ISSN 1471-0498

DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

DECENTRALIZED DYNAMICS AND FAST CONVERGENCE IN THE ASSIGNMENT GAME

Bary S.R. Pradelski

Number 700
February 2014
Revised November 2014

Manor Road Building, Manor Road, Oxford OX1 3UQ

Decentralized Dynamics and Fast Convergence in the Assignment Game*

Bary S. R. Pradelski†

February 27, 2014
This version: September 30, 2014

Abstract

We study decentralized learning dynamics for the classic assignment game with transferable utility. At random points in time firms and workers match, break up, and re-match in the search for better opportunities. We propose a simple learning process in which players have no knowledge about other players' payoffs or actions and they update their behavior in a myopic fashion. Behavior fluctuates according to a random variable that reflects current market conditions: sometimes the firms exhibit greater price stickiness than the workers, and at other times the reverse holds. We show that this stochastic learning process *converges in polynomial time* to the core. While convergence to the core is known for some types of decentralized dynamics this paper is the first to prove fast convergence, a crucial feature from a practical standpoint. The proof relies on novel results for random walks on graphs, and more generally suggests a fruitful connection between the theory of random walks and matching theory.

JEL classifications: C71, C73, C78, D83

Keywords: assignment games, core, evolutionary game theory, matching markets, convergence time, random walks

*I thank Pierre Tarrès, Peyton Young, and Heinrich Nax for their guidance and support throughout this research. Further, I thank Yakov Babichenko, Doyne Farmer, Sergiu Hart, Thomas Norman, and three anonymous referees for suggesting a number of improvements to earlier versions. I am also grateful for comments by participants at seminars at the Hebrew University (Israel), the Technion (Israel) and the University of Oxford (UK). I would like to acknowledge the generous support of the Oxford-Man Institute of Quantitative Finance, the Office of Naval Research grant # N00014-09-1-0751, and the Air Force Office of Scientific Research grant # FA9550-09-1-0538.

†University of Oxford, Eagle House, Walton Well Road, Oxford OX2 6ED, UK. Email: bary.pradelski@oxford-man.ox.ac.uk

1 Introduction

Labor markets have become more fluid – workers change jobs more frequently and market participation increases. The sheer size of the market often hinders firms and workers from learning about the preferences, past actions, and utilities of other market participants. They may often only be able to observe current market prices along with the general market sentiment driven by macro-economic conditions, such as a crisis or a boom.

A crucial feature of a behavioral dynamic is the expected time it takes to converge. A dynamic is said to exhibit *fast convergence* if the expected time to convergence is polynomial in the number of players. The main contribution of this paper is to propose a dynamic for the assignment game with transferable utility that exhibits fast convergence and adheres to characteristics assumed for completely decentralized markets in the related literature (Chen et al. 2011, Nax and Pradelski 2012, Biró et al. 2012, Klaus and Payot 2013). In particular, the dynamic has the following four properties:

Behavioral decentralization. Firms and workers bilaterally form and break matches, and the order of their interactions is random and asynchronous.

Information decentralization. A player has no information about the value other players attach to him. For example, a worker may be able to know the salary of a job before taking it on, but he will not know the utility the firm derives from his employment.

Matching decentralization. The behavior of agents differs when matched or single. Current matchings influence future actions.

Bounded rationality. Players adjust their behavior using simple adaptive rules that are based on limited information. This is in the spirit of much of the learning literature on *completely uncoupled* dynamics.[1]

Several recent studies consider the convergence and convergence time for a family of decentralized dynamics. On the one hand, Chen et al. (2011), Nax and Pradelski (2012), Biró et al. (2012), Klaus and Payot (2013), and Nax et al. (2013) show that certain dynamics converge to the core and satisfy the outlined criteria of decentralization and bounded rationality. They do not attempt to show fast convergence. On the other hand, Bayati et al. (2005, 2008), Azar et al. (2009), Celis et al. (2010), and Bayati et al. (2014) study the convergence time. Fast convergence is only obtained for dynamics where beliefs and bids are updated simultaneously, in particular violating the assumption of behavioral decentralization introduced above. In fact, Azar et al. (2009) show that their dynamic may take exponentially long to converge if behavioral decentralization holds.

From a technical perspective, we develop new results for random walks on graphs which may be of use more generally. In particular, to our knowledge, we are the first to use random walk arguments in the context of assignment and matching markets. In contrast to earlier applications of random walks in economics where the moving particle is usually information a player receives or a player himself who moves on a graph we propose an

[1]In an *uncoupled* dynamic players don't know other players' payoffs (Hart and Mas-Colell 2003). In a *completely uncoupled dynamic* players don't know other players' current or past behavior, or their payoffs (Foster and Young 2006).

auxiliary connection. By creating an appropriate duality we can describe the dynamic process governing which players are currently matched and which are single by a family of random walks on different graphs.

The next section discusses the related literature. In section 3 we define the underlying model for bilateral markets with transferable utility and introduce our dynamic model. We also provide an example. Section 4 states our main results and discusses its assumptions. Section 5 provides an intuitive outline of the proof, a discussion of the random walk argument and the proof. Section 6 concludes.

2 Related literature

Matching theory is a flourishing area of study in cooperative game theory and mechanism design.[2] The most common example for a model with transferable utility (TU) is the labor market. Shapley and Shubik (1972) introduce the *assignment game* where two sides of the market, firms and workers, form bilateral coalitions and share the surplus (transferable utility) according to a sharing rule.[3] They show that there always exists an assignment such that the total sum of utilities is maximized and the division of surplus is robust to coalitional deviations. The set of states with these properties is called the *core* (see section 3.1 for a formal introduction). One of the first algorithms algorithm for the assignment game (Kuhn 1955) has time complexity of $O(N^4)$, where N is the number of players. Edmonds and Karp (1972) give a tighter bound of $O(N^3)$ which – to our knowledge – is the best-known bound for a centralized algorithm.[4]

A second large strand in the matching literature considers matching with non-transferable utility (NTU) (see Gale and Shapley 1962 and Roth and Sotomayor 1990 for a survey). *Transferable utility* describes the existence of a medium of exchange (such as money) and assumes that players' utilities are linear in that medium. The exchange is often called a *side payment*. Under *non-transferable utility* there is no such medium of exchange. Thus, in general, transferable utility models assume *cardinal* utilities, whereas non-transferable utilities assume *ordinal* utilities. For decentralized dynamics in NTU matching markets there exist random adjustment processes leading to stable outcomes in the two-sided matching market (Roth and Vande Vate 1990).[5] The dynamic considered by Roth and Vande Vate (1990) constitutes the iterative resolving of *blocking pairs*.[6] Ackermann et al.

[2]See, for example, Roth and Sotomayor (1990) and Peleg and Sudholter (2007) for introductions.

[3]Remarkably, Boehm-Bawerk (1923) previously studied a related model for particular examples.

[4]The structure of the core has since been widely studied. Crawford and Knoer (1981) and Kelso and Crawford (1982) tailor the model to a labor market and characterize the core allocations that are likely to emerge under different market institutions (such as which side of the market makes offers). For subsequent literature along this lines, see Demange et al. (1986), Shimer (2007, 2008), and Elliott (2010, 2011). This links with a growing literature on dynamic non-cooperative behavior for cooperative game solutions (see, for example, Agastya 1997, 1999, Arnold and Schwalbe 2002, Rozen 2013, Newton 2010, 2012, Sawa 2011).

[5]See Chung (2000), Diamantoudi et al. (2004), Kojima and Uenver (2008), Klaus and Klijn (2007) for subsequent results on more general families of games.

[6]Players from each side of the market constitute a *blocking pair* if each prefers the other over his current match (Gale and Shapley 1962).

(2011) recently showed by example that the *better response dynamic* by Roth and Vande Vate (1990) may take exponentially long to converge. They also extend their result to the *best response dynamic*.

2.1 Random processes to the core and aspiration adaptation

An open question has been to prove corresponding results for TU games. Chen et al. (2011), Biró et al. (2012), Nax and Pradelski (2012), and Klaus and Payot (2013) show that a random updating rule leads to core outcomes in finite time with probability one from any initial state.[7]

In joint work (Nax and Pradelski 2012) we propose an evolutionary, *completely uncoupled*[8] dynamic. Completely uncoupled dynamics have two advantages: First, they provide a more realistic model of economic behavior in situations where little information is available to the players or they may not even be aware to be part of a 'game'. Second, given the absence of information about other players' actions or payoffs, completely uncoupled dynamics allow to avoid the treatment of questions arising from strategic agents. This is justified by the mere fact that there is no useful information to base strategic play on. The behavioral nature of the dynamic in Nax and Pradelski (2012) follows aspiration adjustment theory (Sauermann and Selten 1962, Selten 1998) and related bargaining experiments on directional and reinforcement learning (e.g., Tietz and Weber 1972, Roth and Erev 1995): players occasionally experiment with new actions. If the new match promises to be more profitable than their current aspiration they engage in the new match. We shall detail this in section 3.

We shall now discuss two separate strands of the literature that have considered the same or closely related underlying market models.

2.2 Belief propagation

Graph matching has been widely studied in mathematics (see, for example, Edmonds 1965, Gibbons 1985). Although there are numerous centralized algorithms to find optimal matchings, decentralized algorithms have only sparked interest more recently. Pearl (1982) and Kim and Pearl (1983) introduce *belief propagation*, a family of message passing algorithms for computing optimal matchings and allocations on graphs. Belief propagation is a decentralized algorithm where only neighboring players (nodes) communicate with each other, but it is deterministic (there is a fixed order of updating of each agent). Further, there is no stopping rule specified and matchings do not influence the dynamic behavior. In a game with a unique optimal matching the algorithm will eventually converge to payoffs enforcing the optimal matching (as the only stable matching). Although

[7]Note that Klaus and Payot (2013) for assignment games and Biró et al. (2012) for roommate problems show almost sure convergence under a particular price rule. But to implement their dynamics a solution needs to be known beforehand.

[8]In an *uncoupled* dynamic players don't know other players' payoffs (Hart and Mas-Colell 2003). In a *completely uncoupled* dynamic players don't know other players' current or past behavior, or their payoffs (Foster and Young 2006).

convergence has only been shown recently for games on bipartite graphs with a unique optimal matching[9] (Bayati et al. 2005, 2008) belief propagation has found wide applications in fields such as artificial intelligence and statistics (see, for example, Pearl 1988 and Richardson and Urbanke 2001). In addition Bayati et al. (2005) show polynomial time convergence.[10]

2.3 Bargaining

Independently, in sociology *exchange network theory* studies "the processes through which network structures affect power distributions, power exercise, and benefits gained in exchange"(Walker 2007, see also for a survey of the field). Exchange networks are closely related to and often coincide with the assignment game (see, for example, Cook and Yamagishi 1992, Braun and Gautschi 2006). The focus of exchange network theory has been on the influence of network structure on the resulting sharing of outputs or power distribution. Many studies in exchange networks focus on particular network structures such as trees or certain line segments in experimental setups. Bienenstock and Bonacich (1992) introduced the game theoretic concept of the core to the sociology literature, arguing that while being in accordance with experimental work, its theoretical richness enhances predictive power in a multiplicity of games. Notably, the study of exchange networks was recently brought into the area of computer science and mathematics through the work of Kleinberg and Tardos (2008). Subsequently convergence has been shown for a particular decentralized dynamic, namely *bargaining* with random encounters where some global information is available to all players (Azar et al. 2009). They show that the algorithm may take exponentially long to converge. The deterministic version of the latter algorithm is shown to converge in polynomial time (see Celis et al. 2010, Bayati et al. 2014) for *one-sided markets*.[11] Similar to belief propagation there is no stopping rule specified and matchings do not influence the dynamic behavior. In a game with a unique optimal matching the algorithm will eventually converge to payoffs enforcing the optimal matching (as the only stable matching).

2.4 A brief discussion

The examples of decentralized dynamic models in the previous sections point to the difficulty to prove results regarding the rate of convergence for random processes of this kind. The examples can be regarded as possible approximations of human behavior. For instance, bargaining is aiming to provide "a dynamical description of how players can find such a [balanced] outcome"(Bayati et al. 2014).

[9]Note that generically the optimal matching is unique. In particular this is the case if the weights of the edges are independent, continuous random variables. Then, with probability 1, the optimal matching is unique.

[10]Salez and Shah (2009) show a tighter bound for the belief propagation algorithm for a class of games where the distribution from which the edge weights are drawn is not too heavy-tailed.

[11]One-sided markets is the general class of one-to-one matchings which includes assignment games as a sub-class. One-sided markets are also known as *roommate markets*.

We shall argue that from a behavioral perspective the present models which exhibit fast convergence, namely belief propagation and bargaining, are not completely satisfactory to model myopic human behavior in a decentralized environment as outlined in the introduction.

First, the order of interactions is fixed, suggesting some central authority. This violates our assumption of *behavioral decentralization*. Second, counterfactuals and private information are observable, that is, the match values (w_{ij}) are known to the players. As outlined earlier it is dubitable that players have knowledge of other players' utility functions and thus the match values (*information decentralization*). Thirdly, matchings are not specified and do not influence behavior. The matching is considered as induced, once unique one-to-one correspondences of best matches between two players exist, suggesting a central authority announcing that point in time. This is in stark contrast to *matching decentralization*. Finally, if a player knows the best alternatives of other players and is able to compute his best alternatives, won't he attempt to act strategically given this information? For instance, he could always send messages over-claiming his alternatives and could thus 'game' the system. This question remains unanswered and it is assumed that players are truthful and report and share according to a fixed rule. Given the high level of information and computational capacity of players a discussion of *bounded rationality* seems necessary.

The above criticisms are addressed by the models of Chen et al. (2011), Biró et al. (2012), Nax and Pradelski (2012), and Klaus and Payot (2013) as discussed in section 2.1. But in order to argue that a behavioral dynamic indeed explains a decentralized 'path to stability' we need to ensure that it exhibits *fast convergence* (that is, the expected time to convergence is polynomial in the number of players). In addition, from a mechanism design standpoint fast convergence is crucial. This question is not addressed by the aforementioned papers.

To prove convergence to the core the proofs in the literature rely on showing the existence of a particular path to the core (Chen et al. 2011, Biró et al. 2012, Nax and Pradelski 2012, and Klaus and Payot 2013). On the other hand, the proof techniques employed to show fast convergence are diverse but all rely on the deterministic nature of the studied dynamics. As discussed above this is achieved by simultaneous updating of all players (Bayati et al. 2005, 2008, Azar et al. 2009, Celis et al. 2010, and Bayati et al. 2014). Thus none of the techniques can not be readily used to consider the speed of convergence for the problem at hand.

3 Model

Our model is in the spirit of Nax and Pradelski (2012) as discussed in section 2.1: matched agents occasionally experiment with higher aspiration levels, while single agents, in the hope of attracting partners, lower their aspiration levels. The main difference between the model in Nax and Pradelski (2012) and our model is a market sentiment. We show that the market sentiment is necessary in order to achieve fast convergence. In the next subsection we shall introduce the well-known assignment game and the dynamic components of our model. Subsection 3.2 explains the dynamic play and how step-by-step transitions work.

3.1 The assignment game

Let $G = (V, W)$ be a weighted, bipartite graph with vertex set $V = \mathcal{F} \cup \mathcal{W}$ and edge weights (or *match values*) $W \ni w_{ij} > 0$, for $i, j \in V$. (It is assumed that if $w_{ij} > 0$ there exists an edge between i and j. Otherwise there is no edge.) Let $w^* = \max_{i,j} w_{ij}$. We shall refer to the nodes $i \in V$ as players or agents and in particular to $i \in \mathcal{F}$ as firms and to $j \in \mathcal{W}$ as workers. We assume, that $|\mathcal{F}| = |\mathcal{W}| = N$. Note that this can always be achieved by adding 'dummy' agents to the smaller side of the market without any edges. A player can be *matched* to at most one player at a given point in time. A player currently not matched to another player is called *single*.

Remark 1. *The edge weights w_{ij} are the amounts two matched players can agree to transfer between each other. That is, suppose that firm i is willing to pay at most $r_i^+(j)$ when matching with worker j. Thus $r_i^+(j)$ describes the* reservation value *above which firm i prefers to remain single over matching with j. Similarly, let $r_j^-(i)$ be worker j's minimal willingness to accept when matching with firm i. Then any price between $r_i^+(j)$ and $r_j^-(i)$ is individually rational and thus the actual amount under negotiation is $r_i^+(j) - r_j^-(i) = w_{ij}$. It follows that w_{ij} is not known to the players since the reservation values are private information. In our model w_{ij} is a hidden variable which we shall use for ease of exposition but which – at no point in time – is known to any player.*

Since we will be interested in dynamic processes we introduce time steps $t = 0, 1, 2, 3, \ldots$.

Matching. The *matching* at time t is described by $M^t = (m_{ij}^t)_{i\in\mathcal{F}, j\in\mathcal{W}} \in \{0,1\}^{\mathcal{F}\times\mathcal{W}}$ such that each row and each column of M^t sums to at most 1.

The *payoff* to player i at the end of step t is given by ϕ_i^t. We assume that matches are *efficient*, that is, if (i, j) are matched in step t, then $\phi_i^t + \phi_j^t = w_{ij}$. Let $\phi_i^t = 0$ if i is single and $\Phi^t = (\phi_i^t)_{i\in V}$.

Remark 2. *The payoff to player i is the share of the transferable amount he receives from matching with some j. Note that the payoff is private information, since, given i and j agree on a price π_{ij}^t, the payoff to firm i is $r_i^+(j) - \pi_{ij}^t$ and to worker j is $\pi_{ij}^t - r_j^-(i)$.*

The *matching market* is described by $[V = \mathcal{F} \cup \mathcal{W}, W, M^t, \Phi^t]$:

- $V = \mathcal{F} \cup \mathcal{W}$ is the set of firms and workers,
- $W = (w_{ij})_{i\in\mathcal{F}, j\in\mathcal{W}} > 0$ is the matrix of match values,
- $M^t = (m_{ij}^t)_{i\in\mathcal{F}, j\in\mathcal{W}} \in \{0,1\}^{\mathcal{F}\times\mathcal{W}}$ is the assignment matrix with 0/1 entires, and row/column sums at most one,
- $\Phi^t = (\phi_i^t)_{i\in V}$ is the payoff vector.

Then, given the matching market, the *assignment game* $G(\nu, V)$ is defined as follows. Let $\nu : S \subseteq V \to \mathbb{R}$ such that

- $\nu(i) = \nu(\emptyset) = 0 \; \forall$ singletons $i \in V$
- $\nu(S) = w_{ij} \; \forall S = (i, j)$ with $i \in \mathcal{F}$ and $j \in \mathcal{W}$
- $\nu(S) = \max_{P \in \text{perm}(S)}\{\nu(i_{P(1)}, j_{P(1)}) + ... + \nu(i_{P(k)}, j_{P(k)})\} \; \forall S = (i_1, j_1, i_2, j_2, ..., i_{k^{\mathcal{F}}}, j_{k^{\mathcal{W}}}) \subseteq V$ and $k = \min\{k^{\mathcal{F}}, k^{\mathcal{W}}\}$

The number $\nu(V)$ specifies the value of an optimal assignment (as defined below). We shall now turn to define desirable properties of outcomes:

- A matching M is *optimal* if $\sum_{i,j \in V} m_{ij} \cdot w_{ij}$ is maximal.
- A payoff vector Φ is *pairwise stable* if $\forall i, j \in V : \phi_i + \phi_j \geq w_{ij}$.

Let a matching M together with the payoffs Φ be in the *core*, if M is optimal and Φ is pairwise stable.

So far we introduced the assignment game without any additional components. We shall now introduce dynamic components particular to our model.

Aspiration levels. At the end of any step $t = 0, 1, 2, \ldots$, a player, $i \in V$ has an *aspiration level*, $d_i^t \geq 0$, which determines the minimal payoff at which he is willing to be matched (if the market sentiment, defined below, is not against him). Let $\mathbf{d}^t = (d_i^t)_{i \in V}$.

Feasibility. A pair of aspiration levels (d_i^t, d_j^t) is *feasible* if $d_i^t + d_j^t < w_{ij}^t$ or if $d_i^t + d_j^t = w_{ij}^t$ and one of the players is single.[12] We shall say that i and j are *feasible* in period t if (d_i^t, d_j^t) is feasible.

Market sentiment. There is a *market sentiment* which is described by a random variable and acts as a correlation device: The market sentiment is negative in periods of economic decline where workers are at risk of losing their jobs. It is positive in periods of economic growth where firms risk losing their employees to competitors. Thus, if the market sentiment is negative, workers are asking for a bit less, if the market sentiment is positive firms are offering a bit more. Denote the market sentiment by $\Sigma^t \in \{\ominus, \oplus\}$ We define $\mathcal{S}^t \in \{\mathcal{F}, \mathcal{W}\}$ to be the set of players against whom the market sentiment has currently turned and who are thus under threat. That is, if $\Sigma^t = \ominus$, then $\mathcal{S}^t = \mathcal{W}$, and if $\Sigma^t = \oplus$, then $\mathcal{S}^t = \mathcal{F}$.

We shall see in the description of the dynamic how the market sentiment influences the transitions in detail.

State. The state at the end of period t is given by the matching, the aspiration levels, and the market sentiment: $[M^0, \mathbf{d}^0, \Sigma^0]$.

Finally we assume that there is a minimal unit of payoff $\delta \in \mathbb{R}^+$, and all edge-weights, aspiration levels, etc. are multiples of δ.

[12] Note that this definition assumes a tie-breaking rule: A player prefers to be matched over being single.

3.2 Dynamic play

Agents are *activated* by independent Poisson arrival processes.[13] Once a player is activated the next discrete time step t starts. The active agent scans the other side of the market, looking for a feasible match and picks one at random.[14,15] (Note that this is similar to the idea of resolving blocking pairs.) The two agents match and the additional surplus is split at random, that is, the price is set at random (under some weak conditions). At the end of the step the two players adapt their aspiration levels to equal their new payoffs. If the active player is single at the end of the step he reduces his aspiration level by δ. The state at the end of step t is then uniquely described by the match matrix M^t, the vector of aspiration levels $\mathbf{d}^t$, and the market sentiment Σ^t.

No better alternative. If player i is activated in period $t+1$, say that j' has *no better alternative* (than i) if for all $i' \neq i$ $d^t_{i'} + d^t_{j'} > w_{i'j}$. Otherwise say that j' *has a better alternative.*

Let $[M^0, \mathbf{d}^0, \Sigma^0]$ be any initial state.[16] Suppose player i is activated in period $t+1$ for $t \geq 0$. Let J' be the set of players j' such that (i, j') is feasible.

- If $J' \neq \emptyset$ player i encounters some player $j' \in J'$ at random, matches with him and updates his aspiration level
$$d_i^{t+1} = d_i^t + P_{ij'}^{t+1}, \tag{1}$$
where $P_{ij'}^{t+1}$ is a random variable with full support on:
$$\begin{cases} [\delta, 2\delta, \cdots, w_{ij'} - d_i^t - d_{j'}^t + \delta] & \text{if } j' \in \mathcal{S}^{t+1} \text{ and } j' \text{ has no better alternative,} \\ [0, \delta, 2\delta, \cdots, w_{ij'} - d_i^t - d_{j'}^t] & \text{else,} \end{cases}$$
where the probability of each event is bounded below by some fixed, global constant.

- If $J' = \emptyset$ there is no new match and if i was previously matched he stays with his previous partner and otherwise remains single and updates his aspiration level
$$d_i^{t+1} = \begin{cases} d_i^t & \text{if } \exists j \text{ such that } m_{ij}^t = 1, \\ d_i^t - \delta & \text{else.} \end{cases} \tag{2}$$

If (i, j') have matched during the period the aspiration level of player j' is $d_{j'}^{t+1} = w_{ij'} - d_i^{t+1}$. The matching M^{t+1} is updated accordingly, that is $m_{ij'}^{t+1} = 1$, $m_{kl}^{t+1} = 0$ for all $k \neq i$ with $l = j'$ and for all $l \neq j'$ with $k = i$. Else $M^{t+1} = M^t$. See figure 1 for illustrations.

[13] In order to make statements on the rate of convergence it is important to assume some value on the maximal rate of a player's Poisson clock. Note that, it is actually not necessary that all players have the same rate but just that they are bounded above by some constant.

[14] The only requirement is that he picks the best match with lower bounded probability.

[15] In Nax and Pradelski (2012) we do not assume that the active agent scans the market but just randomly encounters another agent. One can, for example, think of publicly available price lists.

[16] Note that we are interested in the long run behavior of the dynamic and our results are independent of the initial state. Thus any initial matching, aspiration levels and market sentiment may be assumed.

Figure 1: Transition diagram for active, matched (left) and active, single agent (right).

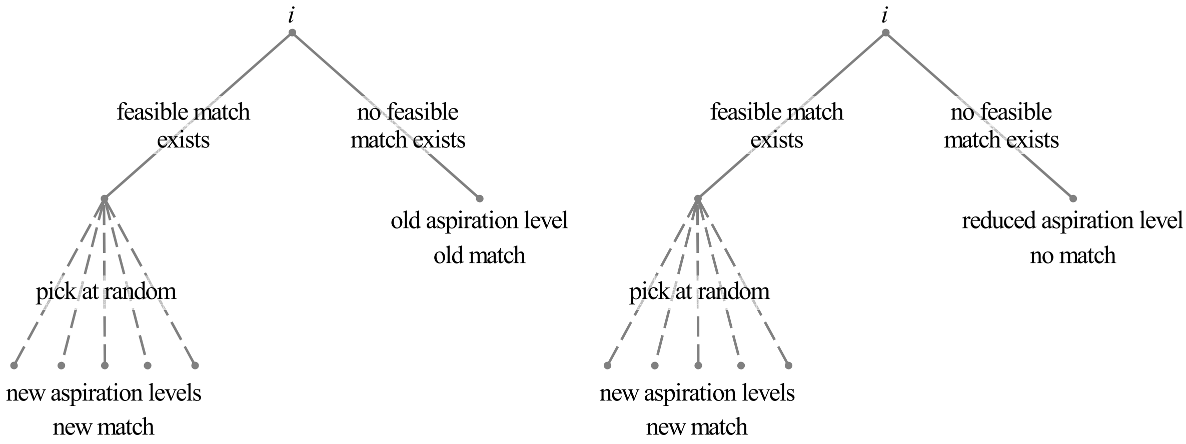

Note that $P_{ij'}^{t+1}$ determines how the additional surplus – taking into account each player's previous aspiration levels – is split. In particular if j' is currently threatened by the market sentiment and has no better alternative the price is set such that in the worst case for j' he is giving up δ of his aspiration and in the best case he takes all the surplus minus δ. On the other hand if j' is currently not threatened by the market sentiment or has a better alternative the price is set such that both players receive a payoff at or above their previous aspiration levels. Finally, the active agent's (i) behavior does not depend on the market sentiment and thus he may only reduce his aspiration level if he is single.

We invite the reader to verify that *behavioral, information, matching decentralization* and *bounded rationality* are all satisfied in above model (compare section 1 for definitions). In order to clarify the information decentralization, consider the following remarks.

Remark 3. *From the standpoint of available information, two players agree on a price but their payoff as well as their aspiration level always remain private information. We defined the price as a random variable within certain bounds (determined by the last period's aspiration levels) for ease of exposition. Given that the payoffs are deducible from the current aspiration levels and the matching, we proceeded to define the algorithm in terms of the aspiration levels which are hidden variables.*

One may argue that players or a central authority may learn (with certainty) other players preferences (and therefore the match values) if they know that players always bid according to their current aspiration levels. This is mitigated by the following remark.

Remark 4. *We could make the following alterations in order to allow any bids such that the prospective payoff would be at least as high as the aspiration level but potentially higher.*[17]

[17] We refrain from stating exact conditions but refer the reader to Nax and Pradelski (2012) who provide conditions similar to the ones needed here.

Random bids above aspiration level ... *The active player submits a random bid vector at or above his aspiration level. That is, he submits an individual bid to each player where a bid is anywhere at or above the aspiration level with positive probability. In particular, the bid vector may be incomplete, that is the active player may decide not to submit a bid to any given player at a given time. All players are confronted with the bid made for them. If a non-active player's aspiration level is such that he would accept the offer he submits a bid such that the match can occur.*

... and downward sticky aspiration level. *An active single reduces his aspiration level only if* all *his bids are* at *his aspiration level and there is no feasible match.*

3.3 Example

Let $V = \mathcal{F} \cup \mathcal{W} = \{i_1, i_2\} \cup \{j_1, j_2\}$ and $w_{11} = 5, w_{12} = w_{21} = 7, w_{22} = 10$. Let $\delta = 1$.

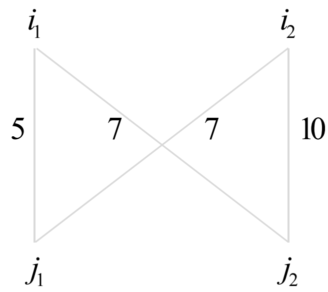

In all figures a gray edge between two players i and j indicates that $w_{ij} > 0$, a black solid edge that they are matched ($m_{ij}^t = 1$), and a black dashed edge that $d_i^t + d_j^t = w_{ij}^t$ and i and j are not matched in period t. We shall call such a dashed edge *tight*. Numbers next to edges indicate match values. The current aspiration levels are shown next to the name of an agent.

One can verify that the optimal matching is unique and such that firm i_1 matches worker j_1 and i_2 matches j_2.

period t: *Current state*

Suppose that, at the end of some period t, (i_2, j_1) are matched, i_1 and j_2 are single, and i_2 and j_2 share a tight edge.

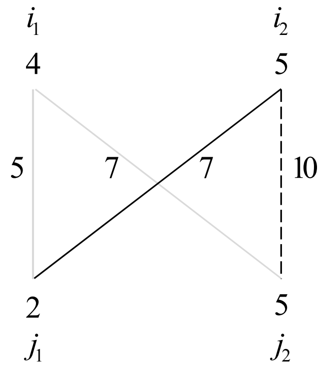

Note that the aspiration levels satisfy $d_i^t + d_j^t \geq w_{ij}$ for all i and j, but the assignment is not optimal and the aspiration levels (and payoffs) are not stable.

period t+1**:** *Activation of single agent* i_1 *and negative market sentiment* $\Sigma^{t+1} = \ominus$

Suppose that i_1 is activated in period $t+1$ and the market sentiment is currently negative. i_1's current aspiration level is too high in order to be feasible with any worker. Hence he remains single and reduces his aspiration level by 1.

i_1 is activated i_1 reduces aspiration level

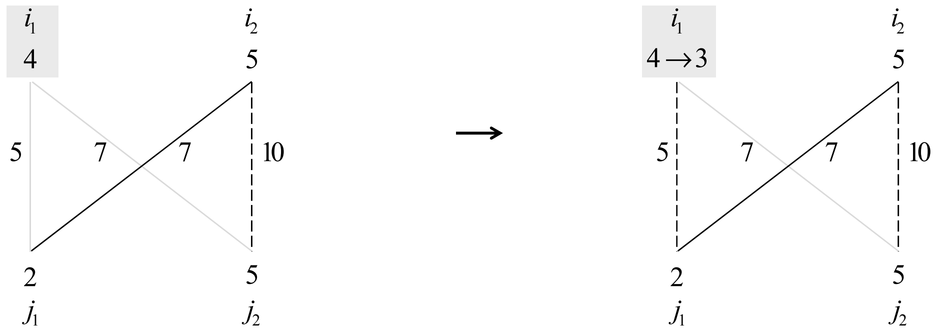

period $t+2$**:** *Activation of matched agent* j_1 *and positive market sentiment* $\Sigma^{t+1} = \oplus$

Suppose that j_1 is activated in period $t+2$ and the market sentiment has changed and it is now positive. (i_1, j_1) is the only feasible pair involving j_1. They thus match and given that i_1 has not better alternative and the market sentiment is currently turned against firms the price, with positive probability, is set such that j_1 can increase his aspiration level by $\delta = 1$.

j_1 is activated

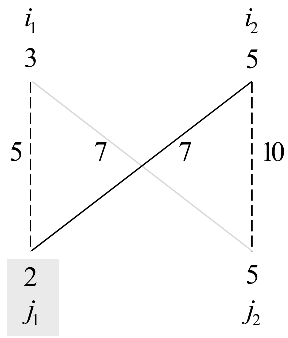

Successful match; j_1 increases aspiration level, i_1 reduces

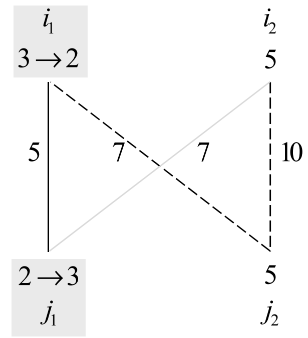

period $t+3$: *Activation of single agent* i_2 *and negative market sentiment* $\Sigma^{t+1} = \ominus$

Suppose that i_2 is activated in period $t+2$ and the market sentiment has changed again is now negative. (i_2, j_2) is the only feasible pair involving i_2. They thus match and although the market sentiment is turned against workers, j_2 has a better alternative (namely i_1) and thus is not accepting less than his aspiration level. They are matching at their current aspiration levels.

i_2 is activated

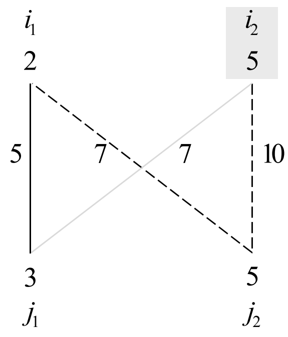

Successful match with j_2 at current aspiration levels

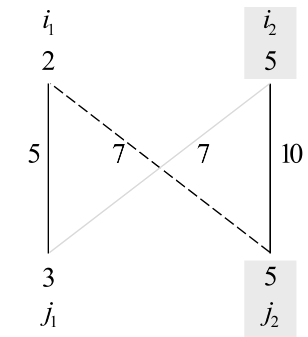

Note that the resulting state is in the core. We invite the reader to convince himself that this state is absorbing.

4 Results

Since the model lives on the discrete grid $\delta \cdot \mathbb{Z}$, the term *generic* needs to be treated with care. We shall give an intuitive definition which will be sufficient for our results to hold.

Definition 5. *A weight matrix* W *is* discrete generic *if there is no sequence of firms and workers* $(1, 1', 2, 2', ..., k, k')$, $i \in \mathcal{F}, i' \in \mathcal{W}$ *such that*

$$w_{11'} + w_{22'} + ... + w_{kk'} = w_{1k'} + w_{21'} + ... + w_{k(k-1)'}. \tag{3}$$

That is, the corresponding graph has no cycles such that there are two alternating partitions which have the same sum of edge weights.[18]

Main Theorem. *Given a discrete generic assignment game, suppose that Σ^t switches at a rate of $\Theta(N^2)$ time steps.*[19] *The dynamic converges to the core in $O(N^3 \cdot \max\{N, \frac{w^*}{\delta}\})$ iterations with total computational complexity of $O(N^5 \cdot \max\{N, \frac{w^*}{\delta}\})$.*

4.1 Discussion

It is important to restrict the games under consideration as we can see by the following counterexample which is not discrete generic (see figure 3). Suppose all gray edges have weight 10 and no other edges exist. Let $\delta = 1$. The numbers below/above the players give their current aspiration level. Note that each player has a better alternative for every other player (as defined above) and for all i, j $d_i + d_j \geq w_{ij}$ holds. Thus regardless of the current matching, aspiration levels will not change in this state. The unique core state is such that players i_3 and j_3 are matched (indicated by black dotted lines in the figure) but their aspiration levels do not permit them to be matched. Therefore the process does not reach the core.

Figure 3: Counterexample – Weight matrix is not discrete generic.

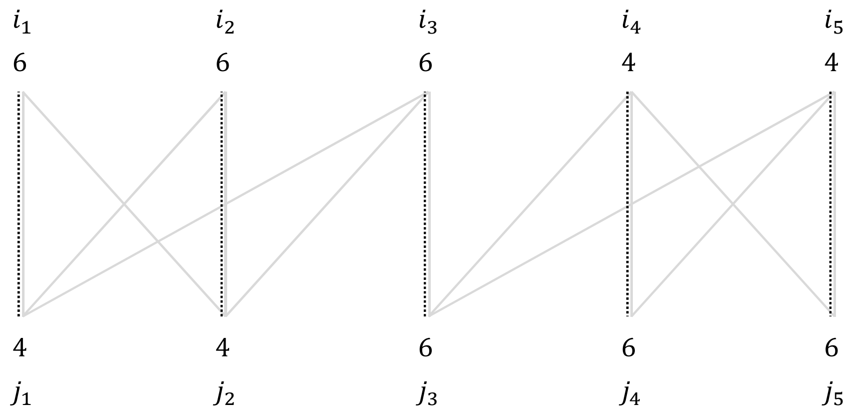

Note that the rate of switching for Σ^t provides a lower bound. That is, a slower polynomial switching rate will slow down the whole process but the process will still converge in polynomial time. On the other hand, if there is no market sentiment our dynamic does not converge. If the market sentiment is always turned against both sides of the market the dynamic converges but may take exponentially long. This can be seen by the following example:

Suppose we have an assignment game with the following edge weights and s, k are both $\Theta(N)$. The unique optimal matching is given by the vertical lines. Suppose that the minimal unit is $\delta = 1$. (See figures 4 and 5.)

[18]Note that, given we consider bipartite graphs, there are no odd cycles.

[19]$f(N) = \Theta(g(N))$ if there exist two constants $k, K \geq 0$ and a positive integer N_0 such that $kg(N) \leq f(N) \leq Kg(N)$ for all $N \geq N_0$.

Figure 4: Counterexample – necessity of alternating market sentiment (1).

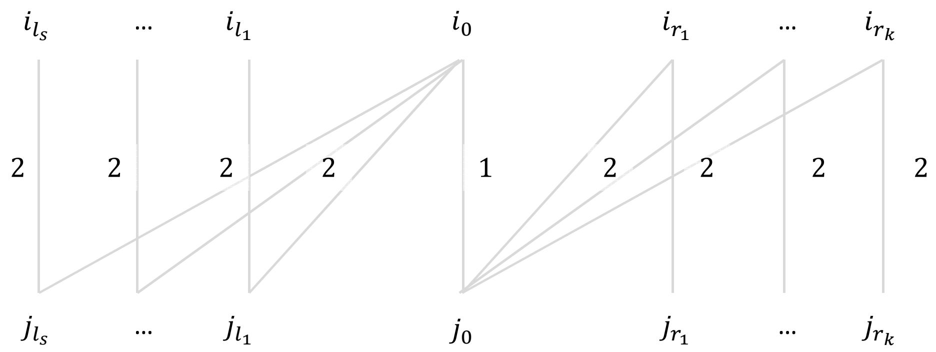

Suppose we are in the following state where the aspiration levels are shown next to each agent (see figure 5). Note that the only edge such that the sum of the aspiration levels exceeds the match value is (i_0, j_0).

Figure 5: Counterexample – necessity of alternating market sentiment (2).

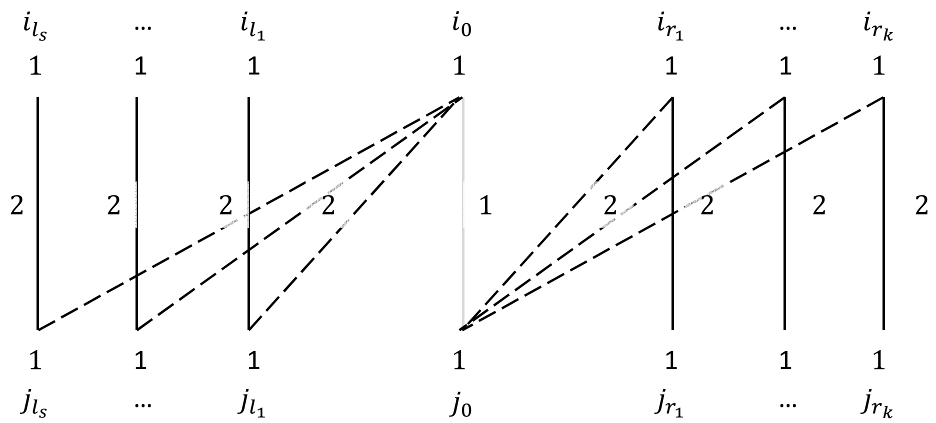

In order for i_0 and j_0 to match at least one of them has to reduce his aspiration level. But this will only be the case if they have no outside option. Thus either all of $j_{r_1}, \ldots, j_{r_s}$ or all of $j_{l_1}, \ldots, j_{l_k}$ have to increase their aspiration levels simultaneously. If both sides of the market are threatened by the market sentiment the sharing between any two (currently matched) agents switches back and forth between sharing equally $(1-1)$ and sharing $0-2$ with the player previously having a tight edge with i_0 or j_0 receiving 2. The probability of them all being at aspiration level 2 at the same time is of the order of $\frac{1}{2^N}$ and therefore the waiting time is exponential and is $O(2^N)$.

5 Proof of the Main Theorem

We shall prove the Main Theorem via several intermediate statements. We first give the intuition and then the formal statements and proofs.

We introduce notation in the spirit of Nax and Pradelski (2012). An edge between player i and j is called *tight* if $d_i^t + d_j^t = w_{ij}$ and *loose* if $d_i^t + d_j^t < w_{ij}$. Given aspiration levels $\mathbf{d}^t$, if for all $i, j, d_i^t + d_j^t \geq w_{ij}$ and for each i, either there exists j such that $d_i^t + d_j^t = w_{ij}$ or $d_i^t = 0$ we shall call $\mathbf{d}^t$ *good*.

5.1 Proof outline

First, we show that once a state is reached such that no loose edges exist all future states won't have loose edges (Lemma 6). We give a bound for the waiting time of $O(N^2 \log N)$. This crude bound results by noting that if two players meet once, they will never again share a loose edge. Thus it suffices that all pairs meet at least once.

Next, note that given we have reached the set of states with no loose edges the sum of aspiration levels is non-increasing in all future states (Lemma 7). Thus we go on to bound the time until the sum of aspiration levels is minimal provided that there are no loose edges (Lemmas 8 and 9). We show that the minimum will be equal to the value of the game, $\nu(V)$. The proof of this result makes use of the market sentiment and the assumption that the game is discrete generic (see Definition 5). Let a *tight component* be a maximal subgraph of players connected by tight edges. Then by discrete genericity all tight components are trees. If all tight components allow perfect matchings we are done. Otherwise we show how in a given tight component which does not allow a perfect matching the sum of aspiration levels can be reduced. We use the market condition to prove that the expected time until the sum of aspiration levels is equal to the value of the game is $O(N^3 \cdot \frac{w^*}{\delta})$. Note that once we are in such a state we will always remain in the set of states with no loose edges and the sum of aspiration levels equal to the value of the game (by Lemmas 6 and 7).

Next, we show by contradiction that from now on for all optimal pairs we have that the edge between them is tight (Lemma 10). It follows that the current aspiration levels are already corresponding to a core allocation. Thus only matchings, break-ups and rematchings are necessary in order to reach the core.

Random walk argument. In order to give a bound on the expected time until matchings occur we create a duality with random walks on trees (Lemma 11). We 'give' each single player a *single token* which is moved around as play unfolds. The idea is that the random walk governing the single tokens will behave in the same way as the underlying model. That is if a player is occupied by a single token he is indeed single. In particular a single player i who matches a player from the other side of the market, j, is passing on the single token to j's previous match. We prove a bound for the meeting time of two random walks on (different) trees. By carefully coupling the movements in our model with the two random walks we show that the expected time until the core is reached from a state as discussed above is $O(N^4)$. We then show that any core state is absorbing.

Finally, it is easy to see that a single iteration has computational complexity $O(N^2)$ (Lemma 12). Putting together the different bounds the Main Theorem will follow.

5.2 Formal proof

Lemma 6. *Starting from any initial state $[M^0, \mathbf{d}^0, \Sigma^0]$, the expected time until the dynamic reaches a state such that for all i, j $d_i + d_j \geq w_{ij}$ is $O(N^2 \log N)$. Once the latter property holds for some $T \geq 0$ it holds for all $t \geq T$.*

Lemma 7. *Let $D^t := \sum_{i \in V} d_i^t$. Suppose the process is in a state as in Lemma 6 at time $T \geq 0$. Then D^t is non-increasing for all $t \geq T$.*

Lemma 8. *Given a state as in Lemma 6, either the expected waiting time until D^t decreases by δ is bounded above by $O(N^2)$ or $D^t = \nu(V)$.*

Lemma 9. *Given a state as in Lemma 6, the expected waiting time until $D^t = \nu(V)$ is $O(N^3 \cdot \frac{w^*}{\delta})$.*

Lemma 10. *Given a state $[M^T, \mathbf{d}^T, \Sigma^T]$ as in Lemma 9 let M be the optimal matching (which is unique by genericity). Then for all $(i, j) \in M$ and for all $t \geq T$, $d_i^t + d_j^t = w_{ij}$.*

Lemma 11. *Given a state as in Lemma 9, the expected time until the core is reached is $O(N^4)$. Further, any core state is absorbing.*

Lemma 12. *One iteration of the dynamic takes $O(N^2)$ operations.*

Proof of Lemma 6. It follows from the updating rule that for any $i \in \mathcal{F}, j \in \mathcal{W}$ and for some T, $d_i^T + d_j^T \geq w_{ij}$ holds, then for all $t \geq T$, $d_i^t + d_j^t \geq w_{ij}$, proving the second part of the Lemma. Now we use an analogy to the Coupon Collector's Problem (see, for example, Feller 1950) to bound the expected time until each pair of players have encountered each other. The Coupon Collector's Problem asks the following question: Suppose you have x coupons, what is the expected time until you have picked each coupon at least once when picking with replacement. The answer is that the number of trials needed is $O(x \log x)$. We thus have that the expected time until each pair of players (N^2) has been selected at least once is $O(N^2 \log N)$. $\square$

Proof of Lemma 7. Given D^t, D^{t+1} can only increase if two players, say i, j, newly match and $d_i^t + d_j^t < w_{ij}$. But given we are in a state as in Lemma 6, for all i, j $d_i^t + d_j^t \geq w_{ij}$ and hence we conclude that $D^{t+1} \leq D^t$. $\square$

Next we shall recall a folk result regarding the existence of perfect matchings on trees (see, for example, Edmonds 1965).

Lemma 13. *A tree C admits a perfect matching (a matching such that each player is matched) if and only if for every vertex v, the graph $C - v$ has exactly one odd sub-component.*

Proof. First, suppose there exists a perfect matching in C. Let u be the vertex matched with v in the perfect matching. Then each component in $C - v$ not containing u must have a perfect matching and is thus even. Finally the component containing u must be odd.

Second, suppose that for every vertex v the graph $C - v$ has exactly one odd component. We shall construct a perfect matching. Given C is a tree and by our assumption each vertex v has a unique neighbor in one component of odd size, say u. We shall show that pairing each v to its neighbor u in the odd component yields a perfect matching. It suffices to show that if u is the neighbor of v in the unique odd component of $C - v$ (say $odd(C - v)$) that it is also the case that v is the neighbor of u in the unique odd component of $C - u$ (say $odd(C - u)$). Note that the component of $C - u$ which contains v is disjoint with $odd(C - v)$. But then it follows that v is part of a component which is comprised of even components in $C - v$. Adding v it follows that v must be in an odd component $odd(C - u)$ which by assumption is unique. □

Note that if the unique optimal matching is not perfect we can transform the game such that the unique optimal matching will be perfect and the core has not changed. To do so add weight zero edges between all players. Note that, two players, i, j, are feasible if one is single and $d_i^t + d_j^t = w_{ij}$, thus in particular they are feasible if $w_{ij} = 0$ and the former conditions hold. Clearly the core does not change through these operations. Suppose from now that the optimal matching is perfect.

Proof of Lemma 8. First, suppose there exists a player i such that he has no tight edge, that is for all j $d_i^t + d_j^t > w_{ij}$. The expected time until i is selected is $O(N)$. If the condition still holds he reduces his aspiration level and no other player increases his aspiration level, hence D has decreased.

Now suppose that for all players i, there exists j such that $d_i^t + d_j^t = w_{ij}$. Let a *tight component*, C^t, be a maximal subgraph of players (nodes) connected by tight ($d_i^t + d_j^t = w_{ij}$) edges. Let $C_1^t, C_2^t, ..., C_k^t$ be all of the tight components. By the genericity assumption the tight components are trees, for otherwise there exists an even cycle such that two alternating partitions have the same sum of edge weights (compare Definition 5).

We shall show that either there exists a perfect matching (a matching such that each player is matched) within each component or there exists at least one component such that the sum of aspiration levels within the component can be reduced independently of the other components.

First, suppose that all components have a perfect matching. It follows that $D^t = \nu(V)$ and we are done.

Now, suppose that there exists a component C^t which has no perfect matching. In the following we describe a series of 'operations' that yield a reduction in the aspiration levels. We shall describe the dynamics and the expected waiting time later.

Suppose, $\#\{i \in C^t : i \in \mathcal{F}\} \neq \#\{i \in C^t : i \in \mathcal{W}\}$, that is, there are either more workers or more firms. A reduction of aspiration levels by δ by the players on the larger side and an increase by δ by the players on the smaller side yields a state where D has decreased.

Note that after the increases and reductions $\mathbf{d}$ is still good, given the reductions are only δ.

Next, suppose $\#\{i \in C^t : i \in \mathcal{F}\} = \#\{i \in C^t : i \in \mathcal{W}\}$. By Lemma 13 a tree C has a perfect matching if and only if for every vertex v, the graph $C - v$ has *exactly* one odd sub-component. Hence, there exists a vertex v such that $C^t - v$ has at least 3 odd sub-components. (Given that the number of firms is equal to the number of workers there must be an odd number of odd sub-components. Thus we can dispense with the case of zero or two odd sub-components.) Without loss of generality, let v be a firm. Denote the odd sub-components resulting from $C - v$ by $C^t_{odd1}, C^t_{odd2}, \ldots$. In each of these sub-components there are either strictly more firms or strictly more workers. In particular given there are at least 3 odd sub-components at least one of them contains more firms than workers. For the latter sub-component aspiration level reductions by δ by all firms and aspiration level increases by δ by all workers yield a decrease of D, while $\mathbf{d}$ remains good.

It is important to note that these adjustments can only take place if the players on the reducing side initially have positive aspiration levels. This might not be the case for all such constructions, but it will always be the case for at least one. This follows by the proof in Nax and Pradelski (2012). They show that the core can be reached by only following transitions as described above or by rematchings within tight component with non-changing aspiration levels. But note that in our model rematchings are not necessary to make transitions as described above possible. Hence, by the proof of Theorem 1 in Nax and Pradelski (2012), if there is no tight component allowing an aspiration level reduction (that is the larger side of a component having all positive aspiration levels), the sum of aspiration levels is optimal, that is $D = \nu(N)$.

We shall now describe how the transitions yielding a reduction in D take place within a (sub-)component. Consider a (sub-)component such that $|\mathcal{F}| > |\mathcal{W}|$ and all firms have positive aspiration levels (the analysis is analogous for the case where there are more workers than firms). By a straightforward argument one sees that there exists at least one firm i which constitutes a leaf of the tree (i.e., only has one edge, say towards worker j). Suppose we are in a phase of a positive market sentiment. Then, in expectation, after N steps worker j is activated and gets matched with firm i and receives an additional δ. (Note that j may be matched with any firm such that he receives an additional δ; we would just follow a different path.) In the meantime, given that workers are currently *not* threatened by the market sentiment no worker has reduced his aspiration level, hence it still holds that firm i has no other tight edge. Players who are not at a leaf in the component have two tight edges and hence no renegotiations can take place. In a next step there is a 'new leaf firm' which can meet its corresponding worker and give up δ. (Since it is a period of a positive market sentiment and thus workers are currently not under threat by the market sentiment such transitions cannot be reversed.) In total this can happen at most N times. Hence after $O(N^2)$ iterations the sum of aspiration levels is reduced by δ when the component has more firms than workers. However, it may be that the only (sub-)component without a perfect matching has more workers than firms. In this case we may have to wait $O(N^2)$ steps until the market sentiment changes and above transitions can take place. It follows that the total time until D is reduced is bounded above by $O(N^2)$. □

Proof of Lemma 9. A crude bound for the maximal aspiration level of players is given by w^*. In the extreme case each player has to reduce his aspiration level w^*/δ times. Thus the number of necessary aspiration level reductions is of the order of $N \cdot \frac{w^*}{\delta}$. The result now follows from Lemma 8. □

Proof of Lemma 10. Let M be the optimal matching. Suppose there exists $(i,j) \in M$, such that $d_i^t + d_j^t > w_{ij}$. By Lemma 6 we know that (in particular) for all $(i',j') \in M$, $d_{i'}^t + d_{j'}^t \geq w_{i'j'}$. Hence

$$\begin{aligned} D = \sum_{(i',j')\in M} d_{i'}^t + d_{j'}^t &= d_i^t + d_j^t + \sum_{(i',j')\neq(i,j),\ (i',j')\in M} d_{i'}^t + d_{j'}^t && (4)\\ &\geq d_i^t + d_j^t + \sum_{(i',j')\neq(i,j),\ (i',j')\in M} w_{i'j'} && (5)\\ &> w_{ij} + \sum_{(i',j')\neq(i,j),\ (i',j')\in M} w_{i'j'} && (6)\\ &= \sum_{(i',j')\in M} w_{i'j'} && (7)\\ &= \nu(V). && (8)\end{aligned}$$

This constitutes a contradiction to the assumption that $D = \nu(V)$ and the result follows. □

In order to prove Lemma 11 we first provide a lemma concerning the expected meeting time of two random walks on two, possibly different, sub-trees of a given tree. Our proof technique mirrors Tetali and Winkler (1991) who give a bound for the meeting time of two random walks on (the same) general graph. They show that the expected meeting time on a general graph on N vertices is at most $\frac{8}{27}N^3$.

Lemma 14. *Let T be an undirected tree on N vertices. Let two tokens x, y be placed on two vertices at time $s = 0$. Suppose that at each time step $s = 1, 2, 3, \dots$ one of the two tokens is activated (which token is chosen may be governed by any probability distribution). The token performs a simple random walk to a neighboring vertex on a truncated tree $T_x^s \subset T$ (or $T_y^s \subset T$). Then, regardless of the rule of activation, if in every time step the (unique) path between x and y is in T_x^s (respectively T_y^s) the expected time $M_{truncated}(x, y)$ until the tokens meet is $O(N^2)$.*

In order to prove the lemma we shall need the following two lemmas from Tetali and Winkler (1991) which can be proved by a straight forward calculation due to the reversibility of the Markov chain for random walks on undirected graphs. Given a connected graph T, let $H_T(x, y)$ be the hitting time from x to y on T, that is the expected time until y is first reached when starting a simple random walk at x on the graph T.

Lemma 15. *(Tetali and Winkler 1991, Lemma 2) Let x, y and z be vertices of a connected undirected graph T. Then*

$$H_T(x,y) + H_T(y,z) + H_T(z,x) = H_T(x,z) + H_T(z,y) + H_T(y,x) \tag{9}$$

Lemma 16. *(Tetali and Winkler 1991, Lemma 3) On any graph T the vertex relation given by*

$$x \leq y \iff H_T(x,y) \leq H_T(y,x) \tag{10}$$

is transitive.

It immediately follows that there exists a vertex v which is *minimal* in this order. Call such a vertex *remote* and note that we have $H_T(x,v) \leq H_T(v,x)$ for every vertex $x \in T$.

Proof of Lemma 14. The proof mirrors Tetali and Winkler (1991). We shall show that $\forall x, y \in T$

$$M_{truncated}(x,y) \leq H_T(x,y) + H_T(y,v) - H_T(v,y). \tag{11}$$

where v is a remote vertex. Define a potential function

$$\begin{aligned} P(x,y) \quad &= \quad H_T(x,y) + \overbrace{H_T(y,v) - H_T(v,y)}^{\geq 0 \text{ since } v \text{ is remote}} &(12)\\ &\overset{\text{Lemma 15}}{=} \quad H_T(y,x) + H_T(x,v) - H_T(v,x). &(13)\end{aligned}$$

Hence, independent of which token is moved we have

$$P(x,y) = 1 + P(\bar{x},y) = 1 + P(x,\bar{y}), \tag{14}$$

where for a function f, $f(\bar{x})$ is defined to be the average of $f(y)$ over all neighbors y of x (note that this defines an expectation). This is the case since, independent of the token moved, the expected hitting time decreases by one for the given token.

Define $\beta(x,y) = M_{truncated}(x,y) - P(x,y)$. Among all pairs maximizing $\beta(\cdot,\cdot)$ choose x^*, y^* at minimal distance from each other. Now suppose inequality (11) does not hold for all $x, y \in T$, that is $\beta(x^*,y^*) > 0$ (note that $x^* \neq y^*$ follows, for otherwise $M_{truncated}(x^*,y^*) = 0 \leq P(x^*,y^*)$). W.l.o.g. assume that the token moved next is on x^*. Then

$$\begin{aligned} M_{truncated}(x^*,y^*) &= P(x^*,y^*) + \beta(x^*,y^*) &(15)\\ &= 1 + P(\bar{x^*},y^*) + \beta(x^*,y^*) &(16)\\ &= 1 + \frac{1}{d(x^*)} \sum_{x_i\colon x_i \text{ neighbor of } x^*} P(x_i,y^*) + \beta(x^*,y^*) &(17)\\ &> 1 + \frac{1}{d(x^*)} \sum_{x_i\colon x_i \text{ neighbor of } x^*} P(x_i,y^*) + \beta(x_i,y^*) &(18)\\ &= 1 + \frac{1}{d(x^*)} \sum_{x_i\colon x_i \text{ neighbor of } x^*} M_{truncated}(x_i,y^*) &(19)\\ &= M_{truncated}(x^*,y^*) &(20)\end{aligned}$$

where $d(x)$ is the degree of node x in T. Inequality (18) holds since for all i, $\beta(x_i, y^*) \leq \beta(x^*, y^*)$ and there exists at least one j such that x_j is closer to y^* than x^* and thus $\beta(x_j, y^*) < \beta(x^*, y^*)$ (since we assumed x^*, y^* to be at minimal distance maximizing $\beta(\cdot,\cdot)$). Equation (19) holds by a similar argument as for equation (14). But $M_{truncated}(x^*, y^*) > M_{truncated}(x^*, y^*)$ constitutes a contradiction and thus proves inequality (11). By Georgakopoulos and Wagner (2013, Corollary 8) the hitting time on a tree is bounded above by the square of the number of edges $(N-1)$. Thus we have

$$M_{truncated}(x, y) \leq H_T(x, y) + H_T(y, v) = O(N^2). \tag{21}$$

□

Lemma 17. *Let T be a tight component (a tree) that allows a perfect matching. Suppose there are exactly two remaining singles i and j. Then after a finite number of transitions where i remained single throughout i matches with some j'.*

Proof. We shall show that the path in the component between i and j must be alternating between unmatched and matched edges. Suppose, by contradiction, that the path between i and j is not alternating between unmatched and matched edges. If i and j share an edge we are done. Else, let i^* be the first firm on the (unique) path from i to j such that i^* is not matched to a worker on the path. (Note that i^* is matched elsewhere.) Let j^* be the worker preceding i^* on the path. We shall construct an instance of Lemma 13. Delete i^* from T. Then the component including i is such that all matches are within the component (since i^* is matched elsewhere).

Since this component does not contain j and contains i and otherwise only contains matched pairs it is of odd size. On the other hand let j^{**} be matched with i^*. Now consider the component containing j^{**} in $T - i^*$. Since j^{**} is matched to i^* it also must be odd. This contradicts Lemma 13 and, by contradiction, establishes the claim that there is an alternating path of unmatched and matched edges between i and j.

Let the unique path be $i, j_1, i_1, j_2, i_2, \ldots, j_k, i_k, j$. We can now construct a path, proving the lemma. Suppose i is activated and matches with j_1, leaving i_1 single. Next i_1 is activated and matches with j_2, leaving i_2 single. Continuing in this manner we see that after k steps i_k is activated and matches with j. □

Proof of Lemma 11. Given $d_i^t + d_j^t \geq w_{ij}$ for all i, j, rematches only occur when at least one player is single. Thus the number of matches is weakly increasing over time and it suffices to bound the expected time until two singles match. It follows immediately from Lemma 10 that each tight component allows a perfect matching. We shall compute the expected time until the number of singles is zero.

We describe the transitions without taking the market sentiment into account since it would only accelerate the process. Note that, since there may be players with aspiration level zero, we cannot necessarily follow a transition pattern as described in the proof of Lemma 8.

We shall construct two simple random walks which we then relate to our process. Let X, Y be two simple random walks on truncated subgraphs of T as defined in Lemma 14 and let $s = 0, 1, 2, 3, ...$ be the time steps at which one of the random walks moves on some truncated tree (such that the path between the two random walks is on the truncated tree). Let the rule of activation for the random walks be such that both random walks always move twice consecutively. That is, for all s odd, if X (respectively Y) is activated in s, then X (respectively Y) is also activated in $s+1$. Define for X (analogously for Y) the following time steps. Initialize $T_0^X = T_0^Y = 0$, then for $i = 1, 2, 3, \ldots$ let

$$T_i^X = \inf\{s > T_{i-1} : \ s \text{ odd, and } X_s \neq X_{s-2}\}, \tag{22}$$

and

$$\widetilde{T} = \inf\{s : \ X_s = Y_s\}. \tag{23}$$

The sequence of time steps at which the random walk X (or Y) has moved away from odd s after two time steps or has collided with Y (respectively X) is given by

$$\mathcal{T}^X = (T_1^X, T_2^X, \ldots, T_I^X, \widetilde{T}), \tag{24}$$

$$\mathcal{T}^Y = (T_1^Y, T_2^Y, \ldots, T_J^Y, \widetilde{T}), \tag{25}$$

where I, J are two random variables describing the lengths of the sequences. Note that for all time-steps $s \notin \mathcal{T}^X \cup \mathcal{T}^Y$ the random walk returned to its previous place after two steps. Define $\widetilde{X}$ and $\widetilde{Y}$ to be the (not simple) random walks of step-size two on truncated trees. Until the two underlying random walks X and Y collide, $\widetilde{X}$ and $\widetilde{Y}$ move at times $\mathcal{T}^X \cup \mathcal{T}^Y$.

We now relate the two random walks, $\widetilde{X}, \widetilde{Y}$, to our process. We shall first consider the case where only two singles remain in a given tight component. Suppose i and j are the only two remaining singles. Note that the two singles will always remain in the *same* tight component. Now consider those time steps where a current single becomes matched and a previously matched player becomes single. (The expected interval between such time steps is upper bounded by $O(N)$.)

Suppose we start the random walk $\widetilde{X}$ on the single firm i and the random walk $\widetilde{Y}$ on the single worker j. Suppose i (or j) gets activated in a given time step. Then $\widetilde{X}$ (or $\widetilde{Y}$) is also activated and moves to a node in the tree at distance two from i, say i' (via j') such that (i', j') were matched in the previous period. This is the condition for the truncated tree on which the random walk $\widetilde{X}$ (or $\widetilde{Y}$) moves, namely the truncated tree starting at i only has branches alternating between tight unmatched and matched edges. Any such node i' is picked uniformly at random amongst the matches of current neighbors. Thus $\widetilde{X}$ (respectively $\widetilde{Y}$) perform the previously defined two-step random walk. The time steps at which this happens are precisely given by $\mathcal{T}^X$ (respectively $\mathcal{T}^Y$). This indicates that player i is matching with player j'. Now if $j' = j$ the two last remaining singles match and the simple random walks X and Y collide. Else i' is now occupied by $\widetilde{X}$.

We are left to give an upper bound for $\mathbb{E}[I+J+1]$. First note that we have $\mathbb{E}[I+J+1] < M_{truncated}(X, Y)$. We shall now argue that

$$\mathbb{E}[I + J + 1] = \frac{1}{4} M_{truncated}(X, Y) \tag{26}$$

and thus the bound we shall give is sharp. If X (or Y) is moved in periods $s+1$ (odd) and $s+2$ we have $\mathbb{P}(X_{s+2} = X_s) = 1/2$. This is the case since on the truncated tree the random walk first moves along a tight edge and then necessarily moves back or along a matched edge. Thus $\mathbb{E}[T_{i+1}^X - T_i^X] = 4$ given the expectation of the geometric distribution with parameter 1/2. Now the result follows.

Overall, given the singles need to be activated in order for the random walks to be moved, the meeting of the two single tokens takes $O(N^3)$ iterations.

Now suppose there are $2k$ ($k \in \{1, \ldots, N\}$) remaining singles. The expected time until there are only $2(k-1)$ singles is bounded by $O(N^3)$. Then for $k = O(N)$ we have a total expected time until the core is reached of $O(N^4)$.[20]

Finally we have to show that any core state is absorbing. We have for all $i \in V$ that i is matched and for all i, j that $d_i^t + d_j^t \geq w_{ij}$ since we are in a core state with the unique perfect matching. Thus there are no feasible pairs of players and the dynamic has reached a fixpoint. □

Proof of Lemma 12. Given that i is activated, suppose there exists a feasible match. The calculation whether a player has no better alternative requires $O(N)$ for each $j' = 1, ..., N$, and thus $O(N^2)$ in total. The case where no feasible match exists for i requires at most $O(N)$ operations. The result now follows. □

Proof of the Main Theorem. Combining Lemma 7, 9, and 11 we find that the expected maximum number of steps to absorption is $O(N^3 \cdot \max\{\log\log N, \frac{w^*}{\delta}\})$ proving the first part of the theorem. By Lemma 12 we know the complexity of each iteration, $O(N^2)$, and the second part of the theorem follows. □

6 Conclusion

In this paper we have shown that agents in large decentralized assignment markets can learn to play core outcomes in polynomial time through simple aspiration adaptation learning rules. Players have no knowledge about other players' past payoffs or actions, their distribution of preferences, or about the value of different matches. Players are randomly activated and thus update their action in an asynchronous, random manner. Players find a feasible match if one exists and the price they agree to is influenced by the market sentiment. While it was previously shown that a variant of this random process is absorbed in the core in finite time with probability one (Chen et al. 2011, Nax and Pradelski 2012, Biró et al. 2012, Klaus and Payot 2013) the dynamic described here converges in polynomial time.

Our result is in the same spirit as Bayati et al. (2014, 2005), Salez and Shah (2009), Celis et al. (2010), Azar et al. (2009), Ackermann et al. (2011), and Kanoria et al. (2011), all of

[20] In line with Tetali and Winkler (1991, Conjecture 2) we conjecture that this bound is not optimal and the total meeting time of k random walks on truncated trees is of the same order as the two-token meeting time. Together with the necessary activation of single tokens this would yield the bound $O(N^3)$.

whom study the rate of convergence of decentralized dynamics for the assignment game. There are two fundamental differences however. First, the informational requirement for players is much higher; in particular players know the match values. Second, all positive results (that is polynomial convergence time) are only shown for deterministic processes with a fixed order of activation. Our result is the first which shows fast convergence for a stochastic decentralized dynamic for the assignment game. To do so we use a duality with random walks on graphs.

The market sentiment turns out to be crucial in our analysis. We have shown that without the market sentiment our dynamic may take exponentially long to converge. This, alongside according results for NTU models (Ackermann et al. 2011) support the conjecture that it is impossible to device a pure better response dynamic which converges fast. It is an open question whether conditions for fast convergence for general classes of games can be formulated. Our proof technique regarding random walks and the use of a market sentiment may be pointers towards a possible path of investigation for other games on networks.

References

Ackermann, H., P. Goldberg, V. Mirrokni, H. Roeglin, and B. Voecking (2011), "Uncoordinated two-sided matching markets." *SIAM Journal on Computing*, 40, 92–106.

Agastya, M. (1997), "Adaptive play in multiplayer bargaining situations." *The Review of Economic Studies*, 64, 411–426.

Agastya, M. (1999), "Perturbed adaptive dynamics in coalition form games." *Journal of Economic Theory*, 89, 207–233.

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

Azar, Y., B. Birnbaum, L. E. Celis, N. R. Devanur, and Y. Peres (2009), "Convergence of local dynamics to balanced outcomes in exchange networks." In *FOCS*.

Bayati, M., C. Borgs, J. Chayes, Y. Kanoria, and A. Montanari (2014), "Bargaining dynamics in exchange networks." *Journal of Economic Theory*, forthcoming.

Bayati, M., D. Shah, and M. Sharma (2005), *Maximum Weight Matching via Max-Product Belief Propagation*. International Symposium on Information Theory.

Bayati, M., D. Shah, and M. Sharma (2008), "Max-product for maximum weight matching: convergence, correctness, and lp duality." *IEEE Transactions on information theory*, 54, 1241–1251.

Bienenstock, E. J. and P. Bonacich (1992), "The core as a solution to exclusionary networks." *Social Networks*, 14, 231–243.

Biró, P., M. Bomhoff, P. A. Golovach, W. Kern, and D. Paulusma (2012), "Solutions for the stable roommates problem with payments." *Lecture Notes in Computer Science*, 7551, 69–80.

Boehm-Bawerk, E. von (1923), *Positive theory of capital (translated by W. Smart)*. G.E. Steckert, New York, (original publication 1891).

Braun, N. and T. Gautschi (2006), "A Nash bargaining model for simple exchange networks." *Social Networks*, 28, 1–23.

Celis, L. E., N. R. Devanur, and Y. Peres (2010), "Local dynamics in bargaining networks via random-turn games." *WINE 2010, LNCS, Springer*, 6484, 133–144.

Chen, B., S. Fujishige, and Z. Yang (2011), "Decentralized market processes to stable job matchings with competitive salaries." KIER Working Papers 749, Kyoto University.

Chung, K.-S. (2000), "On the existence of stable roommate matching." *Games and Economic Behavior*, 33, 206–230.

Cook, K. S. and T. Yamagishi (1992), "Power in exchange networks: a power-dependence formulation." *Social Networks*, 14, 245–265.

Crawford, V. P. and E. M. Knoer (1981), "Job matching with heterogeneous firms and workers." *Econometrica*, 49, 437–540.

Demange, G., D. Gale, and M. Sotomayor (1986), "Multi-item auctions." *Journal of Political Economy*, 94, 863–872.

Diamantoudi, E., L. Xue, and E. Miyagawa (2004), "Random paths to stability in the roommate problem." *Games and Economic Behavior*, 48, 18–28.

Edmonds, J. (1965), "Paths, trees, and flowers." *Canadian Journal of Mathematics*, 17, 49–467.

Edmonds, J. and R. M. Karp (1972), "Theoretical improvements in algorithmic efficiency for network flow problems." *Journal of the ACM*, 19, 248–264.

Elliott, M. L. (2010), "Inefficiencies in networked markets." Working Paper.

Elliott, M. L. (2011), "Search with multilateral bargaining." Working Paper.

Feller, W. (1950), *An introduction to probability theory and ist applications*. John Wiley.

Foster, D. P. and H. P. Young (2006), "Learning to play Nash equilibrium without knowing you have an opponent." *Theoretical Economics*, 1, 341–367.

Gale, D. and L. S. Shapley (1962), "College admissions and stability of marriage." *American Mathematical Monthly*, 69, 9–15.

Georgakopoulos, A. and S. Wagner (2013), "Hitting times, cover costs, and the Wiener index of a tree." Working Paper.

Gibbons, A. (1985), *Algorithmic Graph Theory*. Cambridge University Press.

Hart, S. and A. Mas-Colell (2003), "Uncoupled dynamics do not lead to Nash equilibrium." *American Economic Review*, 93, 1830–1836.

Kanoria, Y., M. Bayati, C. Borgs, J. Chayes, and A. Montanari (2011), "Fast convergence of natural bargaining dynamics in exchange networks." In *Proc. ACM-SIAM Symp. On Discrete Algorithms*, 1518–1537.

Kelso, A. S. and V. P. Crawford (1982), "Job matching, coalition formation, and gross substitutes." *Econometrica*, 50, 1483–1504.

Kim, J. H. and J. Pearl (1983), "A computational model for causal and diagnostic reasoning in inference systems." *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, 83, 190–193.

Klaus, B. and F. Klijn (2007), "Paths to stability for matching markets with couples." *Games and Economic Behavior*, 58, 154–171.

Klaus, B. and F. Payot (2013), "Paths to stability in the assignment problem."

Kleinberg, J. and E. Tardos (2008), "Balanced outcomes in social exchange networks." In *43rd ACM Symposium on Theory of Computing*.

Kojima, F. and M. U. Uenver (2008), "Random paths to pairwise stability in many-to-many matching problems: a study on market equilibrium." *International Journal of Game Theory*, 36, 473–488.

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

Nax, H. H. and B. S. R. Pradelski (2012), "Evolutionary dynamics and equitable core selection in assignment games." Working paper.

Nax, H. H., B. S. R. Pradelski, and H. P. Young (2013), "Decentralized dynamics to optimal and stable states in the assignment game." In *Proceedings of the 52nd IEEE Conference on Decision and Control, December 10-13, 2013. Florence, Italy*, 2391–2397.

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

Newton, J. (2012), "Recontracting and stochastic stability in cooperative games." *Journal of Economic Theory*, 147, 364–381.

Pearl, J. (1982), "Reverend bayes on inference engines: A distributed hierarchical approach." *Proceedings of the Second National Conference on Artificial Intelligence*, 82, 133–136.

Pearl, J. (1988), *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA: Morgan Kaufmann.

Peleg, B. and P. Sudholter (2007), *Introduction to the Theory of Cooperative games.* Springer-Verlag.

Richardson, T. and R. Urbanke (2001), "The capacity of low-density parity check codes under message-passing decoding." *IEEE Transactions on information theory*, 47, 599–618.

Roth, A. E. and I. Erev (1995), "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term." *Games and Economic Behavior*, 8, 164–212.

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

Roth, A. E. and H. Vande Vate (1990), "Random paths to stability in two-sided matching." *Econometrica*, 58, 1475–1480.

Rozen, K. (2013), "Conflict leads to cooperation in Nash Bargaining." *Journal of Economic Behavior and Organization*, 87, 35–42.

Salez, J. and D. Shah (2009), "Belief propagation: an asymptotically optimal algorithm for the random assignment problem." *Mathematics of Operations Research*, 34, 468–480.

Sauermann, H. and R. Selten (1962), "Anspruchsanpassungstheorie der Unternehmung." *Zeitschrift fuer die gesamte Staatswissenschaft*, 118, 577–597.

Sawa, R. (2011), "Coalitional stochastic stability in games, networks and markets." Working Paper, Department of Economics, University of Wisconsin-Madison.

Selten, R. (1998), "Aspiration adaptation theory." *Journal of Mathematical Psychology*, 42, 191–214.

Shapley, L. S. and M. Shubik (1972), "The assignment game 1: The core." *International Journal of Game Theory*, 1, 111–130.

Shimer, R. (2007), "Mismatch." *American Economic Review*, 97, 1074–1101.

Shimer, R. (2008), "The probability of finding a job." *American Economic Review*, 98, 268–273.

Tetali, P. and P. Winkler (1991), "On a random walk problem arising in self-stabilizing token management." In *Proceedings of the tenth annual ACM symposium on Principles of distributed computing*, PODC '91, 273–280, ACM, New York, NY, USA, URL `10.1145/112600.112623`.

Tietz, A. and H. J. Weber (1972), "On the nature of the bargaining process in the Kresko-game." *H. Sauermann (Ed.) Contribution to experimental economics*, 7, 305–334.

Walker, H. A. (2007), *Exchange network theory.* Ritzer, George (ed). Blackwell Encyclopedia of Sociology. Blackwell Publishing.