



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Efficient Dynamic Barter Exchange

Ross Anderson, Itai Ashlagi, David Gamarnik, Yash Kanoria

To cite this article:

Ross Anderson, Itai Ashlagi, David Gamarnik, Yash Kanoria (2017) Efficient Dynamic Barter Exchange. *Operations Research* 65(6):1446-1459. <https://doi.org/10.1287/opre.2017.1644>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2017, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Efficient Dynamic Barter Exchange

 Ross Anderson,^a Itai Ashlagi,^b David Gamarnik,^c Yash Kanoria^d

^a Google, Mountain View, California 94043; ^b Management Science and Engineering, Stanford University, Stanford, California 94305; ^c Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; ^d Decision, Risk and Operations Division, Columbia Business School, New York, New York 10027

Contact: rander@google.com (RA); iashlagi@stanford.edu (IA); gamarnik@mit.edu (DG); ykanoria@gsb.columbia.edu,

 <http://orcid.org/0000-0002-7221-357X> (YK)

Received: May 20, 2014

Revised: December 17, 2015; October 7, 2016

Accepted: April 17, 2017

Published Online in Articles in Advance:
September 14, 2017

Subject Classifications: networks/graphs: matchings; probability:stochastic model applications; queues: algorithms

Area of Review: Games, Information, and Networks

<https://doi.org/10.1287/opre.2017.1644>

Copyright: © 2017 INFORMS

Abstract. We study dynamic matching policies in a stochastic marketplace for barter, with agents arriving over time. Each agent is endowed with an item and is interested in an item possessed by another agent homogeneously with probability p , independently for all pairs of agents. Three settings are considered with respect to the types of allowed exchanges: (a) only two-way cycles, in which two agents swap items, (b) two-way or three-way cycles, (c) (unbounded) chains initiated by an agent who provides an item but expects nothing in return.

We consider the average waiting time as a measure of efficiency and find that the cost outweighs the benefit from waiting to thicken the market. In particular, in each of the above settings, a policy that conducts exchanges in a greedy fashion is near optimal. Further, for small p , we find that allowing three-way cycles greatly reduces the waiting time over just two-way cycles, and conducting exchanges through a chain further reduces the waiting time significantly. Thus, a centralized planner can achieve the smallest waiting times by using a greedy policy, and by facilitating three-way cycles and chains, if possible.

Funding: The second author acknowledges the research support of the National Science Foundation [Grant SES-1254768].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/opre.2017.1644>.

Keywords: barter • matching • market design • random graphs • dynamics • kidney exchange • platform • control policy

1. Introduction

Thousands of incompatible patient-donor pairs enroll at kidney exchange clearinghouses every year around the world in order to swap donors. Kidney exchange is just one example of a marketplace in which agents arrive and exchanges take place over time. Online platforms that enable the exchange of goods (e.g., homes for vacation, used goods, or services) and platforms that allow users to find matches (dating, labor, etc.) can be viewed as dynamic marketplaces for exchanges.

Marketplaces for barter use different matching technologies in order to overcome the rare coincidence of wants (Jevons 1876). For example, while kidney exchanges were first conducted in two-way cycles (Roth et al. 2005), most transplants in kidney exchange are now conducted through chains initiated by altruistic donors (Anderson et al. 2015); see Figure 2. Some three-way cycles are conducted, but cycles are kept short because of logistical constraints. Chains, by contrast, can be longer, as each pair can receive a kidney before it donates a kidney. In many marketplaces for matching (such as dating), only bilateral matches take place.

The *policy* employed by the clearinghouse, which determines which exchanges to implement, and when,

also affects the efficiency of the marketplace. One natural policy is the *greedy policy*, where the clearinghouse conducts exchanges immediately as the opportunity arises. Alternatively, the clearinghouse can adopt a *batching policy* where it accumulates a number of agents in order to thicken the market before it identifies a set of exchanges to conduct. For example, clearinghouses for kidney exchange in the United States have gradually moved to small batches.¹ More sophisticated policies are also possible. This paper is concerned with the effect matching policies have on efficiency of dynamic marketplaces that enable different matching technologies.


To illustrate the trade-off between serving agents quickly and thickening the market, consider a marketplace in which only two-way cycles may be implemented. Further, consider the situation depicted in Figure 1. Suppose agents a , c , and b arrive in that order, and the potential two-way cycles are as in the figure. Under a greedy policy, one of these possible exchanges would immediately be executed, say (b, c) . A different policy \mathcal{P} may not execute any exchange and when d arrives it may implement exchanges (a, b) and (c, d) , so that all four agents are matched. On the other hand, under a greedy policy, a and d remain unmatched and will therefore wait for other agents to complete an

Figure 1. Marketplace with Agents $a, c, b,$ and d Arriving in This Order



Note. Waiting for agent d to arrive ensures that both (a, b) and (c, d) make exchanges.

Figure 2. (Color online) A Compatibility-Graph Representation of the Potential Trades in a Market



Notes. Each circle node is an agent and the rectangle node, a , is an “altruistic donor” who is willing to provide an item for free. A link from agent n_i to agent n_j means that n_j is willing to accept the good that agent n_i has. This graph contains a two-way cycle $(n_1 \rightarrow n_2 \rightarrow n_1)$, three-way cycles including $(n_2 \rightarrow n_3 \rightarrow n_5 \rightarrow n_2)$, and multiple chains beginning from a including $(a \rightarrow n_1 \rightarrow n_2 \rightarrow n_3 \rightarrow n_4 \rightarrow n_7)$.

exchange, which may take a long time. Thus, with the benefit of foresight, it is possible to do better than the greedy policy. Without foresight, it is unclear whether the clearinghouse should operate in a greedy fashion, or adopt a more sophisticated policy.

We address the problem of efficient centralized exchange by studying a stylized dynamic model. Each period a single agent arrives with a single indivisible item that she wishes to exchange. Our model has a homogeneous and independent stochastic demand structure, in which every agent is willing to exchange her item for any other agent’s item with probability p ; see Figure 2. Agents in our model prefer to match as early as possible and are indifferent between alternate feasible exchanges. We therefore adopt the average waiting time of agents as a measure of efficiency. (In various contexts, it may be appropriate to consider some other cost function that is nonlinear in the waiting time. However, for this work we choose the simple linear cost function.) Agents depart upon matching.

Exchanges are conducted through (directed) cycles, or chains, and a policy determines which potential exchanges to conduct. We consider three settings distinguished by the types of possible/allowed exchanges: (i) two-way cycles, (ii) two-way and three-way cycles, and (iii) unbounded chains.

Our main contributions are the following. First, we find the greedy policy to be approximately optimal in the first and third settings, and approximately optimal in a class that includes batching policies in the second setting. Second, we find that conducting greedy exchanges through a chain (even a single one), results in significantly lower waiting times than when only two-way and three-way cycles are feasible, which in turn results in significantly lower waiting times than when only two-way cycles are feasible.

More precisely, we show that as $p \rightarrow 0$, the average waiting time under the greedy policy scales as $\Theta(1/p^2)$ for the setting based on two-way cycles, as $\Theta(1/p^{3/2})$ for the setting based on two-way and three-way cycles, and as $\Theta(1/p)$ for the setting based on chains. Furthermore, these average waiting times are essentially optimal in the first and third settings, and optimal in a class that includes batching policies in the second setting. We remark here that a small value of p is reasonable in many practical contexts, since agents are often interested in only a small fraction of the items offered by other agents. (For instance, kidney-exchange clearinghouses see a substantial fraction of highly sensitized patients who have a 1%–5% probability of matching.) Our results imply that in each setting, for all $p \in (0, 1)$, the waiting time under the greedy policy is within a constant factor of the waiting time using any other batch size. Simulation experiments suggest that a batch of size one is, in fact, truly optimal in each setting for any p .

The near optimality of the greedy policy implies that the cost outweighs the benefit from waiting to thicken the market in our setting with ex ante homogeneous agents. A very crude intuition for this is as follows. Minimizing the expected waiting time is equivalent to minimizing the expected number of agents who are waiting. Given an exchange opportunity, the platform should immediately execute the exchange, reducing the number of agents waiting. In the worst case, the served agents will soon be replaced by similar agents who arrive next if the next arrivals do not match immediately in turn; if the next arrivals *do* match immediately, this will only further reduce the number of agents waiting in the system. (Clearly, this reasoning is grossly incomplete given that our agents are homogeneous only ex ante.)

The two-way cycles setting turns out to be the simplest of the three settings considered; a short but non-trivial analysis allows us to prove tight bounds for this case. A simplifying feature of this case is that the state of the system under the greedy policy is just the number of currently waiting agents. A key challenge in the three-cycles and chains settings is that the compatibility graph between currently waiting agents, conditional on running the greedy policy so far, is not a directed Erdős–Rényi graph and has a complex distribution. It is sparser in terms of compatibilities in a very specific way: there are no possible exchanges, since the greedy policy would already have executed them. We develop methods to control the number of agents in the system despite this complexity. Another contribution is the technique we develop to prove lower bounds on average waiting times: this technique involves proof by contradiction, and is used in the case of three cycles and chains.

1.1. Related Work

The first strand of research that is related to our work is the literature that investigated efficiency in kidney exchange. Roth et al. (2007) found that in a static large pool, in which compatibilities are determined by blood types, three-way exchanges increase the number of matches but there is no need for exchanges of size 4 or larger. Ünver (2010) analyzes a dynamic kidney-exchange model in which compatibilities are determined by blood types. He finds a myopic mechanism that matches patient-donor pairs using only two-way and three-way cycles to be optimal. The finite number of types together with deterministic compatibility essentially creates a thick marketplace and therefore Ünver's results are closely related to the static large market. However, as mentioned in the introduction, most transplants are conducted through chains, which Ashlagi et al. (2012) associate with the lack of thickness in kidney exchange and the large percentage of highly sensitized patients.

Closely related to our paper are the works by Ashlagi et al. (2013) and Akbarpour et al. (2017) who study dynamic matching markets with an underlying stochastic network and preferences are based on compatibilities. Ashlagi et al. (2013) compare batching policies in a finite-horizon model with hard-to-match and easy-to-match agents (motivated by empirical observations in kidney exchange pools), by counting the number of matches. They find a very small benefit from small batches over greedy policies but find that chains lead to many more matches than two-way or three-way cycles. (Dickerson et al. 2012, also demonstrate the benefit of chains, using simulations in dynamic kidney-exchange pools.) When batches are large, more agents can be matched, yet their model does not account for the waiting costs that come with large batches. Another difference is that hard-to-match agents in their model match with a probability that is inversely proportional to the pool size (here the pool size is the arrival rate times the time horizon), whereas in our model compatibility (or the probability of matching) is independent of the pool size.

In a concurrent work, Akbarpour et al. (2017) study a model similar to ours and focus on analyzing policies for two-way exchanges. Agents in their model can depart and the goal of the clearinghouse is to minimize the loss rate (the fraction of unmatched agents). They study the benefits from knowing agents' departure times, and find that one can benefit significantly by making agents wait instead of matching agents greedily. This particular conclusion is the opposite of our own, driven directly by knowledge of departure times by the clearinghouse (and no waiting cost). When agents' departure times are unknown to the clearinghouse, Akbarpour et al. (2017) find the greedy policy to be almost optimal. We remark that this finding

is almost identical to our own for the case of two-way cycles, due to a close relationship between their model and our own, since in both formulations the objective translates to minimizing the expected number of agents waiting; see Appendix EC.4 in the online appendix. In fact, following the same analysis as in our Theorem 1, we obtain a tighter result (Theorem EC.1 in the online appendix) than the bounds on the performance of the greedy policy in Akbarpour et al. (2017, Theorem 3.10).

Periodic matching has been studied in other (non-barter) markets as well. Mendelson (1982) analyzes a clearinghouse that periodically searches for outcomes in a dynamic market with sellers and buyers who arrive according to a stochastic process and trade indivisible goods. He studies the behavior of prices and quantities resulting from periodic trading. Budish et al. (2015) find that some very small batching increases efficiency over continuous-time trading in financial markets as firms compete over price rather than over speed. These works study (in)efficiencies resulting from prices, while our work focuses on the waiting times in a homogeneous environment.

Another strand of research is the literature on online matching motivated by online advertising. Karp et al. (1990) initiated this line of research in an adversarial setting for bilateral matchings, and recent papers model the underlying compatibility graph to be stochastic (Goel and Mehta 2008, Feldman et al. 2009, Manshadi et al. 2012, Jaillet and Lu 2013). These studies focus on who to match and not when to match.

Some papers in the queuing literature study models under which both customers and servers arrive sequentially and have to match (Caldentey et al. 2009, Adan and Weiss 2012). Agents in our model can be viewed as both servers and customers and, in addition, our compatibility graph is stochastic.

We now say a few words to situate our work from a technical perspective. Our model and analysis bring together the rich literature on (static) random graph models, e.g., Bollobás (2001), Janson et al. (2011), and the rich literature on queuing systems, e.g., Kleinrock (1975), Asmussen (2003). In our model, the queue of waiting agents has a graph structure (i.e., the compatibility graph). Our stochastic model of compatibilities mirrors the canonical Erdős–Rényi model of a directed (static) random graph (but the dynamics make it much more complex). Comparing with common models of queuing systems, our system is peculiar in that the queuing system does not contain “servers” per se. Instead, the queue, in some sense, serves itself by executing exchanges that the compatibility graph allows. (Gurvich and Ward 2014, study optimal control in a related setting where jobs—analagous to our agents—can be served by matching with other compatible jobs.

A key difference is that they consider a fixed number of job types, whereas in our setting compatibility is stochastically drawn for each pair of agents, which leads to an unbounded number of possible agent types.) Nodes form cycles or chains with other nodes. Each time an exchange is executed, the corresponding agents/nodes leave the system. As a result, it turns out that for any reasonable policy the system is stable, irrespective of the rate of arrival of agents. If we speed up the arrival rate of agents, the entire system speeds up by the same factor, and waiting time reduces by the same factor. Thus, without loss of generality we consider an arrival rate of 1, with one agent arriving in each time slot.

Finally, there is a large literature on trading in markets without monetary transfers in which agents are endowed with a single good, often referred to as “housing markets” (Shapley and Scarf 1974). In a housing market one considers a group of agents, each of whom owns a house and has preferences over the set of houses. Shapley and Scarf studied the core of the market, and described the well-known top-trading cycles algorithm, which finds an element in the core. This literature has grown to be very mature, analyzing core properties (e.g., Roth and Postlewaite 1977), incentives, and design (e.g., Roth 1982, Abdulkadiroğlu and Sönmez 1998, Pycia and Ünver 2017). Some applied studies are Wang and Krishna (2006) on timeshare exchanges (which allow people to trade a “week” of vacation they own) and Dur and Ünver (2012) on tuition exchange (which allows dependents of faculty members to attend other colleges for no tuition). This literature focuses on static marketplaces.

1.2. Notational Conventions

We conclude with a summary of the mathematical notation used in the paper. Throughout, \mathbb{R} (\mathbb{R}_+) denotes the set of reals (nonnegative reals). We write that $f(p) = O(g(p))$, where $p \in (0, 1]$, if there exists $C < \infty$ such that $|f(p)| \leq Cg(p)$ for all $p \in (0, 1]$. We adapt the $\Theta(\cdot)$ and $\Omega(\cdot)$ notations analogously. We write that $f(p) = o(g(p))$ where $p \in (0, 1]$, if for any $C > 0$ there exists $p_0 > 0$ such that we have $|f(p)| \leq Cg(p)$ for all $p \leq p_0$. We adapt the $\omega(\cdot)$ notation analogously.

We let $\text{Bernoulli}(p)$, $\text{Geometric}(p)$, and $\text{Bin}(n, p)$ denote a Bernoulli variable with mean p , a geometric variable with mean $1/p$, and a binomial random variable that is the sum of n independent identically distributed (i.i.d.) Bernoulli(p) random variables (r.v.s). We write $X \stackrel{d}{=} \mathcal{D}$ when the random variable X is distributed according to the distribution \mathcal{D} . We let $\text{ER}(n, p)$ be a *directed* Erdős–Rényi random graph with n nodes where every two nodes form a directed edge with probability p , independently for all pairs. We let $\text{ER}(n, M)$ be the closely related *directed* Erdős–Rényi random graph with n nodes and M directed edges,

where the set of edges is selected uniformly at random from all subsets of exactly M directed edges. The two models are almost indistinguishable and, as is common in the literature on random graphs (Janson et al. 2011), we will use one model or the other depending on the context. We let $\text{ER}(n_L, n_R, p)$ denote a bipartite *directed* Erdős–Rényi random graph with two sides. This graph contains n_L nodes on the left, n_R nodes on the right, and a directed edge between every pair of nodes containing one node from each side is formed independently with probability p . (Edges occur both from left to right, and from right to left.) Given a Markov chain $\{X_t\}$ defined on a state space \mathcal{X} and given a function $f: \mathcal{X} \rightarrow \mathbb{R}$, for $x \in \mathcal{X}$, we use the shorthand

$$\mathbb{E}_x[f(X_t)] \triangleq \mathbb{E}[f(X_t) \mid X_0 = x].$$

1.3. Organization

The rest of the paper is organized as follows. We describe our model formally in Section 2 and state the main results of the paper in Section 3. In Section 4, we describe simulation results supporting our theoretical findings.

We prove our main results for cycles of length two only in Section 5.1, our results for two- and three-way cycles (technically the most challenging) in Section 5.3, and our results for chains in Section 5.2. Section 6 concludes.

2. Model

Consider the following infinite-horizon model of a barter exchange where each agent arrives with an item that she wants to exchange for another item. In each period $t = 0, 1, 2, \dots$ one new agent arrives to the marketplace.² The item of the new agent v is of interest to each of the waiting agents independently with probability p and, independently, the new agent v is interested in the item of each waiting agent independently with probability p . Agents are indifferent between compatible items and wish to exchange as early as possible, their cost of waiting being proportional to the waiting time. Agents leave the marketplace once they complete an exchange.

We study matching policies in three different settings distinguished by how agents can exchange items. In the first two settings (cyclic exchange) for each $k = 2, 3$, at most k agents can exchange items in a cyclic fashion. In the third setting (chain) agents can exchange items through a single chain: in the initial period $t = 0$, a single altruistic donor arrives who is willing to give her item away without getting anything in return. (In all subsequent periods, regular agents arrive who want to exchange their item for another.) Each agent in the chain receives a compatible item from one agent and gives to another. The chain advances over time and after each period there is exactly one agent who has received an item but has yet to give her item. We refer

to this agent as a *bridge agent* (note that at any time there is exactly one bridge agent in the marketplace) and we refer to the exchanges in a given time period as a chain segment. A *policy* is a mapping from the history of exchanges and the state of the marketplace to a set of feasible exchanges involving nonoverlapping sets of agents.

We adopt the average waiting time in steady state as the measure of the efficiency of a policy (the waiting of an agent is the difference between her departure time and her arrival time). This is equivalent to maximizing the social welfare in our model when each agent is given equal importance, since all agents have the same linear cost of waiting. In our model, minimizing the average waiting time is equivalent to minimizing the average number of agents in the marketplace, since these two quantities are proportional to each other by Little's law (Little 1961).

Remark 1. Our model, while simplistic in its compatibility structure (which is described by a single parameter p), has several advantages. It avoids a “market size” parameter altogether (faster arrival of agents simply leads to an inverse rescaling of time).³ Further, studying steady-state behavior allows performance to be quantified exclusively in terms of waiting times. The alternative approach of studying a finite time horizon, as in Ashlagi et al. (2013), involves end-of-period effects that make it necessary to simultaneously consider both the waiting times and the number of matches, and thus hinder performance comparisons.

Remark 2. We do not model agent departures due to death/time-out. However, in our formulation, minimizing waiting leads to minimizing the expected number of agents waiting in the system by Little's law (Little 1961) and hence aligns with minimizing the fraction of agents who die without being served (if deaths occur). See Appendix EC.4 in the online appendix for a formal treatment.

It is convenient to think about a state of the marketplace at any time in terms of a *compatibility graph*, which is a directed graph with each agent represented by a node, and a directed edge (i, j) representing that agent j wants the item of agent i . Let $\mathcal{G}(t) = (\mathcal{V}(t), \mathcal{E}(t))$ denote the directed graph of compatibilities observed before time t . When a new agent arrives, a directed edge is formed (in each direction) with probability p between the arriving agent v and each other agent that currently exists in the system, independently for all agents and directions. A cyclic exchange corresponds to a directed cycle in $\mathcal{G}(t)$ and a chain segment is a directed path in $\mathcal{G}(t)$ starting from the bridge node/agent.

One natural policy that will play a key role in our results is the *greedy policy*. The greedy policy attempts to match the maximum number of nodes upon each arrival.

Definition 1 (Greedy Policy). The *greedy policy* for each of the settings is defined as follows:

- **Cycle Removal.** At the beginning of each time period the compatibility graph does not contain cycles with length at most k . Upon the arrival of a new node, if a cycle with length at most k can be formed with the newly arrived node, it is removed, with a uniformly random cycle being chosen if multiple cycles are formed. Clearly, at the beginning of the next time period the compatibility graph again does not contain any cycles with length at most k . The procedure is described in Figure 4.

- **Chain Removal.** There is one bridge node in the system at the beginning of every time period. This bridge node does not have any incoming or outgoing edges. Upon the arrival of a new node at the beginning of a new time interval, the greedy policy identifies the longest chain segment originating from the bridge node (breaking ties uniformly at random) and removes these nodes from the system and the last node in the chain becomes a bridge node. Note that such a chain has a positive length if and only if the bridge node has a directed edge from it to the new node. Observe that the new bridge node has indegree and outdegree zero, and so the process can repeat itself. This procedure is described in Figure 3.

In each of the settings, under the greedy policy the state evolves as a Markov chain. Further, this Markov chain is irreducible since an empty graph is reachable from any other state. For Markov chains that are positive recurrent one can study the average waiting time.


Definition 2 (Periodic Markov Policies). We call a policy a *periodic Markov policy* if it employs τ homogeneous first-order Markov policies in round robin for some $\tau \in \mathbb{N}$, and the market state is irreducible and positive recurrent under the policy.⁴

In other words, a periodic Markov policy implements a heterogeneous first-order Markov chain, where the transition matrices repeat cyclically every τ rounds. The subsequence of states starting with the state at time l and then including states at time intervals of τ , i.e., times $t = l, l + \tau, l + 2\tau, \dots$ forms an irreducible and positive recurrent first-order Markov chain. Call this the l th “outer” Markov chain. Without loss of generality, this Markov chain is aperiodic. (If it is periodic with period τ' , then redefine τ as per $\tau \leftarrow \tau\tau'$, and the outer Markov chains will now be aperiodic.) It follows that the market state has a period of τ . Define

$$W_l \equiv \text{Expected number of nodes in the system in the steady state of the } l\text{th outer Markov chain.}$$


Thus, W_l is the expected number of nodes in the system at times that are $l \bmod \tau$ in steady state. Then we define

Figure 3. (Color online) An Illustration of Chain Matching Under the Greedy Policy



Notes. Initially, as shown on the left, h is the head of the chain (the *bridge agent*), and nodes w_1 , w_2 , and w_3 are waiting to be matched. First, node a_1 arrives, and his good is acceptable to both w_1 and w_3 but no one has a good acceptable to a_1 . As h 's good is not acceptable to a_1 , it is not possible to advance the chain. Then node a_2 arrives. His good is acceptable to w_2 and he is able to accept the good from h . The longest possible chain is shown above via dashed edges in the center. The chain is formed, h , a_2 , and w_2 are removed, and w_3 becomes the new head of the chain (bridge agent). Edges incident to the matched nodes are removed, as well as edges going in to w_3 . Note that in this case, the longest chain is not unique: w_1 could have been selected instead of w_3 .

Figure 4. (Color online) An Illustration of Cycle Matching Under the Greedy Policy, with a Maximum Cycle Length of 3



Notes. Initially, as shown on the left, nodes n_1 , n_2 , n_3 , and n_4 are all waiting. Node n_5 arrives, but no directed cycles can be formed. Then n_6 arrives, forming the three-way cycle $n_6 \rightarrow n_2 \rightarrow n_4 \rightarrow n_6$. On the right, the three-way cycle is removed, along with the edges incident to any node in the three-way cycle. Note that when n_6 arrives, a six-way cycle is also formed; but under our assumptions, the maximum-length cycle that can be removed is a three-way cycle.

the average waiting time for a periodic Markov policy as

$$W = \frac{1}{\tau} \sum_{l=0}^{\tau-1} W_l. \quad (1)$$

Note that this is the average number of nodes in the original system over a long horizon in steady state. By Little's law, this is identical to the average waiting time for agents who arrive to the system in steady state, which is our measure of efficiency.

Remark 3. We state our results formally for the broad class of periodic Markov policies, though our bounds extend also to other general policies that lead to a stationary/periodic and ergodic system in the $t \rightarrow \infty$ limit.

3. Main Results

We consider three different settings: (a) bilateral matches (two-way cycles), (b) two-way cycles and three-way cycles, and (c) unbounded chains initiated by altruistic donors. In each setting we look for a policy that minimizes expected waiting time in steady state.

3.1. Bilateral Matching

Our first result considers the case in which agents can exchange only through bilateral matches, i.e., through two-way cycles.

Theorem 1. *In the setting where only two-way cycles may be executed, the greedy policy achieves an average waiting time of $\ln 2/p^2 + o(1/p^2)$. This is optimal, in the sense that for every periodic Markov policy, the average waiting time is at least $\ln 2/(-\ln(1 - p^2)) = \ln 2/p^2 + o(1/p^2)$.*

We provide some intuition regarding the optimality of the greedy policy in this setting in the introduction. The scaling of $1/p^2$ follows from the fact that the prior probability of having a two-way cycle between a given pair of nodes is p^2 , and so an agent needs $\Theta(1/p^2)$ options in order to find another agent with whom a mutual swap is desirable. This result is technically the simplest to establish, but the insight obtained is nevertheless important. We prove Theorem 1 in Section 5.1, using the fact that in steady state, the arrival rate must match the departure rate and hence two-way cycles must be removed at a rate of $1/2$.

3.2. Two-Way Cycles and Three-Way Cycles

Our second result considers the case in which two-way and three-way cycles are feasible.

Definition 3. A *batching policy* with batch size N waits for N new arrivals, and then executes a set of feasible exchanges that maximize the number of agents served from among those currently waiting. The agents who are not served become a part of the next batch.

Note that the greedy policy is a special case of batching with a batch size of 1.

Theorem 2. *Under the cycle removal setting with $k = 3$, the average waiting time under the greedy policy is $O(1/p^{3/2})$. Furthermore, there exists a constant $C < \infty$ such that, for any batching policy, the average waiting time is at least $1/(Cp^{3/2})$.*

Theorem 2 says that we can achieve a much smaller waiting time with $k = 3$, i.e., two-way and three-way

cycle removal, than the removal of two-way cycles only (for small p), since $1/p^{3/2}$ is much smaller than $1/p^2$. Further, for $k = 3$, the greedy policy is again near optimal in the sense that no batching policy can beat the greedy policy by more than a constant factor.

The following fact may provide some intuition for the $\Theta(1/p^{3/2})$ scaling of average waiting time (recall that the average number of nodes is the same as the average waiting time, using Little's law). In a static directed Erdős–Rényi graph with (small) edge probability p , one needs the number of nodes n to grow as $\Omega(1/p^{3/2})$ in order to cover, with high probability, a fixed fraction (e.g., 50%) of the nodes with node disjoint two-way and three-way cycles.⁵ Our rigorous analysis leading to Theorem 2 shows that this coarse calculation in fact leads to the correct scaling of the average number of nodes in the dynamic system under the greedy policy, and that no batching policy can do better.

We provide a partial sketch of the proof of Theorem 2 in Section 5.3, deferring the full proof to Appendix EC.2 in the online appendix. The proof overcomes a multitude of technical challenges arising from the complex distribution of the compatibility graph at a given time, and introduces several new ideas. Our lower bound in this case applies not just to batching policies but to the more general class of *monotone policies*, which we define in Appendix EC.2 in the online appendix. Roughly, under a monotone policy, the presence of a compatibility (i, j) does not affect the exchanges that are executed until either i or j is served.

An example of a nonmonotone policy is one that prioritizes nodes that have a low indegree. However, we remark that we could not construct any candidate policy in our homogeneous model of compatibility that violates monotonicity but should do well in terms of average waiting time. That being the case, we conjecture (but are unable to prove) that our lower bound on average waiting time applies to arbitrary and not just monotone policies.

Our result leaves open the case of larger cycles, i.e., $k > 3$, under the greedy policy, arbitrary monotone policies, and arbitrary general policies. Based on intuition similar to the above, we conjecture that under the cycle removal setting with general k , the greedy policy achieves an average waiting time of $\Theta(p^{-k/(k-1)})$, and furthermore for every policy the average waiting time is lower bounded by $\Omega(p^{-k/(k-1)})$.

3.3. Unbounded Chains Initiated by Altruistic Donors

Our final result concerns performance in the chain removal setting.

Theorem 3. *In the chain removal setting, the greedy policy (see Definition 1) achieves an average waiting time of $O(1/p)$. Further, there exists a constant $C < \infty$ such that even if we allow the removal of cycles of arbitrary length*

in addition to chains, for any periodic Markov policy (see Definition 2) the average waiting time is at least⁶ $1/(Cp)$.

Thus, unbounded chains initiated by altruistic donors lead to a further large reduction in waiting time relative to the case of two-way and three-way cycles, for small p , since $1/p$ is much smaller than $1/p^{3/2}$. In fact, as stated in the theorem, the removal of cycles of arbitrary length (and chains), with any policy, cannot lead to better scaling of waiting time than that achieved with chains alone. Finally, the greedy policy is near optimal among all periodic Markov policies for chain removal. (In Remark 4, we argue that this should also hold in the related setting where chains, two-way cycles, and three-way cycles are all allowed.)

The intuition for the $\Theta(1/p)$ scaling of waiting time is somewhat involved. Since an agent finds the item of another agent acceptable with probability p , it is not hard to argue that no policy can sustain an expected waiting time that is $o(1/p)$; see our proof of the lower bound in Theorem 3 for a formalization of this intuition. On the other hand, under a greedy policy, the chain advances each time a new arrival can accept the item of the bridge agent, which occurs typically at $\Theta(1/p)$ intervals. Intuitively, if there are many agents waiting, then, typically, the next time there is an opportunity to advance the chain, we should be able to identify a long chain that will eliminate more agents than the number of agents who arrived since the last advancement. Our proof shows that this is indeed the case.

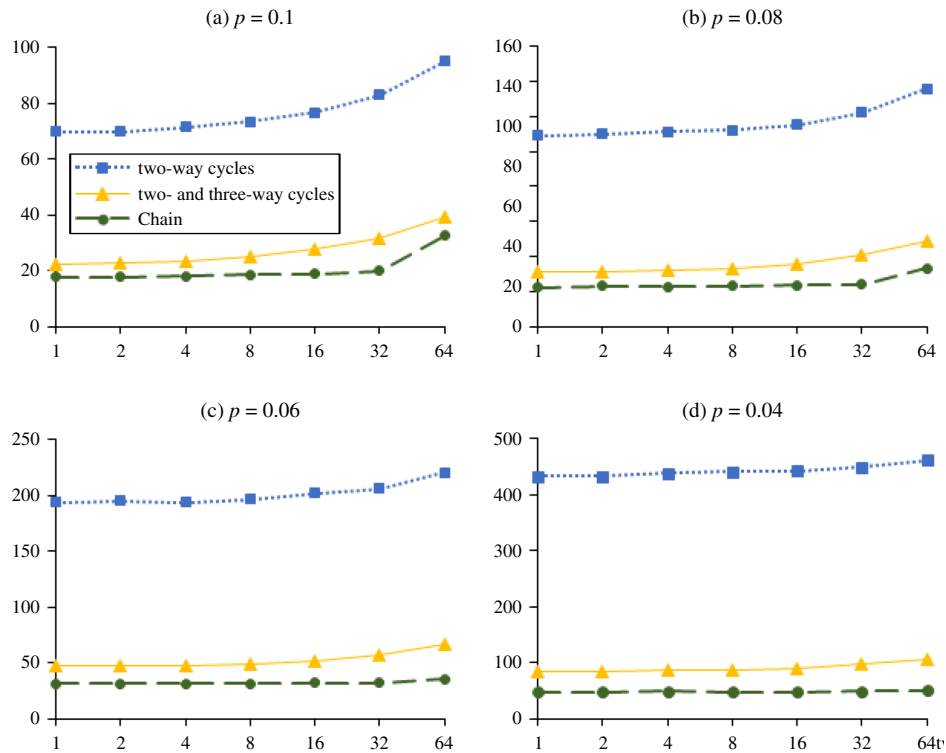
The proof of Theorem 3 is technically challenging. We provide a sketch of the proof in Section 5.2, and defer the full proof to Appendix EC.3 in the online appendix.

4. Computational Experiments

We conducted simulation experiments of our model for each of the matching technologies (two-way cycles, two- and three-way cycles, and chains). First, for each of the matching technologies we calculated the average waiting time under a batching policy for a variety of batch sizes. For each scenario, we simulated a time horizon with 16,000 arriving nodes, and measured the average number of nodes in the system after the arrival of the 2,000th node (the first 2,000 arrivals serve as a “warm-up” period). We conducted three trials for each scenario simulated.

Figure 5(a) illustrates that when $p = 0.1$, the greedy policy, which corresponds to the batching policy of the batch size $x = 1$, performs the best among all simulated batch sizes. In addition, observe the significant difference between average waiting times in the chain setting on the one hand and the two-way cycles on the other hand.⁷ Figures 5(b)–(d) provide similar results for the cases of $p = 0.08, 0.06$, and 0.04 . Moreover, under the

Figure 5. (Color online) Average Waiting Time versus Batch Size Under Batching Policies for Two Way Cycle Removal, Two- and Three-Way Cycle Removal, and Chain Removal



Notes. Batch sizes considered are 1, 2, 4, 8, 16 and 64. Waiting time is on the vertical axis whereas batch size is on the logarithmically-scaled horizontal axis.


greedy policy, for each p we simulated, the difference between the average waiting time and the predicted waiting time $\ln(2)/p^2$ was at most 1.3.

We conducted simulations of the greedy policy with two-way and three-way cycles for $p = 0.04$. We ran the system for 16,000 time steps, and recorded the total number of nodes, the total number of edges, and the number of nodes with each possible indegree and each possible outdegree, after every 100 time steps, skipping the first 2,000 steps (which should have allowed the market to reach steady-state). Though our analysis is not sufficiently fine-grained to capture it, our theoretical intuition suggests that the residual compatibility graph should have the same degree distribution to a leading order as an Erdős–Rényi random graph with the same edge probability. This intuition is confirmed as shown in Figure 6. For the purposes of this figure, we consider only those snapshots of the system in which the number of nodes was between 80 and 86, and compare the average indegree distribution and outdegree distribution with those for an Erdős–Rényi graph with 83 nodes and edge probability 0.04. Similar results were obtained for other ranges of the number of nodes. Further, we expect the number of nodes to remain close to $\sqrt{\ln(3/2)/p^{3/2}} \approx 79.6$; see Conjecture 1. Consistent with this estimate, our simulations reveal an average number of nodes of 84.7 with a standard deviation of 7.3. An Erdős–Rényi random graph

is similarly found to provide a good approximation in the chain setting, as well as the two-way cycle setting, under a greedy policy with $p = 0.04$. We omit the degree distribution charts in the interest of brevity. The distribution of nodes in the system in the case of chains is shown in Figure 7; the number of nodes varies in a wider range in this case (since the chain advances in bursts), but typically remains within a factor of 3 of the average value. In the case of two-way cycles, we find that the number of nodes remains close to $\ln(2)/p^2 \approx 433.2$, as predicted by Theorem 1 (we find a mean of 434.9 with a standard deviation of 22.9).

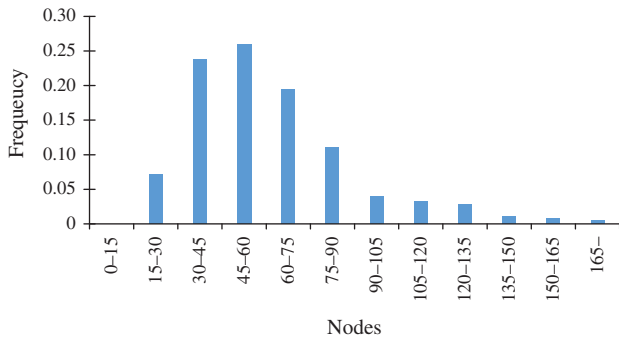
Based on the conjecture that the compatibility graph closely resembles an Erdős–Rényi graph under the greedy policy, we further conjecture that the greedy policy is in fact near optimal in the case of two- and three-way cycles (not just up to constant factors). Our intuition mirrors the proof of optimality of the greedy policy in the case of two-way cycles. First, recall that two-way cycles play almost no role when three-way cycles are permitted. In steady state, a three-way cycle must be formed with probability at least $1/3$ in each period. The probability of formation of a three-way cycle is uniquely determined by the number of directed edges in the residual graph, and this in turn is uniquely determined by the number of nodes if the residual graph resembles an Erdős–Rényi graph. Now

Figure 6. (Color online) Degree Distribution Under the Greedy Policy with Two-Way and Three-Way Cycles and $p = 0.04$



Notes. The solid lines show the average number of waiting nodes with given indegrees and outdegrees when the system contains 80–86 nodes. For comparison, the dashed lines show the average number of nodes with different degrees in a directed Erdős–Rényi random graph with 83 nodes and edge probability 0.04.

Figure 7. (Color online) Distribution of the Number of Nodes in the System Under the Greedy Policy with a Chain and $p = 0.04$



the greedy policy requires the probability of three-way cycle formation to be exactly $1/3$ in each period and no larger, since it executes the three-way cycle immediately. Hence, the greedy policy should be optimal (at least among policies such that the residual graph resembles an Erdős–Rényi graph in terms of average degree).⁸

Finally, we briefly remark on the mixing time of the Markov chain under the greedy policy. In a model similar to our two-way cycle setting, Akbarpour et al. (2017) show that their Markov chain mixes in a time that is equivalent to $O((1/p^2) \log(1/p))$ in our parametrization. Though we do not study this formally, we believe that in our two-way cycle setting under a greedy policy, the Markov chain for even time steps (and also the one for odd time steps) mixes in time of the same order, i.e., $O((1/p^2) \log(1/p))$. In particular, this is a small multiple of the average waiting time, which is $O(1/p^2)$ in that setting. We also believe that the Markov chains under greedy for the two-way and three-way cycles setting, and for the chains setting, each mix in a time that is a small multiple of the average waiting time in each case. Our numerical findings appear to suggest that this holds, in that the degree distribution closely resembles that of an Erdős–Rényi graph (after a short

initial period; see Figure 6), and the number of nodes in the system returns often to each typical value. Formally proving such results would be quite involved in these settings, however, since the state of the market is the entire compatibility graph, as compared to just the number of waiting nodes in Akbarpour et al. (2017).

5. Proof Ideas for the Main Results

In the following sections we give the main ideas behind the proofs. The proof of Theorem 1 is relatively simple and provided in its entirety in Section 5.1. All other proofs are deferred to the appendices.

5.1. Bilateral Matching

In this section we consider the case in which exchanges are done through two-way cycles only. We refer the reader to the introduction for rough intuition regarding the optimality of the greedy policy. Our proof follows a different argument based on the likelihood of two-way cycle formation. We first sketch the proof of Theorem 1, which provides tight bounds.

The lower bound is straightforward: in steady state, any policy must remove a cycle at least once every two periods on average, and so the rate of departures matches the rate of arrivals. As a result, we must have that an incoming node should form a two-way cycle with *some* existing node with probability at least $1/2$ on average. Note that the probability that the incoming node forms a two-way cycle with a particular existing node is p^2 . Therefore the probability of formation of some two-way cycle is exactly $1 - (1 - p^2)^{|V(t)|}$. This immediately leads to a lower bound of $\log(2)/(-(-1 - p^2)) = \log(2)(1 + o(1))/p^2$ on the expected number of nodes in the system, under any policy in steady state.

Now consider the greedy policy. The number of nodes in the system under the greedy policy behaves as a simple random walk with a downward step occurring exactly when the new node forms a two-way cycle.

Thus, the random walk has a negative drift when the probability of the new node forming a cycle with an existing node exceeds $1/2$ by even a little. This ensures that $|\mathcal{V}(t)|$ does not grow much beyond the minimum level needed to ensure that new arrivals form two-way cycles with a probability at least half, and so the performance of the greedy policy closely matches the lower bound.

Proof of Theorem 1. We first compute the expected steady-state waiting time under the greedy policy. Observe that for all $t \geq 0$,

$$|\mathcal{V}(t+1)| = \begin{cases} |\mathcal{V}(t)| + 1 & \text{with probability } (1-p^2)^{|\mathcal{V}(t)|}, \\ |\mathcal{V}(t)| - 1 & \text{with probability } 1 - (1-p^2)^{|\mathcal{V}(t)|}. \end{cases}$$

Let $\epsilon > 0$ be arbitrary. If $|\mathcal{V}(t)| > (1 + \epsilon) \ln(2)/p^2$, then there exists a sufficiently small $p = p(\epsilon)$ such that for all $p > p(\epsilon)$,

$$\mathbb{P}(|\mathcal{V}(t+1)| = |\mathcal{V}(t)| + 1) = (1-p^2)^{|\mathcal{V}(t)|} \leq \frac{1}{2^{1+\epsilon}}.$$

Let $q = 1/2^{1+\epsilon} < 1/2$, and let X_t be a sequence of i.i.d. random variables with distribution

$$X_t = \begin{cases} 1 & \text{with probability } q, \\ -1 & \text{with probability } 1 - q. \end{cases}$$

Let $S_0 = 0$ and for $t \geq 1$, $S_{t+1} = (S_t + X_t)^+$, and so S_t is a birth-death process. Letting $r = q/(1-q) < 1$, in steady-state $\mathbb{P}(S_\infty = i) = r^i(1-r)$ for $i = 0, 1, \dots$, and so

$$\mathbb{E}[S_\infty] = r/(1-r) = q/(1-2q) = \frac{1}{2^{1+\epsilon} - 2}.$$

We can couple the random walk $|\mathcal{V}(t)|$ with S_t such that $|\mathcal{V}(t)| \leq (1 + \epsilon) \ln(2)/p^2 + S_t$ for all t . (When $|\mathcal{V}(t)| < (1 + \epsilon) \ln(2)/p^2$, then we simply use the fact that the number of nodes can increase by at most 1 per time step. On the other hand, if $|\mathcal{V}(t)| \geq (1 + \epsilon) \ln(2)/p^2$ then X_t stochastically dominates the step size $|\mathcal{V}(t+1)| - |\mathcal{V}(t)|$ by construction.) This yields

$$\mathbb{E}[|\mathcal{V}(\infty)|] \leq (1 + \epsilon) \frac{\ln(2)}{p^2} + \mathbb{E}[S_\infty] \leq (1 + \epsilon) \frac{\ln(2)}{p^2} + \frac{1}{2^{1+\epsilon} - 2}.$$

Thus for every $\epsilon > 0$, we have

$$\lim_{p \rightarrow 0} \frac{\mathbb{E}[|\mathcal{V}(\infty)|] - \ln(2)/p^2}{1/p^2} \leq \epsilon \ln(2).$$

As ϵ is arbitrary, the result follows.

Now we establish the lower bound on $|\mathcal{V}(\infty)|$. Let v be a newly arriving node at time t , and \mathcal{W} be the nodes currently in the system that are waiting to be matched. Let I be the indicator that at the arrival time of v (just before cycles are potentially deleted), no two-way cycles between v and any node in \mathcal{W} exist. Let \tilde{I} be

the indicator that at the arrival time of v , no two-way cycles that will eventually be removed that are between v and any node in \mathcal{W} exist (in particular, \tilde{I} depends on the future). Thus $\tilde{I} \geq I$ a.s. Let \tilde{V}_t be the number of vertices in the system before time t such that the cycle that eventually removes them has not yet arrived. We let \tilde{V}_∞ be the distribution of \tilde{V}_t when the system begins in steady state. By stationarity,

$$0 = \mathbb{E}[\tilde{V}_{t+1} - \tilde{V}_t] = \mathbb{E}_{\tilde{V}_\infty} [2\tilde{I} - 1],$$

giving $\mathbb{E}[\tilde{I}] = 1/2$. Intuitively, in steady state, the expected change in the number of vertices not yet “matched” must be zero. Thus we obtain

$$\begin{aligned} \frac{1}{2} &= \mathbb{E}[\tilde{I}] \geq \mathbb{E}[I] = \mathbb{E}[\mathbb{E}[I \mid |\mathcal{V}(\infty)|]] \\ &= \mathbb{E}[(1-p^2)^{|\mathcal{V}(\infty)|}] \geq (1-p^2)^{\mathbb{E}[|\mathcal{V}(\infty)|]}, \end{aligned}$$

by Jensen’s inequality. Taking logarithms on both sides and rearranging terms, we get

$$\mathbb{E}[|\mathcal{V}(\infty)|] \geq \frac{\log(1/2)}{\log(1-p^2)} = \frac{\log(2)}{-\log(1-p^2)}. \quad \square$$

5.2. Chain Removal

In this section we provide the main ideas in the proof of Theorem 3 (the formal proof is given in Appendix EC.3 in the online appendix). First we give the high-level intuition. Recall that at any time period there is a single bridge node in the marketplace and under the greedy policy, the chain advances once the arriving agent can accept the item of the bridge agent. The basic idea in establishing that the greedy policy achieves an $O(1/p)$ average waiting time is to show that when there are more than C/p nodes in the system, on average, the next chain segment will contain more nodes than were added in the interim. This “negative drift” in the number of nodes is crucial in establishing the bound (we use a Lyapunov argument based on Proposition EC.2 in the online appendix to infer a bound on the expected waiting time). The lower bound, proved by contradiction, is based on the idea that the waiting time for a node must be at least the time for the node to get an indegree of one.

5.2.1. Sketch of Proof of the Upper Bound. To provide further intuition we first introduce the following notation. Let the compatibility graph at time t be $\mathcal{G}(t) = (\mathcal{V}(t), \mathcal{E}(t), h(t))$. Here $h(t)$ is a special node not included in $\mathcal{V}(t)$ that is the bridge node, which can form outgoing edges only. We denote by $\mathcal{G}(\infty) = (\mathcal{V}(\infty), \mathcal{E}(\infty), h(\infty))$ the steady-state version of this graph (more precisely, it is a random variable drawn from the steady-state distribution of $\mathcal{G}(t)$).

According to the greedy policy, whenever $h(t)$ forms a directed edge to a newly arriving node, a largest possible chain segment starting from $h(t)$ is executed. Note

that $h(t)$ has indegree and outdegree of zero from the time it becomes a bridge node until it forms a directed edge to the new arriving node. We refer to this period between chain segments as an *interval*. Let τ_i for $i = 1, 2, \dots$, denote the length of the i th interval, and so $\tau_i \sim \text{Geometric}(p)$. Let $T_0 = 0$ and $T_i = \sum_{j=1}^i \tau_j$ for $i = 1, 2, \dots$, be the time at the end of the i th interval. Additionally, let \mathcal{A}_i be the set of nodes that arrived during the i th interval $(T_{i-1}, T_i]$, in particular $|\mathcal{A}_i| = \tau_i$, and let \mathcal{W}_i be the set of nodes that were “waiting” at the start of the i th interval (excluding the bridge node), namely, after time T_{i-1} . Thus, right before the chain is advanced, every node in the graph is either in \mathcal{W}_i , \mathcal{A}_i or it is $h(t)$ itself.

The idea of establishing the upper bound on the average waiting time under the greedy policy is as follows. We use a Lyapunov function argument (using Proposition EC.2 in the online appendix) to argue that for some C , if there are at least C/p nodes in the graph at the start of an interval $(T_i, T_i + \tau_{i+1}]$, then the number of nodes deleted in the subsequent chain segment is on average greater than the number of nodes τ_{i+1} that arrive in the corresponding interval; i.e., we have a negative drift on the number of nodes. We lower bound the number of nodes removed in the i th interval by the length of a longest path starting from the bridge node in the bipartite graph formed by one vertex set being \mathcal{A}_i (the newly arrived nodes) and the other vertex set being \mathcal{W}_i along with the bridge node (the nodes waiting after the previous chain segment), while retaining only edges between the two vertex sets (in particular those leading to a bipartite structure). We lower bound the expected size of a longest path on this subgraph (Corollary EC.3 in the online appendix).

Remark 4. We conjecture that the performance of a greedy policy that executes two- and three-way cycles in addition to chains will be almost identical to that of a greedy policy with chains only, and, in particular, the expected waiting time will still be $O(1/p)$.

The reasoning is as follows. It is not hard to see that the greedy policy that executes only chains misses two- and three-way cycle opportunities for only an $O(p)$ fraction of nodes. Also, the lower bound in Theorem 3 means that the waiting time must be $\Omega(1/p)$, even if some cycles are removed.

5.2.2. Sketch of Proof of the Lower Bound. The lower bound in Theorem 3 is proved by contradiction, which appears to be a novel approach. The waiting time for a node must be at least the time for the node to get an indegree of one. Using this, if the steady-state average waiting time is $W = o(1/p)$, then by Little’s law when a typical node v arrives there are only $o(1/p)$ nodes in the system, and so v is likely not to have any incoming edges connecting with any of these existing nodes. After κW steps, the number of newly arrived nodes is

$W = o(1/p)$, and so v is likely not to connect with any of these nodes either. If we can show that v will be in the system for more than κW steps with high enough probability (i.e., with probability exceeding $1/\kappa$), this will contradict that the expected waiting time in steady state is W . Our proof accomplishes this for $\kappa = 3$.

5.3. Two- and Three-Way Cycles

In this section we sketch the proof of Theorem 2. The proof is far more involved than in the case of only two-way cycles, especially the upper bound. The rigorous proof is in Appendix EC.2 in the online appendix. We also state a conjecture that would refine the result in Theorem 2 in Section 5.3.3.

5.3.1. Sketch of the Proof of the Upper Bound. The proof of the upper bound of $O(1/p^{3/2})$ on waiting time under the greedy policy relies on a delicate combinatorial analysis of three-way cycles in the random graph formed by nodes present in the system in steady state and those arriving over a certain time interval. We consider a time interval $T = \Theta(1/p^{3/2})$ and assume that the system starts with at least order $\Theta(1/p^{3/2})$ nodes in the underlying graph. We establish a negative drift in the system and then, as in the case of chain removal, rely on the Lyapunov function technique in order to establish the required upper bound.

It is very difficult to control the distribution of the residual compatibility graph under a greedy policy. In particular, it is not an Erdős–Rényi random graph. For instance, by definition of the greedy policy it does not contain any two-way or three-way cycles. We do not obtain any control on the edge distribution. Instead we show a negative drift in the number of nodes in the market, regardless of the existing set of edges between the previously waiting nodes. In particular, the negative drift holds even in the intuitively “hard” case where the starting graph contains no edges at all, which is the case we focus on in this proof sketch. Notice that there is no possibility of obtaining a negative drift with a single arrival, even if the number of previously waiting nodes is a large constant multiple of $1/p^{3/2}$. A three-way cycle cannot possibly be formed (since there are no edges between existing nodes), and the probability of forming a two-way cycle is $O(p^2 \cdot 1/p^{3/2}) = O(p^{1/2}) = o(1)$. In fact, it turns out that $T = \Omega(1/p^{3/2})$ arrivals are needed to ensure a negative drift when the starting graph contains no edges. Let \mathcal{W} be the set of previously waiting nodes and let \mathcal{A} be the set of new arrivals.

Consider a new arrival $v \in \mathcal{A}$. Let us estimate the probabilities of v forming a two-way or three-way cycle with other nodes from among the previously waiting nodes and new arrivals. First, note that the probability of two-way cycle formation is $O((T + 1/p^{3/2})p^2) = O(p^{1/2}) = o(1)$ assuming $T = O(1/p^{3/2})$. The probability of three-way cycle formation involving only new

arrivals is $O(T^2 p^3)$ (to form a three-way cycle with the arriving node, two of the $O(T)$ nodes are needed, as well as three directed edges between these three nodes to create a cycle). Moreover, the probability of three-way cycle formation involving one other new arrival and one previously waiting node is $O(p^3 \cdot T \cdot 1/p^{3/2}) = O(T p^{3/2}) = O(T^2 p^3)$ for $T = O(1/p^{3/2})$. Thus, the overall probability of v being part of some two-way or three-way cycle is $O(T^2 p^3) + o(1)$. Any cycle that is removed must involve at least two new arrivals, and removes at most one previously waiting node. To obtain a negative drift, one must have that the expected number of new arrivals that remain after T steps is less than the expected number of previously waiting nodes that are removed. In particular, each new arrival must be removed with probability at least $2/3$, implying that $T^2 p^3 = \Omega(1)$. Thus, we need $T = \Omega(1/p^{3/2})$ to have any hope of obtaining a negative drift.

To make our proof work, we establish a negative drift when $|\mathcal{W}| \geq C^3 p^{3/2}$ (there are many previously waiting nodes in the market) by considering $T = 1/(C p^{3/2})$. By choosing C large enough, we ensure that most of the new arrivals are removed via a three-way cycle containing one previously waiting node and one new arrival, under a greedy policy (call the number of such three-way cycles N). As per our calculations above, two-way cycles play a very small role. The probability that a node in \mathcal{A} is part of some three-way cycle involving only new arrivals is only about $1/C^2$, which is small for large C . Thus, only a few nodes in \mathcal{A} leave via cycles containing only other nodes in \mathcal{A} ; see event \mathcal{E}_2 in the proof. On the other hand, each node is in expectation part of at least $C^3/C = C^2$ cycles consisting of one node in \mathcal{W} and one other node in \mathcal{A} . Performing a more careful analysis, one can check that each new arrival should be removed within a time of about $1/(C^2 p^{3/2})$ after arrival. Thus, we expect about $1/(C^3 p^{3/2})$ new arrivals to remain after T periods. In our proof we show that this number is very unlikely to exceed $2/(C^2 p^{3/2})$, i.e., a fraction $2/C$ of the arrivals \mathcal{A} ; see event \mathcal{E}_1 in the proof. Most of the arrivals that were served by time T were served as part of three-way cycles that involved a node in \mathcal{W} , with every two such nodes leading to a decrease by at least one of the overall number of nodes in the system between time 0 and time T . Thus, the overall number of nodes in the market is typically smaller after T periods than it was at the beginning, for large enough C .

5.3.2. Sketch of the Proof of the Lower Bound. For the lower bound, we adopt a similar approach to the one for chain removal. We prove by contradiction a matching lower bound (up to constants) for batching policies. The rough idea is as follows: if the steady-state expected waiting time is small (in this case smaller than $1/(C p^{3/2})$ for appropriate C), then a typical new arrival sees a small number of nodes currently in the system,

and so typically does not form a two- or three-way cycle with previously waiting nodes or even the next few arrivals. Thus, the typical arrival node has a long waiting time, which contradicts our initial assumption of a small expected waiting time.

5.3.3. A Conjecture That Would Refine Theorem 2. The following characterization of the performance of the greedy policy is obtained if we assume that n_t concentrates, and that the typical number of edges in a compatibility graph at time t with n_t nodes is close to what it would have been under an $ER(n_t, p)$ graph. (We believe these assumptions are valid. We also believe that the greedy policy is very close to optimal, this being the second part of the statement below.)

Conjecture 1. *For cycle removal with $k = 3$, the expected waiting time in steady state under a greedy policy scales as $\sqrt{\ln(3/2)}/p^{3/2} + o(1/p^{3/2})$, and no periodic Markov policy (including nonbatching policies) can achieve an expected waiting time that scales better than this.*

Here the constant $\sqrt{\ln(3/2)}$ results from requiring (under our assumptions) that a newly arrived node forms a triangle with probability $1/3$.

Our simulation results (see Figure 5 in Section 4) are consistent with this conjecture: the predicted expected waiting time for a greedy policy from the leading term $\sqrt{\ln(3/2)}/p^{3/2}$ is $W = 84.7$ for $p = 0.04$, $W = 47.1$ for $p = 0.06$, $W = 30.7$ for $p = 0.08$, and $W = 22.4$ for $p = 0.1$. The degree distribution observed in the residual compatibility is also very close to that of a directed Erdős–Rényi graph; see Figure 6. If proved, this conjecture would be a refinement of Theorem 2. A proof would require a significantly more sophisticated analysis for both the upper bound and the lower bound.

6. Conclusion

This paper was concerned with *when* a centralized planner should match agents in a marketplace for barter. In particular, we studied how centralized matching policies, as well as matching technologies, affect agents' waiting times in a dynamic market with a homogeneous stochastic demand structure. We found the greedy policy to be approximately optimal when either of two-way cycles, two- and three-way cycles, or chains are feasible. Moreover, three-way cycles and chains lead to a large improvement in waiting times relative to two-way cycles only.

An important marketplace for barter is kidney exchange, which allows incompatible patient-donor pairs to exchange donors. Exchanges in this market take place over time and occur through two-way cycles, three-way cycles, and chains. While our model is stylized and abstracts away from many important details in kidney exchange, our findings are consistent with

computational experiments and practice (see Endnote 1 and Anderson et al. 2015).

While our approach is potentially applicable to exchanges involving cycles longer than three, the technical details appear extremely challenging and we leave the question of sharp characterization of performance under four-way and longer exchanges for future research.

We next discuss several implications of our work. Consider first a setting in which multiple clearinghouses compete with each other. When agents can enroll in more than one clearinghouse, there is an incentive for clearinghouses to complete exchanges quickly to avoid agents completing an exchange in a different clearinghouse. A priori one may worry that this greedy-like behavior will harm social welfare, yet, our findings suggests that this is not the case, as greedy nearly optimizes agents' average waiting time.

Our work has implications also for decentralized marketplaces, where agents find their own exchanges. Such marketplaces typically involve only bilateral exchanges. (There are exceptions, such as multiway swaps of players between teams in the National Basketball Association in the United States.) Our finding that three-way cycles and chains can substantially reduce waiting times suggests that in some contexts, switching to a central matching authority may lead to substantial gains in efficiency.

Our work is only a first step in analyzing dynamic matching policies and leaves many questions open. One natural question is whether the greedy policy remains optimal under other natural settings. Akbarpour et al. (2017) allow agents to perish before being matched and seek to minimize the loss rate. A greedy policy is not optimal if the clearinghouse knows the perishing times, but otherwise it is optimal (see Appendix EC.4 in the online appendix where we establish a close connection between minimizing the loss rate and minimizing the average waiting time). Allowing for heterogeneous agents or goods can potentially lead to different results. For example, if Bob is interested only in Alice's item but there is no exchange in which they both participate, it may be beneficial to make Alice wait in hopes of finding an exchange that can allow Bob to get Alice's item. Thus, when chains or cycles with at least three agents are permitted, some waiting may improve efficiency in the presence of heterogeneity (some evidence for this is given by Ashlagi et al. 2013). Allowing for *preferences* over compatible matches, Baccara et al. (2015) have shown that a non-greedy policy is optimal in a bipartite marketplace.

In kidney exchange, patient-donor pairs can roughly be classified as easy to match and hard to match. (Whether a pair is easy to match or not depends on the blood types, the donor's antigens, and the patient's antibodies.) Many hospitals internally match

their easy-to-match pairs, and enroll their harder-to-match pairs in centralized multihospital clearinghouses (Ashlagi and Roth 2014). An important question is how much the waiting times of hard-to-match pairs will improve as the percentage of easy-to-match pairs grows.

Acknowledgments

The authors thank Ramesh Johari, Assaf Zeevi, participants at the Stanford Market Design workshop, participants at the BIRS Asymptotics of Large-Scale Interacting Networks workshop, the editors, and anonymous referees for their valuable comments to improve the paper.

Endnotes

¹The Alliance for Paired Donation moved from monthly to daily matching, the National Kidney Registry now conducts daily matches after experimenting with longer batches, and the United Network for Organ Sharing have moved from monthly matching to weekly matching.

²One can instead consider a stochastic model of arrivals, e.g., Poisson arrivals in continuous time. In our setting, such stochasticity would leave the behavior of the model essentially unchanged, and indeed, each of our main results extend easily to the case of Poisson arrivals at rate 1.

³Previous models involving stochastic compatibilities (Ashlagi et al. 2012, 2013) require p to scale in a particular way with "market size."

⁴Formally, the sequence of states every τ periods (starting from period l) forms a Markov chain, and we require that this Markov chain is irreducible and positive recurrent. (This turns out not to depend on l .)

⁵The expected total number of three-way cycles is nearly n^3p^3 , and the expected number of node disjoint three-way cycles is of the same order for $n^3p^3 \leq n$. We need $n^3p^3 \sim n$ in order to cover a given fraction of nodes with node disjoint three-way cycles, leading to $n \geq 1/p^{3/2}$. For $n \sim 1/p^{3/2}$, the number of two-way cycles is $n^2p^2 \sim 1/p = o(n)$; i.e., very few nodes are part of two-way cycles.

⁶The lower bound holds even for policies that allow a chain to be implemented in a different fashion where agents first give an item before they receive an item. Moreover, the lower bound holds even if there is a constant number of coexisting chains. It is possible, however, to reduce agents' average waiting time by having, for example, a large number of chains in the system.

⁷We see that the difference between waiting times under chain removal and cycle removal with $k = 3$ is less pronounced. One reason for this could be that there are long intervals between consecutive times when a chain can be advanced, leading to a poor constant factor for chain removal. Using nonmaximal chains may shorten these intervals by retaining multiple possible ways that the chain can grow at any time, and this may improve the constant factor.

⁸This line of argument does not work in the case of chains. We are unsure whether the greedy policy is truly optimal with chains, or just optimal up to some constant factor.

References

- Abdulkadiroğlu A, Sönmez T (1998) Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66(3):689–701.
- Adan I, Weiss G (2012) Exact FCFS matching rates for two infinite multitype sequences. *Oper. Res.* 60(2):475–489.
- Akbarpour M, Li S, Oveis Gharan S (2017) Thickness and information in dynamic matching markets. Available at SSRN: <http://doi.org/10.2139/ssrn.2394319>.

- Anderson R, Ashlagi I, Gamarnik D, Rees M, Roth AE, Sönmez T, Ünver MU (2015) Kidney exchange and the alliance for paired donation: Operations research changes the way kidneys are transplanted. *Interfaces* 45(1):26–42.
- Ashlagi I, Roth AE (2014) Free riding and participation in large scale, multi-hospital kidney exchange. *Theoret. Econom.* 9(3):817–863.
- Ashlagi I, Jaillet P, Manshadi V (2013) Kidney exchange in dynamic sparse heterogenous pools. Preprint arXiv:1301.3509.
- Ashlagi I, Gamarnik D, Rees M, Roth A (2012) The need for (long) chains in kidney exchange. Technical report, National Bureau of Economic Research, Cambridge, MA.
- Asmussen S (2003) *Applied Probability and Queues* (Springer, New York).
- Baccara M, Lee S, Yariv L (2015) Optimal dynamic matching. SSRN 2641670.
- Bollobás B (2001) *Random Graphs*, Vol. 73 (Cambridge University Press, Cambridge, UK).
- Budish EB, Cramton P, Shim JJ (2015) The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quart. J. Econom.* 130(4):1547–1621.
- Caldentey R, Kaplan E, Weiss G (2009) FCFS infinite bipartite matching of servers and customers. *Adv. Appl. Probab.* 41(3):695–730.
- Dickerson JP, Procaccia AD, Sandholm T (2012) Optimizing kidney exchange with transplant chains: Theory and reality. van der Hoek W, Padgham L, Conitzer V, Winikoff M, eds. *Proc. Internat. Conf. Autonomous Agents Multiagent Systems, AAMAS '12*, Vol. 2 (IFAAMA, Richland, SC), 711–718.
- Dur UM, Ünver MU (2012) Tuition exchange. *Proc. 2nd Internat. Workshop Matching Under Preferences, MATCH-UP*, 127.
- Feldman J, Mehta A, Mirrokni VS, Muthukrishnan S (2009) Online stochastic matching: Beating $1-1/e$. *Proc. 50th Annual IEEE Sympos. Foundations Comput. Sci. FOCS '09* (IEEE Computer Society, Washington, DC), 117–126.
- Goel G, Mehta A (2008) Online budgeted matching in random input models with applications to adwords. Teng S-H, ed. *Proc. Nineteenth Annual ACM-SIAM Sympos. Discrete Algorithms SODA '08* (SIAM, Philadelphia), 982–991.
- Gurvich I, Ward A (2014) On the dynamic control of matching queues. *Stochastic Systems* 4(2):479–523.
- Jaillet P, Lu X (2013) Online stochastic matching: New algorithms with better bounds. *Math. Oper. Res.* 39(3):624–646.
- Janson S, Luczak T, Rucinski A (2011) *Random Graphs*, Vol. 45 (John Wiley & Sons, Hoboken, NJ).
- Jevons W (1876) *Money and the Mechanism of Exchange* (D. Appleton and Co., New York).
- Karp RM, Vazirani UV, Vazirani VV (1990) An optimal algorithm for on-line bipartite matching. *Proc. 22nd Annual ACM Sympos. Theory Comput., STOC '90* (ACM, New York), 352–358.
- Kleinrock L (1975) *Theory, Queueing Systems*, Vol. 1 (John Wiley & Sons, New York).
- Little J (1961) A proof for the queuing formula: $L = \lambda w$. *Oper. Res.* 9(3):383–387.
- Manshadi V, Gharan S, Saberi A (2012) Online stochastic matching: Online actions based on offline statistics. *Math. Oper. Res.* 37(4):559–573.
- Mendelson H (1982) Market behavior in a clearing house. *Econometrica* 50(6):1505–1524.
- Pycia M, Ünver U (2017) Incentive compatible allocation and exchange of discrete resources. *Theoret. Econom.* 12(1):287–329.
- Roth A, Postlewaite A (1977) Weak versus strong dominance in a market with indivisible goods. *J. Math. Econom.* 4(2):131–137.
- Roth AE (1982) Incentive compatibility in a market with indivisible goods. *Econom. Lett.* 9(2):127–132.
- Roth AE, Sönmez T, Ünver MU (2005) Pairwise kidney exchange. *J. Econom. Theory* 125(2):151–188.
- Roth AE, Sönmez T, Ünver MU (2007) Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *Amer. Econom. Rev.* 97(3):828–851.
- Shapley L, Scarf H (1974) On the core and indivisibility. *J. Math. Econom.* 1(1):23–37.
- Ünver MU (2010) Dynamic kidney exchange. *Rev. Econom. Stud.* 77(1):372–414.
- Wang Y, Krishna A (2006) Timeshare exchange mechanisms. *Management Sci.* 52(8):1223–1237.

Ross Anderson is a software engineer at Google. He received his PhD from the Operations Research Center at the Massachusetts Institute of Technology. His research interests are in applied probability, optimization, and healthcare applications.

Itai Ashlagi is an assistant professor in management science and engineering at Stanford University. He is interested in topics on the border of microeconomic theory and operations research and in particular topics in market design and game theory.

David Gamarnik is a professor of operations research at the Sloan School of Management, Massachusetts Institute of Technology. His interests include analysis of stochastic models and applications, algorithms and optimization methods.

Yash Kanoria is the Sidney Taurel Associate Professor of Business in the Graduate School of Business, Columbia University. His research focuses on the design and operations of marketplaces, especially matching markets and online platforms.