

Competing Dynamic Matching Markets

Sanmay Das* John P. Dickerson†
*Washington University in St. Louis
{sanmay,zhuoshuli}@wustl.edu

Zhuoshu Li* Tuomas Sandholm†
†Carnegie Mellon University
{dickerson,sandholm}@cs.cmu.edu

ABSTRACT

While dynamic matching markets are usually modeled in isolation, assuming that every agent to be matched enters that market, in many real-world settings there exist *rival* matching markets with overlapping pools of agents. We extend a framework of dynamic matching due to Akbarpour et al. [2] to characterize outcomes in cases where two such rival matching markets compete with each other. One market matches quickly while the other builds market thickness by matching slowly. We give an analytic bound on the loss—the expected fraction of unmatched vertices—of this two-market environment relative to one in which all agents enter either one market or the other, and numerically quantify its exact loss, demonstrating that rival markets increase overall loss compared to a single market that builds thickness. We then look at two competing *kidney exchanges*, where patients with end-stage renal failure swap willing but incompatible donors, and show that matching with rival barter exchanges performs qualitatively the same as matching with rival matching markets—that is, rival markets increase global loss.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms

Economics, Experimentation, Theory

Keywords

Matching markets, dynamic matching, greedy matching, trade frequency, kidney exchange

1. INTRODUCTION

In matching problems, a central clearinghouse pairs agents with other agents, transactions, or contracts. Most classical matching problems—matching medical residents to hospitals, matching students to schools—are static, where agents

and items exist at the same time, are matched, and then the market disappears; however, many real-world matching problems are dynamic, with agents arriving and departing over time in a persistent market.

Furthermore, many dynamic matching applications involve *multiple* competing clearinghouses with overlapping sets of participants. For example, a lonely graduate student may register on two dating websites (e.g., Match.com and OkCupid), or choose to only register on one. Thus, a member of both sites can be matched to any member of either site, while single-site members can only be matched to members of their specific dating market. The clearinghouses then compete on a metric like total number of matches. It is also common for patient-donor pairs in kidney exchange to register on multiple exchanges, an application we explore in detail later.

In this paper, we explore, in a dynamic matching setting, how rival clearinghouses affect global social welfare in terms of total agents matched relative to a world in which all agents enter exactly one market, which can optimize for how to match them independently.

1.1 Our contribution

This paper’s major contribution is the extension of a recent framework of dynamic matching due to Akbarpour et al. [2] to two rival matching markets with overlapping pools. Specifically, we formalize a two-market model where agents enter one market or both markets; they can then be matched to other agents who have joined the same market or both markets. The markets adhere to different matching policies, with one matching greedily and the other building market thickness through a patient policy. We provide an analytic lower bound on the *loss*, or the expected fraction of vertices who enter and leave the pool without finding a match, of the two-market model and show that it is higher than running a single “patient” market. We also provide a quantitative method for determining the loss of the two-market model.

Our work draws motivation from kidney exchange, an instantiation of barter exchange where patients paired with willing but medically incompatible donors swap those donors with other patients. In the United States, multiple fielded kidney exchanges exist, and patient-donor pairs are entered simultaneously into one or more of these markets, based on geographical location, travel preferences, home transplant center preferences, or other logistical reasons. Individual kidney exchange clearinghouses have incentive to compete on number of matches performed within their specific pools; yet, fragmenting the market across multiple exchanges operating under different matching policies may lower global

welfare. In this paper, we provide the first experimental evidence on dynamic kidney exchange graphs showing that this may indeed be the case.

1.2 Related work

Most related to our work is a recent paper by Akbarpour et al. [2], which presents a general framework for bilateral dynamic matching in a single market and analyzes the efficacy of a variety of matching policies over time. We build directly on that framework and delay a more in-depth review of that work until Section 2.

Dynamic matching in a *single* market has been explored in many domain-specific applications. Some examples are given below for both one-sided and two-sided traditional matching markets, as well as for barter exchanges; this list is not exhaustive.

One-sided markets. In these settings, only one side (the *agents*) has preferences over the other (the *items*). Waiting lists are used in many applications as a mechanism for allocating the items, which are scarce resources, to agents. Both agents and items arrive over time, and an agent’s priority for an arriving item can be set by a variety of factors. Examples of waiting list applications include public housing assignment [19, 20, 24] and cadaveric organ allocation [9, 33, 35]. In a two-period dynamic housing allocation problem, agents can either apply for a public good (e.g., a house) in the first stage and receive priority in that stage, or opt out in the first stage and receive priority in the second stage [1]. Other variants of the dynamic housing allocation problem have also been addressed where, e.g., agents arrive and depart and, upon departure, an agent’s allocated item is then given to an existing agent in the waiting pool [10, 23].

Two-sided markets. In two-sided markets, participating agents belong to one of two disjoint sets (e.g., “firms” or “workers”), but an agent on either side will have preferences over those on the other. In online labor marketplaces like oDesk, employers and applicants arrive and depart over time and are interested in finding an acceptable match [5, 18]. In the dynamic school choice problem, schools exist permanently and indefinitely, but students arrive and depart periodically in a discrete time model [22]. Students matched to a school at one time period may be matched to a new school at a different time period. Schools and students have preference orderings over each other, based on the utility provided to one side by being allocated an element or elements of the other. Finally, generalizations of the online bipartite matching problem as originally introduced by Karp et al. [21] have recently seen great real-world impact in Internet ad allocation [25, 26].

We note that our work does not assume a bipartite structure in the matching graph—as in [1, 5, 10, 17–20, 22–24, 33–35] and much of the static matching mechanism design literature—and involves more than a single market.

Barter exchange. In barter exchange, agents can directly swap goods with other agents in cycles of length greater than or equal to two. One fielded example is kidney exchange [29], where patients with end-stage renal failure and willing but incompatible paired donors swap those donors with other patients. Ünver [34] was the first to address dynamic kidney exchange, where patient-donor pairs arrive and depart over time, with recent follow-up work by Ashlagi et al. [6] and Anderson et al. [3]. All three papers look at matching policies that aim to maximize (discounted) social welfare. Partic-

ularly relevant to real-world kidney exchanges are *batching* policies, where a market clearing occurs at a fixed interval; some theoretical and empirical explorations of this class of policy has been performed [3, 4, 6, 8]. Learning approaches have also been used to determine more complex matching policies that adhere to specific data distributions [13, 14].

To our knowledge, no work in the general barter exchange or kidney exchange literature has addressed multiple competing exchanges, a problem that is especially relevant in the US now, and, as kidney exchanges move to international swapping, will soon become relevant worldwide.

Interacting mechanisms have been studied in a variety of domains like auctions [12, 27], adaptations of settings from the classical multi-agent systems literature [32], and in two-sided networks that typically exhibit winner-take-all dynamics, where only one or a few large players (e.g., credit card companies, computer operating systems, HMOs) prevail due to network effects [15]. Recent work by Ostrovsky [28] generalizes traditional two-sided matching to a supply chain model with interconnected markets represented as nodes in a path, such that an “upstream” neighbor’s supply overlaps with its “downstream” neighbor’s demand. Relatively little work focuses on markets competing based on variable scheduling or clearing policies, with notable exceptions in cloud or grid computing [7] and in financial markets [11]. To the best of our knowledge, no work looks at competition between two markets in a general framework of dynamic bilateral matching, as this paper does.

2. GREEDY AND PATIENT EXCHANGES

We begin by restating some of the most important results of Akbarpour et al. [2], which will serve as the foundation for our model of competing exchanges. Akbarpour et al. analyze “greedy” and “patient” matching policies—and interpolations between the two—by building stochastic continuous-time bilateral matching models of exchanges running these policies, then measuring the efficacy of the policies in terms of discounted social welfare.

More specifically, an exchange is running in the continuous-time interval $[0, T]$, with agents arriving according to a Poisson process with rate parameter $m \geq 1$. The exchange determines whether potential bilateral transactions between agents are either acceptable or unacceptable. The probability of an *acceptable* transaction existing between any pair of distinct agents is defined as d/m , $0 \leq d \leq m$, and is independent of any other pair of agents in the market. Each agent a remains in the market for a *sojourn* $s(a)$ drawn independently from an exponential distribution with rate parameter $\lambda = 1$; the agent becomes *critical* immediately before her sojourn ends, and this criticality is known to the exchange. An agent leaves either upon being matched successfully by the exchange or upon becoming critical and remaining unmatched, at which point she perishes.

At any time $t \geq 0$, the network of acceptable transactions among agents forms a random graph $G_t = (A_t, E_t)$, where the agents in the exchange at time t form the vertex set A_t , and the acceptable transactions between agents forms the edge set E_t . We assume $A_0 = \emptyset$. Let A_t^n denote the set of agents who enter the exchange at time t , such that with probability 1, $|A_t^n| \leq 1$ for any $t \geq 0$. Finally, let $A = \cup_{t \leq T} A_t^n$.

Akbarpour et al. [2] present a parameterized space of online *matching policies*, with a focus specifically on two:

Patient and Greedy. (In the next section, we will present a novel model of two overlapping exchanges, one running the Patient policy and the other running the Greedy policy.) As described above, vertex arrivals are treated as a continuous-time stochastic process. These policies behave as follows.

Greedy. The Greedy matching algorithm attempts to match each entering agent immediately by selecting one of its neighbors (if a neighbor exists at the time of entry) uniformly at random. One obvious consequence of this is that the remaining graph of unmatched agents at any instant is always empty. We refer to a market running this policy as the *Greedy market* or simply *Greedy* for the rest of the paper.

Patient. The Patient matching algorithm attempts to match each agent only at the instant she becomes critical. As with Greedy, if a critical agent has multiple neighbors, only one is selected uniformly at random. We refer to a market running the Patient policy as a *Patient market* or simply *Patient* when appropriate.

If the random graph model is Erdős-Rényi [16] when not considering arrivals, departures, and matching, then the remaining graph at any instant is also Erdős-Rényi with parameter d/m ; furthermore, d is the average degree of the agents. Both the Patient and Greedy policies maintain this observation.

The main result of Akbarpour et al. [2] is that waiting to thicken the market can be substantially more important than increasing the speed of transactions. Formally, the Patient exchange dramatically reduces the number of agents who perish (and thus leave the exchange without finding a match) compared to the Greedy exchange.

In the Akbarpour et al. [2] paper, an agent a receives zero utility if she perishes, or $u(a) = 0$. If she is matched, she receives a utility of 1 discounted at rate δ , or $u(a) = e^{-\delta s(a)}$. In this work, we focus on the special case of $\delta = 0$ in this paper (i.e., we only consider whether or not an agent is matched), and leave the $\delta \neq 0$ case for future research. Let $\text{ALG}(T) := \{a \in A : a \text{ is matched by ALG by time } T\}$. Then, in this model, the loss of an algorithm ALG is defined as the ratio of the expected number of perished agents to the expected size of A , as shown in Equation 1.

$$\mathbf{L}(\text{ALG}) = \frac{E[|A - \text{ALG}(T) - A_T|]}{E(|A|)} = \frac{E[|A - \text{ALG}(T) - A_T|]}{mT} \quad (1)$$

At any time $t \in [0, T]$, let $Z_{g,t}, Z_{p,t}$ represent the size of the pools under the Greedy and Patient matching policies, respectively. Then, Akbarpour et al. [2] proved that the Markov chain on $Z_{\cdot,t}$ has a unique stationary distribution under either of those policies. Furthermore, let $\pi_g, \pi_p : \mathbb{N} \rightarrow \mathbb{R}^+$ be the unique stationary distribution of the Markov chain on $Z_{g,t}, Z_{p,t}$, respectively, and let $\xi_g := \mathbb{E}_{Z_g \sim \pi_g}[Z_g], \xi_p := \mathbb{E}_{Z_p \sim \pi_p}[Z_p]$ be the expected size of the pool under the stationary distribution under Greedy and Patient. Then, the following observations can be made.

Loss of Greedy. If a Greedy exchange is run for a sufficiently long time, then $\mathbf{L}(\text{Greedy}) \approx \frac{\xi_g}{m}$. The intuition here is that the Greedy pool is (almost) always an empty graph. Equation (2) formalizes the loss.

$$\mathbf{L}(\text{Greedy}) = \frac{1}{mT} \mathbb{E} \left[\int_0^T Z_{g,t} dt \right] = \frac{1}{mT} \int_0^T \mathbb{E}[Z_{g,t}] dt \quad (2)$$

Loss of Patient. If a Patient exchange is run for a sufficiently long time, at any point in time it is an Erdős-Rényi

random graph. So once an agent becomes critical, she has no acceptable transaction with probability $(1 - d/m)^{Z_{p,t}-1}$. Thus, $\mathbf{L}(\text{Patient}) \approx \frac{\xi_p(1-d/m)^{\xi_p-1}}{m}$. Equation (3) formalizes the loss of a Patient market.

$$\begin{aligned} \mathbf{L}(\text{Patient}) &= \frac{1}{mT} \mathbb{E} \left[\int_0^T Z_{p,t} (1 - d/m)^{Z_{p,t}-1} dt \right] \\ &= \frac{1}{mT} \int_0^T \mathbb{E} \left[Z_{p,t} (1 - d/m)^{Z_{p,t}-1} \right] dt \end{aligned} \quad (3)$$

3. OVERLAPPING EXCHANGES

The key result of Akbarpour et al. [2] is that a greedy dynamic matching market leads to significantly lower global social welfare than a patient matching market with full knowledge of criticality. The central question of this paper is what happens in a situation where a greedy exchange and a patient exchange exist *simultaneously* and compete with each other to match some shared portion of the population. Agents in this overlapping subset of the population join both exchanges simultaneously and accept the first match offer from either of the constituent exchanges.

Drawing on Section 2, we model this in a similar stochastic, continuous-time framework as follows. Agents arrive at the Competing market (a model for the whole system, incorporating both the Greedy and Patient exchanges) at some rate m according to a Poisson process. For each agent, the probability of entering *both* the Greedy exchange and the Patient exchange is γ , the probability of entering the Greedy exchange alone is $(1-\gamma)\alpha$, and the probability of entering the Patient exchange alone is $(1-\gamma)(1-\alpha)$, where $\gamma, \alpha \in [0, 1]$. The probability that a bilateral transaction between each pair of agents is acceptable remains d/m , *conditioned on both agents being mutually "visible" to an exchange*. The agents' rates of perishing, received utility for being (un)matched, and other settings are otherwise the same as in Section 2.

We analyze the Competing market as three separate evolving pools:

Greedy_c is the pool consisting of agents who enter the Greedy exchange only (with probability $\alpha(1-\gamma)$).

Patient_c is the pool consisting of agents who enter the Patient exchange only (with probability $(1-\alpha)(1-\gamma)$).

Both_c is the pool consisting of agents who enter both exchanges (with probability γ).

We use $\hat{Z}_{g,t}, \hat{Z}_{p,t}$ and $\hat{Z}_{b,t}$ to denote the size of Greedy_c, Patient_c and Both_c, respectively, at any time t . Similar to an exchange running a single Greedy or Patient matching policy, the Markov chain on $\hat{Z}_{\cdot,t}$ also has a unique stationary distribution. Let $\hat{\pi} : \mathbb{N} \rightarrow \mathbb{R}^+$ be the unique stationary distribution of the Markov chain on $\hat{Z}_{\cdot,t}$, and let $\hat{\xi} := \mathbb{E}_{\hat{Z}_{\cdot,t} \sim \hat{\pi}}[\hat{Z}_{\cdot,t}]$ be the expected size of the pool under the stationary distribution. Using this, we will define the loss of Greedy_c, $\hat{\mathbf{L}}(\text{Greedy}_c)$, the loss of Patient_c, $\hat{\mathbf{L}}(\text{Patient}_c)$, and the loss of Both_c, $\hat{\mathbf{L}}(\text{Both}_c)$.

First, note that the graph formed by the agents in Greedy_c is empty, so the loss—as in Equation (2)—can be approximated by $\hat{\mathbf{L}}(\text{Greedy}_c) \approx \frac{\hat{\xi}_g}{m}$.

Next, we consider the agents in Both_c. If an edge exists between an agent in Both_c and an existing agent in Greedy_c or another agent in Both_c, she will be matched immediately by the Greedy exchange (and thus does not contribute

to the loss). Similar to the Greedy_c case, at any point in time t , the Both_c pool is an empty graph; thus, any unmatched agents who become critical in Both_c will only be matched to agents in Patient_c. Thus, these leftover agents in Both_c have no acceptable transactions with probability $(1 - d/m)^{\hat{Z}_{p,t}}$. Since each agent becomes critical with rate 1, letting Competing market run for a sufficiently long time results in $\hat{\mathbf{L}}(\text{Both}_c) \approx \frac{\hat{\xi}_b(1-d/m)^{\hat{\xi}_p}}{m}$, where $\hat{\xi}_b, \hat{\xi}_p$ are the previously defined expected sizes of Both_c and Patient_c.

Finally, we consider the Patient_c pool. At any time t , the agents who remain in Patient_c potentially have acceptable transactions with only the agents in Both_c and the agents in Patient_c. Hence, in $\hat{Z}_{p,t}$, once an agent is critical, she has no acceptable transactions with probability $(1 - d/m)^{\hat{Z}_{p,t} + \hat{Z}_{b,t} - 1}$. Similarly, each agent becomes critical with rate 1; thus, if we allow the Competing market a sufficiently long execution window, $\hat{\mathbf{L}}(\text{Patient}_c) \approx \frac{\hat{\xi}_p(1-d/m)^{\hat{\xi}_p + \hat{\xi}_b - 1}}{m}$.

Because the three pools of agents—Greedy_c, Patient_c, and Both_c—are disjoint (although they may be connected via possible transactions in the ways listed above), we can define the total loss of the Competing market as follows.

$$\mathbf{L}(\text{Competing}) \approx \frac{\hat{\xi}_g + \hat{\xi}_p(1-d/m)^{\hat{\xi}_p + \hat{\xi}_b - 1} + \hat{\xi}_b(1-d/m)^{\hat{\xi}_p}}{m}. \quad (4)$$

A more precise version of Equation (4) follow as Equation (5); we will make use of this form in Section 5.

$$\begin{aligned} \mathbf{L}(\text{Competing}) &= \frac{1}{mT} \mathbb{E} \left[\int_0^T \hat{Z}_{p,t} (1-d/m)^{\hat{Z}_{p,t} + \hat{Z}_{b,t} - 1} \right. \\ &\quad \left. + \hat{Z}_{b,t} (1-d/m)^{\hat{Z}_{p,t}} + \hat{Z}_{g,t} dt \right] \\ &= \frac{1}{mT} \int_0^T \mathbb{E} [\hat{Z}_{p,t} (1-d/m)^{\hat{Z}_{p,t} + \hat{Z}_{b,t} - 1} \\ &\quad + \hat{Z}_{b,t} (1-d/m)^{\hat{Z}_{p,t}} + \hat{Z}_{g,t}] dt \end{aligned} \quad (5)$$

Unfortunately, we do not have a closed form expression for the stationary distribution or the expected size of the pool under the stationary distribution. We note that each of $\hat{\xi}_g, \hat{\xi}_p$, and $\hat{\xi}_b$ can be approximated well using Monte Carlo simulations—thus, Equation (4) can be solved numerically. We do this in Section 5.1 for two parameterizations of the rival market setting.

4. A BOUND ON TOTAL LOSS

While we do not have a closed form for the exact expected loss of the Competing market as described by Equation (4), we can provide bounds on the overall loss. In this section, we give one such bound for the global loss under the constraint that Greedy_c is more likely to receive agents than the overlapping Both_c exchange. Formally, this occurs when $\gamma \leq 0.5$ and $\alpha \geq \frac{\gamma}{1-\gamma}$. We also impose some loose requirements on the arrival rate of vertices to the exchange and the probability of an acceptable transaction existing between two agents; intuitively, the exchange cannot be “too small” or “too sparse,” which we formalize below. Under these assumptions, we use the bound to prove Theorem 1, which states that a single Patient market outperforms the Competing market.

THEOREM 1. *Assume $\gamma \leq 0.5$, $m > 10d$, and $\alpha(1-\gamma) \geq \max \left\{ \gamma, \frac{1}{2}e^{-d/2}(1+3d) \right\}$. Then, as $m \rightarrow \infty$ and $T \rightarrow \infty$, almost surely*

$$\mathbf{L}(\text{Competing}) > \mathbf{L}(\text{Patient}).$$

PROOF. We prove the theorem by giving a lower bound on $\hat{\mathbf{L}}(\text{Greedy}_c)$, the loss of only the greedy portion of the Competing market. In our model, the fraction of agents entering only the Greedy_c side of the market is $\alpha(1-\gamma)$; for notational simplicity, we use $x := \alpha(1-\gamma)$ in this proof. Similarly, the fraction of agents entering Both_c is γ ; again, for notational simplicity, we use $y := \gamma$ throughout this proof.

As before, let $\hat{Z}_{g,t}$ be the size of Greedy_c at any $t \in [0, T]$, and $\hat{\tau}$ the expected size of the Greedy_c pool. Similarly, let $\hat{Z}_{b,t}$ be the size of Both_c at any $t \in [0, T]$, and $\hat{\eta}$ the expected size of the Both_c pool. That is,

$$\hat{\tau} := \mathbb{E}_{t \sim \text{unif}[0, T]} [\hat{Z}_{g,t}] \quad \text{and} \quad \hat{\eta} := \mathbb{E}_{t \sim \text{unif}[0, T]} [\hat{Z}_{b,t}].$$

By assumption, $\alpha(1-\gamma) \geq \gamma$; that is, the arrival rate of Greedy_c is greater than or equal to the arrival rate of Both_c. In this case, $\hat{\tau} \geq \hat{\eta}$; the Greedy matching policy removes vertices from both Both_c and Greedy_c, while the Patient matching policy removes vertices from only Both_c, which means the matching rate for Both_c is greater than the matching rate for Greedy_c.

From Akbarpour et al. [2], we know the expected rate of perishing of the individual Greedy exchange is equal to the pool size because the Greedy matching policy does not react to the criticality of an agent at any time t in its pool and each critical agent will perish with probability 1. Therefore, we can draw directly on Equation (2) to write

$$\hat{\mathbf{L}}(\text{Greedy}_c) = \frac{1}{xmT} \mathbb{E} \left[\int_{t=0}^T dt \hat{Z}_{g,t} \right] = \frac{\hat{\tau}}{xm}. \quad (6)$$

We know x and m , so lower bounding $\hat{\tau}$ will result in an analytic lower bound on $\hat{\mathbf{L}}(\text{Greedy}_c)$. Following the ideas of Akbarpour et al. [2], we do this by lower bounding the probability that an agent a does not *ever* have an acceptable transaction for the duration of her sojourn $s(a)$. Because these agents cannot be matched by any matching policy, this directly gives a lower bound on $\hat{\mathbf{L}}(\text{Greedy}_c)$. Toward this end, fix an agent $a \in A$ who enters Greedy_c at time $t_0 \in \text{unif}[0, T]$ and draws a sojourn $s(a) = t$. Let $f_{s_a}(t)$ be the probability density function at t of $s(a)$. Then we can write the probability that a will never have a neighbor (i.e., possible match) as

$$\begin{aligned} \mathbb{P}[N(a) = \emptyset] &= \int_{t=0}^{\infty} f_{s_a}(t) \mathbb{E} \left[(1-d/m)^{\hat{Z}_{g,t_0} + \hat{Z}_{b,t_0}} \right] \\ &\quad \mathbb{E} \left[(1-d/m)^{AG_{t_0, t_0+t}^n + AB_{t_0, t_0+t}^n} \right] dt, \end{aligned}$$

where AG_{t_0, t_0+t}^n (resp. AB_{t_0, t_0+t}^n) denotes the set of agents who enter Greedy_c (resp. Both_c) in time interval $[t_0, t_0+t]$. The first expectation captures the probability that agent a has no matching at the moment of entry and the second expectation considers the probability that no new agents that can match with a arrive during her sojourn.

Using Jensen's inequality, we have

$$\begin{aligned} \mathbb{P}[N(a) = \emptyset] &\geq \int_{t=0}^{\infty} e^{-t} (1-d/m)^{\mathbb{E}[\hat{Z}_{g,t_0} + \hat{Z}_{b,t_0}]} \\ &\quad (1-d/m)^{\mathbb{E}[|AG_{t_0,t+t_0}^n + AB_{t_0,t+t_0}^n|]} dt \\ &= \int_{t=0}^{\infty} e^{-t} (1-d/m)^{\hat{\tau} + \hat{\eta}} (1-d/m)^{(x+y)mt} dt. \end{aligned}$$

From the assumptions in the theorem statement, $\frac{d}{m} < \frac{1}{10}$, so $1-d/m \geq e^{-d/m-d^2/m^2}$. Also, as described earlier, $\hat{\tau} \geq \hat{\eta}$ (when $\gamma \leq 0.5$ and $\alpha \geq \frac{\gamma}{1-\gamma}$, as assumed). Therefore,

$$\begin{aligned} \hat{\mathbf{L}}(\text{Greedy}_c) &\geq \mathbb{P}[N(a) = \emptyset] \\ &\geq e^{-(\hat{\tau} + \hat{\eta})(d/m + d^2/m^2)} \\ &\quad \times \int_{t=0}^{\infty} e^{-t - (x+y)td - (x+y)d^2t^2/m} dt \\ &\geq \frac{1 - (\hat{\tau} + \hat{\eta})(1+d/m)d/m}{1 + (x+y)d + (x+y)d^2/m} \\ &\geq \frac{1 - 2\hat{\tau}(1+d/m)d/m}{1 + (x+y)d + (x+y)d^2/m}, \end{aligned} \quad (7)$$

where the third inequality is obtained from the fact that $e^{-z} \geq 1 - z$ when $z \geq 0$, here $z = (\hat{\tau} + \hat{\eta})(d/m + d^2/m^2)$.

Combining Equation (6) and Equation (7),

$$\hat{\mathbf{L}}(\text{Greedy}_c) = \frac{\hat{\tau}}{xm} \geq \frac{1 - 2\hat{\tau}(1+d/m)d/m}{1 + (x+y)d + (x+y)d^2/m},$$

which gives us a lower bound for $\hat{\tau}$,

$$\hat{\tau} \geq \frac{xm}{1 + (3x+y)d + (3x+y)d^2/m}.$$

Thus, as $m \rightarrow \infty$, we get,

$$\begin{aligned} \hat{\mathbf{L}}(\text{Greedy}_c) &\geq \frac{1}{1 + (3x+y)d + (3x+y)d^2/m} \\ &\geq \frac{1}{1 + 3d}. \end{aligned}$$

We are interested in bounding the total loss of the Competing market, which is $\mathbf{L}(\text{Competing}) = x\hat{\mathbf{L}}(\text{Greedy}_c) + (1-\alpha)(1-\gamma)\hat{\mathbf{L}}(\text{Patient}_c) + y\hat{\mathbf{L}}(\text{Both}_c)$. By definition, both $\hat{\mathbf{L}}(\text{Patient}_c) \geq 0$ and $\hat{\mathbf{L}}(\text{Both}_c) \geq 0$, and by Equation (4), $\hat{\mathbf{L}}(\text{Greedy}_c) \geq \frac{1}{1+3d}$. Thus,

$$\mathbf{L}(\text{Competing}) \geq \frac{x}{1+3d}.$$

Akbarpour et al. [2] showed that running an *individual* Patient market results in exponentially small loss $\mathbf{L}(\text{Patient}) < \frac{1}{2}e^{-d/2}$. Thus, as $T, m \rightarrow \infty$, we can get,

$$\mathbf{L}(\text{Competing}) > \mathbf{L}(\text{Patient}). \quad (8)$$

□

We note that the result of Theorem 1 holds for only a section of the possible parameterizations of a Competing market—specifically, when $\gamma \leq 0.5$ and $\alpha \geq \frac{\gamma}{1-\gamma}$. In the next section, we will give numerical results showing that this result—that the loss of the Competing market is greater than the loss of an individual Patient exchange—appears to hold for a vastly larger space of values of γ and α . Indeed, experimentally, we will see that the loss of the Competing market is sometimes greater than the loss of an individual *Greedy* exchange, which itself is substantially greater than the loss of an individual Patient exchange.

5. EXPERIMENTAL VALIDATION

In this section, we provide experimental validation of the theoretical results presented in Sections 3 and 4. Section 5.1 quantifies the loss due to competing markets as described by Equation (4), while Section 5.2 expands the model to *kidney exchange* and draws from realistic data to quantify the loss of competing kidney exchange clearinghouses.

5.1 Dynamic matching

In Section 3, we gave a method for computing the expected loss due to competing markets as Equation (4); however, we were unable to derive closed forms for the expected size of the competing, patient, and greedy pools (ξ_b , ξ_p , and ξ_g , respectively) under the stationary distribution. These quantities can be estimated using Monte Carlo simulation for different entrance rates m . We do that now.

Figures 1 and 2 simulate agents entering the Greedy_c , Both_c , and Patient_c markets according to a Poisson process with rate parameter $m = 1000$ and remaining for a sojourn drawn from an exponential distribution with rate parameter $\lambda = 1$. An agent chooses to enter Both_c with probability γ , only Greedy_c with probability $\alpha(1-\gamma)$, and only Patient_c with probability $(1-\alpha)(1-\gamma)$, as in the theory above. We vary $\alpha \in \{0, 0.1, \dots, 1\}$ and $\gamma \in \{0, 0.1, \dots, 1\}$, and plot the global loss realized for each of these parameter settings.




Figure 1: Average loss (y-axis) as the overlap between markets γ increases (x-axis), with entrance rate parameter $m = 1000$ and $d = 20$, for different values of α . The loss of individual Patient and Greedy markets are shown as thick black and thick dashed bars, respectively.

Immediately obvious is that running a single Patient market results in dramatically less loss than competing markets, for all different values of α and γ . Furthermore, we see that the loss of a single Greedy market is also dramatically higher than the loss of a single Patient market, as predicted by Akbarpour et al. [2]. Indeed, from Equation (3) we would expect the single Patient market to have essentially zero loss, so these experiments show that adding in a rival Greedy market increases loss. In fact, as the left side of Figure 1 and the right side of Figure 2 show, it is the case that if





Figure 2: Average loss (y-axis) as the probability α of entering Patient_c or Greedy_c changes (x-axis), with entrance rate parameter $m = 1000$ and $d = 20$, for different values of the market overlap γ . The loss of individual Patient and Greedy markets are shown as thick black and thick dashed bars, respectively.

the markets do not overlap substantially (i.e., γ is low) and agents are more likely to enter the greedy side of the market (i.e., α is near 1), then the loss of the competing market is worse than running a single Greedy market! This is due in part to the decrease in market thickness on the Patient_c side of the market—a behavior we will see exacerbated below and in the kidney exchange experiments of Section 5.2.

Figure 3 decreases the rate parameter of the entrance Poisson process to $m = 100$, while holding the probability of an acceptable transaction between two agents at that of Figures 1 and 2 (so $d = 2$, leading to $2/100 = 2\%$). With fewer participants in the market overall, all the qualitative results of the $m = 1000$ markets above are amplified. The individual Greedy market’s loss is now 5.9% worse than the individual Patient market (as opposed to 3.3% in the $m = 1000$ case); both individual markets’ losses are substantially higher as well. Similarly, the parameter settings for which the competing market scenario has higher loss than either individual market are much broader than the $m = 1000$ case, which is a product of market thinness.

5.2 Dynamic kidney exchange

In this section, we expand our matching model to one of *barter exchange*, where agents endowed with items participate in directed, cyclic swaps of size greater than or equal to two. One recently-fielded barter application is *kidney exchange*, where patients with kidney failure swap their willing but incompatible organ donors with other patients. We focus on that application here. Dynamic barter exchange generalizes the matching model presented above, so we would not expect the earlier theoretical results to adhere exactly. Interestingly, as we show in Sections 5.2.1 and 5.2.2, the qualitative ranking of matching policy loss (with a patient market outperforming a greedy market, both of which outperform two rival markets) remains.

This section’s experiments draw from two kidney exchange




Figure 3: Average loss as the probability α of entering Patient_c or Greedy_c (top) or the overlap between the two markets γ (bottom) changes, with entrance rate parameter $m = 100$ and $d = 2$. The loss of individual Patient and Greedy markets are shown as thick black and thick dashed bars, respectively.

compatibility graph distributions. One distribution, which we call SAIDMAN, was designed to mimic the characteristics of a nationwide exchange in the United States in steady state [31]. Yet, kidney exchange is still a nascent concept in the US, so fielded exchange pools do not adhere to this model. With this in mind, we also include results performed on a dynamic pool generator that mimics the United Network for Organ Sharing (UNOS) nationwide exchange, drawing data from the first 193 match runs of that exchange. We label the distribution derived from this as UNOS.

Formally, we represent a kidney exchange pool with n patient-donor pairs as a directed *compatibility graph* $G = (V, E)$, such that a directed edge exists from patient-donor pair $v_i \in V$ to patient-donor pair $v_j \in V$ if the donor at v_i can give a kidney to the patient at v_j . Edges exist or do not exist due to the medical characteristics (blood type, tissue type, relation, and many others) of the patient and potential

donor, as well as a variety of logistical constraints. Our generators take care of these details; for more information on how edge existence checking is done in the SAIDMAN and UNOS distributions, see Saidman et al. [31] or Dickerson and Sandholm [14], respectively. Importantly, under either distribution, there is no longer a constant probability “ d/m ” of an acceptable transaction existing between any two agents.

Vertices arrive via a Poisson process with rate parameter $m = 100$ and depart according to an exponential clock with rate parameter $\lambda = 1$ as before, and choose to enter either exchange or both with the previously-defined probabilities γ and α . However, a “match” now only occurs when a vertex forms either a 2-cycle or 3-cycle with one or two other vertices, respectively.¹ Section 5.2.1 performs experiments on 2-cycles alone, which adheres more closely to the theoretical setting above (2-cycles can be viewed as a single undirected edge between two vertices), while Section 5.2.2 expands this to both 2- and 3-cycles.

Code to replicate the experiments in this section is available at github.com/JohnDickerson/KidneyExchange. This codebase includes our experimental framework, dynamic exchange simulator, and graph generators but, due to privacy concerns, does not include the real match runs from the UNOS kidney exchange.

5.2.1 Kidney exchange with 2-cycles only

We now present results for dynamic matching under competing Patient_c and Greedy_c kidney exchanges, both of which use only 2-cycles. Figure 4 and Figure 5 show losses incurred in our parameterized market when run on SAIDMAN-generated and UNOS-generated pools, respectively.

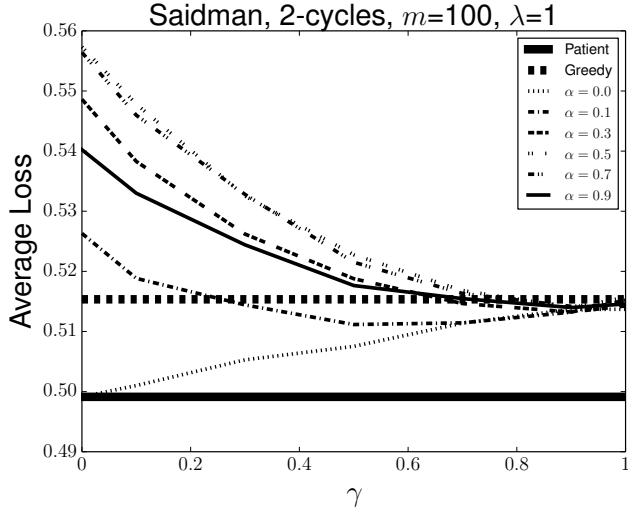


Figure 4: Average loss under various values of γ and α for the Saidman distribution with 2-cycles only.

¹In fielded kidney exchange, cycles longer than some short cap L (e.g., $L = 3$ at the UNOS exchange and many others) are typically infeasible to perform due to logistical constraints, and thus are not allowed. We adhere to that constraint here. Fielded exchanges also realize gains from *chains*, where a donor without a paired patient enters the pool and triggers a directed path of transplants through the compatibility pool. We do not include chains in this work.




Figure 5: Average loss under various values of γ and α for the UNOS distribution with 2-cycles only.

While the barter exchange environment under either the SAIDMAN or UNOS distributions clearly breaks the structural properties of the stationary distribution of the underlying Markov process used in our theoretical results, the qualitative results of these experiments align with the traditional dynamic matching results of Section 5.1. The overall loss realized by UNOS is substantially higher than that realized by SAIDMAN because, in general, UNOS-generated graphs are more sparse than those from the SAIDMAN family. Similarly, in either distribution there exist “highly-sensitized” vertex types that are extremely unlikely to find a match with another randomly selected vertex, and thus almost certainly create loss. Indeed, both Figure 4 and 5 exhibit higher loss than the similarly-parameterized Figure 3 of Section 5.1.

5.2.2 Kidney exchange with both 2- and 3-cycles

We now extend our experiments to allow for “matches” that include both 2- and 3-cycles. Unlike Section 5.1 or 5.2.1, where a matched edge was chosen uniformly at random from the set of all acceptable transactions between a distinguished vertex and its neighbors, in these results we may wish to distinguish a potential match from others (for example, by choosing a 3-cycle before a 2-cycle, as the former results in a larger myopic decrease in the market’s loss). Thus, given a set of possible 2- and 3-cycle matches, we consider two matching policies: UNIFORM selects a cycle at random from the set of possible matches, regardless of cycle cardinality, while UNIFORM3 selects a 3-cycle randomly (if one exists), otherwise a random 2-cycle.

Figures 6 and 7 show results for the SAIDMAN and UNOS distributions, respectively, under the UNIFORM match selection policy. Intuitively, one might expect the loss of a matching policy run in the 2- and 3-cycle case to be less than the same policy run in the 2-cycle case alone, as the set of possible matches weakly increases in the former case. We see this behavior when comparing the SAIDMAN results of Figure 6 to the earlier 2-cycle-only SAIDMAN results of Figure 4, witnessing a drop in global loss of around 4% for any parameter setting. We see a similar decrease in loss when comparing the new UNOS results of Figure 7 to those

in the 2-cycle case shown in Figure 5.

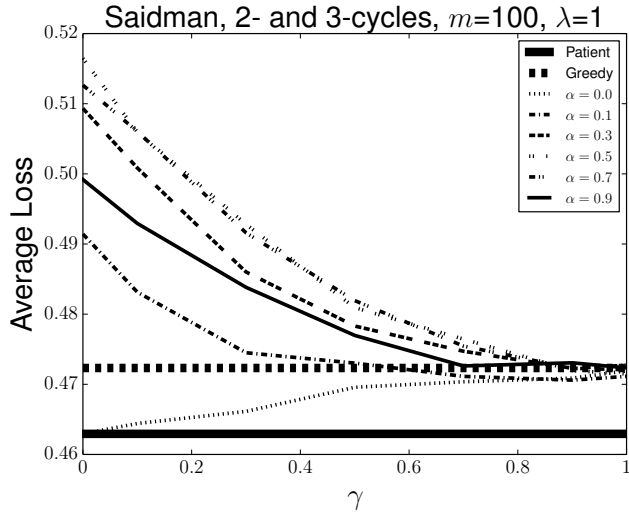


Figure 6: Average loss under various values of γ and α for the Saidman distribution with both 2- and 3-cycles, under the Uniform matching policy.

rival markets hurting global loss relative to either individual market for nearly all settings of γ and α .

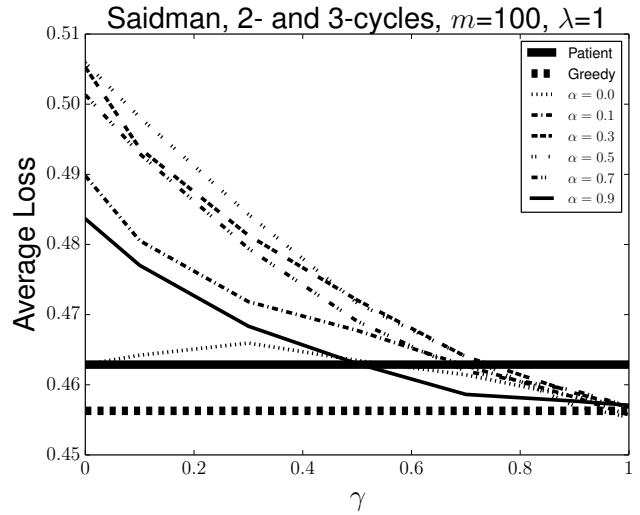


Figure 8: Average loss under various values of γ and α for the Saidman distribution with both 2- and 3-cycles, under the Uniform3 matching policy.




Figure 7: Average loss under various values of γ and α for the UNOS distribution with both 2- and 3-cycles, under the Uniform matching policy.

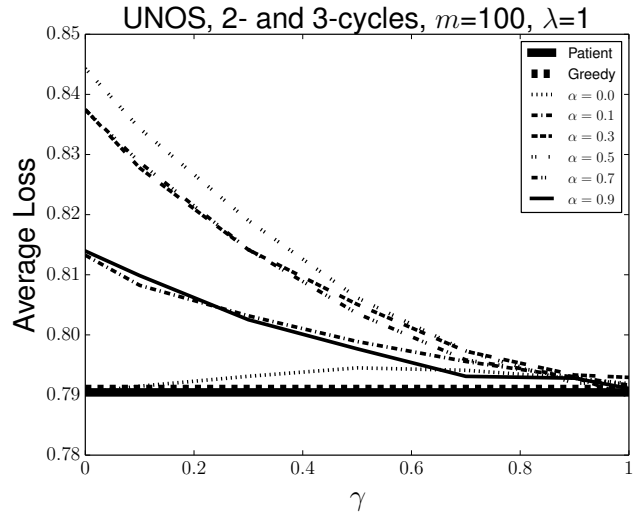


Figure 9: Average loss under various values of γ and α for the UNOS distribution with both 2- and 3-cycles, under the Uniform3 matching policy.

We now consider the UNIFORM3 matching policy, which would likely be closer to how a fielded exchange would act. Figures 8 and 9 show results for the SAIDMAN and UNOS families of compatibility graphs, respectively. The loss of the individual Patient market does not change in either distribution, which is likely a byproduct of the thicker markets induced by its match cadence. Curiously, the loss of the individual Greedy market drops dramatically—to around the Patient loss in the UNOS case, and *below* Patient in the SAIDMAN case. This large drop in Greedy loss is likely due in part to Greedy now “poaching” larger 3-cycles from the leftover market from which the Patient policy draws. The other qualitative results of earlier sections are repeated, with

6. CONCLUSION & FUTURE RESEARCH

Our main goal is to study the impact of competition between exchanges in a dynamic matching setting. In this paper, we extended the recent dynamic matching model of Akbarpour et al. [2] to two rival matching markets with overlapping pools. Specifically, we formalized a two-market model where agents enter one market or both markets; they can then potentially be matched to other agents who have joined the same market or both markets. The markets, called Greedy and Patient, adhere to different matching policies. We provided an analytic lower bound on the *loss* of the

two-market model and showed that it is higher than running a single Patient market. We also provided a quantitative method for determining the loss of the two-market model. We supported these theoretical results with extensive simulation. We also looked at competing *kidney exchanges*, and provided (to our knowledge) the first experimental quantification of the loss in global welfare in a setting with two clearinghouses using realistic kidney exchange data drawn from a generator due to Saidman et al. [31] and another based on the United Network for Organ Sharing (UNOS) program.

We see competing dynamic matching markets as fertile ground for future research, with a trove of both theoretical and practical questions to answer. First, the model of Akbarpour et al. [2] discounts the utility of a match by the time the matching agent has already waited in the pool; this is well motivated in a variety of settings, including kidney exchange. Our results in this paper assume a discount factor of zero, so it would be valuable to consider the impact on discounted loss for non-zero cases. Second, in our model the choice of market to enter is exogenously determined for each agent. In reality, agents with different levels of knowledge, wealth, etc. may make strategic decisions on which markets to enter. Thus, one could approach this dynamic matching problem from a game-theoretic point of view. Similarly, taking network effects (where more popular exchanges have an easier time attracting agents, lower operating costs, higher probabilities of two agents forming an acceptable transaction, and other advantages) into account would make these models more applicable to many real-world settings. Finally, we only looked at two overlapping markets; generalizing this to any number of overlapping markets would also be of interest.

In terms of barter exchange and, specifically, kidney exchange, the question of how clearinghouses interact is a timely one. In the United States and, eventually, elsewhere, multi-center and single-center exchange clearinghouses are already competing, each drawing from some (often overlapping) subset of the full set of patient-donor pairs available. Indeed, the dynamic barter exchange problem *in a single market* is still not fully understood (barring very promising recent work due to Anderson et al. [3]). We saw in Section 5.2.2 that including 3-cycles in the matching process results in lower loss, even when two markets overlap, compared to including only 2-cycles (a result that has been shown repeatedly in the static [30] and dynamic [3] single clearinghouse setting), so extending the theoretical underpinnings of our framework to a more general setting would be of great value. Finally, it is curious that the UNIFORM3 policy had such a large effect on the loss of the individual Patient and Greedy exchanges compared to the UNIFORM policy; further exploration of different matching policies (including those that use a strong prior to consider possible future states of the pool when matching now) would be helpful in making policy recommendations to fielded exchanges.

Acknowledgments. This material was funded by NSF grant IIS-1320620, NSF CAREER Award IIS-1414452, and by a National Defense Science & Engineering Graduate (NDSEG) Fellowship. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by NSF grant OCI-1053575; specifically, it used the Blacklight system at the Pittsburgh Supercomputing Center (PSC).

7. REFERENCES

- [1] A. Abdulkadiroğlu and S. Loertscher. Dynamic house allocations, 2007. Working paper.
- [2] M. Akbarpour, S. Li, and S. O. Gharan. Dynamic matching market design. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, page 355, 2014.
- [3] R. Anderson, I. Ashlagi, D. Gamarnik, and Y. Kanoria. A dynamic model of barter exchange. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2015.
- [4] E. Anshelevich, M. Chhabra, S. Das, and M. Gerrior. On the social welfare of mechanisms for repeated batch matching. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 60–66, 2013.
- [5] N. Arnosti, R. Johari, and Y. Kanoria. Managing congestion in decentralized matching markets. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, page 451, 2014.
- [6] I. Ashlagi, P. Jaillet, and V. H. Manshadi. Kidney exchange in dynamic sparse heterogeneous pools. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 25–26, 2013.
- [7] I. Ashlagi, M. Tennenholtz, and A. Zohar. Competing schedulers. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2010.
- [8] P. Awasthi and T. Sandholm. Online stochastic optimization in the large: Application to kidney exchange. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 405–411, 2009.
- [9] D. Bertsimas, V. F. Farias, and N. Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1):73–87, 2013.
- [10] F. Bloch and D. Cantala. Markovian assignment rules. *Social Choice and Welfare*, 40(1):1–25, 2013.
- [11] E. B. Budish, P. Cramton, and J. J. Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response, 2015. Working paper.
- [12] R. Burguet and J. Sákóvics. Imperfect competition in auction designs. *International Economic Review*, 40(1):231–247, 1999.
- [13] J. P. Dickerson, A. D. Procaccia, and T. Sandholm. Dynamic matching via weighted myopia with application to kidney exchange. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 1340–1346, 2012.
- [14] J. P. Dickerson and T. Sandholm. FutureMatch: Combining human value judgments and machine learning to match in dynamic environments. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2015.
- [15] T. Eisenmann, G. Parker, and M. W. Van Alstyne. Strategies for two-sided markets. *Harvard Business Review*, 84(10):92, 2006.
- [16] P. Erdős and A. Rényi. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5:17–61, 1960.
- [17] H. Hosseini, K. Larson, and R. Cohen. Matching with dynamic ordinal preferences. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

[18] S. V. Kadam and M. H. Kotowski. Multi-period matching, 2014. Working paper.

[19] E. H. Kaplan. *Managing the demand for public housing*. PhD thesis, Massachusetts Institute of Technology, 1984.

[20] E. H. Kaplan. Tenant assignment models. *Operations Research*, 34(6):832–843, 1986.

[21] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*, pages 352–358, 1990.

[22] J. Kennes, D. Monte, and N. Tumennasan. The day care assignment: A dynamic matching problem. *American Economic Journal: Microeconomics*, 6(4):362–406, 2014.

[23] M. Kurino. *Essays on dynamic matching markets*. PhD thesis, University of Pittsburgh, 2009.

[24] J. D. Leshno. Dynamic matching in overloaded waiting lists, 2015. Working paper.

[25] A. Mehta. Online matching and ad allocation. *Theoretical Computer Science*, 8(4):265–368, 2012.

[26] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized online matching. *Journal of the ACM*, 54(5), 2007.

[27] D. Monderer and M. Tennenholtz. K-price auctions: Revenue inequalities, utility equivalence, and competition in auction design. *Economic Theory*, 24(2):255–270, 2004.

[28] M. Ostrovsky. Stability in supply chain networks. *American Economic Review*, pages 897–923, 2008.

[29] A. Roth, T. Sönmez, and U. Ünver. Kidney exchange. *Quarterly Journal of Economics*, 119(2):457–488, 2004.

[30] A. Roth, T. Sönmez, and U. Ünver. Efficient kidney exchange: Coincidence of wants in a market with compatibility-based preferences. *American Economic Review*, 97:828–851, 2007.

[31] S. L. Saidman, A. Roth, T. Sönmez, U. Ünver, and F. Delmonico. Increasing the opportunity of live kidney donation by matching for two and three way exchanges. *Transplantation*, 81(5):773–782, 2006.

[32] Y. Shoham and K. Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2008.

[33] X. Su and S. A. Zenios. Patient choice in kidney allocation: A sequential stochastic assignment model. *Operations Research*, 53:443–455, 2005.

[34] U. Ünver. Dynamic kidney exchange. *Review of Economic Studies*, 77(1):372–414, 2010.

[35] S. A. Zenios. Optimal control of a paired-kidney exchange program. *Management Science*, 48(3):328–342, 2002.

APPENDIX

A. ADDITIONAL EXPERIMENTS

In this section, we provide additional results supporting the dynamic kidney exchange experiments of Section 5.2. Figure 10 corresponds to the 2-cycle-only experiments of Figures 4 and 5 in the body of the paper; instead of varying the market overlap parameter γ on the x-axis, they vary the probability α of entering either the Greedy_c or Patient_c

market, while holding γ constant for a variety of values. Similarly, Figure 11 corresponds to the 2- and 3-cycle UNIFORM matching policy experiments of Figures 6 and 7. Finally, Figure 12 corresponds to the UNIFORM3 matching policy results shown in Figures 8 and 9.




Figure 10: 2-cycles-only experiments, paired with Figure 4 (left) and Figure 5 (right).




Figure 11: 2- and 3-cycle Uniform experiments, paired with Figure 6 (left) and Figure 7 (right).




Figure 12: 2- and 3-cycle Uniform3 experiments, paired with Figure 8 (left) and Figure 9 (right).