

CROSCUTTING AREAS

On Matching and Thickness in Heterogeneous Dynamic Markets

 Itai Ashlagi,^a Maximilien Burq,^b Patrick Jaillet,^{b,c} Vahideh Manshadi^d

^aDepartment of Management Science and Engineering, Stanford University, Stanford, California 94305; ^bOperations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; ^cDepartment of Electrical Engineering and Computer Science and Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; ^dSchool of Management, Yale University, New Haven, Connecticut 06511

Contact: ashlagi@stanford.edu (IA); burq.maximilien@gmail.com (MB); jaillet@mit.edu (PJ); vahideh.manshadi@yale.edu,

 <http://orcid.org/0000-0001-9103-7797> (VM)

Received: March 4, 2016

Accepted: September 29, 2018

Published Online in Articles in Advance:
June 27, 2019

Subject Classifications: networks/graphs: matchings; networks/graphs: stochastic; probability: stochastic model applications; probability: markov processes

Area of Review: Games, Information, and Networks

<https://doi.org/10.1287/opre.2018.1826>

Copyright: © 2019 INFORMS

Abstract. We study dynamic matching in an infinite-horizon stochastic market. Although all agents are potentially compatible with each other, some are hard to match and others are easy to match. Agents prefer to be matched as soon as possible, and matches are formed either bilaterally or indirectly through chains. We adopt an asymptotic approach and compute tight bounds on the limit of waiting time of agents under myopic policies that differ in matching technology and prioritization. We find that when hard-to-match agents arrive less frequently than easy-to-match ones, (i) bilateral matching is almost as efficient as chains (waiting times scale similarly under both, though chains always outperform bilateral matching by a constant factor), and (ii) assigning priorities to hard-to-match agents improves their waiting times. When hard-to-match agents arrive more frequently, chains are much more efficient than bilateral matching, and prioritization has no impact. Furthermore, somewhat surprisingly, we find that in a heterogeneous market and under bilateral matching, increasing the arrival rate of hard-to-match agents has a nonmonotone effect on waiting times. This behavior is in contrast with that of a homogeneous dynamic market, where increasing arrival rate always improves waiting time, and it highlights fundamental differences between heterogeneous and homogeneous dynamic markets.

Funding: P. Jaillet acknowledges support from the National Science Foundation [Grant OR-1029603] and Office of Naval Research [Grants N00014-12-1-0033 and N00014-15-1-2083].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/opre.2018.1826>.

Keywords: dynamic matching • barter exchange • kidney exchange • market design

1. Introduction

This paper is concerned with the problem of matching in a dynamic marketplace, where heterogeneous agents arrive over time to the market looking to exchange an indivisible item for another compatible item. A key feature of the market is its exogenous thickness, as determined by the types of agents and their arrival rates to the marketplace. For example, in kidney exchange, some patient–donor pairs are very hard to match, whereas others are very easy to match. In online labor markets, employers have different qualification requirements and workers have different skills.

Efficiency is determined by the matching policy and the matching technology. The former determines which exchanges to be implemented and when, and in particular which priorities, to assign to different types of agents. The latter determines the forms of matches that can take place. For example, whereas kidney exchanges were first conducted through bilateral exchanges (two-way cycles; Roth et al. 2005), multihospital platforms are now facilitating many of

their transplants through chains initiated by altruistic donors (Anderson et al. 2015). In many matching markets, such as dating, only bilateral matches take place.

We are interested in the behavior of simple *myopic policies* under different matching technologies and different thickness levels of the market. Myopic policies form matches as soon as they become available, but may vary with respect to how they prioritize agents in the events of ties. Our framework will allow to discuss policy questions such as: What is the effect of prioritizing different types of agents? How does disproportional change in arrival of different types influence market efficiency? What is the impact of merging matching marketplaces with different thickness levels on different types?

Two comments are in place. First, restricting attention to myopic policies is motivated by current practices in kidney exchange platforms in the United States. Ashlagi et al. (2017) uses simulations based on empirical data from multiple exchange programs to show matching myopically is nearly harmless.

Anderson et al. (2017) arrive at a similar conclusion in theoretical work. Although they consider a stylized model with homogeneous agents, their result can be generalized to our heterogeneous model. (Although this is not the focus of this paper, see Section 1.1 for further details regarding their model and Section EC.8 in the online appendix for a generalization of their result.)

Second, the literature on dynamic matching in sparse environments has focused on homogeneous agents (Akbarpour et al. 2014, Anderson et al. 2017). The motivation for this paper stems from the heterogeneity of agents in the marketplace.

For our purposes, we propose a simple infinite-horizon model with two types of agents, easy to match (E) and hard to match (H). Agents of each type T arrive to the marketplace according to an independent Poisson process with rate λ_T . Each agent arrives with an indivisible item that she wishes to exchange. We assume a stochastic demand structure, where each agent of type T finds the item of any other agent compatible independently with probability p_T . A key feature of the model is that p_H is significantly smaller than p_E . Agents are indifferent between compatible items but prefer to be matched as early as possible. Moreover, agents in our model depart the market after being matched. We therefore adopt the average waiting time of agents in steady state as a measure for efficiency; more precisely, we focus on the average waiting time of H agents, because the waiting time of E agents is negligible compared with that of H agents. For a more detailed discussion, see Section 2. Although our model is highly stylized, it captures some important features observed in kidney exchange pools (see Section 2.1, where we provide a brief background that further motivates this study).

Two settings are considered, distinguished by how matchings are formed (feasible exchanges): bilateral (two-way cycles) and chains. Our main findings are the following. First, we find that market thickness plays a crucial role in the desired matching technology; when easy-to-match agents arrive more frequently to the market than hard-to-match ones, the average waiting time of H agents scales similarly under chains and bilateral matchings. But there is a sharp increase in the average waiting time of H agents as soon as hard-to-match agents arrive more frequently, highlighting the importance of chains in marketplaces with a majority of hard-to-match agents. Second, we find that, under bilateral matching, increasing arrival rates of hard-to-match agents may negatively affect hard-to-match agents by increasing their waiting times. Under chains, however, increasing arrival rates always shortens waiting times. Third, the impact of prioritization in bilateral matching also depends on the market composition; when hard-to-match agents are the minority type, assigning them

priority improves their waiting times. However, when they are in majority, such prioritization has no significant impact. To be precise, we are only able to prove that prioritizing H agents leads to shorter or equal waiting times (of H agents); however, numerically we confirm that such prioritization indeed leads to strictly shorter waiting times.

Next we describe our results more formally under the bilateral and chain settings. In our analysis, we compute the average waiting time of H agents under various myopic policies as $p_H \rightarrow 0$.

1.1. Bilateral Matching

Two myopic policies are considered for bilateral matching, differing in the types of agents they prioritize. Although it appears natural to prioritize hard-to-match agents, it is also interesting to consider the prioritization of easy-to-match agents, as these may have better outside options. (In practice, agents may leave the market for a variety of reasons including finding an outside option.) Under a stochastic departure model, shorter waiting times correspond to fewer departures because both quantities are proportional with the market size. We leave the rigorous treatment of a model with departure as an open question. We find that, regardless of how agents are prioritized, when $\lambda_H < \lambda_E$, waiting time scales with $1/p_H$, and when $\lambda_H > \lambda_E$, waiting time scales with $1/p_H^2$. When easy-to-match agents arrive more frequently, prioritizing H agents results in shorter waiting times than prioritizing E agents. However, when $\lambda_H > \lambda_E$, the average waiting time in the limit is identical under both types of priorities.

We further provide comparative statics for the case in which H agents are prioritized. Increasing λ_E always decreases waiting times. However, the average waiting time is nonmonotone when increasing λ_H . It has an increasing trend up to a certain threshold, which depends on λ_E , and then it decreases. (Note that in a homogeneous model with only H agents, Little's law implies that increasing λ_H always decreases waiting times.) These findings have two main implications: (i) thickening the market by increasing arrival rates of hard-to-match agents can result in longer waiting times depending on the existing arrival rates, and (ii) merging two marketplaces with different compositions, that is, different ratios between the two arrival rates, may not be beneficial for both.

1.2. Chain Matching

Under the chain setting, we consider policies termed ChainMatch(d) for markets endowed with d altruistic donors who initiate chains that continue indefinitely. In a chain, each agent is matched by (receives an item from) some agent and matches another. Whenever the last agent of a chain can match a new arriving

agent, the policy forms a new *chain segment*, which is a maximal sequence of matches resulting from a *local search*, in which the next matched agent is selected randomly while breaking ties in favor of H agents (so the policy does not always identify the longest possible chain segment, which requires a global search and may be computationally hard). We prove an upper bound on the average waiting time that scales with $1/p_H$ for all positive arrival rates. We also find that even in the regime $\lambda_E > \lambda_H$, where the waiting time scales similarly under both matching technologies, chains result in lower waiting times than bilateral matching.

We provide comparative statics over the arrival rates of both types. We show (analytically for $p_E = 1$ and numerically for $p_E < 1$), that the average waiting time decreases when the arrival rate of either type increases. When $p_E = 1$, we further find that the average waiting time is independent of the constant d . Similar patterns hold numerically when $p_E < 1$. Finally, we are able to compute the average chain segment (which plays an important operational role for example in kidney exchange). An increase in λ_E or d decreases the average length of a chain segment. In contrast, increasing λ_H has the opposite effect.

Next we provide brief intuition for some of the main findings, beginning with why the market composition and the desired matching technology are tightly connected. Under the bilateral setting, when easy-to-match agents arrive more frequently, almost all hard-to-match agents will be matched with easy-to-match ones resulting in a scaling of $1/p_H$; on the other hand, when hard-to-match agents arrive more frequently, many of them will have to match with each other, resulting in a scaling of $1/p_H^2$, which is the inverse of the probability that two H agents can match each other. In contrast, matching through chains does not require such “coincidence of wants” between pairs of H agents even when H agents are the majority (Roth et al. 2007). This results in a waiting time that scales with $1/p_H$ regardless of the composition. We further find that the heterogeneity in the marketplace may lead to nontrivial effects when increasing participation.

The intuition for why, in the bilateral setting, H agents may be harmed when attracting more H agents to the market is the following: When $\lambda_H < \lambda_E$, increasing λ_H reduces the chance that an existing H agent will match with the next arriving E agent. However, when $\lambda_H > \lambda_E$, increasing λ_H implies, on one hand, that more H agents must match with each other, and, on the other hand, it reduces the time to form H - H matches. The first effect is initially stronger, but the second effect dominates once λ_H is sufficiently large. A similar effect happens in kidney exchange where O–A blood type patient–donor pairs that cannot match with each other compete to match

with scarce pairs with blood type O donors. Note, however, that in our setting, all agents can potentially match with each other; in particular, this effect extends to sets of pairs that are blood type compatible with each other, like O–O pairs, some of which are much harder to match than others. We elaborate and provide intuition for other results throughout this paper.

Understanding the impact of market composition by providing comparative statics requires us to not only compute the scaling of asymptotic behavior of average waiting time, but also to characterize the exact limits. Such exact characterization in a heterogeneous model is particularly challenging, as we need to analyze two-dimensional Markov chains. For bilateral matching policies, we directly analyze the underlying two-dimensional spatially nonhomogeneous random walks. One of the main challenges in our analysis is the need to jointly bound the distribution in both dimensions, because applying methods such as Lyapunov functions or analyzing marginal probability distributions would not result in tight bounds. In doing so, we prove two auxiliary lemmas on concentration bounds for a general class of two-dimensional random walks that can be of interest for studying similar random walks that may arise in other applications. For chain policies, we first couple the underlying Markov process with a one-dimensional process where no E agent joins the market. Analysis of the resulting one-dimensional Markov chain presents new challenges, as transitions between nonneighboring states happen because of the possibility of forming arbitrarily long chain segments. However, we show that the chain-segment formation process exhibits a memoryless property, which proves helpful in computing the waiting time limits.

1.3. Related Work

A close stream of related papers study dynamic matching in models, in which agents’ preferences are based on compatibility, that is, agents are indifferent between whom they match with (Ünver 2010, Akbarpour et al. 2014, Anderson et al. 2017).

The impact of the matching technology is addressed in markets comprised of only easy-to-match agents (Ünver 2010; with multiple coarse types) or only hard-to-match ones (Anderson et al. 2017). Ünver (2010) finds that short cycles are sufficient for efficiency. Anderson et al. (2017) consider markets in which all agents are ex ante symmetric and hard to match. They study the waiting-time scaling behavior of myopic policies that attempt to match each agent upon arrival in three settings of exchanges, two-way exchanges, two- and three-way exchanges, and chains, and they find that moving from two-way or three-way exchanges to chains significantly reduces the average

waiting time. See also Dickerson et al. (2012b), who demonstrate the benefit of chains using simulations in dynamic kidney exchange pools. Our paper bridges the gap by looking at a model with both hard- and easy-to-match agents and thus allowing for different levels of thickness in the market. Ding et al. (2015) study a similar two-type model in a static setting and quantify the effectiveness of matching through chains taking a novel random walk approach.

The papers above also find that, by and large, myopic policies are near optimal: Ünver (2010) analyzes a kidney exchange model with different types and deterministic compatibility structure across types and finds that matching upon arrival is near optimal, even though some waiting with certain types to facilitate three-way exchanges adds some benefits (see also Gurvich and Ward 2014, who study a similar compatibility-based inventory control model). Anderson et al. (2017) consider a homogeneous model without departures (similar to our model with $\lambda_E = 0$) and find that there is little benefit from waiting before matching under both matching technologies of short cycles and chains (the waiting-time scales with the same factor with or without waiting before matching). Akbarpour et al. (2014) consider a homogeneous model with departures and find that the optimality gap of the policy that matches without waiting remains constant as the match probability decreases. Moreover, using data-driven simulations, Ashlagi et al. (2017) study the impact of match-run frequency and show that among policies that match periodically (e.g., every week or every day), high matching frequencies perform best. Nonmyopic policies have also been studied; for example, Dickerson et al. (2012a) studied forward-looking policies by casting the dynamic matching problem as a high-dimensional dynamic program and developed a heuristic to overcome the curse of dimensionality. This paper builds on these findings and analyzes only myopic policies that search for a match upon arrival of a new agent.

We elaborate on the relation between our paper and that by Anderson et al. (2017), which is closest to our paper. Studying myopic policies under a homogeneous setting resulted in valuable insights. Some insights, however, do not carry over to heterogeneous settings like kidney exchange (see Section 2.1). For instance, merging markets is often sought as a solution to improve efficiency. A homogeneous model predicts that increasing arrival rates (or merging markets) will always decrease waiting times. In contrast, we find that merging heterogeneous markets may not decrease waiting times for both markets. The homogeneous model by Anderson et al. (2017) predicts very infrequent but very long chain segments. Our model predicts shorter chain segments,

which better fits empirical evidence. (Chain segments typically consist of only a few pairs.) Furthermore, we remark that some questions cannot be addressed in a homogeneous setting. For instance, kidney exchange programs attempt to attract easy-to-match pairs (Ashlagi and Roth 2014), but the impact of such an increase cannot be investigated in a homogeneous model. As another example, exchange programs usually assign high priority to hard-to-match pairs; the effect of such prioritization cannot be studied in a homogeneous model. Overall it is natural and important to study richer models to address relevant policy questions.

Another stream of related research considers models of agents' preferences that do not depend only on compatibility. These papers find that policies that match without waiting are inefficient (Doval 2014, Kadam and Kotowski 2018, Baccara et al. 2015, Fershtman and Pavan 2015) because some waiting can improve the quality of matches (for related results in queueing models, see Bloch and Cantala 2014, Leshno 2014).

Our work is also related to the problem of matching multiclass customers to multiclass servers studied in queueing literature (e.g., Caldentey et al. 2009, Adan and Weiss 2012).

In our model, an agent can be thought of as a customer–server pair, and the compatibility between any two agents is probabilistic; thus, we will not have a finite number of queues.

Finally, our work is related to the online matching literature that studies online matching in which the underlying graph is bipartite and agents on one side of the graph are all present in the market and only agents on the other side arrive over time (Karp et al. 1990, Goel and Mehta 2008, Feldman et al. 2009, Manshadi et al. 2011, Jaillet and Lu 2013).

1.4. Organization

In Section 2, we introduce the model, policies, and the underlying stochastic processes. In Section 2.1, we provide a brief background on kidney exchange further motivating our framework and study. In Section 3, we present the main theoretical results, and Section 4 complements the results with numerical experiments. Section 5 outlines the main proof ideas and techniques along with the details of Markov chains induced by each policy. Section 6 concludes. For the sake of brevity, we include only proofs of selected results in the main text. The detailed proofs of the rest of the statements are deferred to clearly marked sections in the online appendix.

2. Model

We study an infinite-horizon dynamic matching market, where each arriving agent is endowed with a single

item she wants to exchange for another item she finds compatible. Agents are indifferent between compatible items and wish to exchange as early as possible, their cost of waiting being proportional to the waiting time.

There are two types of agents, H and E , referred to as hard to match and easy to match, respectively. Beginning at time $t = 0$, agents of type $T \in \{H, E\}$ arrive to the market according to an independent Poisson process with rate $\lambda_T > 0$.

Any agent of type H (E) finds the item of any other agent compatible independently with probability p_H (p_E). Our analysis is asymptotic in $p_H \rightarrow 0$, whereas p_E is a fixed constant. So, on average, an H agent finds significantly fewer items compatible than an E agent. We say that an agent j is matched by agent i if agent j receives agent i 's item. An agent leaves the market only when she is matched, that is, she receives a compatible item.

We study matching policies in two different settings, distinguished by how agents can exchange items. In the first setting, two agents can exchange items bilaterally in a cyclic fashion. In the second setting, agents exchange items through chains; at time $t = 0$, there are d special agents called *altruistic agents* who are willing to give an item without getting anything in return. (All other agents that will arrive to the market are regular agents who want to exchange their item for another item.) Note that having altruistic agents is an intrinsic property of the market in the sense that some markets do not have access to such agents. Each agent in a chain receives a compatible item from one agent and gives to the next. At any given time, there are exactly $d \geq 1$ agents who are either altruistic or received an item but have not given their item. The latter are called *bridge agents*. We sometimes refer to altruistic agents also as bridge agents. The transactions between two bridge agents in a given chain is called a chain segment. We assume that matches in a chain segment are conducted instantaneously. 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). In our model, the average waiting time of one type of agents is equivalent to the average number of agents of that type in the marketplace divided by the arrival rate of that type, because these two quantities are proportional to each other by Little's law.

It is convenient to think about the 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 from i to j means that agent j finds agent i 's item compatible. Let $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$ denote the (observed) compatibility graph at time t . When a new agent arrives, directed edges are formed in each direction independently, and with probabilities corresponding to the agents' types, between the arriving agent and each agent in the marketplace. A bilateral exchange is a directed cycle of length two in the compatibility graph, and a chain segment is a directed path in this graph starting from a bridge or altruistic agent.

We study the following myopic policies, which attempt to match agents upon arrival.

Definition 1 [BilateralMatch(T) for $T \in \{H, E\}$]. Upon arrival of a new agent, if a cycle of length 2 can be formed with the newly arrived agent, it is removed. If more than one such cycle exists, priority is assigned to cycles with agents of type T . Further ties are broken uniformly at random.

Definition 2 [ChainMatch(d)]. There are d bridge or altruistic agents in the market at any given time. We describe first the policy for $d = 1$. Consider a new arriving agent i_1 . If i_1 does not have an incoming edge from the bridge agent, then no matches happen. Otherwise, a chain segment begins with matching i_1 by the bridge agent and advances as follows. First, we search for an unmatched H agent that has an incoming edge from i_1 . If there is one or more such H agents, we select one uniformly at random; otherwise, if no such an H exists, we search for an unmatched E agent that has an incoming edge from i_1 (again breaking ties uniformly at random). This process repeats itself immediately from the selected agent (selected agents cannot be reselected) until we reach an agent that cannot match any other agent, forming a disjoint path. All agents in the disjoint path leave the market except the last agent who then becomes a bridge agent.

When there are $d > 1$ altruistic/bridge agents, if there is at least one directed edge from one of them to the newly arrived agent, one of such edges is selected uniformly at random. As the process moves forward, each altruistic agent eventually gives her item to an arriving agent and starts a chain.

Under the ChainMatch(d) policy, upon arrival of a new agent, a maximal chain segment (path) is identified through local search originating from a bridge agent. Our local search chain-segment formation process bears similarity to phase 1 of the two-phase clearing procedure of Ding et al. (2015). Note that the chain segment has a positive length if and only if at least one bridge/altruistic agent has a directed edge to the new agent.

For brevity, we often refer to BilateralMatch(E), BilateralMatch(H), and ChainMatch(d), by \mathcal{B}_H , \mathcal{B}_E ,

and $\mathcal{C}(d)$, respectively. All the policies above are Markov policies, and thus define a continuous-time Markov chain (CTMC). The following observation will allow us to ignore the edges within the market when analyzing the underlying stochastic processes.

Observation 1. For each policy \mathcal{B}_H , \mathcal{B}_E , and $\mathcal{C}(d)$, we can construct a two-dimensional CTMC where states correspond to the number of waiting H and E agents and that has a consistent evolution with that of the market under the corresponding policy.

The observation is immediate for the bilateral policies \mathcal{B}_H and \mathcal{B}_E ; because of their myopic behavior, there are no two-length cycles in the market except with a new arriving agent, implying that the corresponding Markov chains can be fully specified using only the set of vertices. For the $\mathcal{C}(d)$ policy, the observation is more subtle. Note that under this policy, there is no outgoing edge from a bridge agent to any waiting agent, again because of the myopic behavior of the policy. The first time we examine whether there is an edge from i to j , we effectively flip a bias coin with probability p_H (p_E) if the agent j is of type H (E). Importantly, we examine at most once whether a directed edge from i to j exists by the definition of the policy, because i either leaves the market or becomes a bridge agent, in which case it will never match to j . Because both the edge formation and the matching policies do not depend on agents' identities (rather only on their types), we can merely keep track of the number of agents of each type.

In the remainder of this paper, for any policy \mathcal{P} , we focus on the state space $\{[H_t^\mathcal{P}, E_t^\mathcal{P}]; t \geq 0\}$, which captures the number of hard- and easy-to-match agents at any time t , and we denote the corresponding transition rate matrix by $Q^\mathcal{P}$.

Given the *self-regulating* dynamic undergoing each matching process, one would expect that all three (irreducible) CTMCs reach steady state. A rigorous statement and proof is provided in Section EC.7 of the online appendix. Hereafter, we are concerned only with steady-state analysis. For policy \mathcal{P} , we denote its steady-state distribution by $\pi^\mathcal{P}$. The random vector $[H^\mathcal{P}, E^\mathcal{P}]$ is the random number of H and E agents in steady state; that is, the vector is distributed according to distribution $\pi^\mathcal{P}$. Finally, we define $w_H^\mathcal{P}$ ($w_E^\mathcal{P}$) to be the average waiting time of type H (E) agents under policy \mathcal{P} . Little's law implies that

$$w_H^\mathcal{P} = \frac{\mathbf{E}[H^\mathcal{P}]}{\lambda_H} \text{ and } w_E^\mathcal{P} = \frac{\mathbf{E}[E^\mathcal{P}]}{\lambda_E}. \tag{1}$$

Because in our model $p_H \rightarrow 0$ while p_E is kept constant, and all policies are myopic, one would expect that $w_H^\mathcal{P}$ is negligible compared with $w_E^\mathcal{P}$. We verify this

claim below using numerical simulations and analytical proofs (see Lemmas EC.1 and EC.2 in the online appendix). We therefore focus on analyzing the average waiting time of H agents under different policies.

In Section 3, we derive asymptotic results ($p_H \rightarrow 0$) for $w_H^\mathcal{P}$ for different sets of parameters λ_H , λ_E , and p_E . We note that $w_H^\mathcal{P}$ is indeed a function of four parameters, and a more precise notation would be $w_H^\mathcal{P}(\lambda_H, \lambda_E, p_H, p_E)$, but we drop these parameters for the sake of brevity.

2.1. Motivating Application: Kidney Exchange

2.1.1. Background. There is a large shortage of kidneys for transplants (as of 2017, the average waiting time was between three and five years in the United States), and many live donors are incompatible with their intended recipients. Kidney exchange allows such patient–donor pairs to swap donors so that each patient can receive a kidney from a compatible donor. There have been efforts to create large platforms to increase opportunities for kidney exchanges (Roth et al. 2004, Nikzad et al. 2017).

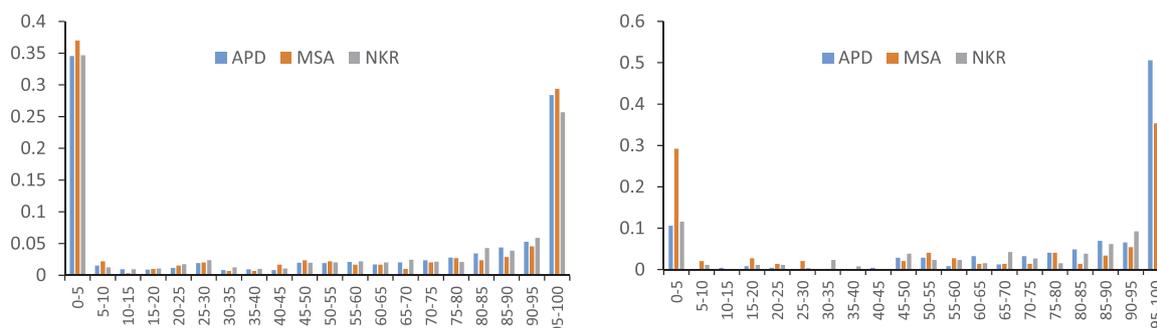
Exchanges are conducted through cycles or chains (for a detailed description of kidney exchange, see Sönmez et al. 2017). Typically, pairs do not give a kidney prior to receiving one. This creates logistical barriers requiring cycles to be limited to two or three pairs. Chains, however, can be organized non-simultaneously, and thus can be longer (Roth et al. 2006, Rees et al. 2009).

For a transplant to take place, the patient needs to be both blood type and tissue type compatible with a donor. The common measure of patient sensitivity is the panel reactive antibody (PRA), which captures the likelihood the patient is tissue type incompatible with a donor chosen at random in the population, based on her antibodies.

Numerous kidney exchange platforms operate in the United States, varying in size, composition, and policies. Some are national platforms (with many participating hospitals) like the Alliance for Paired Donation (APD) and the National Kidney Registry (NKR). Others are regional or even single-center programs like Methodist Hospital in San Antonio (MSA).

2.1.2. Data. Next we provide some figures about the pool composition. Kidney exchange platforms are selected to have a large fraction of highly sensitized patients (Ashlagi et al. 2012). Figure 1 (left panel) plots the PRA distributions of patients enrolled at the NKR, APD, and MSA. Most patients are either highly sensitized (PRA above 95) or have low sensitization (PRA below 5). Note that blood type compatibility is not incorporated in this aggregate PRA distribution. Figure 1 (right panel) provides the same distributions

Figure 1. (Color online) PRA Distributions of Patients Enrolled at the NKR (January 2012 to December 2014), APD (January 2007 to August 2016), and MSA (July 2013 to February 2017)



Note. The left panel shows all patients, and the right panel shows patients belonging to blood type-compatible pairs.

for patients belonging to blood type-compatible pairs (e.g., O–O patient–donor pairs), who can match with each other if they are tissue type compatible. These distributions can be roughly viewed as bimodal; note that among blood type-compatible pairs, there are more patients with high-sensitization than low-sensitization.

The percentage of high-PRA patients also varies across programs outside the United States. In Australia, 42% of registered candidates have a PRA greater than 90 (Ferrari et al. 2012), and in the United Kingdom, 46% of patients have a PRA greater than 85 (Johnson et al. 2008), whereas in Canada, only 36% of patients have a PRA of 80 or more (Malik and Cole 2014). In the Netherlands, Glorie et al. (2014) estimate that 30% of patients have a PRA above 97.

Like the PRA distribution, the pool compositions also vary with respect to blood type distributions of patient–donor pairs. Ashlagi et al. (2017) report that O–O pairs make up 26% of the MSA pool but only 20.81% of the APD pool; the percentages of pairs that contain an O donor in the APD and MSA pools are 34.4% and 50.3%, respectively.

These platforms also differ in size. During the period of the data, the MSA and APD had enrollment rates of roughly 180 pairs per year, whereas the NKR had an enrollment rate of about 360 pairs per year. Access to altruistic donors also varies, with roughly 1, 8, and 50 altruistic donors per year for the MSA, APD, and NKR, respectively.

2.1.3. Matching. Although more than 80% of the transplants at the NKR and the APD have been conducted through chains (Anderson et al. 2015, National Kidney Registry 2017), some platforms (such as the MSA and those in Belgium and the Czech Republic) match their pairs mostly through cycles because of limited access to altruistic donors. In countries like France, Poland, and Portugal, chains are infeasible because altruistic donations are not permitted (Biro et al. 2017).

Exchange platforms in the United States adopt typically myopic-like matching policies that periodically search for matches. The APD, MSA, and NKR search for exchanges on a daily basis, and United Network for Organ Sharing searches for exchanges biweekly. There is some concern that this behavior is inefficient (and arguably a result of competition). However, numerical simulations by Ashlagi et al. (2017) suggest that in steady state, there is essentially no harm from frequent matching (though having multiple small platforms does harm efficiency). Moreover, MSA is not facing any competition. However, some countries, such as Canada, the United Kingdom, the Netherlands, and Australia, search for exchanges every three or four months (Ferrari et al. 2014).

Matching policies at most platforms assign high weights to highly sensitized patients (easy-to-match patients match quickly; Ashlagi et al. 2017, National Kidney Registry 2017). We note, however, that MSA and the NKR assign high priority to compatible pairs, which are very easy to match. (This is because such pairs could choose to go through a direct transplant if they are not matched quickly.) Platforms typically have multiple desiderata. However, implicit first-order-related goals are to reduce waiting times and facilitate many transplants (National Kidney Registry 2017).

2.1.4. Policy. Various challenges arise from variation across kidney exchange pools with respect to their compositions and even operational issues: What priorities should be assigned to different types of patients? What is the impact of attracting more easy-to-match pairs and even compatible pairs? See also Agarwal et al. (2018) and Sönmez et al. (2017) for incentive schemes toward thickening the pool with such pairs. How important is it to incorporate chains and attract altruistic donors?

There are also several initiatives to merge kidney exchange platforms to increase efficiency and matching opportunities for highly sensitized patients

(e.g., for merging the Austrian and the Czech Republic programs, see Böhmig et al. 2017; for Israel and Cyprus, see Siegel-Itzkovich 2017; furthermore, for augmenting national programs through global kidney exchange, see Nikzad et al. 2017). It is natural to study what the impact of merging programs is on different types of patients.

This paper does not intend to model the details of kidney exchange. However, our stylized model does capture some important features in kidney exchange and will hopefully generate some useful insights.

3. Main Results

We analyze the average waiting time under the myopic policies defined in Section 2. For bilateral matching policies, we identify a stark threshold in the scaling of waiting time when moving from the regime where a majority of arrivals are hard-to-match agents to the regime where the majority of arrivals are easy to match. Such a contrast does not exist when agents are matched through chains. We further study the impact of arrival rates of the two types on the market performance under the three policies.

3.1. Bilateral Matching

This section considers the setting in which agents match only through bilateral exchanges, that is, through two-way cycles.

Theorem 1. *Under the BilateralMatch(H) policy and in steady state, the average waiting time w_H^{BH} satisfies the following:*

- If $\lambda_H < \lambda_E$, then $\lim_{p_H \rightarrow 0} p_H w_H^{\text{BH}} = \frac{\ln\left(\frac{\lambda_E}{\lambda_E - \lambda_H}\right)}{p_E \lambda_H}$.
- If $\lambda_H > \lambda_E$, then $\lim_{p_H \rightarrow 0} p_H^2 w_H^{\text{BH}} = \frac{\ln\left(\frac{2\lambda_H}{\lambda_H + \lambda_E}\right)}{\lambda_H}$.

Theorem 1 provides not only the scaling laws on w_H^{BH} but also the associated constants. The following corollaries provide comparative statics with respect to λ_H .

Corollary 1. *Consider the BilateralMatch(H) policy and fix λ_E . The limiting average waiting time w_H^{BH} increases with λ_H in the interval $\lambda_H < \lambda_E$.*

Corollary 2. *Consider the BilateralMatch(H) policy and fix λ_E . The limiting average waiting time w_H^{BH} increases with λ_H in the interval $\lambda_E < \lambda_H < x^* \lambda_E$ and decreases in the interval $\lambda_H > x^* \lambda_E$, where $x^* \approx 2.18$ is the unique solution of*

$$(x + 1) \ln(2 - 2/(x + 1)) = 1. \tag{2}$$

The above theorem and corollaries provide several messages on the impact of thickness on the performance of bilateral matching. First, the main factor in the asymptotic behavior of w_H^{BH} is which type of agent has a larger arrival rate. Some intuition for the scaling factors is the following. Agents' average waiting

time is inversely proportional to the probability of a bilateral match occurring. Under a myopic bilateral policy, no existing pair of agents in the market can match with each other. For an arriving H agent, the probability of forming a bilateral match with an existing E agent is $p_E p_H$, and with an existing H agent is p_H^2 . When $\lambda_H < \lambda_E$, almost all H agents are matched with E agents, resulting in an average waiting time that scales with $1/p_E p_H$. When H agents arrive more frequently than E agents, there are simply not enough E agents to match with H agents. So a nonnegligible fraction of H agents match with each other and thus the scaling of the average waiting time increases to $1/p_H^2$.

Second, the arrival rates affect the average waiting times directly and not necessarily monotonically. Increasing the arrival rate of E agents always decreases the average waiting time. But this is not the case with H agents. When $\lambda_H < \lambda_E$, the average waiting time of H agents increases with λ_H . So in this regime, where almost all H agents match bilaterally with E agents, increasing λ_H reduces the chance of an arbitrary H agent matching with the next E agent.

When $\lambda_H > \lambda_E$, there is a nonmonotone behavior of the waiting time when increasing λ_H . Increasing λ_H has two effects: (i) more H agents must match bilaterally with their same type, which is a negative effect, and (ii) for an existing H agent, it reduces the time to match with another H agent, which is a positive effect. After a certain threshold, the positive effect from having more H agents dominates the negative effect.

The key insight from the above discussion is that in a heterogeneous market, increasing the arrival rate does not always result in improving the waiting time because H agents impose a negative externality on other H agents under certain market compositions. This cannot be captured in a homogeneous model with only hard-to-match agents (the model studied by Anderson et al. 2017). Finally, we comment on the impact of p_E on the waiting time. When $\lambda_H < \lambda_E$, w_H^{BH} is decreasing in p_E . On the other hand, when $\lambda_H > \lambda_E$, w_H^{BH} is independent of p_E . The intuition is that in the former, all H agents match with E agents, and in the latter, the dominant factor in the average waiting time is due to two-way exchanges between H agents, which is independent of p_E .

The proof of Theorem 1 amounts to analyzing the underlying two-dimensional continuous-time spatially nonhomogeneous random walk. The description of the random walk is presented in Section 5.1, along with a heuristic that helps us guess the right constants and build intuition on the behavior of the random walk. The main idea behind the proof is establishing concentration results for a two-dimensional CTMC where the steady-state distribution decays

geometrically when moving away from the expectation. These concentration results allow us to establish matching lower and upper bounds on $w_H^{\mathcal{B}_H}$ (the proof is outlined in Section 5.1 with details deferred to Section EC.2 of the online appendix). We note that one of the main challenges in our analysis is the need to jointly bound the distribution in both dimensions, because analyzing marginal probability distributions would not result in tight bounds. As a byproduct of our analysis, in Section 5.2, we state two auxiliary lemmas on concentration bounds for a general class of two-dimensional random walks. The corollaries follow from basic analysis of the corresponding constants (as a function of λ_H). Both corollaries are proved in Section EC.2.2 of the online appendix.

Theorem 2. *Under the BilateralMatch(E) policy and in steady state, the average waiting time $w_H^{\mathcal{B}_E}$ satisfies the following:*

- If $\lambda_H < \lambda_E$, then $\frac{\ln\left(\frac{\lambda_E}{\lambda_E - \lambda_H}\right)}{p_E \lambda_H} \leq \lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{B}_E} \leq \frac{\ln\left(\frac{2\lambda_E}{\lambda_E - \lambda_H}\right)}{p_E \lambda_H}$.
- If $\lambda_H > \lambda_E$, then $\lim_{p_H \rightarrow 0} p_H^2 w_H^{\mathcal{B}_E} = \frac{\ln\left(\frac{2\lambda_H}{\lambda_H + \lambda_E}\right)}{\lambda_H}$.

Comparing results of Theorems 1 and 2, we observe that when $\lambda_H < \lambda_E$, the average waiting time of H agents is larger or the same when prioritizing E agents rather than H agents. (Numerical simulations presented in Section 4.2 suggest that prioritizing E agents results in a strictly larger average waiting time.) Nevertheless, the scaling remains the same. However, when $\lambda_H > \lambda_E$, prioritizing E agents does not impact the waiting time of H agents. The intuition is as follows. When $\lambda_H > \lambda_E$, the number of H agents waiting in the market scales as $1/p_H^2$, suggesting that the chance that an E agent does not match immediately upon arrival vanishes. Therefore, assigning priority to E agents is redundant. We note that for neither policy BilateralMatch(H) nor BilateralMatch(E) are we able to derive the exact waiting time behavior when $\lambda_E = \lambda_H$. However, our simulation results, presented in Section EC.10 of the online appendix, suggest that under both policies, the waiting time scales with $1/p_H$.

The proof of Theorem 2 also requires analysis of the underlying two-dimensional continuous-time spatially nonhomogeneous random walk and, in most parts, follows a structure similar to that of the proof of Theorem 1. A detailed description of the random walk is presented in Section 5.3. The proof of the upper and lower bounds is presented in Section EC.3 of the online appendix, where establishing the upper bound requires new ideas beyond the concentration results: we couple the Markov process underlying policy \mathcal{B}_E with another process in which an E agent that cannot form a match upon arrival turns into an H agent. In Section 5.3, we provide a rough intuition on why we cannot close the gap between our upper and lower

bounds on $w_H^{\mathcal{B}_E}$ for the regime $\lambda_H < \lambda_E$. In Section 5.3, we also provide a heuristic argument that leads us to guess that the exact limit is $\frac{\ln\left(\frac{\lambda_E + \lambda_H}{\lambda_E - \lambda_H}\right)}{p_E p_H}$.

3.2. Chain Matching

In this section, we analyze the ChainMatch(d) policy, under which agents match myopically through chains.

3.2.1. Waiting-Time Behavior.

Theorem 3. *Let $d \geq 1$ be a constant (independent of p_H). Under the ChainMatch(d) policy and in steady state, the average waiting time $w_H^{\mathcal{C}(d)}$ satisfies*

$$\lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{C}(d)} \leq \frac{\ln\left(\frac{\lambda_H}{\lambda_E(1-(1-p_E)^d)} + 1\right)}{\lambda_H}.$$

The above theorem establishes an upper bound on $w_H^{\mathcal{C}(d)}$ that scales with $1/p_H$. In Section EC.9 of the online appendix, we also establish a lower bound on $w_H^{\mathcal{C}(d)}$ that scales with $1/p_H$. A stronger result is obtained for the special case in which $p_E = 1$.

Proposition 1. *Let $p_E = 1$, and let $d \geq 1$ be a constant (independent of p_H). Then,*

$$\lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{C}(d)} = \frac{\ln\left(\frac{\lambda_H}{\lambda_E} + 1\right)}{\lambda_H}.$$

Consequently, $\lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{C}(d)}$ decreases with λ_E and λ_H .

First, we discuss the intuition behind Proposition 1, which states that when $p_E = 1$, any constant number of altruistic agents will result in the same behavior of $w_H^{\mathcal{C}(d)}$. The positive impact of having $d > 1$ altruistic agents stems from the increase in probability of starting a new chain segment. When an H agent arrives, the probability that she finds one of the bridge agents acceptable is $1 - (1 - p_H)^d$, which vanishes as $p_H \rightarrow 0$. When an E agent arrives, she will always be matched by one of the bridge agents and proceed to advance the chain segment, and thus there is no advantage in having more than one bridge agent. We also point out that we suspect the result of Proposition 1 holds for any finite p_H , and one can prove such a result using a coupling argument similar to the ones we used in our proofs. However, given that in this paper we are concerned with asymptotic analysis, we do not include such a result.

For $p_E < 1$, we give a heuristic argument (in Section EC.5 of the online appendix) in which we analyze a related three-dimensional random walk by artificially assuming that chain segments advance according to an independent Poisson process with a very high rate μ . (Recall that under the $\mathcal{C}(d)$ policy, chain

segments are formed and executed instantaneously upon arrivals.) The heuristic provides an estimated waiting time that scales as $\ln\left(\frac{\lambda_H + \lambda_E}{\lambda_H(1 - (1 - p_H)^d) + \lambda_E}\right) / (p_H \lambda_H)$. In the limit when p_H approaches zero, the constant becomes $\ln\left(\frac{\lambda_H + \lambda_E}{\lambda_E}\right) / \lambda_H$, which is consistent with Proposition 1. Numerical simulations that are aligned with the result of the heuristic argument are presented in Section 4.5.

The heuristic argument in Section EC.5 of the online appendix and simulation results of Section 4.5 both suggest that $\lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{C}(d)}(p_E) = \lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{C}(d)}(p_E = 1)$. (With a slight abuse of notation, we denote the expected waiting time of H agents under $\mathcal{C}(d)$ policy with parameter p_E by $w_H^{\mathcal{C}(d)}(p_E)$.) This means that, in the limit where $p_H \rightarrow 0$, (1) the variability of $w_H^{\mathcal{C}(d)}(p_E)$ with respect to d or p_E is negligible compared with a term of order $\Theta(1/p_H)$, and (2) for any $0 < p_E \leq 1$, $w_H^{\mathcal{C}(d)}(p_E)$ is decreasing in both λ_E and λ_H . Our simulation results verify these behaviors (see Section 4.3).

The main intuition for why $w_H^{\mathcal{C}(d)}(p_E)$ does not vary with d or p_E can be summarized as follows. Under ChainMatch(d), each E agent, immediately after being matched, initiates a subsegment that includes a sequence of H agents, that is, has the form $E - H - \dots - H$. Note that when $p_E < 1$, a chain segment can consist of multiple subsegments that are initiated by some of the E agents waiting in the market. Denote the expected number of H agents in a subsegment by $\mathbf{E}[\Sigma_E]$. Observe that with high probability, an existing H agent is matched through such a subsegment, because the probability that an arriving H agent starts a subsegment of the form $H - H - \dots - H$ is vanishing. Furthermore, in steady state, the number of H agents who join the market (i.e., they do not match immediately), $\lambda_H(1 - p_H)^d$, must equal the number of agents who match through a subsegment. Thus, the departure rate of H agents from the market equals $\lambda_E \mathbf{E}[\Sigma_E] + o(1)$, where $o(1)$ corresponds to the event that an arriving H agent starts a subsegment. As the result of this balance equation, we have $\mathbf{E}[\Sigma_E] = \lambda_H / \lambda_E + o(1)$; that is, $\mathbf{E}[\Sigma_E]$ does not depend on p_E or d (up to a negligible additive factor of order $o(1)$).

Therefore, regardless of p_E and d , in expectation, each E agent “helps” to match the same number of H agents. The only difference between the cases $p_E = 1$ and $p_E < 1$ is the timing in which E agents form subsegments: When $p_E = 1$, an E agent forms a subsegment immediately upon arrival. On the other hand, when $p_E < 1$, an arriving E agent will join the market with probability $(1 - p_E)^d$. In this case, such an E agent will start a subsegment (of the form $E - H - \dots - H$) after staying in the market for a

random duration, which we denote by X . We can think of X as the delay in matching the H agents in the subsegment that starts with the E agent. With high probability, X remains a constant as p_H approaches zero (this follows from Lemma EC.4 in Section EC.9 of the online appendix, where we show that the expected waiting time of an E agent is a constant). Therefore, the delay caused by the E agent joining the market remains a constant. Consequently, the waiting time of H agents in a market with $p_E < 1$ is within $o(1/p_H)$ of its counterpart in a market with $p_E = 1$. In the online appendix, Section EC.6, we provide further details for the above argument and build on this intuition to provide another heuristic argument to show that $\lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{C}(d)}(p_E) = \lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{C}(d)}(p_E = 1)$.

Finally, we comment on the chain-segment formation process. The ChainMatch(d) policy forms chain segments employing a local search process, and indeed our analysis relies on such a chain-segment formation process. This raises the question of how much the waiting time would improve if we employed a global search (that searched for the longest possible chain segment). A precise comparison is beyond the scope of our work; however, we make the following remarks: (1) At the end of Section 4.3, we numerically study this questions, and we see that advancing chains locally results in a small loss in comparison with policies that search globally for the longest possible chain segment. (2) The lower bound on the waiting time of any anonymous Markovian policy (see Anderson et al. 2017 and Proposition EC.4 in Section EC.8 of the online appendix) implies that the scaling of H agent waiting time cannot be smaller than $1/p_H$ (unless the policy makes E agents wait for a very long time, that is, proportional to $1/p_H$). Theorem 3 shows that the local search method already achieves such a scaling.

Under the ChainMatch(d) policy, the length of a chain segment triggered by a newly arrived agent is unrestricted. As a result, the underlying CTMC is significantly more complicated to analyze than those that arise from bilateral policies, and we need other techniques to prove Theorem 3. To bound $w_H^{\mathcal{C}(d)}$, we couple the underlying Markov chain with a one-dimensional chain, in which E agents that are not matched upon arrival leave the market immediately (Lemma 3). A key property used in the analysis of the coupled one-dimensional chain is that chain-segment formation exhibits a memoryless property. This is because of the local search process used to advance a chain segment, which randomly selects the next agent among all possible agents (favoring H agents). The proof is presented in Section 5.4. Finally, we note that for the special case $p_E = 1$, the original CTMC is a

one-dimensional chain for which we can prove matching upper and lower bounds on the limit of $w_H^{\mathcal{C}(d)}$.

Theorems 1 and 3 together highlight the importance of having altruistic agents that can initiate chains. In the regime $\lambda_H > \lambda_E$, comparing $w_H^{\mathcal{B}_H}$ and $w_H^{\mathcal{C}(d)}$ is straightforward, as the former scales as $1/p_H^2$ but the latter scales only as $1/p_H$. The following corollary (proven in Section EC.4 of the online appendix) states that in the regime $\lambda_H < \lambda_E$, where both $w_H^{\mathcal{B}_H}$ and $w_H^{\mathcal{C}(d)}$ scale as $1/p_H$, ChainMatch(d) performs better.

Corollary 3. *For any $\lambda_H, \lambda_E, p_E$, and d , if $\lambda_H < \lambda_E$, then $\lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{C}(d)} < \lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{B}_H}$.*

In Section 4.4, we further compare BilateralMatch(H) to ChainMatch(d) to understand the importance of attracting easy-to-match agents in markets that have limited access to altruistic agents.

3.2.2. Chain-Segment Length. We analyze here the expected lengths of chain segments formed under the ChainMatch(d) policy. Although we focus on the average waiting time to measure efficiency, lengths of chain segments also play a significant role in the operational efficiency of the market. In kidney exchange, for example, executing a chain segment takes time and bears the risk of match failures. These practical considerations motivate extending the analysis to the limiting behavior of chain segments.

First, we define the chain segment length. Let $[H_k^{\mathcal{C}(d)}, E_k^{\mathcal{C}(d)}]$ denote the (discrete-time) Markov chain embedded in the CTMC $[H_t^{\mathcal{C}(d)}, E_t^{\mathcal{C}(d)}]$ resulting from observing the system at arrival epochs. Note that each time an agent arrives, the Markov chain advances in discrete time from k to $k + 1$. Define

$$L_k = H_k + E_k - H_{k+1} - E_{k+1} + 1,$$

and let L be its corresponding random variable in steady state. If the arriving agent cannot be matched by the bridge agent, she will join the market, and therefore $L_k = 0$; otherwise, a chain segment of length $L_k \geq 1$ will be formed. The following proposition characterizes the chain-segment length in the limit.

Proposition 2. *Under the ChainMatch(d) policy and in steady state,*

$$\lim_{p_H \rightarrow 0} \mathbb{E}[L \mid L \geq 1] = \frac{\lambda_H + \lambda_E(1 - p_E)^d}{\lambda_E(1 - (1 - p_E)^d)} + 1.$$

The proof is presented in Section EC.4 of the online appendix. We note that the expected chain length is decreasing in both λ_E and d , but increasing in λ_H . Intuitively, with more E agents or more bridge agents,

chain segments will be formed at a higher rate and thus be shorter (for a fixed λ_H). However, increasing λ_H does not significantly impact the frequency of chain-segment formation, but given that more H agents join the market within two consecutive chain segments, the length of the chain segment grows.

4. Numerical Studies

In this section, we present a set of numerical simulations that complement the theoretical results of the previous section. In Section 4.1, we look at how merging markets with different compositions affect each market. Section 4.2 explores the impact of giving priorities when using the bilateral matching policy. Section 4.3 presents comparative statics for chain matching when $p_E < 1$, and Section 4.4 highlights the advantage of having chains. Finally, Section 4.5 compares our theoretical bounds (for cases for which we do not have matching upper and lower bounds) to heuristic guesses and simulations.

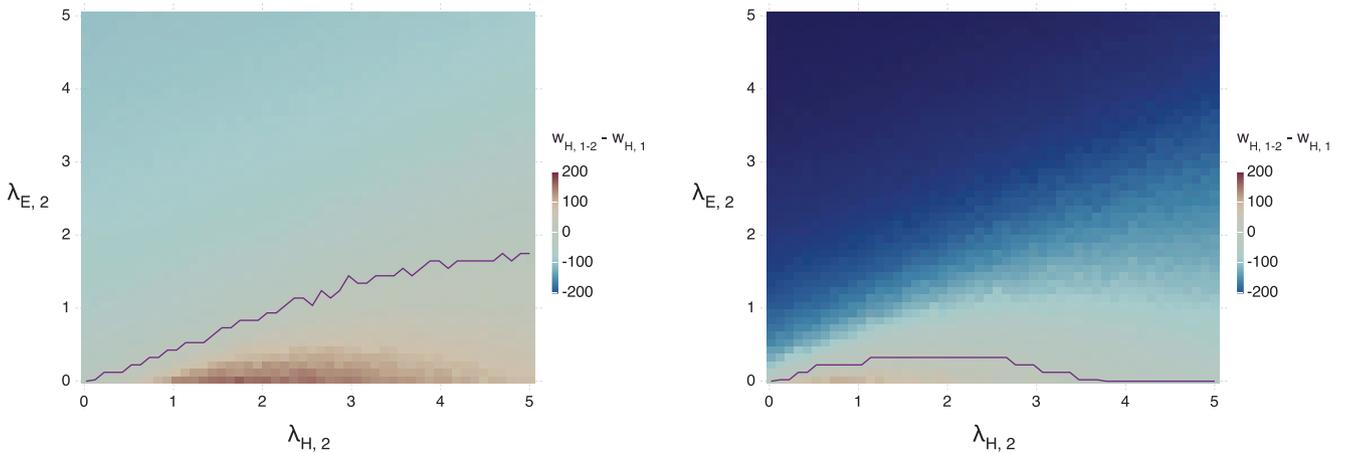
All simulations in this section are conducted by first computing the average number of agents in the market and then applying Little’s law (1). To compute the number of agents, we simulate the discrete-time Markov chain embedded in the corresponding CTMC resulting from observing the system at arrival epochs. We denote by T the number of arrivals (not counting the d initial altruistic agents in the case of $\mathcal{C}(d)$). To remove the transient behavior, the numbers reported correspond to the time average over the second half of the simulation.

4.1. Merging Markets

We consider here the effects from merging two markets, with arrival rates $(\lambda_{H,1}, \lambda_{E,1})$ and $(\lambda_{H,2}, \lambda_{E,2})$ under bilateral exchanges using the BilateralMatch(H) policy. This expands Theorem 1, which provides comparative statics in the limit when p_H tends to zero.

We consider two numerical examples to illustrate these effects. In both examples, the arrival rates to the first market are kept fixed, whereas the arrivals rates to the second market vary. For any pair of arrivals, we compare the waiting time $w_{H,1}$ of H agents in the first market with the average waiting time $w_{H,1-2}$ in the merged market. The results are plotted in Figure 2. Consistent with our prediction, merging can result in one of the markets being worse off. Note that this can happen even if the majority type is the same for both markets (e.g., when $\lambda_{H,1} > \lambda_{E,1}$ and $\lambda_{H,2} > \lambda_{E,2}$). This highlights the effect of arrival rates beyond their impact on the scaling factor. We note that the constants computed in Theorem 1 allow us to determine whether the first market is better or worse off for any $(\lambda_{H,2}, \lambda_{E,2})$, and to compute the boundary separating the two regions, in the limit $p_H \rightarrow 0$.

Figure 2. (Color online) Change in the Waiting Time for H Agents of the First Market, $w_{H,1-2} - w_{H,1}$, as a Function of $(\lambda_{H,2}, \lambda_{E,2})$, for $p_E = 0.5, p_H = 0.02$, After $T = 10^5$ Iterations



Notes. The left plot corresponds to $\lambda_{H,1} = 1, \lambda_{E,1} = 1.3$, and right plot corresponds to $\lambda_{H,1} = 1.3, \lambda_{E,1} = 1$. The purple line separates the region where the waiting time increases after merging (below the line) and the region where it decreases (above the line). Merging the first market with a second market can harm the waiting times of H agents in the first market regardless of whether it has more easy- than hard-to-match agents (left plot) or the opposite (right plot).

4.2. Impact of Priorities in Bilateral Matching

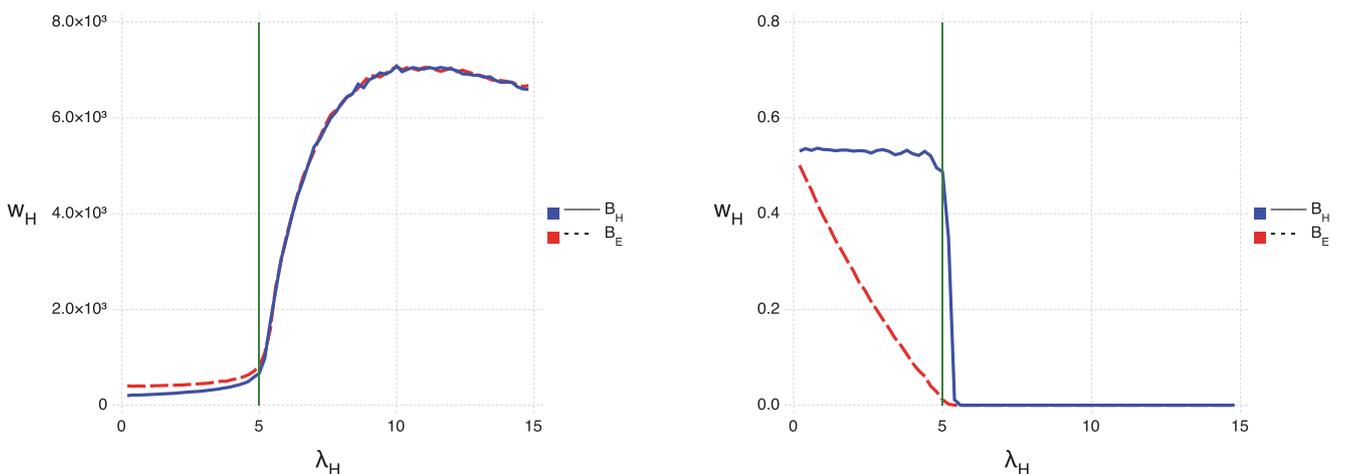
We compare here the average waiting time of H agents under the $\text{BilateralMatch}(H)$ and $\text{BilateralMatch}(E)$ policies. From Theorems 1 and 2, it follows that (i) when $\lambda_H > \lambda_E$, asymptotically, the average waiting time of H agents is the same under both policies, but (ii) when $\lambda_H < \lambda_E$, the average waiting time of H agents under \mathcal{B}_H is at most the average waiting time under \mathcal{B}_E . However, numerical simulations suggest that the average waiting time of H agents is indeed strictly smaller under \mathcal{B}_H than under \mathcal{B}_E (Figure 3, left). For instance, in simulation setting of Figure 3, when $\lambda_H = 4$ and $\lambda_E = 5$, we have $w_H^{\mathcal{B}_E} = 534$, whereas $w_H^{\mathcal{B}_H} = 388$. The average waiting times of E agents are plotted in Figure 3 (right).

The main insight is that the benefit from assigning priority to hard-to-match agents varies based on the composition of the market. Furthermore, our qualitative insights can be useful in understanding the trade-offs that may arise in markets where easy-to-match agents have outside options. For example, when $\lambda_H > \lambda_E$, there is no trade-off from prioritizing E agents. This issue arises in kidney exchange, where very easy-to-match patient–donor pairs (such as compatible pairs) may choose to get transplanted elsewhere.

4.3. Comparative Statics in Chain Matching with $p_E < 1$

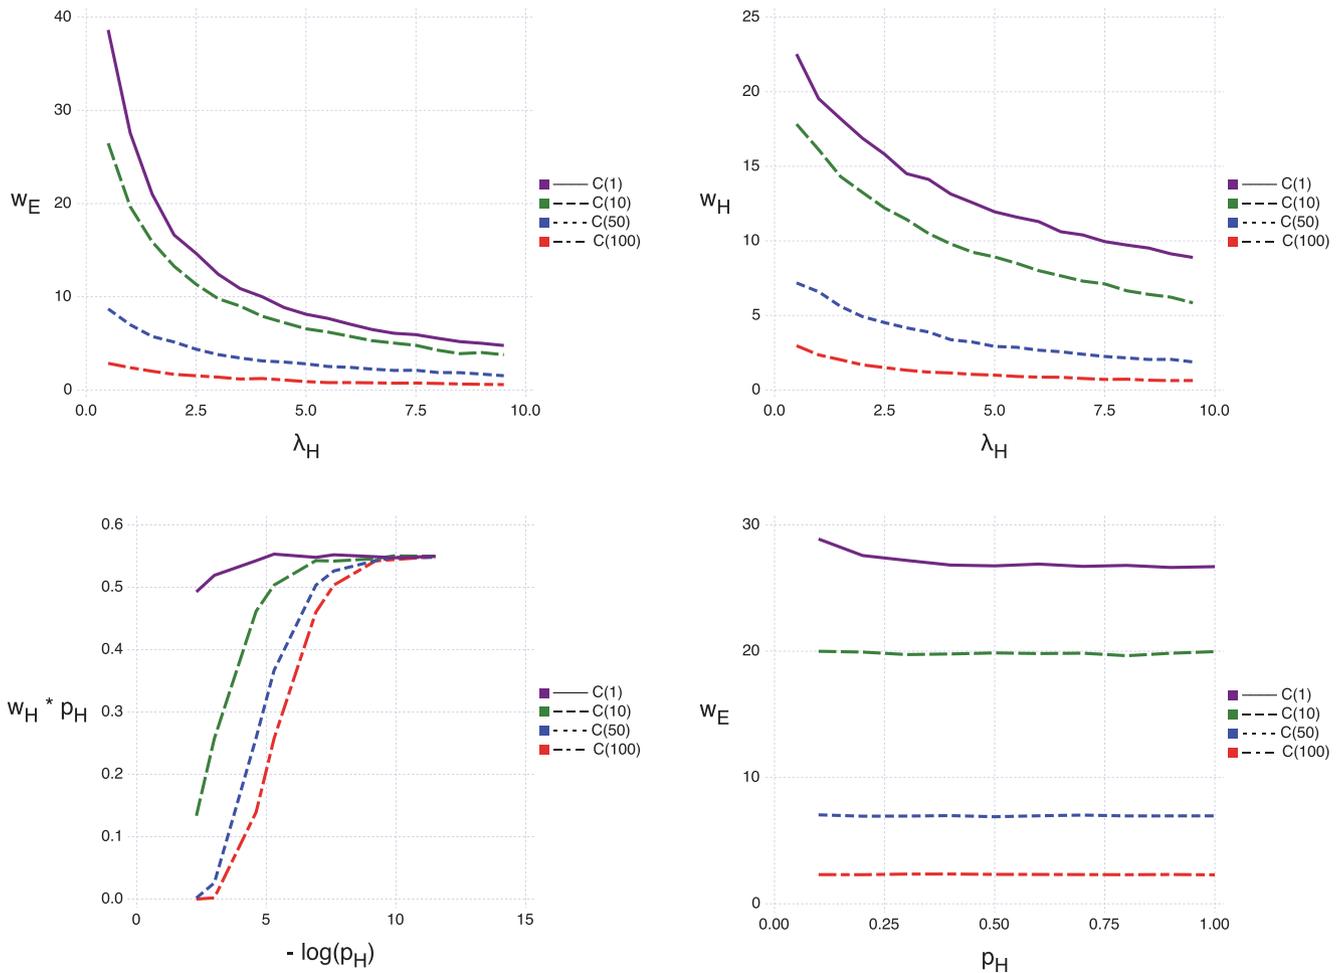
We run simulations using $\text{ChainMatch}(d)$ to numerically explore the effects that varying λ_E, λ_H , and d can have

Figure 3. (Color online) Comparison of the Average Waiting Time of H/E Agents Under the $\text{BilateralMatch}(H)$ and $\text{BilateralMatch}(E)$ Policies



Note. The figure shows w_H (left) and w_E (right) for \mathcal{B}_H and \mathcal{B}_E , as a function of λ_H , for fixed $\lambda_E = 5, T = 2 \cdot 10^6, p_E = 0.5$, and $p_H = 0.002$.

Figure 4. (Color online) Comparative Statics of the Average Waiting Time of H Agents Under the ChainMatch(d) Policy with Respect to Arrival Rates, Number of Altruistic Agents, and Compatibility Probabilities



Notes. The top left shows $w_H^{(d)}$ as a function of λ_E , for varying values of d , for fixed $\lambda_H = 2, T = 10^5, p_E = 0.5$, and $p_H = 0.02$. The top right shows $w_H^{(d)}$ as a function of λ_H , for varying values of d , for fixed $\lambda_E = 2, T = 10^5, p_E = 0.5$, and $p_H = 0.02$. The bottom left shows normalized waiting times (i.e., $p_H w_H^{(d)}$) in the case of chains as a function of $-\log(p_H)$, for varying values of d , for fixed $\lambda_H = 2, \lambda_E = 1, T = 10^5$, and $p_E = 0.5$. The bottom right shows $w_H^{(d)}$ as a function of p_E , for different values of d , for fixed $\lambda_E = 1, \lambda_H = 2, T = 10^6$, and $p_H = 0.02$.

on $w_H^{(d)}$. We find that $w_H^{(d)}$ decreases as the arrival rate of either type increases (Figure 4, top left and top right). Moreover, the value of an additional altruistic agent also diminishes with increasing λ_E, λ_H , or d . Furthermore, as p_H decreases, the impact of d vanishes (Figure 4, bottom left). Recall from Proposition 1 that when $p_E = 1, w_H^{(d)}$ (i) decreases in λ_E , (ii) decreases in λ_H , and (iii) does not depend on d . Although proving (i)–(iii) for $p_E < 1$ remains an open problem, the top left, top right, and bottom left panels of Figure 4 numerically show (i)–(iii), respectively, when $p_E = 0.5$.

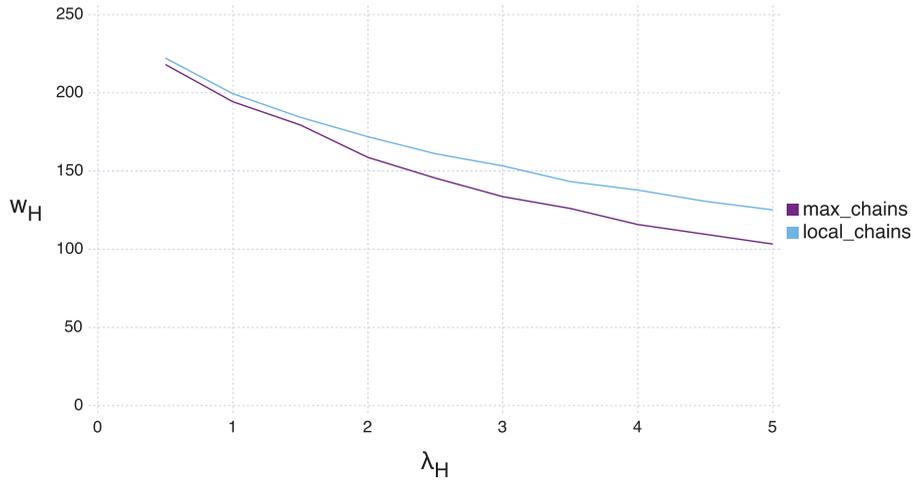
Furthermore, the bottom right plot in Figure 4 plots $w_H^{(d)}$ when p_E ranges from 0.1 to 1, and it shows that $w_H^{(d)}$ does not vary with p_E . (We refer the reader to Section EC.6 of the online appendix for a heuristic

argument on why, in the limit, $w_H^{(d)}$ does not vary with p_E and d .)

Next, in Figure 5, we study the loss from employing a local search for forming chain segments rather than looking for the maximum-length path at each chain segment formation. For this, we define a new policy, *Max-Chains*, that, upon starting a chain segment, searches for the chain segment that maximizes lexicographically the number of H agents matched, while breaking ties over matching more agents over all.

We observe that the benefit of using *Max-Chains* is small when λ_H is small compared with λ_E , and it increases as λ_H increases. If we consider $\lambda_E/2 \leq \lambda_H \leq 2\lambda_E$ as the practical range relevant to the kidney exchange programs, our simulations suggest that the loss ranges between 5% and 15%.

Figure 5. (Color online) Waiting Times w_H for Chains Conducted with Local Search ($\mathcal{C}(d)$) and Max-Chains as a Function of λ_H , for Fixed $\lambda_E = 2$, $T = 10^5$, $p_E = 0.5$, and $p_H = 0.002$

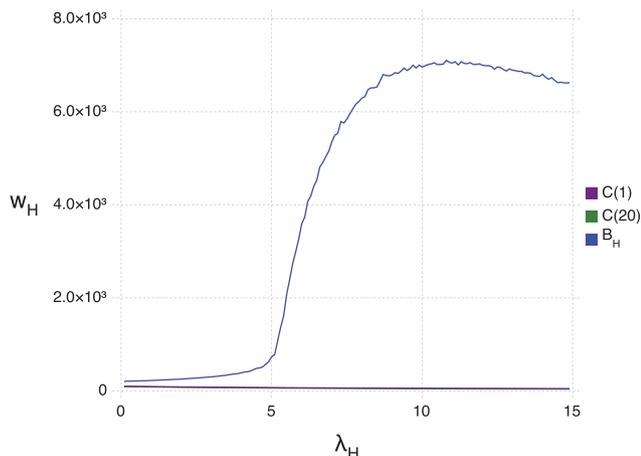


4.4. Impact of the Matching Technology: Bilateral vs. Chain Matching

Theorems 1 and 3 imply that for *any* arrival rates (λ_H, λ_E) , matching through chains even with only one initial altruistic agent [i.e., under ChainMatch(1)] results in shorter average waiting time for H agents. The theoretical gap is significant when $\lambda_H > \lambda_E$. We run numerical simulations for a variety of parameters to examine these differences (see Figure 6).

To further highlight the benefit of matching through chains, we consider the following scenario: Suppose market 1 has rates $(\lambda_{H,1}, \lambda_{E,1})$ with $\lambda_{H,1} < \lambda_{E,1}$ and is endowed with d altruistic agents and employs policy ChainMatch(d). Now consider a second market with arrival rates $(\lambda_{H,2}, \lambda_{E,2})$ that does not have any altruistic agents and therefore employs BilateralMatch(H).

Figure 6. (Color online) Comparison of the Average Waiting Time of H Agents Under the BilateralMatch(H) Policy and the ChainMatch(d) Policy with 1 and 20 Altruistic Agents



Note. The figure shows w_H , for \mathcal{B}_H , $\mathcal{C}(1)$, and $\mathcal{C}(20)$, as a function of λ_H , for fixed $\lambda_E = 5$, $T = 2 \cdot 10^6$, $p_E = 0.5$, and $p_H = 0.02$.

Further suppose $\lambda_{H,1} = \lambda_{H,2} = \lambda_H$. How many more E agents does market 2 need to attract to be able to compete with market 1 in terms of average waiting times of H agents? In the limit $p_H \rightarrow 0$, by Theorems 1 and 3, for this to happen, it is necessary that

$$\frac{\ln\left(\frac{\lambda_H}{\lambda_{E,1}(1-(1-p_E)^d)} + 1\right)}{\lambda_H} \geq \frac{\ln\left(\frac{\lambda_{E,2}}{\lambda_{E,2}-\lambda_H}\right)}{\lambda_H p_E},$$

which is equivalent to

$$\lambda_{E,2} \geq \frac{\lambda_H(\lambda_H + \lambda_{E,1}(1-(1-p_E)^d))^{p_E}}{(\lambda_H + \lambda_{E,1}(1-(1-p_E)^d))^{p_E} - (\lambda_{E,1}(1-(1-p_E)^d))^{p_E}}.$$

Note that the above condition is only a necessary condition and valid in the limit $p_H \rightarrow 0$. In the case where $p_E = 1$, Proposition 1 makes this also a sufficient condition, and it simplifies to $\lambda_{E,2} \geq \lambda_H + \lambda_{E,1}$. In Table 1, we report the numerical values for $\lambda_{E,2}$ such that in simulations $w_{H,2}^{\mathcal{B}_H} = w_{H,1}^{\mathcal{C}(d)}$.

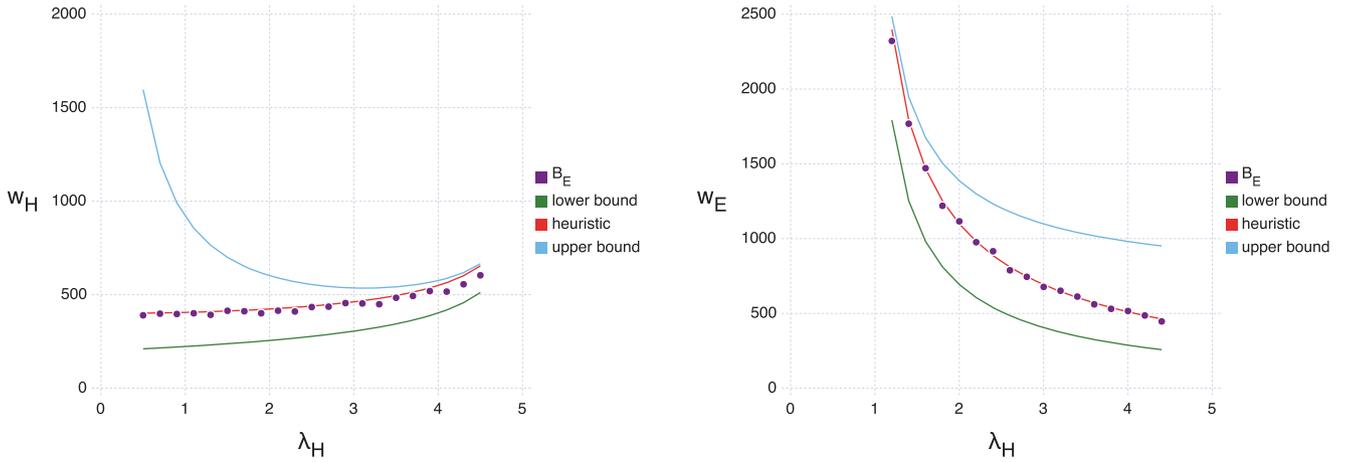
4.5. Theoretical Bounds vs. Heuristics vs. Simulation

In two cases, our theoretical results yield bounds that are not tight. However, in each of these cases, we generate a heuristic guess for the exact behavior. We plot here the simulation results, our heuristically generated guess (described later in Section 5.3 and

Table 1. $\lambda_{E,2}$ as a Function of p_E and d , for $p_H = 0.02$, $\lambda_H = 1$, and $\lambda_{E,1} = 2$, $T = 10^6$

p_E	0.1	0.3	0.5	0.9	1.0
$d = 1$	20.75	8.45	5.4	3.3	3.0
$d = 10$	27.15	10.25	6.55	3.9	3.6
$d = 50$	66.05	24.8	15.1	9.0	8.15

Figure 7. (Color online) Numerical Study of the Tightness of Theoretical Bounds on the Average Waiting Time of H Agents Under the BilateralMatch(E) Policy



Notes. The left panel shows $w_H^{\beta_E}$ as a function of λ_H , for $\lambda_E = 5$, $T = 10^5$, $p_E = 0.5$, and $p_H = 0.002$. The right panel shows $w_H^{\beta_E}$ as a function of λ_E , for $\lambda_H = 1$, $T = 10^5$, $p_E = 0.5$, and $p_H = 0.002$.

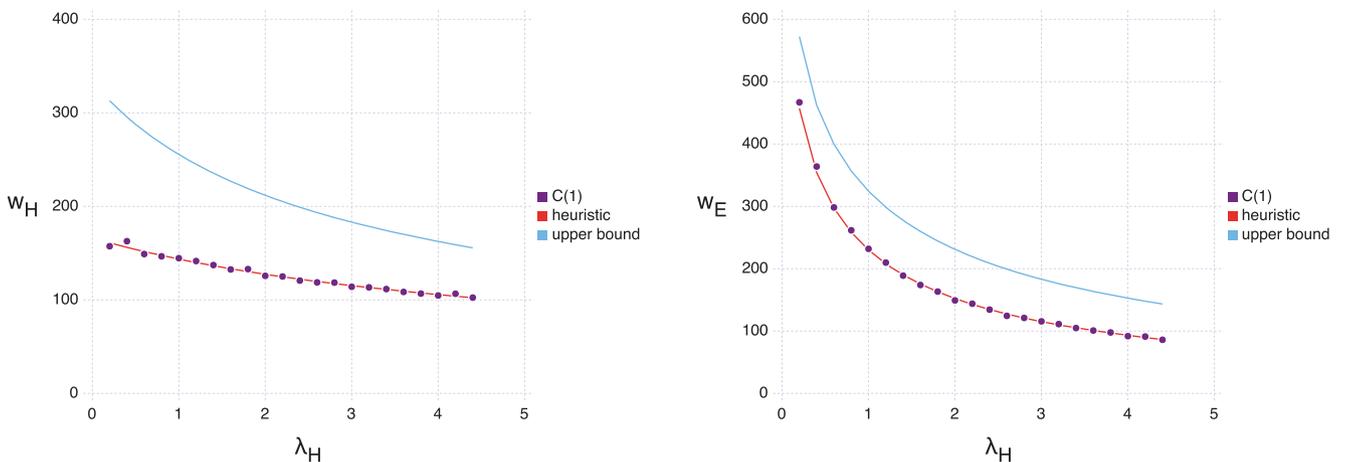
Section EC.5 of the online appendix), and the theoretical bounds for a variety of parameters. The first case is under the policy BilateralMatch(E) when $\lambda_H < \lambda_E$. Figure 7 shows that our heuristic analysis (described in Section 5.3) results in a guess of $\frac{\ln(\frac{\lambda_E + \lambda_H}{\lambda_E - \lambda_H})}{p_E p_H}$ that coincides with the simulation results. The figure further illustrates the behavior of our theoretical bounds for different parameters.

The second case is under the policy ChainMatch(d) when $p_E < 1$. Here too, Figure 8 shows that our heuristic guess $\ln\left(\frac{\lambda_H + \lambda_E}{\lambda_H(1 - (1 - p_H)^d) + \lambda_E}\right) / p_H$ (described in Section EC.5 of the online appendix) coincides with the numerical simulations.

5. Proof Ideas and Outline of Analysis

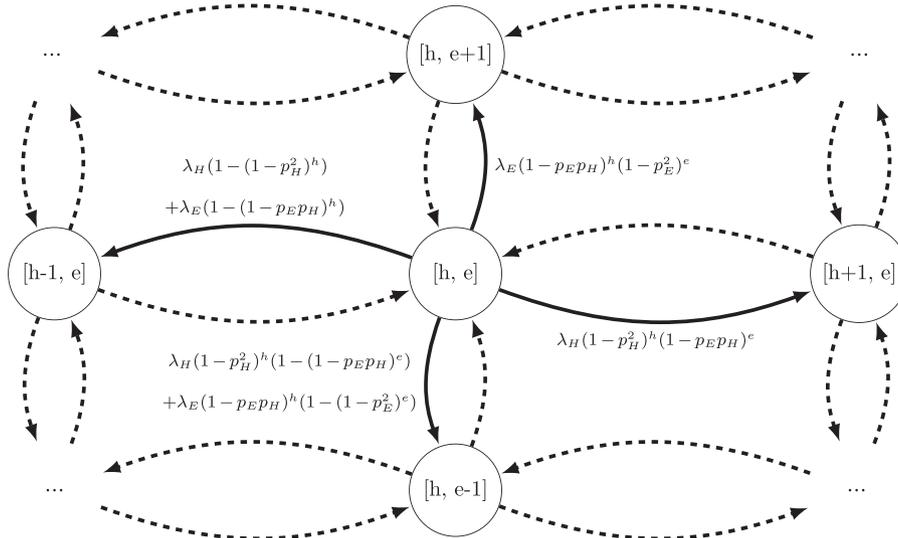
The analysis of each policy follows a similar pattern, although technically, analyzing the bilateral setting and the chain setting is very different. For bilateral policies, we first offer a heuristic that will help in guessing the value of $E[H^{\beta}]$ (which is proportional to the average waiting time) and then proceed to rigorously analyze $E[H^{\beta}]$. For the chain policy, we first couple the underlying Markov chain with a one-dimensional chain whose number of H agents serves as an upper bound on the number of H agents under the ChainMatch(d) policy. We then proceed to analyze the expected number of H agents in the coupled chain. In all three settings, the main idea is to prove

Figure 8. (Color online) Numerical Study of the Tightness of Theoretical Bounds on the Average Waiting Time of H Agents Under the ChainMatch(d) Policy



Notes. The left panel shows $w_H^{\epsilon(1)}$ as a function of λ_H , for $\lambda_E = 3$, $T = 10^5$, $p_E = 0.5$, and $p_H = 0.002$. The right panel shows $w_H^{\epsilon(1)}$ as a function of λ_E , for $\lambda_H = 3$, $T = 10^5$, $p_E = 0.5$, and $p_H = 0.002$.

Figure 9. Transition Rates (Solid Arrows) Under the CTMC Induced by the \mathcal{B}_H Policy



that $H^{\mathcal{P}}$ is concentrated around $\mathbf{E}[H^{\mathcal{P}}]$ without directly computing the steady-state distribution, and based on the exponential decay of the tail distribution when moving away from the expected value.

We often use the following notations to avoid terms that vanish in the limit $p_H \rightarrow 0$. Let $f, g : [0, 1] \rightarrow \mathbb{R}$. We write that $f = o(g)$ if $\lim_{p_H \rightarrow 0} \frac{f(p_H)}{g(p_H)} = 0$ and write that $f = O(g)$ if $\limsup_{p_H \rightarrow 0} \frac{f(p_H)}{g(p_H)} < \infty$.

5.1. The BilateralMatch(H) Policy

In this section, we analyze the policy \mathcal{B}_H , which forms myopically bilateral exchanges while prioritizing H agents. Under this policy, the evolution of the number of H and E agents in the market can be modeled by a CTMC $[H_t, E_t] \in \mathbb{N}^2$ with the following transition rates:

$$Q^{\mathcal{B}_H}([h, e], [h + 1, e]) = \lambda_H(1 - p_H^2)^h(1 - p_E p_H)^e, \quad (3a)$$

$$Q^{\mathcal{B}_H}([h, e], [h - 1, e]) = \lambda_H(1 - (1 - p_H^2)^h) + \lambda_E(1 - (1 - p_E p_H)^h), \quad (3b)$$

$$Q^{\mathcal{B}_H}([h, e], [h, e + 1]) = \lambda_E(1 - p_E p_H)^h(1 - p_E^2)^e, \quad (3c)$$

$$Q^{\mathcal{B}_H}([h, e], [h, e - 1]) = \lambda_H(1 - p_H^2)^h(1 - (1 - p_E p_H)^e) + \lambda_E(1 - p_E p_H)^h(1 - (1 - p_E^2)^e). \quad (3d)$$

The rates are computed based on the Poisson thinning property, simple counting arguments, and our assumption that edges are formed independently.

- Rightward rate [Equation (3a)]: moving from $[h, e]$ to $[h + 1, e]$ happens when an H agent arrives and cannot form a cycle with any of the existing H agents (with probability $(1 - p_H^2)^h$) or with any of the existing E agents (with probability $(1 - p_E p_H)^e$).

- Leftward rate [Equation (3b)]: moving from $[h, e]$ to $[h - 1, e]$ happens when an H agent arrives and forms a cycle with at least one of the existing H agents (probability $(1 - (1 - p_H^2)^h)$) or an E agent arrives and forms a cycle with at least one of the existing H agents (probability $(1 - (1 - p_E p_H)^h)$).

- Upward rate [Equation (3c)]: moving from $[h, e]$ to $[h, e + 1]$ happens when an E agent arrives and cannot form a cycle with any of the existing H agents (probability $(1 - p_E p_H)^h$) or with any of the existing E agents (probability $(1 - p_E^2)^e$).

- Downward rate [Equation (3d)]: moving from $[h, e]$ to $[h, e - 1]$ happens when an H agent arrives and cannot form a cycle with any of the existing H agents (probability $(1 - p_H^2)^h$) but can form a cycle with an existing E agent (probability $(1 - (1 - p_E p_H)^e)$), or an E agent arrives that cannot form a cycle with any of the existing H agents (probability $(1 - p_E p_H)^h$) but can form a cycle with an existing E agent (probability $(1 - (1 - p_E^2)^e)$).

Note that the process is a two-dimensional, continuous-time, spatially nonhomogeneous random walk. Figure 9 illustrates this random walk along with its transition rates. Also observe that the leftward and downward rates (3b) and (3d) depend on the priority assigned to H agents, and these rates will change when prioritizing E agents, as we will see in the Section 5.3. However, fixing the priority, changing the tie-breaking rule between agents of the same type (e.g., favoring agents with longer waiting times instead of selecting one at random), does not change the transition rates.

In Section EC.7 of the online appendix, we prove that the above (irreducible) CTMC is positive recurrent, and therefore reaches steady state. This is intuitive given the above transition rates and the “self-regulating” behavior of the process. The larger

the market, the larger the probability that an arriving agent can form a cycle.

The drifts are given in (3a)–(3d), and therefore,

$$\mathbf{E}\left[\lambda_H(1-p_H^2)^{H^{\mathfrak{B}_H}}(1-p_E p_H)^{E^{\mathfrak{B}_H}} - \lambda_H(1-(1-p_H^2)^{H^{\mathfrak{B}_H}}) - \lambda_E(1-(1-p_E p_H)^{E^{\mathfrak{B}_H}})\right] = 0, \quad (4a)$$

$$\mathbf{E}\left[\lambda_E(1-p_E p_H)^{H^{\mathfrak{B}_H}}(1-p_E^2)^{E^{\mathfrak{B}_H}} - \lambda_H(1-p_H^2)^{H^{\mathfrak{B}_H}} \cdot (1-(1-p_E p_H)^{E^{\mathfrak{B}_H}}) - \lambda_E(1-p_E p_H)^{H^{\mathfrak{B}_H}} \cdot (1-(1-p_E^2)^{E^{\mathfrak{B}_H}})\right] = 0. \quad (4b)$$

Assuming that the random variables $H^{\mathfrak{B}_H}$ and $E^{\mathfrak{B}_H}$ are very concentrated around their expectations, a reasonable approximation is to move the expectation inside the functions and solve the above system of nonlinear equations, and thus obtain approximations for $\mathbf{E}[H^{\mathfrak{B}_H}]$ and $\mathbf{E}[E^{\mathfrak{B}_H}]$:

- For $\lambda_H < \lambda_E$, if we plug

$$[H^{\mathfrak{B}_H}, E^{\mathfrak{B}_H}] = \left[\frac{\ln\left(\frac{\lambda_E}{\lambda_E - \lambda_H}\right)}{p_E p_H}, \frac{-\ln(2)}{\ln(1-p_E^2)} \right]$$

into (4a) and (4b), the right-hand sides will be $O(p_H)$ terms.

- For $\lambda_H > \lambda_E$, if we plug

$$[H^{\mathfrak{B}_H}, E^{\mathfrak{B}_H}] = \left[\frac{\ln\left(\frac{2\lambda_H}{\lambda_H + \lambda_E}\right)}{p_H^2}, 0 \right]$$

into (4a) and (4b), the right-hand sides will be $O(p_H^2)$ terms.

This heuristic exercise provides us the correct value of $\mathbf{E}[H^{\mathfrak{B}_H}]$ in both cases. To establish this value rigorously and prove Theorem 1, we show, in the following two propositions, that $H^{\mathfrak{B}_H}$ is highly concentrated around its mean.

Proposition 3 (Lower Bound). *Under \mathfrak{B}_H and in steady state,*

- if $\lambda_H < \lambda_E$, there exists a constant c_1 such that

$$\mathbf{P}\left[H^{\mathfrak{B}_H} \leq \frac{1}{p_E p_H} \left(\ln\left(\frac{\lambda_E}{\lambda_E - \lambda_H}\right) - c_1 p_H^{1/4} \right)\right] \leq o(p_H);$$

- if $\lambda_H > \lambda_E$, there exists a constant c_2 such that

$$\mathbf{P}\left[H^{\mathfrak{B}_H} \leq \frac{1}{p_H^2} \left(\ln\left(\frac{2\lambda_H}{\lambda_E + \lambda_H}\right) - c_2 \sqrt{p_H} \right)\right] \leq o(p_H).$$

Note that although we do have closed-form formulas for c_1, c_2 (similarly for c_3, c_4 , defined in the next proposition), these values are not informative. We refer the reader to the proofs for more details.

Proposition 4 (Upper Bound). *Under \mathfrak{B}_H and in steady state, for any $k \geq 0$,*

- if $\lambda_H < \lambda_E$, there exists a function $\gamma(p_H) = 1 - \sqrt{p_H} + o(\sqrt{p_H})$ and a constant c_3 such that

$$\mathbf{P}\left[H^{\mathfrak{B}_H} \geq \frac{1}{p_E p_H} \left(\ln\left(\frac{\lambda_E}{\lambda_E - \lambda_H}\right) + c_3 \sqrt{p_H} \right) + k\right] \leq \frac{\gamma(p_H)^k}{1 - \gamma(p_H)};$$

- if $\lambda_H > \lambda_E$, there exists a function $\gamma'(p_H) = 1 - \sqrt{p_H} + o(\sqrt{p_H})$ and a constant c_4 such that

$$\mathbf{P}\left[H^{\mathfrak{B}_H} \geq \frac{1}{p_H^2} \left(\ln\left(\frac{2\lambda_H}{\lambda_E + \lambda_H}\right) + c_4 \sqrt{p_H} \right) + k\right] \leq \frac{\gamma'(p_H)^k}{1 - \gamma'(p_H)}.$$

Note that in both cases in Proposition 4, if $k = p_H^{-3/4}$, then the right-hand side becomes $o(p_H^2)$. The proof of Theorem 1 is a straightforward application of these propositions, and the details are presented in Section EC.2.1 of the online appendix. To prove these propositions, we derive exponentially decaying bounds on tails of the steady-state distribution of $H^{\mathfrak{B}_H}$ and $E^{\mathfrak{B}_H}$. In the next subsection, we present two auxiliary lemmas that establish such bounds for a general class of two-dimensional continuous-time random walks that includes the random walk defined above. The proof of Propositions 3 and 4 amount to applying these lemmas with appropriately defined parameters. The proofs are presented in Sections EC.2.4 and EC.2.3 of the online appendix, respectively.

5.2. Concentration Bounds for a General Class of Two-Dimensional Random Walks

In the analysis of both $\text{BilateralMatch}(H)$ and $\text{BilateralMatch}(E)$ policies, we repeatedly bound the left tail or the right tail of the steady-state distribution of the number of H agents in the market. These bounds rely on certain properties of the corresponding two-dimensional continuous-time random walks, which allow us to establish exponential decay on each tail of the steady-state distribution. To avoid repeating these concentration results for each particular setting, we take a unifying approach and state the following two auxiliary lemmas that establish concentration results for a general class of two-dimensional random walks under certain conditions. These lemmas may be useful in other applications that give rise to similar random walks.

Lemma 1 (Lower Bound). *Let $[X_t, Y_t] \in \mathbb{N}^2$ be a positive recurrent continuous-time random walk with transition rate matrix Q , and let $[X, Y]$ be a corresponding random vector following its steady-state distribution. Suppose the following conditions exist:*

1. a set $S \subset \mathbb{N}$ and a constant $\epsilon > 0$ such that $\mathbf{P}[Y \notin S] \leq \epsilon$;
2. a nonincreasing function $f: \mathbb{N} \mapsto (0, \infty)$ such that $\forall y \in S, Q([x, y], [x+1, y]) \geq f(x)$;
3. a nondecreasing function $g: \mathbb{N} \mapsto (0, \infty)$ such that $\forall y \in S, Q([x, y], [x-1, y]) \leq g(x)$.

Then, for all $\rho < 1$ and $\eta \in \mathbb{N}$ such that $\frac{g(\eta+1)}{f(\eta)} < \rho$, and any $k > 0$, we have

$$\mathbf{P}[X \leq \eta - k] \leq \eta \epsilon \left(1 + \frac{1}{f(\eta) - g(\eta + 1)} \right) + \frac{\rho^k}{1 - \rho}.$$

Proof of Lemma 1. Let $\pi(x, y)$ be the joint distribution of $[X, Y]$, and let $\pi_X(x) = \sum_{y \geq 0} \pi(x, y)$ be the marginal distribution of X . In steady state, conservation of flow implies the following:

$$\begin{aligned} & \sum_{y \in S} \pi(x + 1, y) Q([x + 1, y], [x, y]) \\ & + \sum_{y \notin S} \pi(x + 1, y) Q([x + 1, y], [x, y]) \\ = & \sum_{y \in S} \pi(x, y) Q([x, y], [x + 1, y]) \\ & + \sum_{y \notin S} \pi(x, y) Q([x, y], [x + 1, y]). \end{aligned}$$

Using Conditions 2 and 3, we upper bound the left-hand side and lower bound the right-hand side, which results in

$$\begin{aligned} g(x + 1) \mathbf{P}[X = x + 1, Y \in S] + \mathbf{P}[X = x + 1, Y \notin S] \\ \geq f(x) \mathbf{P}[X = x, Y \in S]. \end{aligned}$$

Let $\pi_S(x) = \mathbf{P}[X = x, Y \in S] = \sum_{y \in S} \pi(x, y)$. Observe that by Condition 1 we have $\pi_X(x) \leq \pi_S(x) + \epsilon$. Using the fact that g is nondecreasing and f is nonincreasing, we get, for $x \leq \eta$,

$$\begin{aligned} \pi_S(x) & \leq \frac{g(x + 1)}{f(x)} \pi_S(x + 1) + \frac{\mathbf{P}[Y \notin S]}{f(x)} \\ & \leq \rho \pi_S(x + 1) + \frac{\epsilon}{f(\eta)}. \end{aligned}$$

We can subtract $\frac{\epsilon/f(\eta)}{1-\rho}$ from both sides and iterate: for all $j \geq 0$,

$$\pi_S(\eta - j) - \frac{\epsilon/f(\eta)}{1 - \rho} \leq \rho^j \left(\pi_S(\eta) - \frac{\epsilon/f(\eta)}{1 - \rho} \right) \leq \rho^j.$$

This allows us to conclude that for any $k > 0$,

$$\begin{aligned} \mathbf{P}[X \leq \eta - k] & = \sum_{i=0}^{\eta-k} \pi_X(i) \leq (\eta - k) \epsilon + \sum_{j=k}^{\eta} \pi_S(\eta - j) \\ & \leq (\eta - k) \epsilon \left(1 + \frac{1/f(\eta)}{1 - \rho} \right) + \sum_{j=k}^{\eta} \rho^j \\ & \leq \eta \epsilon \left(1 + \frac{1}{f(\eta)(1 - \rho)} \right) + \frac{\rho^k}{1 - \rho}. \end{aligned}$$

□

Lemma 2 (Upper Bound). Let $[X_t, Y_t] \in \mathbb{N}^2$ be a positive recurrent continuous-time random walk with transition rate matrix Q , and let $[X, Y]$ be a corresponding random vector following its steady-state distribution. Suppose the following conditions exist:

1. a mapping $S : \mathbb{N} \mapsto 2^{\mathbb{N}}$ and two constants $c \in \mathbb{R}^+$, $\delta \in (0, 1)$ such that $\mathbf{P}[Y \notin S(x)] \leq c\delta^x$;
2. two functions $f, g : \mathbb{N} \mapsto (0, \infty)$ such that $\forall y \in S$, $Q([x, y], [x + 1, y]) \leq f(x)$ and $Q([x, y], [x - 1, y]) \geq g(x)$.

Then, for all $\eta > 0$ and $\rho \in [\delta, 1)$ such that $\forall x \geq \eta$, $\frac{f(x)}{g(x+1)} \leq \rho$, and $\frac{\delta^x}{g(x+1)} \leq \frac{\rho^x}{g(\eta+1)}$, and for any $k > 0$, we have

$$\mathbf{P}[X \geq \eta + k] \leq \frac{\rho^k}{1 - \rho} \left(1 + c + \frac{c(k + 1)}{g(\eta + 1) - f(\eta)} \right).$$

Note that the above conditions are weaker than that of Lemma 1 (where f is nonincreasing, g is nondecreasing, and $\frac{f(\eta)}{g(\eta+1)} \leq \rho$). We will need this for the proofs of Propositions 3 and EC.2, where the corresponding function g is not monotone.

The proof of Lemma 2 follows similar arguments to that of Lemma 1 and is deferred to Section EC.1 of the online appendix.

5.3. The BilateralMatch(E) Policy

The policy \mathcal{B}_E forms myopically bilateral exchanges while prioritizing E agents. The transition rates of the underlying CTMC are as follows:

$$Q^{\mathcal{B}_E}([h, e], [h + 1, e]) = \lambda_H(1 - p_H^2)^h(1 - p_E p_H)^e, \quad (5a)$$

$$\begin{aligned} Q^{\mathcal{B}_E}([h, e], [h - 1, e]) & = \lambda_H(1 - p_E p_H)^e(1 - (1 - p_H^2)^h) \\ & + \lambda_E(1 - p_E^2)^e(1 - (1 - p_E p_H)^h), \end{aligned} \quad (5b)$$

$$Q^{\mathcal{B}_E}([h, e], [h, e + 1]) = \lambda_E(1 - p_E p_H)^h(1 - p_E^2)^e, \quad (5c)$$

$$\begin{aligned} Q^{\mathcal{B}_E}([h, e], [h, e - 1]) & = \lambda_H(1 - (1 - p_E p_H)^e) \\ & + \lambda_E(1 - (1 - p_E^2)^e). \end{aligned} \quad (5d)$$

The rates are computed similarly to those under the BilateralMatch(H) policy. Observe that prioritizing E results in different leftward and downward rates (5b) and (5d) than the corresponding rates under BilateralMatch(H). In particular, note that in the leftward rate (moving from $[h, e]$ to $[h - 1, e]$), the probability that an arriving E agent matches an existing H agent depends now on the current number of E agents. This dependency does not exist in BilateralMatch(H). This makes the analysis of BilateralMatch(E) more difficult because we need to compute tight bounds also on the number of E agents in the market. Although we are able to prove such bounds in the case $\lambda_H > \lambda_E$, we are not able to do so in the case $\lambda_H < \lambda_E$.

As before, we set the expected drifts at steady state in both dimensions to zero, resulting in the following system of equations:

$$\begin{aligned} & \mathbf{E}[\lambda_H(1 - p_H^2)^{H^{\mathcal{B}_E}}(1 - p_E p_H)^{E^{\mathcal{B}_E}} \\ & \quad - \lambda_H(1 - p_E p_H)^{E^{\mathcal{B}_E}}(1 - (1 - p_H^2)^{H^{\mathcal{B}_E}}) \\ & \quad - \lambda_E(1 - p_E^2)^{E^{\mathcal{B}_E}}(1 - (1 - p_E p_H)^{H^{\mathcal{B}_E}})] = 0, \end{aligned} \quad (6a)$$

$$\begin{aligned} & \mathbf{E}[\lambda_E(1 - p_E p_H)^{H^{\mathcal{B}_E}}(1 - p_E^2)^{E^{\mathcal{B}_E}} - \lambda_H(1 - (1 - p_E p_H)^{E^{\mathcal{B}_E}}) \\ & \quad - \lambda_E(1 - (1 - p_E^2)^{E^{\mathcal{B}_E}})] = 0. \end{aligned} \quad (6b)$$

As in the heuristic analysis for $\text{BilateralMatch}(H)$, we can obtain the following approximations for $\mathbf{E}[H^{\mathcal{B}_E}]$ and $\mathbf{E}[E^{\mathcal{B}_E}]$:

- For the case $\lambda_H < \lambda_E$, if we plug

$$\left[\frac{\ln\left(\frac{\lambda_E + \lambda_H}{\lambda_E - \lambda_H}\right)}{p_E p_H}, \frac{\ln\left(\frac{\lambda_E + \lambda_H}{2\lambda_E}\right)}{\ln(1 - p_E^2)} \right]$$

into (6a) and (6b), the right-hand sides will be $O(p_H)$ terms.

- For the case, $\lambda_H > \lambda_E$, if we plug

$$\left[\frac{\ln\left(\frac{2\lambda_H}{\lambda_H + \lambda_E}\right)}{p_H^2}, 0 \right]$$

into (6a) and (6b), the right-hand sides will be $O(p_H^2)$ terms.

As stated in Theorem 2, for the case $\lambda_H > \lambda_E$, the constant for the limit of $\frac{\mathbf{E}[H^{\mathcal{B}_E}]}{\lambda_H}$ coincides with the solution given by the above heuristic. For the case, $\lambda_H < \lambda_E$, the constant resulting from the above heuristic argument lies in between the constants of the lower and upper bounds we can prove (in Theorem 2), that is,

$$\frac{\ln\left(\frac{\lambda_E}{\lambda_E - \lambda_H}\right)}{p_E} \leq \frac{\ln\left(\frac{\lambda_E + \lambda_H}{\lambda_E - \lambda_H}\right)}{p_E} \leq \frac{\ln\left(\frac{2\lambda_E}{\lambda_E - \lambda_H}\right)}{p_E}.$$

In Figure 7, we numerically show that $\frac{\ln\left(\frac{\lambda_E + \lambda_H}{\lambda_E - \lambda_H}\right)}{p_E}$ is indeed the right constant.

The proof of the case $\lambda_H > \lambda_E$, and the lower bound when $\lambda_H < \lambda_E$ in Theorem 2, follows similar steps as that of Theorem 1, and it uses the concentration results of the lemmas stated in the previous subsection. The difficulty in closing the gap between our lower and upper bounds for the case $\lambda_H < \lambda_E$ comes from the dependency of the leftward rate on the current number of E agents [i.e., the second term in (5b)]. Our bounds on the right tail of the distribution of the number of E agents are not tight enough to result in matching lower and upper bounds. Closing this gap remains an open question. A notable difference is that in (3a) and (3b), knowing that E is bounded above by

a constant (independent of p_H) is enough to get matching upper and lower bounds (up to a vanishing term). This, however, is not the case in (5b). To prove the upper bound in the case $\lambda_H < \lambda_E$, we couple the Markov process underlying policy \mathcal{B}_E with another process in which an E agent that cannot form a match upon arrival turns into an H agent. See Section EC.3.2.

5.4. The ChainMatch(d) Policy

This section proves Theorem 3 and Proposition 1. As we could establish only an upper bound for the average waiting time when $p_E < 1$, we refer the reader to Section EC.5 of the online appendix for a heuristic analysis that leads us to guess the constant that we can numerically verify to be the correct one (see Figure 8).

Instead of directly analyzing the $\text{ChainMatch}(d)$ policy under our setting, we consider a modified setting, in which an E agent that does not match immediately upon arrival is removed from the system. We refer to this new setting under the policy $\text{ChainMatch}(d)$ as $\hat{\mathcal{C}}(d)$. Observe that $H_t^{\hat{\mathcal{C}}(d)}$ is a one-dimensional CTMC with the following transition rates:

$$Q^{\hat{\mathcal{C}}(d)}(h, h + 1) = \lambda_H(1 - p_H)^d, \quad (7a)$$

$$\begin{aligned} Q^{\hat{\mathcal{C}}(d)}(h, h - i) &= (\lambda_H(1 - (1 - p_H)^d) + \lambda_E(1 - (1 - p_E)^d)) \\ &\quad \cdot (1 - p_H)^{h-i} \prod_{j=0}^{i-1} (1 - (1 - p_H)^{h-j}), \\ &\quad i \in \{1, 2, \dots, h\}. \end{aligned} \quad (7b)$$

The first expression, (7a), corresponds to rate at which an H agent arrives but cannot be matched by a bridge agent. The second expression, (7b), corresponds to the rate at which an agent arrives, is matched by a bridge agent, and forms a chain segment of length i . Observe that the case $i = 0$ is possible and corresponds to an arriving agent that can be matched by a bridge agent but cannot continue the chain further. In that case, the CTMC does not transition, and we consider the chain segment to have length 1. In Section EC.7 of the online appendix, we show that the above CTMC reaches steady state.

We introduce some notation to simplify (7b). Set $\Lambda = \lambda_H(1 - (1 - p_H)^d) + \lambda_E(1 - (1 - p_E)^d)$, which is the rate at which a new chain segment (possibly of length 1) starts, regardless of the current state, and let S_h be the random number of agents removed from the system, starting from state h . Using the notation from Section 3.2.2, note that $S_h + 1$ corresponds to the length of the chain segment L_k for the one-dimensional Markov chain. For any $i \leq h$, we can write

$$\begin{aligned} Q^{\hat{\mathcal{C}}(d)}(h, h - i) &= \Lambda \mathbf{P}[S_h = i] \\ &= \Lambda(1 - p_H)^{h-i} \prod_{j=0}^{i-1} (1 - (1 - p_H)^{h-j}). \end{aligned} \quad (8)$$

Observe that we have

$$\mathbf{P}[S_h \geq k] = \prod_{j=0}^{k-1} (1 - (1 - p_H)^{h-j}). \quad (9)$$

The proof proceeds by showing that $\mathbf{E}[H^{\hat{\mathcal{C}}(d)}]$ serves as an upper bound for $\mathbf{E}[H^{\mathcal{C}(d)}]$ (Lemma 3) and then computing the limit of $\mathbf{E}[H^{\hat{\mathcal{C}}(d)}]$ (Proposition 5). Before that, we make the following crucial observation: the process of chain-segment formation under $\hat{\mathcal{C}}(d)$ exhibits a memoryless property; that is, for any state h and any $\tilde{i} \leq i \leq h$,

$$\begin{aligned} \mathbf{P}[S_h = i] &= (1 - p_H)^{h-i} \prod_{j=0}^{i-1} (1 - (1 - p_H)^{h-j}) \quad (10) \\ &= \prod_{j=0}^{\tilde{i}-1} (1 - (1 - p_H)^{h-j}) \left[(1 - p_H)^{(h-\tilde{i})-(i-\tilde{i})} \right. \\ &\quad \left. \cdot \prod_{j=0}^{(i-\tilde{i})-1} (1 - (1 - p_H)^{(h-\tilde{i})-j}) \right] \\ &= \mathbf{P}[S_h \geq \tilde{i}] \mathbf{P}[S_{h-\tilde{i}} = i - \tilde{i}]. \end{aligned}$$

In other words, the event of forming a chain segment of length i can be decomposed into two independent events: forming a chain segment of length at least \tilde{i} and then forming a chain segment of length $i - \tilde{i}$ starting with $h - \tilde{i}$ agents in the market. This heavily relies on the fact that chain segments proceed in a local search (one by one) fashion and the independence assumption. Indeed, the chain-segment formation in the original two-dimensional chain $\mathcal{C}(d)$ has a similar property.

We now show that $\mathbf{E}[H^{\hat{\mathcal{C}}(d)}]$ is an upper bound for $\mathbf{E}[H^{\mathcal{C}(d)}]$.

Lemma 3. *The expected number of H agents in steady state under $\hat{\mathcal{C}}(d)$ satisfies*

$$\mathbf{E}[H^{\mathcal{C}(d)}] \leq \mathbf{E}[H^{\hat{\mathcal{C}}(d)}].$$

Proof of Lemma 3. The proof is based on a coupling argument. Consider two copies of the arrival process, one under the setting of $\mathcal{C}(d)$ and one under $\hat{\mathcal{C}}(d)$. Let $[H_k^{\mathcal{C}(d)}, E_k^{\mathcal{C}(d)}]$ and $H_k^{\hat{\mathcal{C}}(d)}$ denote the embedded discrete-time Markov chain resulting from observing the two dynamic systems at arrival epochs. We prove a stronger result: at any step k , $H_k^{\mathcal{C}(d)} \leq H_k^{\hat{\mathcal{C}}(d)}$. We prove this using the following coupling:

1. Upon arrival of an H agent, we flip a biased coin with probability $(1 - p_H)^d$. If the coin flip is heads, the H agent cannot start a chain segment, and both $H_{k+1}^{\mathcal{C}(d)}$ and $H_{k+1}^{\hat{\mathcal{C}}(d)}$ increment by one. If the coin flip is tails, the H agent starts a chain segment in both systems. Suppose that $[H_k^{\mathcal{C}(d)}, E_k^{\mathcal{C}(d)}] = [h, e]$ and $H_k^{\hat{\mathcal{C}}(d)} = \tilde{h}$, and let $[L_{[h,e]}^H, L_{[h,e]}^E]$ denote the random number of H and E

agents in the chain segment formed under $\mathcal{C}(d)$ at state $[h, e]$. Similarly, let $S_{\tilde{h}}$ be the length of chain segment formed under $\hat{\mathcal{C}}(d)$ at state (\tilde{h}) . We distinguish between three cases:

a. $\tilde{h} \geq h$ and the event $\{S_{\tilde{h}} < (\tilde{h} - h)\}$ occurs. We let $[L_{[h,e]}^H, L_{[h,e]}^E]$ be realized independently of $S_{\tilde{h}}$.

b. $\tilde{h} \geq h$ and the event $\{S_{\tilde{h}} \geq (\tilde{h} - h)\}$ occurs. In this case, the memoryless property of $\hat{\mathcal{C}}(d)$ in (10) can be rewritten as $\mathbf{P}[S_{\tilde{h}} = i \mid S_{\tilde{h}} \geq (\tilde{h} - h)] = \mathbf{P}[S_h = i - (\tilde{h} - h)]$. This divides the chain-segment formation into two independent events: a subchain segment of length $(\tilde{h} - h)$ is formed, and then a subchain segment of length ξ , where ξ is a random variable drawn from the distribution of S_h . Now we focus on the chain-segment formation under $\mathcal{C}(d)$. Because H agents get a higher priority, the chain segment can be computed in steps. Starting with $[h, e]$ agents, we first look for a subchain segment L_1^H consisting of only H agents. When this chain segment cannot be continued further with only H agents, we look for an E agent to continue the chain. If this happens (with probability $(1 - (1 - p_E)^e)$), we look for a second subchain segment L_2^H of only H agents, etc. Note that the first subchain segment L_1^H also has the same distribution as S_h . We can therefore set $L_1^H = \xi$. All further subchain segments L_i^H are realized independently.

c. $\tilde{h} < h$. We let $[L_{[h,e]}^H, L_{[h,e]}^E]$ and $S_{\tilde{h}}$ be realized independently.

2. Upon arrival of an E agent, we flip a biased coin with probability $(1 - p_E)^d$. If the flip is heads, the E agent cannot start a chain segment in either system, and we have $H_{k+1}^{\mathcal{C}(d)} = H_k^{\mathcal{C}(d)}$ and $H_{k+1}^{\hat{\mathcal{C}}(d)} = H_k^{\hat{\mathcal{C}}(d)}$. On the other hand, if the flip is tails, the E agent starts a chain segment in both systems. The chain-segment formation in this case is exactly the same as the one for an H arrival.

Having the above coupling, we finish the proof by induction. The base case $k = 0$ is trivial: $H_0^{\mathcal{C}(d)} = H_0^{\hat{\mathcal{C}}(d)} = 0$. Suppose $H_k^{\mathcal{C}(d)} \leq H_k^{\hat{\mathcal{C}}(d)}$ holds for k . We show that it also holds for $k + 1$: If an H/E arrival does not start a chain segment, then, by coupling construction, $H_{k+1}^{\mathcal{C}(d)} \leq H_{k+1}^{\hat{\mathcal{C}}(d)}$. If an H arrival does start a chain segment, then we are either in Case 1a or 1b. In the former, the length of the chain segment in $\hat{\mathcal{C}}(d)$ was not even long enough to bring the number of H agents back to $H_k^{\mathcal{C}(d)} = h$; therefore, $H_{k+1}^{\mathcal{C}(d)} \leq H_{k+1}^{\hat{\mathcal{C}}(d)}$ holds. In the latter case, again by coupling construction, $L_{[h,e]}^H \geq S_h + (\tilde{h} - h)$, which implies that $H_{k+1}^{\mathcal{C}(d)} \leq H_{k+1}^{\hat{\mathcal{C}}(d)}$ holds. A similar argument holds if an E arrival starts a chain segment. \square

The next proposition computes $\mathbf{E}[H^{\hat{\mathcal{C}}(d)}]$ in the limit. Together with Lemma 3, this completes the proof of Theorem 3.

Proposition 5. Under $\hat{c}(d)$ and in steady state, the expected number of H agents satisfies

$$\lim_{p_H \rightarrow 0} p_H \mathbf{E}[H^{\hat{c}(d)}] = \ln \left(1 + \frac{\lambda_H}{\lambda_E(1 - (1 - p_E)^d)} \right).$$

Proof of Proposition 5. Let π be the steady-state probability distribution. By the conservation of flow from state h to $h + 1$, we have

$$\pi(h)\lambda_H(1 - p_H)^d = \sum_{k \geq 1} \pi(h + k) \left(\sum_{i \leq h} Q^{\hat{c}(d)}(h + k, i) \right).$$

Note that $\sum_{i \leq h} Q^{\hat{c}(d)}(h + k, i)$ is the total leftward flow starting from state $h + k$ and ending at state $i \leq h$. Using (8) and (9), we have

$$\sum_{i \leq h} Q^{\hat{c}(d)}(h + k, i) = \Lambda \mathbf{P}[S_{h+k} \geq k],$$

and therefore,

$$\pi(h)\lambda_H(1 - p_H)^d = \Lambda \sum_{k \geq 1} \pi(h + k) \mathbf{P}[S_{h+k} \geq k]. \quad (11)$$

Observe that applying definition (9), we have $\mathbf{P}[S_{h+k} \geq k] = \mathbf{P}[S_{h+k} \geq k - 1] \mathbf{P}[S_{h+1} \geq 1]$. Therefore, we can rewrite (11) as

$$\begin{aligned} \pi(h)\lambda_H(1 - p_H)^d &= \Lambda \left(\pi(h + 1) \mathbf{P}[S_{h+1} \geq 1] + \mathbf{P}[S_{h+1} \geq 1] \right. \\ &\quad \left. \cdot \sum_{k \geq 2} \pi(h + k) \mathbf{P}[S_{h+k} \geq k - 1] \right). \end{aligned} \quad (12)$$

Similarly, we write the conservation of flow from state $h + 1$ to $h + 2$:

$$\begin{aligned} \pi(h + 1)\lambda_H(1 - p_H)^d &= \Lambda \sum_{k \geq 1} \pi(h + 1 + k) \mathbf{P}[S_{h+k+1} \geq k] \\ &= \Lambda \sum_{k' \geq 2} \pi(h + k') \mathbf{P}[S_{h+k'} \geq k' - 1], \end{aligned} \quad (13)$$

where the last step follows from a change of variable $k' = k + 1$. Note that the summation in the right-hand side of (13) also appears in the second term of the right-hand side of (12). Substituting $\sum_{k' \geq 2} \pi(h + k') \mathbf{P}[S_{h+k'} \geq k' - 1]$ with $\pi(h + 1)\lambda_H(1 - p_H)^d / \Lambda$ in (12) gives that

$$\begin{aligned} \pi(h)\lambda_H(1 - p_H)^d &= \pi(h + 1) \mathbf{P}[S_{h+1} \geq 1] \\ &\quad \cdot (\Lambda + \lambda_H(1 - p_H)^d). \end{aligned} \quad (14)$$

We can now compute $\mathbf{E}[H^{\hat{c}(d)}]$ by proving an upper and lower bound separately. We use the fact that for states far enough from the expectation, the distribution decays geometrically.

We start with the upper bound. Let $\eta = \ln \left(1 + \frac{\lambda_H}{\lambda_E(1 - (1 - p_E)^d)} \right) / p_H + 1 / \sqrt{p_H}$. We know from (9) that $\mathbf{P}[S_{h+1} \geq 1] = 1 - (1 - p_H)^{h+1}$. This implies that for $h \geq \eta$,

$$\begin{aligned} \mathbf{P}[S_{h+1} \geq 1] &\geq 1 - (1 - p_H)^{\eta+1} \\ &= 1 - e^{(\eta+1) \ln(1 - p_H)} \\ &= 1 - \frac{\lambda_E(1 - (1 - p_E)^d)}{\lambda_H + \lambda_E(1 - (1 - p_E)^d)} (1 - \sqrt{p_H}) \\ &\quad + o(p_H), \end{aligned}$$

where we used the Taylor expansion $\ln(1 - x) = -x - x^2/2 - x^3/3 - \dots$

Using (14) for $h \geq \eta$, we have

$$\begin{aligned} \frac{\pi(h + 1)}{\pi(h)} &= \frac{\lambda_H(1 - p_H)^d}{\mathbf{P}[S_{h+1} \geq 1](\Lambda + \lambda_H(1 - p_H)^d)} \\ &\leq \frac{\lambda_H(1 - p_H)^d}{\lambda_H + \lambda_E(1 - (1 - p_E)^d)\sqrt{p_H}} + o(\sqrt{p_H}) \\ &= 1 - c\sqrt{p_H} + o(\sqrt{p_H}) =: \delta, \end{aligned} \quad (15)$$

where $c = \frac{\lambda_E(1 - (1 - p_E)^d)}{\lambda_H}$. Having (15), we upper bound $\mathbf{E}[H^{\hat{c}(d)}]$ as follows:

$$\begin{aligned} \mathbf{E}[H^{\hat{c}(d)}] &= \sum_{h \leq \eta + p_H^{-3/4}} \pi(h) + \sum_{h \geq \eta + p_H^{-3/4} + 1} \pi(h) \\ &\leq \eta + p_H^{-3/4} + \pi(\eta) \frac{\delta^{p_H^{-3/4} + 1}}{(1 - \delta)} \\ &= \frac{\ln \left(1 + \frac{\lambda_H}{\lambda_E(1 - (1 - p_E)^d)} \right)}{p_H} + o(1/p_H). \end{aligned}$$

Similarly, we lower bound $\mathbf{E}[H^{\hat{c}(d)}]$: Let $\hat{\eta} = \ln \left(1 + \frac{\lambda_H}{\lambda_E(1 - (1 - p_E)^d)} \right) / p_H - 1 / \sqrt{p_H}$. We can find \hat{c} such that for $h \leq \hat{\eta}$,

$$\frac{\pi(h)}{\pi(h + 1)} \leq 1 - \hat{c}\sqrt{p_H} + o(\sqrt{p_H}).$$

The above inequality combined with Markov inequality enables us to lower bound $\mathbf{E}[H^{\hat{c}(d)}]$ as follows:

$$\begin{aligned} \mathbf{E}[H^{\hat{c}(d)}] &\geq (\eta - p_H^{-3/4}) \left(1 - \sum_{h=0}^{\eta - p_H^{-3/4}} \pi(h) \right) \\ &= \frac{\ln \left(1 + \frac{\lambda_H}{\lambda_E(1 - (1 - p_E)^d)} \right)}{p_H} + o(1/p_H). \quad \square \end{aligned}$$

Finally, note in the special case $p_E = 1$, an arriving E agent is matched immediately by a bridge agent, implying that $E_t^{c(d)} = 0$ and $H_t^{c(d)} = H_t^{\hat{c}(d)}$. Consequently Proposition 5 implies the limit stated in Proposition 1.

6. Final Comments

In matching markets where monetary transfers are not allowed, exogenous thickness increases exchange opportunities (Roth 2008). Using a simple dynamic model with heterogeneous agents, we find a tight connection between market thickness and the desired matching technology; matching through chains is significantly more efficient than (simple) bilateral matching only when the market is sufficiently thin. Furthermore, increasing the arrival rate of hard-to-match agents may have, under bilateral matching, an adverse effect on such agents.

An important dynamic matching market is kidney exchange, which enables incompatible patient–donor pairs to exchange donors. Although our stylized model abstracts away from many details in this market, our findings may provide some useful insights into policy issues. When merging markets, which is an ongoing effort in various countries (see Section 2.1), or attracting different types of pairs, there may be negative effects on some pairs. This effect is well known for pairs with type O patients and non-type O donors who compete to match with scarce type O donors in the pool (Roth et al. 2007). Our findings suggest that this negative effect extends also to blood type-compatible pairs (like O–O), many of which have very highly sensitized patients.

Understanding these externalities is a key step toward aligning incentives toward cooperation between the relevant players (Ashlagi and Roth 2012). Our findings further provide some insights about trade-offs from prioritizing different types of pairs.

Next we discuss some limitations and possible extensions. One interesting challenge is to quantify the exact loss from restricting attention to myopic policies that do not wait before matching, rather than finding the optimal Markovian policy that may make some agents wait to increase matching opportunities.

As is shown in Anderson et al. (2017), one can show that our policies achieve the same scaling as the best anonymous Markovian policy (see Proposition EC.4 in Section EC.8 of the online appendix), but characterizing the best constants is an open question. Another interesting direction is to extend the model to allow departures. For example, Akbarpour et al. (2014) allow agents to depart prior to being matched and consider the match rate as the measure for efficiency. Finally, our focus has been on marketplaces, in which any pair of agents has a nonzero probability of forming a match. We found that the composition of the market crucially impacts the efficiency of the market. An interesting direction for future research would be to extend this study to two-sided marketplaces and, in particular, explore what features determine waiting times, for example, whether it is more beneficial to be on the short side or have a large ex ante match probability.

Acknowledgments

The authors thank the editors and anonymous reviewers for their valuable suggestions that helped to improve this paper.

References

- Adan I, Weiss G (2012) Exact FCFS matching rates for two infinite multitype sequences. *Oper. Res.* 60(2):475–489.
- Agarwal N, Ashlagi I, Azevedo E, Featherstone C, Karaduman O (2018) Market failure in kidney exchange. Working Paper No. 24775, National Bureau of Economic Research, Cambridge.
- Akbarpour M, Li S, Oveis Gharan S (2014) Dynamic matching market design. Working paper, Stanford University, Stanford, CA.
- Anderson R, Ashlagi I, Kanoria Y, Gamarnik D (2017) Efficient dynamic barter exchange. *Oper. Res.* 65(6):1446–1459.
- Anderson R, Ashlagi I, Gamarnik D, Rees M, Roth A, Sönmez T, Ünver M (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 A (2012) New challenges in multi-hospital kidney exchange. *Amer. Econom. Rev.* 102(3):354–359.
- 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.
- Ashlagi I, Bingaman A, Burq M, Manshadi V, Melcher M, Murphey M, Roth AE, Rees M (2017) The effect of match-run frequencies on the number of transplants and waiting times in kidney exchange. *Amer. J. Transplantation* 18(5):1177–1186.
- Ashlagi I, Roth AE (2014) Free riding and participation in large scale, multi-hospital kidney exchange. *Theoret. Econom.* 9(3):817–863.
- Baccara M, Lee S, Yariv L (2015) Optimal dynamic matching. Working paper, Olin School of Business, Washington University, St. Louis.
- Biro P, Burnapp L, Bernadette H, Hemke A, Johnson R, van de Klundert J, Manlove D (2017) *First Handbook of the COST Action CA15210: European Network for Collaboration on Kidney Exchange Programmes (ENCKEP)* (European Cooperation in Science and Technology, Brussels).
- Bloch F, Cantala D (2014) Dynamic allocation of objects to queuing agents: The discrete model Working paper, Université Paris 1 and Paris School of Economics, Paris.
- Böhmig GA, Fronek J, Slavcev A, Fischer GF, Berlakovich G, Viklicky O (2017) Czech-Austrian kidney paired donation: First European cross-border living donor kidney exchange. *Transplant Internat.* 30(6):638–639.
- Caldentey R, Kaplan EH, Weiss G (2009) FCFS infinite bipartite matching of servers and customers. *Adv. Appl. Prob.* 41(3):695–730.
- Dickerson JP, Procaccia AD, Sandholm T (2012a) Dynamic matching via weighted myopia with application to kidney exchange. *Proc. Twenty-Sixth AAAI Conf. Artificial Intelligence (AAAI, Toronto)*, 1340–1346.
- Dickerson JP, Procaccia AD, Sandholm T (2012b) Optimizing kidney exchange with transplant chains: Theory and reality. *Proc. 11th Internat. Conf. Autonomous Agents Multiagent Systems* (International Foundation for Autonomous Agents and Multiagent Systems, Valencia, Spain), 808–815.
- Ding Y, Ge D, He S, Ryan CT (2015) A non-asymptotic approach to analyzing kidney exchange graphs. *Proc. 16th ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 257–258.
- Doval L (2014) A theory of stability in dynamic matching markets. Technical report, California Institute of Technology, Pasadena.
- Feldman J, Mehta A, Mirrokni VS, Muthukrishnan S (2009) Online stochastic matching: Beating 1-1/e. *Proc. 50th Annual IEEE Sympos. Foundations Comput. Sci.* (Institute of Electrical and Electronics Engineers, Piscataway, NJ), 117–126.

- Ferrari P, Weimar W, Johnson RJ, Lim W, Tinckam K (2014) Kidney paired donation: Principles, protocols and programs. *Nephrology Dialysis Transplantation* 30(8):1276–1285.
- Ferrari P, Fidler S, Holdsworth R, Woodroffe C, Tassone G, Watson N, Cantwell L, et al. (2012) High transplant rates of highly sensitized recipients with virtual crossmatching in kidney paired donation. *Transplantation* 94(7):744–749.
- Fershtman D, Pavan A (2015) Re-matching, experimentation, and cross-subsidization. Working paper, Northwestern University, Evanston, IL.
- Glorie KM, van de Klundert JJ, Wagelmans AP (2014) Kidney exchange with long chains: An efficient pricing algorithm for clearing barter exchanges with branch-and-price. *Manufacturing Service Oper. Management* 16(4):498–512.
- Goel G, Mehta A (2008) Online budgeted matching in random input models with applications to Adwords. *Proc. 19th Annual ACM-SIAM Sympos. Discrete Algorithms* (Society for Industrial and Applied Mathematics, San Francisco), 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.
- Johnson R, Allen J, Fuggle S, Bradley JA, Rudge C (2008) Early experience of paired living kidney donation in the United Kingdom. *Transplantation* 86(12):1672–1677.
- Kadam S, Kotowski M (2018) Multiperiod matching. *Internat. Econom. Rev.* 59(4):1927–1947.
- Karp RM, Vazirani Uv, Vazirani VV (1990) An optimal algorithm for on-line bipartite matching. *Proc. 22nd Annual ACM Sympos. Theory Comput.* (ACM, New York), 352–358.
- Leshno J (2014) Dynamic matching in overloaded waiting lists. Working paper, University of Chicago Booth School of Business, Chicago.
- Malik S, Cole E (2014) Foundations and principles of the Canadian living donor paired exchange program. *Canadian J. Kidney Health Disease* 1:6.
- Manshadi VH, Oveis-Gharan S, Saberi A (2011) Online stochastic matching: Online actions based on offline statistics. *Proc. 22nd Annual ACM-SIAM Sympos. Discrete Algorithms* (ACM, San Francisco), 1285–1294.
- Nikzad A, Akbarpour M, Rees MA, Roth AE (2017) Financing transplants' costs of the poor: A dynamic model of global kidney exchange. Working paper, Stanford University, Stanford, CA.
- National Kidney Registry (2017) Paired exchange results quarterly report. Technical report, National Kidney Registry, Babylon, NY.
- Rees MA, Kopke JE, Pelletier RP, Segev DL, Rutter ME, Fabrega AJ, Rogers J, et al. (2009) A non-simultaneous extended altruistic donor chain. *New England J. Medicine* 360(11):1096–1101.
- Roth AE (2008) What have we learned from market design? *Econom. J.* 118(527):285–310.
- Roth AE, Sönmez T, Ünver MU (2004) Kidney exchange. *Quart. J. Econom.* 119(2):457–488.
- 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.
- Roth AE, Sönmez T, Ünver MU, Delmonico FL, Saidman SL (2006) Utilizing list exchange and nondirected donation through chain kidney paired donations. *Amer. J. Transplantation* 6(11):2694–2705.
- Siegel-Itzkovich J (2017) Israel and Cyprus sign accord for sharing kidneys from live donors. *Jerusalem Post* (June 20), <http://www.jpost.com/Israel-News/Israel-and-Cyprus-sign-accord-for-sharing-kidneys-from-live-donors-497422>.
- Sönmez T, Ünver MU, Yenmez MB (2017) Incentivized kidney exchange. Technical report, Department of Economics, Boston College, Boston.
- Ünver MU (2010) Dynamic kidney exchange. *Rev. Econom. Stud.* 77(1):372–414.

Itai Ashlagi is an assistant professor in the Department of Management Science and Engineering at Stanford University. He is interested in game theory and the design and analysis of marketplaces, and especially in matching markets. His work influenced the practice of kidney exchange, for which he has become a Franz Edelman laureate. He was the recipient of an Outstanding Paper Award at the 2009 ACM Conference on Electronic Commerce.

Maximilien Burq is a PhD student at the Operations Research Center at the Massachusetts Institute of Technology. His research interests include online learning, dynamic optimization, and market design with applications to health-care (kidney exchange) and the sharing economy (carpooling).

Patrick Jaillet is the Dugald C. Jackson Professor in the Department of Electrical Engineering and Computer Science and a member of the Laboratory for Information and Decision Systems at the Massachusetts Institute of Technology. He is also codirector of the MIT Operations Research Center. His research interests include online learning and optimization with applications to the internet economy, transportation and logistics, market analytics, and machine learning. He is a fellow of INFORMS.

Vahideh Manshadi is an assistant professor of operations at the Yale School of Management, with secondary appointment in the Department of Statistics and Data Science. Her research interests include sequential and dynamic decision making with applications to the design of online and matching marketplaces as well as information platforms and networks.