**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2013

Decentralized dynamics to optimal and stable states in the assignment game

Nax, Heinrich H ; Pradelski, Bary S R ; Young, H Peyton

Abstract: Payoff-driven adjustment dynamics lead to stable and optimal outcomes in decentralized two-sided assignment markets. Pairs of agents from both sides of the market randomly encounter each other and match if 'profitable'. Very little information is available, in particular agents have no knowledge of others' preferences, their past actions and payoffs or the value of the different matches. This process implements optimal and stable - i.e. core - allocations even though agents interact asynchronously and randomly, and there is no central authority enforcing matchings or sharing rules.

DOI: <https://doi.org/10.1109/cdc.2013.6760238>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-191930>

Conference or Workshop Item

Accepted Version

Originally published at:

Nax, Heinrich H; Pradelski, Bary S R; Young, H Peyton (2013). Decentralized dynamics to optimal and stable states in the assignment game. In: 2013 IEEE 52nd Annual Conference on Decision and Control (CDC), Firenze, 10 December 2013 - 13 December 2013. Institute of Electrical and Electronics Engineers, 2391-2397.

DOI: <https://doi.org/10.1109/cdc.2013.6760238>

Decentralized dynamics to optimal and stable states in the assignment game*

Heinrich H. Nax,[†] Bary S. R. Pradelski[‡] & H. Peyton Young[§]

Abstract

Payoff-driven adjustment dynamics lead to stable and optimal outcomes in decentralized two-sided assignment markets. Pairs of agents from both sides of the market randomly encounter each other and match if 'profitable'. Very little information is available, in particular agents have no knowledge of others' preferences, their past actions and payoffs or the value of the different matches. This process implements optimal and stable – i.e. core – allocations even though agents interact asynchronously and randomly, and there is no central authority enforcing matchings or sharing rules.

Keywords – assignment games, cooperative games, core, distributed optimization, evolutionary game theory, learning, linear programming, matching markets

I. Introduction

Internet and communication technologies have created environments where many agents repeatedly interact with each other despite having little information about the structure of the game and other agents' actions. Indeed, even after many encounters, agents may learn little or nothing about the preferences and past actions of other participants. In this paper we propose a dynamic model that incorporates these features and explore its performance and stability properties.

We shall be interested in decentralized bilateral environments with transferable utility (TU) where agents from each side repeatedly interact and submit share demands at which they are currently willing to be matched. Areas in which such models have found applications in-

clude artificial intelligence [1], routing unmanned vehicles (e.g., UAVs) [2], error correcting codes [3], double auctions [4], and online markets such as matching buyers and sellers of goods, for matching workers and firms, and for matching hotels with clients.¹

In this paper, we provide a random, asynchronous and uncoupled dynamic that converges to the core. This is of interest in mechanism design for the development of algorithms for decentralized optimization tasks. Our work complements a number of existing contributions from mathematics, computer science, sociology and economics.²

In economics, TU matching markets have traditionally been analyzed using game-theoretic methods based on submission of bid-ask lists to a central authority [10]. There is a sizable literature on centralized matching algorithms and centralized auction mechanisms based on the *assignment game* [10]; see, for example, [11, 12, 13]. While decentralized dynamics [14, 15, 16, 17] have been studied for the related class of non-transferable utility (NTU) matching problems as introduced by [18], it is an open problem to characterize the asymptotic behavior of such dynamics for TU games.

In [19] we propose the following behavioral model for assignment games: pairs of agents from the two market sides randomly encounter each other and enter a new match if their match is 'profitable', which they can see from their current bids and offers; moreover, depending on whether an agent is currently matched or single, aspirations are adjusted up or down, thus, driving bids, offers and match prices. Importantly, the players do not have enough information or cognitive capacity to optimize the value of the matches. We propose simple payoff-based adjustment rules that lead to core-stable outcomes in such environments.³

The result fits into a wider literature showing how cooperative game solutions can be understood as outcomes

*Manuscript created August 26, 2013. This research was supported by the United States Air Force Office of Scientific Research Grant FA9550-09-1-0538 and the Office of Naval Research Grant N00014-09-1-0751. Heinrich Nax also acknowledges support of the Agence Nationale de Recherche project NET; Bary Pradelski of the Oxford-Man Institute of Quantitative Finance.

[†]Paris School of Economics, École Normale Supérieure.

[‡]Oxford-Man Institute, University of Oxford.

[§]Department of Economics, University of Oxford.

¹Examples for online platforms with a two-sided bid-ask structure include www.priceline.com's Name-Your-Own-Price[®] and www.HireMeNow.com's Name-Your-Own-Wage[™].

²See, for example, [5, 1, 6, 7] for mathematics and computer science and [8, 9] for sociology.

³See also [20] for a decentralized dynamic without aspiration adjustments and [21] for a generalization thereof to roommate problems.

of a dynamic learning process [22, 23, 24, 25]. The existing strand of this literature is based on noisy best-response dynamics and the theory of large deviations [26, 27, 28]. An alternative approach to learning, and the one taken in this paper, is based on the idea that agents adjust their aspiration levels depending on their realized payoffs. There is no presumption about attempts to play a best-reply. Instead, such rules are *completely uncoupled* [29] from information about other players' utilities and actions.⁴ Families of such rules lead to Nash equilibrium in generic non-cooperative games (see, for example, [31, 29, 32, 33, 34, 35, 36]). In this paper, we analyze a completely uncoupled procedure for assignment games that implements the core – that is, optimal and stable outcomes – with probability one after finite time. This is achieved even though there is no central authority arranging prices or matches.⁵

II. TU matching markets

In this section we shall introduce the conceptual framework for analyzing matching markets with transferable utility (TU).

II.A. The model

II.A.1. Static components. The population $N = \{F \cup W\}$ consists of firms $F = \{f_1, \dots, f_m\}$ and workers $W = \{w_1, \dots, w_n\}$.⁶

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 to firm i .

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

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

$$\alpha_{ij} = (p_{ij}^+ - q_{ij}^-)_+. \quad (1)$$

It will be convenient to assume that all values p_{ij}^+ and q_{ij}^- can be expressed as multiples of some minimal unit

⁴This definition is a strengthening of *uncoupled* rules introduced by [30].

⁵See [19] for the full proofs.

⁶The two sides of the market could also, for example, represent tasks and machines, or areas and UAVs.

of currency δ , e.g., “dollars”.

II.A.2. Dynamic components. Let $t = 0, 1, 2, \dots$ be the time periods.

Assignment. For all agents $(i, j) \in F \times W$, let $a_{ij}^t \in \{0, 1\}$. If (i, j) is matched then $a_{ij}^t = 1$, if (i, j) is unmatched then $a_{ij}^t = 0$. 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 , one pair of players is drawn at random and each one makes a bid for the other. We assume that the two players' bids are such that the resulting payoff to each player is at least equal to his aspiration level, and with positive probability is exactly equal to his aspiration level.

Formally, firm $i \in F$ encounters $j \in W$ and submits a random bid $b_i^t = p_{ij}^t$, where p_{ij}^t is the maximal amount i is currently willing to pay if matched with j . Similarly, worker $j \in W$ submits $b_j^t = q_{ij}^t$, where q_{ij}^t is the minimal amount j is currently willing to accept if matched with i . A bid is separable into two components; the current (deterministic) aspiration level and a random variable that represents an exogenous shock to the agent's aspiration level. Specifically let P_{ij}^t, Q_{ij}^t be independent random variables that take values in $\delta\mathbb{N}_0$ where 0 has positive probability.⁷ We thus have for all i, j

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

$$\text{and} \quad q_{ij}^t = (q_{ij}^- + d_j^{t-1}) + Q_{ij}^t \quad (3)$$

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. Note that if the random variable is zero, the agent bids exactly according to his current aspiration level.

Profitability. A pair of bids (p_{ij}^t, q_{ij}^t) is *profitable* if a player, in expectation, receives a (strictly) higher payoff if the match is formed.

Note that, two players whose match value is zero will never match.

⁷Note 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.

Random matches. Agents are *activated* by independent Poisson arrival processes. An active agent randomly encounters one agent from the other side of the market drawn from a distribution with full support. The two agents

- reveal current bids for each other, and
- identify if these are profitable.

The two agents match if the bids are profitable. (Details 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".⁸ Namely, there is an integer k such that, if $p_{ij}^t + q_{ij}^t = 2k\delta$ then $\pi_{ij}^t = k\delta$, and if $p_{ij}^t + q_{ij}^t = (2k+1)\delta$ then $\pi_{ij}^t = k\delta$ with probability 0.5 and $\pi_{ij}^t = (k+1)\delta$ with probability 0.5. This implies that when a pair is matched we have $p_{ij}^t = q_{ij}^t$. 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.

Payoffs. In Z^t the *payoff* to firm i , ϕ_i^t , is $p_{ij}^t - \pi_{ij}^t$ if (i, j) are matched and 0 if i is single. Similarly, the *payoff* to worker j , ϕ_j^t , is $\pi_{ij}^t - q_{ij}^t$ if (i, j) are matched and 0 if j is single.

II.B. Assignment games

We can now formally define the assignment game as introduced in [10].

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, $W = \{w_1, \dots, w_n\}$ is a set of n workers, $\alpha = (\alpha_{ij})_{i \in \{1, 2, \dots, m\}, j \in \{1, 2, \dots, n\}}$ is the matrix of match values, and $\mathbf{A} = (A_{ij})_{i \in \{1, 2, \dots, m\}, j \in \{1, 2, \dots, n\}}$ is the assignment matrix with 0/1 values and row/column sums at most one. Let $\mathcal{A}$ be the set of all possible assignments.

Cooperative assignment game. Given $[F, W, \alpha]$, the *cooperative assignment game* $G(v, N)$ is defined as follows. Define $v : S \subseteq N \rightarrow \mathbb{R}$ such that

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

⁸It 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. This we assume to not favor either market side.

where the maximum is taken over all sets $\{(i_1, j_1), \dots, (i_k, j_k)\}$ consisting of disjoint pairs of firms and workers in S . The number $v(N)$ is 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 $\mathbf{d}^t$ are the aspiration levels. Let Ω be the set of all states.

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

Pairwise stability. An aspiration level $\mathbf{d}^t$ is *pairwise stable* if $\forall i, j$ with $A_{ij} = 1$, $p_{ij}^+ - d_i^t = q_{ij}^- + d_j^t$, 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. (Shapley & Shubik, 1972) 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 optimal and $\mathbf{d}$ is pairwise stable.

Shapley & Shubik formulate this result in terms of pay-offs. Subsequent literature has investigated the structure of the assignment game core, which turns out to be very rich (see, for example, [37, 38]).

III. Evolving play

A fixed population of agents, $N = F \cup W$, repeatedly plays the assignment game $G(v, N)$ by making random encounters with potential partners, by submitting bids and by adjusting them dynamically as the game evolves. Agents become activated spontaneously according to independent Poisson arrival processes. Our results hold independent of whether the rates differ across agents and time (for example, single agents might become active at a faster rate). The distinct times at which one agent becomes active will be called *periods*.

III.A. Behavioral algorithm

We provide the behavioral algorithm in pseudo code.

Algorithm 1 Random behavioral dynamics

Initialization. Initially $t = 0$ and set $Z^0 = [A^0, d^0]$ to any permissible state, e.g., $Z^0 = [A^0, d^0] = [0_{m,n}, \vec{0}]$.

Updating rule. For $t + 1, t \geq 0$:

1. A unique agent, i , becomes active (activated by an independent Poisson arrival process).
2. He randomly encounters one agent from the other side of the market, j , drawn from a distribution with full support.⁹

If the match is profitable given current bids, then
the new match forms,

else
the match does not form.

end

3. **If a new match (i, j) is formed, then**

the price governing the new match π_{ij}^{t+1} , is the average (subject to rounding) of p_{ij}^t and q_{ij}^t . The aspiration levels of the newly matched pair (i, j) are adjusted according to their newly realized payoffs: $d_i^{t+1} = p_{ij}^+ - \pi_{ij}^{t+1}$ and $d_j^{t+1} = \pi_{ij}^{t+1} - q_{ij}^-$.

else

if the active agent i is matched, then
 i remains with his previous partner and keeps his previous aspiration level.

else
 i remains single and with positive probability reduces his aspiration level to

$$d_i^{t+1} = (d_i^t - X_i^{t+1})_+, \quad (4)$$

where X_i^{t+1} is an independent random variable taking values in $\delta \cdot \mathbb{N}_0$, such that $X_i^{t+1} = \delta$ with positive probability.

end

end

4. All other aspiration levels and matches remain fixed. If i or j are newly matched their previous partners are now single.
-

III.B. Example

Let $N = \{f_1, f_2\} \cup \{w_1, w_2, w_3\}$, $p_{1j}^+ = 40, 31, 20$ 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 $\delta = 1$.

⁹In terms of the distribution of encounters, priority could be given to those involving single agents; or any distribution with full support can be used.

III.B.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. The illustrations below show the states at the end of the respective periods. 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.) Note that some of the bids for players which are currently not matched may exceed the respective aspiration levels. 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, 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).

III.B.2. 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 he remains single independent of his encounter (f_2 in the illustration) and, with positive probability, reduces his aspiration level by 1.

III.B.3. 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 and w_3 meet,

and f_2 bids 30 for w_3 and w_3 bids 29 for f_2 (hence the match is profitable), in which case 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.)

III.B.4. Period $t + 3$: activation of matched 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 independent of his encounter (f_2 in the illustration) and, with positive probability, reduces his aspiration level by 1.

The resulting state is in the core.¹⁰

IV. 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 pay-offs.

Theorem 1. *Given an assignment game $G(v, N)$, from any initial state $Z^t \in \Omega$, the process is absorbed into the core in finite time with probability 1.*

For the detailed proof see [19]. We shall outline the main ideas of the proof. See also [20, 21] for models without aspiration levels for the assignment game and general graphs.

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

¹⁰Note 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.

positive probability. Since the state space is finite this will suffice to prove the theorem. 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 equations (2,3) for the bids, that for all $i, j, p_{ij}^t = p_{ij}^+ - d_i^{t-1}$ and $q_{ij}^t = q_{ij}^- + d_j^{t-1}$. Recall that any two agents, with positive probability, encounter each other in any period. It shall be understood in the proof that the relevant agents in any period encounter each other. We can then say that a pair of aspiration levels d_i^t, d_j^t is *profitable* if either $d_i^t + d_j^t < \alpha_{ij}$, or $d_i^t + d_j^t = \alpha_{ij}$ and i, j are single. Restricting attention to this particular class of paths will permit a more transparent analysis of the transitions, which can be described solely in terms of the aspiration levels. It now will suffice to establish 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.¹¹

Claim 1's proof is based on the fact that any random sequence of activations may occur. A straightforward argument can then be applied. We shall give a more detailed outline of the proof of Claim 2, since it is more involved.

Outline 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

¹¹Note 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.

Figure 1. Example alternating path.

any such state, we show that, within a bounded number of periods and with positive probability (bounded below), one of the following holds:

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

(B) 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.

In general, say an edge is *tight* if $d_i + d_j = \alpha_{ij}$. 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 1 illustrates a maximal alternating path starting at f_1 . Unmatched tight edges are indicated by dashed, matched tight edges by solid lines.

With this definitions at hand the proof follows the following logic: Starting in a state $[A, \mathbf{d}]$ 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. The elimination of paths is based on integer programming arguments [39, 40], carefully “shifting” one δ along a paths. 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. The resulting state must be in the core and is absorbing because single agents (demanding zero) 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 is absorbed into the core in finite time with probability 1.

V. Conclusion

In this paper we have shown that agents in large decentralized assignment games 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 algorithm leads to the core with probability one. The proof uses integer programming arguments, but the players do 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.

Acknowledgments

We thank Peter Biró, Gabrielle Demange, Gabriel Kreindler, Jonathan Newton, Tom Norman and four anonymous referees for suggesting a number of improvements.

References

- [1] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA: Morgan Kaufmann, 1988.
- [2] M. Alighanbary and J. P. How, “Decentralized task assignment for unmanned aerial vehicles,” *Conference on Decision and Control and European Control Conference '05*, vol. 1, pp. 5668–5673, 2005.
- [3] T. Richardson and R. Urbanke, “The capacity of low-density parity check codes under message-passing decoding,” *IEEE Transactions on information theory*, vol. 47, pp. 599–618, 2001.
- [4] D. P. Bertsekas and D. A. Castanon, “A forward/reverse auction algorithm for asymmetric assignment problems,” *Computational Optimization and Applications*, vol. 1, pp. 277–297, 1993.
- [5] M. Bayati, D. Shah, and M. Sharma, “Max-product for maximum weight matching: convergence, correctness, and lp duality,” *IEEE Transactions on information theory*, vol. 54, no. 3, pp. 1241–1251, 2008.
- [6] J. Kleinberg and E. Tardos, “Balanced outcomes in social exchange networks,” in *43rd ACM Symposium on Theory of Computing*, 2008.
- [7] M. Bayati, C. Borgs, J. Chayes, Y. Kanoria, and A. Montanari, “Bargaining dynamics in exchange networks,” 2011. Working Paper.

- [8] H. A. Walker, *Exchange network theory*. Ritzer, George (ed). Blackwell Encyclopedia of Sociology. Blackwell Publishing, 2007.
- [9] K. S. Cook and T. Yamagishi, "Power in exchange networks: a power-dependence formulation," *Social Networks*, vol. 14, pp. 245–265, 1992.
- [10] L. S. Shapley and M. Shubik, "The assignment game 1: The core," *International Journal of Game Theory*, vol. 1, no. 1, pp. 111–130, 1972.
- [11] V. P. Crawford and E. M. Knoer, "Job matching with heterogeneous firms and workers," *Econometrica*, vol. 49, pp. 437–540, 1981.
- [12] A. S. Kelso and V. P. Crawford, "Job matching, coalition formation, and gross substitutes," *Econometrica*, vol. 50, pp. 1483–1504, 1982.
- [13] G. Demange, D. Gale, and M. Sotomayor, "Multi-item auctions," *Journal of Political Economics*, vol. 94, pp. 863–872, 1986.
- [14] A. E. Roth and H. Vande Vate, "Random paths to stability in two-sided matching," *Econometrica*, vol. 58, pp. 1475–1480, 1990.
- [15] K.-S. Chung, "On the existence of stable roommate matching," *Games and Economic Behavior*, vol. 33, pp. 206–230, 2000.
- [16] E. Diamantoudi, L. Xue, and E. Miyagawa, "Random paths to stability in the roommate problem," *Games and Economic Behavior*, vol. 48, no. 1, pp. 18–28, 2004.
- [17] B. Klaus and F. Klijn, "Paths to stability for matching markets with couples," *Games and Economic Behavior*, vol. 58, no. 1, pp. 154–171, 2007.
- [18] D. Gale and L. S. Shapley, "College admissions and stability of marriage," *American Mathematical Monthly*, vol. 69, pp. 9–15, 1962.
- [19] H. H. Nax, B. S. R. Pradelski, and H. P. Young, "The evolution of core stability in decentralized matching markets," 2012. Department of Economics Discussion Paper Series 607, University of Oxford.
- [20] B. Chen, S. Fujishige, and Z. Yang, "Decentralized market processes to stable job matchings with competitive salaries," 2012. Working Paper.
- [21] P. Biro, M. Bomhoff, P. A. Golovach, W. Kern, and D. Paulusma, "Solutions for the stable roommates problem with payments," *Lecture Notes in Computer Science*, vol. 7551, pp. 69–80, 2012.
- [22] M. Agastya, "Perturbed adaptive dynamics in coalition form games," *Journal of Economic Theory*, vol. 89, pp. 207–233, 1999.
- [23] T. Arnold and U. Schwalbe, "Dynamic coalition formation and the core," *Journal of Economic Behavior and Organization*, vol. 49, pp. 363–380, 2002.
- [24] K. Rozen, "Conflict leads to cooperation in Nash Bargaining," 2010. Working Paper.
- [25] J. Newton, "Recontracting and stochastic stability in cooperative games," *Journal of Economic Theory*, vol. 147, no. 1, pp. 364–381, 2012.
- [26] D. Foster and H. P. Young, "Stochastic evolutionary game dynamics," *Theoretical Population Biology*, vol. 38, pp. 219–232, 1990.
- [27] H. P. Young, "The evolution of conventions," *Econometrica*, vol. 61, no. 1, pp. 57–84, 1993.
- [28] M. Kandori, G. J. Mailath, and R. Rob, "Learning, mutation, and long run equilibria in games," *Econometrica*, vol. 61, no. 1, pp. 29–56, 1993.
- [29] D. P. Foster and H. P. Young, "Learning to play Nash equilibrium without knowing you have an opponent," *Theoretical Economics*, vol. 1, pp. 341–367, 2006.
- [30] S. Hart and A. Mas-Colell, "Uncoupled dynamics do not lead to Nash equilibrium," *American Economic Review*, vol. 93, pp. 1830–1836, 2003.
- [31] P. S. Sastry, V. V. Phansalkar, and M. Thathachar, "Decentralized learning of Nash equilibria in multi-person stochastic games with incomplete information," *IEEE Transactions on Systems, Man and Cybernetics*, vol. 24, no. 5, pp. 769–777, 1994.
- [32] F. Germano and G. Lugosi, "Global Nash convergence of Foster and Young's regret testing," *Games and Economic Behavior*, vol. 60, pp. 135–154, 2007.
- [33] J. R. Marden, H. P. Young, G. Arslan, and J. Shamma, "Payoff-based dynamics for multi-player weakly acyclic games," *SIAM Journal on Control and Optimization*, vol. 48, pp. 373–396, 2009.
- [34] H. P. Young, "Learning by trial and error," *Games and Economic Behavior*, vol. 65, pp. 626–643, 2009.
- [35] B. S. R. Pradelski and H. P. Young, "Learning efficient Nash equilibria in distributed systems," *Games and Economic Behavior*, vol. 75, pp. 882–897, 2012.
- [36] M. S. Stankovic, K. H. Johansson, and D. M. Stipanovic, "Distributed seeking of Nash equilibria in mobile sensor networks," *IEEE Transactions on Automatic Control*, vol. 57, no. 4, pp. 904–919, 2011.
- [37] A. E. Roth and M. Sotomayor, "Two-sided matching," *Handbook of Game Theory with Economic Applications*, vol. 1, pp. 485–541, 1992.
- [38] M. Sotomayor, "Some further remarks on the core structure of the assignment game," *Mathematical Social Sciences*, vol. 46, pp. 261–265, 2003.
- [39] H. W. Kuhn, "The Hungarian Method for the assignment problem," *Naval Research Logistic Quarterly*, vol. 2, pp. 83–97, 1955.
- [40] M. L. Balinski, "Integer programming: Methods, uses, computation," *Management Science*, vol. 12, pp. 253–313, 1965.