

6-10-2013

The Evolution of Core Stability in Decentralized Matching Markets

Heinrich H. Nax

Paris School of Economics

Bary S. R. Pradelski

University of Oxford, Oxford-Man Institute, bary.pradelski@oxford-man.ox.ac.uk

H. Peyton Young

University of Oxford

Follow this and additional works at: <http://services.bepress.com/feem>

Recommended Citation

Nax, Heinrich H.; Pradelski, Bary S. R.; and Young, H. Peyton, "The Evolution of Core Stability in Decentralized Matching Markets" (June 10, 2013). *Fondazione Eni Enrico Mattei Working Papers*. Paper 801.
<http://services.bepress.com/feem/paper801>

The Evolution of Core Stability in Decentralized Matching Markets

Heinrich H. Nax, Bary S. R. Pradelski & H. Peyton Young*

June 1, 2012

This version: February 24, 2013

Abstract

Decentralized matching markets 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, completely uncoupled learning process in such environments that leads to stable and efficient outcomes. Agents on each side of the market make bids for potential partners and are matched if their bids are mutually profitable. Matched agents occasionally experiment with higher bids if on the buy-side (or lower bids if on the sell-side), while single agents, in the hope of attracting partners, lower their bids if on the buy-side (or raise their bids if on the sell-side). This simple and intuitive learning process implements core allocations even though agents have no knowledge of other agents' strategies, pay-offs, 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

*We thank Itai Arieli, Gabrielle Demange, Gabriel Kreindler and Tom Norman for suggesting a number of improvements to an earlier draft, and are grateful to participants at the 23rd International Conference on Game Theory at Stony Brook University, the Paris Game Theory Seminar, the AFOSR MUIR 2013 meeting at MIT, and the 18th Coalitions and Network Workshop at the University of Warwick. This research was supported by the United States Air Force Office of Scientific Research Grant FA9550-09-1-0538. Heinrich Nax also acknowledges support of the Agence Nationale de Recherche project NET, Bary Pradelski of the Oxford-Man Institute of Quantitative Finance.

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. In this paper we propose a dynamic model that incorporates these features and explore its convergence and welfare properties. We see this as a first step towards 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 prices at which they are willing to be matched. Examples include online platforms for matching buyers and sellers of goods, for matching workers and firms, for matching hotels with clients, and for matching men and women.¹ 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 is different in character: agents make bids that are conditional on the characteristics of those with whom they wish to be matched, and a profitable (not necessarily optimal) set of matches is realized at each point in time. There is no presumption that agents or a central authority know anything about others' preferences, or that they can deduce such information from prior rounds. Instead, the agents, through trial-and-error, look for profitable matches and adjust their bids dependent on whether being matched or being single.

Rules of this type have a long history in the psychology literature (Thorndike [1898], Hoppe [1931], Estes [1950], Bush & Mosteller [1955], Herrnstein [1961]). To the best of our knowledge, however, such a framework has not previously been used in the study of matching markets in cooperative games.² The approach seems especially well-suited to modeling behavior in large decentralized matching markets, where agents have little information about the overall game and about the identity of the other market participants. We show that a class of learning rules with simple adjustment dynamics of this type implements the core with probability one after finite time. The main contribution of the paper is to show that this can be achieved even though agents have no knowledge of other agents' strategies or preferences, 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.

¹An example is www.priceline.com's Name-Your-Own-Price[®]; www.HireMeNow.com's Name-Your-Own-Wage[™] uses a similar reverse auction mechanism for temporary employment.

²For a review of other mechanisms in the literature see Sandholm [2008].

2. *Related literature*

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). For subsequent literature, see Crawford & Knoer [1981], Kelso & Crawford [1982], Demange & Gale [1985], Demange, Gale & Sotomayor [1986], Shimer [2007, 2008], Elliott [2010, 2011].³ These algorithms have been successfully applied in situations where agents engage in a formal application process, such as students seeking admission to universities, doctors applying for hospital residencies, or transplant patients looking for organ donors.⁴

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. Examples include markets matching buyers with sellers or firms with workers. These constitute a special class of cooperative games with transferable utility (Shapley & Shubik [1972]). 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.

In our model, there is a simple clearing mechanism, “the Matchmaker”, whose function is to match agents with mutually profitable bids and offers who are currently “active”. Neither the players nor the Matchmaker have enough information to optimize the value of the matches. This limited role is what distinguishes our Matchmaker from a central authority governing a traditional matching environment as in, for example, the National Resident Matching Program (Roth & Peranson [1999]). We shall show that simple adjustment rules by the agents lead to efficient and stable outcomes without any centralized information about which matches are best.

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, 2012], 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.⁵ In each period a player is activated 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 provided the game has a nonempty interior core.⁶

³Shimer [2007, 2008] and Elliott [2010, 2011] explore empirical and network elements of matching.

⁴See Roth [1984], Roth & Peranson [1999] for discussions of the US medical resident market, and Roth, Sönmez & Ünver [2005] for the kidney exchange market.

⁵Newton [2012] nests the models of Agastya [1997, 1999] and Rozen [2010a, 2010b] as special cases.

⁶The interior of the core is said to be nonempty if the core is of maximal dimension. This is not

The main difference between existing learning models and ours is the amount of information available to market participants.⁷ The approach we take here requires considerably less information on the part of the agents: players know nothing about the other players' current or past behavior, or their 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. Adaptive rules of this type are said to be *completely uncoupled* (Foster & Young [2006]).⁸ In recent years it has been shown that there are families of such rules that lead to equilibrium behavior in generic non-cooperative games (Karandikar, Mookherjee, Ray & Vega-Redondo [1998], Foster & Young [2006], Germano & Lugosi [2007], Marden, Young, Arslan & Shamma [2009], Young [2009], Pradelski & Young [2012]). Here we shall demonstrate that a very simple rule of this form leads to stability and optimality in two-sided matching markets.

3. Matching markets with transferable utility

In this section we shall introduce the conceptual framework for analyzing matching markets with transferable utility; in the next section we introduce the learning process itself. The population $N = F \cup W$ consists of firms $F = \{f_1, \dots, f_m\}$ and workers $W = \{w_1, \dots, w_n\}$.⁹ They interact by submitting bids and offers to “the Matchmaker”, whose function is to propose matches between firms and workers whose bids and offers are mutually profitable.

3.1 Static components

Willingness to pay. Each firm i has a *willingness to pay*, $p_{ij}^+ \geq 0$, for being matched to worker j .

Willingness to accept. Each worker j has a *willingness to accept*, $q_{ij}^- \geq 0$, 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 Matchmaker.

It will be convenient to assume that all values p_{ij}^+ and q_{ij}^- can be expressed as multiples of some minimal unit of currency δ , e.g., “dollars”. At the end of section 5 (corollary 2), we shall show that all the results extend to continuous space.

guaranteed (and not likely) in many applications.

⁷Moreover, the core of an assignment game typically has an empty interior, so that the aforementioned results cannot be applied directly to the present set-up.

⁸This definition is a strengthening of *uncoupled* rules introduced by Hart & Mas-Colell [2003].

⁹The two sides of the market could also, for example, represent buyers and sellers, or men and women in a (monetized) marriage market.

3.2 Dynamic components

Let $t = 0, 1, 2, \dots$ be the time periods.

Assignment. For all agents $(i, j) \in F \times W$, let $a_{ij}^t \in \{0, 1\}$.

$$\text{If } (i, j) \text{ is } \begin{cases} \text{matched} & \text{then } a_{ij}^t = 1, \\ \text{unmatched} & \text{then } a_{ij}^t = 0. \end{cases} \quad (1)$$

If for a given agent $i \in N$ there exists j such that $a_{ij}^t = 1$ we shall refer to that agent as *matched*; otherwise i is *single*.

Aspiration level. At the end of any period t , a player has an *aspiration level*, d_i^t , which determines the minimal payoff at which he is willing to be matched. Let $\mathbf{d}^t = \{d_i^t\}_{i \in F \cup W}$.

Bids. In any period t , each agent submits conditional bids for players on the other side of the market to the Matchmaker. We assume that these bids are such that the resulting payoff to a player (if he is matched) is at least equal to his aspiration level, and with positive probability is exactly equal to his aspiration level. Moreover, every pair of players submit bids to be matched with each other in any given period with positive probability.

Formally, firm $i \in F$ submits a vector of random bids $b_i^t = (p_{i1}^t, \dots, p_{in}^t)$, where p_{ij}^t is the maximal amount i is currently willing to pay if matched with $j \in W$. Similarly, worker $j \in W$ submits $b_j^t = (q_{1j}^t, \dots, q_{mj}^t)$, where q_{ij}^t is the minimal amount j is currently willing to accept if matched with $i \in F$. The bids are separable into two components; the current aspiration level beyond firm i 's (worker j 's) willingness to pay (accept) and a random variable P_{ij}^t (Q_{ij}^t):

$$\text{for all } i, j, \quad p_{ij}^t = (p_{ij}^+ - d_i^{t-1}) - P_{ij}^t \quad \text{and} \quad q_{ij}^t = (q_{ij}^- + d_j^{t-1}) + Q_{ij}^t \quad (2)$$

Consider, for example, worker j 's bid for firm i . The amount q_{ij}^- is the minimum that j would ever accept to be matched with i , while d_j^{t-1} is his previous aspiration level over and above the minimum. Thus Q_{ij}^t is j 's attempt to get even more in the current period. We assume that P_{ij}^t, Q_{ij}^t are independent random variables that take values in $\delta\mathbb{N}_0$ where 0 has positive probability.¹⁰ Note that if the random variable is zero, the agent bids exactly according to his current aspiration level. We shall use the convention $p_{ij}^t = -\infty$ ($q_{ij}^t = \infty$) if firm i (worker j) does not bid for worker j (firm i) in the current period.

Tie-breaking. A firm (worker) prefers to be matched at p_{ij}^+ (q_{ij}^-) rather than being single.

Profitability. A pair of bids (p_{ij}^t, q_{ij}^t) is *profitable* if $p_{ij}^t > q_{ij}^t$ or if $p_{ij}^t \geq q_{ij}^t$ and i and j are single.

Matchmaker. At each moment in time, at most one player is *active*. The Matchmaker observes

¹⁰Note that $\mathbb{P}[P_{ij}^t = 0] > 0$ and $\mathbb{P}[Q_{ij}^t = 0] > 0$ are trivial assumptions, since we can adjust p_{ij}^+ and q_{ij}^- in order for it to hold.

- the current bids and which agent is currently active,
- who is currently matched with whom and which bids are profitable.

The Matchmaker then matches the active agent to some agent (if one exists) with whom the bids are profitable. (Details about the Matchmaker and about how players are activated are specified in the next section.)

Prices. When i is matched with j given bids $p_{ij}^t \geq q_{ij}^t$, the resulting *price*, π_{ij}^t , is the average of the players' bids subject to "rounding". Namely, there is an integer k such that

$$\begin{array}{ll} \text{if } p_{ij}^t + q_{ij}^t = 2k\delta & \text{then } \pi_{ij}^t = k\delta, \\ \text{if } p_{ij}^t + q_{ij}^t = (2k+1)\delta & \text{then } \begin{cases} \pi_{ij}^t = k\delta & \text{with probability 0.5,} \\ \pi_{ij}^t = (k+1)\delta & \text{with probability 0.5.} \end{cases} \end{array} \quad (3)$$

This implies that when a pair is matched we have

$$p_{ij}^t = q_{ij}^t. \quad (4)$$

Note that when a new match forms that is profitable (as defined earlier), neither of the agents is worse off, and if one agent was previously matched both agents are better off in expectation due to the rounding rule.¹¹

3.3 Assignment games

We are now in a position to formally define matching markets and assignment games.

Match value. Assume that utility is linear and separable in money. The *value* of a match $(i, j) \in F \times W$ is the potential surplus

$$\alpha_{ij} = (p_{ij}^+ - q_{ij}^-)_+. \quad (5)$$

Matching market. The *matching market* is described by $[F, W, \alpha, \mathbf{A}]$:

- $F = \{f_1, \dots, f_m\}$ is a set of m firms (or men or sellers),
- $W = \{w_1, \dots, w_n\}$ is a set of n workers (or women or buyers),
- $\alpha = \begin{pmatrix} \alpha_{11} & \dots & \alpha_{1n} \\ \vdots & \alpha_{ij} & \vdots \\ \alpha_{m1} & \dots & \alpha_{mn} \end{pmatrix}$ is the matrix of match values.
- $\mathbf{A} = \begin{pmatrix} a_{11} & \dots & a_{1n} \\ \vdots & a_{ij} & \vdots \\ a_{m1} & \dots & a_{mn} \end{pmatrix}$ is the assignment matrix with 0/1 values and row/column sums at most one.

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

¹¹It is not necessary for our result to assume the price to be the average of the bids. We only need that the price, with positive probability, is different from a players bid when bids strictly cross.

Cooperative assignment game. Given $[F, W, \alpha]$, the *cooperative assignment game* $G(v, N)$ is defined as follows. Let $N = F \cup W$ and define $v : S \subseteq N \rightarrow \mathbb{R}$ 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 $j \in W$,
- $v(S) = \max\{v(i_1, j_1) + \dots + v(i_k, j_k)\}$ for every $S \subseteq N$,

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

States. The *state* at the end of period t is given by $Z^t = [\mathbf{A}^t, \mathbf{d}^t]$ where $\mathbf{A} \in \mathcal{A}$ is an assignment and \mathbf{d}^t is the aspiration level vector. Denote the set of all states by Ω .

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

Pairwise stability. An aspiration level \mathbf{d}^t is *pairwise stable* if $\forall i, j$ with $a_{ij} = 1$,

$$p_{ij}^+ - d_i^t = q_{ij}^- + d_j^t, \tag{6}$$

and $p_{i'j}^+ - d_{i'}^t \leq q_{i'j}^- + d_j^t$ for every alternative firm i' and $q_{ij'}^- + d_{j'}^t \geq p_{ij'}^+ - d_i^t$ for every alternative worker j' .

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

Shapley & Shubik [1972] show that the core of any assignment game is always non-empty and coincides with the set of pairwise stable aspiration levels that are supported by optimal assignments. (In Shapley & Shubik [1972] this is formulated in terms of payoffs, as we now proceed to define.) Subsequent literature has investigated the structure of the assignment game core, which turns out to be very rich.¹²

Payoffs. Given $[\mathbf{A}^t, \mathbf{d}^t]$ the *payoff* to firm i / worker j is

$$\phi_i^t = \begin{cases} p_{ij}^+ - \pi_{ij}^t & \text{if } i \text{ is matched to } j, \\ 0 & \text{if } i \text{ is single.} \end{cases} \quad \phi_j^t = \begin{cases} \pi_{ij}^t - q_{ij}^- & \text{if } j \text{ is matched to } i, \\ 0 & \text{if } j \text{ is single.} \end{cases} \tag{7}$$

In our framework, $[\mathbf{A}, \mathbf{d}]$ is in the core if all $a_{ij} = 0$ or 1, all $\phi_i \geq 0$ and the following conditions hold:¹³

- (i) $\forall i \in F, \sum_{j \in W} a_{ij} \leq 1$ and $\forall j \in W, \sum_{i \in F} a_{ij} \leq 1$,
- (ii) $\forall i, j \in F \times W, \phi_i + \phi_j \geq \alpha_{ij}$,
- (iii) $\forall i \in F, \sum_{j \in W} a_{ij} < 1 \Rightarrow \phi_i = 0$ and $\forall j \in W, \sum_{i \in F} a_{ij} < 1 \Rightarrow \phi_j = 0$.
- (iv) $\forall i, j \in F \times W, a_{ij} = 1 \Rightarrow \phi_i + \phi_j = \alpha_{ij}$.

¹²See, for example, Roth & Sotomayor [1992], Balinski & Gale [1987], Sotomayor [2003].

¹³These are the feasibility and complementary slackness conditions for the associated linear program and its dual (see, for example, Balinski [1965]).

4. *Evolving play*

A fixed population of agents, $N = F \cup W$, repeatedly plays the assignment game $G(v, N)$ by submitting bids to the Matchmaker and by adjusting them dynamically as the game evolves. Agents become activated spontaneously according to independent Poisson arrival processes. For simplicity we shall assume that the arrival rates are the same for all agents, but our results also hold when the rates differ across agents (for example, single agents might become active at a faster rate than matched agents). The distinct times at which one agent becomes active will be called *periods*.

4.1. *Behavioral dynamics*

The essential steps and features of the learning process are as follows. At the start of period $t + 1$:

1. A unique agent becomes active.
- 2a. If a profitable match exists given the current bids, the Matchmaker selects a randomly drawn profitable match with the active agent.
- 2b. If no profitable match exists, the Matchmaker rejects the bids.
- 3a. If a new match (i, j) is formed, the price is the average of the two bids (subject to rounding). The bids of i and j next period are at least their realized payoffs this period.
- 3b. If no new match is formed, the active agent, if he was previously matched, keeps his previous bid and stays with his previous partner. If he was previously single, he remains single and lowers his aspiration level with positive probability.

We shall now describe the process in more detail, distinguishing the cases where the active agent is currently *matched* or *single*. Let Z^t be the state at the end of period t (and the beginning of period $t + 1$), and let i be the unique *active* agent.

I. *The active agent is currently matched*

Let J' be the set of players with whom i can be profitably matched, that is,

$$J' = \{j' : p_{ij'}^t > q_{ij'}^t\}. \quad (8)$$

If $J' \neq \emptyset$, some agent $j' \in J'$ is drawn uniformly at random by the Matchmaker, and is matched with i .¹⁴ As a result, i 's former partner is now single (and so is j' 's former partner if j' was matched in period t). The price governing the new match, $\pi_{ij'}^{t+1}$, is the average (subject to rounding) of $p_{ij'}^t$ and $q_{ij'}^t$.

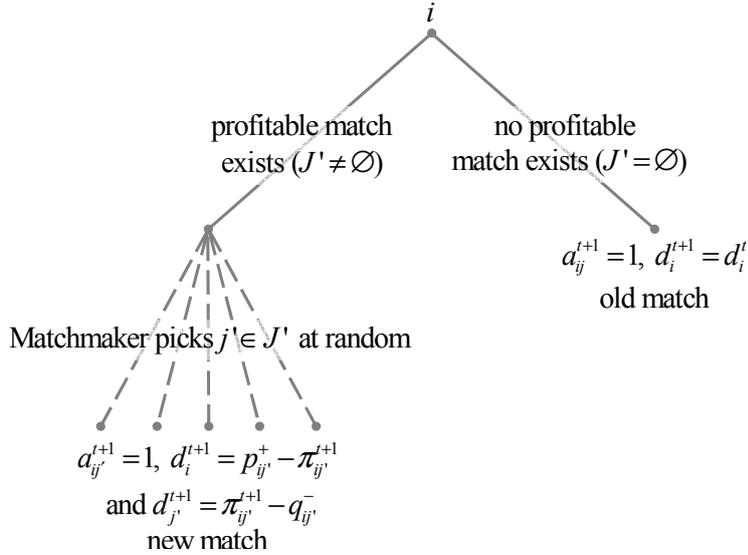
¹⁴Instead of a uniform random draw from the profitable matches, priority could be given to those involving single agents; or any distribution with full support on the profitable matches can be used.

At the end of period $t + 1$, the aspiration levels of the newly matched pair (i, j') are adjusted according to their newly realized payoffs:

$$d_i^{t+1} = p_{ij'}^+ - \pi_{ij'}^{t+1} \quad \text{and} \quad d_{j'}^{t+1} = \pi_{ij'}^{t+1} - q_{ij'}^- \quad (9)$$

All other aspiration levels and matches remain fixed. If $J' = \emptyset$, i remains matched with his previous partner and keeps his previous aspiration level. See Figure 1 for an illustration.

Figure 1: Transition diagram for active, matched agent (period $t + 1$).



II. The active agent is currently single

Let J be the set of players with whom i can be profitably matched, that is,

$$J = \{j : j \text{ single, } p_{ij}^t \geq q_{ij}^t\} \cup \{j : j \text{ matched and } p_{ij}^t > q_{ij}^t\}. \quad (10)$$

If $J \neq \emptyset$, some agent $j \in J$ is drawn uniformly at random by the Matchmaker, and is matched with i . If j was matched in period t his former partner is now single. The price governing the new match, π_{ij}^{t+1} , is the average (subject to rounding) of p_{ij}^t and q_{ij}^t .

At the end of period $t + 1$, the aspiration levels of the newly matched pair (i, j) are adjusted to equal their newly realized payoffs:

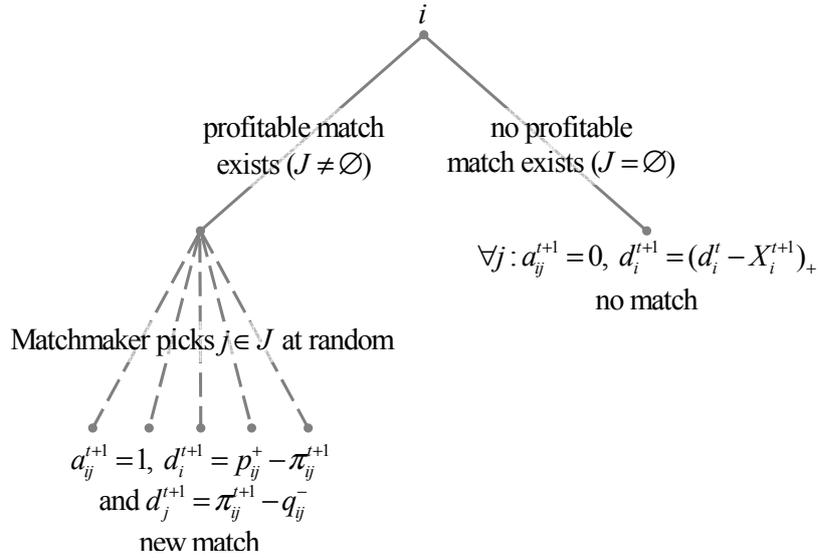
$$d_i^{t+1} = p_{ij}^+ - \pi_{ij}^{t+1} \quad \text{and} \quad d_j^{t+1} = \pi_{ij}^{t+1} - q_{ij}^-. \quad (11)$$

All other aspiration levels and matches remain as before. If $J = \emptyset$, i remains single and, with positive probability, reduces his aspiration level,

$$d_i^{t+1} = (d_i^t - X_i^{t+1})_+, \quad (12)$$

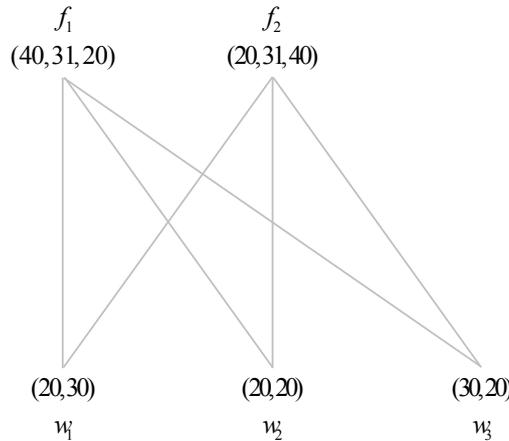
where X_i^{t+1} is an independent random variable taking values in $\delta \cdot \mathbb{N}_0$, such that $\mathbb{E}[X_i^t] > C$ (where $C > 0$ is a constant independent of δ), and δ occurs with positive probability. See Figure 2 for an illustration.

Figure 2: Transition diagram for active, single agent (period $t + 1$).



4.2. Example

Let $N = F \cup W = \{f_1, f_2\} \cup \{w_1, w_2, w_3\}$, $p_{1j}^+ = 40, 31, 20$ and $p_{2j}^+ = 20, 31, 40$ for $j = 1, 2, 3$, and $q_{i1}^- = 20, 30$, $q_{i2}^- = 20, 20$ and $q_{i3}^- = 30, 20$ for $i = 1, 2$.

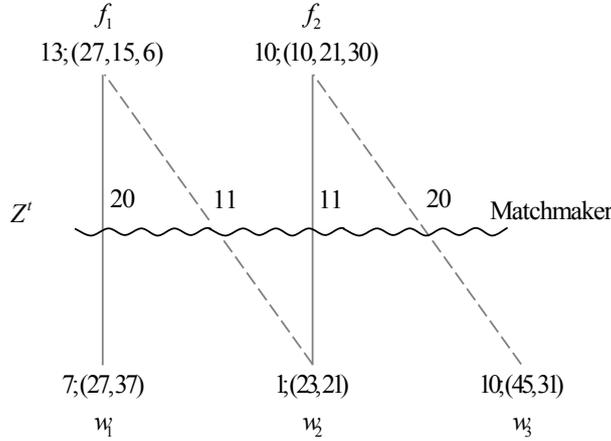


Then one can compute the match values: $\alpha_{11} = \alpha_{23} = 20$, $\alpha_{12} = \alpha_{22} = 11$, and $\alpha_{ij} = 0$ for all other pairs (i, j) . Let $\delta = 1$.

period t : Current state

Suppose that, in some period t , (f_1, w_1) and (f_2, w_2) are matched and w_3 is single. In the illustrations below, the current aspiration level and bid vector of each agent is shown next to the name of that agent, and the values α_{ij} are shown next to the edges (if positive). Solid edges indicate matched pairs, and dashed edges indicate unmatched pairs. (Edges with value zero are not shown.) The wavy line indicates that no player can see the bids or the status of the players on the other side of the market.

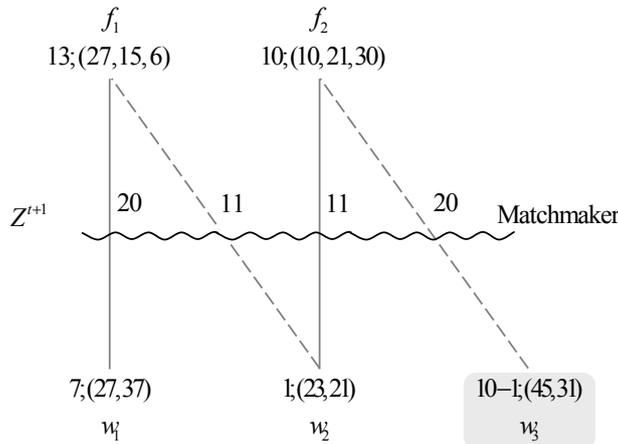
Note that some of the bids for players which are currently not matched may exceed the respective match values. For example f_2 , at the beginning of the period, was willing to pay 30 for w_3 , but w_3 was asking for 31 from f_2 , 1 above the minimum bid not violating his aspiration level. Further, note that, some matches can never occur. For example f_1 is never willing to pay more than 20 for w_3 , but w_3 would only accept a price above 30 from f_1 .



Note that the aspiration levels satisfy $d_i^t + d_j^t \geq \alpha_{ij}$ for all i and j , but the assignment is not optimal (firm 2 should match with worker 3).

period $t + 1$: *Activation of single agent w_3*

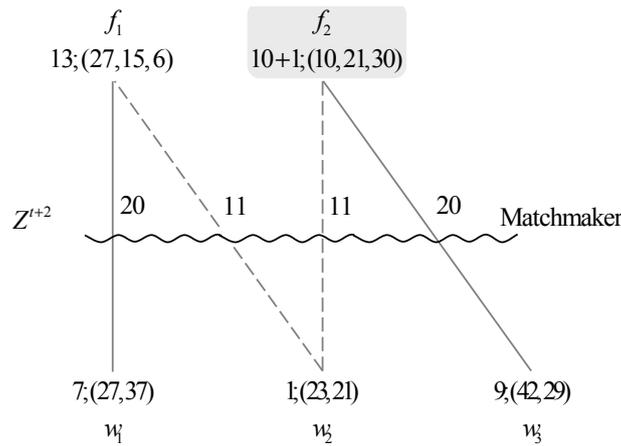
w_3 's current aspiration level is too high in the sense that he has no profitable matches. Hence, independent of the specific bids he makes, he remains single and, with positive probability, reduces his aspiration level by 1.



period $t + 2$: *Activation of matched agent f_2*

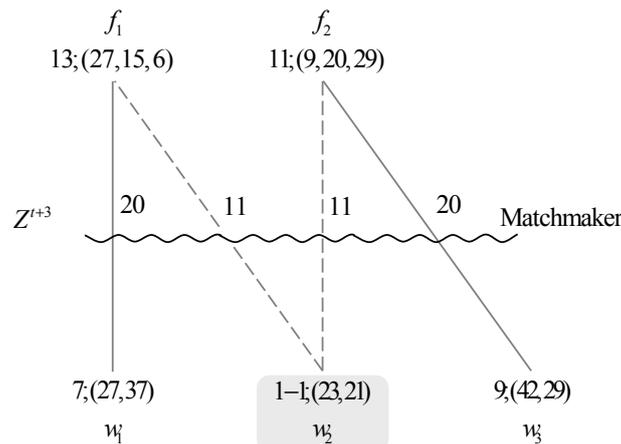
f_2 's only profitable match, under any possible bid, is with w_3 . With positive probability f_2 bids 30 for w_3 and w_3 bids 29 for f_2 (hence the match is profitable), and the match forms. With probability 0.5 the price is set to 29 such that f_2 raises his aspiration level by one unit (11) and w_3 keeps his aspiration level (9), while with probability 0.5 the price is set to 30, f_2 keeps his aspiration level (10) and w_3 raises his aspiration level by one unit

(10). (Thus in expectation the active agent f_2 gets a higher payoff than before.)



period $t + 3$: *Activation of single agent w_2*

w_2 's current aspiration level is too high in the sense that he has no profitable matches (under any possible bids). Hence he remains single and, with positive probability, reduces his aspiration level by 1.



The resulting state is in the core.¹⁵

5. Core stability

Recall that a state Z^t is defined by an assignment \mathbf{A}^t and aspiration levels \mathbf{d}^t that jointly determine the payoffs. Further Z^t is in the core, \mathbf{C} , if conditions (i)-(iv) are satisfied.

Theorem 1. *Given an assignment game $G(v, N)$, from any initial state $Z^t = [\mathbf{A}^0, \mathbf{d}^0] \in \Omega$, the process is absorbed into the core in finite time with probability 1.*

¹⁵Note that the states Z^{t+2} and Z^{t+3} are both in the core, but Z^{t+3} is absorbing whereas Z^{t+2} is not.

Throughout the proof we shall omit the time superscript since the process is time-homogeneous. The general idea of the proof is to show a particular path leading into the core which has positive probability. It will simplify the argument to restrict our attention to a particular class of paths with the property that the realizations of the random variables P_{ij}^t, Q_{ij}^t are always 0 and the realizations of X_i^t are always δ . (Recall that P_{ij}^t, Q_{ij}^t determine the gaps between the bids and the aspiration levels, and X_i^t determines the reduction of the aspiration level by a single agent.) One obtains from equation (2) for the bids:

$$\text{for all } i, j, \quad p_{ij}^t = p_{ij}^+ - d_i^{t-1} \quad \text{and} \quad q_{ij}^t = q_{ij}^- + d_j^{t-1} \quad (13)$$

Recall that every two agents post bids for each other with positive probability in any given period. We shall therefore construct a path along which the relevant agents in any period post bids for each other in that period. Jointly with equation (5), we can then say that a pair of aspiration levels (d_i^t, d_j^t) is *profitable* if

$$\text{either } d_i^t + d_j^t < \alpha_{ij}, \quad \text{or} \quad d_i^t + d_j^t = \alpha_{ij} \text{ and both } i \text{ and } j \text{ are single.} \quad (14)$$

Restricting attention to this particular class of paths will permit a more transparent analysis of the transitions, which we can describe solely in terms of the aspiration levels.

We shall proceed by establishing the following two claims.

Claim 1. There is a positive probability path to aspiration levels \mathbf{d} such that $d_i + d_j \geq \alpha_{ij}$ for all i, j and such that, for every i , either there exists a j such that $d_i + d_j = \alpha_{ij}$ or else $d_i = 0$.

Any aspiration levels satisfying Claim 1 will be called *good*. Note that, even if aspiration levels are good, the assignment does not need to be optimal and not every agent with a positive aspiration level needs to be matched. (See the period- t example in the preceding section.)

Claim 2. Starting at any state with good aspiration levels, there is a positive probability path to a pair (\mathbf{A}, \mathbf{d}) where \mathbf{d} is good, \mathbf{A} is optimal, and all singles' aspiration levels are zero.¹⁶

Proof of Claim 1.

Case 1. Suppose the aspiration levels \mathbf{d} are such that $d_i + d_j < \alpha_{ij}$ for some i, j .

Case 1a. i and j are not matched with each other.

With positive probability, either i or j is activated and i and j become matched. The new aspiration levels are set equal to the new payoffs. Thus the sum of the aspiration levels is equal to the match's value α_{ij} .

Case 1b. i and j are matched with each other.

¹⁶Note that this claim describes an absorbing state in the core. It may well be that the core is reached while a single's aspiration level is more than zero. The latter state, however, is transient and will converge to the corresponding absorbing state.

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

Therefore, there is a positive probability path along which \mathbf{d} increases monotonically until $d_i + d_j \geq \alpha_{ij}$ for all i, j .

Case 2. Suppose the aspiration levels \mathbf{d} are such that $d_i + d_j \geq \alpha_{ij}$ for all i, j .

We can suppose that there exists a single agent i with $d_i > 0$ and $d_i + d_j > \alpha_{ij}$ for all j , else we are done. With positive probability, i is activated. Since no profitable match exists, he lowers his aspiration level by δ . In this manner, a suitable path can be constructed along which \mathbf{d} decreases monotonically until the aspiration levels are good. Note that at the end of such a path, the assignment does not need to be optimal and not every agent with a positive aspiration level needs to be matched. (See the period- t example in the preceding section.)

□

Proof of Claim 2.

Suppose that the state (\mathbf{A}, \mathbf{d}) satisfies Claim 1 (\mathbf{d} is good) and that some single exists whose aspiration level is positive. (If no such single exists, the assignment is optimal and we have reached a core state.) Starting at any such state, we show that, within a bounded number of periods and with positive probability (bounded below), one of the following holds:

The aspiration levels are good, the number of single agents with positive aspiration level decreases, and the sum of the aspiration levels remains constant. (15)

The aspiration levels are good, the sum of the aspiration levels decreases by $\delta > 0$, and the number of single agents with a positive aspiration level does not increase. (16)

In general, say an edge is *tight* if $d_i + d_j = \alpha_{ij}$ and *loose* if $d_i + d_j = \alpha_{ij} - \delta$. Define a *maximal alternating path* P to be a maximal-length path that starts at a single player with positive aspiration level, and that alternates between unmatched tight edges and matched tight edges. Note that, for every single with a positive aspiration level, at least one maximal alternating path exists. Figure 3 (left panel) illustrates a maximal alternating path starting at f_1 . Unmatched tight edges are indicated by dashed lines, matched tight edges by solid lines and loose edges by dotted lines.

Without loss of generality, let f_1 be a single firm with positive aspiration level.

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

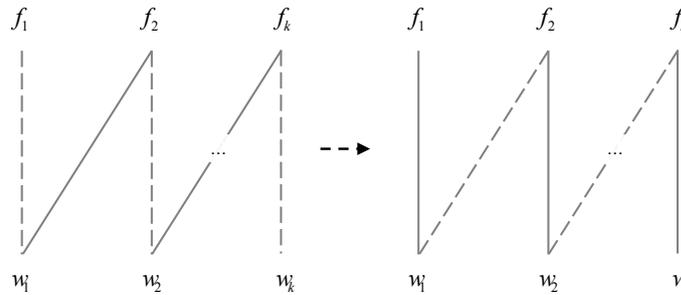
Case 1a. All firms on the path have a positive aspiration level.

We shall demonstrate a sequence of adjustments leading to a state as in (15).

Let $P = (f_1, w_1, f_2, w_2, \dots, w_{k-1}, f_k, w_k)$. Note that, since the path is maximal and of odd length, w_k must be single. With positive probability, f_1 is activated. Since no profitable match exists, he lowers his aspiration level by δ . With positive probability, f_1 is activated again next period, he snags w_1 and with probability 0.5 he receives the residual δ . At this point the aspiration levels are unchanged but f_2 is now single. With positive probability, f_2 is activated. Since no profitable match exists, he lowers his aspiration level by δ . With positive probability, f_2 is activated again next period, he snags w_2 and with probability 0.5 he receives the residual δ . Within a finite number of periods a state is reached where all players on P are matched and the aspiration levels are as before. (Note that f_k is matched with w_k without a previous reduction by f_k since w_k is single and thus their bids are profitable.)

In summary, the number of matched agents has increased by two and the number of single agents with positive aspiration level has decreased by at least one. The aspiration levels did not change, hence they are still good. (See Figure 3 for an illustration.)

Figure 3: Transition diagram for Case 1a.



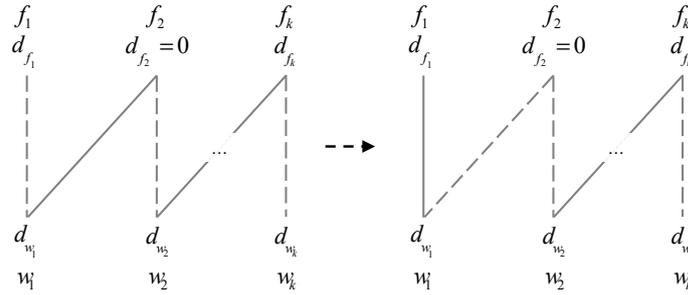
Case 1b. At least one firm on the path has aspiration level zero.

We shall demonstrate a sequence of adjustments leading to a state as in (15).

Let $P = (f_1, w_1, f_2, w_2, \dots, w_{k-1}, f_k, w_k)$. There exists a firm $f_i \in P$ with current aspiration level zero (f_2 in the illustration), hence no further reduction by f_i can occur. (If multiple firms on P have aspiration level zero, let f_i be the first such firm on the path.) Apply the same sequence of transitions as in Case 1a up to firm f_i . At the end of this sequence the aspiration levels are as before. Once f_{i-1} snags w_{i-1} , f_i becomes single and his aspiration level is still zero.

In summary, the number of single agents with a positive aspiration level has decreased by one because f_1 is no longer single and the new single agent f_i has aspiration level zero. The aspiration levels did not change, hence they are still good. (See Figure 4 for an illustration.)

Figure 4: Transition diagram for Case 1b.



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

Case 2a. All firms on the paths have a positive aspiration level.

We shall demonstrate a sequence of adjustments leading to a state as in (16).

With positive probability f_1 is activated. Since no profitable match exists, he lowers his aspiration level by δ . Hence, all previously tight edges starting at f_1 are now loose.

We shall describe a sequence of transitions under which a given loose edge is eliminated (by making it tight again), the matching does not change and the sum of aspiration levels remains fixed. Consider a loose edge between a firm, say f'_1 , and a worker, say w'_1 . Since all maximal alternating paths starting at f_1 are of even length, the worker has to be matched to a firm, say f'_2 . With positive probability w'_1 is activated, snags f'_1 , and with probability 0.5 f'_1 receives the residual δ . (Such a transition occurs with strictly positive probability whether or not f'_1 is matched because aspiration levels are strictly below the match value of (w'_1, f'_1) .) Note that f'_2 and possibly f'_1 's previous partner, say w''_1 , are now single. With positive probability f'_2 is activated. Since no profitable match exists, he lowers his aspiration level by δ . (This occurs because all firms on the maximal alternating paths starting at f_1 have an aspiration level at least δ .) With positive probability, f'_2 is activated again, snags w'_1 , and with probability 0.5 w'_1 receives the residual δ . Finally, with positive probability f'_1 is activated. Since no profitable match exists, he lowers his aspiration level by δ . If previously matched, f'_1 is activated again in the next period and matches with w''_1 . At the end of this sequence the matching is the same as at the beginning. Moreover, w'_1 's aspiration level went up by δ while f'_2 's aspiration level went down by δ and all other aspiration levels stayed the same. The originally loose edge between f'_1 and w'_1 is now tight.

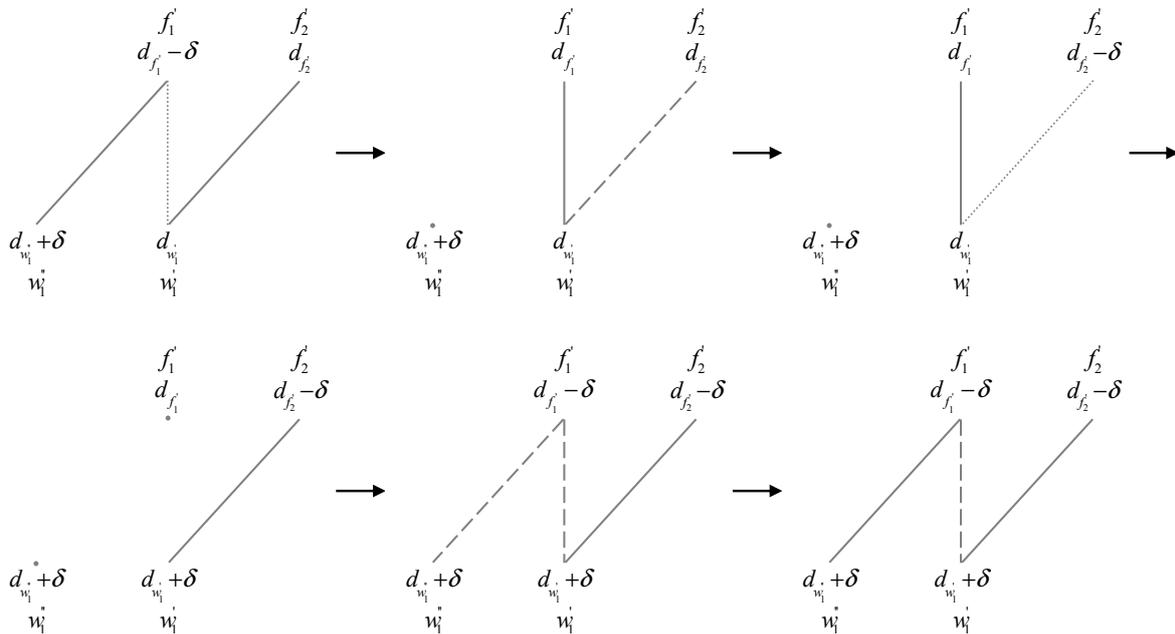
We iterate the latter construction for $f'_1 = f_1$ until all loose edges at f'_1 have been eliminated. However, given f'_2 's reduction by δ there may be new loose edges connecting f'_2 to workers. In this case we repeat the preceding construction for f'_2 until all of the loose edges at f'_2 have been eliminated. If any agents still exist with loose edges we repeat the construction again. This iteration eventually terminates given the following observation. Any worker on a maximal alternating path who previously increased his aspiration level cannot still be connected to a firm by a loose edge. Similarly, any firm that previously reduced its aspiration level cannot now be matched to a worker with a loose edge because such a worker increased his aspiration level. Therefore the preceding

construction involves any given firm (or worker) at most once. It follows that, in a finite number of periods, all firms on maximal alternating paths starting at f_1 have reduced their aspiration level by δ and all workers have increased their aspiration level by δ .

In summary, the number of aspiration level reductions outnumbers the number of aspiration level increases by one (namely by the firm f_1), hence the sum of the aspiration levels has decreased. The number of single agents with a positive aspiration level has not increased. Moreover the aspiration levels are still good. (See Figure 5 for an illustration.)

Note that the δ -reductions may lead to new tight edges, resulting in new maximal alternating paths of odd or even lengths.

Figure 5: Transition diagram for Case 2a.



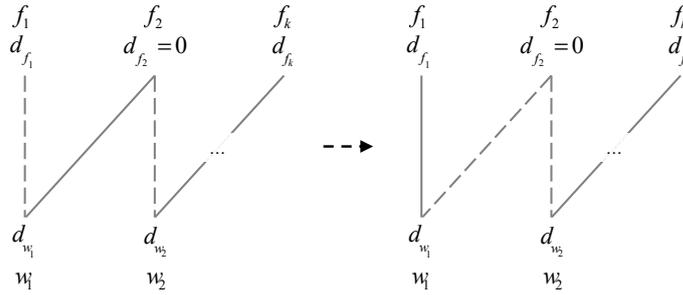
Case 2b. At least one firm on the path has aspiration level zero.

We shall demonstrate a sequence of adjustments leading to a state as in (15).

Let $P = (f_1, w_1, f_2, w_2, \dots, w_{k-1}, f_k)$. There exists a firm $f_i \in P$ with current aspiration level zero (f_2 in the illustration), hence no further reduction by f_i can occur. (If multiple firms on P have aspiration level zero, let f_i be the first such firm on the path.) With positive probability f_1 is activated. Since no profitable match exists, he lowers his aspiration level by δ . With positive probability, f_1 is activated again next period, he snags w_1 and with probability 0.5 he receives the residual δ . Now f_2 is single. With positive probability f_2 is activated, lowers, snags w_2 , and so forth. This sequence continues until f_i is reached, who is now single with aspiration level zero.

In summary, the number of single agents with a positive aspiration level has decreased. The aspiration levels did not change, hence they are still good. (See Figure 6 for an illustration.)

Figure 6: Transition diagram for Case 2b.



Let us summarize the argument. Starting in a state $[\mathbf{A}, \mathbf{d}]$ with good aspiration levels \mathbf{d} , we successively (if any exist) eliminate the odd paths starting at firms/workers followed by the even paths starting at firms/workers, while maintaining good aspiration levels. This process must come to an end because at each iteration either the sum of aspiration levels decreases by δ and the number of single agents with positive aspiration levels stays fixed, or the sum of aspiration levels stays fixed and the number of single agents with positive aspiration levels decreases. Finally, single agents (with aspiration level zero) successively match at aspiration level zero until all agents on the smaller side of the market are matched. The resulting state must be in the core and is absorbing because single agents cannot reduce their aspiration level further and no new matches can be formed. Since an aspiration level constitutes a lower bound on a player's bids we can conclude that the process Z^t is absorbed into the core in finite time with probability 1. \square

We have so far shown that the core is absorbed when we operate on the $\delta \cdot \mathbb{N}_0$ grid. The following corollary states that the result also holds in a continuous space in which our price rounding assumption vanishes.

Corollary 2. *Let $p_{ij}^+, q_{ij}^- \in \mathbb{R}$ and let $X_i^t, P_{ij}^t, Q_{ij}^t$ be independent random variables taking values in \mathbb{R}^+ such that the expectation of X_i^t is positive and there exists a constant c such that for all $\epsilon > 0$, $\mathbb{P}[P_{ij}^t < \epsilon] > c > 0$, and $\mathbb{P}[Q_{ij}^t < \epsilon] > c > 0$.*

Define the assignment game $G(v, N)$ as above. From any initial state $[\mathbf{A}^0, \mathbf{d}^0] \in \Omega$, the process is absorbed into the core in finite time with probability 1.

Proof. The conditions of the corollary are satisfied in the earlier setup for any $\delta > 0$. Hence for $\delta \rightarrow 0$ absorption into the core follows. To see that absorption occurs in finite time, note that δ only influences the convergence time when players are single and reduce their aspiration level. By (12) the latter reductions are bounded away from zero and the result follows. \square

6. Conclusion

In this paper we have shown that agents in large decentralized matching markets can learn to play stable and efficient outcomes through a trial-and-error learning process. We assume that the agents have no information about the distribution of others' preferences, their past actions and payoffs, or about the value of different matches. Nevertheless the learning process leads to the core with probability one. The proof uses integer programming arguments (Kuhn [1955], Balinski [1965]), but the Matchmaker does not “solve” an integer programming problem. Rather, a path into the core is discovered in finite time by a random sequence of adjustments by the agents.

A crucial feature of our model is that the Matchmaker has no knowledge of match values, hence standard matching procedures cannot be used. In fact, the role of the Matchmaker can be eliminated entirely, and the process can be interpreted as a purely evolutionary process with no third party at all. As before, let agents be activated by independent Poisson clocks. Suppose that an active agent randomly encounters one agent from the other side of the market drawn from a distribution with full support. The two players enter a new match with positive probability if their match is potentially profitable, which they can see from their current bids and offers. If the two players are already matched with each other, they remain so. If both are single, they agree to be matched if their bid and offer cross. If at least one agent is matched (but with someone else), they agree to be matched if their bid and offer *strictly* cross. This is essentially the same process as the one described above, and the same proof shows that it leads to the core in finite time with probability one.

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. L. Balinski, “Integer Programming: Methods, Uses, Computations”, *Management Science* 12, 253-313, 1965.
- M. L. 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.
- 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 multi-player 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", *Journal of Economic Theory* 147(1), 364-381, 2012.

- B. S. R. Pradelski & H. P. Young, “Learning Efficient Nash Equilibria in Distributed Systems”, *Games and Economic Behavior* 75, 882-897, 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 & E. Peranson, “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design”, *The American Economic Review* 89, 756-757, 1999.
- A. E. Roth, T. Sönmez & U. 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.
- R. Shimer, “Mismatch”, *The American Economic Review* 97, 1074-1101, 2007.
- R. Shimer, “The Probability of Finding a Job”, *The American Economic Review* 98 (Papers and Proceedings), 268-273, 2008.
- 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.