

Uncoordinated Two-Sided Matching Markets*

Heiner Ackermann[†]

Fraunhofer Institute for Industrial Mathematics
Kaiserslautern, Germany
heiner.ackermann@itwm.fraunhofer.de

Paul W. Goldberg[‡]
Dept. of Computer Science
University of Liverpool, U.K.
P.W.Goldberg@liverpool.ac.uk

Vahab S. Mirrokni
Google Research
New York, NY.
mirrokni@google.com

Heiko Röglin
Dept. of Computer Science
University of Bonn
heiko@roeglin.org

Berthold Vöcking[†]
Dept. of Computer Science
RWTH Aachen University
voecking@cs.rwth-aachen.de

November 7, 2010

Abstract

Various economic interactions can be modeled as two-sided markets. A central solution concept to these markets are *stable matchings*, introduced by Gale and Shapley. It is well known that stable matchings can be computed in polynomial time, but many real-life markets lack a central authority to match agents. In those markets, matchings are formed by actions of self-interested agents. Knuth introduced uncoordinated two-sided markets and showed that the uncoordinated better response dynamics may cycle. However, Roth and Vande Vate showed that the random better response dynamics converges to a stable matching with probability one, but did not address the question of *convergence time*.

In this paper, we give an *exponential lower bound* for the convergence time of the random better response dynamics in two-sided markets. We also extend the results for the better response dynamics to the *best response* dynamics, i.e., we present a cycle of best responses, and prove that the random best response dynamics converges to a stable matching with probability one, but its convergence time is exponential. Additionally, we identify the special class of *correlated matroid two-sided markets* with real-life applications for which we prove that the random best response dynamics converges in expected polynomial time.

*A preliminary version appeared in the Proceedings of the 9th ACM Conference on Electronic Commerce (EC).

[†]Heiner Ackermann and Berthold Vöcking are supported in part by the Ultra High-Speed Mobile Information and Communication Research Cluster (UMIC) established under the excellence initiative of the German government.

[‡]Paul W. Goldberg is supported by EPSRC Grant EP/G069239/1.

1 Introduction

One main function of many markets is to match agents of different kinds to one another, for example men and women, students and colleges [7], interns and hospitals [17, 18], and firms and workers. Gale and Shapley [7] introduced *two-sided markets* to model these problems. A two-sided market consists of two disjoint groups of agents. Each agent has some preferences about the agents on the other side and can be matched to one of them. A matching is *stable* if it does not contain a *blocking pair*, that is, a pair of agents from different sides who can deviate from this matching and both benefit. Gale and Shapley [7] showed that stable matchings always exist and can be found in polynomial time. Besides their theoretical appeal, two-sided matching models have proved useful in the empirical study of many labor markets such as the National Resident Matching Program (NRMP). Since the seminal work of Gale and Shapley, there has been a significant amount of work in studying two-sided markets, especially on extensions to many-to-one matchings and preference lists with ties [11, 19, 5, 4]. See for example, the book by Knuth [12], the book by Gusfield and Irving [9], or the book by Roth and Sotomayor [19].

In many real-life markets, there is no central authority to match agents, and agents are self-interested entities. This motivates the study of *uncoordinated two-sided markets*, first proposed by Knuth [12]. Uncoordinated two-sided markets can be modeled as a game among agents of one side, which we call the *active* side. The strategy of each active agent is to choose one agent from the *passive* side. Stable matchings correspond to Nash equilibria of the corresponding games. In order to understand the behavior of the agents in these uncoordinated markets, it is interesting to consider better response dynamics among agents and to analyze whether uncoordinated agents reach a stable matching and if so how long it takes. In this regard, Knuth showed that a sequence of better responses of agents can cycle and posed a question concerning the convergence of this dynamics. Consider the following *random better response dynamics*: at each step, pick a blocking pair of agents uniformly at random and let the agents in this pair match to each other. Roth and Vande Vate [20] proved that the random better response dynamics converges to a stable matching with probability one. This is done by showing that given any matching, there is a polynomial-length sequence that leads from that matching to a stable one. This implies an exponential-time upper bound on the expected time taken by this dynamics. However, convergence time was not their main focus. We believe that studying this question is crucial for understanding the behavior of uncoordinated agents as it corresponds to the question of how long an uncoordinated market needs to stabilize.

Our first result in this paper is an *exponential lower bound* for the convergence time of this better response dynamics in uncoordinated two-sided markets. Both Knuth's cycle [12], and Roth and Vande Vate's proof [20] hold only for the better response dynamics, and not for the *best response dynamics*. We extend the results in [12, 20] to best responses. That is, we illustrate a cycle of best responses of agents, and then, using a potential function argument, we show that starting from any matching, there exists a short sequence of best responses of agents to a stable matching. As a corollary of the latter result, we obtain that every sequence of best responses starting with the *empty* matching reaches a stable matching after a polynomial number of steps. Hence, when starting with the empty matching, no central coordination is needed to reach a stable matching quickly if

agents play only best responses. In contrast to this, we show an exponential lower bound for the convergence time of the *random best response dynamics* when arbitrary starting configurations are allowed.

The above lower bounds show that the decentralized game theoretic approach for stable matchings does not converge in polynomial time. This motivates studying special cases of two-sided markets for which the convergence time is polynomial. In this regard, we consider a natural class of *correlated two-sided markets*, which are inspired from real-life one-sided market games in which players have preferences about a set of markets, and the preferences of markets are correlated with the preferences of players. In a correlated two-sided market, there is a payoff associated with every possible pair of active and passive agent. Both active and passive agents are interested in maximizing their payoff, that is, an agent i prefers an agent j to an agent j' if the payoff associated with pair (i, j) is larger than the payoff associated with pair (i, j') . Two illustrative examples of these markets are market sharing games [8], and distributed caching games [6, 15]. These markets have also been studied for finding stable geometric configurations with applications in VLSI design [10]. This special class of two-sided markets is shown to be a potential game in [2] and complexity related questions are studied in [1]. For the stable roommates problem, Lebedev et al. [13] and Mathieu [14] consider instances with *acyclic* preference lists. It turns out that the classes of acyclic and of correlated instances coincide [1]. Lebedev et al. show that for acyclic instances the best response dynamics cannot cycle, that from every state there exists a short sequence of best responses leading to a Nash equilibrium, and that the random best response dynamics converges in expected polynomial time. Mathieu shows that in the worst case the best response dynamics can take an exponential number of steps to reach a stable matching. We extend the result that the random best response dynamics converges quickly to *correlated matroid two-sided markets*, in which each active agent can propose and can be matched to several passive agents.

2 Preliminaries and Notations

In this section, we define the problems and notations that are used throughout the paper. **Two-sided Markets.** A two-sided market consists of two disjoint groups of agents \mathcal{X} and \mathcal{Y} , e.g., women and men. Each agent has a preference list over the agents of the other side. An agent $i \in \mathcal{X} \cup \mathcal{Y}$ can be assigned to one agent j in the other side. Then she gets *payoff* $p_i(j)$. If the preference list of agent i is (a_1, a_2, \dots, a_n) , we say that agent i has payoff $k \in \{0, \dots, n-1\}$ if she is matched to agent a_{n-k} . Also, we say that an agent has payoff -1 if she is unmatched. Given a matching M , we denote the payoff of an agent i in matching M by $p_i(M)$.

Given a matching M , an agent $x \in \mathcal{X}$ and an agent $y \in \mathcal{Y}$ form a *blocking pair* if $\{x, y\} \notin M$ and $p_x(y) > p_x(M)$ and $p_y(x) > p_y(M)$. Given a matching M and a blocking pair (x, y) in M , we say that a matching M' is obtained from M by *resolving* the blocking pair (x, y) if the following holds: $\{x, y\} \in M'$, any partners with whom x and y are matched in M are unmatched in M' , and all other edges in M and M' coincide. A matching is *stable* if it does not contain a blocking pair.

Uncoordinated Two-sided Markets. We model the uncoordinated two-sided market $(\mathcal{X}, \mathcal{Y})$ as a game $G(\mathcal{X}, \mathcal{Y})$ among agents of the *active* side \mathcal{X} . The strategy of each active

agent $x \in \mathcal{X}$ is to choose one agent y from the *passive* side \mathcal{Y} . The goal of each active agent $x \in \mathcal{X}$ is to maximize her payoff $p_x(y)$. Given a strategy vector of active agents, an active agent x obtains payoff $p_x(y)$ if she proposes to y , and if she is the *winner* of y . Agent x is the winner of y if y ranks x highest among all active agents who currently propose to him. Additionally, passive agent y obtains $p_y(x)$ if x is the winner of y . We say that a strategy vector is a *pure Nash equilibrium* if none of the active agents can increase their payoff by unilaterally changing their strategy. Hence, stable matchings in an uncoordinated two-sided market $(\mathcal{X}, \mathcal{Y})$ correspond to *pure Nash equilibria* of the corresponding game $G(\mathcal{X}, \mathcal{Y})$ and vice versa.

Consider two agents $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. If a blocking pair (x, y) is resolved, we say that x plays a *better response*. If there does not exist a blocking pair (x, y') with $p_x(y') > p_x(y)$, then we say that x plays a *best response* when the blocking pair (x, y) is resolved. In the *random better response dynamics* at each step a blocking pair is chosen uniformly at random and resolved. In the *random best response dynamics* at each step an active agent from \mathcal{X} is chosen uniformly at random and allowed to play a best response. This is equivalent to choosing uniformly at random a blocking pair that constitutes a best response by an agent from \mathcal{X} .

Throughout the paper, we use *women* or *players* as active agents, and *men* or *resources* as passive agents in the corresponding market game. Notice that the random better response dynamics treats men and women symmetrically, so that in this context there is in fact no distinction between the two sides. On the other hand, the asymmetry affects the random best response dynamics, and the lower bound on its convergence time (Theorem 6) does not apply to a dynamics where members of either side get chosen at random to make best responses.

Correlated Two-sided Markets. In general, there are no dependencies between the preference lists of agents. Correlated two-sided markets are examples in which the preference lists are correlated. Assume that there is a payoff $p_{x,y} \in \mathbb{N}$ associated with every pair (x, y) of agents $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ such that $p_x(y) = p_y(x) = p_{x,y}$. The preference lists of both active and passive agents are then defined according to these payoffs, e.g., a passive agent y prefers an active agent x to an active agent x' if $p_{x,y} > p_{x',y}$. We assume that for every agent i , the payoffs associated to all pairs including agent i are pairwise distinct. Then the preference lists are uniquely determined by the ordering of the payoffs.

Many-to-One Two-Sided Markets. In a many-to-one two-sided market, the *strategy space* $\mathcal{F}_x \subseteq 2^{\mathcal{Y}}$ of every player $x \in \mathcal{X}$ is a collection of subsets of resources, that is, every player $x \in \mathcal{X}$ can propose to a subset $S_x \in \mathcal{F}_x$ of resources. Each resource $y \in \mathcal{Y}$ has a strict preference list over the set of players in \mathcal{X} . Given a vector of strategies $S = (S_1, \dots, S_n)$ for the players from $\mathcal{X} = \{1, \dots, n\}$, a resource y is matched to the winner x of y , that is, the most preferred player who proposes to y . The goal of each player $x \in \mathcal{X}$ is to maximize the total payoff of the resources that she wins. More formally, given a strategy vector S , let $T_x(S) \subseteq S_x$ be the set of resources that agent x wins. The goal of each player x is to maximize $\sum_{y \in T_x(S)} p_x(y)$.

Matroid Two-Sided Markets. A matroid two-sided market is a many-to-one two-sided market in which for each player x , the family \mathcal{F}_x of subsets of resources corresponds to the independent sets of a *matroid*. In other words, in a matroid two-sided market for every player $x \in \mathcal{X}$, the set system $(\mathcal{Y}, \mathcal{F}_x)$ is a matroid. This means, that for every player

	m_1	m_2	m_3	\dots	m_{n-2}	m_{n-1}	m_n
w_1	1	2	3	\dots	$n-2$	$n-1$	n
w_2	n	1	2	\dots	$n-3$	$n-2$	$n-1$
w_3	$n-1$	n	1	\dots	$n-4$	$n-3$	$n-2$
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
w_{n-1}	3	4	5	\dots	n	1	2
w_n	2	3	4	\dots	$n-1$	n	1

Figure 1: The weights of the edges in our construction.

$x \in \mathcal{X}$ it holds that: $\emptyset \in \mathcal{F}_x$; if $A \in \mathcal{F}_x$ and $B \subseteq A$, then also $B \in \mathcal{F}_x$; and if $A, B \in \mathcal{F}_x$ with $|A| < |B|$, then there must be a $b \in B$ such that $A \cup \{b\} \in \mathcal{F}_x$. Such matroid two-sided markets arise naturally if, for example, every employer is interested in hiring a fixed number of workers or if the workers can be partitioned into different classes and a certain number of workers from each class is to be hired. We define *correlated matroid two-sided markets* analogously to the singleton case, that is, there is a payoff $p_{x,y} \in \mathbb{N}$ associated with every pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $p_x(y) = p_y(x) = p_{x,y}$.

3 Better Response Dynamics

In this section, we consider the random better response dynamics and present instances for which it takes with high probability an exponential number of steps to reach a stable matching. This contrasts with the result due to Roth and Vande Vate [20] that from every matching there exists a polynomial sequence of better responses leading to a stable matching. In the following, we assume that $\mathcal{X} = \{w_1, \dots, w_n\}$ and $\mathcal{Y} = \{m_1, \dots, m_n\}$ for some $n \in \mathbb{N}$. We present our instances using an edge-weighted bipartite graph with an edge for each pair of woman and man. A woman w prefers a man m to a man m' if the weight of the edge $\{w, m\}$ is *smaller* than the weight of $\{w, m'\}$. On the other hand, a man m prefers a woman w to a woman w' if the weight of the edge $\{m, w\}$ is *larger* than the weight of the edge $\{m, w'\}$. The weights of the edges in the bipartite graph are depicted in Figure 1. Before we analyze the number of better responses needed to reach a stable matching, we prove a structural property of the instances we construct.

Lemma 1. *For the family of two-sided markets that is depicted in Figure 1, a matching M is stable if and only if it is perfect and every woman has the same payoff in M .*

Proof. First we show that every perfect matching M in which every woman has the same payoff is stable. One crucial property of our construction is that whenever a woman w and a man m are matched to each other, the sum $p_w(m) + p_m(w)$ of their payoffs is $n - 1$. In order to see this, assume that the edge between w and m has weight $l + 1$. Then there are l men whom woman w prefers to m , i.e., $p_w(m) = n - 1 - l$. Furthermore, there are $n - 1 - l$ women whom man m prefers to w , i.e., $p_m(w) = l$. This implies $p_w(m) + p_m(w) = n - 1$, regardless of l . We consider the case that every woman has payoff k and hence every man has a payoff of $n - 1 - k$ in M . Assume that there exists a blocking pair (w, m) . Currently w has payoff k , m has payoff $n - 1 - k$, and w and

m are not matched to each other. Since (w, m) is a blocking pair, $p_w(m) > k$ and hence $p_m(w) = n - 1 - p_w(m) < n - 1 - k = p_m(M)$, contradicting the assumption that (w, m) is a blocking pair. This implies that every state in which all women have the same payoff is stable.

Now we have to show that a state M in which not every woman has the same payoff cannot be a stable matching. We can assume that M is a perfect matching as otherwise it obviously cannot be stable. Let M be a perfect matching and define $l(M)$ to be the lowest payoff that one of the women receives, i.e., $l(M) = \min\{p_w(M) \mid w \in \mathcal{X}\}$. Furthermore, by $L(M)$ we denote the set of women receiving payoff $l(M)$, i.e., $L(M) = \{w \in \mathcal{X} \mid p_w(M) = l(M)\}$. We claim that there exists at least one woman in $L(M)$ who forms a blocking pair with one of the men.

First we consider the case that the lowest payoff is unique, i.e., $L(M) = \{w\}$. Let m be the man with $p_w(m) = l(M) + 1$. We claim that (w, m) is a blocking pair. To see this, let M' denote the matching obtained from M by resolving (w, m) . We have to show that the payoff $p_m(M)$ of man m in matching M is smaller than his payoff $p_m(M')$ in M' . Due to our construction $p_m(M') = n - 1 - p_w(m)$ and $p_m(M) = n - 1 - p_{w'}(m)$, where w' denotes m 's partner in M . Due to our assumption, w is the unique woman with the lowest payoff in M . Hence, $p_{w'}(m) = p_{w'}(M) > p_w(M) = p_w(m) - 1$. This implies $p_m(M') \geq p_m(M)$, which in turn implies $p_m(M') > p_m(M)$ since $w \neq w'$, and hence, (w, m) is a blocking pair.

It remains to consider the case that the woman with the lowest payoff is not unique. We claim that also in this case we can identify one woman in $L(M)$ who forms a blocking pair. Let $w^{(1)} \in L(M)$ be chosen arbitrarily and let $m^{(1)}$ denote her partner in M . Let $m^{(2)}$ denote the man with $p_{w^{(1)}}(m^{(2)}) = p_{w^{(1)}}(m^{(1)}) + 1$ and let $w^{(2)}$ denote the woman matched to $m^{(2)}$ in M . If the payoff of $w^{(2)}$ in M is larger than the payoff of $w^{(1)}$ in M , then by the same arguments as for the case $|L(M)| = 1$ it follows that $(w^{(1)}, m^{(2)})$ is a blocking pair. Otherwise, if $p_{w^{(1)}}(M) = p_{w^{(2)}}(M)$, we continue our construction with $w^{(2)}$. To be more precise, we choose the man $m^{(3)}$ with $p_{w^{(2)}}(m^{(3)}) = p_{w^{(2)}}(m^{(2)}) + 1$ and denote by $w^{(3)}$ his partner in M . Again either $w^{(3)} \in L(M)$ or $(w^{(2)}, m^{(3)})$ is a blocking pair. In the former case, we continue the process analogously, yielding a sequence $m^{(1)}, m^{(2)}, m^{(3)}, \dots$ of men. If the sequence is finite, a blocking pair exists. Now we consider the case that the sequence is not finite. Let $j \in \{1, \dots, n\}$ be chosen such that $m^{(1)} = m_j$. Due to the weights shown in Figure 1, it holds $m^{(i)} = m_{(j-i \bmod n)+1}$ for $i \in \mathbb{N}$. Hence, in this case, every man appears in the sequence, and hence, every woman has the same payoff $l(M)$. \square

Let us quote a couple of Chernoff bounds from [16] that we will use in the proofs below. Theorem 4.1 of [16] states that for independent Poisson trials X_1, \dots, X_n , with $\Pr[X_i = 1] = p_i$ for $i \in \{1, \dots, n\}$ and $0 < p_i < 1$, and for $X = \sum_i X_i$, $\mu = \mathbf{E}[X] = \sum_i p_i$, and any $\delta > 0$, we have

$$\Pr[X > (1 + \delta)\mu] < \left[\frac{e^\delta}{(1 + \delta)^{(1+\delta)}} \right]^\mu. \quad (1)$$

Theorem 4.2 of [16] states that for $0 < \delta \leq 1$,

$$\Pr[X < (1 - \delta)\mu] < \exp(-\mu\delta^2/2). \quad (2)$$

Now we can prove that with high probability the number of better responses needed to reach a stable matching is exponential.

Theorem 2. *There exists a family of two-sided markets I_1, I_2, I_3, \dots with corresponding matchings M_1, M_2, M_3, \dots such that, for $n \in \mathbb{N}$, I_n consists of n women and n men and a sequence of random better responses starting in M_n with probability $1 - 2^{-\Omega(n)}$ needs $2^{\Omega(n)}$ steps to reach a stable matching.*

Proof. We consider the instances shown in Figure 1. In Lemma 1, we have shown that in any stable matching all women have the same payoff. For a given matching M , we are interested in the most common payoff among the women and denote by $\chi(M)$ the number of women receiving this payoff, i.e.,

$$\chi(M) = \max_{i \in \{0, \dots, n-1\}} |\{w \in \mathcal{X} \mid p_w(M) = i\}| .$$

In the following, we show that whenever $\chi(M)$ is at least $15n/16$, then $\chi(M)$ is more likely to decrease than to increase. This yields a biased random walk that takes with high probability exponentially many steps to reach $\chi(M) = n$. If the most common payoff is unique, which is always the case if $\chi(M) > n/2$, then we denote by $\mathcal{X}'(M)$ the set of women receiving this payoff and by $\mathcal{Y}'(M)$ the set of men matched to women from $\mathcal{X}'(M)$.

Let $\delta = 15/16$ and assume that $\chi(M) \geq \delta n$. First, we consider the case that the current matching M is not perfect, i.e., there exists at least one unmatched woman w and at least one unmatched man m . We call a blocking pair *good* if for the matching M' obtained from resolving it, $\chi(M') \leq \chi(M) - 1$. On the other hand, we call a blocking pair *bad* if $\chi(M') = \chi(M) + 1$ or if M' is a perfect matching. Let us count the number of good and of bad blocking pairs. Let k denote the most common payoff. Both the unmatched woman w and the unmatched man m form a blocking pair with each person who prefers her/him to his/her current partner. Since the current payoff of the women in $\mathcal{X}'(M)$ is k , at most k of these women do not improve their payoff by being matched to the unmatched man m . Analogously, since the payoff of the men in $\mathcal{Y}'(M)$ is $n - 1 - k$, at most $n - 1 - k$ of these men do not improve their payoff by being matched to the unmatched woman w . This implies that the number of good blocking pairs is at least¹ $\max\{\delta n - k, \delta n - n + 1 + k\} \geq (\delta - 1/2)n$. On the other hand, there can be at most $(1 - \delta)n + 1$ bad blocking pairs. This follows easily because only women from $\mathcal{X} \setminus \mathcal{X}'(M)$ can form bad blocking pairs and each of these women forms at most one bad blocking pair as there is only one man who is at position $n - k$ in her preference list. Furthermore, there exists at most one blocking pair that makes the matching perfect.

The aforementioned arguments show that for a matching M with $\chi(M) \geq \delta n$ and sufficiently large n , the ratio of good blocking pairs to bad blocking pairs is bounded from below by

$$\frac{(\delta - 1/2)n}{(1 - \delta)n + 1} \geq \frac{7}{2} .$$

¹Let us remark that it is also correct to use the sum of $\delta n - k$ and $\delta n - n + 1 + k$ instead of the maximum as a lower bound for the number of good blocking pairs. This leads to slightly different constants in the following arguments but does not make any qualitative difference.




Figure 2: Transition probabilities of the random walk.

This implies that the conditional probability of choosing a good blocking pair under the condition that either a good or a bad blocking pair is chosen is bounded from below by $7/9$.

If a good blocking pair is chosen, χ decreases by at least 1 and the matching obtained is again not perfect (otherwise the blocking pair would be bad by definition). If a bad blocking pair is chosen, χ increases by 1 or the matching obtained is perfect. In any other case, χ remains unchanged. If the matching obtained is perfect, after the next step again a matching M'' is obtained that is not perfect. For this matching M'' , we have $\chi(M'') \leq \chi(M) + 2$. Since we are interested in proving a lower bound, we can pessimistically assume that the current matching is not perfect and that whenever a bad blocking pair is chosen, χ increases by 2. Hence, we can obtain a lower bound on the number of better responses needed to reach a stable state, i.e., a state M with $\chi(M) = n$, by considering a random walk on the set $\{\lceil \delta n \rceil, \lceil \delta n \rceil + 1, \dots, n\}$ that starts at $\lceil \delta n \rceil$, terminates when it reaches n , and has the transition probabilities shown in Figure 2. (Note that if χ ever drops below $\lceil \delta n \rceil$ we wait for it to return to that value, and re-start the random walk.)

This is a biased random walk. If we start with an arbitrary matching M satisfying $\chi(M) \leq \delta n$, then, with probability $1 - 2^{-\Omega(n)}$, it takes $2^{\Omega(n)}$ steps to reach state n . To see this, let us consider an even simpler random walk \mathcal{W} on the set $\{0, 1, 2, \dots\}$: if the particle is at 0, it deterministically moves to 2; at any other position $i > 0$, it moves to $i - 1$ with probability $7/9$ and to $i + 2$ with probability $2/9$. If the particle starts at 0, the time it takes to reach $n - \lceil \delta n \rceil$ for the first time is a lower bound on the time it takes to reach a stable matching in the instance we constructed.

For $i \in \mathbb{N}$, let $t(i)$ denote the position of the particle in the random walk \mathcal{W} after i steps. For $i, j \in \mathbb{N}$ and $j < i$, let $\mathcal{F}_{i,j}$ denote the event that

$$t(i) \geq n - \lceil \delta n \rceil \quad \text{and} \quad t(j) = 0 \quad \text{and} \quad \forall k \in \{j + 1, \dots, i - 1\}: t(k) > 0.$$

If the particle is at position $n - \lceil \delta n \rceil$ after i steps, then for one $j \in \{0, \dots, i - (n - \lceil \delta n \rceil)/2\}$, the event $\mathcal{F}_{i,j}$ must occur. Let $R_{i,j}$ denote the number of times the particle increases its position by 2 during the steps $j + 1, \dots, i$, and let $L_{i,j}$ denote the number of times the particle decreases its position by 1. Then the event $\mathcal{F}_{i,j}$ can only occur if $2R_{i,j} - L_{i,j} \geq n - \lceil \delta n \rceil$. As $L_{i,j} = i - j - R_{i,j}$, this is equivalent to $3R_{i,j} \geq n - \lceil \delta n \rceil + i - j$, which can only occur if $R_{i,j} \geq (i - j)/3 = 3\mathbf{E}[R_{i,j}]/2$. As $\mathcal{F}_{i,j}$ can only occur if $i - j = \Omega(n)$, it follows by a direct application of (1) that $\Pr[\mathcal{F}_{i,j}] \leq 2^{-cn}$, for some constant c .

Now we can upper-bound the probability that the particle reaches $n - \lceil \delta n \rceil$ during the first z steps by taking a union bound over values of i, j of the probabilities of the $\mathcal{F}_{i,j}$:

$$\Pr[\exists i, j \in \{0, \dots, z\}: \mathcal{F}_{i,j}] \leq z^2 \cdot 2^{-cn}.$$

Hence, for $z = 2^{cn/3}$, the probability that the random walk reaches position $n - \lceil \delta n \rceil$ during the first z steps, is bounded from above by $2^{-cn/3}$, which concludes the proof. \square

4 Best Response Dynamics

In this section, we study the best response dynamics in two-sided markets. First we show that this dynamics can cycle. Let us remark again, that we use *women* to denote active agents and *men* to denote passive agents.

Theorem 3. *There exists a two-sided market with three women and three men in which the best response dynamics can cycle.*

Proof. We denote by w_1, w_2, w_3 the women and we denote by m_1, m_2, m_3 the men. We choose the following preference lists for women and men:

$$\begin{array}{c|ccc} w_1 & m_2 & m_3 & m_1 \\ w_2 & m_1 & m_2 & m_3 \\ w_3 & m_3 & m_1 & m_2 \end{array} \quad \begin{array}{c|ccc} m_1 & w_1 & w_3 & w_2 \\ m_2 & w_2 & w_1 & w_3 \\ m_3 & w_1 & w_2 & w_3 \end{array}$$

We describe a state by a triple (x, y, z) , meaning that the first woman is matched to the man m_x , the second woman to man m_y , and the third woman to man m_z . A value of -1 indicates that the corresponding woman is unmatched. The following sequence of states constitutes a cycle in the best response dynamics:

$$\begin{aligned} (-1, 2, 3) &\rightarrow (3, 2, -1) \rightarrow (3, 1, -1) \rightarrow (3, -1, 1) \\ &\rightarrow (2, -1, 1) \rightarrow (-1, 2, 1) \rightarrow (-1, 2, 3) . \end{aligned}$$

\square

Roth and Vande Vate [20] show that from every matching there exists a polynomial sequence of better responses leading to a stable matching. We show that this is also true for the best response dynamics. The following approach is not based on [20], and we cannot see any simple modification of their algorithm to restrict it to best responses.

Theorem 4. *For every two-sided market with n women and m men and every matching M , there exists a sequence of at most $2nm$ best responses starting in M and leading to a stable matching.*

Proof. We divide the sequence of best responses into two phases. In the first phase, only matched women are allowed to change their matches. If no matched woman can improve her payoff anymore, then the second phase starts. In the second phase, all women are allowed to play best responses in an arbitrary order. In the first phase, we use the potential function

$$\Phi(M) = \sum_{x \in X} (m - p_x(M)) ,$$

where X denotes the set of matched women. This potential function decreases with every best response of a matched woman by at least 1 because this woman increases her payoff and the set X can only become smaller. Since Φ is bounded from above by nm , the first

phase terminates after at most nm best responses in a state in which no matched woman can improve her payoff.

Now consider the second phase. We claim that if we start in a state M' in which no matched woman can improve her match, then every sequence of best responses terminates after at most nm steps in a stable matching. This will be established using the potential function consisting of the sum of all the men's payoffs. We will see that this potential is monotonically increasing.

Assume that we start in a state M' in which no matched woman can improve her match and that an unmatched woman plays a best response and gets matched with a man y , leading to state M'' . If y was not already matched, then the sum of the men's payoffs has increased, and it remains the case that no matched woman can improve on her match. Alternatively, if y was already matched, note first that y 's payoff has increased (or he would not have accepted the new match). Also we have maintained the property that no matched woman can improve on her match (the set of matched women has gained one member and lost one, and note that the members who have remained in that set are not able to propose successfully to y). In either case, note that matched men never become unmatched and therefore they can only improve their payoffs. With every best response one man increases his payoff by at least 1. This concludes the proof of the theorem as each of the m men can increase his payoff at most n times. \square

From the previous proof, the following corollary is immediate.

Corollary 5. *For every two-sided market with n women and m men and every matching M in which no matched woman can improve, every sequence of best responses starting in M has length at most nm . In particular, this is true if M is the empty matching.*

Let us remark that this result is not true for the better response dynamics, as from the empty matching any other matching M is reachable by a sequence of better responses.

Finally, we show that Theorem 2 is also valid for the random best response dynamics. The following result is in a sense weaker than Theorem 2 in that here, both the preference lists and the starting configuration need to be chosen by an adversary. (In Theorem 2 it was just the preference lists that are chosen by an adversary, and after that, nearly all initial matchings run into the same problem.)

Theorem 6. *There exists a family of two-sided markets I_1, I_2, I_3, \dots and corresponding matchings M_1, M_2, M_3, \dots such that, for $n \in \mathbb{N}$, I_n consists of n women and n men and a sequence of random best responses starting in M_n with probability $1 - 2^{-\Omega(n)}$ needs $2^{\Omega(n)}$ steps to reach a stable matching.*

Proof. We will assume throughout that 8 divides n , to avoid rounding when we refer to various multiples of $n/8$.

For every large enough $n \in \mathbb{N}$, we construct an instance I_n with n women and n men in which the preference lists and the initial state M_n are chosen as shown in Figure 3. That is, every woman w_i with $i \in \{2, \dots, n\}$ prefers man m_{i-1} to man m_i whom she prefers to every other man. Woman w_1 prefers the men $m_{7n/8}, \dots, m_{n-1}$ to man m_1 whom she prefers to every other man. Man m_1 prefers woman w_1 to woman w_2 whom he prefers to every other woman. Every man m_i with $i \in \{2, \dots, n-1\}$ prefers woman w_i to woman




Figure 3: Nodes in the upper and lower row correspond to women and men, respectively. The figure also shows the initial state and the preference lists. The lists are only partially defined, but they can be completed arbitrarily. This matching belongs to \mathcal{M} , where the right-hand block is empty, i.e. $l = n$.




Figure 4: Matching from \mathcal{M} .

w_{i+1} whom he prefers to woman w_1 whom he prefers to every other woman. Man m_n prefers woman w_n to all other women.

Let \mathcal{M} denote the set of matchings that contain the edges

$$(w_1, m_1), \dots, (w_{j-2}, m_{j-2}), (w_j, m_{j-1}), \dots, (w_k, m_{k-1}), \\ (w_{k+1}, m_{k+1}), \dots, (w_l, m_l), (w_{l+2}, m_{l+1}), \dots, (w_n, m_{n-1})$$

for some $j < k < l$ with $n/16 \leq k - j \leq n/4$, $k < n/4$, and $l \geq 5n/8$ (cf. Figure 4). We claim that if one starts in a matching that belongs to \mathcal{M} , then with probability $1 - 2^{-cn}$, for an appropriate constant $c > 0$, another matching from \mathcal{M} is reached after $\Theta(n)$ many steps. Since no matching from \mathcal{M} is stable, this implies the theorem.

If the current matching belongs to \mathcal{M} , then there are at most three women who are able to change the matching via a best response. Woman w_{j-1} can propose to man m_{j-1} , woman w_{k+1} can propose to man m_k , and, if $l < n$, woman w_{l+1} can propose to man m_{l+1} .

Define a *block* to be a maximal-length sequence of matches of the form $(w_r, m_{r-1}), \dots, (w_s, m_{s-1})$, where $1 \leq r \leq s \leq n$. When a change occurs to the matching, the effect is to either

- cause one end of a block to move 1 to the right, or
- to eliminate a block of length 1, or to merge two blocks.



Figure 5: w_1 proposes to m_k if $\frac{7n}{8} \leq k < n$.




Figure 6: A new diagonal is introduced.

We continue by showing two things. First, changes of the second kind are (exponentially) unlikely to take place. Second, by the time the old block leaves the right-hand end of the sequence, a new block has been generated at the left-hand side of the sequence with an expected length of $n/8$.

In \mathcal{M} there is a block of length $\Omega(n)$, possibly another block to its right, and a gap of length $\Omega(n)$ between them. When a change of the first kind above takes place, the midpoint of the affected block moves $1/2$ to the right. Hence any block is altered at most $2n$ times by these changes. On the i -th occasion that the midpoint of the left-hand block moves to the right, associate with that event a Bernoulli trial X_i that is 1 if the left-hand end of the block moves, and 0 if the right-hand end moves. Each value has probability $1/2$. Let X be the sum of the X_i . Provided that the block has length $\Omega(n)$, the probability that the block vanishes before reaching the right-hand side corresponds to X exceeding its expected value by some factor greater than 1, which by (1) is exponentially small in n . By a similar argument, the gap between the two blocks is equally unlikely to vanish (which would cause the blocks to merge).

Next we analyse the mechanism whereby new blocks are generated at the left-hand side of the gadget at the time that an existing block is exiting the right-hand side — a process that is designed to (with high probability) return us to a matching in \mathcal{M} . When the right end of the first block has reached man $m_{7n/8}$, i.e., $m_{7n/8}$ is unmatched, then with probability exponentially close to 1, the second block has already vanished (see Figure 5) because the initial distance between the two blocks is at least $3n/8$ and only with probability $2^{-\Omega(n)}$ does it decrease to $n/8$ before the second block vanishes (applying (1) as we did above). Now consider the case that the second block has vanished and the right end of the first block lies in the interval $\{7n/8, \dots, n-1\}$. Then, woman w_1 has an incentive to change her match since she prefers m_k with $k \in \{7n/8, \dots, n-1\}$ to m_1 . Once she has changed her strategy, a new block of diagonals can be created on the left end of the gadget (see Figure 6). In particular, woman w_1 will only return to m_1 if no man m_k with $k \in \{7n/8, \dots, n-1\}$ is unmatched, that is, she will only return to m_1 if the right end of the first block has reached man m_n . At this stage, at any step there is equal probability that the new block is extended by 1, or that the right-hand block's right-hand side moves 1 to the right. Hence, when that right-hand block has reached m_{n-1} , the length of the new block (to the left) has the distribution of the random variable X of Lemma 7. (Each time the right-hand block moves one to the right, the left-hand block's length increase is a geometric random variable with parameter $1/2$.) By Lemma 7 it follows that the length of the new block lies with high probability in the interval $[n/16, n/4]$. Only with exponentially small probability the left end of the block has not passed man $m_{5n/8}$ when the right end has reached man m_n because this would imply that the length of the block has increased from at most $n/4$ to $3n/8$. If none of these exponentially unlikely failures events occurs, we are again in a matching from \mathcal{M} . \square

In the following lemma, we use the notion of a *geometric random variable with parameter* $1/2$. Such a random variable Z describes, in a sequence of Bernoulli trials with success probability $1/2$, the number of failures before the first success is obtained, that is, for $i \in \{0, 1, 2, \dots\}$, $\Pr[Z = i] = (1/2)^{i+1}$.

Lemma 7. *Assume n is a multiple of 8. Let X be the sum of $n/8$ geometric random variables with parameter $p = 1/2$. There exists a constant $c > 0$ such that*

$$\Pr[X \notin [n/16, n/4]] \leq 2e^{-cn} .$$

Proof. For a series of independent Bernoulli trials with success probability $1/2$, the random variable X describes the number of failures before the $(n/8)$ -th success is obtained. For $a \in \mathbb{N}$, let Y_a be a binomially distributed random variable with parameters a and $1/2$. Now, suppose we generate an observation of X by flipping a coin until we obtain $n/8$ successes (and X is the number of failures). If the observed value of X is greater than $n/4$, then the first $3n/8$ coin flips contained less than $n/8$ successes. Hence observe that

$$\Pr[X > n/4] = \Pr[Y_{3n/8} < n/8] .$$

This in turn is equal to

$$\Pr\left[Y_{3n/8} < \frac{2}{3}\mathbf{E}[Y_{3n/8}]\right] \leq e^{-cn} ,$$

where the last inequality follows, for an appropriate constant $c > 0$, from (2). Furthermore, by a similar argument

$$\begin{aligned} \Pr[X < n/16] &= \Pr[Y_{3n/16} > n/8] \\ &= \Pr\left[Y_{3n/16} > \frac{4}{3}\mathbf{E}[Y_{3n/16}]\right] \leq e^{-cn} . \end{aligned}$$

□

A natural question about the proof of Theorem 6 is whether it is really necessary to divide the agents into active and passive agents and to require the active agents to play only best responses. Here we sketch why both assumptions are necessary for the particular construction in the proof. Assume that instead we consider the random better response dynamics or a symmetric best response dynamics in which one agent from $\mathcal{X} \cup \mathcal{Y}$ is chosen uniformly at random and allowed to play a best response. As long as the right-hand end of the initial block of diagonals has not reached man $m_{7n/8}$, the better, best, and symmetric best response dynamics behave identically.

However, the generation of the new block of diagonals does not work as desired for the better and the symmetric best response dynamics. To see this assume that we have reached a situation in which the first couple of diagonals have been inserted at the left-hand end. Then either w_1 is unmatched or she is matched to a man m_i with $i \in \{7n/8, \dots, n-1\}$. In the latter case it takes in expectation a constant number of steps until w_1 is again unmatched. As soon as w_1 is unmatched, the best response of m_1 is to get matched to w_1 . Hence, both in the better response dynamics and in the symmetric best response dynamics

the blocking pair $\{w_1, m_1\}$ will be resolved after a constant number of steps. This means that a new vertical edge is inserted at the beginning, which removes the leftmost diagonal. From that moment on, the block of diagonals can shrink as new vertical edges can replace the diagonal edges on its left-hand end. Hence, the block has a significant probability of vanishing before it reaches $m_{7n/8}$. At the same time another block of diagonals can be generated so that after the initial block of diagonals has reached the right-hand end, there is not only a single new block, but multiple disjoint blocks whose total length will likely be much smaller than $\Omega(n)$.

5 Correlated Two-Sided Markets

In this section, we show that, in contrast to general two-sided markets, the convergence time of the random better and best response dynamics in correlated two-sided markets is polynomial.

In the context of a 2-sided market, a *strategy* of a player is the (possibly empty) set of members of the opposite side that she is matched with. A *strategy vector* is the set of all the players' strategies, associated with a particular matching. In a *potential game* there is a function Φ that maps strategy vectors to the real numbers, in such a way that a better response by any player results in a new strategy vector with a smaller value of Φ . The existence of such a function Φ is a sufficient condition for termination of sequences of better responses, and in addition, one can often study the rate of decrease of Φ to obtain quantitative bounds on the number of steps required.

In [2], it is shown that correlated two-sided markets are potential games. Correlated two-sided markets have previously been considered by Abraham et al. [1], Lebedev et al. [13], and Mathieu [14]. In the latter two publications these markets are defined in a different way and they are called *acyclic*. It is, however, shown by Abraham et al. that the classes of acyclic and of correlated markets coincide. Lebedev et al. [13] prove that every correlated market has a unique stable matching, that the best response dynamics cannot cycle, and that from every state a short sequence of best responses to a Nash equilibrium exists. They also conclude that the random best response dynamics converges quickly (the expected number of steps is at most quadratic in the number of players). Finally, let us mention that Mathieu [14] also presented an exponential lower bound on the convergence time if an adversary selects the next player to play a best response.

5.1 Correlated Matroid Two-Sided Markets

In matroid two-sided markets, we consider a restricted class of better responses, so-called *lazy better* responses, introduced in [3]. Let a vector of strategies $S = (S_1, \dots, S_n)$ be given and denote by $S \oplus S_x^*$ for $S_x^* \in \mathcal{F}_x$ the state S except that player x plays S_x^* instead of S_x . Assume that a player $x \in \mathcal{X}$ plays a better response and changes her strategy from S_x to S'_x . We call this better response *lazy* if it can be decomposed into a sequence of strategies $S_x = S_x^0, S_x^1, \dots, S_x^k = S'_x$ such that $|S_x^{i+1} \setminus S_x^i| = 1$ and the payoff of player x in state $S \oplus S_x^{i+1}$ is strictly larger than her payoff in state $S \oplus S_x^i$ for all $i \in \{0, \dots, k-1\}$. That is, a lazy better response can be decomposed into a sequence of additions and exchanges of single resources such that each step strictly increases the payoff of the corresponding

player.

In [3], it is observed that for matroid strategy spaces, there does always exist a best response that is lazy. In particular, the best response that exchanges the least number of resources is lazy, and in singleton games, every better response is lazy. On the other hand, the following example demonstrates that even in simple matroid markets there can be non-lazy best responses: Let $\mathcal{X} = \{x\}$, $\mathcal{Y} = \{y_1, y_2, y_3, y_4\}$, and $\mathcal{F}_x = \{S \subset \mathcal{Y} \mid |S| \leq 2\}$. Then $(\mathcal{Y}, \mathcal{F}_x)$ is a so-called uniform matroid. Let $p_x(y_1) = 1$, $p_x(y_2) = p_x(y_3) = 2$, and $p_x(y_4) = 3$, and let $\{y_1, y_2\}$ be the current strategy of player x . Then $\{y_3, y_4\}$ is a best response that is not lazy. The only lazy best response is $\{y_2, y_4\}$.

In [2], it is shown that correlated matroid two-sided markets are potential games with respect to the lazy better response dynamics. Furthermore, it is shown that the restriction to lazy better responses is necessary as even the best response dynamics can cycle in correlated matroid two-sided markets.

Theorem 8. *In correlated matroid two-sided markets, the random lazy best response dynamics converges to a stable matching in expected polynomial time.*

Proof. Let p denote the highest possible payoff that can be achieved. As long as no pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ with $p_y(x) = p$ is contained in the matching, there exists one player whose best response would result in such a pair. This follows since players allocate optimal bases and an optimal basis of a matroid must contain the most valuable element. Since this player is allowed to play a best response with probability at least $1/n$ in each step, it takes $O(n)$ best responses until a pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ with $p_y(x) = p$ is contained in the matching in expectation.

After an edge $(x, y) \in \mathcal{X} \times \mathcal{Y}$ with $p_y(x) = p$ is contained in the matching, player x will never leave resource y again because she only plays lazy best responses. Furthermore, x cannot be displaced from y since no player is strictly preferred to x by resource y . Hence, the assignment of x to y can be fixed and we can modify the strategy space of x by contracting its matroid by removing y . By this contraction, we obtain another matroid two-sided market in which the rank of x 's matroid is decreased by 1. Now we can inductively apply the same argument to this game. \square

Theorem 9. *In correlated two-sided markets (where each agent is matched to at most one other agent) the best and better response dynamics converge in expected polynomial time.*

Proof. At any point in the process, let p be the highest payoff that can be obtained by players who are not already making a best response. There is a probability at least $1/n^2$ that a player who is not making a best response, will be chosen and will make a best response, with payoff p . When this happens, that player (and her chosen resource) become inactive, and since there are n players, this gives us a polynomial expected time taken for all players to become inactive, at which point the process has converged. \square

6 Acknowledgments

We thank Fuhito Kojima and Yuval Peres for pointing out related work. We thank the referees for proposing improvements to the presentation of these results.

References

- [1] David J. Abraham, Ariel Levavi, David Manlove, and Gregg O'Malley. The stable roommates problem with globally ranked pairs. *Internet Mathematics*, 5(4):493–515, 2008.
- [2] Heiner Ackermann, Paul W. Goldberg, Vahab S. Mirrokni, Heiko Röglin, and Berthold Vöcking. A unified approach to congestion games and two-sided markets. *Internet Mathematics*, 5(4):439–457, 2008.
- [3] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. *Theoretical Computer Science*, 410(17):1552–1563, 2009.
- [4] Federico Echenique and Jorge Oviedo. A theory of stability in many-to-many matching markets. *Theoretical Economics*, 1(2):233–273, 2006.
- [5] Tamás Fleiner. A fixed-point approach to stable matchings and some applications. *Mathematics of Operations Research*, 28(1):103–126, 2003.
- [6] Lisa Fleischer, Michel X. Goemans, Vahab S. Mirrokni, and Maxim Sviridenko. Tight approximation algorithms for maximum general assignment problems. In *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 611–620, 2006.
- [7] David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
- [8] Michel X. Goemans, Erran L. Li, Vahab S. Mirrokni, and Marina Thottan. Market sharing games applied to content distribution in ad-hoc networks. *IEEE Journal on Selected Areas in Communications*, 24(5):1020–1033, 2006.
- [9] Dan Gusfield and Robert W. Irving. *The Stable Marriage Problem: Structure and Algorithms*. MIT Press, 1989.
- [10] Christopher Hoffman, Alexander E. Holroyd, and Yuval Peres. A stable marriage of Poisson and Lebesgue. *Annals of Probability*, 34(4):1241–1272, 2006.
- [11] Alexander S. Kelso and Vincent P. Crawford. Job matchings, coalition formation, and gross substitute. *Econometrica*, 50:1483–1504, 1982.
- [12] Donald E. Knuth. *Mariages stables et leurs relations avec d'autres problèmes combinatoires*. Les Presses de l'Université de Montréal, 1976.
- [13] Dmitry Lebedev, Fabien Mathieu, Laurent Viennot, Anh-Tuan Gai, Julien Reynier, and Fabien de Montgolfier. On using matching theory to understand p2p network design. *CoRR*, abs/cs/0612108, 2006.
- [14] Fabien Mathieu. Upper bounds for stabilization in acyclic preference-based systems. In *Proceedings of the 9th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 372–382, 2007.

- [15] Vahab S. Mirrokni. *Approximation Algorithms for Distributed and Selfish Agents*. PhD thesis, Massachusetts Institute of Technology, 2005.
- [16] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- [17] Alvin E. Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, 92:991–1016, 1984.
- [18] Alvin E. Roth. The national residency matching program as a labor market. *Journal of American Medical Association*, 275(13):1054–1056, 1996.
- [19] Alvin E. Roth and Marilda A. Oliveira Sotomayor. *Two-sided Matching: A study in game-theoretic modeling and analysis*. Cambridge University Press, 1990.
- [20] Alvin E. Roth and John H. Vande Vate. Random paths to stability in two-sided matching. *Econometrica*, 58(6):1475–1480, 1990.