i's former partner is now unassigned. If there exists no such j', i and j remain matched with their previous demands.

(*ii.*) *Single agent*

A single agent ratchets down his demand in the hope of attracting a partner. (See Figure 2 for an illustration.)

If selected, the single agent i demands d_i^t. If there exists at least one $j \in N$ such that $d_i^t + d_j^t \leq \alpha_{ij}$, such a j is drawn uniformly at random by the Center and accepts the proposed match with i with probability $p \in (0,1)$. If no such j exists, i reduces his demand to $d_i^{t+1} = d_i^t - \delta$ with positive probability, otherwise continues to demand d_i^t. A special case is when $d_i^t = 0$, in which case $d_i^{t+1} = 0$ with probability 1.

Example

Let $N = F \cup W = \{f_1, f_2, f_3\} \cup \{w_1, w_2, w_3\}$, $\alpha_{11} = \alpha_{22} = \alpha_{33} = \alpha_{21} = \alpha_{23} = 1$ and $\alpha_{ij} = 0$ for all other (i,j) (we omit the connections worth zero). Now suppose that (f_2, w_2) and (f_3, w_3) are assigned and f_1 and w_1 are unassigned in period t. Let $\delta = 0.1$.

Figure 1: Demand increase by a matched agent.

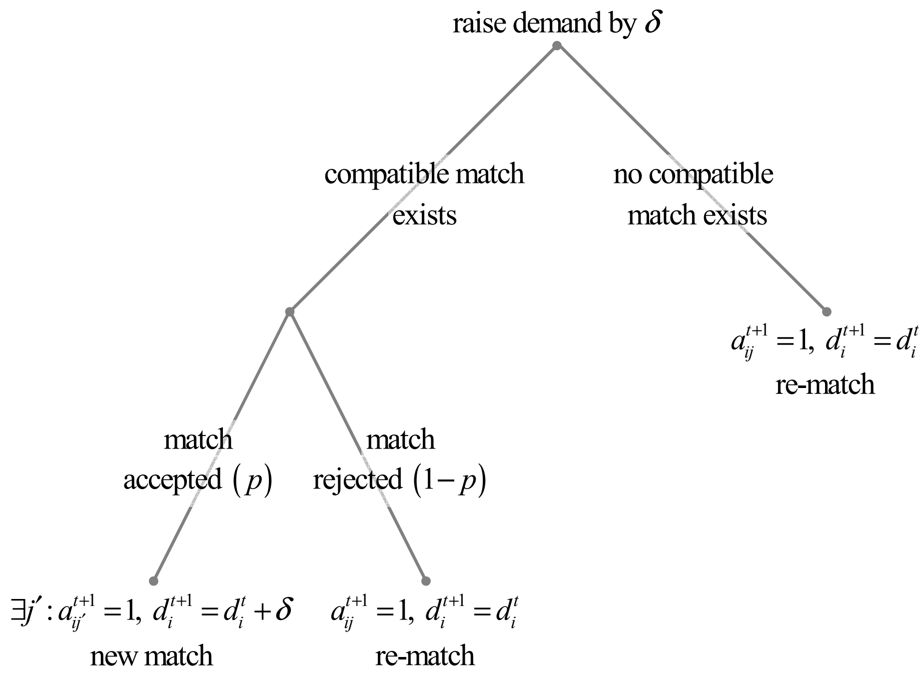

The current demand of each agent is shown next to the name of that agent, and the values α_{ij} are shown next to the edges. Solid edges indicate current matches.

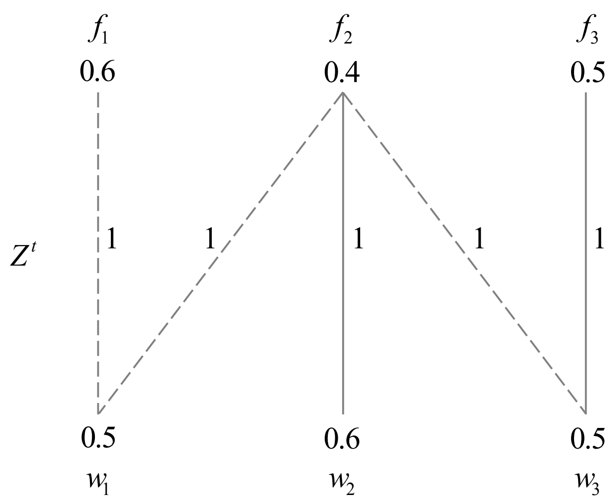

Figure 2: Demand decrease by a single agent.

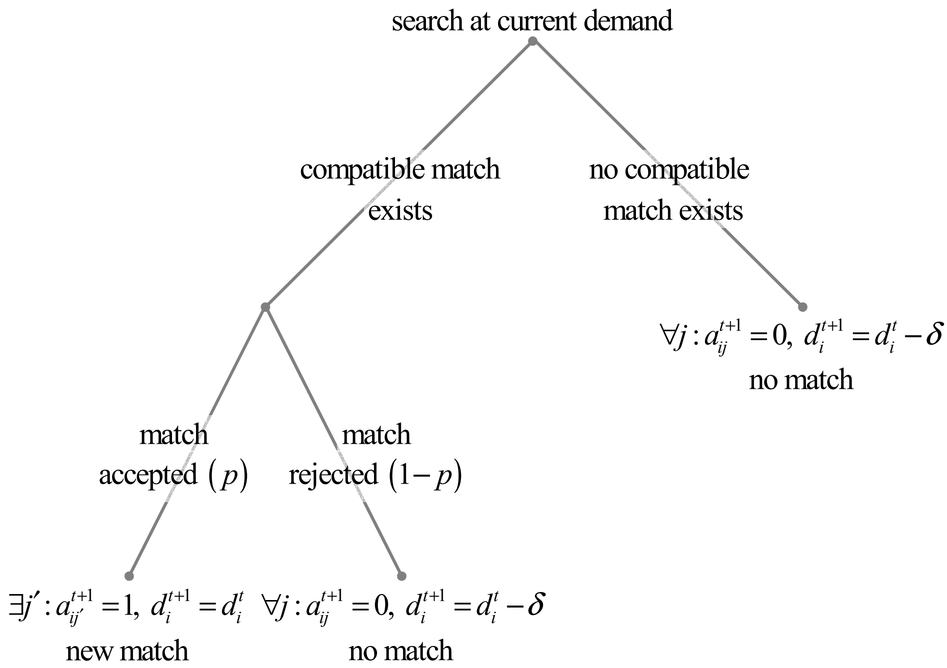

(*i.*) *Increases by matched agents.*

Suppose that f_2 probes $+0.1$, which is incompatible with w_2. f_2's new demand is compatible for w_1 and w_3. With positive probability, f_2 is matched at his new demand with w_1.

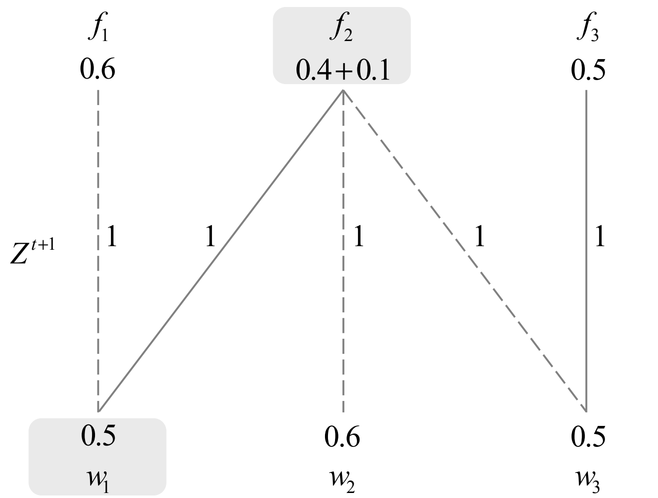

Next suppose that f_3 probes +0.1. No feasible match exists at the higher demand 0.6, hence f_3 returns to his previous demand and his previous partner w_3.

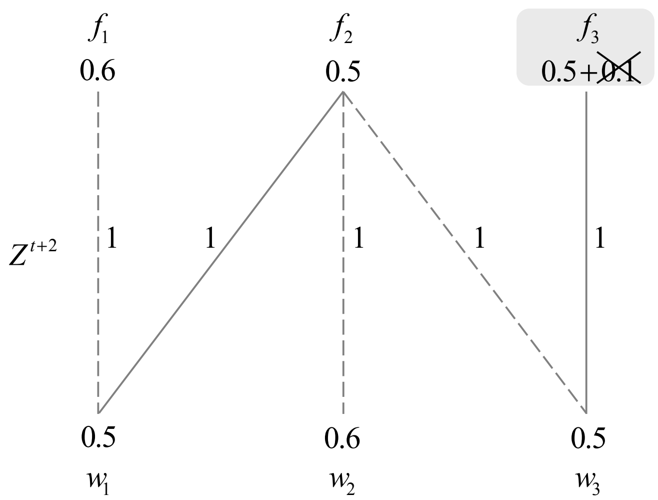

(*ii.*) *Decreases by single agents.*

Now suppose that the single agent f_1 reduces his demand by 0.1 to 0.5. With positive probability, he is then matched with the single agent w_1.

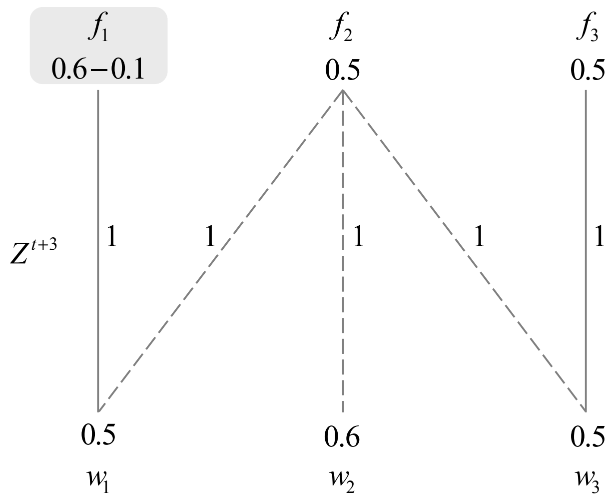

Next suppose that w_2 reduces his demand by 0.1 to 0.5. With positive probability, he is then matched with the single agent w_1, and the resulting state is core-stable.

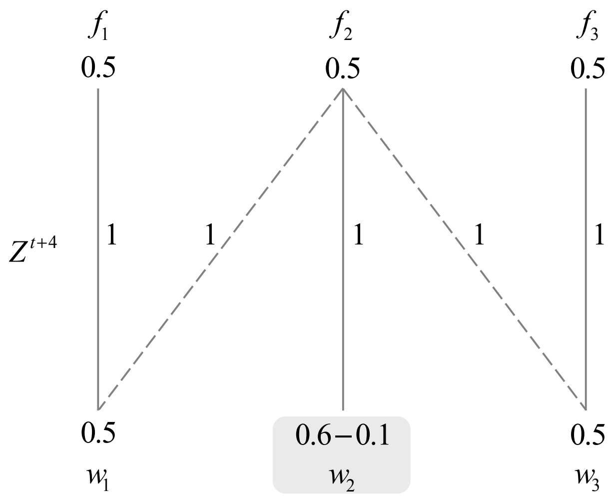

5. Stability

Recall that a state Z is defined by an assignment A and demands d that jointly determine payoffs $\boldsymbol{\phi}$. Recall that an outcome $[\mathbf{A}, \boldsymbol{\phi}]$ is in the core, $\mathbf{C}$, if conditions (i)-(iv) are satisfied. We shall write $\tilde{\mathbf{C}}$ for the subset of the core in which everyone is assigned.[11]

Theorem 1. *Given an assignment game $G(v, N)$, and $0 < \gamma < 1$, there exists a time $T_\gamma < \infty$ such that $[\mathbf{A}^t, \boldsymbol{\phi}^t] \in \tilde{\mathbf{C}}$ with probability at least $1 - \gamma$.*

Proof. Start with Z^t, demands d^t and a partial matching $\mathbf{A}^t$.

Claim 1. There is a positive probability path to demands d such that $d_i + d_j \geq \alpha_{ij}$ for all i, j.

Suppose $d_i + d_j < \alpha_{ij}$.

Case 1.) i is single (or j is single).

With positive probability, either i or j is activated and i, j match. As a result, next period, one of them has higher realized surplus which becomes new demand.

Case 2.) i and j are matched with each other.

In this case, $d_i + d_j = \alpha_{ij}$ because whenever two players are matched the entire surplus is allocated.

Case 3.) i and j are both matched but not with each other.

In this case, with positive probability, either i or j is activated and i, j match. As a result, next period, one of them has higher realized surplus which becomes his new demand, and other's surplus has not decreased.

Therefore, a suitable path can be constructed under which d increases monotonically until the claim is satisfied.

Claim 2. There is a positive probability path to demands d such that $d_i + d_j \geq \alpha_{ij}$ for all i, j and, for every i, either there exists a j such that $d_i + d_j = \alpha_{ij}$ or else $d_i = 0$.

The first part of the claim follows from claim 1. To establish the second part, let $\mathbf{A}$ be the current assignment. If $\mathbf{A}$ we are done. Else $\mathbf{A}$ is incomplete and there exists at least one single i. Assume $d_i > 0$, else we are done. With positive probability, i is selected and searches unsuccessfully, and lowers his demand by δ. Therefore, a suitable path can be constructed under which d decreases monotonically until the claim is satisfied. Note that at the end of this path not every agent with a positive demand needs to be matched.

Claim 3. There is a positive probability path to a pair $(d, \mathbf{A})$ such that $d_i + d_j \geq \alpha_{ij}$ for all i, j and $\mathbf{A}$ is complete.

Note that if the latter holds, then $\mathbf{A}$ and d are complementary dual variables so both are optimal and $\mathbf{A}$ is stable. (Recall complementary slackness holds.)

By claims 1 and 2, we know that $d_i + d_j \geq \alpha_{ij}$ for all i, j. Furthermore, for every i, there exists at least one j such that $d_i + d_j = \alpha_{ij}$ or else $d_i = 0$. Let $\mathbf{A}$ be the current

[11]Recall there may be dummy agents i with $\alpha_{ij} = 0$ for all $j \in N$. Such agents need not be matched in the core, and there are optimal allocations where not everyone is matched.

assignment which is incomplete (or else we are done).

Without loss of generality, we can assume there exists a single firm, say f_1, such that $d_{f_1} > 0$. Otherwise all single agents with demand 0 match with positive probability and the claim follows. Say an edge is *tight* if $d_i + d_j = \alpha_{ij}$. Starting at f_1, we shall construct a maximal path P of tight edges which alternate between unmatched and tight (dashed), and matched edges (solid) which are also tight by construction. We shall show that there exists a positive probability path of bounded length to a state which still satisfies the dual feasibility conditions and $|\mathbf{A}|$ increased.

Case 1.) Starting at f_1, there exists a maximal alternating path P of odd length.

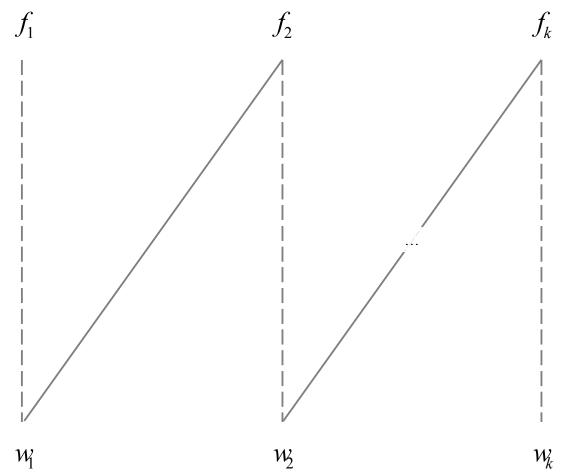

Let $P = (f_1, w_1, f_2, w_2, ..., w_{k-1}, f_k, w_k)$. First note that, since the path is maximal and of odd length, w_k must be unmatched. Now, with positive probability, f_1 is selected, searches and snags w_1. It follows that f_2 is now single. Again with positive probability, f_2 is selected, searches and snags w_2. Hence after k periods (k being the number of firms on path P) a state is reached where all players on P are matched and the demands have remained constant. Thus the number of matched pairs has increased by one.

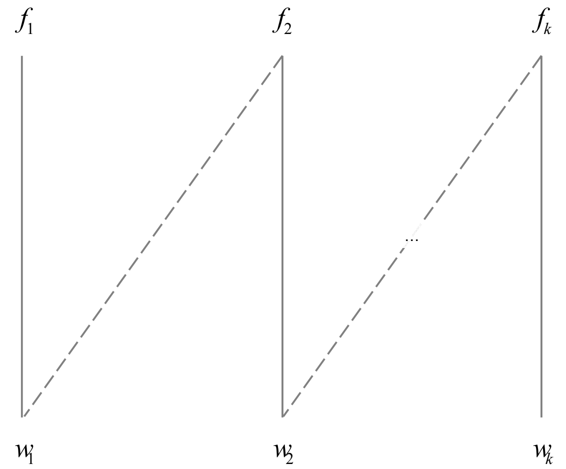

Case 2.) Starting at f_1, all maximal alternating paths are of even length.

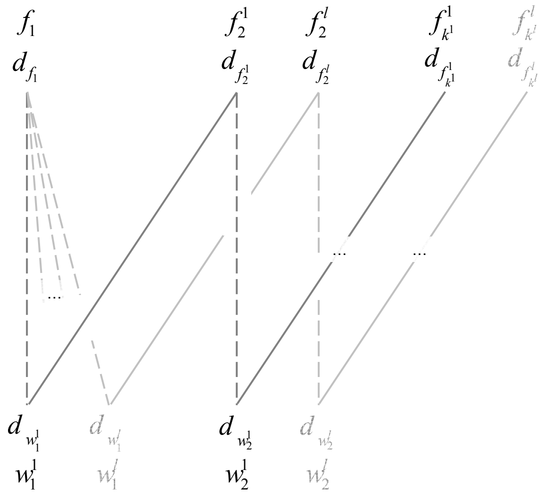

Let $\mathbf{P} = P^1, P^2, ..., P^l$ be the family of maximal alternating paths starting at f_1. On any path $P^r \in \mathbf{P}$ where $P^r = (f_1^r, w_1^r, f_2^r, w_2^r, ..., f_{k-1^r}^r, w_{k^r}^r, f_{k^r}^r)$, first note that, since the path is maximal and of even length, all other agents must be matched. With positive probability, f_1 is selected, searches unsuccessfully and reduces his demand by δ. f_1 remains single. Next with positive probability, w_1^r is selected, can increase his demand and match with f_1. f_1 is matched but w_1^r's former partner f_2^r is single. With positive probability, given f_2^r has a positive demand, f_2^r will reduce and be rematched with w_1^r, allowing w_2^r to increase and snag f_2^r. Now, reiterate the latter sequence of transitions up to firm $f_{k^r}^r$ and then for all $r \in 1, 2, ..., l$. Hence, after a finite number of periods, a state is reached where all players on every path $P \in \mathbf{P}$ have been unmatched, rematched, and are finally again matched with their original partners. In that state, the number of reductions by δ outnumbers the number of increases by one. Note that the new demand vector remains dual feasible.

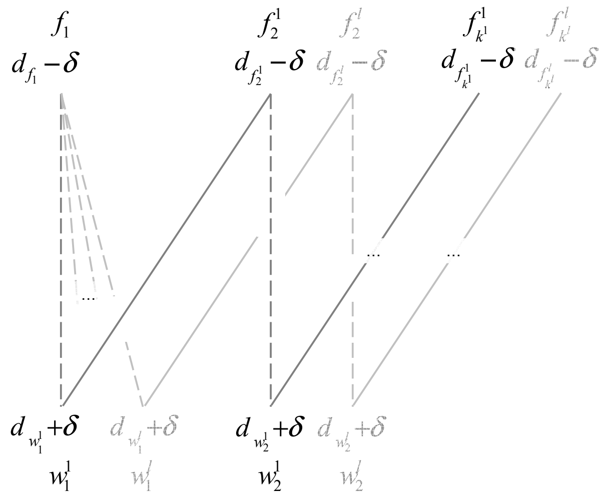

As long as all firms on paths in $\mathbf{P}$ have positive demands, the same path reoccurs with positive probability.

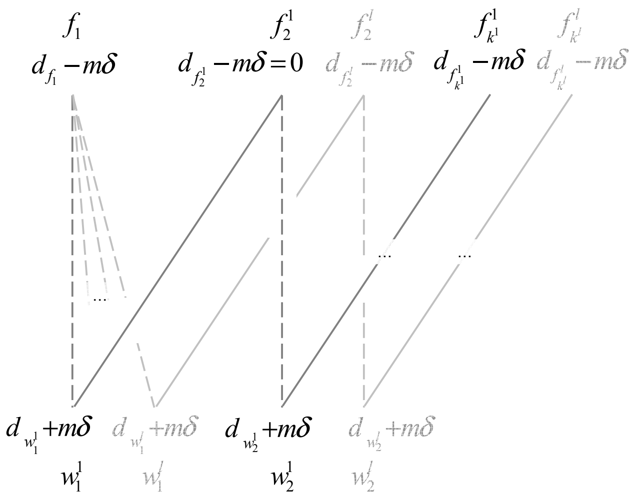

Let $m \cdot \delta$ ($m \geq 0$) be the minimal demand of any firm on $P \in \mathbf{P}$ at the beginning of the described transitions. After m-many rounds of reductions by the firms, a firm f_i^r on P exists with current demand zero (f_2^1 in the illustration) and hence no further reduction by f_i^r can occur. (If multiple firms in $\mathbf{P}$ have demand zero after m rounds, choose any firm which on its alternating path is the first firm with demand zero.) Hence after m rounds of reductions, with positive probability, the single agent f_1 searches successfully at current demand $d_{f_1} - m\delta$ and snags w_1^r. Then f_2^r becomes single. Again f_2^r, with positive probability, snags w_2^r until f_i^r is reached who is now single with demand zero. Note that all other paths in $\mathbf{P}$ did not change in the meantime and dual feasibility still holds since demands did not change. The described sequence of transitions led from an alternating path of even length with a single demanding > 0 to no alternating paths (in $\mathbf{P}$) and the only single agent in $\mathbf{P}$ demanding zero. Note that from the latter together with the fact that families of even alternating path only have one single agent, it follows that all other such families $\mathbf{P}'$ are disjunct with $\mathbf{P}$.

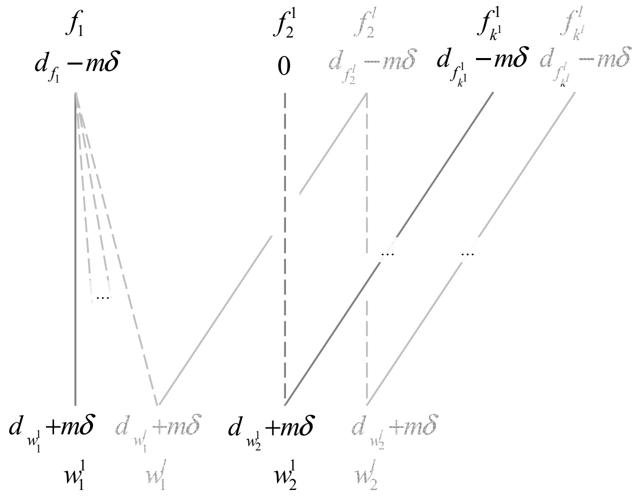

Thus by iterating case 2 for all disjunct families of even alternating paths, a state is reached where all single agents demand zero while preserving dual feasibility. Finally, with positive probability the single agents in the state (if there are any) find each other and are matched (at demands zero) with positive probability. The result is a complete core matching which is an absorbing state.

Consequently claims 1-3 jointly give a finite, positive probability path in which the process reaches a termination at state variables $(\mathbf{A}, d) \in \tilde{\mathbf{C}}$. The theorem follows.

$\square$

Readers familiar with the Hungarian method will note that our proof mirrors the adjustment chain of the algorithm introduced in Kuhn [1955]. It should be noted here that adjustment chains off the Hungarian paths, too, have positive probability because there is no central market authority enforcing any paths with global significance. With positive probability, however, agents – even though without knowledge of others' strategies, payoffs, or the structure of the game – behave as if implementing these adjustment and, thus, eventually implement the core by themselves.

6. Conclusion

In this paper we have demonstrated that agents in large decentralized matching markets can learn to play stable and efficient outcomes through a trial-and-error learning process. The agents need have no information about the distribution of others' preferences, their past behavior, or about the values of different partnerships. Core allocations are achieved by experiments on the singles market and resulting transition chains, even though the underlying behavior is nonstrategic, mirror key features of the Hungarian primal-dual method. A direction for further research would be to test experimentally how agents actually do adjust their bids and offers in markets of the type discussed above. Prior experimental work has shown that in fairly large matching games (seven on each side), agents frequently do converge to core outcomes (Corominas-Bosch [2004], Charness, Corominas-Bosch & Frechette [2007]). Convergence to the core is achieved relatively quickly as agents learn each others' preferences. It would be interesting to see whether this result holds up in situations with many more agents and when less information is made available. It would also be useful to examine their revision procedures as the game is repeated. Is there an asymmetry in the size of upward and downward revisions? Does giving agents more information about the others' preferences change their behavior? The answers to these questions will suggest ways in which the theoretical learning model can be aligned more closely with empirical behavior.

References

M. Agastya, "Adaptive Play in Multiplayer Bargaining Situations," *Review of Economic Studies 64*, 411-26, 1997.

M. Agastya, "Perturbed Adaptive Dynamics in Coalition Form Games," *Journal of Economic Theory 89*, 207-233, 1999.

T. Arnold & U. Schwalbe, "Dynamic coalition formation and the core," *Journal of Economic Behavior and Organization 49*, 363-380, 2002.

M. Balinski & D. Gale, "On the Core of the Assignment game," in *Functional Analysis, Optimization and Mathematical Economics*, L. J. Leifman (ed.), Oxford University Press, 274-289, 1987.

R. Bush & F. Mosteller, *Stochastic Models of Learning,* Wiley, 1955.

G. Charness, M. Corominas-Bosch & G. Fréchette, "Bargaining and network structure: An experiment," *Journal of Economic Theory 136*, 28-65, 2007.

M. Corominas-Bosch, "Bargaining in a network of buyers and sellers," *Journal of Economic Theory 115,* 35-77, 2004.

V. P. Crawford & E. M. Knoer, "Job Matching with Heterogeneous Firms and Workers," *Econometrica 49*, 437-540, 1981.

G. Demange & D. Gale, "The strategy of two-sided matching markets," *Econometrica 53,* 873-988, 1985.

G. Demange, D. Gale & M. Sotomayor, "Multi-item auctions," *Journal of Political Economics 94,* 863-872, 1986.

M. L. Elliott, "Inefficiencies in networked markets," *working paper,* Stanford University, 2010.

M. L. Elliott, "Search with multilateral bargaining," *working paper,* Stanford University, 2011.

W. Estes, "Towards a statistical theory of learning," *Psychological Review 57*, 94-107, 1950.

D. Foster & H. P. Young, "Regret testing: Learning to play Nash equilibrium without knowing you have an opponent," *Theoretical Economics 1*, 341-367, 2006.

D. Gale & L. S. Shapley, "College admissions and the stability of marriage," *American Mathematical Monthly 69*, 9-15, 1962.

F. Germano & G. Lugosi, "Global Nash convergence of Foster and Young's regret testing", *Games and Economic Behavior 60*, 135-154, 2007.

S. Hart & A. Mas-Colell, "Uncoupled Dynamics Do Not Lead to Nash Equilibrium," *American Economic Review 93*, 1830-1836, 2003.

R. J. Herrnstein, "Relative and absolute strength of response as a function of frequency of reinforcement", *Journal of Experimental Analysis of Behavior 4*, 267-272, 1961.

F. Hoppe, "Erfolg und Mißerfolg", *Psychologische Forschung 14,* 1-62, 1931.

R. Karandikar, D. Mookherjee, D. Ray & F. Vega-Redondo, "Evolving Aspirations and Cooperation", *Journal of Economic Theory 80*, 292-331, 1998.

A. S. Kelso & V. P. Crawford, "Job Matching, Coalition Formation, and Gross Substitutes", *Econometrica 50,* 1483-1504, 1982.

H. W. Kuhn, "The Hungarian Method for the assignment problem", *Naval Research Logistic Quarterly 2,* 83-97, 1955.

J. R. Marden, H. P. Young, G. Arslan, J. Shamma, "Payoff-based dynamics for multiplayer weakly acyclic games", *SIAM Journal on Control and Optimization 48,* special issue on "Control and Optimization in Cooperative Networks", 373-396, 2009.

J. Newton, "Non-cooperative convergence to the core in Nash demand games without random errors or convexity assumptions", Ph.D. thesis, University of Cambridge, 2010.

J. Newton, "Recontracting and stochastic stability in cooperative games", mimeo, University of Cambridge, 2011.

B. S. R. Pradelski & H. P. Young, "Learning Efficient Nash Equilibria in Distributed Systems", *Games and Economic Behavior*, http://dx.doi.org/10.1016/j.geb.2012.02.017, 2012.

A. E. Roth, "The Evolution of the Labor Markets for Medical Interns and Residents: A Case Study in Game Theory", *Journal of Political Economy 92,* 991-1016, 1984.

A. E. Roth, T. Sönmez & Ünver, "Pairwise kidney exchange", *Journal of Economic Theory 125*, 151-188, 2005.

A. E. Roth & M. Sotomayor, *Two-Sided Matching: A Study in Game Theoretic Modeling and Analysis*, Cambridge University Press, 1990.

A. E. Roth & M. Sotomayor, "Two-sided matching", in *Handbook of Game Theory with Economic Applications*, Volume 1, R. Aumann & S. Hart (eds.), 485-541, 1992.

K. Rozen, "Conflict Leads to Cooperation in Nash Bargaining", *mimeo,* Yale University, 2010a.

K. Rozen, "Conflict Leads to Cooperation in Nash Bargaining: Supplemental Result on Evolutionary Dynamics", *web appendix,* Yale University, 2010b.

T. Sandholm, "Computing in Mechanism Design", *New Palgrave Dictionary of Economics*, 2008.

R. Sawa, "Coalitional stochastic stability in games, networks and markets", *working paper,* University of Wisconsin-Madison, 2011.

L. S. Shapley & M. Shubik, "The Assignment Game I: The Core", *International Journal of Game Theory 1*, 111-130, 1972.

M. Sotomayor, "Some further remark on the core structure of the assignment game", *Mathematical Social Sciences 46,* 261-265, 2003.

E. Thorndike, "Animal Intelligence: An Experimental Study of the Associative Processes in Animals", *Psychological Review 8*, 1898.

H. P. Young, "Learning by trial and error", *Games and Economic Behavior 65*, 626-643, 2009.