

Equilibrium Behavior in Competing Dynamic Matching Markets

Zhuoshu Li¹, Neal Gupta², Sanmay Das¹ and John P. Dickerson²

¹ Washington University in St. Louis

² University of Maryland

zhuoshuli@wustl.edu, ngupta@cs.umd.edu, sanmay@wustl.edu, john@cs.umd.edu

Abstract

Rival markets like rideshare services, universities, and organ exchanges compete to attract participants, seeking to maximize their own utility at potential cost to overall social welfare. Similarly, individual participants in such multi-market systems also seek to maximize their individual utility. If entry is costly, they should strategically enter only a subset of the available markets. All of this decision making—markets competitively adapting their matching strategies and participants arriving, choosing which market(s) to enter, and departing from the system—occurs dynamically over time. This paper provides the first analysis of equilibrium behavior in dynamic competing matching market systems—first from the points of view of individual participants when market policies are fixed, and then from the points of view of markets when agents are stochastic. When compared to *single* markets running social-welfare-maximizing matching policies, losses in overall social welfare in competitive systems manifest due to both market fragmentation and the use of non-optimal matching policies. We quantify such losses and provide policy recommendations to help alleviate them in fielded systems.

1 Introduction

In centralized dynamic matching problems, agents arrive to a market and are either paired with other agents or left unpaired until they depart. A successfully matched agent receives some individual one-time utility, while the central market accumulates utility over many matches. Examples of such markets are myriad: online dating [Bragdon *et al.*, 2010], rideshare [Hall *et al.*, 2017], organ allocation [Bertsimas *et al.*, 2013], online labor [Arnosti *et al.*, 2014], housing allocation [Bloch and Cantala, 2013], job search [Das and Tsitsiklis, 2010], refugee placement [Jones and Teytelboym, 2017], and general barter markets like kidney exchange [Dickerson and Sandholm, 2015]. Many matching applications involve two or more markets that *compete* in the same space for a common set of agents who choose to enter one or more of the competing markets based on their expected utility from

getting matched. Match.com, OkCupid, and Tinder compete on matching policies (e.g., match speed, profile complexity) to segment the market, while users choose to enter these markets according to perceived (monetary or temporal) cost and likelihood of finding a partner. Similarly, Uber, Lyft, and Didi Chuxing compete internationally and locally via driver/rider incentives and geographic preferences. Finally, organ exchanges compete nationally—and, increasingly, internationally—to attract participants.

In this paper, we explore two complementary research questions in a stylized model of multiple competing dynamic matching markets. Building on an important dynamic matching model introduced by Akbarpour, Gharan, and Li 2017, Das *et al.* 2015 demonstrated that market fragmentation (where significant fractions of the population go to different markets) can have significant social costs by reducing the thickness of matching markets. This work considered models where market policies were fixed (one market matched greedily and the other was maximally patient), and agents arrived to one or both markets stochastically. In reality, we expect at least some proportion of sophisticated agents to be strategic about which market they choose to enter, and we also expect markets to compete with each other for volume, quality of agents, and so on.

This paper begins to address questions about the equilibrium behavior of multi-market systems under these two types of strategic agency. First, we analyze models where individual market participants have agency. These participants can be of different types (short-lived or long-lived) and may choose entrance into the market system such that their individual utility is maximized. Different types of agents may have different preferences, and we analyze equilibrium behavior in both continuous (Section 3) and discrete (Section 4) time settings. We show that even with just two types of agents, strategic market choice can induce market fragmentation—while there are some pooling equilibria where all strategic agents choose the same market (which is socially preferable), as the proportion of agents who are assigned to a particular market increases (these agents may be constrained by geography or cost, for example), separating equilibria become significantly more likely, with short-lived agents choosing the patient market and long-lived agents the greedy market. This is because the patient market is typically thicker, giving a higher probability of matching during

an agent’s sojourn, and short-lived agents suffer less penalty because the market attempts to match them sooner relative to arrival. Unfortunately, the fragmentation comes at significant social cost in reduced thickness.

Second, in Section 5, we prescribe agency to the markets themselves, allowing them to choose overall matching policies (defined by the frequency at which they decide to match) strategically to maximize their overall utility. In this case, the agents are stochastic in their choice to join one or the other market, or to enter both markets. We quantify via best response dynamics the social welfare loss of this competitive marketplace under a variety of initial conditions, and compare that loss to the lower bound provided by a single market running an optimal matching policy.

Overall, our results demonstrate the serious concern of a “race to the bottom” when multiple matching markets compete. This is due to both fragmentation and the choice of socially suboptimal matching policies by individual markets. When agents choose markets strategically, differences in their types and utilities can lead to preferences for one or the other markets and induce separating equilibria and fragmented markets. Even when agents do not have market choice, if markets can choose their matching policies, individual markets may be incentivized to match as early as possible an inefficient fraction of the time in the race to match more agents. The intersection of differential impact on different types and competing matching platforms raises important ethical issues in allocation and regulation. Such discussions can be informed by our models. Further, our models can also provide the foundation for future models that consider situations where *both* agents and markets can be strategic. Section 6 concludes with some recommendations for policymakers derived from our results.

2 Preliminaries

Akbarpour *et al.* 2017 show that platforms may maximize the number of matches achieved by being *patient* instead of trying to match new pairs immediately. The intuition is that waiting for the market to become thicker can be helpful. However, any individual would prefer to take the earliest match that becomes available to them. Das *et al.* 2015 consider what happens when rival patient and greedy markets compete with each other under stochastic assumptions about individual agent behavior, and show that such market fragmentation can have significant social costs.

We start with a brief overview of the Akbarpour *et al.* model as applied to our setting. We have two rival dynamic matching markets. Each market runs in time interval $[0, T]$, with agents arriving according to a Poisson process with rate parameter k . The market determines whether potential bilateral transactions between agents are acceptable. The probability of an acceptable transaction existing between any pair of distinct agents is defined as p , and is independent of any other pair of agents in the market. Each agent a remains in the market for a sojourn $s(a)$; the agent becomes *critical* immediately before her sojourn ends, and this criticality time is known to the market. An agent leaves either upon being matched successfully by the market or upon becoming criti-

cal and remaining un-matched, at which point she perishes. An agent a receives zero utility if she perishes, or $u(a) = 0$. If she is matched, she receives a utility 1 discounted at rate δ , or $u(a) = e^{-\delta s(a)}$.

Another critical issue is when the market tries to *clear* or actively find matches. Two extremes are:

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*.

The market can also choose a clearing rule that interpolates between the Patient and Greedy clearing rules (the so-called Patient(α) clearing rule), which allows tuning of the matching rate. Specifically, a market matching with the Patient(α) strategy draws an exponential random variable C_v with rate parameter $1/\alpha$ for each vertex v . If vertex v entering at time t becomes critical at time $t_c < t + C_v$, she matches at t_c , as in the Patient matching algorithm. Otherwise the vertex matches at time $t + C_v$. Note that when $\alpha \rightarrow 0$ we will have $C_v \xrightarrow{p} 0$, which corresponds to a Greedy matching algorithm.

2.1 MODEL I: Strategic Agents

Our first model considers two types of agents in terms of length of life, *short-lived* and *long-lived*. Short-lived agents come into the markets with a length of life T_s and long-lived agents have a length of life T_l , where $T_s < T_l$. Each agent (who is aware of her own type) decides which market to enter upon arrival. A fraction θ of agents are short-lived and the remaining $1 - \theta$ fraction are long-lived. We allow a ϕ fraction of random-choice agents (*random* agents) to choose either market with 0.5 probability. The remaining $1 - \phi$ fraction of agents are *strategic*. For these models, we restrict attention to models in which one Greedy and one Patient market compete. The action space for agents is the market choice, $B = \{Greedy, Patient\}$. We want to analyze the equilibrium strategies of strategic agents given the setting of θ and ϕ . Here, the market choice becomes a tradeoff between matching probability and utility. That is, entering a Patient market may give an agent a higher matching probability but lower utility as the agent has a higher expected sojourn time; in contrast, immediate matching from a Greedy market provides a higher utility but may lower the probability of matching since the market is not thick enough.

We investigate the behavior of strategic agents in the two-market MODEL I in both continuous time (Section 3) and discrete time (Section 4) models.

2.2 MODEL II: Strategic Markets

Our second, complementary, direction is to model the situation where agent behaviors are stochastic, but markets themselves make strategic decisions. We define each market’s

utility as the aggregate utility of the (non-strategic) agents it matches (it is reasonable to assume that the market can capture some fraction of this utility). We follow the model of Das *et al* 2015 for assigning agents to one or both of the two competing markets. A γ_1 fraction of agents are assigned to both markets; the market which successfully matches the agent first will receive utility from the match. The remaining agents are only assigned to one market: a γ_2 fraction enter the first market, while a $1 - \gamma_2$ fraction enter the second market.

The action chosen by a market is its choice of market-clearing rule, parameterized by the matching rate α described above. The market-clearing rule choice involves a tradeoff: if Market 1 chooses a fast matching rate, it will match more agents assigned to both markets, but will match fewer agents which are only assigned to Market 1. The relative market sizes, parameterized by γ_1 and γ_2 are factors in the optimal choice. We investigate equilibrium behavior via simulation of two markets in MODEL II in Section 5.

3 Strategic Agents in Continuous Time

We consider two markets operating simultaneously, one Greedy and one Patient. For simplicity, we assume that lengths of life T_s and T_l for short-lived and long-lived agents are constants that are fixed across the same type of agents.¹

The markets run in the continuous-time interval $[0, T]$. Agents arrive according to a Poisson process with rate parameter $k \geq 1$ ($k = 100$ in our simulations). The type of each arriving agent is stochastic; with probability θ , the arriving agent is a short-lived type; and with probability $1 - \theta$, she is a long-lived type. Parameter ϕ controls whether an agent is random or strategic, that is, with probability ϕ , she is a random agent and w.p. $1 - \phi$ she is a strategic agent. Upon arrival, the agent needs to decide which market to enter. Random agents choose a market uniformly at random and strategic agents choose a market based on comparing the expected utilities of entering each market.

We first consider agents entering the Greedy market. As the Greedy market matches agents immediately upon entry, the probability of an agent having acceptable transactions immediately after entering at any time t is $m_{g,e}(t) = (1 - (1 - p)^{Z_{g,t}})$, where $Z_{g,t}$ represents the size of the pool under the Greedy matching policy at time t . To be noticed, t here is an infinitesimal time. Since entry occurs stochastically in continuous time, only one agent enters exactly at time t . Therefore, as long as there exist any acceptable transactions, the entering agent will be matched immediately. Once the moment of entry has passed, an agent can only be matched at the point in time when some other agent enters the market. The probability of an agent who was not matched at entry having an acceptable transaction at the time of entry of some other agent is $m_{g,s}(t) = \frac{(1 - (1 - p)^{Z_{g,t}})}{Z_{g,t} - 1}$. Denote the probability of an agent entering the Greedy market at any point in time t as $P[\text{Entry}_g^t]$. Thus, the expected utility of an agent

for choosing the Greedy market $U_{type,g}(t)$ at time t given she knows her type is

$$U_{type,g}(t) = m_{g,e}(t) + \int_0^{T_{type}} P[\text{Entry}_g^t] m_{g,s}(t + s(a)) e^{-\delta s(a)} ds(a), \quad (1)$$

where $type \in \{\text{short}, \text{long}\}$.

Now consider agents entering the Patient market. The Patient market attempts to match agents at the instant they become critical. The probability of an agent having acceptable transactions during their stay (before perishing) at any time t is $m_{p,s}(t) = \Lambda_{p,t} \frac{(1 - (1 - p)^{Z_{p,t}})}{Z_{p,t} - 1}$, where $\Lambda_{p,t}$ is the rate of perishing in the Patient market and $Z_{p,t}$ is the size of the pool under the Patient matching policy at time t . The probability of an agent having acceptable transactions at the instant she becomes critical is $m_{p,c}(t) = (1 - (1 - p)^{Z_{p,t}})$. Denote the probability that some agent in the Patient market becomes critical at any time t as $P[\text{Exit}_p^t]$. The expected utility of an agent for choosing patient market $U_{type,p}(t)$ at time t given her type is

$$U_{type,p}(t) = \int_0^{T_{type} - \epsilon} P[\text{Exit}_p^t] m_{p,s}(t + s(a)) e^{-\delta s(a)} ds(a) + m_{p,c}(t + T_{type}) e^{-\delta T_{type}}, \quad (2)$$

where ϵ is an infinitesimal amount of time right before an agent perishes and $type \in \{\text{short}, \text{long}\}$.

Equations (1) and (2) clarify the tradeoffs agents face. In general, while the patient market may give a higher probability of finding a match, the fact that an agent typically has to wait longer diminishes her expected utility. Since agents start with the same utility and it diminishes at the same rate, this means that short-lived agents will have a relatively higher preference for the Patient market compared with long-lived agents (who have to wait longer until the point in time when they are most likely to get matched, the time of criticality, in the patient market). Since there are positive externalities to entering a market and making it thicker, we may expect that the market-choice game may have both pooling and separating equilibria, where either both types of agents enter one market or short-lived agents enter the Patient market while long-lived agents enter the Greedy market.

Since the equations above do not admit closed-form solutions, we use empirical game-theoretic analysis to find equilibria in the game with strategic market-choice for each type. The strategy space is $B = \{\text{Greedy}, \text{Patient}\}$. For different values of θ and ϕ , we compute the utilities of strategic short-lived and long-lived agents if they choose the Greedy market or the Patient market respectively using Monte Carlo simulations holding the strategies of the other agents fixed, and ascertain whether or not pooling or separating equilibria exist in different regions of the θ, ϕ space. As conjectured, we do see an overall pattern of pooling and separating equilibria in different regions. Figure 1a shows an example of the results when the fraction of random agents is $\phi = 0.4$. These results can be broken up into three regions: The red region represents pooling equilibria where both long-type and short-type strategic agents choose the Greedy market; long-lived and random agents are the majority in this region, thus the Greedy market can be thick enough.² The yellow region represents pooling

¹We also ran experiments where T_s and T_l are sampled from two exponential distributions with different rate parameters λ_s and λ_l , truncated so that $T_s < 1$ and $T_l \geq 1$. The results were qualitatively very similar to the case where T_s and T_l are constants.

²Note that, in this region, both types of agents choosing the Pa-

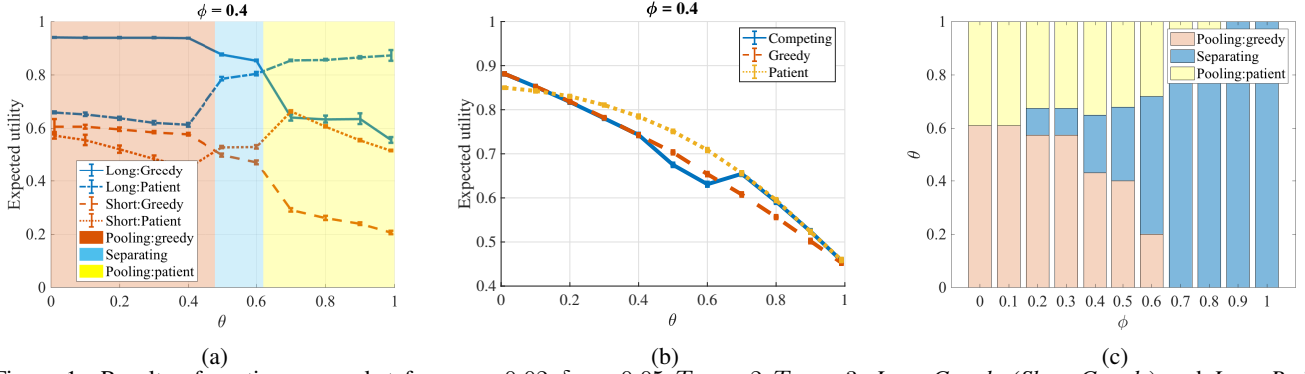


Figure 1: Results of continuous market for $p = 0.02, \delta = 0.05, T_s = 2, T_l = 3$. *Long:Greedy* (*Short:Greedy*) and *Long:Patient* (*Short:Patient*) show the expected utility of a strategic long-type (short-type) agent if she chooses Greedy and Patient respectively.

equilibria where both types choose the Patient market; in this range of settings, we have more short-lived agents and the Greedy market is not thick enough as the short-lived agents perish too soon.³ In the blue region, we find a separating equilibrium exists: strategic short-lived types choose the Patient market ($U_{s,p} > U_{s,g}$) and long-lived types choose the Greedy market ($U_{l,g} > U_{l,p}$).

Figure 1b shows overall social welfare in the Competing system with a single Greedy market and a single Patient market under different settings of θ when the fraction of random agents $\phi = 0.4$. We can see the market fragmentation caused by competition, separating equilibria ($\theta \in [0.4, 0.6]$), lowers the social welfare when compared to a single market. This pattern holds across the whole range of ϕ .

Finally Figure 1c shows the range of separating and pooling equilibria as a function of ϕ , the proportion of random agents. As the proportion of random agents increases, the portion of the θ domain covered by separating equilibria increases, since the thicknesses of the two markets are determined almost entirely exogenously, and the main consideration is an optimization of utility rather than equilibrium considerations of what other strategic agents are doing.

4 Strategic Agents in Discrete Time

While the model of Section 3 uses essentially the same models of utility as prior work, we are restricted by the lack of analytical tractability. We now consider a discrete time version of MODEL I that captures the same basic intuitions and can be used more directly in modeling strategic market choice. We believe this model is more amenable for further work on these questions. Now, agents enter the market in “batches”, that is, $k \geq 1$ agents enter the Competing market at each time step t . Short-lived agents live for T_s time steps and long-lived agents live for T_l time steps ($T_l > T_s$), where T_s and T_l are fixed constants for each type. At each time step t , each market operates as follows: o_1 : agents enter $\rightarrow o_2$: market clears $\rightarrow o_3$: agents perish. Random agents choose a market to enter uniformly at random and strategic agents choose a

market is also an equilibrium, albeit one with overall lower social welfare.

³Similarly, in this region, both types of agents choosing the Greedy market is also an equilibrium, lower in social welfare than the Patient pooling equilibrium.

market based on comparing the expected utilities of entering each market.

We first analyze the utility of agents choosing the Greedy market. As the Greedy policy will match agents immediately after they enter the market, the probability of an agent having acceptable transactions immediately after entering is $m_{g,e}(t) = (1 - (1 - p)^{Z_{g,t}})$, where $Z_{g,t}$ represents the size of the pool under the Greedy matching policy at time t . The market will run a maximum matching algorithm at each time step during o_2 as $k \geq 1$ agents enter the market at the same time. This means that agents may be unmatched even if they have potential acceptable transactions. We define the probability of being matched in the maximum matching given the agent has acceptable transactions as $\chi_g(t)$ in the Greedy market at time t . The probability of an agent having acceptable transactions when they stay in the market (that is, not at their time-step of entry) at time t is $m_{g,s}(t) = (1 - (1 - p)^{k_{g,t}})$, where $k_{g,t}$ is the number of agents entering to the Greedy pool at time t . Thus, the expected utility of an agent for choosing Greedy market $U_{type,g}$ at time t given she knows her type is

$$\begin{aligned}
 U_{type,g}(t) &= m_{g,e}(t)\chi_g(t) + (1 - m_{g,e}(t))\chi_g(t) \\
 &\quad \left[e^{-\delta} m_{g,s}(t+1)\chi_g(t+1) + \right. \\
 &\quad \left. \sum_{s(a)=2}^{T_{type}-1} e^{-\delta s(a)} m_{g,s}(t+s(a))\chi_g(t+s(a)) \right. \\
 &\quad \left. \prod_{j=1}^{s(a)-1} (1 - m_{g,s}(t+j))\chi_g(t+j) \right], \tag{3}
 \end{aligned}$$

where $type \in \{short, long\}$ and $T_{type} \geq 3$. We have two special cases, where $U_{type,g}(t) = m_{g,e}(t)\chi_g(t)$ when $T_{type} = 1$; and when $T_{type} = 2$, $U_{type,g}(t) = m_{g,e}(t)\chi_g(t) + (1 - m_{g,e}(t))m_{g,s}(t+1)\chi_g(t+1)e^{-\delta}$.

We next consider the expected utility of agents choosing the Patient market. The Patient market will match agents only at the instant they become critical. The probability of an agent having acceptable transactions when they stay in the Patient market at each time step t is $m_{p,s}(t) = (1 - (1 - p)^{\Lambda_{p,t}})$, where $\Lambda_{p,t}$ is the number of agents becoming critical in the Patient market at time t . As there may be more than one agent becoming critical at each time t , the Patient market will also run a maximum matching at o_2 . We define the probability of being matched in the maximum matching given the agent has acceptable transactions as $\chi_p(t)$. The probability of an

agent having acceptable transactions when she is critical is $m_{p,c}(t) = (1 - (1 - p)^{Z_{p,t}})$, where $Z_{p,t}$ is the size of the pool under the Patient matching policy at time t . Thus, the expected utility of an agent for choosing Patient market $U_{type,p}$ at time t given she knows her type is

$$U_{type,p}(t) = m_{p,s}(t)\chi_p(t) + \left[\sum_{s(a)=1}^{T_{type}-2} e^{-\delta s(a)} m_{p,s}(t+s(a))\chi_p(t+s(a)) \prod_{j=0}^{s(a)-1} (1 - m_{p,s}(t+j)\chi_p(t+j)) \right] + \prod_{j=0}^{T_{type}-2} (1 - m_{p,s}(t+j)\chi_p(t+j)) m_{p,c}(t+T_{type}-1)\chi_p(t+T_{type}-1)e^{-\delta(T_{type}-1)}, \quad (4)$$

where $type \in \{short, long\}$ and $T_{type} \geq 3$. We also have two special cases, where $U_{type,p}(t) = m_{p,c}(t)x_p(t)$ when $T_{type} = 1$; and when $T_{type} = 2$, $U_{type,p}(t) = m_{p,s}(t)\chi_p(t) + (1 - m_{p,s}(t)\chi_p(t))m_{p,c}(t+1)\chi_p(t+1)e^{-\delta}$.

At any time $t \in [0, T]$, $Z_{g,t}$, $Z_{p,t}$ represent the sizes of the pools under the Greedy and Patient matching policies, respectively. The Markov chain on $Z_{\cdot,t}$ has a unique stationary distribution under either of those policies. Let $\pi_g, \pi_p : \mathbb{N} \rightarrow \mathbb{R}^+$ be the unique stationary distributions 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 sizes of the pool under the stationary distribution under Greedy and Patient. After mixing, we represent the expected sizes of the pools at any time as ξ_g, ξ_p respectively. Similarly, $k_{g,t}, \Lambda_{p,t}, \chi_g(t)$ and $\chi_p(t)$ also can be represented by expected values k_g, Λ_p, χ_g and χ_p . We use Monte Carlo simulations to estimate $\xi_g, \xi_p, k_g, \Lambda_p, \chi_g$ and χ_p respectively. This then allows us to numerically compute the expected utilities in Equation (3) and (4) and derive the equilibria for different parameter settings.

Due to space limitations, we defer the results of the discrete model to a longer version of this paper. The results are qualitatively very similar to those from the continuous-time model, but the additional analytical tractability of the model presented here makes it promising for future development of models of competing markets.

5 Strategic Markets

In the previous two sections, we assumed all markets operated with *fixed* matching policies, and strategic agents entered that system in a way that maximized their individual expected utility. Here, under MODEL II, we prescribe agency onto the markets themselves, allowing them to strategically *adjust* their matching policies under best response dynamics to maximize their expected aggregate utility. We investigate equilibrium behavior in this model, and measure overall social welfare loss relative to a single-market baseline.

5.1 Experimental Setup

We are interested in modeling the behavior of a two-market system where the markets respond to each other under best response dynamics. Formally, at any time period, one market observes the matching rate of its competitor and then chooses, *for the next time period*, its own matching rate that will yield

maximum payoff for perpetuity,⁴ even though the market will change its best response within a short span of time T_R . Best response dynamics have been shown to mimic many settings where agents operate reactively or with bounded expertise [Weibull, 1995], and can be used in some cases to find equilibria [Monderer and Shapley, 1996].

We simulated the long-term utilities for two markets M_1 and M_2 with Patient(α^1) and Patient(α^2) matching policies, respectively, for $(\alpha^1, \alpha^2) \in \mathbb{R}^{\geq 0} \times \mathbb{R}^{\geq 0}$, for $T = 250$ periods, and 100 trials. We estimated the best response functions BR_1 and BR_2 for markets M_1 and M_2 , respectively by simulating two markets with overlaps $\gamma_1 \in \{0.1, 0.2, \dots, 0.9\}$ over a grid of patience parameters α^1, α^2 . As a reminder, higher values of α correspond to more patience—i.e., matching less frequently—and higher values of γ_1 indicate higher overlap—i.e., more agents entering both markets.

We assumed the markets have bounded rationality in their computations of best response functions. From the set S of all Monte Carlo simulations, we took $X = 2500$ bootstrap samples of size $n = 50$, $\{S'_i\}_{i=1}^X$ where $S'_i \subset S$. Each bootstrap sample represents simulations that a boundedly rational market would run. Thus, given a single bootstrap sample $S'_i \subset S$:

$$BR_i(\alpha^1) = \arg \max_{\alpha^2} \mathbb{E}_{s \in S'_i} [u_{M_2}(\alpha^1, \alpha^2)] \\ BR_i(\alpha^2) = \arg \max_{\alpha^1} \mathbb{E}_{s \in S'_i} [u_{M_1}(\alpha^1, \alpha^2)]$$

Under best response dynamics, the matching rate will now change over time, so we let α_t^1 and α_t^2 denote the matching rates at time t of market M_1 and M_2 , respectively. We iterated best responses until convergence or cycles occurred over initial conditions of α_0^1, α_0^2 values.

5.2 Experimental Results

In general, we observe two main phenomena for the best response dynamics. First, we observe convergence to the Patient strategy under appropriate initial conditions (α_0^1, α_0^2) for any constituent in the competing market system. Second, for markets with sufficient overlap, and sufficiently low initial values of (α^1, α^2) , we observe convergence to a (Greedy, Greedy) equilibrium or (α^1, α^2) parameters very close to (Greedy, Greedy). No other phenomena occur in more than 5% of bootstrap samples.

To simplify the description of results, we refer to convergence to (Greedy, Greedy), or cycles or equilibria involving solely $0 \leq \alpha^{\{1,2\}} \leq 1/100$, as fast matching. For the parameter range chosen for the simulations (specifically $d = 1$), $0 < \alpha^{\{1,2\}} \leq 1/100$ rarely impacts the matching choice. Furthermore, the social welfare for these outcomes only differ by at most 0.3%. We describe the notable effects of the parameter choices on best response dynamics below.

Market overlap. The impact of the market overlap γ_1 on the best response dynamics can be characterized by the effect on the range of initial matching rates (α_0^1, α_0^2) which converged to a fast matching outcome in “many” bootstrap samples—here, we use a cutoff of 25%. Figure 2 visualizes this behavior for increasing values of market overlap γ_1 .

When market overlap $\gamma_1 \leq 0.4$, less than 0.1% of bootstrap samples converge to a fast matching outcome for any chosen

⁴For a formal overview of best response dynamics, see, for example, the book by [Nisan *et al.*, 2007].

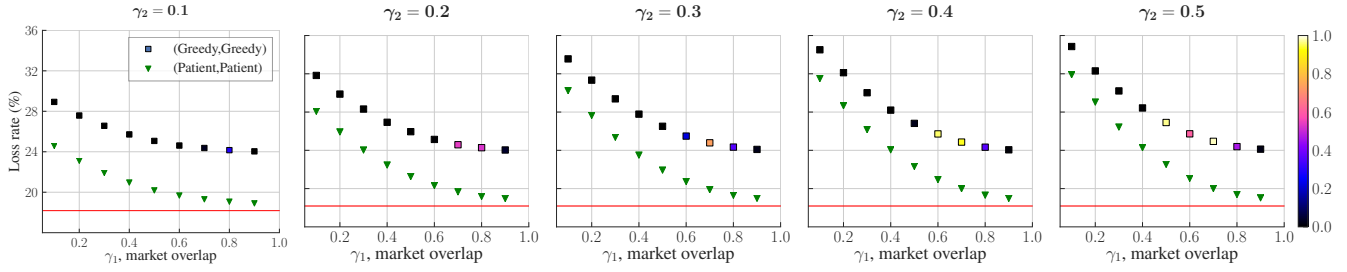


Figure 2: Simulation results for $d = 1, k = 100, p = 0.02$. γ_2 denotes the fraction of vertices, not in both markets, that enter only market M_1 . The red line denotes the loss rate with a single Patient market. The green triangles denote the loss rate of a (Patient, Patient) equilibrium (an equilibrium in essentially all bootstrap samples). The squares denote the loss rate of a (Greedy, Greedy) equilibrium and the color of squares denotes the proportion of bootstrap samples which reach a fast matching outcome (defined in §5.2) from initial conditions (Greedy, Greedy).

initial matching rates. This is expected, as with $\delta = 0$, a faster matching rate increases utility u_{M_1} of M_1 primarily when M_1 successfully matches an agent that enters both markets before M_2 can. When $\gamma_1 \in [0.4, 0.8]$, the range of initial conditions that converged to a fast matching outcome rose to a peak at or before $\gamma_1 = 0.7$, then fell off. Surprisingly, when $\gamma_1 = 0.9$, no initial conditions converged to fast matching in more than 4.4% of bootstrap samples.

Market asymmetry. We investigate the impact of γ_2 , which controls the balance of agents entering only market M_1 (which occurs as $\gamma_2 \rightarrow 1$) or market M_2 ($\gamma_2 \rightarrow 0$). When $\gamma_2 \leq 0.1$, no conditions outside of fast matching converged to fast matching in more than 25% of samples.

5.3 Welfare Loss

We now measure the impact of competition on global social welfare. As with MODEL I, we define social welfare as the discounted total number of matches; here, however, we set $\delta = 0$. As before, we compute the distribution of social welfare for a range of γ_1 and γ_2 with respect to our bootstrap samples of Monte Carlo simulations. Figure 2 shows social welfare for outcomes (Patient, Patient) and (Greedy, Greedy)—which approximates the loss rate of fast matching—as well as the proportion of bootstrap samples that converge to fast matching from initial conditions (Greedy, Greedy). Note that, just as in MODEL I, all experimental outcomes are strictly worse (i.e., result in lower social welfare) than that of a single Patient market.

As expected, when the overlap γ_1 increases, the expected loss rate decreases due to a larger network of potential matches. However, as γ_1 increases, some initial conditions also become more likely to result in a fast matching outcome. For example, under two equally-sized markets ($\gamma_2 = 0.5$), the Greedy loss rate of 26.9% is higher than the loss rate of 24.3% for $\gamma_1 = 0.4$, where the only outcome that occurs with meaningful probability is (Patient, Patient). This additional equilibrium occurs in 96.4% of bootstrap samples. The additional overall welfare loss of 4.4%-5.0% incurred by a fast matching outcome for the same initial parameters is shown in Figure 2.

We also observe the effect on social welfare of a thicker market as market asymmetry increases; at its most extreme ($\gamma_2 = 0$ and $\gamma_2 = 1$), all vertices are effectively in a single market. As such, as γ_2 moves toward its bounds, again there are stronger network effects on social welfare.

Welfare losses arise both from matching speed and market

fragmentation. As a baseline, the loss rate that occurs from a *single* Patient market—one with no competition—under the same model parameters is 18.2%. As shown in Figure 2, all other market conditions result in greater overall loss. In the succeeding section we explore policy options that could help a central planner mitigate this loss due to competition.

6 Policy Implications & Future Directions

Our results indicate that, left to themselves, matching markets that compete with each other can cause significant social welfare losses through fragmentation (§3 and §4) and sub-optimal matching policies (§5). Our results are a proof-of-concept support of the “race to the bottom” seen in many real competing matching market systems. For example, in the US, multiple kidney exchanges compete over patient-donor pairs and/or hospitals. Two of the largest US exchanges are the National Kidney Registry (NKR) and the United Network for Organ Sharing (UNOS). NKR matches in an essentially greedy fashion. UNOS started by matching once per month, then moved to twice per month, weekly, and now 2+ times per week in part to reduce the “failure rate” caused by competition with the fast-matching NKR. We see this behavior replicated in our model, and can quantify social welfare loss as well. Combinations of analytic and simulation results of this nature have set policy in kidney exchanges before (e.g., Dickerson *et al* 2012 and Dickerson and Sandholm 2015 have set parts of UNOS policy), and our model could help inform this debate.

While our research can inform policy discussions, it is important to have a separate conversation about the ethics of different regulatory and policy changes and how these can impact different populations (for example, by better-serving short-lived patients at the expense of long-lived ones), and our research is not intended to be prescriptive on those issues. That said, since we cannot use money directly to match supply and demand in a matching market, the market/policy designer’s toolkit must consider other options. For example, in kidney exchange, one could offer increased priority in the future on the deceased-donor waitlist (living donor kidney grafts typically survive 10-15 years before another transplant is needed) if they were to go to a patient market rather than a greedy [Ashlagi and Roth, 2014; Hajaj *et al.*, 2015; Toulis and Parkes, 2015; Sönmez and Unver, 2017].

Acknowledgments

Zhuoshu Li and Sanmay Das are grateful for support from the US National Science Foundation under Award #1527037.

References

- [Akbarpour *et al.*, 2017] Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. Thickness and information in dynamic matching markets, 2017. Working paper. Initial version appeared at ACM Conference on Economics and Computation (EC-14).
- [Arnosti *et al.*, 2014] Nick Arnosti, Ramesh Johari, and Yash Kanoria. Managing congestion in decentralized matching markets. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, page 451, 2014.
- [Ashlagi and Roth, 2014] Itai Ashlagi and Alvin E Roth. Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics*, 9(3):817–863, 2014.
- [Bertsimas *et al.*, 2013] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. Fairness, efficiency, and flexibility in organ allocation for kidney transplantation. *Operations Research*, 61(1):73–87, 2013.
- [Bloch and Cantala, 2013] Francis Bloch and David Cantala. Markovian assignment rules. *Social Choice and Welfare*, 40(1):1–25, 2013.
- [Bragdon *et al.*, 2010] Joseph H Bragdon, Ali Hortaçsu, and Dan Ariely. Matching and sorting in online dating. *American Economic Review*, 100(1):130–163, 2010.
- [Das and Tsitsiklis, 2010] Sanmay Das and John N. Tsitsiklis. When is it important to know you’ve been rejected? a search problem with probabilistic appearance of offers. *Journal of Economic Behavior and Organization*, 74:104–122, 2010.
- [Das *et al.*, 2015] Sanmay Das, John P. Dickerson, Zhuoshu Li, and Tuomas Sandholm. Competing dynamic matching markets. In *Proceedings of the Conference on Auctions, Market Mechanisms and Their Applications (AMMA)*, 2015.
- [Dickerson and Sandholm, 2015] John P. Dickerson and Tuomas Sandholm. FutureMatch: Combining human value judgments and machine learning to match in dynamic environments. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 622–628, 2015.
- [Dickerson *et al.*, 2012] John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. Optimizing kidney exchange with transplant chains: Theory and reality. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 711–718, 2012.
- [Hajaj *et al.*, 2015] Chen Hajaj, John P. Dickerson, Avinatan Hassidim, Tuomas Sandholm, and David Sarne. Strategy-proof and efficient kidney exchange using a credit mechanism. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 921–928, 2015.
- [Hall *et al.*, 2017] Jonathan V. Hall, John J. Horton, and Daniel T. Knoepfle. Labor market equilibration: Evidence from Uber, 2017. Working paper Revision requested from American Economic Review.
- [Jones and Teytelboym, 2017] Will Jones and Alex Teytelboym. The refugee match, 2017. Working paper. Revision requested from Refugee Survey Quarterly.
- [Monderer and Shapley, 1996] Dov Monderer and Lloyd S Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- [Nisan *et al.*, 2007] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic Game Theory*. Cambridge University Press Cambridge, 2007.
- [Sönmez and Unver, 2017] Tayfun Sönmez and Utku Unver. Market design for living-donor organ exchanges: An economic policy perspective. Technical report, Boston College Department of Economics, 2017.
- [Toulis and Parkes, 2015] Panos Toulis and David C. Parkes. Design and analysis of multi-hospital kidney exchange mechanisms using random graphs. *Games and Economic Behavior*, 91(0):360–382, 2015.
- [Weibull, 1995] Jörgen Weibull. *Evolutionary Game Theory*. MIT Press, 1995.