Aleksei Yu. Kondratev, Alexander S. Nesterov

RANDOM PATHS TO POPULARITY IN TWO-SIDED MATCHING

Random Paths To Popularity In Two-Sided Matching

Aleksei Yu. Kondratev[1,2] and Alexander S. Nesterov[3]

[1] Higher School of Economics, Ulitsa Soyuza Pechatnikov 16, 190121 St.Petersburg, Russia
akondratev@hse.ru

[2] St.Petersburg Institute for Economics and Mathematics RAS, Ulitsa Serpuhovkaya 38, 190013 St.Petersburg, Russia
kondratev.aleksei@yandex.ru

[3] Higher School of Economics, Ulitsa Soyuza Pechatnikov 16, 190121 St.Petersburg, Russia
asnesterov@hse.ru

Abstract

We study practically relevant aspects of popularity in two-sided matching where only one side has preferences. A matching is called *popular* if there does not exist another matching that is preferred by a simple majority. We show that for a matching to be popular it is necessary and sufficient that no coalition of size up to 3 decides to exchange their houses by simple majority. We then constructively show that a market where such coalitions meet at random converges to a popular matching whenever it exists.

Keywords. two-sided matching, popular matching, random paths, house allocation, assignment problem

1 Introduction

Various real-life economic situations can be modeled as two-sided matching markets where agents have preferences over indivisible objects and such that each agent gets at most one object. These situations include housing markets, assigning students to primary schools, job placement for graduates, and so forth.

Among different notions of efficiency for these matching markets, recent literature highlights the concept of *popularity*. A matching is called *popular* if majority of agents weakly prefers it over any other matching.[1] Popularity has mainly served as a normative property as it is a natural non-Paretian selection from a (possibly very large) set of efficient matchings.

The seminal paper by [1] that introduced popularity for the house allocation problem proposed a simple characterization of popular matchings. A matching is popular if and only if (1) no agent gets what we call *a bad house*, that is each agent gets either his best house among all houses (called his *first house*), or the best house among all houses that are not someone's best (called his *second house*), and (2) all first houses are allocated among agents that deem them as the best.[2]

To better understand the concept of popularity consider the following example.

Example 1. Let there be $n = 4$ agents $A = \{1, 2, 3, 4\}$ that have the following preferences (Table 1) over $m = 4$ houses $H = \{a, b, c, d\}$:

---

[1] One can also see popular matchings as weak Condorcet winners in a voting problem where the candidates are all possible matchings.

[2] Note that no agent should get a house that is worse than his first house but better than his second house as each such house is the first house for some other agent.

Table 1. Preference profile and matching μ_1

1	2	3	4
<u>a</u>	a	<u>d</u>	d
d	<u>c</u>	b	<u>b</u>
b	b	c	a
c	d	a	c

In this profile, the set of first houses is $FH = \{a, d\}$ and set of second houses is $SH = \{b, c\}$. Hence, by characterization in [1] there are only two popular matchings $\mu_1 = (1a, 2c, 3d, 4b)$ and $\mu_2 = (1a, 2c, 3b, 4d)$. Each other feasible matching either assigns some agent his bad house, or does not distribute first houses among agents that prefer them most (or both, as in matching $\mu_3 = (1d, 2a, 3b, 4c)$ where agent 4 gets his bad house c and the first house d is not assigned to agents 3 or 4 that value it the most).

The subsequent literature (see [7, 15] and the subsection below) focused mainly on issues relevant for centralized markets. In contrast to that, in this paper we shift the attention to popularity in decentralized markets.

Our contribution, which we describe below in detail, is three-fold: we provide a new characterization of popular matchings based on popularity within small groups of agents, propose a new efficient algorithm to find a popular matching, and use this algorithm to show that the sequence of "popular" exchanges in random small groups converges to a popular matching.

Our first result (Theorem 1) shows that a popular matching only needs to be popular locally: the matching is popular if and only if no group of up to three agents decides (by simple majority) to exchange their houses, keeping the matching of other agents intact.[3] The original characterization in [1] directly follows from our result (Corollary 1).

For an illustration consider again Example 1 and a popular matching $\mu_1 = (1a, 2c, 3d, 4b)$. We need to check that in each triple of agents, when we only consider the houses owned by this triple, then each such (reduced) matching is popular within this triple. The following Table 2 illustrates the popularity within each triple; each reduced matching is popular by characterization in [1] as each agent either gets his first house or his second house, and each first house is given to the agent that prefers it the most.

Table 2. Popular matching μ_1 reduced to each possible triple of agents

1	2	3		1	2	4		1	3	4		2	3	4
<u>a</u>	a	<u>d</u>		<u>a</u>	a	<u>b</u>		<u>a</u>	<u>d</u>	d		<u>c</u>	<u>d</u>	d
d	<u>c</u>	c		b	<u>c</u>	a		d	b	<u>b</u>		b	b	<u>b</u>
c	d	a		c	b	c		b	a	a		d	c	c

Our second result (Theorem 2) is an algorithm to find a popular matching. The algorithm begins with an arbitrary matching and then modifies it using local exchanges among one, two or three agents. The algorithm has two parts. In the first part the algorithm assigns each agent his first house unless this house already belongs to an agent that prefers it the most. In the second part the algorithm (forcibly) gives each agent who has a bad house his second house, and the owner of that house gets his first house. We illustrate how the algorithm works using the same preference profile from Example 1.

Let the initial matching be $\mu = (1b, 2d, 3a, 4c)$. We now show how the first part of the algorithm turns μ into μ' pictured below (Table 3), and the second part of the algorithm turns μ' into μ''.

---

[3]This result can also be interpreted from the cooperative standpoint. If for each coalition we take the majority rule as the solution concept, then for a matching to be in the core it is enough to check coalitions of size up to three. The analogous result for the marriage market states that the set of pairwise stable matchings coincides with the core [23].

Table 3. Initial matching μ, matching μ', and matching μ''

1	2	3	4
a	a	d	d
d	c	b	b
<u>b</u>	b	c	a
c	<u>d</u>	<u>a</u>	<u>c</u>

1	2	3	4
<u>a</u>	a	<u>d</u>	d
d	c	b	b
b	<u>b</u>	c	a
c	d	a	<u>c</u>

1	2	3	4
<u>a</u>	a	d	<u>d</u>
d	<u>c</u>	<u>b</u>	b
b	b	c	a
c	d	a	c

First part of the algorithm. Agent 1 owns house b, his first house is a owned by agent 3 whose first house is d. We implement the following exchange: agent 1 gets house a, agent 3 gets house d, and house d's previous owner agent 2 gets the leftover house b. In the resulting matching $\mu' = (1a, 2b, 3d, 4c)$ all first houses are assigned to agents that prefer them the most and thus the first part of the algorithm is completed.

Second part of the algorithm. After first part agents 1 and 3 get their first houses, agents 2 and 4 get bad houses. We take an arbitrary agent with a *bad* house, e.g. agent 2, and give him his *second* house, c, and the owner of that house c, agent 4, gets her *first* house, d. The previous owner of house d agent 3 gets the leftover house b, which is his second house. In the resulting matching $\mu'' = (1a, 2c, 3b, 4d)$ no agent gets a bad house, and each first house is assigned to an agent that prefers it the most. Thus we arrived to a popular matching and the second part of the algorithm is completed.

Despite being greedy our algorithm is computationally efficient. The speed of the algorithm is quadratic in the number of agents, as is the speed of the algorithm in [1].[4]

Importantly, the algorithm employs only exchanges in small groups and each exchange makes the matching within this group (reduced) popular. Indeed, in the first part of the algorithm agents 1,2 and 3 end up in a (reduced) popular matching μ' and they prefer this matching over the original matching μ. In the second part of the algorithm agents 2,3 and 4 also end up in a (reduced) popular matching μ'' and the majority of them, agent 2 and agent 4, prefer this matching over the previous matching μ'.

Based on this result our paper also suggests a positive rationale behind popularity: we show that arbitrary locally popular improvements as in our algorithm lead to a globally popular matching. Thus one may expect that popularity is likely to be eventually observed in realistic situations. Specifically, we consider a decentralized market where agents meet in arbitrary groups and exchange their houses when this is supported by majority of them. Our third result (Corollary 2) shows that this market eventually converges to a popular matching whenever it exists.

This finding is analogous to the result in about convergence in a marriage market. There, one matching is modified locally by a blocking pair of a man and a woman that prefer each other over their current matches. As this man and this woman match, their previous partners become unmatched, and these changes constitute a new matching. Then a new blocking pair is considered, a new matching is formed, and so forth, and [24] show that the sequence of these matchings lead to a stable matching.

Another closely related paper is [2] that considers the popularity-improvement paths from an arbitrary matching. The main finding is that, given a popular matching exists, it can be attained by at most two steps using an efficient algorithm.

1.1 Background

The house allocation problem where agents exchange indivisible objects (houses) without money was first introduced in [25], the assignment problem where all houses are initially commonly owned was first studied in [11]. The concept of popularity was first introduced by

---

[4]The paper [1] proposes an $O(n + m')$ algorithm, where m' is the total length of all preferences, i.e. up to $m' = |A| \cdot |H|$, where $|A|$ is the number of agents and $|H|$ is the number of houses.

[10] for the marriage problem [9], where popularity coincides with stability, and was applied to house allocation problem only recently by [1]. The characterization in [1] also allows ties.[5]

Existence and multiplicity of popular matchings was studied from several sides. First, [17] shows that a popular matching is likely to exist whenever preferences are uniformly random and the number of houses is approximately 1.42 times larger than the number of agents. For settings where a popular matching does not exist, [14] studied how to minimally augment the preference profile so that the existence is guaranteed; this problem is, in general, NP-hard.

Another way to ensure popularity is to consider mixed matchings, i.e. lotteries over matchings, and a straightforward generalization of the popularity property; [13] show that a popular mixed matching always exists and propose an efficient algorithm to find one. The recent literature has studied compatibility of popularity in mixed matchings with various fairness and incentive properties ([4], [6]).

As an alternative approach, [18] proposes least-unpopularity criteria to find the "most" popular matching; finding his least-unpopular matchings is, in general, NP-hard.

The problem of counting the number of popular matchings has been addressed in [19] for the case of strict preferences and in [21] and [3] for the case of weak preferences. Popularity with agents having different weights has been studied in [20].

2 The Model

Let A be a set of agents and H be a (larger) set of houses, $|H| \geq |A|$. Each agent $a \in A$ is endowed with a strict preference relation $\succ_a$ over the set of houses $H \cup \{\emptyset\}$ (i.e. $\succ_a$ is a linear order), and a prefers each house $h \in H$ over having no house, $h \succ_a \emptyset$.[6] The collection of individual preferences of all agents $\succ = (\succ_a)_{a \in A}$ is referred to as **the preference profile**. The triple $(A, H, \succ)$ constitutes the two-sided matching problem (aka house allocation problem), or simply **a problem**. In what follows we assume that the sets A and H are fixed and the problem is given by the preference profile $\succ$.

A solution to the problem is **a matching** μ – a mapping from $A \cup H \cup \emptyset$ on itself: by definition agent $a \in A$ is said to be matched to a house $h \in H$ in matching μ if $\mu(a) = h$ and also $\mu(h) = a$. If some agent or house remain unmatched, we say that they are matched to $\emptyset$. Let $\mathcal{M}$ denote the set of all possible matchings.

For any two matchings $\mu, \mu' \in \mathcal{M}$ and a subset of agents $B \subset A$ define **pairwise comparison** $PC_B(\mu, \mu')$ as the number of agents in B that strictly prefer their house in μ over their house in μ'.

A matching $\mu \in \mathcal{M}$ is called **popular (among set A)** if there does not exist another matching $\mu' \in \mathcal{M}$ such that μ' is preferred over μ by simple majority within entire set of agents A: $PC_A(\mu', \mu) > PC_A(\mu, \mu')$.

For each agent a let us call his most preferred house in H as a's **first house**: $FH(a) = h$ such that for each $h' \in H$ and $h' \neq h$ it holds that $h \succ_a h'$. **The set of all first houses** is denoted as $FH(A) = \{FH(a)\}_{a \in A}$. For each house h let us call agents for whom h is the first house as h's **first agents**: $FA(h) = \{a \in A | h = FH(a)\}$.

For each agent a let us call his most preferred house among all non-first houses as a's **second house**: $SH(a) = h$ such that for each $h' \in H \setminus FH(A)$ and $h' \neq h$ it holds that $h \succ_a h'$. **The set of all second houses** is denoted as $SH(A) = \{SH(a)\}_{a \in A}$. For each house h let us call agents for whom h is the second house as h's **second agents**: $SA(h) = \{a \in A | h = SH(a)\}$.

---

[5] This setting was further generalized to the case with ties and matroid constraints by [12] and to the case with two-sided preferences and one-sided ties by [8] (the latter problem turns out to be NP-hard). The many-to-one matching problem, where each house has a capacity was studied in [26], and the many-to-many problem was studied by [22].

[6] All results remain true when agents have short preference lists with last resort.

Note that sets $FH(A)$ and $SH(A)$ are disjoint, i.e. no agent's second house can be a first house for any other agent.

3 Characterization of Popular Matching

Note that a matching cannot be popular if at least one agent is unmatched. Therefore throughout the paper we can focus only on full matchings, $\mu(A) \subset H$.

Our first main result characterizes the popular matching as a matching that is popular among each triple of agents.

For a profile $\succ$, we say that a matching μ is *popular among each three agents* if for each three agents $a, b, c \in A$ there does not exist a matching $\mu' \in \mathcal{M}$ same as μ for each other agent $a' \notin \{a, b, c\}$ $\mu'(a') = \mu(a')$ and such that it wins μ in pairwise comparison within this triple of agents $PC_{\{a,b,c\}}(\mu', \mu) > PC_{\{a,b,c\}}(\mu, \mu')$.

Theorem 1. A matching is popular if and only if it is popular among each three agents.

Proof. The "only if" part is straightforward: each popular matching μ is popular among each triple of agents. For a contradiction, assume that there is a triple of agents $a, b, c \in A$ and another matching μ' same as μ for all other agents and such that it is preferred over μ: $PC_{\{a,b,c\}}(\mu', \mu) > PC_{\{a,b,c\}}(\mu, \mu')$. Then μ cannot be popular among all agents since all other agents are indifferent and thus: $PC_A(\mu', \mu) - PC_A(\mu, \mu') = PC_{\{a,b,c\}}(\mu', \mu) - PC_{\{a,b,c\}}(\mu, \mu') > 0$.

The "if" part we also prove by contradiction. For a contradiction, assume that there is a matching μ that is popular among each triple of agents, but it loses in pairwise comparison to some other matching μ': $PC_A(\mu', \mu) > PC_A(\mu, \mu')$. Consider all agents that have different houses in these two matchings, denote the set of these agents as $A_1 = \{a \in A : \mu(a) \neq \mu'(a)\}$. (In what follows we will change the notation of these agents for convenience).

We partition all agents into those who participate in a trading cycle, i.e. exchange their matched houses among themselves, and trading chains, i.e. those that are matched in μ' to a previously empty house or whose house in μ becomes empty in μ'.

We first deal with chains. Consider an arbitrary agent $b_1 \in A_1$ that received a previously empty house $\mu'(b_1) \notin \mu(A)$, $\mu(\mu'(b_1)) = \emptyset$. If b_1's house is empty in μ', $\mu'(\mu(b_1)) = \emptyset$, then we get a chain of size 1. Otherwise there is some agent b_2 such that $\mu'(b_2) = \mu(b_1)$. If b_2's house is empty in μ', $\mu'(\mu(b_2)) = \emptyset$, then we get a chain of size 2. Otherwise, we continue in the same way until we find the last agent in the chain. Similarly, determine chains for each agent that receives a previously empty house. Denote the set of agents participating in a chain as B_1.

We then deal with cycles. Consider an arbitrary agent not from any chain $a_1 \in A_1 \setminus B_1$, $\mu(a_1) \neq \mu'(a_1)$. Consider agent a_2 that owns house $\mu'(a_1)$, $a_2 = \mu(\mu'(a_1))$. Agent a_2 also does not belong to any chain, $a_1 \in A_1 \setminus B_1$ and as $\mu(a_2) = \mu'(a_1)$, then $a_2 \neq a_1$. If the two agents just exchanged their houses, $\mu'(a_2) = \mu(a_1)$, then we get a trading cycle $(\mu(a_1), a_1, \mu'(a_1), a_2)$ of length 2. Otherwise, if $\mu'(a_2) \neq \mu(a_1)$, then consider agent $a_3 = \mu(\mu'(a_2))$. Since $\mu(a_3) = \mu'(a_2) \neq \mu(a_1)$, then $a_2 \neq a_3$, $a_1 \neq a_3$ и $a_3 \in A_1$.

And so forth until we get a cycle of length at least 2 and at most $|A_1 \setminus B_1|$. In the same way we find all trading cycles among all other agents.

Thus, the set A_1 and the set of corresponding houses $\mu(A_1) \cup \mu'(A_1)$ is partitioned into trading chains of size at least 1 and cycles of size at least 2.

By assumption $PC_A(\mu', \mu) > PC_A(\mu, \mu')$, there is at least one trading chain or one trading cycle such that more than half of its agents prefer μ' over μ. Formally, if A_{TC} denotes the set of agents in this chain or cycle, $PC_{A_{TC}}(\mu', \mu) > PC_{A_{TC}}(\mu, \mu')$.

If A_{TC} form a cycle, then we can find two neighbouring agents $i, j \in A_{TC}$, $j = \mu(\mu'(i))$, that both prefer μ' over μ. If this trading cycle is of length 2, then consider a new matching

μ'' that is identical to μ for each agent except $a = \{i,j\}$ and same as μ' for these pair $\mu''(a) = \mu'(a)$. Then by adding one other arbitrary agent we get a triple of agents that prefer μ'' over μ by majority – contrary to our premise. If this trading cycle is of length more than 2, then consider the next neighbouring agent $l = \mu(\mu'(j))$. Consider now a new matching μ'' that is identical to μ for each agent except $a = \{i,j,l\}$ and $\mu''(i) = \mu'(i)$, $\mu''(j) = \mu'(j)$, and $\mu''(l) = \mu(i)$. The triple of agents i,j,l prefers μ'' over μ by majority: $PC_{\{i,j,l\}}(\mu'',\mu) > PC_{\{i,j,l\}}(\mu,\mu'')$, contrary to our premise.

If A_{TC} forms a chain of length 1, $A_{TC} = \{a_1\}$, then consider a new matching μ'' constructed as before: μ'' is identical to μ for each agent except for a_1, $\mu''(a_1) = \mu'(a_1)$. A triple of agents a_1 and two arbitrary agents a_2, a_3 prefers μ'' over the original matching μ: $PC_{\{a_1,a_2,a_3\}}(\mu'',\mu) > PC_{\{a_1,a_2,a_3\}}(\mu,\mu'')$, contrary to our premise.

If A_{TC} forms a chain of length 2, then both agents in A_{TC} are better off in μ' compared to μ. By adding one other arbitrary agent we get a triple of agents that prefers a similarly constructed μ'' over μ by majority, contrary to our premise.

If the length of the chain is above 2, then either (1) we can find two neighbouring agents $i,j \in A_{TC}$, $j = \mu(\mu'(i))$, that both prefer μ' over μ, or (2) the chain begins and ends with agents that are better off in μ' compared to μ (and agents in between interchange). In case (1) we take the triple of these agents i,j and the previous owner of j's house $l = \mu(\mu'(j))$ (if j's house was empty, then take an arbitrary l). This triple i,j,l prefers a similarly constructed μ'' over μ by majority, contrary to our premise.

In case (2) we take the triple of agents as the first agent in the chain a_1, $\mu(\mu'(a_1)) = \emptyset$, the last agent a_k, $\mu'(\mu(a_k)) = \emptyset$, and the one before the last a_{k-1}. The triple a_1, a_{k-1}, a_k prefers a similarly constructed μ'' over μ by majority, contrary to our premise. ■

As an immediate corollary we get the characterization of popular matchings from [1].

Corollary 1. A matching is popular if and only if (1) each agent gets either his first house or his second house, and (2) each first house is matched with one of its first agents.

Proof. The "if" part is straightforward since it is enough to check only triples of agents. In each such triple only an agent a with a second house can become better off, but each better house $f \succ_a SH(a)$ is already matched to one of its first agents $b = \mu(f) \in FA(f)$, making a better off requires making b worse off, which cannot be supported by majority.

We prove the "only if" part by contradiction. Let condition (2) be violated: some first house f is not allocated to one of its first agents. Then each f's first agent $a \in FA(f)$, the owner of f $b = \mu(f)$ and the owner of b's first house $c = \mu(FH(b))$ form a triple for which μ is not popular.

Hence, in any popular matching, each agent gets his first house, second house, or a bad house.

Let condition (1) be violated: some agent a_1 gets a bad house t in matching μ, there is a triple of agents a_1, the owner of a_1's second house $a_2 = \mu(SH(a_1))$, and the owner of a_2's first house $a_3 = \mu(FH(a_2))$ for whom μ is not popular. ■

4 The Algorithm and Random Paths to Popularity

We represent the sequence of matchings as a finite Markov chain. The set space is the set of matchings $\mathcal{M}$. The transition probabilities between the states depend on how many agents become better off in one state compared to the other. Specifically, for each matching $\mu \in \mathcal{M}$ we consider all "neighbouring" matchings $\mu' \in \mathcal{M}$ that is matchings where at most three agents are matched to different house than in μ. If $k = 1,2,3$ agents are matched differently in μ and μ', then we say that μ and μ' are connected by a k-way exchange. If the k-way exchange makes more than half of these k agents better off, then the transition probability is positive, otherwise the transition probability is zero.

Next we present our second main result.

Theorem 2. Let μ be an arbitrary matching for $(A, H, \succ)$, $|A| = n$. Let a popular matching exist. Then there exists a finite sequence of matchings $\mu_0, \mu_1, \ldots, \mu_l$ such that $\mu = \mu_0$, and μ_l is popular, $l \leq (n^2 - n + 2)/2$, and for each μ_i, $i = 0, \ldots, l-1$ there is a blocking coalition of size up to 3 such that μ_{i+1} is obtained from μ_i by satisfying this coalition.

Sketch of the proof. We propose a simple finite algorithm that does it only by using one-,two- and three-way exchanges.

The algorithm has two stages. In the first stage it matches each first house to some of its first agents. This is done in a greedy serial dictatorship fashion. According to a fixed order each agent a takes his first house f unless this house is already matched to one of its other first agents (in this case no exchange takes place and we proceed to the next agent in the order). In the same time, the agent owning house f takes his own first house g and the owner of this house $\mu(g)$ takes the house of agent a. This three-way exchange is supported by at least two agents a and $\mu(f)$, and, possibly, also by agent $\mu(g)$.

In the second stage of the algorithm we use another simple greedy procedure where owners of bad houses are forcibly given their second houses. Each agent a owning some bad house t takes his second house s, while the owner of s takes his first house f,[7] and house t goes to the owner of f. This three-way exchange is supported by at least two agents a and $\mu(s)$, but the exchange also might be "bad" if t is a bad house for both agents $\mu(s)$ and $\mu(f)$. A bad exchange like that leads to the same situation as before: out of three agents one − agent $\mu(s)$ − owns his first house f, one − agent a − owns his second house s and one − agent $\mu(f)$ owns a bad house t. Next we agent $\mu(f)$ is given his second house and we continue until the procedure stops. It remains to show that this sequence of bad exchanges is finite.

The finiteness follows from that the sequence of bad exchanges eventually arrives to a house that was in the sequence earlier (due to finiteness of A). At this step k of the sequence we have $k - 2$ agents that have the same $\lfloor \frac{k-2}{2} \rfloor$ as their first and second houses. By the Hall's theorem (applied to the characterization in Corollary 1) these houses can only be matched to these agents. Thus agent k that gets house t after a series of bad exchanges can only start a new sequence of bad exchanges but the other $k - 1$ agents remain untouched with their matched houses until the end of the algorithm. Thus the procedure converges to some matching and, by Theorem 1, this matching is popular. ∎

Since our algorithm is finite we immediately get the convergence result that the set of absorbing states coincides with the set of popular matchings.

Corollary 2. For any initial matching, the random sequence of 1,2 and 3-way exchanges converges with probability one to a popular matching whenever such matching exists.

The restriction to the groups of up to three agents is not compulsory as the same algorithm works when groups of larger size are also allowed. The convergence result also holds for exchanges of arbitrary sizes.

Corollary 3. For any initial matching, the random sequence of arbitrary exchanges converges with probability one to a popular matching whenever such matching exists.

The original result in [24] was partially motivated by the example in [16] where he shows that a sequence of blocking pairs might have an infinite cycle and might never converge to stability. The same is true in our setting: even when a popular matching exists, the sequence of popular exchanges might have cycles. To see that let us consider the same preference profile as in Example 1:

---

[7]Note that s cannot be owned by his first agent, otherwise s does not qualify as a second house for agent a.

Table 4. Cycle with 4 matchings: $\mu_1, \mu_2, \mu_3, \mu_4$

1	2	3	4
<u>a</u>	a	d	<u>d</u>
d	c	b	b
b	<u>b</u>	<u>c</u>	a
c	d	a	c

1	2	3	4
a	<u>a</u>	d	<u>d</u>
d	c	<u>b</u>	b
b	b	c	a
<u>c</u>	d	a	c

1	2	3	4
a	<u>a</u>	<u>d</u>	d
d	c	b	b
<u>b</u>	b	c	a
c	d	a	<u>c</u>

1	2	3	4
a	<u>a</u>	d	d
<u>d</u>	c	b	<u>b</u>
b	b	<u>c</u>	a
c	d	a	c

We begin with matching $\mu_1 = (1a, 2b, 3c, 4d)$. If agents 1,2,3 meet and decide to exchange their houses by majority we get matching $\mu_2 = (1c, 2a, 3b, 4d)$. Next, if agents 1,3,4 meet and do the same, we get matching $\mu_3 = (1b, 2a, 3d, 4c)$, which can be again changed by majority to matching $\mu_4 = (1d, 2a, 3c, 4b)$.[8] Finally, if agents 1,2,4 meet, we again get matching μ_1.

Note that in this example there is a path that leads to a popular matching as was previously shown in the Introduction. In fact, as Theorem 2 demonstrates such a path exists in any instance.

5   Conclusions

In the current paper we propose a novel characterization of "global" popularity via "local" popularity, and also show that locally popular exchanges lead to a globally popular matching.

One important open question is about the convergence speed of popular markets. To answer this question one may need to design a more efficient algorithm: our greedy algorithm does many unnecessary steps, for instance when it repeatedly runs the same chains. We cannot simply avoid these steps as then we cannot build a triple that blocks the current matching. However, it might be possible if we use alternative algorithms.

Another open question is about popular markets in instances when popular matchings do not exist. Perhaps, these markets converge to some stationary probabilistic distribution over the set of matchings, and it is reasonable to deem the more probable matchings as more popular. Both questions are interesting but hard.

References

[1] Abraham, D. J., Irving, R. W., Kavitha, T., Mehlhorn, K. (2007). Popular matchings. SIAM Journal on Computing, 37(4), 1030-1045.

[2] Abraham, D. J., & Kavitha, T. (2010). Voting paths. SIAM Journal on Discrete Mathematics, 24(2), 520-537.

[3] Acharyya, R., Chakraborty, S., Jha, N. (2014, June). Counting popular matchings in house allocation problems. In International Computer Science Symposium in Russia (pp. 39-51). Springer, Cham.

[4] Aziz, H., Brandt, F., & Stursberg, P. (2013, October). On popular random assignments. In International Symposium on Algorithmic Game Theory (pp. 183-194). Springer, Berlin, Heidelberg.

[5] Bhattacharya, S., Hoefer, M., Huang, C. C., Kavitha, T., Wagner, L. (2015, July). Maintaining near-popular matchings. In International Colloquium on Automata, Languages, and Programming (pp. 504-515). Springer, Berlin, Heidelberg.

---

[8] Agents 1,3,4 have the same preferences over houses b, c, d that they own and we get a Condorcet cycle, which can also infinitely repeat itself.

[6] Brandt, F., Hofbauer, J., & Suderland, M. (2017, May). Majority graphs of assignment problems and properties of popular random assignments. In Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (pp. 335-343). International Foundation for Autonomous Agents and Multiagent Systems.

[7] Cseh, A. (2017). Popular matchings. In Endriss, U. (editor). Trends in Computational Social Choice. AI Access.

[8] Cseh, A., Huang, C. C., Kavitha, T. (2017). Popular matchings with two-sided preferences and one-sided ties. SIAM Journal on Discrete Mathematics, 31(4), 2348-2377.

[9] Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. The American Mathematical Monthly, 69(1), 9-15.

[10] Gardenfors, P. (1975). Match making: assignments based on bilateral preferences. Systems Research and Behavioral Science, 20(3), 166-173.

[11] Hylland, A., & Zeckhauser, R. (1979). The efficient allocation of individuals to positions. Journal of Political economy, 87(2), 293-314.

[12] Kamiyama, N. (2017). Popular matchings with ties and matroid constraints. SIAM Journal on Discrete Mathematics, 31(3), 1801-1819.

[13] Kavitha, T., Mestre, J., Nasre, M. (2011). Popular mixed matchings. Theoretical Computer Science, 412(24), 2679-2690.

[14] Kavitha, T. (2014). A size-popularity tradeoff in the stable marriage problem. SIAM Journal on Computing, 43(1), 52-71.

[15] Klaus, B., Manlove, D. F., Rossi, F. (2016). Matching under Preferences. In Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A. (editors). Handbook of Computational Social Choice. Cambridge University Press.

[16] Knuth, D. E. (1976). Marriages Stable. Universitй de Montrйal Press, Translated as "Stable Marriage and Its Relation to Other Combinatorial Problems,". CRM Proceedings and Lecture Notes, American Mathematical Society.

[17] Mahdian, M. (2006, June). Random popular matchings. In Proceedings of the 7th ACM conference on Electronic commerce (pp. 238-242). ACM.

[18] McCutchen, R. M. (2008, April). The least-unpopularity-factor and least-unpopularity-margin criteria for matching problems with one-sided preferences. In Latin American Symposium on Theoretical Informatics (pp. 593-604). Springer, Berlin, Heidelberg.

[19] McDermid, E., Irving, R. W. (2011). Popular matchings: structure and algorithms. Journal of combinatorial optimization, 22(3), 339-358.

[20] Mestre, J. (2014). Weighted popular matchings. ACM Transactions on Algorithms (TALG), 10(1), 2.

[21] Nasre, M. (2014). Popular Matchings: Structure and Strategic Issues. SIAM Journal on Discrete Mathematics, 28(3), 1423-1448.

[22] Paluch, K. (2014). Popular and clan-popular b-matchings. Theoretical Computer Science, 544, 3-13.

[23] Roth, A. E., & Sotomayor, M. (1992). Two-sided matching. Handbook of game theory with economic applications, 1, 485-541.

[24] Roth, A. E., & Vate, J. H. V. (1990). Random paths to stability in two-sided matching. Econometrica: Journal of the Econometric Society, 1475-1480.

[25] Shapley, L., & Scarf, H. (1974). On cores and indivisibility. Journal of mathematical economics, 1(1), 23-37.

[26] Sng, C. T., Manlove, D. F. (2010). Popular matchings in the weighted capacitated house allocation problem. Journal of Discrete Algorithms, 8(2), 102-116.

APPENDIX

Proof of Theorem 2.

The first part of the algorithm.

Let μ be the arbitrary initial matching where each agent is endowed with some house: for each $a \in A$ $\mu(a) \neq \emptyset$. Let us fix some ordering of agents $A = \{a_1, \ldots, a_n\}$.

For steps $k = 1, \ldots, n$ we make the following exchanges.

If in step k house $\mu(a_k)$ is the best house for agent a_k, then proceed to step $k+1$ without changing the current matching μ. Otherwise, consider house $h \neq \mu(a_k)$ that is the best house of agent a_k. If this house h is empty, $\mu(h) = \emptyset$, then we give it to agent a_k in the new matching $\mu'(a_k) = h$. Otherwise, consider the owner of h, $\mu(h)$.

If h is the best for its owner $\mu(h)$, then proceed to the next step $k+1$ without changing the current matching μ. Otherwise, consider the best house for agent $\mu(h)$: $h' \neq h$. If $h' = \mu(a_k)$ or $\mu(h') = \emptyset$ then make the mutually beneficial two-way exchange: $\mu'(a_k) = h, \mu'(\mu(h)) = h'$. Otherwise, if $\mu(h') \notin \{a_k, \mu(h), \emptyset\}$ we make the three-way exchange: $\mu'(a_k) = h, \mu'(\mu(h)) = h', \mu'(\mu(h')) = \mu(a_k)$. This exchange is beneficial for at least two of the three agents.

After each of the above exchanges the number of agents that own their best houses goes up, and each agent gets his best house unless it is taken by some other agent. Thus after $x \leq n$ exchanges we get a new matching μ where each agent gets either his first house, his second house, or a bad house.

Denote the number of agents who get a bad house by $\beta(\mu)$. At least $x \geq 1$ agents get their first house, therefore $\beta(\mu) \leq n - x \leq n - 1$. Note that $n - \beta(\mu)$ agents get either a first house or a second house.

The second part of the algorithm.

We will make exchanges that *weakly* decrease the number of agents with a bad house $\beta(\mu)$.

Consider some agent $\mu(t)$ that gets a bad house t. If his second house s is free, we give him s: $\mu'(\mu(t)) = s$ and decrease $\beta(\mu)$ by one. Otherwise there is some agent $\mu(s)$ that owns s, and s might be his bad house or his second house (but not his first house from the definition of second house). We now study these two cases.

1. Let s be a bad house for $\mu(s)$. Denote the second house of $\mu(s)$ as h. If $h = t$ or empty, then make the two-way exchange decreasing $\beta(\mu)$ by 2. Otherwise, make the three-way exchange $\mu'(\mu(t)) = s, \mu'(\mu(s)) = h, \mu'(\mu(h)) = t$, decreasing $\beta(\mu)$ by 1, 2 or 3 depending on how the owner of h ranks t.

2. Let s be the second house for $\mu(s)$. Let f be the first house for agent $\mu(s)$. From the first part of the algorithm we know that f is also the first house of his owner $\mu(f)$. Make the following three-way exchange: $\mu'(\mu(t)) = s, \mu'(\mu(s)) = f, \mu'(\mu(f)) = t$. If t is the second house for agent $\mu(f)$, then $\beta(\mu)$ decreases by one.

Thus $\beta(\mu)$ is only constant if house s is the second house for both $\mu(t)$ and $\mu(s)$, house f is the first house for both $\mu(s)$ and $\mu(f)$, and house t is a bad house for both agents $\mu(t)$ and $\mu(f)$. Denote such exchange as *bad*. We show now that a sequence of these bad exchanges in which $\beta(\mu)$ remains constant is finite.

Table 5. Current matching μ before and after a bad three-way exchange which keeps $\beta(\mu)$ a constant

$\mu(t)$	$\mu(s)$	$\mu(f)$	$\mu(t)$	$\mu(s)$	$\mu(f)$
	f	f̲		f̲	f
s	s̲		s̲	s	
t̲		t	t		t̲

2.1 Let f be the first house also for agent $\mu(t)$. For convenience denote $f = f_1, s = s_1, \mu(t) = 1, \mu(s) = 2, \mu(f) = 3$. By Hall's theorem the second house for agent 3 cannot be

the same as s_1, $s_3 \neq s_1$ (otherwise three agents have the same first house and the same second house, and thus a popular matching does not exist). After the bad exchange among agents 1,2,3 the bad house t is matched to agent 3. Consider another chain of three agents that starts with the bad house t. Denote $\mu(s_3) = 4$. Note that $f_4 \neq f_1$ (otherwise four agents have the same first house, two of them have the same second house, and the other two of them also have the same second house, and thus a popular matching does not exist). Denote $\mu(f_4) = 5$. By Hall's theorem $s_5 \notin \{s_1, s_3\}$ (otherwise, similar to the previous arguments the popular matching does not exist). After the bad exchange between agents 3,4,5 the bad house is matched with agent 5, and so forth.

Table 6. Current matching μ before and after two bad three-way exchanges

1	2	3	4	5	1	2	3	4	5
f_1	f_1	$\underline{f_1}$	f_4	$\underline{f_4}$	f_1	$\underline{f_1}$	f_1	$\underline{f_4}$	f_4
s_1	$\underline{s_1}$	s_3	$\underline{s_3}$	s_5	$\underline{s_1}$	s_1	$\underline{s_3}$	s_3	s_5
$\underline{t}$		t		t	t		t		$\underline{t}$

Note that in this case $\beta(\mu) \leq n - 2$. In each such bad exchange two new agents enter the chain, these agents own their first and second houses. Then, we need not more than $(n - \beta(\mu))/2$ bad exchanges and one additional exchange to reduce $\beta(\mu)$. Hence, the total number of exchanges reducing $\beta(\mu)$ is not more than

$$\frac{n - \beta(\mu)}{2} + 1 \leq n - \beta(\mu).$$

2.2 Let the first house f_1 for agent $\mu(t)$ be different from house f. Denote $f = f_2, s = s_1, \mu(t) = 1, \mu(s) = 2, \mu(f) = 3$. After one bad exchange agent 3 would be matched to house t.

Assume that the second house for agent 3 $s_3 \neq s_1$ − we did not meet s_3 earlier in the chain. Consider agent 4 that owns his second house s_3. Assume that agent 4's first house f_4 was note previously in the chain: $f_4 \neq f_1, f_2$.

Consider the next agent 5 and so on: we get a chain of agents such that each two neighbours have either the same first house or the same second house. Eventually we arrive to some agent k that has the same first or second house as earlier in the chain.

Let agent k be the first agent in the chain such that his first house has already appeared in the chain. In this case k is even. Then after $(k - 2)/2$ bad exchanges in one direction agent $k - 1$ gets the bad house t. Then agent $k - 1$ reverses the direction of bad exchanges such that agent $k + 1$ gets the bad house t, which happens after not more than than $k/2$ bad exchanges. And we see that not less than k agents get their first or second houses, $k \leq n - \beta(\mu)$.

Now agent $k + 1$ continues the bad exchanges. By Hall's theorem, in each such bad exchange two new agents enter the chain, these agents own their first and second houses. Then, we need not more than $(n - \beta(\mu) - k)/2$ bad exchanges and one additional exchange to reduce $\beta(\mu)$. Hence, for even k the total number of exchanges reducing $\beta(\mu)$ is not more than

$$\frac{k - 2}{2} + \frac{k}{2} + \frac{n - \beta(\mu) - k}{2} + 1 = \frac{n - \beta(\mu)}{2} + \frac{k}{2} \leq n - \beta(\mu).$$

Similarly, let agent k be the first agent in the chain such that his second house has already appeared in the chain. In this case k is odd. Then after $(k - 1)/2$ bad exchanges in one direction agent k gets the bad house t. Then agent k reverses the direction of bad exchanges such that agent 1 gets house f_1 and agent $k + 1$ gets the bad house t, which happens after not more than than $(k - 1)/2$ bad exchanges. And we see that not less than k agents get their first or second houses, $k \leq n - \beta(\mu)$.

Now agent $k + 1$ continues the bad exchanges. By Hall's theorem, in each such bad exchange two new agents enter the chain, these agents own their first and second houses. Then, we need not more than $(n - \beta(\mu) - k)/2$ bad exchanges and one additional exchange to reduce $\beta(\mu)$. Hence, for odd k the total number of exchanges reducing $\beta(\mu)$ is not more than

$$\frac{k-1}{2} + \frac{k-1}{2} + \frac{n - \beta(\mu) - k}{2} + 1 = \frac{n - \beta(\mu)}{2} + \frac{k}{2} \leq n - \beta(\mu).$$

Eventually, in the second part of the algorithm after at most $n - \beta(\mu)$ exchanges we decrease $\beta(\mu)$. In the worst case $\beta(\mu) = n - 1$ agents have a bad house, therefore, including the first part of the algorithm the upper bound is $1 + 1 + 2 + \ldots + (n - 1) = (n^2 - n + 2)/2$. ∎