

On Matching and Thickness in Heterogeneous Dynamic Markets*

Itai Ashlagi, Maximilien Burq, Patrick Jaillet[†], Vahideh Manshadi

May 7, 2018

Abstract

We study dynamic matching in an infinite-horizon stochastic market. While 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.

Further, somewhat surprisingly, we find that in a heterogeneous market and under bilateral matching, increasing the arrival rate of hard-to-match agents has a non-monotone 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.

*First Draft: Feb 2015.

[†]Jaillet acknowledges the research support of the National Science Foundation grant OR-1029603 and the Office of Naval Research grants N00014-12-1-0033 and N00014-15-1-2083.

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 while 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, while kidney exchanges were first conducted through bilateral exchanges (2-way cycles) (Roth et al. (2005)), multi-hospital 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. Moreover, a similar conclusion is also arrived in theoretical work Anderson et al. (2017).¹ While they consider a stylized model with homogeneous agents, their result can

¹See Subsection 1.1 for further details.

be generalized to our heterogeneous model.² Second, the literature on dynamic matching in sparse environments has focused on homogeneous agents (Anderson et al. (2017); Akbarpour et al. (2014)). 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.³ While 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 (2-way cycles), and chains. Our main findings are the following. First, we find that market thickness plays a crucial role on 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, 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.⁴ However, when they are in majority,

²This is not the focus of this paper, but for completeness, we show this in Appendix H.

³More precisely, we focus on the average waiting time of H agents, because the waiting time of E agents is negligible compared to that of H agents. For a more detailed discussion, see Section 2.

⁴To be precise, theoretically we are only able to prove that prioritizing H agents leads to shorter or equal waiting

such prioritization has no significant impact.

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 $p_H \rightarrow 0$.

Bilateral matching. Two myopic policies are considered for bilateral matching, differing in the type of agents they prioritize; While 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.⁵ 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 non-monotone 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, (ii) merging two marketplaces with different compositions, i.e. different ratios between the two arrival rates, may not be beneficial for both.

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

times (of H agents), however, numerically we confirm that such prioritization indeed leads to strictly shorter waiting times.

⁵ In reality, agents may leave the market without being matched due to various reasons such as finding outside options. 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.

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. 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 non-trivial 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.⁶ We elaborate and provide intuition for other results

⁶A similar effect happens in kidney exchange where O-A 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

throughout the 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 2-dimensional Markov chains. For bilateral matching policies, we directly analyze the underlying 2-dimensional spatially non-homogeneous 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 2-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 1-dimensional process where no E agent joins the market. Analysis of the resulting 1-dimensional Markov chain presents new challenges as transitions between non-neighboring states happen due to 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.

1.1 Related work

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

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.⁷ 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

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.

⁷The findings by Ünver (2010) thus provide a rationale for the static large market results (see, e.g. Roth et al. (2007)).

three settings of exchanges, 2-ways, 2 and 3-ways, and chains, and find that moving from 2-ways or 3-ways to chains significantly reduces the average waiting time.⁸ 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.⁹

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.¹⁰ Anderson et al. (2017) consider a homogeneous model without departures (similar to our model with $\lambda_E = 0$) and finds that there is little benefit from waiting before matching under both matching technologies of short cycles and chains.¹¹ Akbarpour et al. (2014) consider a homogeneous model with departures and finds 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.¹² This paper builds on these findings, and only analyzes myopic policies that search for a match upon arrival of a new agent.

We elaborate on the relation to Anderson et al. (2017), which is closest to our paper. Studying myopic policies under a homogeneous setting resulted valuable insights. Some insights, however, do not carry over to heterogeneous settings like kidney exchange (See Subsection 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-

⁸See also Dickerson et al. (2012b) that demonstrate the benefit of chains using simulations in dynamic kidney exchange pools.

⁹Ding et al. (2015) study a similar two-type model in a static setting and quantifies the effectiveness of matching through chains taking a novel random walk approach.

¹⁰See also Gurvich et al. (2014), who study a similar compatibility-based inventory control model.

¹¹The waiting-time scales with the same factor with or without waiting before matching.

¹²Non-myopic policies have also been studied, for example Dickerson et al. (2012a) study forward-looking policies by casting the dynamic matching problem as a high-dimensional dynamic program, and develop a heuristic to overcome the curse of dimensionality.

segments. Our model predicts shorter chain-segments, which fits better empirical evidence (chain-segments typically consist of only a few pairs). Further, 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; effect of such prioritization cannot be studied in a homogeneous model. Overall it is natural and important to study richer models in order 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 (Baccara et al. (2015); Fershtman and Pavan (2015); Doval (2014); Kadam and Kotowski (2014)) since some waiting can improve the quality of matches.¹³

Our work is also related to the problem of matching multi-class customers to multi-class servers studied in queueing literature (e.g., Caldentey et al. (2009); Adan and Weiss (2012)). In our model, an agent can be thought as a pair of customer-server, 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 study 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)).

1.2 Organization

In Section 2 we introduce the model, policies, and the underlying stochastic processes. In Subsection 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,

¹³See also related results in queueing models Leshno (2014); Bloch and Cantala (2014).

we only include proofs of selected results in the main text. The detailed proofs of the rest of the statements are deferred to clearly marked appendices.

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 by 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 $p_H \rightarrow 0$, while 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*, i.e., 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).¹⁴ 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 by 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

¹⁴Having altruistic agents is an intrinsic property of the market in the sense that some markets do not have access to such agents.

policy is a mapping from the history of exchanges and the state of the marketplace to a set of feasible exchanges involving non-overlapping 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 since 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.¹⁵ 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 2-dimensional CTMC where states correspond to the number of waiting H and E agents, and which 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 ; due to their myopic behavior there are no 2-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 due to 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 since i either leaves the market or becomes a bridge agent, in which case it will never match to j . Since both the edge formation and the matching policies do not depend on agents' identities (rather only on their types) we can merely

¹⁵Our local search chain-segment formation process bears similarity to Phase 1 of the two-phase clearing procedure of Ding et al. (2015).

keep track of the number of agents of each type.

In the remainder of the 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) CTMC's reach steady-state. A rigorous statement and proof is provided in Appendix G. 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, i.e., 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

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

Since in our model $p_H \rightarrow 0$ while p_E is kept constant, and all policies are myopic, one would expect that $w_E^{\mathcal{P}}$ is negligible compared to $w_H^{\mathcal{P}}$. We verify this claim below using numerical simulations and analytical proofs (see Figure 3 and Lemmas 4 and 5). Therefore we 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 set 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

Background. There is a large shortage of kidneys for transplants¹⁶ 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);

¹⁶As for 2017, the average waiting time is between 3-5 years in the U.S.

Nikzad et al. (2017).

Exchanges are conducted through cycles or chains.¹⁷ Typically pairs do not give a kidney prior to receiving one. This creates logistical barriers requiring cycles to be limited to 2 or 3 pairs. Chains, however, can be organized non-simultaneously, 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 U.S., 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).

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) plots the PRA distributions of patients enrolled at the NKR, APD and MSA. Most patients are either highly sensitized (PRA above 95) or low sensitized (below 5 PRA). Note that blood-type compatibility is not incorporated in this aggregate PRA distribution. Figure 1(right) provides the same distributions 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 highly sensitized patients than low sensitized ones.

Percentage of high PRA patients also varies across programs outside the U.S. In Australia, 42% of registered candidates have a PRA greater than 90 (Ferrari et al. (2012)) and in the UK, 46% of patients have a PRA greater than 85 (Johnson et al. (2008)), while in Canada, only 36% of pairs 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.

Similar to PRA distribution, the pool compositions also vary with respect to blood-type distri-

¹⁷See Sönmez et al. (2017) for a detailed description of kidney exchange.




Figure 1: PRA distributions of patients enrolled at NKR (1/2012-12/2014), APD (1/2007-8/2016), and MSA (7/2013-2/2017). Left: all patients. Right: patients belonging to blood-type compatible pairs.

butions of patient-donor pairs. Ashlagi et al. (2017) report that O-O pairs make 26% at the MSA pool but only 20.81% of the APD pool; the percentage 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, MSA and APD had an enrollment rate of roughly 180 pairs per year, while 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 at the MSA, APD and NKR, respectively.

Matching. While more than 80% of the transplants at the NKR and the APD have been conducted through chains (Anderson et al. (2015); NKR (2017)), some platforms (such as MSA, Belgium, Czech Republic) match their pairs mostly through cycles due to short access to altruistic donors. In countries like France, Poland and Portugal, chains are infeasible since altruistic donations are not permitted (Biro et al. (2017)).

Exchange platforms in the U.S. adopt typically myopic-like matching policies that periodically search for matches.^{18,19} The APD, MSA, and NKR search for exchanges on a daily basis and UNOS searches for exchanges bi-weekly.²⁰ However, some countries, such as Canada, United Kingdom,

¹⁸Based on personal communication with numerous platforms.

¹⁹Generally speaking, a myopic policy is one that upon matching does not explicitly account for possible future matches.

²⁰There 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.

the Netherlands and Australia, search for exchanges every 3 or 4 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); NKR (2017))). We note, however, that MSA and NKR assign high priority to compatible pairs, which are very easy-to-match.²¹ Platforms typically have multiple desiderata. However, implicit first order related goals are to reduce waiting times and facilitate many transplants (NKR (2017)).

Policy. Various challenges arise due to variation across kidney exchange pools with respect to their compositions and even operational issues: What priorities to assign to different types of patients? What is the impact of attracting more easy-to-match pairs and even compatible pairs?²² How important is it to incorporate chains and attract altruistic donors?

There are also several initiatives to merge kidney exchange platforms in order to increase efficiency and matching opportunities for highly sensitized patients (for example, see Böhmig et al. (2017) for merging the Austrian and the Czech Republic programs, Siegel-Itzkovich (2017) for Israel and Cyprus; further, Nikzad et al. (2017) look at augmenting national programs through global kidney exchange.). It is natural to study what is the impact of merging programs on different types of patients.

This paper does not intend to model the details in 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

²¹Such pairs could choose to go through a direct transplant if they are not matched quickly.

²²See for example Agarwal et al. (2017) and Sönmez et al. (2017) for incentive schemes towards thickening the pool with such pairs.

under the three policies.

3.1 Bilateral matching

This section considers the setting, in which agents match only through bilateral exchanges, i.e. through 2-way cycles.

Theorem 1. *Under the $BilateralMatch(H)$ policy and in steady-state, the average waiting time $w_H^{\mathcal{B}H}$ satisfies the following.*

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

Theorem 1 provides not only the scaling laws on $w_H^{\mathcal{B}H}$ 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 $w_H^{\mathcal{B}H}$ 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 $w_H^{\mathcal{B}H}$ 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*

$$(x + 1) \ln(2 - 2/(x + 1)) = 1. \quad (2)$$

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 $w_H^{\mathcal{B}H}$ is which type of agents 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 to occur. 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 . So a non-negligible 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 to match with the next E agent. When $\lambda_H > \lambda_E$, there is a non-monotone 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 since 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 in Anderson et al. (2017)).

Finally, we comment on the impact of p_E on the waiting time. When $\lambda_H < \lambda_E$, $w_H^{B_H}$ is decreasing in p_E . On the other hand, when $\lambda_H > \lambda_E$, $w_H^{B_H}$ 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 2-ways between H agents, which is independent of p_E .

The proof of Theorem 1 amounts to analyzing the underlying 2-dimensional continuous-time spatially non-homogeneous random walk. The description of the random walk is presented in Subsection 5.1 (Figure 9), 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 2-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 Subsection 5.1 with details deferred to Appendix B). 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 Subsection 5.2, we state two auxiliary lemmas on concentration bounds for a general class of 2-dimensional random walks. The corollaries follow from basic analysis of the corresponding constants (as a function of λ_H). Both corollaries are proved in Appendix B.2.

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 equal when prioritizing E agents rather than H agents (numerical simulations presented in Subsection 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.²³

The proof of Theorem 2 also requires analysis of the underlying 2-dimensional continuous-time spatially non-homogeneous random walk, and, in most parts, follows a similar structure to the proof of Theorem 1. A detailed description of the random walk is presented in Subsection 5.3. The proof of the upper and lower bounds is presented in Appendix C, 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

²³We note that for neither policy $BilateralMatch(H)$ nor $BilateralMatch(E)$, we are able to derive the exact waiting time behavior when $\lambda_E = \lambda_H$. However, our simulation results, presented in Appendix J, suggest that under both policies, the waiting time scales with $1/p_H$.

into an H agent.²⁴ In Subsection 5.3, we also provide a heuristic argument that leads us to guess that the exact limit is $\frac{\ln\left(\frac{\lambda_E + \lambda_H}{\lambda_E - \lambda_H}\right)}{p_E p_H}$ (See Figure 7 in Subsection 4.5).

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 $w_H^{C(d)}$ satisfies*

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

The above theorem establishes an upper-bound on $w_H^{C(d)}$ that scales with $1/p_H$. In Appendix I, we also establish a lower-bound on $w_H^{C(d)}$ 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 $d \geq 1$ be a constant (independent of p_H). Then*

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

Consequently, $\lim_{p_H \rightarrow 0} p_H w_H^{C(d)}$ 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 $w_H^{C(d)}$. 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

²⁴In Subsection 5.3 we provide a rough intuition on why we cannot close the gap between our upper and lower bounds on $w_H^{B_E}$ for the regime $\lambda_H < \lambda_E$.

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.

For $p_E < 1$, we give a heuristic argument (in Appendix E), in which we analyze a related 3-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 $\mathcal{C}(d)$ 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 Subsection 4.5 (see Figure 8).

The heuristic argument in Appendix E and simulation results of Figure 8 both suggest that $\lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{C}(d)}(p_E) = \lim_{p_H \rightarrow 0} p_H w_H^{\mathcal{C}(d)}(p_E = 1)$.²⁵ 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 to a term of order $\Theta(1/p_H)$, (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 and Figure 4; in particular, note that Figure 4 Bottom right illustrates that $w_H^{\mathcal{C}(d)}(p_E)$ does not vary with p_E).

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 sub-segment that includes a sequence of H agents, i.e., it has the form $E - H - \dots - H$.²⁶ Denote the expected number of H agents in a sub-segment by $\mathbf{E}[\Sigma_E]$. Observe that with high probability an existing H agent is matched through such a sub-segment, because the probability that an arriving H agent starts a sub-segment of the form $H - H - \dots - H$ is vanishing. Further, 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 sub-segment. 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 sub-segment. As the result of this balance equation, we have $\mathbf{E}[\Sigma_E] = \lambda_H / \lambda_E + o(1)$, i.e.,

²⁵With a slight abuse of notation, we denoted the expected waiting time of H agents under $\mathcal{C}(d)$ policy with parameter p_E by $w_H^{\mathcal{C}(d)}(p_E)$.

²⁶Note that when $p_E < 1$, a chain-segment can consist of multiple sub-segments that are initiated by some of the E agents waiting in the market.

$\mathbf{E}[\Sigma_E]$ 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 sub-segments: when $p_E = 1$, an E agent forms a sub-segment 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 sub-segment (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 sub-segment that starts with the E agent. With high probability, X remains a constant as p_H approaches 0 (this follows from Lemma 7 in Appendix I, 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 Appendix F, we provide further details for the above argument and build on this intuition to provide another heuristic argument to show that $\lim_{p_H \rightarrow 0} p_H w_H^{C(d)}(p_E) = \lim_{p_H \rightarrow 0} p_H w_H^{C(d)}(p_E = 1)$.

Finally we comment on the chain-segment formation process; *ChainMatch*(d) policy forms chain-segments employing a local search process and indeed our analysis relies on such chain-segment formation process. This begs the question of how much the waiting time improves if we employed a global search (that searches for the longest possible chain-segment). A precise comparison is beyond the scope of our work, however, we make the following remarks: (1) In Figure 5 of Subsection 4.3, we numerically study this questions, and we see that advancing chains locally results in a small loss in comparison to 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 9 in Appendix H) 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, i.e., 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. In order

to bound $w_H^{\mathcal{C}(d)}$, we couple the underlying Markov chain with a 1-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 1-dimensional chain is that chain-segment formation exhibits a memoryless property.²⁷ This is due to 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 5.4. Finally, we note that for the special case $p_E = 1$, the original CTMC is a 1-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 only scales as $1/p_H$. The following corollary (proven in appendix D) 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 Subsection 4.4 we further compare $BilateralMatch(H)$ to $ChainMatch(d)$ in order 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 length of chain-segments formed under the $ChainMatch(d)$ policy.

While we focus on the average waiting time to measure efficiency, length of chain-segments also play a significant role on the operational efficiency of the market. In kidney exchange for example, executing a chain-segment takes time and bears the risk of match failures.²⁸ 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.²⁹

²⁷This is different from the Markov property of the overall CTMC under $\mathcal{C}(d)$.

²⁸In this stylized model, we abstract away from both of these effects.

²⁹Note that every 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,*

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

The proof is presented in Appendix D. 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 Subsection 4.1 we look at how merging markets with different compositions affect each market. Subsection 4.2 explores the impact of giving priorities when using the bilateral matching policy. Subsection 4.3 presents comparative statics for chain matching when $p_E < 1$, and subsection 4.4 highlights the advantage of having chains. Finally, Subsection 4.5 compares our theoretical bounds (for cases for which we do not have matching upper and lower bounds) to heuristics guesses and simulations.

All simulations in this section are conducted by first computing the average number of agents in the market; then applying Little’s law (1). In order 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 T the number of arrivals (not counting the d initial altruistic agents in the case of $\mathcal{C}(d)$). In order 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 while 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.³⁰

4.2 Impact of priorities in bilateral matching

We compare here the average waiting time of H agents under the *BilateralMatch(H)* and *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 $w_H^{\mathcal{B}_E} = 534$ while $w_H^{\mathcal{B}_H} = 388$. 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

³⁰We note that the constants computed in Theorem 1 allow us to determine whether market one is better off or worse off for any $(\lambda_{H,2}, \lambda_{E,2})$, and to compute the boundary separating the two regions, in the limit $p_H \rightarrow 0$.

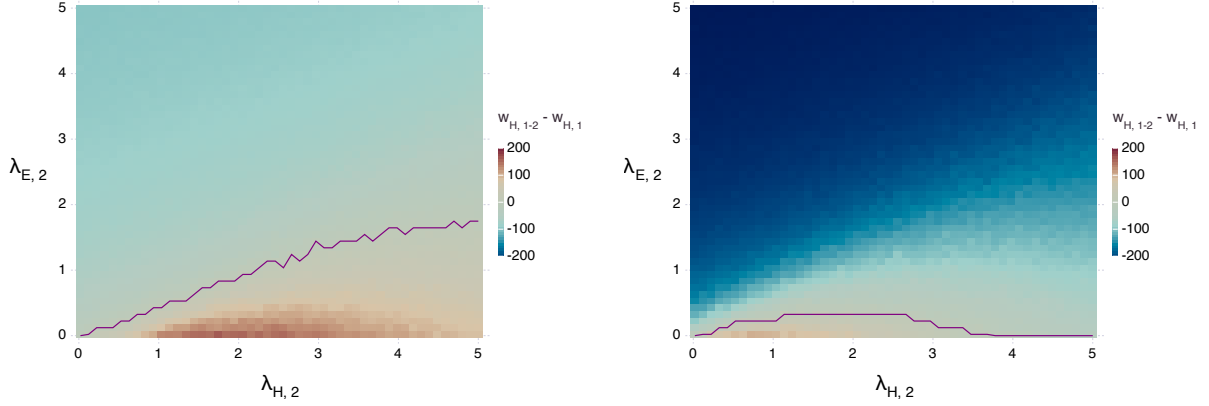


Figure 2: Change in the waiting time for H agents of the first market: $w_{H,1-2} - w_{H,1}$, as a function of $(\lambda_{H,2}, \lambda_{E,2})$, for $p_E = 0.5$, $p_H = 0.02$, after $T = 10^5$ iterations. Left subplot corresponds to $\lambda_{H,1} = 1$, $\lambda_{E,1} = 1.3$, right subplot corresponds 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).

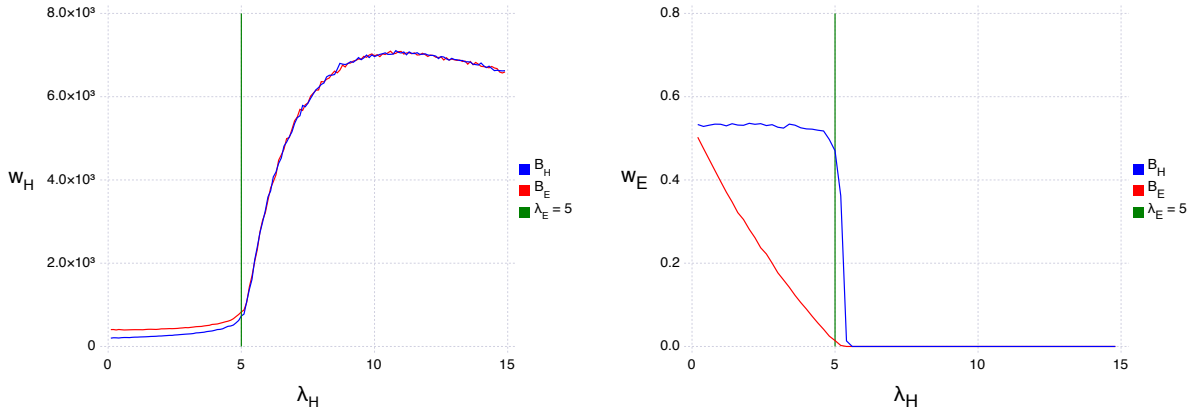


Figure 3: Comparison of w_H (left) and w_E (right) for \mathcal{B}_H , \mathcal{B}_E , as a function of λ_H , for a fixed $\lambda_E = 5$, $T = 2 \cdot 10^6$, $p_E = 0.5$ and $p_H = 0.002$.

on the composition of the market. Further, our qualitative insights can be useful in understanding the tradeoffs that may arise in markets where easy-to-match agents have outside options. For example, when $\lambda_H > \lambda_E$, there is no tradeoff 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 $ChainMatch(d)$ to numerically explore the effects that varying λ_E , λ_H , and d can have on $w_H^{C(d)}$. We find that $w_H^{C(d)}$ decreases as the arrival rate of either types 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 . Further, as p_H decreases, the impact of d vanishes (Figure 4 bottom left). Recall from Proposition 1 that when $p_E = 1$, $w_H^{C(d)}$ (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, Figure 4 top left, top right, and bottom left show numerically (i)-(iii), respectively, when $p_E = 0.5$.

Further, the bottom right plot in figure 4 plots $w_H^{C(d)}$ when p_E ranges from 0.1 to 1, and it shows that $w_H^{C(d)}$ does not vary with p_E (we refer the reader to Appendix F, for a heuristic argument on why - in the limit - $w_H^{C(d)}$ does not vary with p_E and d).

Next 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 to 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 to 15%.

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 $(\lambda_{H,2}, \lambda_{E,2})$

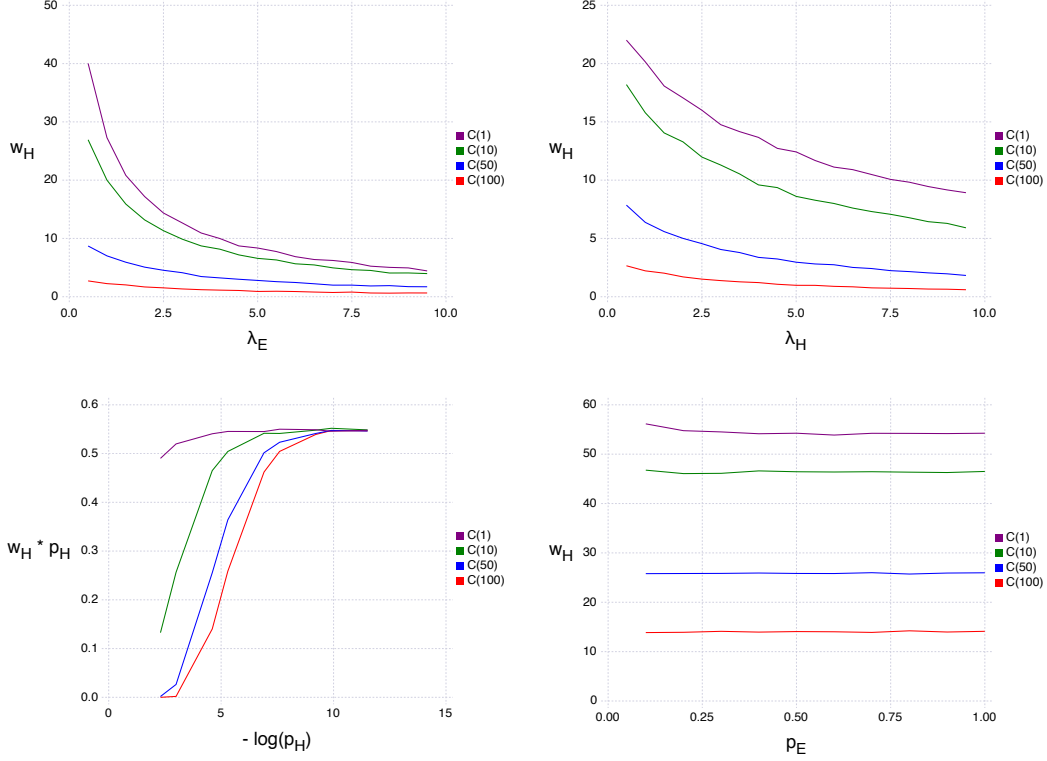


Figure 4: Top left: $w_H^{C(d)}$ as a function of λ_E , for varying values of d , for fixed $\lambda_H = 2$, $T = 10^5$, $p_E = 0.5$, $p_H = 0.02$. Top right: $w_H^{C(d)}$ as a function of λ_H , for varying values of d , for fixed $\lambda_E = 2$, $T = 10^5$, $p_E = 0.5$, $p_H = 0.02$. Bottom left: Normalized waiting times (i.e., $p_H w_H^{C(d)}$) in the case of chains as a function of $-\log(p_H)$, for varying values of d , for fixed $\lambda_H = 2$, $\lambda_E = 1$, $T = 10^5$, $p_E = 0.5$. Bottom right: $w_H^{C(d)}$ as a function of p_E , for different values of d , for fixed $\lambda_E = 1$, $\lambda_H = 2$, $T = 10^6$ and $p_H = 0.02$.

that does not have any altruistic agents and therefore employs $BilateralMatch(H)$. 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 term of average waiting times of H agents? In the limit $p_H \rightarrow 0$, by Theorems 1 and 3, for this to happens it is necessary that:

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




Figure 5: Waiting times w_H for chains conducted with local-search ($\mathcal{C}(d)$) and *Max-Chains* as a function of λ_H , for fixed $\lambda_E = 2$, $T = 10^5$, $p_E = 0.5$, $p_H = 0.002$.




Figure 6: Comparison of w_H , for \mathcal{B}_H , $\mathcal{C}(1)$, and $\mathcal{C}(20)$, as a function of λ_H , for a fixed $\lambda_E = 5$, $T = 2 * 10^6$, $p_E = 0.5$ and $p_H = 0.02$.

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 $w_{H,2}^{\mathcal{B}_H} = w_{H,1}^{\mathcal{C}(d)}$.

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

4.5 Theoretical bounds vs heuristics vs. simulation

In two cases, our theoretical results yield bounds which are not tight. However, in each of these cases we generate a heuristic guess for the exact behaviour. We plot here the simulation results, our heuristically generated guess (described later in Section 5.3 and Appendix E) 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 $\frac{\ln\left(\frac{\lambda_E + \lambda_H}{\lambda_E - \lambda_H}\right)}{p_E p_H}$ that coincides with the simulation results. The figure further illustrates the behavior of our theoretical bounds for different parameters.




Figure 7: Left: $w_H^{B_E}$ as a function of λ_H , for $\lambda_E = 5$, $T = 10^5$, $p_E = 0.5$, $p_H = 0.002$. Right: $w_H^{B_E}$ as a function of λ_E , for $\lambda_H = 1$, $T = 10^5$, $p_E = 0.5$, $p_H = 0.002$.

The second case is under the policy *ChainMatch(d)* when $p_E < 1$. Here too, Figure 8 shows that 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 Appendix E) coincides with the numerical simulations.




Figure 8: Left: $w_H^{\mathcal{C}(1)}$ as a function of λ_H , for $\lambda_E = 3$, $T = 10^5$, $p_E = 0.5$, $p_H = 0.002$. Right: $w_H^{\mathcal{C}(1)}$ as a function of λ_E , for $\lambda_H = 3$, $T = 10^5$, $p_E = 0.5$, $p_H = 0.002$.

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 are very different. For bilateral policies, we first offer a heuristic that will help guessing the value of $\mathbf{E}[H^{\mathcal{P}}]$ (which is proportional to the average waiting time) and then proceed to rigorously analyze $\mathbf{E}[H^{\mathcal{P}}]$. For the chain policy, we first couple the underlying Markov chain with a 1-dimensional chain whose number of H agents serve as an upper-bound on the number of H agents under $ChainMatch(d)$ 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 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.

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

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

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

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

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

- Rightward rate (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 $(1 - p_H^2)^h$) nor with any of the existing E agents (with probability $(1 - p_{EPH})^e$).
- Leftward rate (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 $1 - (1 - p_{EPH})^h$).
- Upward rate (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 $(1 - p_{EPH})^h$) nor with any of the existing E agents (probability $(1 - p_E^2)^e$).
- Downward rate (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 agent (probability $(1 - p_H^2)^h$) but can form a cycle with an existing E agents (probability $1 - (1 - p_{EPH})^e$), or an E agent arrives that cannot form a cycle with any of the existing H agents (probability $(1 - p_{EPH})^h$) but can form a cycle with an existing E agent (probability $1 - (1 - p_E^2)^e$).

Note that the process is a 2-dimensional continuous-time spatially non-homogeneous random walk. Figure 9 illustrates this random walk along with its transition rates. Also observe that the




Figure 9: Transition rates (solid arrows) under the CTMC induced by the \mathcal{B}_H policy.

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 Subsection 5.3. However, fixing the priority, changing the tie-breaking rule between agents of the same type (for example, favoring agents with longer waiting times instead of selecting one at random) does not change the transition rates.

In Appendix G, 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. Note that in steady-state, the expected drift for this CTMC in both horizontal

and vertical dimension is zero. The drifts are given in (3a)-(3d), and therefore,

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

$$\begin{aligned} \mathbf{E} [& \lambda_E (1 - p_{EPH})^{H^{\mathcal{B}_H}} (1 - p_E^2)^{E^{\mathcal{B}_H}} - \lambda_H (1 - p_H^2)^{H^{\mathcal{B}_H}} (1 - (1 - p_{EPH})^{E^{\mathcal{B}_H}}) \\ & - \lambda_E (1 - p_{EPH})^{H^{\mathcal{B}_H}} (1 - (1 - p_E^2)^{E^{\mathcal{B}_H}})] = 0 \end{aligned} \quad (4b)$$

Assuming that the random variables $H^{\mathcal{B}_H}$ and $E^{\mathcal{B}_H}$ 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 $\mathbf{E}[H^{\mathcal{B}_H}]$ and $\mathbf{E}[E^{\mathcal{B}_H}]$.

- For $\lambda_H < \lambda_E$, if we plug $[H^{\mathcal{B}_H}, E^{\mathcal{B}_H}] = \left[\frac{\ln\left(\frac{\lambda_E}{\lambda_E - \lambda_H}\right)}{p_{EPH}}, \frac{-\ln(2)}{\ln(1 - p_E^2)} \right]$ into (4a)-(4b), the right-hand sides will be $O(p_H)$ terms.
- For $\lambda_H > \lambda_E$, if we plug $[H^{\mathcal{B}_H}, E^{\mathcal{B}_H}] = \left[\frac{\ln\left(\frac{2\lambda_H}{\lambda_H + \lambda_E}\right)}{p_H^2}, 0 \right]$ into (4a)-(4b), the right-hand sides will be $O(p_H^2)$ terms.

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

Proposition 3. [Lower-bound] *Under \mathcal{B}_H and in steady-state,*

- *If $\lambda_H < \lambda_E$, there exists a constant c_1 such that:*

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

- *If $\lambda_H > \lambda_E$, there exists a constant c_2 such that:*

$$\mathbf{P} \left[H^{\mathcal{B}_H} \leq \frac{1}{p_H^2} \left(\ln \left(\frac{2\lambda_H}{\lambda_E + \lambda_H} \right) - c_2 \sqrt{p_H} \right) \right] \leq o(p_H).^{31}$$

³¹While 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 \mathcal{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:

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

- 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:

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

Note that in both cases in Proposition 4, if $k = p_H^{-3/4}$ then the right-hand sides become $o(p_H^2)$. The proof of Theorem 1 is a straightforward application of these propositions and the details are presented in Appendix B.1. To prove these propositions we derive exponentially decaying bounds on tails of the steady-state distribution of $H^{\mathcal{B}_H}$ and $E^{\mathcal{B}_H}$. In the next subsection we present two auxiliary lemmas that establish such bounds for a general class of 2-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 Appendix B.4 and B.3, respectively.

5.2 Concentration bounds for a general class of 2-dimensional random walks

In the analysis of both $BilateralMatch(H)$ and $BilateralMatch(E)$ 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 2-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 2-dimensional random walks under certain conditions. These lemmas maybe 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 $[X, Y]$ be a corresponding random vector following its steady-state distribution. Suppose the following exist:

Condition 1. A set $S \subset \mathbb{N}$ and a constant $\epsilon > 0$ such that $\mathbf{P}[Y \notin S] \leq \epsilon$.

Condition 2. A non-increasing function $f : \mathbb{N} \mapsto (0, \infty)$ such that $\forall y \in S, Q([x, y], [x+1, y]) \geq f(x)$.

Condition 3. A non-decreasing 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:

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

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

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

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 non-decreasing and f is non-increasing, we get for $x \leq \eta$:

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

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 exist:*

Condition 1. A mapping $S : \mathbb{N} \mapsto 2^{\mathbb{N}}$ and two constants $c \in \mathbb{R}^+$, $\delta \in (0, 1)$ such that

$$\mathbf{P}[Y \notin S(x)] \leq c\delta^x.$$

Condition 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:*³²

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

The proof of Lemma 2 follows similar arguments to that of Lemma 1 and is deferred to Appendix A.

³²Note that the above conditions are weaker than that of Lemma 1 (where f is non-increasing, g is non-decreasing and $\frac{f(\eta)}{g(\eta+1)} \leq \rho$). We will need this for the proofs of Propositions 3 and 7 where the corresponding function g is not monotone.

5.3 The *BilateralMatch(E)* policy

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

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

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

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

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

The rates are computed similarly to those under the *BilateralMatch(H)*. 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 since we need to compute tight bounds also on the number of E agents in the market. While 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.

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

$$\mathbf{E}[\lambda_E(1 - p_{EPH})^{H^{\mathcal{B}_E}}(1 - p_E^2)^{E^{\mathcal{B}_E}} - \lambda_H(1 - (1 - p_{EPH})^{E^{\mathcal{B}_E}}) - \lambda_E(1 - (1 - p_E^2)^{E^{\mathcal{B}_E}})] = 0. \quad (6b)$$

Similar to the heuristic analysis for *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)-(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)-(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), i.e.,

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

In Figure 7 (Section 4.5), 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 number of E agent are not tight enough to result in a 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 subsection C.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 Appendix E for a heuristic analysis

that leads us to guess the constant that we can numerically verify to be the correct one (See Figure 8 of Subsection 4.5).

Instead of directly analyzing the $ChainMatch(d)$ 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 $ChainMatch(d)$ by $\hat{\mathcal{C}}(d)$. Observe that $H_t^{\hat{\mathcal{C}}(d)}$ is a 1-dimensional CTMC with the following transition rates:

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

$$Q^{\hat{\mathcal{C}}(d)}(h, h-i) = \left(\lambda_H(1-(1-p_H)^d) + \lambda_E(1-(1-p_E)^d) \right) (1-p_H)^{h-i} \prod_{j=0}^{i-1} (1-(1-p_H)^{h-j}),$$

$$i \in \{1, 2, \dots, h\}. \quad (7b)$$

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 .³³ In Appendix G, 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 .³⁴ For any $i \leq h$ we can write

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

³³Observe that the case $i = 0$ is possible, and corresponds to an arriving agent that can receive from 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.

³⁴Note that using the notation from Section 3.2.2, $S_h + 1$ corresponds to the length of the chain-segment L_k for the 1-dimensional Markov chain.

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 $\tilde{i} \leq i \leq h$:

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

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 2-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. 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 head, 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 tail, 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 realised independently.

- (c) $\tilde{h} < h$: we let $[L_{[h,e]}^H, L_{[h,e]}^E]$ and $S_{\tilde{h}}$ be realized independently.³⁵

³⁵This case is only defined here for the sake of completeness, the induction will ensure that this never happens.

2. Upon arrival of an E agent we flip a biased coin with probability $(1 - p_E)^d$. If the flip is head, the E agent cannot start a chain-segment in either system and we have: $H_{k+1}^{C(d)} = H_k^{C(d)}$, and $H_{k+1}^{\hat{C}(d)} = H_k^{\hat{C}(d)}$; on the other hand, if the flip is tail, 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: $H_0^{C(d)} = H_0^{\hat{C}(d)} = 0$. Suppose $H_k^{C(d)} \leq H_k^{\hat{C}(d)}$ 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 $H_{k+1}^{C(d)} \leq H_{k+1}^{\hat{C}(d)}$. 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 $\hat{C}(d)$ was not even long enough to bring the number of H agents back to $H_k^{C(d)} = h$; therefore $H_{k+1}^{C(d)} \leq H_{k+1}^{\hat{C}(d)}$ holds. In the latter case, again by coupling construction $L_{[h,e]}^H \geq S_h + (\tilde{h} - h)$, which implies that $H_{k+1}^{C(d)} \leq H_{k+1}^{\hat{C}(d)}$ holds. A similar argument holds if an E arrival starts a chain-segment.

□

The next proposition computes $\mathbf{E}[H^{\hat{C}(d)}]$ 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. 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:

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

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

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

where the last step follows from a change of variable $k' = k+1$. Note that the summation in the RHS of (13) also appears in the second term of RHS 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

$$\pi(h)\lambda_H(1-p_H)^d = \pi(h+1) \mathbf{P}[S_{h+1} \geq 1] \left(\Lambda + \lambda_H(1-p_H)^d \right). \quad (14)$$

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 $h \geq \eta$,

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

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:

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

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:

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

□

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. While our stylized model abstracts away from many details in this market, our findings may provide some useful insights to 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 O patients and non-O donors who compete to match with scarce 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 are a key step towards aligning incentives towards cooperation between the relevant players (Ashlagi and Roth (2012)). Our findings further provide some insights about tradeoffs 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 in order to increase matching opportunities.³⁶ Another interesting direction is to extend the model to allow departures.³⁷ Finally, our focus has been on marketplaces, in which any pair of agents have a non-zero 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, 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.

References

- Adan, I. and G. Weiss (2012). Exact FCFS Matching Rates for Two Infinite Multitype Sequences. *Operations Research* 60(2), 475–489.
- Agarwal, N., I. Ashlagi, E. Azevedo, C. Featherstone, and O. Karaduman (2017). Market failure in kidney exchange. Technical report, mimeo.
- Akbarpour, M., S. Li, and S. Oveis Gharan (2014). Dynamic matching market design. Working paper.
- Anderson, R., I. Ashlagi, D. Gamarnik, M. Rees, A. Roth, T. Sönmez, and M. Ünver (2015). Kidney exchange and the alliance for paired donation: Operations research changes the way kidneys are transplanted. *Interfaces* 45(1), 26–42.
- Anderson, R., I. Ashlagi, Y. Kanoria, and D. Gamarnik (2017). Efficient dynamic barter exchange. Forthcoming in *Operations Research*.

³⁶Similar to Anderson et al. (2017) we can show our policies achieve the same scaling as the best anonymous Markovian policy (see Proposition 9 in Appendix H) but characterising the best constants is an open question.

³⁷For example, Akbarpour et al. (2014) allow agents to depart prior to being matched and consider the match rate as the measure for efficiency.

- Ashlagi, I., A. Bingaman, M. Burq, V. Manshadi, M. Melcher, M. Murphey, A. E. Roth, and M. Rees (2017). The effect of match-run frequencies on the number of transplants and waiting times in kidney exchange. Forthcoming in *American Journal of Transplantation*.
- Ashlagi, I., D. Gamarnik, M. Rees, and A. Roth (2012). The need for (long) chains in kidney exchange. Technical report, National Bureau of Economic Research.
- Ashlagi, I. and A. Roth (2012). New Challenges in Multi-hospital Kidney Exchange. *American Economic Review, Papers and Proceedings* 102(3), 354–359.
- Ashlagi, I. and A. E. Roth (2014). Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics* 9(3), 817–863.
- Baccara, M., S. Lee, and L. Yariv (2015). Optimal dynamic matching. Working paper.
- Biro, P., L. Burnapp, H. Bernadette, A. Hemke, R. Johnson, J. van de Klundert, and D. Manlove (2017). First Handbook of the COST Action CA15210: European Network for Collaboration on Kidney Exchange Programmes (ENCKEP). http://www.enckep-cost.eu/assets/content/57/handbook1_28july2017-20170731121404-57.pdf. [Online].
- Bloch, F. and D. Cantala (2014). Dynamic allocation of objects to queuing agents: The discrete model.
- Böhmig, G. A., J. Fronek, A. Slavcev, G. F. Fischer, G. Berlakovich, and O. Viklicky (2017). Czech-austrian kidney paired donation: first european cross-border living donor kidney exchange. *Transplant International* 30(6), 638–639.
- Caldentey, R., E. H. Kaplan, and G. Weiss (2009). Fcfs infinite bipartite matching of servers and customers. *Adv. Appl. Prob* 41(3), 695–730.
- Dickerson, J. P., A. D. Procaccia, and T. Sandholm (2012a). Dynamic Matching via Weighted Myopia with Application to Kidney Exchange. *Proc of the 6th AAAI Conference on Artificial Intelligence*, 1340–1346.

- Dickerson, J. P., A. D. Procaccia, and T. Sandholm (2012b). Optimizing Kidney Exchange with Transplant Chains: Theory and Reality. *Proc of the eleventh international conference on autonomous agents and multiagent systems*.
- Ding, Y., D. Ge, S. He, and C. T. Ryan (2015). A non-asymptotic approach to analyzing kidney exchange graphs. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pp. 257–258. ACM.
- Doval, L. (2014). A theory of stability in dynamic matching markets. Technical report, mimeo.
- Feldman, J., A. Mehta, V. S. Mirrokni, and S. Muthukrishnan (2009). Online stochastic matching: Beating $1-1/e$. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 117–126.
- Ferrari, P., S. Fidler, R. Holdsworth, C. Woodroffe, G. Tassone, N. Watson, L. Cantwell, G. Bennett, A. Thornton, F. T. Christiansen, et al. (2012). High transplant rates of highly sensitized recipients with virtual crossmatching in kidney paired donation. *Transplantation* 94(7), 744–749.
- Ferrari, P., W. Weimar, R. J. Johnson, W. Lim, and K. Tinckam (2014). Kidney paired donation: principles, protocols and programs. *Nephrology Dialysis Transplantation* 30(8), 1276–1285.
- Fershtman, D. and A. Pavan (2015). Re-matching, experimentation, and cross-subsidization. Working paper.
- Glorie, K. M., J. J. van de Klundert, and A. P. Wagelmans (2014). Kidney exchange with long chains: An efficient pricing algorithm for clearing barter exchanges with branch-and-price. *Manufacturing & Service Operations Management* 16(4), 498–512.
- Goel, G. and A. Mehta (2008). Online budgeted matching in random input models with applications to adwords. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pp. 982–991.
- Gurvich, I., A. Ward, et al. (2014). On the dynamic control of matching queues. *Stochastic Systems* 4(2), 479–523.

- Jaillet, P. and X. Lu (2013). Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research* 39(3), 624–646.
- Johnson, R., J. Allen, S. Fuggle, J. A. Bradley, C. Rudge, et al. (2008). Early experience of paired living kidney donation in the united kingdom. *Transplantation* 86(12), 1672–1677.
- Kadam, S. and M. Kotowski (2014). Multi-period matching. Technical report, Mimeo.
- Karp, R. M., U. V. Vazirani, and V. V. Vazirani (1990). An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing (STOC)*, pp. 352–358.
- Leshno, J. (2014). Dynamic Matching in Overloaded Waiting Lists. Working paper.
- Malik, S. and E. Cole (2014). Foundations and principles of the canadian living donor paired exchange program. *Canadian Journal of Kidney Health and Disease* 1, 6. PMID: 25780601.
- Manshadi, V. H., S. Oveis-Gharan, and A. Saberi (2011). Online stochastic matching: online actions based on offline statistics. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1285–1294.
- Meyn, S. P. and R. L. Tweedie (1993). Stability of markovian processes iii: Foster–lyapunov criteria for continuous-time processes. *Advances in Applied Probability* 25(03), 518–548.
- Nikzad, A., M. Akbarpour, M. A. Rees, and A. E. Roth (2017). Financing transplants’ costs of the poor: A dynamic model of global kidney exchange. Technical report, Stanford University Working Paper.
- NKR (2017). Paired exchange results quarterly report. Technical report.
- Prieto-Rumeau, T. and O. Hernández-Lerma (2016). Uniform ergodicity of continuous-time controlled markov chains: A survey and new results. *Annals of Operations Research* 241(1), 249–293.
- Rees, M. A., J. E. Kopke, R. P. Pelletier, D. L. Segev, M. E. Rutter, A. J. Fabrega, J. Rogers, O. G. Pankewycz, J. Hiller, A. E. Roth, T. Sandholm, M. U. Ünver, and R. A. Montgomery (2009).

- A non-simultaneous extended altruistic donor chain. *New England Journal of Medicine* 360, 1096–1101.
- Roth, A. E. (2008). What have we learned from market design? *Economic Journal* 118, 285–310.
- Roth, A. E., T. Sönmez, and M. U. Ünver (2004). Kidney exchange. *Quarterly Journal of Economics* 119, 457–488.
- Roth, A. E., T. Sönmez, and M. U. Ünver (2005). Pairwise kidney exchange. *Journal of Economic Theory* 125, 151–188.
- Roth, A. E., T. Sönmez, and M. U. Ünver (2007). Efficient kidney exchange: coincidence of wants in markets with compatibility-based preferences. *American Economic Review* 97, 828–851.
- Roth, A. E., T. Sönmez, M. U. Ünver, F. L. Delmonico, and S. L. Saidman (2006). Utilizing list exchange and nondirected donation through chain kidney paired donations. *American Journal of Transplantation* 6, 2694–2705.
- Siegel-Itzkovich, J. (2017). Israel and Cyprus sign accord for sharing kidneys from live donors. <http://www.jpost.com/Israel-News/Israel-and-Cyprus-sign-accord-for-sharing-kidneys-from-live-donors-497422>. [Online].
- Sönmez, T., M. U. Ünver, and M. B. Yenmez (2017). Incentivized kidney exchange. Technical report, Boston College Department of Economics.
- Sönmez, T., U. Ünver, et al. (2017). Market design for living-donor organ exchanges: An economic policy perspective. *Oxford Review of Economic Policy* 33(4).
- Ünver, M. U. (2010). Dynamic Kidney Exchange. *Review of Economic Studies* 77(1), 372–414.

A Proof of Lemma 2

Proof of Lemma 2. Let $\pi(x, y)$ be the state joint distribution of $[X, Y]$, and let $\pi_X(x) = \sum_{y \geq 0} \pi(x, y)$ be the marginal distribution of X . For a given $x \geq 0$, we have:

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

Using Condition 2, we upper-bound the LHS and lower-bound the RHS, which results in having:

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

Let $\pi_S(x) = \mathbf{P}[X = x, Y \in S(x)] = \sum_{y \in S(x)} \pi(x, y)$. Observe that $\pi_X(x) \leq \pi_S(x) + c\delta^x \leq \pi_S(x) + c\rho^x$ from Condition 1 and for any $\rho \in [\delta, 1)$. Assuming that for $x \geq \eta$, $\frac{f(x)}{g(x+1)} \leq \rho$, we get:

$$\pi_S(x+1) \leq \frac{f(x)}{g(x+1)}\pi_S(x) + \frac{\mathbf{P}[Y \notin S(x)]}{g(x+1)} \leq \rho\pi_S(x) + \frac{c\delta^x}{g(x+1)} \leq \rho\pi_S(x) + \frac{c\rho^x}{g(\eta+1)},$$

where the last inequality results from the assumption $\frac{\delta^x}{g(x+1)} \leq \frac{\rho^x}{g(\eta+1)}$. We can now prove by induction that:

$$\pi_S(\eta + i) \leq \rho^i \left(\pi_S(\eta) + i \frac{c\rho^{\eta-1}}{g(\eta+1)} \right)$$

This allows us to conclude:

$$\begin{aligned}
\mathbf{P}[X \geq \eta + k] &= \sum_{i=\eta+k}^{\infty} \pi_X(i) \\
&\leq \sum_{i=k}^{\infty} \pi_S(\eta + i) + \sum_{i=k}^{\infty} c\rho^{\eta+i} \\
&\leq \sum_{i=k}^{\infty} \rho^i \pi_S(\eta) + \frac{c\rho^{\eta-1}}{g(\eta+1)} \sum_{i=k}^{\infty} i\rho^i + \sum_{i=k}^{\infty} c\rho^{\eta+i} \\
&\leq \frac{\rho^k}{1-\rho} \left(1 + c\rho^\eta + \frac{c\rho^{\eta-1}}{g(\eta+1)} \frac{k+1}{1-\rho} \right) \\
&\leq \frac{\rho^k}{1-\rho} \left(1 + c + \frac{c(k+1)}{g(\eta+1) - f(\eta)} \right)
\end{aligned}$$

□

B Missing proofs for *BilateralMatch(H)* Policy

In the proofs of this section and of the next ones we will use the following facts: for any bounded, non-negative function $\xi : \mathbb{R}_+ \mapsto \mathbb{R}_+$ and any constant $u > 0$, in the limit where $p_H \rightarrow 0$,

Fact 1. For $\eta = \frac{(\ln u) + \xi(p_H)}{p_H^2}$, we have $(1 - p_H^2)^\eta = \frac{e^{-\xi(p_H)}}{u} + O(p_H^2)$.

Fact 2. For any constants $p_E, r > 0$ and for $\eta = \frac{(\ln u) + \xi(p)}{p_H^2}$, we have $(1 - p_E p_H)^\eta = o(p_H^r)$.

Fact 3. For any constant p_E , and for $\eta = \frac{(\ln u) + \xi(p_H)}{p_E p_H}$, we have $(1 - p_H^2)^\eta = 1 - O(p_H)$.

Fact 4. For any constant p_E and for $\eta = \frac{(\ln u) + \xi(p)}{p_E p_H}$, we have $(1 - p_E p_H)^\eta = \frac{e^{-\xi(p)}}{u} + O(p_H)$.

B.1 Proof of Theorem 1

Proof of Theorem 1. We first upper-bound $\mathbf{E}[H^{\mathcal{B}_H}]$. Let

$$v(\lambda_H, \lambda_E, p_E, p_H) = \begin{cases} \frac{1}{p_E p_H} \left(\ln \left(\frac{\lambda_E}{\lambda_E - \lambda_H} \right) + c_3 \sqrt{p_H} \right) + p_H^{-3/4} & \text{when } \lambda_H < \lambda_E, \\ \frac{1}{p_H^2} \left(\ln \left(\frac{2\lambda_H}{\lambda_E + \lambda_H} \right) + c_4 \sqrt{p_H} \right) + p_H^{-3/4} & \text{when } \lambda_H > \lambda_E. \end{cases}$$

Where c_3 and c_4 are the constants from Proposition 4. Using the equality $\mathbf{E}[H^{\mathcal{B}_H}] = \sum_{i=1}^{\infty} \mathbf{P}[H^{\mathcal{B}_H} \geq i]$, we have:

$$\begin{aligned} \mathbf{E}[H^{\mathcal{B}_H}] &= \sum_{i=1}^{v(\lambda_H, \lambda_E, p_E, p_H) - 1} \mathbf{P}[H^{\mathcal{B}_H} \geq i] + \sum_{i=v(\lambda_H, \lambda_E, p_E, p_H)}^{\infty} \mathbf{P}[H^{\mathcal{B}_H} \geq i] \\ &\leq v(\lambda_H, \lambda_E, p_E, p_H) + \sum_{j=p_H^{-3/4}}^{\infty} \frac{\gamma(p_H)^j}{1 - \gamma(p_H)} \\ &\leq v(\lambda_H, \lambda_E, p_E, p_H) + \frac{\gamma''(p_H)^{p_H^{-3/4}}}{(1 - \gamma''(p_H))^2}. \end{aligned}$$

Where we denote $\gamma'' = \max(\gamma, \gamma')$ and we used the result from Proposition 4 : $\mathbf{P}[H^{\mathcal{B}_H} \geq v(\lambda_H, \lambda_E, p_E, p_H) + j] \leq \frac{\gamma''(p_H)^j}{1 - \gamma''(p_H)}$.

Applying the fact that $\gamma''(p_H)^{p_H^{-3/4}} = (1 - \sqrt{p_H} + o(\sqrt{p_H}))^{p_H^{-3/4}} = o(p_H^2)$, and some algebra we get the following upper-bound on $\mathbf{E}[H^{\mathcal{B}_H}]$:

- If $\lambda_H < \lambda_E$, then $\mathbf{E}[H^{\mathcal{B}_H}] \leq \frac{\ln\left(\frac{\lambda_E}{\lambda_E - \lambda_H}\right)}{p_E p_H} + o\left(\frac{1}{p_H}\right)$.
- If $\lambda_H > \lambda_E$, then $\mathbf{E}[H^{\mathcal{B}_H}] \leq \frac{\ln\left(\frac{2\lambda_H}{\lambda_H + \lambda_E}\right)}{p_H^2} + o\left(\frac{1}{p_H^2}\right)$.

Now we proceed to lower-bound $\mathbf{E}[H^{\mathcal{B}_H}]$: Applying Markov inequality to $\mathbf{E}[H^{\mathcal{B}_H}]$ and using Proposition 3, we get the following lower-bound on $\mathbf{E}[H^{\mathcal{B}_H}]$:

- If $\lambda_H < \lambda_E$, then $\mathbf{E}[H^{\mathcal{B}_H}] \geq \frac{\ln\left(\frac{\lambda_E}{\lambda_E - \lambda_H}\right)}{p_E p_H} + o\left(\frac{1}{p_H}\right)$.
- If $\lambda_H > \lambda_E$, then $\mathbf{E}[H^{\mathcal{B}_H}] \geq \frac{\ln\left(\frac{2\lambda_H}{\lambda_H + \lambda_E}\right)}{p_H^2} + o\left(\frac{1}{p_H^2}\right)$.

This completes the proof. □

B.2 Proof of Corollaries 1 and 2

Proof of Corollary 1. Define $x = \lambda_H/\lambda_E$, and $f(x) = \frac{\ln\left(\frac{1}{1-x}\right)}{x}$. Note that the constant $\frac{\ln\left(\frac{\lambda_E}{\lambda_E - \lambda_H}\right)}{p_E \lambda_H} = \frac{f(x)}{\lambda_E p_E}$. It is easy check that $f'(x) > 0$ in $x \in (0, 1)$, and therefore $f(x)$ is increasing in $x \in (0, 1)$. □

Proof of Corollary 2. Define $y = \lambda_H/\lambda_E$, and $g(y) = \frac{\ln\left(\frac{2y}{1+y}\right)}{y}$. Note that the constant $\frac{\ln\left(\frac{2\lambda_H}{\lambda_H + \lambda_E}\right)}{\lambda_H} = \frac{g(y)}{\lambda_E}$. It is easy to check that $g'(y) > 0$ when $y \in (1, y^*)$, and $g'(y) < 0$ in $y > y^*$ where y^* is the

solution of $g'(y) = 0$:

$$g'(y^*) = 0 \Leftrightarrow \frac{1}{y^* + 1} = \ln \left(\frac{2y^*}{1 + y^*} \right) \Leftrightarrow (y^* + 1) \ln(2 - 2/(y^* + 1)) = 1,$$

□

B.3 Proof of Proposition 4

Instead of proving Proposition 4, we prove a stronger result. This will be useful later on to prove an upper bound for E agents (see Lemma 4 in B.4).

Proposition 6. *Under \mathcal{B}_H , if $\lambda_H < \lambda_E$, for any non-negative bounded function $\xi(p_H)$, for all $k \geq 0$:*

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

where $\gamma(p_H) = \frac{\lambda_H}{\lambda_E - (\lambda_E - \lambda_H)e^{-\xi(p_H)}} + O(p_H)$.

If $\lambda_H > \lambda_E$, for any non-negative bounded function $\xi'(p_H)$, for all $k \geq 0$:

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

where $\gamma'(p_H) = \frac{e^{-\xi'(p_H)}}{2 - e^{-\xi'(p_H)}} + O(p_H^2)$.

Proof. We wish to apply Lemma 2 with $[X(t), Y(t)] = [H^{\mathcal{B}_H}(t), E^{\mathcal{B}_H}(t)]$, in the special case where $S(h) = \mathbb{N}$ for all h , and $c = \delta = 0$. This implies that we need to find functions $f(\cdot)$ and $g(\cdot)$ such that for all $e \in \mathbb{N}$, $f(h) \geq Q^{\mathcal{B}_H}([h, e], [h + 1, e])$, and $g(h) \leq Q^{\mathcal{B}_H}([h, e], [h - 1, e])$.

Let $f(h) = \lambda_H(1 - p_H^2)^h$ and $g(h) = \lambda_H(1 - (1 - p_H^2)^h) + \lambda_E(1 - (1 - p_E p_H)^h)$. We have

$$\frac{f(h)}{g(h+1)} = \frac{\lambda_H(1 - p_H^2)^h}{\lambda_H(1 - (1 - p_H^2)^{h+1}) + \lambda_E(1 - (1 - p_E p_H)^{h+1})}.$$

Case $\lambda_H < \lambda_E$: Take $\eta = \frac{1}{p_E p_H} \left(\ln \left(\frac{\lambda_E}{\lambda_E - \lambda_H} \right) + \xi(p_H) \right)$. Facts 3 and 4 imply respectively that

$(1 - p_H^2)^\eta = 1 + O(p_H)$ and $(1 - p_E p_H)^\eta = \frac{\lambda_E - \lambda_H}{\lambda_E} e^{-\xi(p_H)} + O(p_H)$. This yields for all $h \geq \eta$:

$$\frac{f(h)}{g(h+1)} \leq \frac{\lambda_H}{\lambda_E - (\lambda_E - \lambda_H)e^{-\xi(p_H)}} + O(p_H) := \gamma(p_H).$$

Case $\lambda_H > \lambda_E$: Taking $\eta = \frac{1}{p_H^2} \left(\ln \left(\frac{2\lambda_H}{\lambda_E + \lambda_H} \right) + \xi'(p_H) \right)$. Facts 1 and 2 imply respectively $(1 - p_H^2)^\eta = \frac{\lambda_E + \lambda_H}{2\lambda_H} e^{-\xi'(p_H)} + O(p_H^2)$ and $(1 - p_E p_H)^\eta = o(p_H^2)$. This yields for all $h \geq \eta$:

$$\frac{f(h)}{g(h+1)} \leq \frac{\frac{1}{2} e^{-\xi'(p_H)}}{1 - \frac{1}{2} e^{-\xi'(p_H)}} + O(p_H^2) := \gamma'(p_H).$$

In both cases, we conclude by applying Lemma 2 with $\rho = \gamma(p_H)$ or $\rho = \gamma'(p_H)$.

□

Proof of Proposition 4. The proof of Proposition 4 is a consequence of Proposition 6:

- In the case where $\lambda_H < \lambda_E$, take $\xi(p_H) = \sqrt{p_H} \frac{\lambda_H}{\lambda_E - \lambda_H}$. This implies that $\gamma(p_H) = \frac{\lambda_H}{\lambda_E - (\lambda_E - \lambda_H)e^{-\xi(p_H)}} + O(p_H) = 1 - \sqrt{p_H} + O(\sqrt{p_H})$.
- In the case where $\lambda_H > \lambda_E$, take $\xi'(p_H) = 2\sqrt{p_H}$. This implies that $\gamma'(p_H) = \frac{e^{-\xi'(p_H)}}{2 - e^{-\xi'(p_H)}} + O(p_H^2) = 1 - \sqrt{p_H} + O(\sqrt{p_H})$.

□

B.4 Proof of Proposition 3

The proof of Proposition 3 requires a concentration bound on the number of E agents, which we state in Lemma 4.

Lemma 4. *Under \mathcal{B}_H , and assuming $p_H \leq p_E$, there exist constants C and $\zeta < 1$ (which only depend on λ_H, λ_E , and p_E) such that for any $k \geq 0$, there exists p such that for any $p_H < p$, we have:*

$$\mathbf{P} \left[E^{\mathcal{B}_H} \geq \frac{1}{\sqrt{p_H}} + k \right] \leq C \zeta^k.$$

Proof. The proof is based on a bound on the right-tail distribution of E agents in the market. To do this, we will apply Lemma 2 with $[X(t), Y(t)] = [E^{\mathcal{B}_H}(t), H^{\mathcal{B}_H}(t)]$. Therefore, we find an

upper-bound $f(e) = \lambda_E(1 - p_E^2)^e \geq Q^{\mathcal{B}H}([h, e], [h, e + 1])$ on the upward transition (3c). Similarly, we would like to find a lower-bound $g(e)$ on the downward transition (3d), but we cannot find one for any $h \in \mathbb{N}$. Therefore we need to restrict our attention to some subset $S(e) \subset \mathbb{N}$. Recall that $Q^{\mathcal{B}H}([h, e], [h, e - 1]) = \lambda_H(1 - p_H^2)^h(1 - (1 - p_{EPH})^e) + \lambda_E(1 - p_{EPH})^h(1 - (1 - p_E^2)^e)$.

Case $\lambda_H < \lambda_E$:

$$S(e) = \left\{ h \in \mathbb{N} \mid h \leq \frac{1}{p_{EPH}} \left(\ln \left(\frac{\lambda_E}{\lambda_E - \lambda_H} \right) + \ln 2 \right) + e \right\}.$$

Applying Proposition 6 with $\xi(p_H) = \ln 2$, we get that $\mathbf{P}[H^{\mathcal{B}H} \notin S(e)] \leq \frac{\gamma(p_H)^e}{1 - \gamma(p_H)} = c\delta^e$, where $c = \frac{\lambda_H + \lambda_E}{\lambda_E - \lambda_H}$, $\delta = \frac{2\lambda_H}{\lambda_E + \lambda_H}$. Fact 4 implies that for all $h \in S(e)$, $(1 - p_{EPH})^h \geq \frac{\lambda_E - \lambda_H}{2\lambda_E}(1 - p_{EPH})^e + O(p_H)$. Keeping only the second term in $Q^{\mathcal{B}H}([h, e], [h, e - 1])$, we get the lower-bound:

$$Q^{\mathcal{B}H}([h, e], [h, e - 1]) \geq \frac{\lambda_E - \lambda_H}{2}(1 - p_{EPH})^e(1 - (1 - p_E^2)^e) + O(p_H) := g(e)$$

This yields:

$$\frac{f(e)}{g(e+1)} = \frac{2\lambda_E}{(\lambda_E - \lambda_H)(1 - (1 - p_E^2)^{e+1})} \left(\frac{1 - p_E^2}{1 - p_{EPH}} \right)^{e+1} + O(p_H).$$

We get for all $e \geq \frac{1}{\sqrt{p_H}}$, and p_H small enough, $\frac{f(e)}{g(e+1)} \leq \delta$. Furthermore, for $\rho = \frac{1+\delta}{2}$ and p_H small enough,

$$\frac{\delta^e}{g(e+1)} \leq \delta^e \frac{(1 - p_{EPH})^{\frac{1}{\sqrt{p_H}} - e - 1}}{g(\frac{1}{\sqrt{p_H}} + 1)} \leq \left(\frac{\delta}{1 - p_{EPH}} \right)^e \frac{1}{g(\frac{1}{\sqrt{p_H}} + 1)} \leq \frac{\rho^e}{g(\frac{1}{\sqrt{p_H}} + 1)}$$

We can apply Lemma 2 which yields the desired bound:

$$\begin{aligned} \mathbf{P} \left[E^{\mathcal{B}H} \geq \frac{1}{\sqrt{p_H}} + k \right] &\leq \frac{\rho^k}{1 - \rho} \left(1 + c + \frac{c(k+1)}{g(\frac{1}{\sqrt{p_H}} + 1) - f(\frac{1}{\sqrt{p_H}})} \right) \\ &\leq \frac{\rho^k}{1 - \rho} \left(1 + c + \frac{c(k+1)(\lambda_H + \lambda_E)}{2\lambda_H g(1/\sqrt{p_H} + 1)} \right) \\ &\leq \frac{\rho^k}{1 - \rho} \left(1 + c + \frac{c(k+1)(\lambda_H + \lambda_E)}{\lambda_H(\lambda_E - \lambda_H)} \right). \end{aligned}$$

Where we used first the fact that $\frac{f(1/\sqrt{p_H})}{g(1/\sqrt{p_H}+1)} \leq \delta = \frac{2\lambda_H}{\lambda_E + \lambda_H}$ and therefore $g(1/\sqrt{p_H} + 1) - f(1/\sqrt{p_H}) \geq \frac{\lambda_E - \lambda_H}{\lambda_E + \lambda_H}$ and second the fact that $g(1/\sqrt{p_H} + 1) = \frac{\lambda_E - \lambda_H}{2} + O(p_H)$.

Case $\lambda_H > \lambda_E$:

$$S(e) = \left\{ h \in \mathbb{N} \mid h \leq \frac{1}{p_H^2} \left(\ln \left(\frac{2\lambda_H}{\lambda_E + \lambda_H} \right) + \ln 2 \right) + e \right\}.$$

Applying Proposition 6 with $\xi(p_H) = \ln(2)$, we have $\mathbf{P}[H^{\mathcal{B}_E} \notin S(e)] \leq \frac{\gamma(p_H)^e}{1 - \gamma(p_H)} = c\delta^e$, with $c = 3/2$ and $\delta = 1/3$. Fact 1 implies that for all $h \in S(e)$, $(1 - p_H^2)^h \geq \frac{\lambda_H + \lambda_E}{4\lambda_H} (1 - p_H^2)^e + O(p_H^2)$. Keeping only the first term in $Q^{\mathcal{B}_H}([h, e], [h, e - 1])$, we get the lower-bound:

$$Q^{\mathcal{B}_H}([h, e], [h, e - 1]) \geq \frac{\lambda_H + \lambda_E}{4\lambda_H} (1 - p_H^2)^e (1 - (1 - p_E p_H)^e) + O(p_H^2) = g(e).$$

This yields:

$$\frac{f(e)}{g(e+1)} \leq \frac{4\lambda_H \lambda_E}{(\lambda_H + \lambda_E)(1 - (1 - p_E p_H)^e)} \left(\frac{1 - p_E^2}{1 - p_H^2} \right)^e + O(p_H^2).$$

Furthermore, with $\rho = 1/2$ we get that for p_H small enough, $\frac{\delta}{1 - p_H^2} \leq \rho$. This leads to

$$\frac{\delta^e}{g(e+1)} \leq \delta^e \frac{(1 - p_H^2)^{p_H^{-1/2} - e}}{g(p_H^{-1/2} + 1)} \leq \left(\frac{\delta}{1 - p_H^2} \right)^e \frac{1}{g(p_H^{-1/2} + 1)} \leq \frac{\rho^e}{g(p_H^{-1/2} + 1)}.$$

Therefore we can apply Lemma 2:

$$\begin{aligned} \mathbf{P}[E^{\mathcal{B}_H} \geq n_E + k] &\leq \frac{\rho^k}{1 - \rho} \left(1 + c + \frac{c(k+1)}{g(p_H^{-1/2} + 1) - f(p_H^{-1/2})} \right) \\ &\leq \frac{\rho^k}{1 - \rho} \left(1 + c + \frac{c(k+1)3}{2g(p_H^{-1/2})} \right) \\ &\leq \frac{\rho^k}{1 - \rho} \left(1 + c + \frac{6c(k+1)}{\lambda_H + \lambda_E} \right) \end{aligned}$$

Where we used the fact that for all $e \geq \frac{1}{\sqrt{p_H}}$, $\frac{f(e)}{g(e+1)} = o(p_H)$. Therefore for p_H small enough, $\frac{f(e)}{g(e+1)} \leq \frac{1}{3}$ and therefore $g(e+1) - f(e) \geq \frac{2}{3}g(e+1)$. Which concludes the proof. □

We can now prove Proposition 3. The main idea is to apply Lemma 1, in two cases separately, where in one case we have few ($\leq \frac{2}{\sqrt{p_H}}$) E agents, and in the other case where we have many. Using Lemma 4, we can exponentially bound the second case.

Proof of Proposition 3. We will apply Lemma 1. Consider $S = \{e \in \mathbb{R}, e \leq p_H^{-1/2} + p_H^{-1/2}\}$. Using Lemma 4 with $k = p_H^{-1/2}$, we get that $\mathbf{P}[E^{\mathcal{B}_H} \notin S] = o(p_H^4)$.

For $e \in S$, we have $(1 - p_E p_H)^e \geq 1 - p_E p_H \left(p_H^{-1/2} + p_H^{-1/2} + o(\sqrt{p_H}) \right) = 1 - 2p_E \sqrt{p_H} + o(\sqrt{p_H})$. Taking $f(h) = \lambda_H (1 - p_H^2)^h (1 - 2p_E \sqrt{p_H}) + o(\sqrt{p_H})$, we have for $e \in S$, $Q^{\mathcal{B}_H}([h, e], [h+1, e]) \geq f(h)$. Let $g(h) = Q^{\mathcal{B}_H}([h, e], [h-1, e]) = \lambda_H (1 - (1 - p_H^2)^h) + \lambda_E (1 - (1 - p_E p_H)^h)$, and note that $f(h)$ is non-increasing and $g(h)$ is non-decreasing. We have:

$$\frac{g(h+1)}{f(h)} = \frac{\lambda_H + \lambda_E - \lambda_H (1 - p_H^2)^{h+1} - \lambda_E (1 - p_E p_H)^{h+1}}{\lambda_H (1 - p_H^2)^h (1 - 2p_E \sqrt{p_H} + o(\sqrt{p_H}))}$$

In the case $\lambda_H < \lambda_E$, Take $\eta = \frac{1}{p_E p_H} \left(\ln \left(\frac{\lambda_E}{\lambda_E - \lambda_H} \right) - c'_1 \sqrt{p_H} \right)$. Using Facts 3 and 4, we have $(1 - p_H^2)^\eta = 1 + O(p_H)$ and $(1 - p_E p_H)^\eta = \frac{\lambda_E - \lambda_H}{\lambda_E} e^{c'_1 \sqrt{p_H}} + O(p_H) = \frac{\lambda_E - \lambda_H}{\lambda_E} (1 + c'_1 \sqrt{p_H} + o(\sqrt{p_H}))$, therefore:

$$\begin{aligned} \frac{g(\eta+1)}{f(\eta)} &= \frac{\lambda_H + \lambda_E - \lambda_H - (\lambda_E - \lambda_H)(1 + c'_1 \sqrt{p_H}) + o(\sqrt{p_H})}{\lambda_H (1 - 2p_E \sqrt{p_H}) + O(p_H)} \\ &= 1 - \left(\frac{\lambda_E - \lambda_H}{\lambda_H} c'_1 - 2p_E \right) \sqrt{p_H} + o(\sqrt{p_H}). \\ &= 1 - \sqrt{p_H} + o(\sqrt{p_H}). \end{aligned}$$

Where we fixed $c'_1 = \frac{\lambda_H(1+2p_E)}{\lambda_E - \lambda_H}$. Using Lemma 1 with $[X(t), Y(t)] = [H^{\mathcal{B}_H}(t), E^{\mathcal{B}_H}(t)]$, $k = p_H^{-3/4}$ and $\rho = \frac{g(\eta+1)}{f(\eta)}$, we get $\rho^k = o(p_H^2)$ and:

$$\begin{aligned} \mathbf{P}[H^{\mathcal{B}_H} \leq \eta - k] &\leq \eta \cdot o(p_H^4) \left(1 + \frac{1}{f(\eta) - g(\eta+1)} \right) + \frac{\rho^k}{1 - \rho} \\ &\leq \eta \cdot o(p_H^4) \left(1 + \frac{1}{f(\eta) \sqrt{p_H} + o(\sqrt{p_H})} \right) + \frac{o(p_H^2)}{\sqrt{p_H} + o(\sqrt{p_H})} \\ &= o(p_H). \end{aligned}$$

Taking $c_1 = c'_1 + p_E$, this enables us to conclude that $\mathbf{P}[H^{\mathcal{B}_H} \leq \frac{1}{p_E p_H} \left(\ln \left(\frac{\lambda_E}{\lambda_E - \lambda_H} \right) - c_1 p_H^{1/4} \right)] \leq o(p_H)$.

In the case $\lambda_H > \lambda_E$, Let $\eta = \frac{1}{p_H^2} \left(\ln \left(\frac{2\lambda_H}{\lambda_H + \lambda_E} \right) - c'_2 p_H^{1/4} \right)$. Using Facts 1 and 2, we have $(1 - p_H^2)^\eta = \frac{\lambda_H + \lambda_E}{2\lambda_H} e^{c'_2 p_H^{1/4}} + O(p_H^2)$ and $(1 - p_E p_H)^\eta = O(p_H^2)$. This implies that:

$$\begin{aligned} \rho &= \frac{g(\eta + 1)}{f(\eta)} = \frac{(\lambda_H + \lambda_E) - \frac{(\lambda_H + \lambda_E)}{2} e^{c'_2 \sqrt{p_H}} + O(p_H^2)}{\frac{(\lambda_H + \lambda_E)}{2} e^{c'_2 \sqrt{p_H}} (1 - 2p_E \sqrt{p_H}) + O(p_H^2)} \\ &= \frac{2 - (1 + c'_2 \sqrt{p_H})}{(1 + c'_2 \sqrt{p_H})(1 - 2p_E \sqrt{p_H})} + o(\sqrt{p_H}) \\ &= 1 - (2c'_2 - 2p_E) \sqrt{p_H} + o(\sqrt{p_H}) \\ &= 1 - \sqrt{p_H} + o(\sqrt{p_H}). \end{aligned}$$

Where we have chosen $c'_2 = \frac{1+2p_E}{2}$. Taking $k = p_H^{-3/4}$, we have $\rho^k = o(p_H^2)$. Applying Lemma 1, we get that for all :

$$\begin{aligned} \mathbf{P}[H^{\mathcal{B}_H} \leq \eta - k] &\leq \eta \cdot o(p_H^4) \left(1 + \frac{1}{f(\eta) - g(\eta + 1)} \right) + \frac{\rho^k}{1 - \rho} \\ &= \eta \cdot o(p_H^3) + \frac{o(p_H^2)}{f(\eta) \sqrt{p_H} + o(\sqrt{p_H})} \\ &= o(p_H) \end{aligned}$$

Taking $c_2 = c'_2 + 1$, this enables us to conclude that $\mathbf{P}[H^{\mathcal{B}_H} \leq \frac{1}{p_H^2} \left(\ln \left(\frac{2\lambda_H}{\lambda_H + \lambda_E} \right) - c_2 \sqrt{p_H} \right)] \leq o(p_H)$. □

C Missing proofs for *BilateralMatch(E)* Policy

The proof of Theorem 2 relies on a concentration result on the left tail of the distribution of $H^{\mathcal{B}_E}$ (Proposition 7), and a coupling argument to upper-bound $\mathbf{E}[H^{\mathcal{B}_E}]$ (Proposition 8).

Proposition 7. [Lower Bound] *Under \mathcal{B}_E in steady-state,*

- *If $\lambda_H < \lambda_E$, there exist a function $\gamma(p_H) = 1 - \sqrt{p_H} + o(\sqrt{p_H})$ and a constant $c_5 \geq 0$ such*

that for all $k \geq 0$:

$$\mathbf{P} \left[H^{\mathcal{B}_E} \leq \frac{1}{p_E p_H} \left(\ln \left(\frac{\lambda_E}{\lambda_E - \lambda_H} \right) - c_5 p_H^{1/4} \right) - k \right] \leq \frac{\gamma(p_H)^k}{1 - \gamma(p_H)}.$$

- If $\lambda_H > \lambda_E$, there exist a function $\gamma'(p_H) = 1 - \sqrt{p_H} + o(\sqrt{p_H})$ and a constant $c_6 \geq 0$ such that for all $k \geq 0$:

$$\mathbf{P} \left[H^{\mathcal{B}_E} \leq \frac{1}{p_H^2} \left(\ln \left(\frac{2\lambda_H}{\lambda_H + \lambda_E} \right) - c_6 \sqrt{p_H} \right) - k \right] \leq \frac{\gamma'(p_H)^k}{1 - \gamma'(p_H)}.$$

Proposition 8. [Upper-bound] Under \mathcal{B}_E and in steady-state,

- If $\lambda_H < \lambda_E$, then $\mathbf{E}[H^{\mathcal{B}_E}] \leq \frac{\ln\left(\frac{2\lambda_E}{\lambda_E - \lambda_H}\right)}{p_E p_H} + o\left(\frac{1}{p_H}\right)$
- If $\lambda_H > \lambda_E$, then $\mathbf{E}[H^{\mathcal{B}_E}] \leq \frac{1}{p_H^2} \ln\left(\frac{2\lambda_H}{\lambda_H + \lambda_E}\right) + o\left(\frac{1}{p_H^2}\right)$.

Proof of Theorem 2. We will first compute a lower bound for $\mathbf{E}[H^{\mathcal{B}_E}]$: Applying Markov inequality to $\mathbf{E}[H^{\mathcal{B}_E}]$ and using Proposition 7 with $k = p_H^{-3/4}$, we get the following lower-bound on $\mathbf{E}[H^{\mathcal{B}_E}]$:

- If $\lambda_H < \lambda_E$, then $\mathbf{E}[H^{\mathcal{B}_E}] \geq \frac{\ln\left(\frac{\lambda_E}{\lambda_E - \lambda_H}\right)}{p_E p_H} + o\left(\frac{1}{p_H}\right)$.
- If $\lambda_H > \lambda_E$, then $\mathbf{E}[H^{\mathcal{B}_E}] \geq \frac{\ln\left(\frac{2\lambda_H}{\lambda_H + \lambda_E}\right)}{p_H^2} + o\left(\frac{1}{p_H^2}\right)$.

Using Proposition 8, we can get the desired upper-bounds for $\mathbf{E}[H^{\mathcal{B}_E}]$. We can conclude using Little's law: $w_H = \frac{\mathbf{E}[H^{\mathcal{B}_E}]}{\lambda_H}$.

□

C.1 Proof of Proposition 7.

In order to prove Proposition 7, we first need a concentration result on $E^{\mathcal{B}_E}$ agents. This is stated in Lemma 5.

Lemma 5. Under \mathcal{B}_E in steady-state, for any $\alpha \in [0, 1]$, let $n_E(\alpha) = \frac{\ln(\alpha/(1+\alpha(1-p_E^2)))}{\ln(1-p_E^2)}$. For any $k \geq 0$, we have:

$$\mathbf{P} [E^{\mathcal{B}_E} \geq n_E(\alpha) + k] \leq \frac{\alpha^k}{1 - \alpha}.$$

Also, $(1 - p_E^2)^{n_E(\alpha)} = \frac{\alpha}{1 + \alpha(1 - p_E^2)}$.

Proof. In the same way that upper-bounding H agents when they get priority is easy because waiting E agents can be ignored, upper-bounding the number of E agents is easy when they get the priority because H agents can be ignored. We can get the following bounds on the transition probabilities: $Q^{\mathcal{B}_E}([h, e], [h, e + 1]) \leq \lambda_E(1 - p_E^2)^e = f(e)$, and $Q^{\mathcal{B}_E}([h, e], [h, e - 1]) \geq \lambda_E(1 - (1 - p_E^2)^e) = g(e)$ which leads to:

$$\frac{f(e)}{g(e + 1)} = \frac{(1 - p_E^2)^e}{1 - (1 - p_E^2)^{e+1}}.$$

Setting $\eta = n_E(\alpha)$ to be the solution to $\frac{f(n)}{g(n+1)} = \alpha$, and applying Lemma 2 with $[X(t), Y(t)] = [E(t), H(t)]$, and $S(e) = \mathbb{N}$ for all h , $c = 0$ and $\delta = 0$, we get the desired result. \square

We can now prove Proposition 7.

Proof of Proposition 7. Notice that $Q^{\mathcal{B}_E}([h, e], [h + 1, e]) = Q^{\mathcal{B}_H}([h, e], [h + 1, e])$, therefore the function $f(h) = \lambda_H(1 - p_H^2)^h (1 - 2p_E\sqrt{p_H}) + o(\sqrt{p_H})$ used in the proof of Proposition 3 remains a lower bound for $e \in S = \{e \leq 2p_H^{-1/2}\}$. Furthermore, $Q^{\mathcal{B}_E}([h, e], [h - 1, e]) \leq Q^{\mathcal{B}_H}([h, e], [h - 1, e])$ therefore the function $g(h) = \lambda_H(1 - (1 - p_H^2)^h) + \lambda_E(1 - (1 - p_E p_H)^h)$ remains an upper bound. Therefore, the proof is exactly the same as the proof for the lower bounds for Proposition 3, except that instead of Lemma 4 we use Lemma 5 in the special case $n_E(\alpha) = p_H^{-1/2}$ with $S = \{e \in \mathbb{R}, e \leq 2p_H^{-1/2}\}$, and $k = p_H^{-1/2}$. We still have, $\mathbf{P}[E^{\mathcal{B}_H} \notin S] = o(p_H^4)$. \square

C.2 Proof of Propostion 8

This proof is based on a coupling argument: instead of analysing the CTMC resulted from \mathcal{B}_E , we analyse the CTMC underlying another process that we call $\widetilde{\mathcal{B}}_E$. $\widetilde{\mathcal{B}}_E$ works similarly to \mathcal{B}_E with one crucial difference: each time an E agent arrives which cannot be matched immediately, it changes type and becomes an H agent and then joins the market. First note that under $\widetilde{\mathcal{B}}_E$, no E agent joins the market, making its underlying CTMC a 1-dimensional Markov chain. Proving that the stochastic process underlying $\widetilde{\mathcal{B}}_E$ is a positive recurrent CTMC, and therefore it reaches steady-state is similar to the proof of positive recurrence of \mathcal{B}_E (See Claim 3) and therefore omitted.

Using a coupling argument we show that in steady-state number of H agents under \mathcal{B}_E can be upper-bounded by the number of H agents under $\widetilde{\mathcal{B}}_E$. Our motivation behind defining $\widetilde{\mathcal{B}}_E$ is that in \mathcal{B}_E , unmatched E agents in the market are competing against H agents over whom they have priority. The idea is that by turning unmatched E agents into H ones, we expect to have more H agents waiting. Let $H^{\widetilde{\mathcal{B}}_E}$ be the random number of H agents under $\widetilde{\mathcal{B}}_E$ in steady-state. First in the next lemma we show that $\mathbf{E}[H^{\mathcal{B}_E}] \leq \mathbf{E}[H^{\widetilde{\mathcal{B}}_E}] + 1$, then we compute an upper-bound on $\mathbf{E}[H^{\widetilde{\mathcal{B}}_E}]$.

Lemma 6. $\mathbf{E}[H^{\mathcal{B}_E}] \leq \mathbf{E}[H^{\widetilde{\mathcal{B}}_E}] + 1$.

Proof. Consider two copies of the arrival process, one under the setting of \mathcal{B}_E and one under $\widetilde{\mathcal{B}}_E$. Let $[H_k^{\mathcal{B}_E}, E_k^{\mathcal{B}_E}]$ and $H_k^{\widetilde{\mathcal{B}}_E}$ denote the embedded (discrete-time) Markov chain resulting from observing the two dynamic systems at arrival epochs. We will prove the following (stronger) result: there exists a coupling such that at any step k , $H_k^{\mathcal{B}_E} + E_k^{\mathcal{B}_E} \leq H_k^{\widetilde{\mathcal{B}}_E} + 1$.

We start by coupling all arrivals: with probability $\frac{\lambda_H}{\lambda_E + \lambda_H}$ ($\frac{\lambda_E}{\lambda_E + \lambda_H}$), arrivals to \mathcal{B}_E and $\widetilde{\mathcal{B}}_E$ at k are both H (E) agents. Three cases can happen:

1. $H_k^{\mathcal{B}_E} + E_k^{\mathcal{B}_E} < H_k^{\widetilde{\mathcal{B}}_E}$; in this case, we let the two chains evolve independently.
2. $H_k^{\widetilde{\mathcal{B}}_E} \leq H_k^{\mathcal{B}_E} + E_k^{\mathcal{B}_E} \leq H_k^{\widetilde{\mathcal{B}}_E} + 1$. Let $h = H_k^{\widetilde{\mathcal{B}}_E}$. We consider two sub-cases:
 - (a) The arrival at $k + 1$ is an H agent. We couple the events that $\widetilde{\mathcal{B}}_E$ and \mathcal{B}_E cannot find a match as follows: we draw two independent Bernoulli random variables $\mathbb{B}_1, \mathbb{B}_2$ with respective parameters of $(1 - p_H^2)^h$ and $(1 - p_H^2)^{H_k^{\mathcal{B}_E}} (1 - p_H p_E)^{E_k^{\mathcal{B}_E}} (1 - p_H^2)^{-h}$. $\mathbb{B}_1 = 1$ corresponds to the event that $\widetilde{\mathcal{B}}_E$ cannot find a match; similarly $\mathbb{B}_1 \mathbb{B}_2 = 1$ corresponds to the event that \mathcal{B}_E cannot find a match.
 - (b) The arrival at $k + 1$ is an E agent. Similarly we couple the events that $\widetilde{\mathcal{B}}_E$ and \mathcal{B}_E cannot find a match as follows: we draw two independent Bernoulli random variables $\mathbb{B}_3, \mathbb{B}_4$ with respective parameters of $(1 - p_E p_H)^h$ and $(1 - p_E p_H)^{H_k^{\mathcal{B}_E}} (1 - p_E^2)^{E_k^{\mathcal{B}_E}} (1 - p_E p_H)^{-h}$. $\mathbb{B}_3 = 1$ corresponds to the event that $\widetilde{\mathcal{B}}_E$ cannot find a match. Similarly $\mathbb{B}_3 \mathbb{B}_4 = 1$ corresponds to the event that \mathcal{B}_E cannot find a match.
3. $H_k^{\mathcal{B}_E} + E_k^{\mathcal{B}_E} > H_k^{\widetilde{\mathcal{B}}_E} + 1$; in this case, we let the two chains evolve independently.

Having the above coupling, we finish the proof by induction. The base case $k = 0$ is trivial: $H_0^{\mathcal{B}_E} + E_0^{\mathcal{B}_E} = H_0^{\widetilde{\mathcal{B}_E}} = 0$. Suppose $H_k^{\mathcal{B}_E} + E_k^{\mathcal{B}_E} \leq H_k^{\widetilde{\mathcal{B}_E}} + 1$, we show that $H_{k+1}^{\mathcal{B}_E} + E_{k+1}^{\mathcal{B}_E} \leq H_{k+1}^{\widetilde{\mathcal{B}_E}} + 1$: In Case (1), note that the number of agents waiting can increase or decrease by at most 1. For both cases, we have $H_{k+1}^{\mathcal{B}_E} + E_{k+1}^{\mathcal{B}_E} \leq H_{k+1}^{\widetilde{\mathcal{B}_E}} + 1$. In Case (2a), we have:

- If $\mathbb{B}_1 = 0$, then $H_{k+1}^{\widetilde{\mathcal{B}_E}} = H_k^{\widetilde{\mathcal{B}_E}} - 1$. Further if $\mathbb{B}_1 = 0$ then $\mathbb{B}_1\mathbb{B}_2 = 0$, and $H_{k+1}^{\mathcal{B}_E} + E_{k+1}^{\mathcal{B}_E} = H_k^{\mathcal{B}_E} + E_k^{\mathcal{B}_E} - 1$. Therefore $H_{k+1}^{\mathcal{B}_E} + E_{k+1}^{\mathcal{B}_E} \leq H_{k+1}^{\widetilde{\mathcal{B}_E}} + 1$
- If $\mathbb{B}_1 = 1$, then $H_{k+1}^{\widetilde{\mathcal{B}_E}} = H_k^{\widetilde{\mathcal{B}_E}} + 1$. Note that number of agents under \mathcal{B}_E can increase by at most one, therefore $H_{k+1}^{\mathcal{B}_E} + E_{k+1}^{\mathcal{B}_E} \leq H_{k+1}^{\widetilde{\mathcal{B}_E}} + 1$.

Similarly in Case (2b), we have:

- If $\mathbb{B}_3 = 0$, then $H_{k+1}^{\widetilde{\mathcal{B}_E}} = H_k^{\widetilde{\mathcal{B}_E}} - 1$. Further if $\mathbb{B}_3 = 0$ then $\mathbb{B}_3\mathbb{B}_4 = 0$, and $H_{k+1}^{\mathcal{B}_E} + E_{k+1}^{\mathcal{B}_E} = H_k^{\mathcal{B}_E} + E_k^{\mathcal{B}_E} - 1$; therefore $H_{k+1}^{\mathcal{B}_E} + E_{k+1}^{\mathcal{B}_E} \leq H_{k+1}^{\widetilde{\mathcal{B}_E}} + 1$
- If $\mathbb{B}_3 = 1$, then $H_{k+1}^{\widetilde{\mathcal{B}_E}} = H_k^{\widetilde{\mathcal{B}_E}} + 1$. Note that number of agents under \mathcal{B}_E can increase by at most one, therefore and $H_{k+1}^{\mathcal{B}_E} + E_{k+1}^{\mathcal{B}_E} \leq H_{k+1}^{\widetilde{\mathcal{B}_E}} + 1$.

Thus, in all possible cases $H_{k+1}^{\mathcal{B}_E} + E_{k+1}^{\mathcal{B}_E} \leq H_{k+1}^{\widetilde{\mathcal{B}_E}} + 1$ finishing the proof. Note that the above induction also implies that Case (3) never occurs. □

Proof of Proposition 8. Observe that the CTMC $H_t^{\widetilde{\mathcal{B}_E}}$ has the following rate matrix:

$$\begin{aligned} Q^{\widetilde{\mathcal{B}_E}}(h, h+1) &= \lambda_H(1 - p_H^2)^h + \lambda_E(1 - p_{EPH})^h \\ Q^{\widetilde{\mathcal{B}_E}}(h, h-1) &= \lambda_H(1 - (1 - p_H^2)^h) + \lambda_E(1 - (1 - p_{EPH})^h) \end{aligned}$$

And let us define:

$$\rho(h) = \frac{Q^{\widetilde{\mathcal{B}_E}}(h, h+1)}{Q^{\widetilde{\mathcal{B}_E}}(h+1, h)} = \frac{\lambda_H(1 - p_H^2)^h + \lambda_E(1 - p_{EPH})^h}{\lambda_H(1 - (1 - p_H^2)^{h+1}) + \lambda_E(1 - (1 - p_{EPH})^{h+1})}.$$

Note that $\rho(h)$ is a decreasing function, and suppose that there exists $\eta \geq 0$ such that $\rho(\eta) < 1 - \sqrt{p_H}$ and let $\pi = \pi(h)_{h \geq 0}$ be the stationary distribution of $H^{\widetilde{\mathcal{B}_E}}$. Then we have for all $i \geq 0$,

$\pi(\eta + i) \leq \rho(\eta)^i \pi(\eta) \leq \rho(\eta)^i$ and for all $k \geq 0$, $\mathbf{P}[H^{\widetilde{\mathcal{B}}_E} \geq \eta + k] \leq \pi(\eta) \sum_{i \geq k} \rho(\eta)^i \leq \frac{\rho(\eta)^k}{1 - \rho(\eta)}$. We then have

$$\begin{aligned} \mathbf{E}[H^{\widetilde{\mathcal{B}}_E}] &= \sum_{h \geq 1} \mathbf{P}[H^{\widetilde{\mathcal{B}}_E} \geq h] \\ \mathbf{E}[H^{\widetilde{\mathcal{B}}_E}] &\leq (\eta + k) + \sum_{h \geq \eta + k} \mathbf{P}[H^{\widetilde{\mathcal{B}}_E} \geq h] \\ \mathbf{E}[H^{\widetilde{\mathcal{B}}_E}] &\leq (\eta + k) + \frac{\rho(\eta)^k}{(1 - \rho(\eta))^2} \end{aligned}$$

Let us consider the case when $\lambda_H < \lambda_E$. Let $\eta = \frac{\ln(1/u) + c\sqrt{p_H}}{p_E p_H}$. For $h \geq \eta$, we have $(1 - p_E p_H)^h \leq u(1 - c\sqrt{p_H} + o(\sqrt{p_H}))$ and $(1 - p_H^2)^h \leq 1 - \ln(1/u)p_H^2 + o(p_H^2)$. We have $\rho(\eta) = \frac{\lambda_H + \lambda_E u(1 - c\sqrt{p_H})}{\lambda_E(1 - u(1 - c\sqrt{p_H}))} + o(\sqrt{p_H})$. Taking $u = \frac{\lambda_E - \lambda_H}{2\lambda_E}$ and $c = \frac{\lambda_E + \lambda_H}{2(\lambda_E - \lambda_H)}$, we get $\rho(\eta) = 1 - \sqrt{p_H} + o(\sqrt{p_H})$. Taking $k = p_H^{-3/4}$ we get:

$$\mathbf{E}[H^{\widetilde{\mathcal{B}}_E}] \leq \eta + o(1/p_H) = \frac{\ln(\frac{2\lambda_E}{\lambda_E - \lambda_H})}{p_E p_H} + o(1/p_H), \quad \lambda_H < \lambda_E. \quad (16)$$

In the case when $\lambda_H > \lambda_E$, let $\eta = \frac{\ln(1/u) + c\sqrt{p_H}}{p_H^2}$. For $h \geq \eta$, we have $(1 - p_E p_H)^h = o(p_H)$ and $(1 - p_H^2)^h \leq u(1 - c\sqrt{p_H}) + o(\sqrt{p_H})$. Taking $u = \frac{\lambda_H + \lambda_E}{2\lambda_H}$ and $c = 1/2$, we get $\rho(\eta) = 1 - \sqrt{p_H} + o(\sqrt{p_H})$. Taking $k = p_H^{-3/4}$ we get:

$$\mathbf{E}[H^{\widetilde{\mathcal{B}}_E}] \leq \eta + o(1/p_H^2) = \frac{\ln(\frac{2\lambda_H}{\lambda_E + \lambda_H})}{p_H^2} + o(1/p_H^2), \quad \lambda_H > \lambda_E. \quad (17)$$

Putting Lemma 6 and upper-bounds given in (16) and (17) together finishes the proof. \square

D Missing proofs for *ChainMatch(d) Policy*

Proof of Proposition 1. As pointed at the end of Subsection 5.4, when $p_E = 1$, 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 the Proposition 1.

Further note that fixing λ_H , it is straightforward to check that function $f(\lambda_E) = \frac{\ln\left(\frac{\lambda_H}{\lambda_E} + 1\right)}{\lambda_H}$ is decreasing for $\lambda_E > 0$. Similarly fixing λ_E , it is straightforward to check that function $g(\lambda_H) = \frac{\ln\left(\frac{\lambda_H}{\lambda_E} + 1\right)}{\lambda_H}$ is decreasing for $\lambda_H > 0$. \square

Proof of Corollary 3. First we observe that $\ln\left(\frac{\lambda_H}{\lambda_E(1-(1-p_E)^d)} + 1\right)$ is decreasing in d . Therefore the worst upper-bound on $w_H^{C(d)}$ is for $d = 1$. Next we show that:

$$\begin{aligned} \ln\left(\frac{\lambda_H}{\lambda_E p_E} + 1\right)^{p_E} &\leq \ln\left(\frac{\lambda_H}{\lambda_E} + 1\right) \\ &\leq \ln\left(\frac{\lambda_H}{\lambda_E - \lambda_H} + 1\right), \end{aligned}$$

where the first inequality holds because function $f(x) := \left(\frac{\lambda_H/\lambda_E}{x} + 1\right)^x$ is increasing in $x \in (0, 1]$. \square

Proof of Proposition 2. In steady-state, for any function $f(\cdot, \cdot)$, $0 = \mathbb{E}[f(H_k, E_k) - f(H_{k+1}, E_{k+1})]$, and in particular for $f(h, e) = h + e$, we get

$$\begin{aligned} 0 &= (\lambda_H(1-p_H)^d + \lambda_E(1-p_E)^d)\mathbb{E}[H_k + E_k - H_{k+1} - E_{k+1} | H_k + E_k < H_{k+1} + E_{k+1}] \\ &\quad + (\lambda_H(1-(1-p_H)^d) + \lambda_E(1-(1-p_E)^d))\mathbb{E}[H_k + E_k - H_{k+1} - E_{k+1} | H_k + E_k \geq H_{k+1} + E_{k+1}] \\ &= -(\lambda_H(1-p_H)^d + \lambda_E(1-p_E)^d) + (\lambda_H(1-(1-p_H)^d) + \lambda_E(1-(1-p_E)^d))\mathbf{E}[L - 1 | L \geq 1]. \end{aligned} \tag{18}$$

This gives us:

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

\square

E Heuristic argument for *ChainMatch(d)* Policy

In what follows we provide a heuristic analysis of the CTMC underlying the *ChainMatch(d)*. This heuristic adds intuition for the behavior of the policy and further establishes what is supposedly the right constant in the case in which our theory only generates an upper bound ($p_E < 1$).

We introduce an auxiliary 3-dimensional CTMC, in which a chain-segment is not formed instantaneously; instead a chain-segment can only advance at certain “tokens” (or times) that arrive according to a Poisson process with rate μ ³⁸. We denote this auxiliary CTMC by $\tilde{\mathcal{C}}(d, \mu)$ and its states by $[H_t^{\tilde{\mathcal{C}}(d, \mu)}, E_t^{\tilde{\mathcal{C}}(d, \mu)}, U_t] \in \mathbb{N}^2 \times \{0, 1\}$. The first two dimensions represent the number of H and E agents. The third dimension U_t indicates whether a chain-segment is being conducted. Initially $U_0 = 0$. Suppose at time t_1 the first agent i arrives that is matched by an altruistic agent. At this time U_{t_1} becomes 1 indicating that a chain-segment formation is in process. The policy waits x unit of time for a token to arrive to find an agent who can be matched by i (note that x is exponentially distributed with rate μ)³⁹. If such an agent does not exist, then chain-segment is stopped; and U_{t_1+x} becomes 0. Otherwise U_{t_1+x} remains 1 and the same process repeats. The transition rates of $\tilde{\mathcal{C}}(d, \mu)$ are:

$$Q^{\tilde{\mathcal{C}}(d, \mu)}([h, e, 0], [h + 1, e, 0]) = \lambda_H(1 - p_H)^d \quad (19a)$$

$$Q^{\tilde{\mathcal{C}}(d, \mu)}([h, e, 0], [h, e + 1, 0]) = \lambda_E(1 - p_E)^d \quad (19b)$$

$$Q^{\tilde{\mathcal{C}}(d, \mu)}([h, e, 0], [h, e, 1]) = \lambda_H(1 - (1 - p_H)^d) + \lambda_E(1 - (1 - p_E)^d) \quad (19c)$$

$$Q^{\tilde{\mathcal{C}}(d, \mu)}([h, e, 1], [h, e, 0]) = \mu(1 - p_H)^h(1 - p_E)^e \quad (19d)$$

$$Q^{\tilde{\mathcal{C}}(d, \mu)}([h, e, 1], [h - 1, e, 1]) = \mu \left(1 - (1 - p_H)^h\right) \quad (19e)$$

$$Q^{\tilde{\mathcal{C}}(d, \mu)}([h, e, 1], [h, e - 1, 1]) = \mu(1 - p_H)^h(1 - (1 - p_E)^e) \quad (19f)$$

Figure 10 illustrates the transitions of $\tilde{\mathcal{C}}(d, \mu)$ in $H_t^{\tilde{\mathcal{C}}(d, \mu)}$ and U_t dimensions along with the

³⁸This process is independent from the Poisson processes guiding the arrivals of H and E agents.

³⁹Note that it is possible that there is a new arrivals of H and/or E agents in the interval $(t_1, t_1 + x]$. In this heuristic, we assume that the arriving H or E agent is deleted. The intuition is that when μ is large, this event happens rarely.




Figure 10: Transitions for the CTMC underlying $\tilde{\mathcal{C}}(d, \mu)$ in the first and third dimensions, i.e., transitions to states with a different number of E agents are not shown. Transition rates are given only for solid arrows.

transition rates. The rate (19a) ((19b)) corresponds to the event that an H (E) agent arrives but cannot be matched by any bridge agent which happens with probability $(1 - p_H)^d$ ($(1 - p_E)^d$). The rate (19c) corresponds to the case an H arrives and starts a chain-segment (occurs with probability $(1 - (1 - p_H)^d)$) or an E agent arrives and starts a chain-segment (occurs with probability $(1 - (1 - p_E)^d)$); similarly (19d) represents the case where the chain-segment cannot advance any further (happens with probability $(1 - p_H)^h(1 - p_E)^e$). The last two rates correspond to adding one more H and E agent to the chain-segment respectively (with probability $(1 - (1 - p_H)^h)$ and $(1 - p_H)^h(1 - (1 - p_E)^e)$ respectively).

Similar to the heuristic argument for $BilateralMatch(H)$ and $BilateralMatch(E)$, we can try to solve the following system of nonlinear equations (that result from setting expected drifts in all three dimensions to be zero):

$$\begin{aligned} \mathbf{E} \left[\left(\lambda_H(1 - (1 - p_H)^d) + \lambda_E(1 - (1 - p_E)^d) \right) \mathbb{I}(U = 0) - \mu(1 - p_H)^{H^{\tilde{\mathcal{C}}(d, \mu)}} (1 - p_E)^{E^{\tilde{\mathcal{C}}(d, \mu)}} \mathbb{I}(U = 1) \right] &= 0 \\ \mathbf{E} \left[\lambda_E(1 - p_E)^d \mathbb{I}(U = 0) - \mu(1 - p_H)^{H^{\tilde{\mathcal{C}}(d, \mu)}} (1 - (1 - p_E)^{E^{\tilde{\mathcal{C}}(d, \mu)}}) \mathbb{I}(U = 1) \right] &= 0 \\ \mathbf{E} \left[\lambda_H(1 - p_H)^d \mathbb{I}(U = 0) - \mu(1 - (1 - p_H)^{H^{\tilde{\mathcal{C}}(d, \mu)}}) \mathbb{I}(U = 1) \right] &= 0, \end{aligned}$$

where $\mathbb{I}(\cdot)$ is the indicator function. We find that $\ln\left(\frac{\lambda_H + \lambda_E}{\lambda_H(1 - (1 - p_H)^d) + \lambda_E}\right) / p_H$ is an approximate solution for $\mathbf{E}[H^{\tilde{\mathcal{C}}(d, \mu)}]$. Further, note that in $\tilde{\mathcal{C}}(d, \mu)$ if an H/E agent arrives while a chain-segment is being formed then the agent will not join the market. However, probability of such an event vanishes as $\mu \rightarrow \infty$ (i.e, forming chain-segments becomes instantaneous). Therefore, it is reasonable to approximate $\mathbf{E}[H^{\mathcal{C}(d)}]$ with $\lim_{\mu \rightarrow \infty} \mathbf{E}[H^{\tilde{\mathcal{C}}(d, \mu)}]$, and thus with $\ln\left(\frac{\lambda_H + \lambda_E}{\lambda_H(1 - (1 - p_H)^d) + \lambda_E}\right) / p_H$. Finally note that:

$$\lim_{p_H \rightarrow 0} \ln\left(\frac{\lambda_H + \lambda_E}{\lambda_H(1 - (1 - p_H)^d) + \lambda_E}\right) = \ln\left(\frac{\lambda_H + \lambda_E}{\lambda_E}\right),$$

which is the constant in Proposition 1.

We emphasize that the analysis above is only a heuristic for guessing the right constant when $p_E < 1$. We refer the reader to Section 4.5 for numerical simulations that show the tightness of this constant.

F Intuition and Another Heuristic for Limit of $w_H^{\mathcal{C}(d)}$

In this section, we provide another heuristic argument that explains why p_E and d do not impact waiting time in the limit.

Let us look at the formation of a chain-segment starting with an E agent; it will have the following structure: $E - H - \dots - H - E - H - \dots - H - E - H - \dots - H$; in other words, the chain-segment can be divided into sub-segments, each starting with an E agent, and following by a number of H agents until the sub-segment gets “stuck”, meaning that there is no H agent that can continue the sub-segment (Recall that priority is given to H agents; therefore we will not search among E agents as long as there is an H agent that can continue the sub-segment). At that stage, the policy will look for an E agent to start a new sub-segment. Let Σ_E be the number of H agents in each sub-segment (between two E agents).

Next, let us look at the formation of a chain-segment starting with an H agent; the chain-segment has the following structure: $H - H - \dots - H - E - H - \dots - H - E - H - \dots - H$, i.e, the first sub-segment starts with an H , but all the subsequent ones start with an E . Let us denote the expected number of H agents in the very first sub-segment by Σ_H .

With the above, let us consider the rate at which H agents join the market and the rate at which they depart. The former is simply $\lambda_H(1-p_H)^d$. The latter is $\lambda_E\mathbf{E}[\Sigma_E]+\lambda_H[1-(1-p_H)^d](\mathbf{E}[\Sigma_H]-1)$. Thus by conservation of H agents, we have:

$$(1-p_H)^d\lambda_H = \lambda_E\mathbf{E}[\Sigma_E] + \lambda_H[1-(1-p_H)^d](\mathbf{E}[\Sigma_H]-1).$$

Note that Proposition 2 - which shows that the expected length of a chain-segment is a constant - implies that $\mathbf{E}[\Sigma_H]$ is a constant, because $\mathbf{E}[\Sigma_H] \leq \mathbf{E}[L|L \geq 1]$, and Proposition 2 states that $\mathbf{E}[L|L \geq 1]$ remains a constant as $p_H \rightarrow 0$. Therefore, we can conclude that:

$$\mathbf{E}[\Sigma_E] = \lambda_H/\lambda_E + o(1),$$

i.e., in the limit $p_H \rightarrow 0$, almost all H agents match through sub-segments initiated by an E , and the length of such sub-segment only depends on the arrival rates. In other words, if we view the number of H agents that an E agent can help match as the “usefulness” of the E , this already suggests that (in the limit) the usefulness of an E does not depend on p_E nor on d .

However, we note that the time at which an E agent initiates a sub-segment changes with p_E :

- When $p_E = 1$, an arriving E agent initiates the first (and only) sub-segment of the chain-segment.
- When $p_E < 1$:
 - With probability $(1 - (1 - p_E)^d)$, the E agent is matched by the bridge agent, and initiates the 1-st sub-segment.
 - With probability $(1 - p_E)^d$, the E agent joins the market, and initiates a sub-segment (that is not the first) at a later time.

Next, we analyze $\mathbf{E}[\Sigma_E]$ and relate it to the expected number of H agents. Recall that we denote the random number of H agents in steady-state by $H^{C(d)}$. Consider the first sub-segment formed upon arrival of an E agent that receives from the bridge agent. We denote this sub-segment by Σ_1 . Observe that

$$\mathbf{E}[\Sigma_1 | H^{C(d)} = h] = \sum_{i=1}^h \prod_{j=0}^{i-1} (1 - (1 - p_H)^{h-j}) := g(H^{C(d)} = h). \quad (21)$$

Note that this equation is the direct consequence of giving priority to H agents. Therefore, $\mathbf{E}[\Sigma_1] = \mathbf{E}[\mathbf{E}[\Sigma_1 | H^{C(d)} = h]] = \mathbf{E}[g(H^{C(d)})]$. Note that $g(H^{C(d)})$ does not depend on p_E or d . We make the following observations:

1. Function $g(\cdot)$ is a monotone increasing function, therefore it has a well-defined inverse function. Further, we have:

Claim 1. For $y > x$, we have $g(y) \geq g(x) + \ln\left(\frac{1}{1-p_H}\right)(y-x)$.

The above claim implies that if $g(y) - g(x) = o(1)$, then $y - x = o(1/p_H)$. The proof of the claim follows the same induction steps as that of Lemma 7; note that function $g(\cdot)$ has the same form as function $f(\cdot)$ in that lemma. Therefore, we omit the proof of the above claim.

2. For p_H small enough, $\mathbf{E}[g(H^{C(d)})] = g(\mathbf{E}[H^{C(d)}]) + o(1)$. Even though we cannot prove this concentration result, in Figure 11, we provide numerical simulation verifying it for different p_E 's.

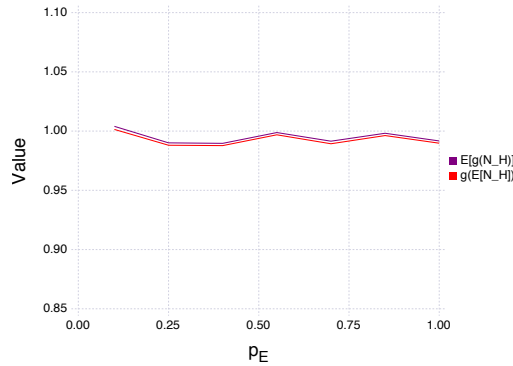


Figure 11: Values of $g(\mathbf{E}[H^{C(d)}])$ and $\mathbf{E}[g(H^{C(d)})]$ as a function of p_E , for $p_H = 0.002$, $\lambda_H = 2$, $\lambda_E = 1$, $d = 1$ and $T = 10^6$.

Finally we claim that for p_H small enough, $\mathbf{E}[\Sigma_E] \approx \mathbf{E}[\Sigma_1]$; this follows from two observations:

1. $\mathbf{E}[\Sigma_1]$ is a constant; to see this, we use Proposition 2 which shows that the expected length of a chain-segment is a constant. Given that $\mathbf{E}[\Sigma_1] \leq \mathbf{E}[L|L \geq 1]$, we can conclude that $\mathbf{E}[\Sigma_1]$ is also a constant (L denotes the chain-segment length as it is defined above Proposition 2). Therefore when the second sub-segment is being formed, the number of H agents in the market will be $H^{C(d)} - \Theta(1) \approx H^{C(d)}$ (remember that $H^{C(d)} = \Theta(1/p_H)$ following the upper-bound of Theorem 3 and the lower-bound of Appendix I).
2. The number of sub-segments is a constant: this follows from the observation that number of E agents in the market is an upper-bound on the number of sub-segments. In Lemma 7, we show that the expected number of E agents is a constant (independent of p_H).

Combining the above observations, for sufficiently small p_H , we have that:

$$g(\mathbf{E}[H^{C(d)}(p_E = 1)]) + o(1) = \lambda_H/\lambda_E + o(1) = g(\mathbf{E}[H^{C(d)}(p_E < 1)]) + o(1), \quad (22)$$

where with a slight abuse of notation we denoted the number of H agents in the market with parameter p_E by $H^{C(d)}(p_E)$. Since $g(\cdot)$ has an inverse and using Claim 1, (22) implies that:

$$H^{C(d)}(p_E = 1) = H^{C(d)}(p_E < 1) + o(1/p_H). \quad (23)$$

We believe that the approximation becomes exact in the limit $p_H \rightarrow 0$. Even though the above argument is not rigorous, we hope that it sheds light on why p_E and d do not impact waiting time in the limit.

G Positive recurrence proofs

To prove existence of a stationary distribution, we use a special case of a result from Meyn and Tweedie (1993), as stated in Prieto-Rumeau and Hernández-Lerma (2016).

Theorem 4 (Meyn and Tweedie (1993)). *Suppose that X_t is an irreducible continuous time Markov*

chain, and suppose that there exist a nonnegative function V on S , a function $w \geq 1$ on S , a finite set $C \subset S$, and constants $c > 0$ and $b \in \mathbb{R}$ such that

$$\sum_{j \in S} q_{ij} V(j) \leq -cw(i) + b \cdot \mathbb{I}_C(i) \text{ for all } i \in S,$$

where \mathbb{I}_C denotes the indicator function of the set C . Then the Markov chain X is ergodic.

It is clear that all four markov chains are irreducible, so our proofs will focus on finding a suitable set C and function V for each case.

G.1 Positive recurrence of \mathcal{B}_H and \mathcal{B}_E .

Claim 2. *The CTMC $[H_t^{\mathcal{B}_H}, E_t^{\mathcal{B}_H}]$ defined in (3a), (3b), (3c), (3d), is positive recurrent.*

Proof. Let $V([h, e]) = h + e$. Observe that for a continuous time random walk, we can write for any state $i = [h, e]$:

$$\begin{aligned} \sum_{j \in \mathbb{N}^2} q_{i,j} V(j) &= \sum_{j \neq i} q_{i,j} (V(j) - V(i)) \\ &= Q([h, e], [h+1, e]) - Q([h, e], [h-1, e]) + Q([h, e], [h, e+1]) - Q([h, e], [h, e-1]). \\ &= \lambda_H (1 - p_H^2)^h (1 - p_{EPH})^e - \lambda_H (1 - (1 - p_H^2)^h) - \lambda_E (1 - (1 - p_{EPH})^h) + \\ &\quad \lambda_E (1 - p_{EPH})^h (1 - p_E^2)^e - \lambda_H (1 - p_H^2)^h (1 - (1 - p_{EPH})^e) - \lambda_E (1 - p_{EPH})^h (1 - (1 - p_E^2)^e) \\ &= 2\lambda_H (1 - p_H^2)^h (1 - p_{EPH})^e - \lambda_H + 2\lambda_E (1 - p_{EPH})^h (1 - p_E^2)^e - \lambda_E. \end{aligned}$$

Take $M(p_H)$ such that $(1 - p_H^2)^{M(p_H)} = \frac{1}{3}$, and let $C = \{[x, y] \text{ s.t. } x + y \leq 2M(p_H)\}$, note that C is finite. For any $[h, e] \notin C$, either $h \geq M(p_H)$ or $e \geq M(p_H)$. In both cases, we have $\sum_{j \in \mathbb{N}^2} q_{i,j} V(j) \leq -\frac{\lambda_H + \lambda_E}{3}$. \square

Claim 3. *The CTMC $[H_t^{\mathcal{B}_E}, E_t^{\mathcal{B}_E}]$ defined in (5a), (5b), (5c), (5d) is irreducible and positive recurrent.*

The proof follows similar ideas as Claim 2.

G.2 Positive recurrence of $\mathcal{C}(d)$ and $\hat{\mathcal{C}}(d)$

Claim 4. *Under $\mathcal{C}(d)$, the number of H and E agents $[H_t^{\mathcal{C}(d)}, E_t^{\mathcal{C}(d)}]$ is a positive recurrent CTMC.*

Proof. Similarly to the proof of Claim 2, let $V([h, e]) = h + e$, and $C = \{[x, y] \text{ s.t. } x + y \leq 2M\}$ for an appropriately chosen M . Consider a state $[h, e] \notin C$, and assume first that $h \geq M$. Denoting $\Lambda = \lambda_H(1 - p_H)^d + \lambda_E(1 - p_E)^d$, we have:

$$\begin{aligned} \sum_{j \in \mathbb{N}^2} q_{i,j} V(j) &= \Lambda - (\lambda_H + \lambda_E - \Lambda) \sum_{[k_H, k_E] \leq [h, e]} (k_H + k_E) \mathbf{P}[k_H(k_E) \text{ H (E) agents get matched}] \\ &\leq \Lambda - (\lambda_H + \lambda_E - \Lambda) \sum_{[k_H] \leq [h, e]} k_H \mathbf{P}[k_H \text{ H agents get matched}] \\ &\leq \Lambda - (\lambda_H + \lambda_E - \Lambda) \sum_{k \leq h} \prod_{i=h}^{h-k} (1 - (1 - p_H)^i) \\ &\leq \Lambda - (\lambda_H + \lambda_E - \Lambda) \sum_{k \leq M/2} (1 - (1 - p_H)^{M/2})^k \end{aligned}$$

Going from the second to the third inequality bounds the number of agents matched with the number of agents matched in the first sub-chain-segment (before matching an E agent). Because the function $M \rightarrow \sum_{k \leq M/2} (1 - (1 - p_H)^{M/2})^k$ tends to infinity as M grows large, we can find M such that $\sum_{j \in \mathbb{N}^2} q_{i,j} V(j) \leq -1$, which concludes the proof.

The case where $h < M$ and $e \geq M$ can be treated similarly using the following observation: every-time the chain cannot match to any H agent, it tries to match an E , and it succeeds with probability $1 - (1 - p_E)^e$ irrespective of h . Put together, these events constitute a set of E agents that are matched sequentially, and the total length of the chain is more than the number of E agents selected in that way. \square

Claim 5. *$H_t^{\hat{\mathcal{C}}(d)}$ is a positive recurrent CTMC.*

The proof follows similar arguments as that of Claim 4.

H Lower Bound

Definition 3. We call a matching policy anonymous Markovian if matching decisions are made at arrival epochs, and only depend on the current compatibility graph $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$, and are anonymous among agents of the same type. In that case, the market \mathcal{G}_t is a continuous-time Markov chain.

We say that an anonymous Markovian policy \mathcal{P} is stable if the resulting Markov chain is ergodic and has a stationary distribution. In line with our previous notation, we denote $w_H^{\mathcal{P}}$ ($w_E^{\mathcal{P}}$) the expected stationary waiting times for H and E agents under policy \mathcal{P} .

Remark. Observe that all the matching policies described in this paper are anonymous Markovian.

Proposition 9. For any $p_E \in [0, 1]$, $\lambda_H > 0, \lambda_E \geq 0$, there exists a constant c such that for any $p_H > 0$, under any stable anonymous Markovian matching policy \mathcal{P} , $w_H^{\mathcal{P}} + w_E^{\mathcal{P}} \geq \frac{c}{p_H}$.

The proof follows ideas used in Anderson et al. (2017). The main intuition is the following: Suppose the market size is too small, then an arriving agent has to wait a *long* time to obtain at least one incoming edge. This long waiting time contradicts the small market size (with Little's law).

Proof. In this proof, we fix a Markovian policy \mathcal{P} , and we will omit the superscript notations. Observe that under \mathcal{P} , the market \mathcal{G}_t only evolves at arrival epochs, and we can analyze the embedded discrete-time Markov chain resulting from observing the system at arrival epochs which we denote by \mathcal{G}_i . Let $n = \mathbf{E}[|\mathcal{V}|] = \mathbf{E}[H] + \mathbf{E}[E]$ be the expected size of the market in steady-state. Let us denote $\theta = \frac{\lambda_H}{\lambda_H + \lambda_E}$, and let \hat{w}_H be the expected number of (discrete) time steps that an H agent arriving in steady-state has to wait before getting matched. Little's law for discrete Markov chains implies that $\hat{w}_H = \mathbf{E}[H]/\theta$.

Note that it is enough to prove that there exists a constant k such that $n \geq k/p_H$ (we then choose $c = \frac{k}{\max(\lambda_H, \lambda_E)}$). Let k be a constant to be defined later. Assume for contradiction that there exists p_H such that $n < k/p_H$. Let a be an H agent entering the market in steady-state. Let \mathcal{V}_a be the set of agents in the market when agent a arrives. Note that we assumed $\mathbf{E}[|\mathcal{V}_a|] = n < k/p_H$. Define the event $E_1 = \{|\mathcal{V}_a| \leq 3n/\theta\}$. By Markov's inequality, $\mathbb{P}[E_1] \geq 1 - \frac{\mathbf{E}[|\mathcal{V}_a|]\theta}{3n} \geq 1 - \theta/3$.

Let \mathcal{A}_a be the first $3n/\theta$ arrivals after a , and let E_2 be the event that at least one agent from $\mathcal{V}_a \cup \mathcal{A}_a$ has an outgoing edge towards a . We have

$$\mathbb{P}[E_2] = \mathbb{P}[\mathbf{Bin}(|\mathcal{V}_a| + |\mathcal{A}_a|, p_H) \geq 1].^{40}$$

Therefore we get:

$$\mathbb{P}[E_2|E_1] \leq \mathbb{P}[\mathbf{Bin}(6n/\theta, p_H) \geq 1] \leq \mathbb{P}[\text{Bin}(6k/\theta p_H, p_H) \geq 1] \leq 6k/\theta.$$

Where the first inequality derives from the definition of E_1 , the second uses the fact that $n \leq k/p_H$ and the third is Markov's inequality.

We now use the fact that if a does not have any edge from either \mathcal{V}_a or \mathcal{A}_a , then she must wait longer than $3n/\theta$ time steps.

$$\hat{w}_H \geq \frac{3n}{\theta} \mathbb{P}[E_2^c] \geq \frac{3n}{\theta} \mathbb{P}[E_2^c|E_1] \mathbb{P}[E_1] \geq \frac{3n}{\theta} (1 - 6k/\theta)(1 - \theta/3) \geq \frac{3n}{\theta} (1 - 6k/\theta)(2/3).$$

Thus we get:

$$n \geq \mathbf{E}[H] = \hat{w}_H \theta \geq 2n(1 - \frac{6k}{\theta}).$$

Therefore for $k = \frac{\theta}{24}$, we obtain a contradiction. □

Observe that similar to Anderson et al. (2017), the above reasoning could be generalized to the case of *periodic* Markovian policies (see Definition 2 of Anderson et al. (2017)) which includes batching policies with constant batch size.

⁴⁰Note that we abuse the notation of $\mathbf{Bin}(\mathbf{u}, \mathbf{p})$ by allowing its parameters to be random variables. In this case, conditional on the event $|\mathcal{V}_a| + |\mathcal{A}_a| = u$, the random variable $\mathbf{Bin}(|\mathcal{V}_a| + |\mathcal{A}_a|, p)$ has a binomial distribution with parameters u and p .

I A Lower Bound on $w_H^{C(d)}$

Here, we apply Proposition 9 to the $ChainMatch(d)$ policy (which is a stable anonymous Markovian matching policy) to prove a lower bound on $w_H^{C(d)}$:

$$w_H^{C(d)} + w_E^{C(d)} \geq \frac{c}{p_H}. \quad (24)$$

In the following lemma, we show that $w_E^{C(d)} < k_E$, where k_E is a constant independent of p_H ; therefore (24) implies that $w_H^{C(d)} \geq \frac{c}{p_H} - k_E$, i.e., there exists a lower-bound on $w_H^{C(d)}$ that scales with $1/p_H$.

Lemma 7. *Under $ChainMatch(d)$ and for $0 < p_E \leq 1$, we have that $w_E^{C(d)} < k_E$, where k_E is a constant independent of p_H .*

Proof. Let $[H^{C(d)}, E^{C(d)}]$ denote is the random number of H and E agents in steady-state. First observe that conditional on $E^{C(d)} = x$, the expected number of E agents matched in a chain-segment is given by:

$$f(x) := \sum_{k=0}^x \prod_{i=0}^k (1 - (1 - p_E)^{x-i}).$$

Note that the above holds because under $ChainMatch(d)$, we give priority to H agents; therefore, when computing the probability E agents in a chain-segment, we can ignore the H agents that are matched between two consecutive E agents.

Function $f(x)$ can be re-written recursively as $f(x) = (1 - (1 - p_E)^x)(1 + f(x - 1))$. We now wish to provide a lower bound of $f(x)$. Let α be a constant to be determined later, and let us assume by induction that $f(x - 1) \geq \alpha(x - 1)$. Then:

$$f(x) \geq (1 - (1 - p_E)^x)(1 + \alpha(x - 1)) = \alpha(x - 1) + 1 - (1 - p_E)^x(1 + \alpha(x - 1)).$$

Let $\xi_\alpha(y) = (1 - p_E)^y(1 + \alpha(y - 1))$. Its derivative satisfies

$$\xi'_\alpha(y) = (1 - p_E)^y (\ln(1 - p_E)(1 + \alpha(y - 1)) + \alpha).$$

Note that in $\xi_\alpha(y)$, y is a continuous variable. Let $\alpha = \ln(\frac{1}{1-p_E})$, for all $y \geq 1$, $\xi'_\alpha(y) \leq 0$. Further, $\xi_\alpha(1) \leq \xi_\alpha(0)$. Thus we have $\xi_\alpha(x) \leq \xi_\alpha(0)$, for $x \in \mathbb{N}$. Therefore we can complete our induction.

$$f(x) \geq \alpha(x-1) + 1 - \xi_\alpha(x) \geq \ln\left(\frac{1}{1-p_E}\right)x.$$

We can now conclude our proof using conservation of E agents in the market:

$$\lambda_E(1-p_E)^d = (\lambda_E(1-(1-p_E)^d) + \lambda_H(1-(1-p_H)^d))\mathbf{E}[f(E^{C(d)})] \geq \lambda_E(1-(1-p_E)^d) \ln\left(\frac{1}{1-p_E}\right)\mathbf{E}[E^{C(d)}]$$

where in the last inequality we use $f(x) \geq \ln(\frac{1}{1-p_E})x$ and $(1 - (1 - p_H)^d) \geq 0$. We take $k_E = \frac{(1-p_E)^d}{(1-(1-p_E)^d)\ln(\frac{1}{1-p_E})}$, which does not depend on p_H . \square

J Scaling of $w_H^{\mathcal{B}_H}$ and $w_H^{\mathcal{B}_E}$ when $\lambda_H = \lambda_E$

In Figure 12, we present our numerical study on how $w_H^{\mathcal{B}_H}$ and $w_H^{\mathcal{B}_E}$ scale with p_H in the case that $\lambda_H = \lambda_E$; the figure plots $p_H w_H^{\mathcal{B}_H}$ and $p_H w_H^{\mathcal{B}_E}$ when p_H ranges from $10^{-3.5}$ to 10^{-7} while $p_E = 0.5$. As the plot shows, for p_H smaller than 10^{-5} , the normalized waiting time $p_H w_H^{\mathcal{B}_H}$ and $p_H w_H^{\mathcal{B}_E}$ both remain constant, which implies the asymptotic scaling is $1/p_H$ under both policies.




Figure 12: Normalized waiting time $w_H \rho_H$ as a function of p_H (x axis is in log scale), for both \mathcal{B}_H , and \mathcal{B}_E . $T = 10^{12}$, $\lambda_H = \lambda_E = 1$, $p_E = 0.5$.