Matching in Dynamic Imbalanced Markets

Itai Ashlagi Afshin Nikzad Philipp Strack*

Abstract

We study dynamic matching in exchange markets with easy- and hard-to-match agents. A greedy policy, which attempts to match agents upon arrival, ignores the positive externality that waiting agents provide by facilitating future matchings. We prove that the trade-off between a "thicker" market and faster matching vanishes in large markets; the greedy policy leads to shorter waiting times and more agents matched than any other policy. We empirically confirm these findings in data from the National Kidney Registry. Greedy matching achieves as many transplants as commonly-used policies (1.8% more than monthly batching), and shorter waiting times (16 days faster than monthly batching).

*We want to thank Mohammad Akbarpour, Aaron Bodoh-Creed, Yeon-Koo Che, Yuichiro Kamada, and Olivier Tercieux for useful comments and suggestions. Ashlagi: MS&E, Stanford, iashlagi@stanford.edu. Nikzad: Economics, University of Southern California, afshin.nikzad@usc.edu. Strack: Economics, Yale University, philipp.strack@gmail.com. Ashlagi acknowledges the research support of the National Science Foundation grant (SES-1254768). Philipp Strack was supported by a Sloan fellowship.

1 Introduction

We study how to optimally match agents in a dynamic random exchange market. Faster matching of agents reduces waiting times, but at the same time makes the market thinner, leaving more agents without a compatible partner. This trade-off naturally arises for kidney exchange platforms that seek to form exchanges between patient-donor pairs, whose patient can not receive the donor's organ.[1] Waiting to match may increase the number of patients receiving a kidney, but this comes at a cost: receiving a transplant earlier not only improves the quality of life for the patient but also leads to substantial savings in dialysis costs for society.[2] In the last decade kidney exchange platforms in the United States gradually moved from matching roughly every month to matching daily.[3] Practitioners are concerned that this behavior, some of which is driven by competition between kidney exchanges, is harmful, especially for the most highly sensitized patients. In contrast, kidney exchange programs in Canada, Australia, and the Netherlands match periodically every 3 or 4 months (Ferrari et al., 2014).

This paper analyzes the trade-off between agents' waiting times and the percentage of matched agents in dynamic markets. We find that, maybe surprisingly, matching greedily minimizes the waiting time and simultaneously maximizes the chances to find a compatible partner for *all* agents in sufficiently large markets. We further quantify the inefficiency associated with other commonly used policies like monthly matching using data from the National Kidney Registry (NKR).

To analyze this question we propose a stochastic compatibility model with easy-to-match and hard-to-match agents. Easy-to-match agents can match with each other with probability $q > 0$ and with hard-to-match-agents with probability $p > 0$, whereas hard-to-match agents can match only with easy-to-match agents with probability $p > 0$. The main focus of our analysis is on the case where the majority of agents are hard-to-match, which is in line with kidney exchange pools. This compatibility model captures two empirical regularities of the patient-donor data from the NKR. First, as the market grows large, the fraction of patient-donor pairs that are matched in a maximal matching does not approach 1, which is a consequence of the imbalance between different pairs' blood types in kidney exchange (Saidman et al., 2006; Roth et al., 2007). Second, as the market grows large, the fraction of agents that cannot be matched in *any* matching goes to zero.[4] Our parsimonious two-type model captures the above regularities and no single-type model can account for both of them (Propositions 1 and 2).

We study a dynamic model based on the two-type compatibility structure in which easy- and

[1]For some early work on kidney exchange in static pools and the importance of creating a thick marketplace, see Roth et al. (2007, 2004).

[2]The savings from a transplant over dialysis is estimated by over \$270,000 per year over the first five years (Held et al., 2016).

[3]The National Kidney Registry (NKR) and the Alliance for Paired Donation (APD) search for matches on a daily basis, whereas the United Network for Organ Sharing (UNOS) searches for matches twice per week.

[4]A patient-donor pair cannot be matched in any matching if it cannot form a (two-way) exchange with any other patient-donor pair due to biological compatibility.

hard-to-match agents arrive to the market according to independent Poisson processes with rates m_E and m_H. Agents depart exogenously at rate d. The market-maker observes the realized compatibilities and decides when to match compatible agents. We evaluate a policy based on two measures: *match rate* and *waiting time*. The match rate is the probability with which an agent is matched. The waiting time is the average difference between the time an agent arrives and the time she leaves, matched or unmatched. Our two-type model captures the potential trade-off between match rates and waiting times that concerns practitioners in an intuitive way: matching quickly could lead to easy-to-match agents being paired with each other thereby making it more difficult for hard-to-match agents to find a partner and thus potentially decreasing the overall match rate.

We start by analyzing the *greedy policy*, which matches every agent upon its arrival if possible. We derive the distribution of the number of hard- and easy-to-match agents waiting in the market in steady state. As the market grows large, many hard-to-match agents will wait in the market for a compatible partner at any point in time. Consequently, almost every easy-to-match agent is matched with a hard-to-match agent immediately upon arrival and the probability that an easy-to-match agent leaves the market unmatched converges to zero (Proposition 3). As their match rate is close to one and their waiting time is close to zero, the greedy policy is asymptotically optimal for easy-to-match agents in large markets. As hard-to-match agents are incompatible with each other and almost every easy-to-match agent is matched with a hard-to-match agent, the greedy policy also maximizes the match rate of hard-to-match agents. Perhaps less intuitively, the greedy policy minimizes the waiting time of hard-to-match agents compared to *any* other policy when the market grows large (Proposition 4). This holds since the greedy policy matches weakly more hard-to-match agents than any other policy. Little's law implies that the average number of hard-to-match agents waiting in the market is proportional to their waiting time and thus that the greedy policy will perform weakly better than any other policy in a sufficiently large market.

The main challenge in the proof is analyzing the steady state distribution of a two-dimensional Markov chain which keeps track of the number of easy- and hard-to-match agents waiting in the market. Using the Lyapunov function method, we show that the stationary distribution is concentrated around the solution of a fixed-point equation that describes the average numbers of easy- and hard-to-match agents waiting in the pool. These concentration bounds allow us to compute the agents' match rate and waiting time.

Next, we quantify the inefficiency associated with batching policies, which are commonly used. A batching policy periodically (e.g., monthly) matches as many agents as possible. We derive the waiting time and match rate under batching policies in large markets. Batching less frequently decreases the match rate and increases the waiting time. Therefore, in a large market, greedy matching dominates any batching policy with a fixed batch length, as it leads to *strictly shorter* waiting times and *strictly higher* match rates. We also compare batching and greedy matching in finite markets, where batching may outperform greedy matching. We find that for parameters in line with our kidney exchange application batching policies need to match more frequently than daily to not be outperformed by greedy matching.

We also analyze the *patient policy* introduced by Akbarpour et al. (2020). This policy assumes that agents' exogenous departure times are observable. It matches an agent at her departure time if possible, and otherwise the agent leaves the market unmatched. We show that the patient policy leads to the same match rate as the greedy policy when the market becomes large. In both policies almost all easy-to-match agents are matched almost upon arrival in a large market. Moreover, hard-to-match agents wait longer (in first order stochastic dominance) under the patient policy compared to the greedy policy. Quantitatively, the waiting time of hard-to-match agents under the greedy policy equals the waiting time under the patient policy multiplied by $\left(1 - \frac{m_E}{m_H}\right)$ where m_E, m_H are the arrival rates of easy- and hard-to-match agents. For example, when a third of agents are easy-to-match ($2m_E = m_H$), hard-to-match agents will wait twice as long under the patient policy.

Finally, we test whether the large-market predictions of our model hold in data from the NKR. This data differs from our assumptions along two dimensions: first, because of blood and tissue types it does not match our stylized two-type compatibility structure. Second, it is unclear that the market is sufficiently large for our results to apply, because only a finite number of agents arrive every year (around 360/year). Nevertheless, the data confirms the predictions of our model (Section 5): As the market becomes large, the waiting times of patient-donor pairs who are "easier" to match approach 0, but the waiting times of "harder" to match pairs do not. Moreover, batching policies result in no improvement to the match rate and lead to longer waiting times relative to greedy matching (c.f. Table 1). Finally, waiting times under the greedy policy are significantly lower than under the patient policy. At the same time, we do not find significant differences between the match rates under greedy and patient policies (Table 1 and Figure 9).

As mentioned earlier, practitioners expressed concerns that kidney exchange platforms in the US are adopting greedy algorithms, arguably due to competition.[5] That is, matching greedily can potentially squeeze the liquidity generated by easy-to-match pairs in an inefficient manner. Our theory combined with simulations suggest that matching greedily is not a real source of inefficiency.

1.1 Related literature

Closely related literature studies dynamic matching on networks when agents' preferences are based on compatibilities, motivated by kidney exchanges. This literature, initiated by Ünver (2010), can be organized into two perspectives. The first perspective seeks to minimize waiting times in models where agent do not depart exogenously (Ünver, 2010; Anderson et al., 2017; Ashlagi et al., 2016). The key finding in this literature is that greedy matching minimizes the average waiting time.

[5]See Ashlagi et al. (2018) who write: "..competition among KPD programs to produce transplants may have incentivized programs to perform match-runs at high frequency, which raises a major concern that such frequent matching may lead to fewer transplants." Gentry and Segev (2015) raise similar concerns motivated by match failures: "...registries forced by competition to perform match runs very frequently cannot take advantage of mathematical optimization, and likely fewer transplants are accomplished nationwide as a result."

The second perspective is concerned with how many agents are matched. Akbarpour et al. (2020) consider a model with exogenous departures, in which each agent is compatible with any other agent with a fixed probability. They find that the patient policy leads to an exponentially smaller loss rate (i.e., fraction of unmatched agents) compared to the greedy policy.[6]

Each of these perspectives studies one of two objectives: minimizing the time until an agent is matched, or minimizing the number of agents that leave the market unmatched. The two perspectives lead to different conclusions about the optimality of the greedy policy and suggest a trade-off between matching agents quickly and matching as many agents as possible. This paper contributes by studying this trade-off and showing that it vanishes in large kidney-exchange markets with asymmetric agents. Technically, our paper is the first to analyze a model with both exogenous departures and heterogeneous agents. We further contribute by analyzing the distribution of waiting times rather than just averages and by providing an analysis of finite markets.

The effectiveness of thickening the market by waiting to increase the number of matches has also been studied in markets other than kidney exchanges. Liu et al. (2018) compare the match rates of greedy and patient policies in ride-sharing markets for matching drivers with passengers and find that the waiting increases market thickness and average match quality but decreases the number of matches. Finally, recent and indirectly related papers study dynamic matching when agents have preferences beyond compatibility and find that greedy policies are sometimes inefficient since some waiting can improve the quality of matches (Baccara et al., 2019; Doval, 2014; Mertikopoulos et al., 2019; Ashlagi et al., 2019; Blanchet et al., 2020).

2 The Compatibility Graph

A kidney exchange pool can be represented by a *compatibility graph* G. Each node in the graph represents an agent (a patient-donor pair), and a link between two nodes exists if and only if the two corresponding agents are *compatible* with each other (so a bilateral exchange between the nodes is feasible). We restrict attention to bilateral exchanges.

A *matching* μ is a set of non-overlapping compatible pairs of agents. Denote by $M(G)$ the set of matchings in G.[7] For every compatibility graph G let $|G|$ denote the number of agents in the graph, and for every matching μ let $|\mu|$ denote the number of agents in that matching.

Define the *(normalized) size of the maximum matching* (SMM) in a graph G to be the fraction of matched agents in a maximum matching:

$$\text{SMM} = \max_{\mu \in M(G)} \frac{|\mu|}{|G|}\,.$$

[6]The differences with Akbarpour et al. (2020) are discussed in detail in Section 6.

[7]The paper restricts attention to matching only pairs of agents and not through chains. For the effect of matching through chains see, e.g., Ashlagi et al. (2011) and Anderson et al. (2017).

Define the *fraction of agents without a partner* (FWP) to be the fraction of agents that are not matched in any matching (thus have no compatible agent):

$$\text{FWP} = \frac{|\{i \in G \colon (i,j) \notin M(G) \text{ for all } j\}|}{|G|}.$$

Figure 1 depicts the expected SMM and FWP for a subset of a given size drawn uniformly at random from the patient-donor population acquired from the National Kidney Registry (NKR), the Alliance for Paired Donation (APD), and the United Network for Organ Sharing (UNOS) and Methodist at San Antonio. This data includes 4992 patient-donor pairs.[8] The data allows to verify whether each patient and donor are virtually compatible, and therefore whether two pairs can match. Two features stand out: first, even as the market grows large, the size of the maximum matching stays bounded away from 1, i.e., SMM < 1. This is a natural consequence of the different blood types (Roth et al., 2007). Second, when the market grows large, the fraction of pairs that have no compatible pairs decreases. Roughly 5.5% of pairs are incompatible with any other pair in this data (FWP).[9] Since compatibility depends only on the characteristics of the patients and donors, it is independent of pool size, and thus in a sufficiently large pool one would expect that the FWP would further decrease to zero.

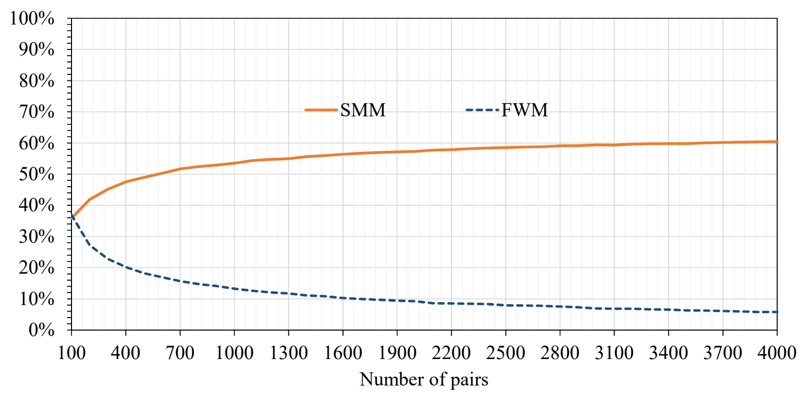

Figure 1: Average percentage of pairs without a compatible partner (dashed) and the percentage matched in a maximum matching (solid). The average for every fixed pool size on the horizontal axis is computed by random sampling from the combined data set from NKR, APD, UNOS and Methodist at San Antonio.

Fact 1. *As the kidney exchange patient-donor pool grows large, the compatibility graph (Figure 1) is such that the size of the maximal matching (SMM) stays bounded away from* 1 *and the fraction*

[8]Each data set includes pairs from a different period of time, but no earlier than 2007. The data from NKR, APD and Methodist was obtained directly and is not publicly available. The data from UNOS can be obtained directly from the Organ Procurement and Transplantation Network (OPTN), a contractor for the US Department of Health and Human Services.

[9]In practice some patients can receive a kidney from blood-type incompatible donors due to advanced technology. For the sake of simplicity we ignore this in our simulations, but it is worth noting that the FWP drops to less than 3% when this form of compatibility is allowed.

of patient-donor pairs without a compatible partner (FWP) becomes small.

The change in both the SMM and FWP measures captures the benefit of a larger market. Since a matching policy in a dynamic environment trades off the benefits of a larger market with the waiting costs incurred by the agents, having a model that accurately represents the SMM and the FWP is important to correctly describe the costs and benefits of waiting to match.

2.1 A Two-Type Compatibility Model

To capture the features of kidney exchange identified in Fact 1 we adopt a stylized and tractable model with random compatibilities. There are two types of agents, *easy-to-match* or *hard-to-match*, denoted by E and H, respectively. There are more hard-to-match than easy-to-match agents. Any pair of hard-to-match and easy-to-match agents are compatible independently with probability $p > 0$, any pair of easy-to-match agents are compatible independently with probability $q > 0$, and no pair of hard-to-match agents are compatible with each other (Figure 2).

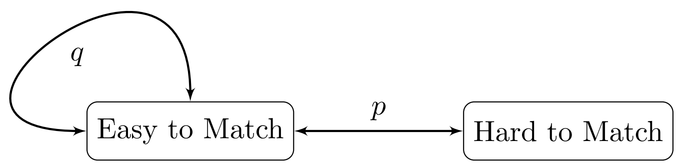

Figure 2: The random compatibility model.

Proposition 1 shows that this simple model is indeed able to capture the features of real kidney exchanges identified in Fact 1.

Proposition 1. *Consider a compatibility graph with m easy-to-match agents and $(1+\lambda)m$ hard-to-match agents where $\lambda > 0$. Compatibilities between pairs of agents are generated as described in Section 2.1. Then, as m grows large* $\text{SMM} = \frac{2}{2+\lambda}$ *and* $\text{FWP} = 0$ *hold with high probability.*[10]

That the size of the maximal matching cannot exceed $\frac{2}{2+\lambda}$ is intuitive: since H agents cannot match with each other and there are more H agents than E agents, some H agents must remain unmatched when the pool is large. An upper bound on the fraction of agents that can be matched equals twice the fraction of E agents $\frac{1}{2+\lambda}$. Furthermore, note that this fraction is achieved whenever there exists a matching in which all E agents are matched with H agents. It follows from a standard result in random graph theory that the probability that such a "perfect matching" exists approaches 1 as the pool grows large. Furthermore, as the pool grows large any H agent will be compatible with some E agent, since compatibilities between agents are drawn independently. Thus, the fraction of agents without a partner converges to 0. The parameter λ of the model measures the degree of imbalance between hard- and easy-to-match agents. So $\lambda = 0$ corresponds to a balanced pool.

[10]We say a sequence of events $E_1, E_2, \ldots$ hold with high probability when $\lim_{k\to\infty} \mathbb{P}[E_k] = 1$.

Figure 1 suggests that the size of the maximal matching in the data converges to roughly 60% when the pool becomes large, implying that $\lambda \approx 1.33$ in the context of our model.

Proposition 1 establishes that our two-type model can match the empirical behavior of the SMM and FWP measures observed in Fact 1. Proposition 2 establishes that no model with a single type can replicate the empirical features of real kidney exchanges that the size of the maximal matching is less than one while the fraction of agents without a partner goes to zero, even when allowing the probability of compatibility between two agents to depend on the market size in arbitrary ways.

Proposition 2. *Consider a model with m homogeneous agents, in which every pair of agents are compatible independently with probability $p(m) > 0$ that may depend on the market size. The following two conditions cannot be satisfied simultaneously:*

$$\lim_{m\to\infty} \mathbb{E}\left[\text{SMM}\right] < 1, \textit{and} \tag{1}$$

$$\lim_{m\to\infty} \mathbb{E}\left[\text{FWP}\right] = 0\,. \tag{2}$$

The proof is constructive. It begins with assuming that every agent has a compatible partner when the pool grows large, i.e., (2) is satisfied. It then introduces an algorithm which constructs a matching that includes all agents with high probability as the pool grows large. This implies that (1) and (2) cannot hold together in *any* random graph model with homogeneous agents.

Economically, this observation implies that heterogeneity of agents plays a major role in kidney exchanges.[11] Our two-type model is arguably the simplest random compatibility model that captures these features of the compatibility graph. It may be useful to illustrate the connection with kidney exchanges while restricting attention to blood-type compatibilities. One may think of A-O and O-A patient-donor pairs as easy- and hard-to-match, respectively. More generally patient-donor pairs with blood types X-O for X≠O, AB-A, and AB-B are blood-type compatible with a pair of the same category and can be considered as easy-to-match, whereas those with blood types O-X for X≠O, A-AB, and B-AB are not, and hence can be considered as hard-to-match. Typical exchange pools have fewer easy-to-match than hard-to-match pairs simply because a patient who is compatible with her intended donor will be transplanted directly. For example, an A patient and her intended O donor may be compatible and would not join the exchange (they only need to join if they are tissue-type incompatible).

3 Dynamic Matching

We embed the static compatibility model from Section 2.1 in a dynamic model that allows to study matching policies in a dynamic setting. We consider an infinite-horizon dynamic model, in which agents can match bilaterally. Easy-to-match agents arrive to the market according to a Poisson

[11] This is consistent with Roth et al. (2007) and Agarwal et al. (2019), who demonstrate that the types of patients and donors play a crucial role for efficiency.

process with rate m, and hard-to-match agents arrive to the market according to an independent Poisson process with rate $(1 + \lambda)m$. We assume that the majority of agents are hard-to-match, that is $\lambda > 0$, unless explicitly stated otherwise. We sometimes refer to m as the *market size*.

An agent that arrives to the market at time t becomes *critical* after Z units of time in the market, where Z is distributed exponentially with mean d, independently between agents. We refer to $1/d$ as the *criticality rate*. The latest an agent can match is the time she becomes critical, $t + Z$; immediately after this time the agent leaves the market unmatched. The planner observes when an agent gets critical and can attempt to match the agent immediately at that time before she departs.[12]

Matching policies. Denote by G_t the compatibility graph induced by the agents that are present at time t. A *dynamic matching policy* selects at any time t a matching $\mu_t \in M(G_t)$, which may be empty. Whenever a non-empty matching is selected, all matched agents leave the market.

Several kidney exchange platforms in the U.S. attempt to match pairs as soon as they arrive to the market (Ashlagi et al., 2018). A tractable approximation of this behavior is a greedy policy.

Definition 1 (Greedy)**.** *In the* greedy *policy an agent is matched upon arrival with a compatible agent if such an agent exists. If she is compatible with more than one agent, H agents are prioritized over E agents and otherwise ties are broken randomly.*

Some platforms identify matches only periodically, allowing the pool to thicken and possibly offer more matching opportunities. For example, UNOS matches twice a week, whereas national platforms in the United Kingdom and the Netherlands identify matches every three months (Biro et al., 2017). This behavior is approximated with the following batching policy.

Definition 2 (Batching)**.** *A* batching *policy executes a maximal match every T units of time. If there are multiple maximal matches, select randomly one that maximizes the number of matched H agents. The parameter T is called the* batch length.[13]

We also consider the patient policy, proposed by Akbarpour et al. (2020), which attempts to match an agent only once she becomes critical. In the context of kidney exchange this means that two patient-donor pairs in the pool are matched only if the condition of one of these pairs is such that it cannot match at a later point in time.

Definition 3 (Patient)**.** *In the* patient *policy an agent that becomes critical is matched with a compatible agent if one exists. If she is compatible with more than one agent, H agents are prioritized over E agents, and otherwise ties are broken randomly.*

[12] As we will show that the asymptotically optimal policy does not condition on this information, thus allowing for this larger set of policies strengthens our results.

[13] We note that every agent leaves the market either because of being matched to another agent or because of getting critical. Thus, in the batching policy, (only) the agents who are in an executed matching are removed from the market at the time the matching is executed.

The patient policy can be viewed as a theoretical benchmark, as predicting the time at which the patient will become too sick to transplant is generally not feasible. Observe that the greedy and patient policies match at most two agents at any given time because no two agents ever arrive or become critical at the same time. The batching policy, however, can match multiple agents when it executes a matching.

Measures for performance. Denote by $\theta_i \in \{E, H\}$ agent i's type, by $\alpha_i \geq 0$ her arrival time, by $\varphi_i \geq 0$ how long she is present in the market, and indicate by $\mu_i \in \{0, 1\}$ whether she is matched. To study the performance of a matching policy we focus on two measures. One is the *match rate* $q_\Theta(m) \in [0, 1]$ of each type $\Theta \in \{E, H\}$, which is the fraction of agents of type Θ that get matched. Formally, we define the match rates for each arrival rate m by

$$q_\Theta(m) = \lim_{t\to\infty} \mathbb{E}\left[\frac{|\{i\colon \mu_i = 1 \text{ and } \alpha_i \leq t \text{ and } \theta_i = \Theta\}|}{|\{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta\}|}\right].$$

The other is the *expected waiting time* (or simply *waiting time*) $w_\Theta(m)$ of agents of type Θ, whether eventually matched or not. Formally, we define the waiting time for each arrival rate m by

$$w_\Theta(m) = \lim_{t\to\infty} \mathbb{E}\left[\frac{\sum_{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta} \varphi_i}{|\{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta\}|}\right].$$

One reason for studying match rate and waiting time is that they together determine the payoff of a risk-neutral expected-utility-maximizer (EU) who assigns a fixed value to being matched and incurs a constant cost while waiting in the market.

We are interested in optimal policies for large pools and denote by $q_\Theta = \lim_{m\to\infty} q_\Theta(m)$ and $w_\Theta = \lim_{m\to\infty} w_\Theta(m)$ the match rate and waiting time when the market becomes (infinitely) large, i.e., when the arrival rate m goes to infinity.[14] We consider the following notion of optimality:

Definition 4 (Asymptotic optimality)**.** *A policy is asymptotically optimal if for every $\epsilon > 0$ there exists m_ϵ such that, when $m \geq m_\epsilon$,* no *type of agent can improve its match rate $q_\Theta(m)$ or expected waiting time $w_\Theta(m)$ by more than ϵ when changing to any other policy.*

This optimality notion is demanding, since it requires the policy to be optimal for every type of agent simultaneously. It is unclear whether an asymptotically optimal policy exists, since a policy that is optimal for H agents might be suboptimal for E agents.

3.1 Results

In this section we present a characterization of the match rates and waiting times associated with the greedy, batching and patient matching policies and discuss its implications.

[14] Throughout, we restrict attention to policies where these limits are well defined. This assumption could be relaxed by considering the lim inf and lim sup in the definitions.

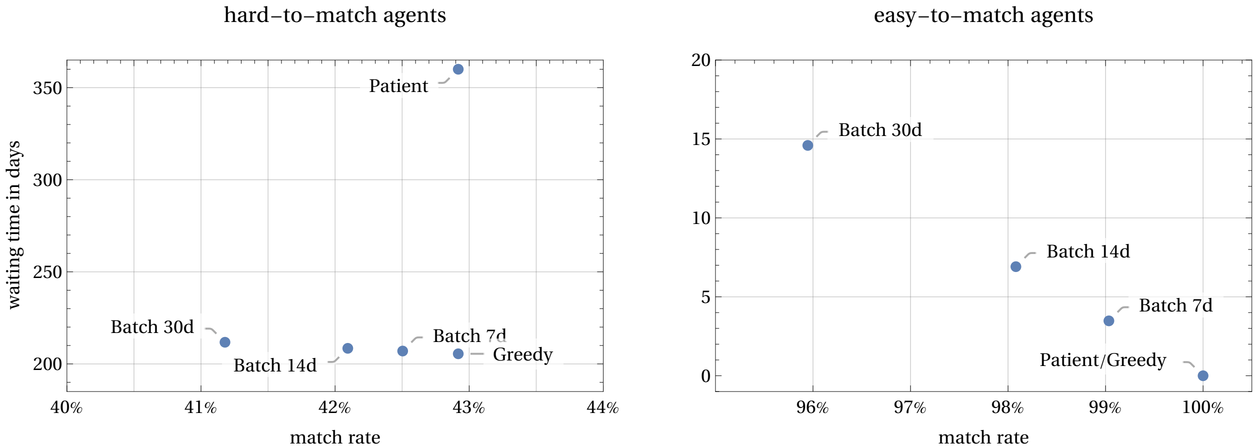

Figure 3: Illustration of Proposition 3 when $\lambda = 1.33$ and d equals 360 days. The blue points represent the predictions of our model for large markets, which are derived in Proposition 3.

Theorem 1. *The greedy policy is asymptotically optimal, whereas the batching policy (for any fixed batch length) and the patient policy are not asymptotically optimal.*

We further compute the match rates and expected waiting times under these policies.

Proposition 3. *As the arrival rate m grows large:*

(i) The match rates of hard- and easy-to-match pairs under the greedy policy are $(q_H^G, q_E^G) = (\frac{1}{1+\lambda}, 1)$, respectively, and their expected waiting times are $(w_H^G, w_E^G) = (\frac{\lambda d}{1+\lambda}, 0)$.

(ii) A batching policy with batch length T achieves match rates of $(q_H^B, q_E^B) = (\frac{1-e^{-T/d}}{(1+\lambda)T/d}, \frac{1-e^{-T/d}}{T/d})$. Furthermore, the expected waiting time for each type Θ is $w_\Theta^B = d(1 - q_\Theta)$. Also, $q_\Theta^B < q_\Theta^G$, whereas q_Θ^B approaches q_Θ^G as T approaches 0. In addition, $w_\Theta^B > w_\Theta^G$, whereas w_Θ^B approaches w_Θ^G as T approaches 0.

(iii) The match rates of hard- and easy-to-match pairs under the patient policy approach $(q_H^P, q_E^P) = (\frac{1}{1+\lambda}, 1)$, respectively, and their expected waiting times approach d and 0, respectively.

Figure 3 illustrates the match rates and waiting times of H and E agents under the different policies as found in Proposition 3. In the figure, the values for λ, d are chosen to match the imbalance and criticality rate in the NKR data ($\lambda = 1.33, d = 360$). As Figure 3 illustrates, each batching policy leads agents to wait longer and get matched with a smaller probability than under the greedy policy. For example, in comparison to greedy, under a monthly batching policy hard-to-match agents wait on average 6 days longer and easy-to-match agents 15 days longer. Hard- and easy-to-match agents get matched with 1.7% and 4% lower probability. Similarly, the patient policy matches equally many agents as the greedy policy, but leads to a substantially longer expected waiting time for hard-to-match agents (155 more days).

Remark 1. *It is important to note that Theorem 1 and Proposition 3 do not imply that batching policies are suboptimal for a fixed market size. For a fixed market size, a batching policy which matches very frequently will achieve (almost) the same outcome as the greedy policy and thus will also be close to optimal in large markets.*[15] *We investigate this in detail in Section 4 and show how frequent batching policies must match to be close to optimal.*

We now provide rough intuition for the differences between greedy, batching and patient matching policies. In Section 3.2 we sketch the argument for the various parts of the results. Section 4 provides a more extensive comparison of these three policies, including a non-asymptotic analysis.

Intuition for the Optimality of Greedy. As there are more hard- than easy-to-match agents, hard-to-match agents will accumulate and a large number of them will be waiting to be matched at any time under any policy.[16] This implies that under greedy matching, easy-to-match agents will have upon arrival, with high probability, a compatible hard-to-match agent and are therefore matched immediately. So every easy-to-match agent is matched with a hard-to-match agent, implying that greedy matching asymptotically achieves the optimal match rate. Intuitively, the market for hard-to-match agents already thickens under the greedy policy and further thickening the market is not beneficial for increasing the match rate.

Under the batching policy each agent waits at least from the time of her arrival until the next time a matching is identified. Thus each agent waits on average at least half the length of the batching interval. Furthermore, each agent becomes critical during that time with strictly positive probability. Thus, easy-to-match agents are worse off under the batching policy than under the greedy policy where they get matched immediately with probability 1. As some easy-to-match agents leave the market unmatched, hard-to-match agents are matched with a smaller probability. Consequently there are, on average, more hard-to-match agents waiting in the market. Little's law, which states that the arrival rate multiplied by the average waiting time equals the average number of waiting agents (Little and Graves, 2008), implies that hard-to-match agents also wait longer under any batching policy than under a greedy matching policy. As both types are worse off, batching policies are not asymptotically optimal.

Under the patient policy, so many hard-to-match agents accumulate that an easy-to-match agent will, with high probability, match with a critical hard-to-match agent almost immediately upon arrival. This implies that the policy asymptotically achieves the optimal match rate. As hard-to-match agents get matched only when they become critical, the distribution and expectation of their waiting time is the same as if they do not match at all. Hence, hard-to-match agents get matched faster under a greedy policy, implying that the patient policy is not asymptotically optimal.

[15] We formally establish this in Proposition 6.

[16] On the other hand, if there are more easy- than hard-to-match agents, then under the greedy policy only a small number of hard-to-match agents would be in the pool at any time; because otherwise, arriving easy-to-match agents will often be matched to hard-to-match agents, which will reduce the number of hard-to-match agents in the pool over time.

Intuitively, when the departure rate is small, increasing the market size and making the market thicker by waiting are loosely speaking substitutes in the sense that they increase the match rate. This is an intuition for why the greedy policy becomes optimal when the market grows large and the benefit of thickening the market vanishes. We note that thickening the market through waiting and increasing the market size exogenously lead to quite different compositions of waiting agents. If the arrival rate increases, under a greedy policy, only hard-to-match agents will accumulate (as in a large market easy-to-match agents will match immediately). In contrast, if the market is thickened by waiting in a batching policy, both hard- and easy-to-match agents accumulate.

3.2 Discussion of Results

In this section we provide a proof sketch for the various parts of Theorem 1 and Proposition 3 as well as additional results on the waiting time distributions. We first establish an upper bound on the performance of *any* policy.

Proposition 4 (Upper bound on the performance of any policy)**.** *For any market size m, and under any policy, the match rate of hard-to-match agents is at most $\frac{1}{1+\lambda}$ and their expected waiting time is at least $\frac{\lambda d}{1+\lambda}$.*

The result on the match rate follows from the fact that H agents cannot match with other H agents. The arrival rate of E agents is only m per unit of time, compared to $(1+\lambda)m$ for H agents. Thus, the strong law of large numbers implies that a fraction of $\frac{\lambda}{1+\lambda}$ of the H agents remain unmatched in the long-run. The result on waiting times is less immediate. We prove it by first deriving a policy-specific lower bound on the match rate; this lower bound holds with equality in policies where every agent who gets critical leaves the market unmatched.[17] Combining this lower bound on the match rate with the previously derived upper bound gives us a lower bound on the average number of H agents present in the market, which, through an application of Little's law, yields a lower bound on their waiting time.

3.2.1 Greedy Policy

Next we analyze the performance of the greedy policy as the market grows large. The following proposition includes the results in the first part of Proposition 3.

Proposition 5 (Performance of the greedy policy)**.** *Consider the greedy policy as the market grows large, i.e., $m \to \infty$. The match rate of H agents converges to $\frac{1}{1+\lambda}$ and their waiting time converges to an exponential distribution with mean $\frac{\lambda d}{1+\lambda}$. The match rate of E agents converges to 1 and their waiting time converges to 0.*

[17] Thus, the lower bound holds with equality in the greedy and batching policies, but not in the patient policy.

We first provide intuition for the waiting time distribution. Consider greedy matching in a deterministic setting where every E agent is compatible with every H agent, agents arrive deterministically, and get critical after precisely d units of time. In this setting E agents will be matched upon arrival with H agents. This means that there will be no E agents waiting in the market, and their waiting time equals zero. Denote by x the steady state number of H agents present in the market. Per unit of time, the number of H agents arriving to the market equals $(1+\lambda)m$, and m of them are matched with E agents. Furthermore, $\frac{x}{d}$ of the waiting agents are expected to depart unmatched per unit of time. At the steady state the number of unmatched departing agents equals the number of unmatched arriving agents. Thus, x solves the balance equation

$$\frac{x}{d} = \lambda\, m \quad \Rightarrow \quad x = \lambda\, m\, d\,. \tag{3}$$

Therefore, if the matching partner for an E agent is chosen at random, each H agent is matched at rate $\frac{m}{\lambda\, m\, d} = \frac{1}{\lambda\, d}$. The time at which a never-departing H agent would be matched is therefore exponentially distributed with rate $\frac{1}{\lambda\, d}$. The time until an H agent becomes critical is exponentially distributed with rate $1/d$. Since the minimum of two exponentially distributed random variables is again exponentially distributed with rate equal to the sum of the rates, the waiting time of an H agent is exponentially distributed with the rate $\frac{1+\lambda}{\lambda\, d}$, and thus with mean $\frac{\lambda\, d}{1+\lambda}$.

The formal proof of Proposition 5 is more complex as it needs to deal with the randomness in compatibilities, arrivals, and criticality times. Our analysis is based on the Lyapunov function method.[18] In our case, a Markov chain tracks the number of easy- and hard-to-match agents in the pool, and the Lyapunov function argument translates into a concentration bound for the number of easy- and hard-to-match agents in the pool. The number of H agents in the pool at any time is, with high probability, not more than an additive factor of $\sqrt{m}\log m$ away from the solution of the balance equation (3). As this distance grows slow relative to the market size, the Markov chain is well-approximated by the dynamics of the deterministic setting described earlier.

3.2.2 Batching Policy

We next sketch the analysis for the match rates and waiting times under the batching policy, as given in the second part of Proposition 3. We start by providing lower and upper bounds on the match rate of E agents. A simple upper bound on the match rate can be derived based on the fact that an arriving agent should wait until the next matching period and may not get matched if she becomes critical before that. We compute $\gamma_{T,d} = \frac{1-e^{-T/d}}{T/d}$ as the probability that an agent does not

[18] Variations of this method are widely used to identify stable points of ordinary differential equations and to analyze steady states of stochastic systems (see for example Khalil, 2009; Brémaud, 2013). The idea is defining a function on the state space of a Markov chain such that the expected change of the function is negative outside of a "small box" and possibly positive inside that box. The fact that the expected change of the function equals zero at the steady state distribution implies a bound on the time the process can spend outside the box. This idea can be used to provide bounds on the expectation of a given function f defined over a Markov chain (Anderson et al., 2017).

become critical before the first matching period after her arrival. This is clearly an upper bound on the match rate of E agents. Then, from the fact that H agents are compatible only with E agents, we imply that the match rate of H agents is at most $\frac{\gamma_{T,d}}{1+\lambda}$.

Providing a lower bound on the match rate of E agents is more involved. The key idea is showing that, every time when a matching is executed, the number of H agents matched is at least the number of E agents who are present in the pool at that time and arrived after the previous matching execution. The proof for this fact is based on a probabilistic analysis argument and uses *augmenting path* techniques from Berge's lemma in matching theory[19]. By this fact, almost every E agent who participates in at least one execution of the matching is matched to an H agent. This translates into an asymptotic lower bound of $\gamma_{T,d}$ on the match rate of E agents, and an asymptotic lower bound of $\frac{\gamma_{T,d}}{1+\lambda}$ on the match rate of H agents. Finally, the claim about the match rates follows immediately from the fact that the upper and the lower bound coincide.

To compute expected waiting times, we first note that $y/d = m(1 - q_E(m))$, where y is the time-average number of E agents in the pool and $1 - q_E(m)$ is the probability of an E agent not getting matched. This holds because, under the batching policy, the number of E agents that get critical equals the number of E agents that do not get matched. We then note that the probability $1 - q_E(m)$ converges to $1 - \gamma_{T,d}$ as m approaches infinity, since the match rate of E agents is $\gamma_{T,d}$. Thus, $d(1 - \gamma_{T,d}) = \frac{y}{m}$. The right-hand side is the expected waiting time of E agents, by Little's law. Hence, the expected waiting time of E agents equals $d(1 - \gamma_{T,d})$. A similar argument proves the claim for the waiting time of H agents.

Batching with vanishingly batch length. We next strengthen our previous analysis and show that a batching policy is asymptotically optimal if and only if the batch length vanishes with the market size.

Proposition 6. *A market size dependent batching policy with batch length T_m is asymptotically optimal if and only if the batch length goes to zero as the market becomes large* $\lim_{m\to\infty} T_m = 0$.

The intuition for the only if direction is that if $T_m > \delta > 0$ for every m, then the probability that a newly arrived agent becomes critical before a matching is executed is positive and does not vanish as the market grows large. Thus, the match rate of E agents would be below their match rate under the greedy policy. To prove the if direction, one cannot just take the limit $T \to 0$ of the expressions for match rate and waiting time obtained in Proposition 3, since the limits with respect to time and market size are not interchangeable. In particular, as T_m approaches 0, the Poisson concentration bounds on the number of arrivals in each period that we use to establish Proposition 3 become arbitrarily weak. Instead, we show that there is always a large number of H agents waiting in the pool, and an arriving E agent would be matched to one of these agents with high probability.

[19]See, e.g., West (2000) for more on Berge's lemma and augmenting paths.

3.2.3 Patient Policy

We next quantify the performance of the patient policy. The following proposition includes the third part of Proposition 3.

Proposition 7 (Performance of the patient policy)**.** *Consider the patient policy when the pool grows large, i.e., $m \to \infty$. The match rate of H agents converges to $\frac{1}{1+\lambda}$ and their waiting time converges to an exponential distribution with mean d. The match rate of E agents converges to* 1 *and their waiting time converges to* 0*.*

To get some intuition for this result consider again the hypothetical case in which every H agent is compatible to every E agent, and agents arrive and get critical deterministically. In steady state there are almost no E agents in the market and the number of H agents in the market is approximately $(1 + \lambda)\, m\, d$. To see why, suppose this is indeed the state of the market at the beginning of time. H agents get critical and attempt to match with E agents at a rate of $\frac{(1+\lambda)\, m\, d}{d} = (1 + \lambda)\, m$. Since this rate is much larger than the arrival rate of E agents m, then E agents are matched almost immediately at the steady state. Thus, their number remains close to zero at any time. Consequently, almost no E agent becomes critical, and almost all matches are initiated due to an H agent becoming critical. Since H agents arrive at rate $(1 + \lambda)m$ and get critical at rate $1/d$, their number in the pool remains close to $(1 + \lambda)md$, and the steady state is maintained. As H agents are the ones that initiate matches, their average waiting time equals the average time until they become critical, d.

4 Comparison of Batching and Greedy Policies in Finite Markets

Our results imply that, in a large market, greedy will outperform batching if the length of the batching interval does not go to zero (Theorem 1 and Proposition 6). It is natural to ask how frequent batching policies must execute matches to (potentially) outperform the greedy policy in a *finite* market. We address this question in two ways. First, we derive an upper bound on the batch length such that *any* batching policy with a larger batch length matches fewer agents and has a higher waiting time than the greedy policy. Second, we run simulations based on the model and real data that indicate that at realistic market sizes and parameters in the context of kidney exchanges, greedy matching dominates batching policies that match less frequently than daily.

4.1 An Analytical Bound on the Performance of Batching in Finite Markets

The next result combines non-asymptotic upper bounds on the performance of batching policies with lower bounds on the performance of the greedy policy to establish a bound on the batch length such that less frequent batching will be dominated by the greedy policy in a *finite* market.

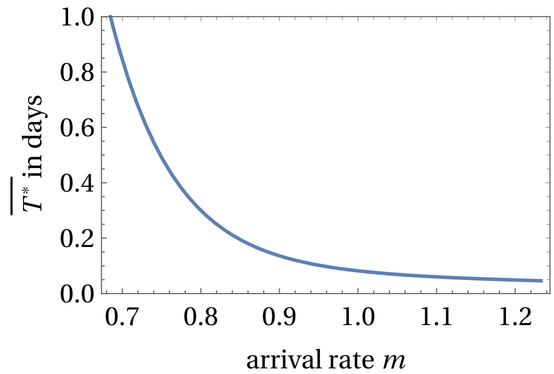

Figure 4: The batch length above which greedy dominates batching for various arrival rates per day, $\lambda = 1.33, p = 0.037$ and average criticality time $d = 360$ days. The bound $\overline{T^*}$ is independent of $q \in [0, 1]$, and is decreasing in p.

Proposition 8. *Let $m > 0$ be an arbitrary* fixed *arrival rate. Define z^* to be the steady-state probability that an easy-to-match agent, upon her arrival, is matched to a hard-to-match agent under the greedy policy. Then, for every agent type (easy- and hard-to-match), the match rate and waiting time of that type under the batching policy are respectively smaller and larger than under the greedy policy if the batch length T satisfies $T > T^*$, where*

$$T^* = \frac{z^* W\left(-\frac{e^{-1/z^*}}{z^*}\right) + 1}{z^*/d} \tag{4}$$

and $W(\cdot)$ is the Lambert W function.[20]

The above proposition involves the exact value for the steady state probability of an easy-to-match agent being matched to a hard-to-match agent upon her arrival. In Online Appendix ii we compute a lower bound on this probability using the Lyapunov function method and combine it with the above result to get an upper bound for T^*, namely $\overline{T^*}$. Therefore, a batching policy that makes matches less frequently than $\overline{T^*}$ is dominated by the greedy policy for the *fixed* arrival rate m. Figure 4 visualizes this bound. Recall that $\lambda = 1.33$ is consistent with the NKR data (see Section 2.1). Using the same data, we can calibrate the compatibility probabilities between the types by defining easy-to-match agents as those who are part of every maximal matching, which leads to empirical compatibilities of $p \approx 0.037, q \approx 0.087$.[21] For these parameters matching less frequently than daily leads to a lower match rate than greedy for any arrival rate of $m \geq 0.7$ per

[20] We recall that the Lambert W function is the inverse of the function $F(w) = we^w$.

[21] It is important to note that the patient-donor pairs identified as hard-to-match are not necessarily blood-type incompatible. Typically, patients of blood-type compatible pairs are much more sensitized, which means they have a lower chance to match with a random blood-type compatible donor. As there are few donors who can match highly sensitized patients, highly sensitized blood-type compatible pairs are typically not part of every maximal matching and our calibration thus identifies them as hard-to-match. The simulation results that use these calibrations are also confirmed by simulations that directly use real compatibility data.

day (roughly twice the size of the NKR). We note that Proposition 8 provides a sufficient but not necessary condition. The next section uses simulations to explore smaller arrival rates.

4.2 Numerical Simulations

We run simulations based on our two-type model to compare greedy and batching for various batch lengths and different parameters. The first set of simulations varies the total arrival rate per day $\bar{m} = m + (1+\lambda)m$, the imbalance λ and compatibility structure (p, q). In line with our calibration to the NKR data, the base case values are set to $\lambda = 1.33$ and $(p, q) = (0.037, 0.087)$. The *total arrival rate* (sum of arrival rates of both types) of agents is set to $\bar{m} = 0.25$, i.e., an agent every 4 days (and thus lower than the arrival rate at the NKR). Agents become critical on average after 360 days. This exercises suggests that even in moderately-sized markets batching needs to be very frequent for the batching policy not to be outperformed by the greedy policy. The results are reported in Figure 5. In each of these simulations a batching policy which matches less frequently than once a day will result in a lower match rate than greedy while matching more frequently than daily will be indistinguishable from greedy matching.[22]

Similar results hold for a wide range of parameters in the two-type model that are consistent with kidney exchange. A natural question is whether it is the case for *all* parameterizations of the model that greedy matching dominates batching in the two-type model. This turns out not to be true and we next provide a counter example.

An example where batching is beneficial. We identify a non-asymptotic setting where batching leads to a higher match rate than greedy. Consider the scenario in which all easy-to-match agents are compatible $q = 1$ and hard- and easy-to-match agents are rarely compatible $p = 0.02$.[23] The total arrival rate equals $\bar{m} = 2$ agents per day and agents become critical after $d = 360$ days. Intuitively, as in this market many easy-to-match agents have no hard-to-match partner, the greedy policy frequently matches easy-to-match agents to themselves; this results in more unmatched hard-to-match agents compared to the batching policy which matches fewer easy-to-match agents to themselves. As shown in Figure 6, when the market is balanced ($\lambda = 0$) a batching policy with a batch length of two days leads to matching about 2.6% more agents than greedy. The parameters in this example are carefully chosen and the benefit from batching vanishes when either the market size grows large—as predicted by Theorem 1—or as the market imbalance grows large.

The effect of imbalance. To see why imbalance has an effect on the optimality of the greedy policy, note that the number of H agents in the pool under greedy is less than $\lambda md - k\sqrt{md}$

[22]Figure 5d also plots results for values of $(p, q) = (0.3, 0.15)$ that are chosen to match the empirical frequency with which blood-type compatible and incompatible pairs can match.

[23]These parameters are not in line with kidney exchange data as blood-type compatible patient-donor pairs are selected to have highly sensitized patients and therefore typically have a low chance to be compatible to each other.

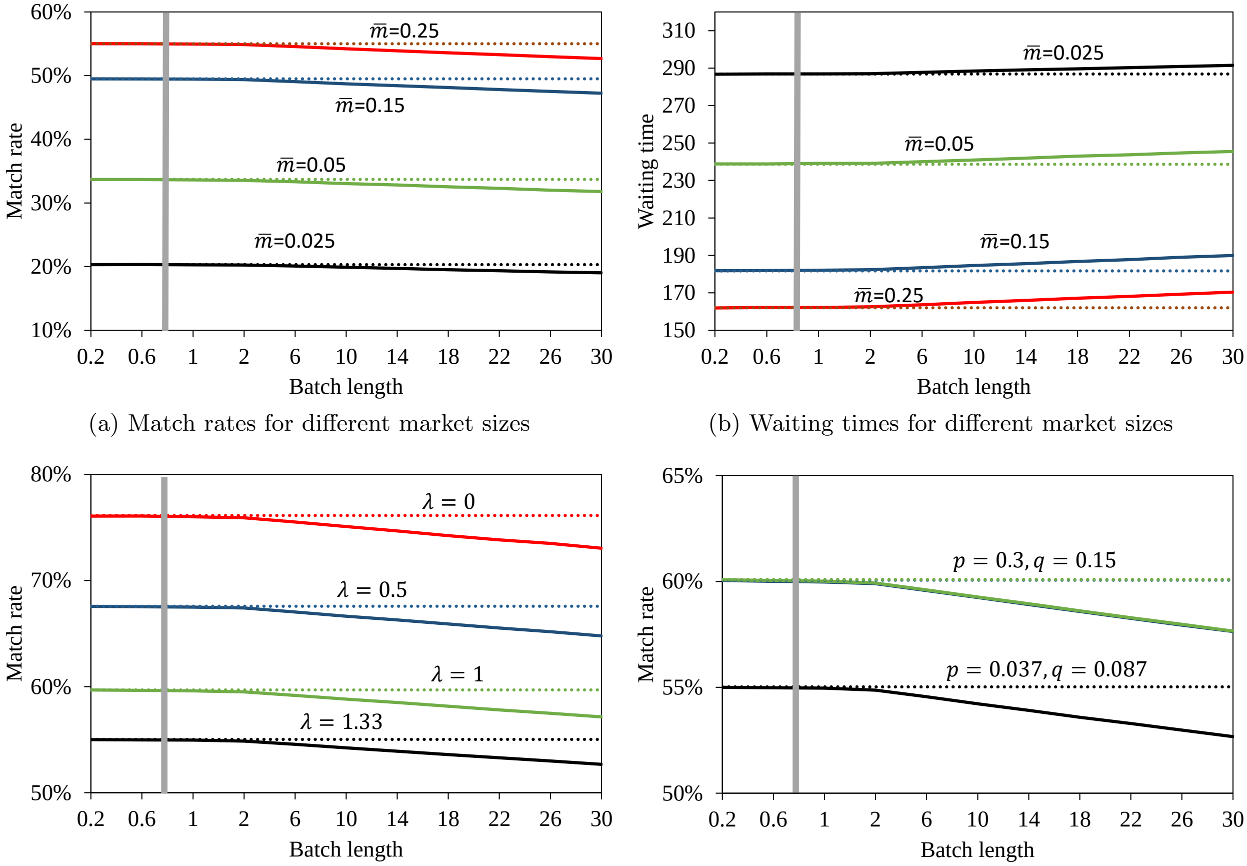

(a) Match rates for different market sizes

(b) Waiting times for different market sizes

(c) Different imbalances

(d) Different densities

Figure 5: Comparision of the greedy (dashed line) and the batching (solid line) policies in terms of the fraction of matched agents and waiting times. Batch lengths are in days. Comparisons are plotted for different total arrival rates $\bar{m}$ (Figures 5a and 5b), different imbalance values λ (5c), and different compatibilities (5d).

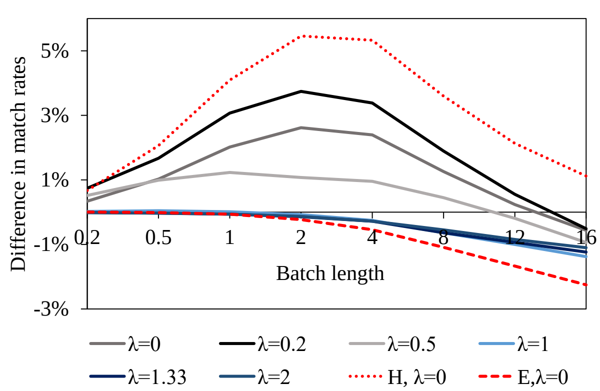

Figure 6: The difference between the match rates of batching and greedy policies for extreme match probabilities $p = 0.02, q = 1$ and various market imbalances.

with probability at most $O(e^{-k})$.[24] When the number of H agents is at least $\lambda md - k\sqrt{md}$, the probability that an arriving E agent is not matched to an H agent is at most $(1-p)^{\lambda md - k\sqrt{md}}$. A union bound then implies that the probability that an arriving E agent is not matched to an H agent is at most of the order of $e^{-k} + e^{-p(\lambda md - k\sqrt{md})}$ for every $k > 0$. Setting $k = \lambda\sqrt{md}/2$ implies that this probability is of the order of $e^{-\lambda\sqrt{md}/2}$, which approaches zero exponentially fast in $\lambda\sqrt{md}$. So, the inefficiency that results from E agents being matched to each other in the greedy policy vanishes if the market is either large (m is large) or imbalanced (λ is large).

5 Empirical Findings

While the greedy policy is optimal in a sufficiently large or imbalanced market, it is ultimately an empirical question whether our results hold in practice. We next complement our theoretical predictions with simulations using kidney exchange data. These simulations indicate that the greedy policy dominates batching in terms of match rate and waiting time, and has a much better waiting time compared to patient and a slightly worse match rate.[25]

The simulations presented here use data from the NKR. The NKR data includes 1881 de-identified patient-donor pairs between July 2007 to December 2014.[26] We use patients' and donors' blood types, antigens and antibodies to verify (virtual) compatibility between each donor and each patient. On average, approximately one patient-donor pair arrives per day to the NKR, and the average criticality time of a pair is estimated to be 360 days.[27] Arrivals of pairs are generated according to a Poisson process with a fixed arrival rate. We vary the arrival rate capturing market sizes from one-tenth to four times the size of the NKR. Pairs become critical according to an independent exponentially distributed random variable with mean equal to 360 (days), based on the empirical estimate. We simulate greedy, patient, and batching policies until 10 million pairs arrive to the market and report match rate and waiting time by taking averages over all or over a predefined subset of pairs of a certain type.

Table 1 reports the fraction of matched pairs and average waiting time. For the batching policy, we report results for weekly, monthly, and bimonthly batching ($T = 7, 30, 60$ days). The patient policy always results in the highest match rate, and the greedy and weekly batching result in a slightly lower match rate (and larger batch lengths result in a lower match rate). Moreover, the average waiting time under greedy matching is the smallest among all policies. Next we compare the greedy policy and batching policies with a finer range of batch lengths. The results are plotted in Figure 7 for three different total arrival rates ($\bar{m}$). In all cases the greedy policy outperformed

[24]This follows from our concentration bounds; see Theorem 4 in the appendix.

[25]Note here that the patient policy constitutes a theoretical benchmark as it is often impossible to observe criticality. The greedy and batching policies do not use any criticality information.

[26]Our focus is on bilateral matching and we therefore omit altruistic donors from the data.

[27]Hazard rates vary only slightly across pair types, such that for the sake of simplicity we aggregate all pairs and use a simple hazard rate model from Agarwal et al. (2019) to estimate criticality rate.

arrivals per day	match rate					waiting time in days				
	Greedy	Patient	Batching 7 days	30 days	60 days	Greedy	Patient	Batching 7 days	30 days	60 days
0.01	10.7%	11.9%	10.4%	9.9%	9.3%	322	355	322	324	326
0.05	22.4%	23.4%	22.2%	21.2%	20.2%	279	324	280	283	288
0.25	34.3%	35.4%	33.8%	32.6%	31.2%	237	298	238	243	248
0.5	38.5%	39.5%	38.0%	36.8%	35.2%	222	290	223	228	233
1	42.0%	43%	41.6%	40.2%	38.6%	209	283	210	215	221
2	45%	45.8%	44.5%	43.1%	41.5%	198	278	200	205	211
4	47.2%	48%	46.8%	45.4%	43.6%	190	274	192	196	207

Table 1: Fraction of pairs matched and average waiting times in days over all pairs in simulations using NKR data.

the batching policy in match rate and waiting time.

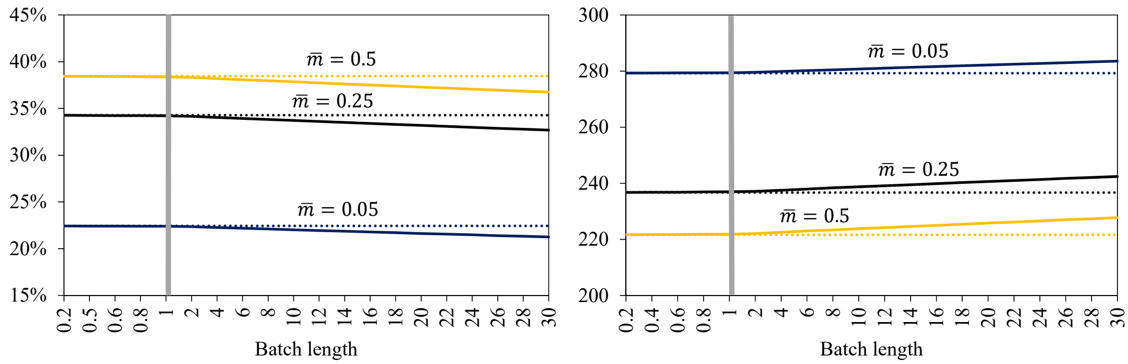

(a) Match rate (b) Waiting time

Figure 7: Match rate (left) and average waiting times (right) under greedy (dashed line) and batching (solid line) policies. Batch length given in days. $\bar{m}$ is the total arrival rate per day.

In the next simulations we address the concern that some types might be harmed by the greedy policy. To do so we compute average waiting times and match rates separately for two types of pairs, *under-demanded* and *over-demanded*. These can be thought of as hard- and easy-to-match, respectively. Under-demanded patient-donor pairs are blood type incompatible with each other (these include blood types patient-donor pairs O-X for X≠O, A-AB, and B-AB).[28] Over-demanded pairs are blood type compatible with each other (but not tissue-type compatible) and include pairs X-O for X≠O, AB-A, and AB-B.[29] Figure 8 reports the results. The waiting times (solid lines) of over-demanded pairs steadily decrease as the market becomes thicker, whereas the average waiting times of under-demanded pairs change only slightly. This finding is in line with the predictions from Proposition 3. Despite the heterogeneity in the data, the theoretical predictions (of the stylized two-type model) are aligned with the experiments when we categorize pairs as either over-demanded

[28] An X-Y patient-donor pair contains a patient with bloodtype X and a donor with bloodtype Y.

[29] Observe that A-O pairs (which are over-demanded) can match potentially match with each other and with O-A pairs (under-demanded) but O-A pairs cannot match with each other.

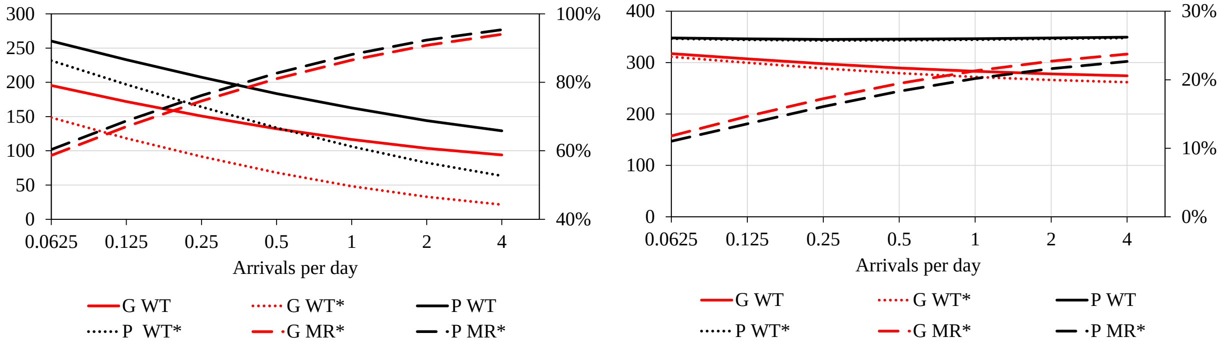

(a) Over-demanded pairs

(b) Under-demanded pairs

Figure 8: Average waiting times (WT) and match rate (MR) in days under greedy (G) and patient (P) policies. The left and right axes are WT and MR. The label (*) excludes pairs who have no match in the data.

or under-demanded. These patterns hold even though patients belonging to over-demanded pairs are, on average, more sensitized than those in under-demanded pairs.[30] We also report the statistics for the set of pairs that have at least one match in the historical data (dotted lines labeled with *).

In the last simulation we run greedy and patient policies under the base case scenario (with an arrival rate of 1 pair per day). For each pair, we compute the average waiting time over the copies of this pair sampled in the simulation as well as the fraction of the copies that are matched (i.e., the empirical probability of getting matched). For each of the 1881 pairs in the NKR data set, this simulation gives an average waiting time and an empirical probability of being matched under both the greedy and patient policies. Figure 9a shows that for each pair, the average waiting time is shorter under the greedy policy than under the patient policy; all of the dots are above the 45° line. This observation suggests that the waiting time distribution under the greedy policy first-order stochastically dominates the waiting time distribution under the patient policy.[31] Figure 9b reports the match rates under the greedy and patient policies. Observe that for most pairs the empirical probabilities of matching under the greedy and patient policies are "close" to the 45 degree line suggesting that the probability of being matched is roughly the same for *every pair*. Interestingly, Figure 9b suggests that under the greedy policy, easy-to-match pairs are slightly worse off because they are matched with slightly lower probability, whereas hard-to-match pairs are better off.

6 A Detailed Discussion of our Modeling Assumptions

From a modeling perspective there are three major differences between our paper and the closely related literature on dynamic matching (Ünver, 2010; Ashlagi et al., 2013; Anderson et al., 2017;

[30]More than 40% of patients in over-demanded pairs have less than a 5% chance of being tissue-type compatible with a random donor. Furthermore, about 10% of over-demanded pairs have no match within this data set, which is why the average waiting times do not drop all the way to zero.

[31]A detailed analysis of the simulation results confirms that this is indeed the case. We omit the details.

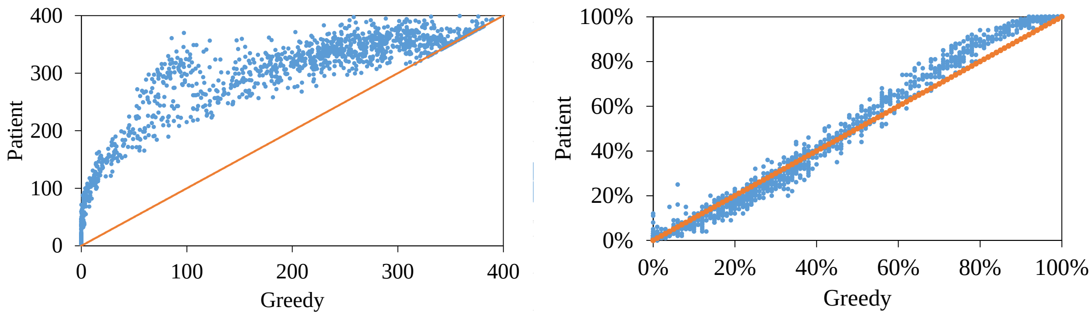

(a) Waiting times

(b) Chance of matching

Figure 9: Averages of waiting times (left) and chance of matching (right) taken over copies for each pair in the data. The axes correspond to the greedy and patient policies.

Akbarpour et al., 2020, 2019; Nikzad et al., 2019):

(i) Compatibilities in our model depend on the agents' types (blood types) and a random component (sensitization of the patient).[32]

(ii) We focus on markets where the compatibility probabilities do not vanish with the market size.

(iii) The objective of expected waiting time and probability of being matched differs from the objectives considered in some of the literature.

The Two-Type Compatibility Model. One may interpret our two-type model as a stylized way of capturing both blood types and randomness due to tissue-type incompatibilities. Intuitively, hard-to-match agents in our model correspond to pairs who cannot match with each other due to blood type incompatibility (for instance O-A blood type patient-donor pairs).[33] Due to the presence of such agents in kidney exchanges, not all agents can be matched even when the market grows large. As we have argued in Section 2.1, no single type model can reproduce this feature of real kidney exchanges while at the same time providing each pair at least one potential match. Furthermore, one motivation for our paper is the concern, sometimes raised by practitioners, that greedy matching might harm hard-to-match agents. This concern can by definition not be addressed in a single-type model where all agents are equally difficult to match.

As a robustness exercise we also ran simulations for our model when there is only a single type, i.e., $\lambda = -1$. We report the result in Section vi in the online appendix. These simulations indicate

[32] Anderson et al. (2017); Akbarpour et al. (2020) are examples of papers that consider a single type model. A notable exception is Ünver (2010) who analyzes matching policies in a deterministic compatibility model with multiple compatibility types, but without agent criticality times or random compatibility.

[33] Similar type of asymmetries across agents also appear in Nikzad et al. (2019). They are concerned with a proposal for global kidney exchange, which incorporates international pairs to domestic kidney exchange pools. Their model takes a reduced form approach where there is a continuum of international pairs who do not get matched to each other and a continuum of domestic pairs who can get matched to each other and to the international pairs. The compatibilities (between measures of pairs) are determined by a "matching function". They do a steady-state analysis to answer whether the savings from dialysis can cover the surgery costs of international pairs.

that the greedy policy remains optimal in large markets even with a single type.

Non-Vanishing Compatibility. Previous literature (Ashlagi et al., 2013; Anderson et al., 2017; Akbarpour et al., 2020, 2019; Nikzad et al., 2019) considered models with vanishing compatibility probabilities, i.e., when the arrival rate of agents equals m, the probability of compatibility of some pairs of agents equals $\frac{c}{m}$. In contrast, the probability of two patient-donor pairs being compatible in our model is independent of the market size. Assuming that the compatibility probability depends on the market size is intended to capture small and sparse kidney exchanges. In contrast, assuming that this probability is unaffected by the market size is natural when considering large kidney exchanges. Whether a given market is approximated by either compatibility model depends on the specific context, and is ultimately an empirical question. As we explained in Section 2.1, the combination of having two types and non-vanishing compatibilities allows us to capture crucial features of the compatibility graph in kidney exchanges.

We provide simulations for our model with vanishing compatibility probabilities in the online appendix. These simulations indicate if the compatibility probability approaches zero at rate $\frac{1}{\sqrt{m}}$, greedy is optimal in a large market while the patient and batching policies with a fixed batch length are not (Section vi in the online appendix). The simulations also show that this is not necessarily true if the compatibility probability vanishes at rate $\frac{1}{m}$. In this case, there is a trade-off and the greedy policy leads to a lower waiting time while the patient policy matches more agents.

Different Objectives. Another difference with some of the literature is the objective we consider. Akbarpour et al. (2020) consider the *loss rate*, denoted by $L_\pi \in [0,1]$ for a policy π, which is the probability that an agent is *not* matched under the policy π. To compare two policies π, π', they consider the *ratio of loss rates* $L_\pi/L_{\pi'}$. We focus on the probability of being matched in a policy and expected waiting time. One reason for studying match rate and waiting time is that they together determine the payoff of a risk-neutral expected-utility-maximizer (EU) who assigns a fixed value to being matched and incurs a constant cost while waiting in the market, whereas the ratio between the loss rates is not related to EU preferences. For example, we can have two policies π, π' which both match almost everyone (and thus achieve a loss rate close to zero) such that the ratio of loss rates $L_\pi/L_{\pi'}$ is arbitrarily large yet policy π matches agents much faster than π'. In the limit every risk-neutral EU maximizer prefers π over π' even though the policy π' is arbitrarily much better according to the ratio of loss rate.[34] In our model, the loss rates will converge to the same non-zero limit in a large market for the greedy and patient policy (as we established in Proposition 1). Thus,

[34]Consider an EU agent, with utility $u_\pi = vp - ct$ under policy π, whose value for being matched is $v > 0$ and her waiting cost per unit of time is $c > 0$. Consider two policies π, π', with match rates p, p' and waiting times t, t'. Suppose $p = 1 - 10^{-k}, p' = 1 - 10^{-2k}$ and $t = t'$, where k is a sufficiently large positive number. Then, $L_\pi/L_{\pi'} = 10^k$, whereas $u_\pi - u_{\pi'} = (10^{-k} - 10^{-2k})v$ approaches 0 as k grows large. Furthermore, when $t' - t > \frac{(10^{-k}-10^{-2k})v}{c}$, it holds that $u_{\pi'} < u_\pi$. Hence, even though the there is large gap between the loss ratios of the policies (in favor of π'), the agent's expected utility would be larger under π.

the greedy and patient policies are not ordered according to the ratio of loss rates, but any decision maker who considers a combination of the loss rates and the waiting time would prefer the greedy over the patient policy in a sufficiently large market. In the kidney exchange context both policies lead to almost the same probability of receiving a match, whereas pairs match much faster under the greedy policy than under the patient policy (see our experiments in Section 5).

7 Conclusion

This paper studies matching policies in a random dynamic market in which some agents are easier to match than others. We show theoretically as well as empirically that the greedy matching policy is arbitrarily close to optimal for all agents in sufficiently large markets. This finding has direct practical implications for kidney exchanges that may not employ greedy matching policies out of concern that greedy matching may harm those patients for whom it is hardest to find a compatible partner (Ferrari et al., 2014). Our simulations further suggest that matching frequently does not reduce the number of transplants even for realistic market sizes.[35]

This paper has some limitations in the context of kidney exchange. We only considered pairwise matchings and ignored frictions that occur in practice. Simulations in Ashlagi et al. (2018) account for such frictions, 3-way cycles, and chains.[36] They find that the greedy policy is optimal among a class of batching policies. We conjecture that this holds also true within our model, and that the benefit of matching in chains or longer cycles vanishes in a sufficiently large markets. It remains an interesting question to study these questions theoretically in small markets.

Throughout, we abstracted away from match qualities and assumed that agents are indifferent to whom they match with. When preferences over match partners play an important role, the greedy policy might not be optimal. For example, if matching easy-to-match agents with each other creates a much larger value than matching an easy-to-match agent to a hard-to-match one, the greedy policy can be suboptimal since almost all easy-to-match agents will match with hard-to-match agents in a large market (for studies with match qualities see Baccara et al., 2019; Li et al., 2019; Mertikopoulos et al., 2019; Blanchet et al., 2020; Aquilina et al., 2020).

We also abstracted away from the incentives agents might have to misreport their private information, such as arrival time, type, and realized compatibilities. One may ask what policy would be optimal if any of these would be the agent's private information. Under greedy the agent would have no incentive to delay reporting her arrival or claiming to be incompatible with agents whom she is compatible with; doing so would lead her to be matched later and less often. However,

[35]Independent of the policy, increasing the market size by merging pools can improve the match rate and waiting times of agents. So, competition between different kidney exchanges can harm the number of transplants but not because of frequent matching.

[36]For many exchanges chains are not a practical concern as numerous programs have very low enrollment of altruistic donors that initiate chains and, in some countries like France, Poland, and Portugal, chains are not even feasible since altruistic donation is illegal (Biro et al., 2017).

easy-to-match agents, after some histories, may prefer to misreport their type since our greedy policy break ties in favor of hard-to-match agents. To satisfy incentive compatibility one would need to break ties uniformly, which does not affect the asymptotic optimality of the greedy policy. Notably, this conclusion holds since agents are indifferent between who they match with.[37]

Appendix

A Proofs for Propositions 1 and 2

Proof of Proposition 1. Note that, as there are more E agents than H agents and H agents cannot match to themselves, $\frac{2}{2+\lambda}$ is an upper bound on the fraction of agents which can be matched for any m. Note, that the size of the maximal matching (SMM) equals $\frac{2}{2+\lambda}$ if the bipartite graph with m easy-to-match agents and m hard-to-match agents on the other side admits a perfect matching. The probability that such a perfect matching exists converges to one as $m \to \infty$ (see for example Theorem 5.1 page 77 in Frieze and Karoński (2015)). This proves the claim about SSM.

The probability that a hard-to-match agent has no partner is given by $(1-p)^m$. Because the compatibilities between hard-to-match and easy-to-match agents are drawn independently, the probability that all hard-to-match agents have at least one partner is given by $(1-(1-p)^m)^{m(1+\lambda)}$. This probability converges to one as $m \to \infty$. The same argument shows that the probability that all easy-to-match agents have at least one partner approaches one as m approaches infinity. □

Proof of Proposition 2. The proof is by contradiction; suppose such $p(m)$ exists. The chance that an agent has no other compatible agents is $(1-p(m))^m$. If $p(m) = O(1/m)$, then for sufficiently large m we have

$$(1-p(m))^m > e^{-2mp(m)} = e^{-O(1)},$$

since $1-\alpha > e^{-2\alpha}$ for $\alpha \in (0, \frac{1}{2})$. Thus, (2) cannot be satisfied. Therefore, suppose that $p(m) = \frac{\omega(m)}{m}$, where $\lim_{m\to\infty} \omega(m) = \infty$. Next, we use this property to show that (1) cannot be satisfied.

The proof is constructive. We propose a simple algorithm that chooses a matching μ with size $|\mu|$ such that $\lim_{m\to\infty} \frac{|\mu|}{m} = 1$. Our algorithm is a greedy algorithm, defined as follows. It orders agents of the graph from 1 to m and visits the agents one by one. When visiting agent i, if there are no agents left that are compatible with agent i, then the algorithm passes agent i and moves to agent $i+1$. Otherwise, the algorithm chooses one of the neighbors of agent i arbitrarily, namely agent j, and adds the pair (i, j) to the matching. The algorithm then visits the next available agent in the ordering. This process continues until the algorithm visits all agents.

We claim that the algorithm produces a matching μ which satisfies $\lim_{m\to\infty} \frac{|\mu|}{m} = 1$. Let $\phi(m)$ be a function that grows faster than $\frac{m}{w(m)}$ but slower than m. Then, during the algorithm, so long

[37] Baccara et al. (2019) study incentive compatible matching policies in a setting with match qualities.

as there are $\phi(m)$ unnvisited agents, the chance that a visited agent has no compatible agents is

$$(1 - w(m)/m)^{\phi(m)} \leq e^{-\frac{w(m)\phi(m)}{m}} = o(1).$$

Hence, so long as there are $\phi(m)$ agents left in the graph, the agent visited by the algorithm will be matched with a probability at least $1 - q(m)$ where $\lim_{m\to\infty} q(m) = 0$. By linearity of expectation, the expected number of unmatched agents by the end of the algorithm is then at most $\phi(m) + (m - \phi(m)) \cdot q(m)$. Noting that $\lim_{m\to\infty} \frac{\phi(m)+(m-\phi(m))\cdot q(m)}{m} = 0$ completes the proof. □

B Preliminaries

Let m_Θ denote the arrival rate of agents of type $\Theta \in \{E, H\}$. We use the terms E *pool* and H *pool* to denote the pools containing only E agents and only H agents, respectively. The *criticality clock* of an agent refers to the exponential random variable that determines the exogenous time that an agent becomes critical. Immediately after the criticality clock of an agent present in the pool *ticks*, she departs the pool. (Throughout the draft, the term *departure* is used to refer to the event of an agent leaving the pool, either matched or unmatched.)

We next describe how the greedy and patient policies break ties between feasible matches. Consider an agent, say a, who arrives to the market under the greedy policy or gets critical under the patient policy. At the time of this event, both policies attempt to match a as follows. First a strict order over all H agents in the market is selected uniformly at random, and a is matched with the first compatible H agent according to the selected order. If such an H agent does not exist, then a strict order over all E agents in the pool is selected uniformly at random, and a is matched with the first compatible E agent in that order if such an agent exists.

B.1 Asymptotic Notions

We say a statement $\mathcal{S}(i)$ holds for *sufficiently large* i if there exists i_0 such that $\mathcal{S}(i)$ holds for all $i > i_0$. Let $E(i)$ be an event parameterized by a positive integer i. We say that $E(i)$ occurs with *high probability as* i *grows large* if $\lim_{i\to\infty} \mathbb{P}[E(i)] = 1$. We often let the parameter i be m, the arrival rate of easy-to-match agents. When this is clearly known from the context, we simply say that $E(m)$ occurs *with high probability* or, briefly, $E(m)$ occurs *whp*.

Furthermore, we say that $E(i)$ occurs *with very high probability as* i *grows large* if there exists $\alpha > 1$ such that $\lim_{i\to\infty} \frac{1-\mathbb{P}[E(i)]}{e^{-(\ln i)^\alpha}} = 0$. We often let the parameter i be m, the arrival rate of easy-to-match agents. When this is clearly known from the context, we simply say that $E(m)$ occurs *with very high probability* or, briefly, $E(m)$ occurs *wvhp*.

For any two functions $f, g : \mathbb{R}_+ \to \mathbb{R}_+$ we adopt the notation $f = o(g)$ when for every positive constant ϵ there exists a constant i_ϵ such that $f(i) \leq \epsilon g(i)$ holds for all $i > i_\epsilon$. We define $f = O(g)$ if there exist positive constants i_0, Δ such that $f(i) \leq g(i)\Delta$ holds for all $i > i_0$.

B.2 Markov Chains

We denote the state space of a Markov chain $\mathcal{X}$ by $V(\mathcal{X})$.

Proposition 9 (Anderson et al., 2017, Proposition EC.4)**.** *Let $\mathcal{X} = (X_0, X_1, X_2, X_3, \ldots)$ be a discrete time positive recurrent Markov chain with a countable state space and steady state distribution η. Also, let $\mathbb{E}_x[\cdot]$ denote the expectation operator conditional on $X_0 = x$. Suppose that there exist real numbers $\alpha, \beta \geq 0$ and $\gamma > 0$, a set $B \subset S$, and functions $U : V(\mathcal{X}) \to \mathbb{R}_+$ and $f : V(\mathcal{X}) \to \mathbb{R}_+$ such that for $x \in V(\mathcal{X}) \backslash B$,*

$$\mathbb{E}_x \left[U(X_1) - U(X_0) \right] \leq -\gamma f(x), \tag{5}$$

and for $x \in B$,

$$f(x) \leq \alpha, \tag{6}$$
$$\mathbb{E}_x \left[U(X_1) - U(X_0) \right] \leq \beta. \tag{7}$$

Then

$$\mathbb{E}_{X \sim \eta} \left[f(X) \right] \leq \alpha + \frac{\beta}{\gamma}. \tag{8}$$

The stochastic processes associated with our matching policies are continuous-time processes. The above proposition, however, is applicable to discrete-time processes. To close this gap, we use the notion of *embedded Markov chain.*

Embedded Markov Chain

Let $\mathcal{X}$ be a continuous-time Markov chain with a countable state space. For any two states of $\mathcal{X}$, namely i, j, let $n_{i,j}$ denote the transition rate from state i to state j. Let N be the *transition rate matrix* for $\mathcal{X}$, i.e., $N_{i,j} = n_{i,j}$ for $i \neq j$, and the entries on the diagonal of N are set so that each row in N sums to 0.

Definition 5. *The embedded Markov chain of $\mathcal{X}$, denoted by $\widehat{\mathcal{X}}$, is a discrete-time Markov chain with the same state space as $\mathcal{X}$. The transition probability from state i to state j in $\widehat{\mathcal{X}}$ is denoted by $\widehat{n}_{i,j}$ and is defined by*

$$\widehat{n}_{i,j} = \begin{cases} \frac{n_{i,j}}{\sum_{k \neq i} n_{i,k}} & \text{if } i \neq j \\ 0 & \text{if } i = j. \end{cases}$$

Fact 2 (Harchol-Balter (2013))**.** *Let $\mathcal{X}$ be an ergodic continuous-time Markov chain with a unique stationary distribution ρ and transition rate matrix N. Then, the embedded Markov chain of $\mathcal{X}$,* *namely $\widehat{\mathcal{X}}$ has a unique steady-state distribution, namely $\widehat{\rho}$. Furthermore, for every state $i \in V(\mathcal{X})$,*

$$\rho(i) = \frac{\widehat{\rho}(i)/N_{i,i}}{\sum_{j\in V(\mathcal{N})} \widehat{\rho}(j)/N_{j,j}}.$$

B.3 Inequalities

Fact 3 (Canonne, 2019)**.** *For a Poisson random variable X with mean μ, it holds that*

$$\mathbb{P}\left[|X-\mu| > z\right] \le 2e^{-\frac{z^2}{\mu+z}}. \tag{9}$$

Chernoff bounds are concentration inequalities that bound the deviations of a weighted sum of Bernouli random variables from its mean. Below we present their multiplicative form.

Fact 4 (Chernoff bound Brmaud 2017)**.** *Let $X_1, \dots, X_n$ be a sequence of n independent random binary variables such that $X_i = 1$ with probability p_i and $X_i = 0$ with probability $1 - p_i$. Let $\alpha_1, \dots, \alpha_n$ be arbitrary real numbers in the unit interval. Also, let $S = \sum_{i=1}^{n} \alpha_i \mathbb{E}\left[X_i\right]$. Then for any ϵ with $0 \le \epsilon \le 1$ we have:*

$$\Pr\left[\textstyle\sum_{i=1}^{n} \alpha_i X_i > (1+\epsilon)S\right] \le e^{-\epsilon^2 S/3},$$
$$\Pr\left[\textstyle\sum_{i=1}^{n} \alpha_i X_i < (1-\epsilon)S\right] \le e^{-\epsilon^2 S/2}.$$

Fact 5. *For reals A, B, at every $x \ne -B$, the function $g(x) = \frac{x+A}{x+B}$ is increasing in x iff $A \le B$.*

Proof. We observe that $g'(x) = \frac{B-A}{(B+x)^2}$. Hence, when $x \ne -B$, $g'(x) \ge 0$ holds iff $A \le B$. ☐

Fact 6. *For every real A, B, at every $x \ne -B/2$, the function $g(x) = \frac{x-A}{2x+B}$ is increasing in x.*

Proof. Observing that $g'(x) = \frac{2A+B}{(B+2x)^2}$ proves the claim. ☐

C Proof of Proposition 4

We fix $\lambda > 0$ throughout the proof.

Definition 6. *For a policy τ and every agent type $\Theta \in \{E, H\}$, let $q_\Theta^\tau(m)$ and $w_\Theta^\tau(m)$ respectively denote the match rate and the expected waiting time of agents of type Θ under the policy τ when the arrival rates of E and H agents are respectively m and $m(1+\lambda)$.*

Lemma 1. *For every agent type $\Theta \in \{E, H\}$, $q_\Theta^\tau(m) \ge 1 - \frac{w_\Theta^\tau(m)}{d}$.*

Proof. Let $\mathbb{E}_t[\cdot], \mathbb{P}_t[\cdot]$ denote the expectation and probability operator conditional on all information the planner has at time t. Consider an agent i with an arbitrary type $\Theta \in \{E, H\}$. Recall that φ_i is the random variable denoting the difference between the time an agent i arrives to the market and the time that she departs (whether matched or not). Let κ_i be the difference between the time an agent i arrives to the market and the time she becomes critical. By definition κ_i is exponentially distributed with mean d. Since the policy does not observe the value of κ_i before the agent becomes critical, then

$$\mathbb{P}_{\alpha_i}[\varphi_i < t \mid \kappa_i = t] = \mathbb{P}_{\alpha_i}[\varphi_i < t \mid \kappa_i \geq t] .$$

This implies that the probability that the agent departs strictly before she becomes critical is given by

$$\begin{aligned}
\mathbb{P}_{\alpha_i}[\varphi_i < \kappa_i] &= \int_0^\infty \mathbb{P}_{\alpha_i}[\varphi_i < \kappa_i \mid \kappa_i = t]\frac{1}{d}e^{-\frac{1}{d}t}dt = \int_0^\infty \mathbb{P}_{\alpha_i}[\varphi_i < t \mid \kappa_i \geq t]\frac{1}{d}e^{-\frac{1}{d}t}dt \\
&= \int_0^\infty (1 - \mathbb{P}_{\alpha_i}[\varphi_i \geq t \mid \kappa_i \geq t])\frac{1}{d}e^{-\frac{1}{d}t}dt = 1 - \int_0^\infty \mathbb{P}_{\alpha_i}[\varphi_i \geq t \mid \kappa_i \geq t]\mathbb{P}_{\alpha_i}[\kappa_i \geq t]\frac{1}{d}dt \\
&= 1 - \int_0^\infty \mathbb{P}_{\alpha_i}[\varphi_i \geq t]\frac{1}{d}dt = 1 - \frac{1}{d}\mathbb{E}[\varphi_i] .
\end{aligned}$$

As the agent always departs the market matched whenever she departs before becoming critical, then a lower bound on the probability that this agent is matched is given by $1 - \frac{1}{d}\mathbb{E}_{\alpha_i}[\varphi_i]$, i.e., $1 - \frac{1}{d}\mathbb{E}_{\alpha_i}[\varphi_i] \leq \mathbb{E}_{\alpha_i}[\mu_i]$. Taking the average over agents of type Θ and using the law of iterated expectations yields that

$$\begin{aligned}
1 - \frac{w_\Theta^\tau(m)}{d} &= 1 - \frac{1}{d}\lim_{t\to\infty}\mathbb{E}\left[\frac{\sum_{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta}\varphi_i}{|\{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta\}|}\right] = \lim_{t\to\infty}\mathbb{E}\left[\frac{\sum_{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta} 1 - \frac{1}{d}\varphi_i}{|\{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta\}|}\right] \\
&= \lim_{t\to\infty}\mathbb{E}\left[\frac{\sum_{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta} 1 - \frac{1}{d}\mathbb{E}_{\alpha_i}[\varphi_i]}{|\{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta\}|}\right] \leq \lim_{t\to\infty}\mathbb{E}\left[\frac{\sum_{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta}\mathbb{E}_{\alpha_i}[\mu_i]}{|\{i\colon \alpha_i \leq t\}|}\right] \\
&= \lim_{t\to\infty}\mathbb{E}\left[\frac{\sum_{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta}\mu_i}{|\{i\colon \alpha_i \leq t\}|}\right] = q_\Theta^\tau(m) .
\end{aligned}$$

□

Proof of Proposition 4. Consider a policy τ, and let A_t and B_t respectively denote the number of hard-to-match and easy-to-match agents that arrive prior to time t under the policy τ. Since hard-to-match agents can be matched only to easy-to-match agents, the match rate of hard-to-match agents is at most $\lim_{t\to\infty}\frac{B_t}{A_t}$, by the Ergodic theorem. By the strong law of large numbers, $\lim_{t\to\infty}\frac{B_t/t}{A_t/t} = \frac{1}{1+\lambda}$. This proves that $q_H^\tau(m) \leq \frac{1}{1+\lambda}$ and establishes the claim about the match rate.

To prove the result for waiting time, we use Lemma 1 to write $q_H^\tau(m) \geq 1 - \frac{w_H^\tau(m)}{d}$. On the other hand, we showed that $\frac{1}{1+\lambda} \geq q_H^\tau(m)$. Hence, $\frac{1}{1+\lambda} \geq 1 - \frac{w_H^\tau(m)}{d}$, which means that $w_H^\tau(m) \geq \frac{d\lambda}{1+\lambda}$. □

D Greedy Matching: Concentration Bound

Here we analyze the stochastic process corresponding to the greedy policy and develop a concentration bound on the number of hard-to-match agents present in the pool. Because one can renormalize the time scale and the arrival rates linearly with a factor of $1/d$, we suppose that $d = 1$ throughout this section. This is without generality, as speeding up or slowing down time does not change the steady state distribution of the number of hard- or easy-to-match agents in the pool. We further assume that $m \geq 1$ throughout this section, i.e., at least one agent arrives per year.

The main technical results established by this analysis are a lower and an upper concentration bound for the number of hard-to-match agents in the greedy policy, developed respectively in Sections D.2 and D.3. Using these concentration bounds we will be able to prove the results on the match rate and waiting time under the greedy policy, which appear in Section E.

D.1 Modeling the Dynamics

We use a two-dimensional continuous-time Markov chain, which we denote by $\mathcal{M}$, to model the dynamics of the market. First we set up some notation before proceeding to the description. For a Markov chain $\mathcal{M}$, we recall that $V(\mathcal{M})$ denotes the state space of $\mathcal{M}$. We represent each state by a pair (x, y) where x, y respectively denote the number of H agents and the number of E agents. In other words, we have

$$V(\mathcal{M}) = \{(x, y) : \ x, y \in \mathbb{Z} \text{ and } x, y \geq 0\}.$$

By definition, the Markov chain is at state (x, y) if there are x hard-to-match and y easy-to-match agents in the pool.

Lemma 2. *$\mathcal{M}$ has a unique stationary distribution.*

The above lemma is proved in the online appendix, Section i. The proof uses standard arguments that bound the expected return time to a fixed state.

Definition 7. *Let π denote the stationary distribution of $\mathcal{M}$, with $\pi_{x,y}$ denoting the probability that π assigns to a state (x, y), and let $\pi_x = \sum_{y=0}^{\infty} \pi_{x,y}$ be the associated marginal distribution of x.*

D.2 A Large Market Lower Concentration Bound for Greedy Matching

In this section we provide the core technical result for the analysis of greedy matching: Theorem 3. This theorem provides a concentration bound for the stochastic process associated with the greedy policy. We state and prove Theorem 3 in Section D.2.4, after the following preliminary analysis.

We first define and analyze a new stochastic process called *simplified greedy* which differs from greedy in that E agents are not matched with each other. We prove that the number of H agents present in the market in the new process is lower (in first-order stochastic dominance) than the number of H agents present under the greedy policy, and use this to prove the concentration bound.

D.2.1 The Simplified Greedy Process and its Associated Markov Chain

Definition 8. *The simplified greedy process is the same as the greedy process with the difference that, in the simplified greedy process, E agents are considered to be incompatible (and therefore are*

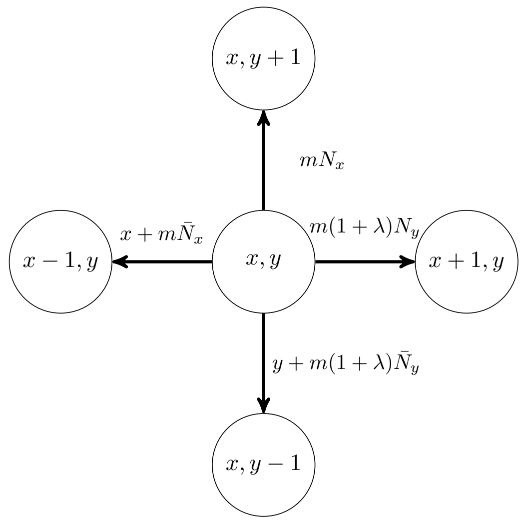

Figure 10: An illustration of the transitions from node (x, y) to its neighbors.

never matched to each other).

We denote by $\mathcal{N}$ the two-dimensional, continuous time Markov chain counting the number of H and E agents present at every point in time under the simplified greedy policy. As before we denote by x the number of H agents and by y the number of E agents. We next describe the transition rates of $\mathcal{N}$. A transition can only happen from a state (x, y) to its (at most) four *neighbors*,

$$\{(x', y') \in \mathbb{Z}_+^2 \colon |x - x'| + |y - y'| = 1\}.$$

See Figure 10 for a visual depiction of the neighbors and transition rates. To simplify the definition of transition rates from a node to its neighbors, we define the following notations: Let $N_x = (1-p)^x$ and $\bar{N}_x = 1 - N_x$. (Thus, N_x is the probability that an E agent is incompatible with x H agents.) For each state (x, y), we denote the transition rates from this state to its neighbor on the top, right, bottom, and left by $u_{x,y}, r_{x,y}, d_{x,y}, l_{x,y}$, respectively. These rates are defined as follows:

- $u_{x,y} = mN_x$ is the transition rate from the node (x, y) to node $(x, y+1)$. This holds because E agents arrive with rate m; after the arrival of an E agent, the number of E agents increases by one if the arriving E agent is not compatible to any H agent present in the pool.

- $r_{x,y} = m(1+\lambda)N_y$ is the transition rate from the node (x, y) to node $(x+1, y)$. This holds because H agents arrive with rate $(1+\lambda)m$; after the arrival of an H agent, the number of H agents increases by one if the arriving H agent is incompatible to all E agents in the pool.

- $d_{x,y} = y + m(1+\lambda)\bar{N}_y$ is the transition rate from the node (x, y) to node $(x, y-1)$. This holds because the number of E agents goes down by one when (i) a new H agent arrives who is compatible to an E agent; this happens with rate $m(1+\lambda)\bar{N}_y$; (ii) an existing E agent becomes critical and departs the pool; this happens with rate y.

- $l_{x,y} = x + m\bar{N}_x$ is the transition rate from the node (x, y) to node $(x-1, y)$. This holds because the number of H agents goes down by one when (i) a new E agent arrives who is compatible to an H agent; this happens with rate $m\bar{N}_x$; (ii) an existing H agent becomes critical and departs the pool; this happens with rate x.

Lemma 3. *$\mathcal{N}$ has a unique stationary distribution.*

Proof. The proof is identical to the proof of Lemma 2, but for the letter $\mathcal{M}$ replaced with $\mathcal{N}$. □

Definition 9. *Let ρ denote the stationary distribution of $\mathcal{N}$, with $\rho_{x,y}$ denoting the probability that ρ assigns to the state (x, y). Also, let $\rho_x = \sum_{y=0}^{\infty} \rho_{x,y}$.*

Next we show that, at the steady state, fewer hard-to-match agents wait in the simplified greedy process $\mathcal{N}$ than in the original greedy process $\mathcal{M}$ in the sense of first-order stochastic dominance.

Lemma 4. *For every $x \geq 0$, $\sum_{i=0}^{x} \pi_i \leq \sum_{i=0}^{x} \rho_i$.*

The proof of the lemma is technical and is deferred to the online appendix, Section i. The proof idea is defining a *coupling* of $\mathcal{M}$ and $\mathcal{N}$ such that, in the coupled process, there are more H agents in the pool under $\mathcal{M}$ than under $\mathcal{N}$ at any time, and fewer E agents.

Definition 10. *Let $\widehat{\mathcal{N}}$ denote the embedded Markov chain corresponding to $\mathcal{N}$. Also, let $\hat{\rho}$ denote its unique stationary distribution.*

In the above definition, we recall that $\hat{\rho}$ exists and is unique by Fact 2.

D.2.2 Applying Proposition 9 to $\widehat{\mathcal{N}}$

We develop a concentration bound for $\hat{\rho}$ using Proposition 9. This bound is parameterized by a number $k > 0$. We let $k > 0$ be an arbitrary real number in the following analysis.

Let $x^* = \lambda m$. Define $B_k = \{(x, y) \in \mathbb{Z}_+^2 : |x - x^*| + |y| < k\sqrt{m}\}$. Let the functions f, U be:

$$f(x, y) = \mathbb{1}_{(x,y) \notin B_k}, \ U(x, y) = e^{\frac{y + |x - x^*|}{\sqrt{m}}},$$

where $\mathbb{1}_{(x,y) \notin B_k}$ is the indicator function that equals 1 if $(x, y) \notin B_k$. We will apply Proposition 9 to $\widehat{\mathcal{N}}$. To this end, we need the following definitions.

Definition 11. *Conditional on the Markov chain $\widehat{\mathcal{N}}$ being at a state (x, y), let the random variable (x_1, y_1) denote the next state that the Markov chain moves to. Define*

$$\Delta(x, y) = \mathbb{E}\left[U(x_1, y_1)\right] - U(x, y).$$

Definition 12. *For every $z \geq 0$, define*

$$H(z) = -e^{\frac{z}{\sqrt{m}}} \left(\frac{1}{\sqrt{m}} \frac{z - 2m(1+\lambda)N_z}{2m(1+\lambda) + z} - \frac{1}{m} \right).$$

Lemma 5. *For every $x > x^*$ and $y \geq 0$, $\Delta(x, y) \leq H(y + x - x^*)$.*

Proof. To shorten notation, let $n = (1+\lambda)m$, $\theta = m + n + x + y$, and $z = y + x - x^*$. Note that $z = y + |x - x^*|$. Using this notation, we can write $U(x, y) = e^{\frac{z}{\sqrt{m}}}$. Recall that we defined $N_\alpha = (1-p)^\alpha$ and $\bar{N}_\alpha = 1 - N_\alpha$ for every real $\alpha > 0$. Then,

$$\begin{aligned}
\Delta(x, y) &= e^{\frac{z+1}{\sqrt{m}}} \left(\frac{mN_x + nN_y}{\theta} \right) + e^{\frac{z-1}{\sqrt{m}}} \left(\frac{x + m\bar{N}_x + y + n\bar{N}_y}{\theta} \right) - e^{\frac{z}{\sqrt{m}}} && (10) \\
&= e^{\frac{z}{\sqrt{m}}} \left(e^{\frac{1}{\sqrt{m}}} \frac{mN_x + nN_y}{\theta} + e^{-\frac{1}{\sqrt{m}}} \left(\frac{x + m\bar{N}_x + y + n\bar{N}_y}{\theta} \right) - 1 \right) \\
&\leq e^{\frac{z}{\sqrt{m}}} \left(\left(1 + \frac{1}{\sqrt{m}} + \frac{1}{m} \right) \frac{mN_x + nN_y}{\theta} + \left(1 - \frac{1}{\sqrt{m}} + \frac{1}{m} \right) \left(\frac{x + m\bar{N}_x + y + n\bar{N}_y}{\theta} \right) - 1 \right) && (11) \\
&= e^{\frac{z}{\sqrt{m}}} \left(\frac{1}{\sqrt{m}} \frac{2mN_x + 2nN_y - m - n - x - y}{\theta} + \frac{1}{m} \right) && (12) \\
&= e^{\frac{z}{\sqrt{m}}} \left(\frac{1}{\sqrt{m}} \frac{2mN_x + 2nN_y - \theta}{\theta} + \frac{1}{m} \right) && (13)
\end{aligned}$$

To see why (10) holds, observe that $\frac{mN_x}{\theta}$ and $\frac{nN_y}{\theta}$ are the transition probabilities of $\widehat{\mathcal{N}}$ from the state (x, y) to the states $(x, y+1)$ and $(x+1, y)$, respectively. When either of these transitions occur, the value of U changes from $e^{\frac{z}{\sqrt{m}}}$ to $e^{\frac{z+1}{\sqrt{m}}}$. Also, $\frac{x+m\bar{N}_x}{\theta}$ and $\frac{y+n\bar{N}_y}{\theta}$ are the transition probabilities of $\widehat{\mathcal{N}}$ from the state (x, y) to the states $(x-1, y)$ and $(x, y-1)$, respectively. Note that these transition probabilities are 0 if $x - 1$ or $y - 1$ are negative. When either of these transitions occur, the value of U changes from $e^{\frac{z}{\sqrt{m}}}$ to $e^{\frac{z-1}{\sqrt{m}}}$. Inequality (11) holds because $e^\alpha \leq 1 + \alpha + \alpha^2$ holds for every $\alpha \in [-1, 1]$. Equations (12) and (13) hold by rearrangement of terms.

Claim 1. *Let a, b, c be positive reals such that $a < b$. The function $g(s) = (1-p)^{b-s} + c(1-p)^{s-a}$ is convex over $[a, b]$.*

Proof. Observe that

$$g''(s) = (1-p)^{-s} \log^2(1-p) \left(c(1-p)^{2s-a} + (1-p)^b \right) \geq 0,$$

which means that g is convex when $p \in (0, 1)$. Also, when $p = 1$, $g(s) = 0$ for every $s \in [a, b]$. □

Claim 2.

$$e^{\frac{z}{\sqrt{m}}} \left(\frac{1}{\sqrt{m}} \frac{2m\left(N_{\lambda m+z} + 1 + \lambda\right) - \theta}{\theta} + \frac{1}{m} \right) \leq H(z)$$

Proof. Recall that $\theta = m + n + x + y$, which means that $\theta = 2m(1+\lambda) + z$. Hence, to prove the claim, which says that

$$e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2m\left(N_{\lambda m+z}+1+\lambda\right)-\theta}{\theta}+\frac{1}{m}\right) \le e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2m(1+\lambda)N_z - z}{2m(1+\lambda)+z}+\frac{1}{m}\right),$$

it suffices to prove that

$$2m\left(N_{\lambda m+z}+1+\lambda\right)-\theta \le 2m(1+\lambda)N_z - z,$$

or equivalently, $2mN_{\lambda m+z} - z \le 2m(1+\lambda)N_z - z$. The latter inequality holds as $N_{\lambda m+z} \le N_z$. □

Claim 3.

$$e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2m\left(N_{\lambda m}+(1+\lambda)N_z\right)-\theta}{\theta}+\frac{1}{m}\right) \le H(z)$$

Proof. Recall that $\theta = 2m(1+\lambda) + z$. Hence, to prove that

$$e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2m\left(N_{\lambda m}+(1+\lambda)N_z\right)-\theta}{\theta}+\frac{1}{m}\right) \le e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2m(1+\lambda)N_z - z}{2m(1+\lambda)+z}+\frac{1}{m}\right),$$

it suffices to prove that

$$2m\left(N_{\lambda m}+(1+\lambda)N_z\right)-\theta \le 2m(1+\lambda)N_z - z,$$

or equivalently, $2mN_{\lambda m} - \theta \le -z$. The latter inequality holds as $\theta - z = 2m(1+\lambda)$ and $N_{\lambda m} \le 1$. □

Recall that $z = y + x - x^* = y + x - \lambda m$. We next complete the proof of the Lemma 5 by showing that

$$\Delta(x,y) \le e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2mN_x + 2nN_y - \theta}{\theta}+\frac{1}{m}\right) \le H(z)\,.$$

The first inequality holds by (13). To prove the second inequality, we observe that

$$\begin{aligned}
& e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2mN_x + 2nN_y - \theta}{\theta}+\frac{1}{m}\right) \\
& \le e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2m\max\{N_{\lambda m+z}+(1+\lambda)N_0, N_{\lambda m}+(1+\lambda)N_z\}-\theta}{\theta}+\frac{1}{m}\right) \qquad (14) \\
& \le \max\left\{e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2m\left(N_{\lambda m+z}+1+\lambda\right)-\theta}{\theta}+\frac{1}{m}\right), e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2m\left(N_{\lambda m}+(1+\lambda)N_z\right)-\theta}{\theta}+\frac{1}{m}\right)\right\} \qquad (15) \\
& \le H(z)
\end{aligned}$$

where (14) holds due to the convexity property established by Claim 1, (15) holds by rearrangement of terms, and the last inequality holds because each of the expressions in the max are at most $H(d)$, by Claim 2 and Claim 3. □

Lemma 6. *For every $x \le x^*$ and $y \ge 0$, $\Delta(x,y) \le H(y + x^* - x)$.*

Proof. Let $z = y + x^* - x$, $n = (1+\lambda)m$, and $\theta = m + n + x + y$. The proof considers two cases: either $x < x^*$ or $x = x^*$. First, we suppose that $x < x^*$. Then,

$$\Delta(x,y) = e^{\frac{z+1}{\sqrt{m}}}\left(\frac{m+x}{\theta}\right) + e^{\frac{z-1}{\sqrt{m}}}\left(\frac{y+n}{\theta}\right) - e^{\frac{z}{\sqrt{m}}} \tag{16}$$

$$= e^{\frac{z}{\sqrt{m}}}\left(e^{\frac{1}{\sqrt{m}}}\frac{m+x}{\theta} + e^{-\frac{1}{\sqrt{m}}}\left(\frac{y+n}{\theta}\right) - 1\right)$$

$$\le e^{\frac{z}{\sqrt{m}}}\left(\left(1 + \frac{1}{\sqrt{m}} + \frac{1}{m}\right)\frac{m+x}{\theta} + \left(1 - \frac{1}{\sqrt{m}} + \frac{1}{m}\right)\left(\frac{n+y}{\theta}\right) - 1\right) \tag{17}$$

$$= e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{m+x-n-y}{\theta} + \frac{1}{m}\right) = e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{-z}{\theta} + \frac{1}{m}\right) \tag{18}$$

$$\le e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{-z+2nN_z}{\theta} + \frac{1}{m}\right) = H(z). \tag{19}$$

To see why (16) holds, observe that $\frac{x+m\bar{N}_x}{\theta}$ and $\frac{N_x}{\theta}$ are the transition probabilities of $\widehat{\mathcal{N}}$ from the state (x,y) to the states $(x-1,y)$ and $(x,y+1)$, respectively; these transition probabilities sum up to $\frac{m+x}{\theta}$. When either of these transitions occur, the value of U changes from $e^{\frac{z}{\sqrt{m}}}$ to $e^{\frac{z+1}{\sqrt{m}}}$. Also, $\frac{nN_y}{\theta}$ and $\frac{y+n\bar{N}_y}{\theta}$ are the transition probabilities of $\widehat{\mathcal{N}}$ from the state (x,y) to the states $(x+1,y)$ and $(x,y-1)$, respectively; these transition probabilities sum up to $\frac{y+n}{\theta}$. (Note that the transition probabilities are 0 if $x-1$ or $y-1$ are negative.) When either of these transitions occur, the value of U changes from $e^{\frac{z}{\sqrt{m}}}$ to $e^{\frac{z-1}{\sqrt{m}}}$. Inequality (17) holds because $e^{\alpha} \le 1 + \alpha + \alpha^2$ holds for every real number in $[-1,1]$, and (19) holds because $N_z \ge 0$.

To complete the proof, it remains to prove the claim for the case of $x = x^*$. In this case, when a transition from (x,y) to $(x,y+1)$ occurs, the value of U changes from $e^{\frac{z}{\sqrt{m}}}$ to $e^{\frac{z+1}{\sqrt{m}}}$ (whereas in the above case, the value of U changes to $e^{\frac{z-1}{\sqrt{m}}}$ when this transition occurs). Accounting for this difference slightly changes the above calculations, but leads to the same conclusion:

$$\begin{aligned}
\Delta(x,y)&e^{\frac{z+1}{\sqrt{m}}}\left(\frac{m+x+nN_y}{\theta}\right) + e^{\frac{z-1}{\sqrt{m}}}\left(\frac{y+n\bar{N}_y}{\theta}\right) - e^{\frac{z}{\sqrt{m}}} \\
&\le e^{\frac{z}{\sqrt{m}}}\left(\left(1 + \frac{1}{\sqrt{m}} + \frac{1}{m}\right)\frac{m+x+nN_y}{\theta} + \left(1 - \frac{1}{\sqrt{m}} + \frac{1}{m}\right)\left(\frac{y+n\bar{N}_y}{\theta}\right) - 1\right) \\
&= e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{m+x+nN_y-y-n\bar{N}_y}{\theta} + \frac{1}{m}\right) \\
&= e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{-z+2nN_z}{\theta} + \frac{1}{m}\right) = H(z),
\end{aligned}$$

where the penultimate equality follows from $x = x^* = n - m$ and $y = z$. □

Lemma 7. *For every* $x, y \geq 0$ $\Delta(x,y) \leq H(y + |x - \lambda m|)$.

Proof. By Lemma 5, if $x > x^*$, then $\Delta(x,y) \leq H(y + x - x^*)$. By Lemma 6, if $x \leq x^*$, then $\Delta(x,y) \leq H(y + x^* - x)$. The claim follows immediately from the two latter bounds and the fact that $x^* = \lambda m$. □

Lemma 8. *For every* $k > 0$ *satisfying* $\sup_{z \geq k\sqrt{m}} H(z) < 0$ *it holds that*

$$\sum_{(x,y) \notin B_k} \hat{\rho}_{x,y} \leq -\frac{\sup\limits_{0 \leq z \leq k\sqrt{m}} H(z)}{\sup\limits_{z \geq k\sqrt{m}} H(z)}.$$

Proof. Applying Proposition 9 on $\widehat{\mathcal{N}}$ directly implies that

$$\sum_{(x,y) \notin B_k} \hat{\rho}_{x,y} \leq \alpha + \frac{\beta}{\gamma} \tag{20}$$

holds if there exist $\alpha, \beta \geq 0$ and $\gamma > 0$ such that

$$\begin{aligned} \forall (x,y) \notin B_k, \Delta(x,y) &\leq -\gamma f(x,y), \\ \forall (x,y) \in B_k, f(x,y) &\leq \alpha, \\ \forall (x,y) \in B_k, \Delta(x,y) &\leq \beta. \end{aligned}$$

Since $f(x,y) = \mathbb{1}_{(x,y) \notin B_k}$ by definition, then we can set $\alpha = 0$. Recall that

$$B_k = \{(x,y) \in \mathbb{Z}_+^2 : |x - x^*| + |y| < k\sqrt{m}\}.$$

Hence, Lemma 7 implies that

$$\sup_{(x,y) \in B_k} \Delta(x,y) \leq \sup_{0 \leq z \leq k\sqrt{m}} H(z).$$

Therefore, we can set

$$\beta = \sup_{0 \leq z \leq k\sqrt{m}} H(z). \tag{21}$$

We note that $\beta > 0$ holds since $H(0) = e^{\frac{1}{\sqrt{m}}} - 1 > 0$.

Finally, we observe that by Lemma 7,

$$\sup_{(x,y) \notin B_k} \Delta(x,y) \leq \sup_{z \geq k\sqrt{m}} H(z).$$

Therefore, we can set

$$\gamma = -\sup_{z \geq k\sqrt{m}} H(z). \tag{22}$$

Since $\sup_{z \geq k\sqrt{m}} H(z) < 0$ holds by assumption, then $\gamma > 0$.

We have set $\alpha = 0$, and have set β, γ by (21) and (22), respectively. This choice of parameters, together with (20), directly proves the claim. $\square$

D.2.3 A Concentration Bound for ρ

Corollary 1 (of Lemma 8)**.** *For every $k > 0$ satisfying $\sup_{z \geq k\sqrt{m}} \{H(z)\} < 0$ it holds that*

$$\sum_{(x,y) \notin B_k} \rho_{x,y} \leq -\frac{\sup\limits_{0 \leq z \leq k\sqrt{m}} H(z)}{\sup\limits_{z \geq k\sqrt{m}} H(z)} \cdot \frac{2(1 + \lambda + k/\sqrt{m})}{2 + \lambda}.$$

Proof. For a subset $S \in \mathbb{Z}_+^2$, let $\rho[S]$ denote $\sum_{(x,y) \in S} \rho_{x,y}$. Similarly, let $\hat{\rho}[S]$ denote $\sum_{(x,y) \in S} \hat{\rho}_{x,y}$. We denote $\mathbb{Z}_+^2 \backslash S$ by $\overline{S}$.

If $\rho[\,\overline{B_k}\,] \leq \hat{\rho}[\,\overline{B_k}\,]$, then the claim holds by the upper bound on $\hat{\rho}[\,\overline{B_k}\,]$ provided by Lemma 8. So, suppose that this is not the case; i.e.,

$$\rho[B_k] < \hat{\rho}[B_k]. \tag{23}$$

Define $w_{x,y} = \frac{1}{m(2+\lambda)+x+y}$. We note that $w_{0,0} \geq w_{x,y}$ holds for every $(x, y) \in \mathbb{Z}_+^2$. Also, define $\underline{w} = w_{\lambda m + k\sqrt{m}, k\sqrt{m}}$. We note that $\underline{w} \leq w_{x,y}$ for every $(x, y) \in B_k$.

By Fact 2 regarding the steady-state distribution of Embedded Markov chains, it holds that

$$\frac{\rho[\,\overline{B_k}\,]}{\rho[B_k]} \leq \frac{\hat{\rho}[\,\overline{B_k}\,]}{\hat{\rho}[B_k]} \cdot \frac{w_{0,0}}{\underline{w}}.$$

This implies that

$$\rho[\,\overline{B_k}\,] \leq \rho[B_k] \frac{\hat{\rho}[\,\overline{B_k}\,]}{\hat{\rho}[B_k]} \cdot \frac{w_{0,0}}{\underline{w}}.$$

The above inequality, together with (23), implies that

$$\rho[\,\overline{B_k}\,] \leq \hat{\rho}[\,\overline{B_k}\,] \cdot \frac{w_{0,0}}{\underline{w}} = \hat{\rho}[\,\overline{B_k}\,] \cdot \frac{2(1 + \lambda + k/\sqrt{m})}{2 + \lambda}.$$

The above bound, together with Lemma 8, implies that

$$\rho[\overline{B_k}] \le -\frac{\sup_{0\le z\le k\sqrt{m}} H(z)}{\sup_{z\ge k\sqrt{m}} H(z)} \cdot \frac{2(1+\lambda+k/\sqrt{m})}{2+\lambda}. \qquad \square$$

Theorem 2. *For every $k > 0$ define*

$$Q(k) = e^k\left(\frac{1}{\sqrt{m}}\frac{k\sqrt{m} - 2m(1+\lambda)N_{k\sqrt{m}}}{2m(1+\lambda)+k\sqrt{m}} - \frac{1}{m}\right),$$
$$R(k) = \sup_{0\le z\le k\sqrt{m}} H(z),$$
$$\Phi(k) = \frac{R(k)}{Q(k)} \cdot \frac{2(1+\lambda+k/\sqrt{m})}{2+\lambda}.$$

Then, $\sum_{(x,y)\notin B_k} \rho_{x,y} \le \Phi(k)$ holds if $Q(k) > 0$.

Proof. First we prove the following claim.

Claim 4. $\sup_{z\ge k\sqrt{m}} H(z) \le -Q(k)$.

Proof. Observe that for every $z \ge k\sqrt{m}$,

$$\begin{aligned}
H(z) &= -e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{z-2m(1+\lambda)N_z}{2m(1+\lambda)+z} - \frac{1}{m}\right) \\
&\le -e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\left(\frac{z}{2m(1+\lambda)+z} - \frac{2m(1+\lambda)N_z}{2m(1+\lambda)+z}\right) - \frac{1}{m}\right) \\
&\le -e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\left(\frac{k\sqrt{m}}{2m(1+\lambda)+k\sqrt{m}} - \frac{2m(1+\lambda)N_z}{2m(1+\lambda)+z}\right) - \frac{1}{m}\right) && (24) \\
&\le -e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\left(\frac{k\sqrt{m}}{2m(1+\lambda)+k\sqrt{m}} - \frac{2m(1+\lambda)N_{k\sqrt{m}}}{2m(1+\lambda)+k\sqrt{m}}\right) - \frac{1}{m}\right) && (25) \\
&\le -e^{k}\left(\frac{1}{\sqrt{m}}\left(\frac{k\sqrt{m}}{2m(1+\lambda)+k\sqrt{m}} - \frac{2m(1+\lambda)N_{k\sqrt{m}}}{2m(1+\lambda)+k\sqrt{m}}\right) - \frac{1}{m}\right) && (26) \\
&= -Q(k).
\end{aligned}$$

where (24) holds by Fact 5, (25) holds by the fact that $z \ge k\sqrt{m}$, and (26) holds since $Q(k) > 0$. $\square$

By the above claim, $\sup_{z\ge k\sqrt{m}} H(z) < 0$. Thus, Corollary 1 applies, which implies that

$$\sum_{(x,y)\notin B_k} \rho_{x,y} \le -\frac{\sup_{0\le z\le k\sqrt{m}} H(z)}{\sup_{z\ge k\sqrt{m}} H(z)} \cdot \frac{2(1+\lambda+k/\sqrt{m})}{2+\lambda} \le \frac{R(k)}{Q(k)} \cdot \frac{2(1+\lambda+k/\sqrt{m})}{2+\lambda},$$

where the last inequality holds by Claim 4. This completes the proof. $\square$

D.2.4 A Concentration Bound for π

Lemma 9. *For every $\alpha \geq 0$ and $z \leq 2\alpha(1+\lambda)\sqrt{m}$, it holds that $H(z) \leq e^{2\alpha(1+\lambda)}\left(\frac{1}{\sqrt{m}} + \frac{1}{m}\right)$.*

Proof. Observe that

$$H(z) = e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2m(1+\lambda)N_z - z}{2m(1+\lambda)+z} + \frac{1}{m}\right) \leq e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}} + \frac{1}{m}\right) \leq e^{2\alpha(1+\lambda)}\left(\frac{1}{\sqrt{m}} + \frac{1}{m}\right), \quad (27)$$

where the last inequality holds because $z \leq 2\alpha(1+\lambda)\sqrt{m}$. □

Lemma 10. *For every $\alpha \geq 3$, $m \geq \max\{36, p^{-2}\}$, and $z \geq 2\alpha(1+\lambda)\sqrt{m}$, $H(z) \leq \frac{1}{\sqrt{m}} - \frac{1}{m}e^{\frac{z}{\sqrt{m}}}$.*

Proof. Observe that

$$H(z) = e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2m(1+\lambda)N_z - z}{2m(1+\lambda)+z} + \frac{1}{m}\right) \leq e^{\frac{z}{\sqrt{m}}}\left(\frac{1}{\sqrt{m}}\frac{2m(1+\lambda)e^{-pz} - z}{2m(1+\lambda)+z} + \frac{1}{m}\right) \quad (28)$$

$$\leq \frac{1}{\sqrt{m}}\frac{2m(1+\lambda)e^{\frac{z}{\sqrt{m}}-pz}}{2m(1+\lambda)+z} - \frac{1}{\sqrt{m}}\frac{ze^{\frac{z}{\sqrt{m}}}}{2m(1+\lambda)+z} + \frac{1}{m}e^{\frac{z}{\sqrt{m}}}$$

$$\leq \frac{1}{\sqrt{m}} + e^{\frac{z}{\sqrt{m}}}\left(-\frac{1}{\sqrt{m}}\frac{z}{2m(1+\lambda)+z} + \frac{1}{m}\right) \quad (29)$$

$$\leq \frac{1}{\sqrt{m}} + e^{\frac{z}{\sqrt{m}}}\left(-\frac{1}{\sqrt{m}}\frac{2\alpha(1+\lambda)\sqrt{m}}{2m(1+\lambda)+2\alpha(1+\lambda)\sqrt{m}} + \frac{1}{m}\right) \quad (30)$$

$$\leq \frac{1}{\sqrt{m}} + e^{\frac{z}{\sqrt{m}}}\left(-\frac{1}{m/\alpha+\sqrt{m}} + \frac{1}{m}\right) \leq \frac{1}{\sqrt{m}} - \frac{1}{m}e^{\frac{z}{\sqrt{m}}} \quad (31)$$

where (28) holds because $1-\alpha \leq e^{-\alpha}$ for every real α, (29) holds because $\frac{z}{\sqrt{m}} - pz \leq 0$ (which holds since $m \geq p^{-2}$), (30) holds by Fact 5, and (31) holds because $\alpha \geq 3$ and $m \geq 36$. □

Corollary 2. *For every $m \geq \max\{36, p^{-2}\}$ and every $z \geq 0$, it holds that $H(z) \leq e^{6(1+\lambda)}\left(\frac{1}{\sqrt{m}} + \frac{1}{m}\right)$.*

Proof. Let $\alpha = 3$. Lemma 9 proves the claim for when $z \leq 2\alpha(1+\lambda)\sqrt{m}$. On the other hand, when $z \geq 2\alpha(1+\lambda)\sqrt{m}$, then Lemma 10 implies that

$$H(z) \leq \frac{1}{\sqrt{m}} - \frac{1}{m}e^{\frac{z}{\sqrt{m}}} < \frac{1}{\sqrt{m}}.$$ □

Theorem 3 (Large market lower concentration bound)**.** *There exist positive constants m_0, k_0, c_0 such that for every $m > m_0$ and $k > k_0$,*

$$\sum_{x=0}^{\lambda m - k\sqrt{m}} \pi_x \leq c_0 m k e^{-k}.$$

Proof. Let $m_0 = \max\{36, p^{-2}\}$ and $k_0 = \max\{6(1+\lambda), \log m\}$. By Lemma 10, for every $m \geq m_0$

and $z \geq k_0\sqrt{m}$, it holds that $H(z) \leq \frac{1}{\sqrt{m}} - \frac{1}{m}e^{\frac{z}{\sqrt{m}}}$. Hence, for every $k > k_0$ and $m > m_0$,

$$H(k\sqrt{m}) \leq \frac{1}{\sqrt{m}} - \frac{1}{m}e^k \leq \frac{1}{\sqrt{m}} - 1 < 0.$$

Consequently, Corollary 1 implies that for every $m > m_0$ and $k > k_0$,

$$\sum_{(x,y)\notin B_k} \rho_{x,y} \leq -\frac{\sup_{0\leq z\leq k\sqrt{m}} H(z)}{\sup_{z\geq k\sqrt{m}} H(z)} \cdot \frac{2(1+\lambda+k/\sqrt{m})}{2+\lambda}. \tag{32}$$

By Corollary 2, for every $m > m_0$ and $z \geq 0$, $H(z) \leq e^{6(1+\lambda)}$ holds, which implies that

$$\sup_{0\leq z\leq k\sqrt{m}} H(z) \leq e^{6(1+\lambda)}. \tag{33}$$

On the other hand, Lemma 10 implies that for every $m > m_0$, $k > k_0$, and $z \geq k\sqrt{m}$, it holds that $H(z) \leq \frac{1}{\sqrt{m}} - \frac{1}{m}e^{\frac{z}{\sqrt{m}}}$, which implies that

$$\sup_{z\geq k\sqrt{m}} H(z) \leq \frac{1}{\sqrt{m}} - \frac{1}{m}e^k \leq -\frac{1}{2m}e^k, \tag{34}$$

where the last inequality holds because $\frac{1}{\sqrt{m}} \leq \frac{1}{2m}e^k$ holds when $m > m_0$. We next observe that (32), (33), and (34) together imply that

$$\sum_{(x,y)\notin B_k} \rho_{x,y} \leq \frac{e^{6(1+\lambda)}}{\frac{1}{2m}e^k} \cdot \frac{2(1+\lambda+k/\sqrt{m})}{2+\lambda} \leq 2e^{6(1+\lambda)}me^{-k}4k = 8e^{6(1+\lambda)}kme^{-k}$$

where the last inequality holds because $\frac{2(1+\lambda+k/\sqrt{m})}{2+\lambda} \leq 4k$ for every $k > k_0$.

The above bound, together with Lemma 4, implies that

$$\sum_{x=0}^{\lambda m-k\sqrt{m}} \pi_x \leq \sum_{x=0}^{\lambda m-k\sqrt{m}} \rho_x \leq \sum_{(x,y)\notin B_k} \rho_{x,y} \leq 8e^{6(1+\lambda)}kme^{-k}.$$

Setting $c_0 = 8e^{6(1+\lambda)}$ concludes the proof. □

We next provide an upper concentration bound for the greedy policy.

D.3 An Upper Concentration Bound for the Greedy Policy

In this section we complement Theorem 3 by providing an upper concentration bound for the greedy process. The main idea is defining a simpler process that is coupled with $\mathcal{M}$, namely $\mathcal{M}_{\mathfrak{u}}$. This process is defined such that the number of unmatched H agents in $\mathcal{M}$ is stochastically dominated by the number of unmatched H agents in $\mathcal{M}_{\mathfrak{u}}$.

D.3.1 Definition of $\mathcal{M}_{\mathfrak{u}}$

Consider a Markov process which is the same as $\mathcal{M}$ but with the following differences:

1. Matches are made only between H agents and E agents.
2. Matches are made greedily only upon the arrival of E agents.
3. E agents do not stay in the market: If upon the arrival of an E agent she is not matched to an H agent, then the E agent departs the market immediately.

The number of H agents in this random process is a one-dimensional Markov chain $\mathcal{M}_{\mathfrak{u}}$. The state space of $\mathcal{M}_{\mathfrak{u}}$ is $V(\mathcal{M}_{\mathfrak{u}}) = \{0, 1, 2, \ldots\}$. The Markov chain is at state x when the number of H agents present in the pool is x. The transition rate from a state x to state $x+1$ is $r_x = m(1+\lambda)$, since H agents arrive at a rate of $m(1+\lambda)$. The transition rate from a state x to state $x-1$ (when it exists) is $l_x = m(1-N_x) + x$. The first summand corresponds to the event that an arriving E agent is compatible to an H agent, and the second corresponds to the departures of H agents.

Since $\mathcal{M}_{\mathfrak{u}}$ is irreducible and positive recurrent, it has a unique stationary distribution (Norris (1997), Theorem 3.5.3), which we denote by $\pi^{\mathfrak{u}}$. Let $\pi^{\mathfrak{u}}_i$ be the probability that $\pi^{\mathfrak{u}}$ assigns to state i.

Lemma 11. *The steady state distribution of the number of H agents in $\mathcal{M}_{\mathfrak{u}}$ stochastically dominates the steady-state distribution of the number of H agents in $\mathcal{M}$.*

The proof of the lemma is technical and is deferred to the online appendix, Section i. The proof idea is defining a *coupling* of $\mathcal{M}$ and $\mathcal{M}_{\mathfrak{u}}$ such that, in the coupled process, there are fewer H agents in the pool under $\mathcal{M}$ than under $\mathcal{M}_{\mathfrak{u}}$ at any time, and more E agents.

Lemma 12. *Suppose that $m \geq 4(1+\lambda)^2$ and $p\lambda \geq \frac{\ln m}{2m}$. Then, for every positive integer k we have*

$$\pi^{\mathfrak{u}}_{\lambda m+k} \leq e^{-\frac{k}{\sqrt{m}}+(3+2\lambda)}.$$

Proof. For notational simplicity, we denote $\pi^{\mathfrak{u}}$ by η throughout this proof. We start by writing the balance equations, according to which, for every positive integer i, $\frac{\eta_i}{\eta_{i-1}} = \frac{r_i}{l_{i-1}}$.

Let $x^* = \lambda m$, and suppose $i \geq x^*$. Then,

$$\frac{\eta_i}{\eta_{i-1}} = \frac{(1+\lambda)m}{i + m(1-N_i)} = \frac{1+\lambda}{i/m + 1 - N_i}.$$

When $i \geq x^* + 2(1+\lambda)\sqrt{m} + me^{-p\lambda m}$, and $m \geq 4(1+\lambda)^2$, we can write

$$\frac{\eta_i}{\eta_{i-1}} = \frac{1+\lambda}{i/m + 1 - N_i} \leq \frac{1+\lambda}{1+\lambda+(i-x^*)/m - e^{-p\lambda m}} \leq \frac{1+\lambda}{1+\lambda+2(1+\lambda)/\sqrt{m}} \tag{35}$$

$$\leq \frac{1}{1+2/\sqrt{m}} \leq 1 - \frac{1}{\sqrt{m}} \tag{36}$$

where (35) uses the fact that $N_i \leq e^{-ip}$ and that $i \geq 2(1+\lambda)\sqrt{m} + me^{-p\lambda m}$, and (36) uses the fact that $\frac{1}{1+\alpha} \leq 1 - \alpha/2$ for every positive $\alpha \leq 1$.

For every positive integer k we have $\frac{\eta_{x^*+k}}{\eta_{x^*}} = \prod_{j=0}^{k-1} \frac{\eta_{x^*+j+1}}{\eta_{x^*+j}}$. Then, for every integer $k \geq 2(1+\lambda)\sqrt{m} + me^{-p\lambda m}$, we can use (36) to write

$$\eta_{x^*+k} \leq \frac{\eta_{x^*+k}}{\eta_{x^*}} \leq e^{-\frac{k-2(1+\lambda)\sqrt{m}-me^{-p\lambda m}}{\sqrt{m}}} \leq e^{-\frac{k}{\sqrt{m}}+(3+2\lambda)}, \tag{37}$$

where the last inequality holds because $p\lambda m \geq \frac{\ln m}{2}$.

On the other hand, for every non-negative integer $k < 2(1+\lambda)\sqrt{m} + me^{-p\lambda m}$, we have

$$\eta_{x^*+k} \leq 1 < e^{-\frac{k-2(1+\lambda)\sqrt{m}-me^{-p\lambda m}}{\sqrt{m}}} \leq e^{-\frac{k}{\sqrt{m}}+(3+2\lambda)}, \tag{38}$$

where the last inequality holds because $p\lambda m \geq \frac{\ln m}{2}$. Finally, (37) and (38) conclude the proof. □

Theorem 4 (Large market upper concentration bound)**.** *Suppose that $m \geq 4(1+\lambda)^2$ and $p\lambda \geq \frac{\ln m}{2m}$. Also, let π denote the steady state distribution of $\mathcal{M}$. Then, for every positive integer k we have*

$$\sum_{j=k}^{\infty} \pi_{\lambda m+j} \leq \frac{m}{\sqrt{m}-1} e^{-\frac{k}{\sqrt{m}}+3+2\lambda}.$$

Proof. By Lemma 12, for every positive integer j we have $\pi^{\mathfrak{u}}_{\lambda m+j} \leq e^{-\frac{j}{\sqrt{m}}+(3+2\lambda)}$. On the other hand, by Lemma 11, the steady-state distribution of the number of H agents in $\mathcal{M}_{\mathfrak{u}}$ stochastically dominates the steady-state distribution of the number of H agents in $\mathcal{M}$. Therefore,

$$\sum_{j=k}^{\infty} \pi_{\lambda m+j} \leq \sum_{j=k}^{\infty} \pi^{\mathfrak{u}}_{\lambda m+j} \leq \sum_{j=k}^{\infty} e^{-\frac{j}{\sqrt{m}}+(3+2\lambda)} = e^{3+2\lambda} \frac{e^{-\frac{k}{\sqrt{m}}}}{1-e^{-\frac{1}{\sqrt{m}}}} \leq e^{3+2\lambda} \frac{e^{-\frac{k}{\sqrt{m}}}}{\frac{1}{\sqrt{m}} - \frac{1}{m}},$$

where the last inequality holds because $e^z \leq 1 + z + z^2$ for every real number in $[-1, 1]$. □

E Match Rate and Waiting Time under the Greedy Policy

E.1 Match Rates under the Greedy Policy

Claim 5. *$q^G_E(m) \geq 1 - O(1/m)$.*

Proof. Fix an E agent, namely a, who has just arrived to the market. Let x_a, y_a respectively denote the size of the H pool and the E pool just before a arrives. By the PASTA property of the Poisson process,[38] the probability distribution of (x_a, y_a) is the same as the steady-state distribution π. Therefore, Theorem 3 implies that for sufficiently large m,

$$\mathbb{P}_\pi\left[x_a < m - 3\log m\sqrt{m}\right] \leq e^{-3\log m} 3c_0 m \log m \leq m^{-1}.$$

This implies that, upon her arrival, agent a has a compatible H agent with probability at least

$$(1 - m^{-1})\left(1 - (1-p)^{\lambda m - 3\log m\sqrt{m}}\right) = 1 - O(m^{-1}). \tag{39}$$

This probability is a lower bound for $q_E^G(m)$. □

Claim 6. $q_H^G(m) \in \left(\frac{1}{1+\lambda} - O(m^{-1/3}), \frac{1}{1+\lambda} + O(m^{-1/3})\right)$.

Proof. Fix $m > 0$. Suppose that the market starts at time 0 when there are no agents in the market. For any $t > 0$, let $m_E(t), m_H(t)$ respectively denote the number of E agents and H agents that arrive from time 0 to time t. Also, let $\psi_H(t)$ denote the number of H agents that are matched from time 0 to time t.

By the Ergodic theorem, $\lim_{t\to\infty} \psi_H(t)/m_H(t) = q_H^G(m)$. Hence, it suffices to prove the claim for the left-hand side of the equality. Since E agents arrive according to a Poisson process with rate m, then, for any $t > 1$, the event $m_E(t) \in [mt - (mt)^{2/3}, mt + (mt)^{2/3}]$ holds with very high probability. This holds by the concentration bound of Fact 3 for the Poisson distribution. Similarly, since H agents arrive according to a Poisson process with rate $(1+\lambda)m$, for any $t > 1$, the event

$$m_H(t) \in [(1+\lambda)mt - ((1+\lambda)mt)^{2/3}, (1+\lambda)mt + m^{2/3}]$$

holds with very high probability due to Fact 3. This implies that, for any $t > 1$,

$$\frac{\psi_H(t)}{m_H(t)} \leq \frac{m_E(t)}{m_H(t)} \leq \frac{mt + (mt)^{2/3}}{(1+\lambda)mt - ((1+\lambda)mt)^{2/3}} \tag{40}$$

holds with very high probability.

Next, we provide a lower bound for $\frac{\psi_H(t)}{m_H(t)}$. Recall from (39) that, upon her arrival, any E agent is matched to an H agent with probability at least $1 - O(m^{-1})$. Therefore,

$$\frac{\psi_H(t)}{m_H(t)} \geq \frac{m_E(t)(1 - O(m^{-1}))}{(1+\lambda)mt + ((1+\lambda)mt)^{2/3}} \geq \frac{(mt - (mt)^{2/3})(1 - O(m^{-1}))}{(1+\lambda)mt + ((1+\lambda)mt)^{2/3}}. \tag{41}$$

[38] PASTA, or Poisson Arrivals See Time Averages, is a well-known property in the queuing literature; e.g., see Harchol-Balter (2013).

Now observe that, for any $t > 1$, (40) and (41) together imply that

$$\frac{\psi_H(t)}{m_H(t)} \in \left(\frac{1}{1+\lambda} - O(m^{-1/3}), \frac{1}{1+\lambda} + O(m^{-1/3})\right)$$

holds with very high probability. Since, by the Ergodic theorem, $\lim_{t\to\infty} \psi_H(t)/m_H(t)$ exists and converges to $q_H^G(m)$ almost surely in any sample path, then we have

$$q_H^G(m) \in \left(\frac{1}{1+\lambda} - O(m^{-1/3}), \frac{1}{1+\lambda} + O(m^{-1/3})\right). \quad \square$$

Lemma 13. *Under the greedy policy, the match rate of hard-to-match agents $q_H^G(m)$ is $\frac{1}{1+\lambda} - o(1)$ and the match rate of easy-to-match agents $q_E^G(m)$ is $1 - o(1)$.*

Proof. The lemma follows immediately from Claim 5 and Claim 6. □

E.2 Distribution of Waiting Time under Greedy Matching

We show that as m approaches infinity, the waiting time for easy-to-match agents converges in distribution to the degenerate distribution at 0, and the waiting time for hard-to-match agents converges to the exponential distribution with rate $1/d + 1/\lambda$.

Lemma 14. *Under the greedy policy, as m approaches infinity, the waiting time of an easy-to-match agent converges in distribution to the degenerate distribution at 0.*

Proof. Fix an E agent, e, and let w_e denote the waiting time for e. For any fixed constant $t > 0$, we will show that $\lim_{m\to\infty} \mathbb{P}[t > w_e] = 1$. This will prove the claim. Recall from (39) that upon her arrival, agent e is matched to an H agent with probability $1 - O(m^{-1})$. Therefore, $\mathbb{P}[w_e = 0] = 1 - O(m^{-1})$, which implies that, $\lim_{m\to\infty} \mathbb{P}[t > w_e] = 1$ holds for any $t > 0$. □

Lemma 15. *As m approaches infinity, the waiting time of hard-to-match agents converges in distribution to the exponential distribution with rate $\frac{1}{d} + \frac{1}{\lambda}$.*

We sketch the proof below. The formal proof is technical and is presented in the online appendix, Section i. We will use $\mathsf{Exp}(x)$ to denote the exponential distribution with rate x.

Proof sketch. We define a new process, namely $\mathcal{P}$, in which there are no easy-to-match agents. Instead, *attach* an exponential clock to each hard-to-match agent which ticks at rate $1/\lambda$. We call this clock the *match clock* of the agent. We consider an agent to be matched if the match clock ticks before the agent becomes critical. Without providing a formal proof in this proof sketch, we suppose that H agents in the new process $\mathcal{P}$ have approximately the same waiting time as in the original process (the greedy policy). Given this assumption, we compute the distribution for the waiting time of a hard-to-match agent h in $\mathcal{P}$.

Consider the agent h and suppose she has entered the pool at time t_0. Note that h is matched if and only if it is matched before her criticality clock ticks. Let t_1, t_2 be random variables such that $t_1 \sim \mathsf{Exp}(1/\lambda), t_2 \sim \mathsf{Exp}(1/d)$. These random variables are interpreted as follows. The agent becomes critical at time $t_0 + t_2$ if she is not matched by then, i.e., if the match clock attached to her has not ticked by then. The agent's match clock ticks at time $t_0 + t_1$. So, the agent is matched if and only if $t_1 < t_2$. Alternatively, we can say the agent is matched if and only if $t_1 = t_{\min}$ where $t_{\min} = \min\{t_1, t_2\}$. Since $t_{\min}$ is distributed according to $\mathsf{Exp}(1/d + 1/\lambda)$, and since $t_{\min}$ equals the waiting time of the agent, the claim is proved. $\square$

E.3 Proof of Proposition 5

Proof of Proposition 5. The claim about the match rate was proved in Lemma 13, where we showed that under the greedy policy, the match rate of hard-to-match agents $q_H^G(m)$ is $\frac{1}{1+\lambda} - O(\frac{1}{(1+\lambda)\sqrt{m}})$ and the match rate of easy-to-match agents $q_E^G(m)$ is $1 - o(1)$. The claim about waiting times was proved in Lemma 14 and Lemma 15, for easy- and hard-to-match agents, respectively. $\square$

F Analysis of the Batching Policy

F.1 Preliminary Graph Theory Results

For every graph G, we let $V(G)$ denote the set of its nodes and $E(G)$ denote the set of its edges. An *independent* set in a graph G is a subset of nodes $S \subseteq V(G)$ such that no two nodes in S are adjacent (i.e., are connected by an edge) in G.

We denote a bipartite graph by $G(X, Y)$ where X, Y denote the set of nodes on each side of G. (That is, $V(G) = X \cup Y$, and both X, Y are independents sets in G.) A *matching* is a set of edges such that no two of the edges have a common node. The size of a matching is the number of the edges that it contains. A *perfect matching* in $G(X, Y)$ is a matching with size $\min\{|X|, |Y|\}$.

Lemma 16. *Let $G(X, Y)$ be a randomly drawn bipartite graph with non-random X, Y being its partitions where $|X| = |Y| = n$. Suppose that the probability that a node $u \in X$ is connected a node $v \in Y$ equals $p \in (0, 1)$ independently across all pairs (u, v). Then, the graph contains a perfect matching with probability at least $1 - n2^{2n}p^{n^2/4}$.*

Proof. By the König-Egerváry Theorem, there exists a perfect matching in G if and only if the size of the maximum independent set is at most n (West, 2000).

For $X' \subseteq X$ and $Y' \subseteq Y$, let $E(X', Y')$ denote the event in which no node in X' is adjacent to a node in Y'. We note that if $E(X', Y')$ happens then $X' \cup Y'$ is an independent set. Also, let the set $\mathcal{E}'$ be the set of all pairs (X', Y') such that $X' \subseteq X, Y' \subseteq Y$, and $|X'| + |Y'| = n + 1$. Therefore, $\bigcup_{(X',Y') \in \mathcal{E}'} E(X', Y')$ is the event that there exists an independent set of size larger than n (which also means that no perfect matching exists). By a union bound, the probability that a perfect matching does not exist in G is then at most

$$\mathbb{P}\left[\bigcup_{(X',Y')\in\mathcal{E}'} E(X',Y')\right] \le \sum_{(X',Y')\in\mathcal{E}'} \mathbb{P}\left[E(X',Y')\right]. \tag{42}$$

Consider (X',Y') with $|X'|+|Y'| = n+1$, and let $i = |X'|$. The probability that $E(X',Y')$ holds then equals $p^{i(n-i+1)}$. Therefore,

$$\begin{aligned}\sum_{(X',Y')\in\mathcal{E}'} \mathbb{P}\left[E(X',Y')\right] &= \sum_{i=1}^{n} \binom{n}{i}\binom{n}{n-i+1}(1-p)^{i(n-i+1)} = \sum_{i=1}^{n} \binom{n}{i}\binom{n}{i-1}(1-p)^{i(n-i+1)} \\ &\le \sum_{i=1}^{n} 2^{2n}(1-p)^{i(n-i+1)} \le n2^{2n}(1-p)^{n^2/4},\end{aligned}$$

where the penultimate inequality follows from the fact that the product of the binomial coefficients $\binom{n}{i}\binom{n}{i-1}$ is bounded by 2^{2n}, as each of the multiplicands is bounded by 2^n. $\square$

Corollary 3 (Corollary of Lemma 16)**.** *Let $G(X,Y)$ be a random bipartite graph with X,Y being its partitions where $|Y| = n$ and $|X| < |Y|$. A node $u \in X$ is adjacent to a node $v \in Y$ with probability p, independently across all pairs (u,v). Then, G contains a matching of size $|X|$ with probability at least $1 - n2^{2n}p^{n^2/4}$.*

Proof. Construct a graph H from G by adding $n - |X|$ dummy nodes to X. Let each dummy node x' and each node $y \in Y$ be adjacent independently with probability p.

Let $\mathfrak{p}$ denote the probability that H contains a matching of size n, and $\mathfrak{q}$ denote the probability that G contains a matching of size $|X|$. As any matching of size $n = |Y|$ in H must cover every node in X, then $\mathfrak{q} > \mathfrak{p}$. By Lemma 16, $\mathfrak{p} \ge 1 - n2^{2n}p^{n^2/4}$. This concludes the proof. $\square$

Definition 13. *In a bipartite graph $G(X,Y)$, a subset $S \subseteq X \cup Y$ is called an (x,y)-independent set if S is an independent set in G such that $|S \cap X| = x$ and $|S \cap Y| = y$.*

Lemma 17. *Let $\alpha, \beta, \gamma > 0$ be arbitrary constants such that $\alpha, \beta \in (0,1)$. Let $G(X,Y)$ be a bipartite graph such that $|Y| = \gamma|X|$ and, furthermore, for every pair of nodes $u \in X$ and $v \in Y$, u is adjacent to v independently with probability $p > 0$. Then, with high probability as $|X|$ grows large, G contains no $(\alpha|X|, \beta|Y|)$-independent set.*

Proof. Consider arbitrary subsets $X' \subseteq X$ and $Y' \subseteq Y$ such that $|X'| = \alpha|X|$ and $|Y'| = \beta|Y|$. The probability that $X' \cup Y'$ is an independent set is $(1-p)^{\alpha\beta|X|\cdot|Y|}$. Hence, the probability that there exists at least one $(|X'|, |Y'|)$-independent set in G is bounded by

$$\binom{|X|}{|X'|}\binom{|Y|}{|Y'|}(1-p)^{\alpha\beta|X|\cdot|Y|} \le 2^{|X|(1+\gamma)}(1-p)^{\alpha\beta\gamma|X|^2}.$$

The right-hand side of the above inequality approaches 0 as $|X|$ approaches infinity. $\square$

F.2 Preliminary Definitions and Lemmas

In the analysis, we suppose that time is indexed by non-negative real numbers. The batching policy makes matches at times iT for every positive integer i; these times are called *matching times*. For every $i \geq 0$ the interval $(iT, (i+1)T]$ is called *period* i.

The batching policy *executes* a matching at the *end* of every period i: at time $(i+1)T$, it finds the largest matching in the pool. If there are several such matchings, it selects the matching among them which has the maximum number of H agents. The policy then executes the selected matching and removes the agents involved in that matching from the pool.

For any matching time t, x_t and y_t respectively denote the number of H agents and the number of E agents in the pool after the execution of the matching at time t. If t is not a matching time, let x_t and y_t denote the number of H and E agents in the pool at time t, respectively.

We note that the sequence $\langle(x_{iT}, y_{iT})\rangle_{i\geq 0}$ is a discrete time Markov chain with a state space $\mathbb{Z}_+^2$. Since this Markov chain is Ergodic, it has a steady-state distribution. Since we are performing a steady-state analysis, we suppose that (x_0, y_0) is drawn from the steady state distribution. This assumption is without loss of generality by the Ergodic theorem for Markov chains.

Lemma 18. *The match rate of agents of type* Θ *equals* $1 - w_\Theta^\tau(m)/d$ *if* τ *is either the batching or greedy policy.*

The proof of Lemma 18 is identical to the proof of Lemma 1 adjusting for the fact that instead of an upper bound on the match rate we know its exact value; i.e., the probability that the agent is matched is given by $1 - \frac{1}{d}\mathbb{E}_{\alpha_i}[\varphi_i] = \mathbb{E}_{\alpha_i}[\mu_i]$.

Lemma 19. *Consider a time interval* (a, b) *and let* $c = b - a$. *Conditional on an agent arriving in the interval* (a, b)*, the criticality time of the agent is larger than* b *with probability* $\gamma_{c,d} = \frac{1-e^{-c/d}}{c/d}$.

Proof. By the properties of the Poisson process, the distribution of the arrival time of an agent conditional on the agent arriving in an interval $[a, b]$ equals the uniform distribution over the interval $[a, b]$. Hence, the chance that the criticality time of the agent is larger than b equals

$$\int_0^c \frac{1}{c} e^{-(c-s)/d} \mathrm{d}s = \frac{d(1 - e^{-c/d})}{c}. \qquad \square$$

F.3 Analysis of Match Rate

We first provide an upper bound for match rate, and then a matching lower bound for it.

Lemma 20. *Under a batching policy with batch length* T, $q_E^B(m) \leq \gamma_{T,d}$ *and* $q_H^B(m) \leq \frac{\gamma_{T,d}}{1+\lambda}$.

Proof. Let $i > 0$ be an arbitrary integer. Conditional on an E agent arriving at a time in the interval $(iT, (i+1)T]$, the agent is present in the pool at time $(i+1)T$ with probability $\gamma_{T,d}$ by Lemma 19. Therefore, the match rate of E agents is at most $\gamma_{T,d}$.

To prove the claim for H agents, observe that H agents can be matched only to E agents. Under the batching policy, only a fraction $\gamma_{T,d}$ of the E agents would not become critical before the first matching time after their arrival. Hence, the match rate of H agents is at most $\frac{\gamma_{T,d}}{1+\lambda}$. □

To provide lower bounds on match rate, we need the following definitions and lemmas.

Definition 14. *For an integer $i \geq 0$, an agent present in the pool at time $(i+1)T$ is called a* new *agent if she has arrived later than time iT.*

Definition 15. *In a graph G whose nodes correspond to E and H agents, an edge (u, v) is a* cross-edge *if u, v are agents of different types.*

Lemma 21. *Let e_i denote the number of new E agents who are present in the pool at time $(i+1)T$. Then, the matching executed at the matching time $(i+1)T$ involves at least e_i cross-edges, whp.*

Proof. We first construct a bipartite graph $G(X, Y)$, where X and Y respectively denote the set of H agents in the pool at time $(i+1)T$ before the matching is executed, and the set of new E agents in the pool at time $(i+1)T$ before the matching is executed.

Claim 7. *Whp, it holds that $|Y| < (1 + \lambda/2)\gamma_{T,d}mT < |X|$.*

Proof. By Lemma 19, conditional on an agent arriving to the pool after time iT, that agent remains in the pool until time $(i+1)T$ with probability $\gamma_{T,d}$, independently. (The independence is due to the independence of the criticality times.) Therefore, the the random variable $|Y|$ is a Poisson random variables with mean $\gamma_{T,d}mT$. This holds because E agents arrive with rate m but are present in the pool in the next batching time after their arrival only with probability $\gamma_{T,d}$. This fact, together with the concentration bound of Fact 3 for the Poisson distribution, implies that $|Y| < (1 + \lambda/2)\gamma_{T,d}mT$ holds whp.

Let X' denote the set of new H agents in the pool at time $(i+1)T$ before the matching is executed. By Lemma 19, $|X'|$ is a Poisson random variables with mean $\gamma_{T,d}(1+\lambda)mT$. (This holds by the same argument for the case of E agents, with the difference that the arrival rate of H agents is $(1+\lambda)m$.) This fact, together with the concentration bound of Fact 3 for the Poisson distribution, implies that $|X'| > (1 + \lambda/2)\gamma_{T,d}mT$ holds whp. Since $X' \subseteq X$, therefore, $|X| > (1 + \lambda/2)\gamma_{T,d}mT$ holds whp. The proof is complete. □

Recall that $e_i = |Y|$. By Claim 7 and Corollary 3, whp there exists a matching of size $|Y|$ in G. Given this fact, the next claim concludes the proof.

Claim 8. *If there exists a matching of size $|Y|$ in G, then the matching executed at time $(i+1)T$ involves at least $|Y|$ cross-edges.*

Proof. Let M denote a matching of size $|Y|$ in G. Let M' denote the maximum matching chosen by the batching policy to be executed at time $(i+1)T$. We construct a graph, G', where $V(G')$ is the set of all of the agents present in the pool at time $(i+1)T$, before the matching is executed, and $E(G') = E(M) \cup E(M')$. Thus, G' must be a union of paths and cycles (West, 2000). Let C, P_e, P_o respectively denote the set of cycles, the set of paths of even length, and the set of paths of odd length in G'.

For a subgraph F of G', let $D(F)$ denote the set of cross-edges of F. For two subgraphs F_1, F_2 of G', let $F_1 \Delta F_2$ denote the subgraph of G' with the set of edges $E(F_1) \cup E(F_2) - (E(F_1) \cap E(F_2))$.

First, we show that for every cycle or path of even length $Z \in C \cup P_e$,

$$|D(Z) \cap E(M)| \leq |D(Z) \cap E(M')|. \tag{43}$$

Suppose not. Then, observe that $M' \Delta Z$ would be a matching with the same size as M' but a larger number of cross-edges. This contradicts the definition of M'. Hence (43) must hold.

Next, we show that for every path of odd length $Z \in P_o$,

$$|D(Z) \cap E(M)| \leq |D(Z) \cap E(M')|. \tag{44}$$

To see why, first note that the first and last edges in Z must belong to M'. Otherwise $M' \Delta Z$ would be a matching with a larger size than M', which would be a contradiction. Now, consider a cross-edge (e, h) belonging to both Z and M, where e, h are respectively E and H agents. Since the first and last edges in Z must belong to M', then there must exist a cross-edge (e', h) belonging to M'. This means that for every cross-edge (e, h) belonging to both Z and M, there exists a cross-edge (e', h) belonging to both Z and M'. Therefore, (44) holds.

Finally, (43) and (44) together imply that the number of cross-edges in M' is at least as large as the number of cross-edges in M, which equals $|Y|$. ☐

This completes the proof of Lemma 21. ☐

Lemma 22. *For a batching policy with batch length T, $q_E^B(m) \geq \gamma_{T,d} - o(1)$ and $q_H^B(m) \geq \frac{\gamma_{T,d}}{1+\lambda} - o(1)$.*

Proof. Recall that e_i denotes the number of E agents who arrived after time iT and are present in the pool at time $(i+1)T$ before the execution of the matching. By Lemma 19, for every integer $i \geq 0$ we have that $\mathbb{E}[e_i] = \gamma_{T,d} mT$. By Lemma 21, the matching executed at the matching time $(i+1)T$ involves at least e_i cross-edges whp. Therefore, the expected number of E agents matched in every executed matching is at least $\gamma_{T,d} mT(1 - o(1))$, which is also a lower bound on the expected number of matched H agents. The following bounds thus hold for the match rates of E and H agents:

$$q_E^B(m) \geq \frac{1}{mT} \gamma_{T,d} mT(1 - o(1)) = \gamma_{T,d}(1 - o(1)),$$
$$q_H^B(m) \geq \frac{1}{(1+\lambda)mT} \gamma_{T,d} mT(1 - o(1)) = \frac{1}{1+\lambda} \gamma_{T,d}(1 - o(1)).$$ ☐

Proposition 10. *For a fixed batching policy with batch length T, the match rates of E agents and H agents as m grows large are $q_E^B = \gamma_{T,d}$ and $q_H^B = \frac{\gamma_{T,d}}{1+\lambda}$, respectively.*

Proof. For every agent type, Lemma 20 provided an upper bound for the match rate of agents of that type, and Lemma 22, provided a matching lower bound. The upper and lower bounds are $\gamma_{T,d}$ for E agents, and $\frac{\gamma_{T,d}}{1+\lambda}$ for H agents. This concludes the proof. □

Recall that $\gamma_{T,d} = \frac{1-e^{-T/d}}{T/d}$. The above proposition directly proves the claim of part ii of Proposition 3 about the match rates under the batching policy.

F.4 Analysis of Waiting Time

Lemma 23. *For every agent type $\Theta \in \{E, H\}$, $w_\Theta^B(m) = d(1 - q_\Theta^B(m))$.*

Proof. The proof follows directly from Lemma 18. By that lemma, $q_\Theta^B(m) = 1 - \frac{w_\Theta^B(m)}{d}$. Rearranging the equality implies that $w_\Theta^B(m) = d(1 - q_\Theta^B(m))$. □

G Proofs for Proposition 3 and Theorem 1

Proof of Proposition 3. We analyzed the match rate and waiting time under the greedy and batching policies respectively in Sections E and F. In particular, the claims about the match rate and waiting time under greedy policy (i.e., part (i) of the proposition) were proved in Sections E.1 and E.2, respectively. The claims about the match rate and waiting time under the batching policy (i.e., part (ii) of the proposition) were proved in Sections F.3 and F.4, respectively. The analysis of the patient policy is deferred to the online appendix, Section v. The claims about the match rate and waiting time under the patient policy (i.e., part (iii) of the proposition) are proved there. □

Proof of Theorem 1. In Proposition 4 we showed that, under any policy, the match rate of hard-to-match agents is at most $\frac{1}{1+\lambda}$ and their expected waiting time is at least $\frac{\lambda d}{1+\lambda}$. On the other hand, in part (i) of Proposition 3 we showed that as m approaches infinity, the match rates of hard- and easy-to-match agents under the greedy policy approach $(q_H^G, q_E^G) = (\frac{1}{1+\lambda}, 1)$, respectively, and their expected waiting times approach $(w_H^G, w_E^G) = (\frac{\lambda d}{1+\lambda}, 0)$. This proves the first part of the theorem about the optimality of the greedy policy. It remains to show that the batching and patient policies are not asymptotically optimal.

Claim 9. *For every $T > 0$, $\frac{1-e^{-T/d}}{T/d} < 1$.*

Proof. For all $z > 0$, $1 - z < e^{-z}$. Setting $z = T/d$ and rearranging the terms proves the claim. □

Recall that by part (ii) of Proposition 3, as m grows large a batching policy with batch length $T > 0$ achieves match rates of $(q_H^B, q_E^B) = (\frac{1-e^{-T/d}}{(1+\lambda)T/d}, \frac{1-e^{-T/d}}{T/d})$, for hard- and easy-to-match agents, respectively. This fact, together with Claim 9, implies that $q_H^B < \frac{1}{1+\lambda} = q_H^G$ and $q_E^B < 1 = q_E^G$. This proves the claim about the sub-optimality of the batching policy.

To show that the patient policy is is not asymptotically optimal, we recall part (iii) of Proposition 3, which shows that the expected waiting time of hard-to-match pairs under the patient policy approaches $w_H^P = d$ as m grows large. On the other hand, under the greedy policy, the expected waiting time of hard-to-match agents approaches $w_H^G = \frac{\lambda d}{1+\lambda}$ as m grows large, which is strictly smaller than w_H^P. Therefore, the patient policy is not optimal. □

References

Agarwal, Nikhil, Itai Ashlagi, Eduardo Azevedo, Clayton R Featherstone, and Ömer Karaduman, "Market failure in kidney exchange," *American Economic Review*, 2019, *109* (11), 4026–70.

Akbarpour, Mohammad, Julien Combe, Yinghua He, Victor Hiller, Robert Shimer, and Olivier Tercieux, "Unpaired Kidney Exchange: Overcoming Double Coincidence of Wants without Money," 2019.

_ , Shengwu Li, and Shayan Oveis Gharan, "Thickness and information in dynamic matching markets," *Journal of Political Economy*, 2020, *128* (3), 783–815.

Anderson, Ross, Itai Ashlagi, David Gamarnik, and Yash Kanoria, "Efficient Dynamic Barter Exchange," *Operations Research*, 2017, *65* (6), 1446–1459.

Aquilina, Matteo, Eric B Budish, and Peter O'Neill, "Quantifying the high-frequency trading "arms race": A simple new methodology and estimates," *Chicago Booth Research Paper*, 2020, (20-16).

Ashlagi, I., M. Burq, P. Jaillet, and V. Manshadi, "On Matching and Thickness in Heterogeneous Dynamic Markets," in "Proceedings of the 2016 ACM Conference on Economics and Computation" ACM 2016, pp. 765–765.

_ , P. Jaillet, and V.H. Manshadi, "Kidney exchange in dynamic sparse heterogenous pools," in "Proceedings of the fourteenth ACM conference on Electronic commerce" ACM 2013, pp. 25–26.

Ashlagi, Itai, Adam Bingaman, Maximilien Burq, Vahideh Manshadi, David Gamarnik, Cathi Murphey, Alvin E Roth, Marc L Melcher, and Michael A Rees, "Effect of match-run frequencies on the number of transplants and waiting times in kidney exchange," *American Journal of Transplantation*, 2018, *18* (5), 1177–1186.

_ , Duncan. S. Gilchrist, Alvin. E. Roth, and Michael. A. Rees, "Nead Chains in Transplantation," *American Journal of transplantation*, 2011, *11* (12), 2780–2781.

_ , Maximilien Burq, Chinmoy Dutta, Patrick Jaillet, Amin Saberi, and Chris Sholley, "Edge Weighted Online Windowed Matching," in "Proceedings of the 2019 ACM Conference on Economics and Computation" ACM 2019, pp. 729–742.

Baccara, M., S. Lee, and L. Yariv, *Theoretical Economics*, 2019, *Forthcoming.*

Biro, P., L. Burnapp, H. Bernadette, A. Hemke, R. Johnson, J. van de Klundert, and D. Manlove, "First Handbook of the COST Action CA15210: European Network for Collaboration on Kidney Exchange Programmes (ENCKEP)," 2017.

Blanchet, Jose H, Martin I Reiman, Viragh Shah, and Lawrence M Wein, "Asymptotically optimal control of a centralized dynamic matching market with general utilities," *arXiv preprint arXiv:2002.03205*, 2020.

Brémaud, Pierre, *Markov chains: Gibbs fields, Monte Carlo simulation, and queues*, Vol. 31, Springer Science & Business Media, 2013.

Brmaud, Pierre, *Discrete Probability Models and Methods*, Springer, 2017.

Canonne, Clement, "A short note on Poisson tail bounds," 2019.

Corless, Robert M, Gaston H Gonnet, David EG Hare, David J Jeffrey, and Donald E Knuth, "On the LambertW function," *Advances in Computational mathematics*, 1996, *5* (1), 329–359.

Doval, Laura, "A theory of stability in dynamic matching markets," Technical Report, mimeo 2014.

Ferrari, Paolo, Willem Weimar, Rachel J Johnson, Wai H. Lim, and Kathryn J. Tinckam, "Kidney paired donation: principles, protocols and programs," *Nephrology Dialysis Transplantation*, 2014, *30* (8), 1276–1285.

Frieze, Alan and Michal Karoński, *Introduction to random graphs*, Cambridge University Press, 2015.

Gentry, Sommer Elizabeth and DL Segev, "The Best-Laid Schemes of Mice and Men Often Go Awry; How Should We Repair Them?," *American Journal of Transplantation*, 2015, *15* (10), 2539–2540.

Harchol-Balter, M., *Performance Modeling and Design of Computer Systems: Queueing Theory in Action*, 1st ed., New York: Cambridge University Press, 2013.

Held, Philip J, Frank McCormick, Akinlolu Ojo, and John P Roberts, "A cost-benefit analysis of government compensation of kidney donors," *American Journal of Transplantation*, 2016, *16* (3), 877–885.

Khalil, Hassan K, "Lyapunov stability," *Control Systems, Robotics and Automation–Volume XII: Nonlinear, Distributed, and Time Delay Systems-I*, 2009, p. 115.

Li, Zhuoshu, Kelsey Lieberman, William Macke, Sofia Carrillo, Chien-Ju Ho, Jason Wellen, and Sanmay Das, "Incorporating compatible pairs in kidney exchange: A dynamic weighted matching model," in "Proceedings of the 2019 ACM Conference on Economics and Computation" 2019, pp. 349–367.

Little, John DC and Stephen C Graves, "Little's law," in "Building intuition," Springer, 2008, pp. 81–100.

Liu, Tracy Xiao, Zhixi Wan, and Chenyu Yang, "The Efficiency of A Dynamic Decentralized Two-sided Matching Market," *Working paper*, 2018.

Mertikopoulos, Panayotis, Heinrich H Nax, and Bary Pradelski, "Quick or cheap? Breaking points in dynamic markets," *Breaking points in dynamic markets (December 30, 2019)*, 2019.

Nikzad, Afshin, Mohammad Akbarpour, Michael A Rees, and Alvin E Roth, "Financing transplants' costs of the poor: A dynamic model of global kidney exchange," 2019.

Norris, J. R., *Markov Chains* Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, 1997.

Roth, Alvin E., Tayfun Sönmez, and M. Utku Ünver, "Kidney exchange," *The Quarterly Journal of Economics*, 2004, *119* (2), 457–488.

_ , _ , and _ , "Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences," *The American economic review*, 2007, pp. 828–851.

Saidman, Susan. L., Alvin. E. Roth, Tayfun Sönmez, M. Utku Ünver, and Francis. L. Delmonico, "Increasing the opportunity of live kidney donation by matching for two-and three-way exchanges," *Transplantation*, 2006, *81* (5), 773–782.

Ünver, M. Utku, "Dynamic Kidney Exchange," *Review of Economic Studies*, 2010, *77(1)*, 372–414.

West, Douglas B., *Introduction to Graph Theory*, 2 ed., Prentice Hall, September 2000.

Online Appendix

i Remaining Proofs from the Analysis of the Greedy Policy

i.1 Proof of Lemma 2

By the ergodic theorem (Norris (1997), Theorem 3.5.3), $\mathcal{M}$ has a unique steady-state distribution if and only if it is irreducible and has at least one positive recurrent state. Irreducibility holds because there is only one communication class; i.e., it is possible to reach from any state to every other state.

To prove that $\mathcal{M}$ has at least one positive recurrent state, we need to show that there exists a state (x_0, y_0) with a finite expected return time. (The expected return time of a state is the expected time that it takes for the Markov chain to return to that state conditional on starting from that state.) We will prove that the state $(x_0, y_0) = (0, 0)$ has a finite expected return time.

To this end, consider an auxiliary stochastic process which is the same process as the greedy process with the difference that it makes no matches. So, the sequence of the agents' arrivals and their criticality times are the same in the auxiliary process and the greedy process, but agents are never matched in the auxiliary process.

We first observe that the expected return time of the state $(0, 0)$ is larger in the auxiliary process than in the greedy process (because making matches only increases the chance of visiting the state $(0, 0)$). In every sample path (i.e., a sequence of arrival times of agents and their criticality times), the auxiliary process returns to the state $(0, 0)$ no sooner than the greedy process. Hence, to prove the claim, we show that the expected return time of the state $(0, 0)$ is finite in the auxiliary process.

To show this, let $\mathcal{M}'$ denote the Markov chain corresponding to the auxiliary process. Hence, every state in $\mathcal{M}'$ is a pair of non-negative integers (x, y). We note that $\mathcal{M}'$ consists of two Markov processes that are run independently: one corresponding to the number of H agents (x) and the other one corresponding to the number of E agents (y). We use $\mathcal{M}'_H$ and $\mathcal{M}'_E$ to denote the Markov processes that keep track of the number of H agents and E agents in $\mathcal{M}'$, respectively.

Theorem 7 of Akbarpour et al. (2020) shows that $\mathcal{M}'_H$ and $\mathcal{M}'_E$ are positive recurrent. Hence, the expected return time of the state 0 is finite in both $\mathcal{M}'_H$ and $\mathcal{M}'_E$. Let R_H and R_E denote the expected return times of the state 0 in $\mathcal{M}'_H$ and $\mathcal{M}'_E$, respectively. By the independence of $\mathcal{M}'_H$ and $\mathcal{M}'_E$, the expected return time of the state $(0, 0)$ in $\mathcal{M}'$ equals $R_E R_H$, which is finite. This proves the claim.

i.2 Proof of Lemma 4

The proof works by coupling the Markov processes $\mathcal{M}$ and $\mathcal{N}$, i.e., constructing a probability space in which the outcomes of these two process are pathwise related to each other.

To define the coupling, we need a few definitions. An *event* is a tuple $e = (a, t, \zeta)$ such that a is an agent, $t \in \mathbb{R}_+$ is the time of the event, and $\zeta \in \{\mathsf{A}, \mathsf{C}\}$ specifies whether the event corresponds to the arrival ($\zeta = \mathsf{A}$) or the criticality ($\zeta = \mathsf{C}$) of agent a. Let $\mathfrak{e} = \langle e_1, e_2, \ldots \rangle$ denote a sequence of events defined as follows. For every agent that arrives to the market in the process $\mathcal{M}$, $\mathfrak{e}$ contains two events corresponding to the arrival time and the criticality time of a (regardless of whether agent a is matched). Hence, for every tuple (a, t, A) in the sequence $\mathfrak{e}$, there exists a tuple (a, t', C) in $\mathfrak{e}$ such that $t' > t$. We emphasize that the sequence $\mathfrak{e}$ merely determines the arrival times and criticality times of the agents (and is not concerned with the matches made between the agents).

We suppose that the events in $\mathfrak{e}$ are ordered according to their times: For every positive integer i, let t_i denote the time that event e_i occurs. We suppose that $t_1 \leq t_2 \leq \ldots$. For notational convenience, we define $t_0 = 0$. In the coupling, the arrival times and criticality times of the agents in both processes $\mathcal{M}$ and $\mathcal{N}$ are defined by the sequence $\mathfrak{e}$.

Let X_t and Y_t respectively denote the sets of agents in the H pool and in the E pool at time t in the process $\mathcal{M}$. Similarly, let X'_t and Y'_t respectively denote the sets of agents in the H pool and in the E pool at time at time t in the process $\mathcal{N}$. In the coupling, both $\mathcal{M}$ and $\mathcal{N}$ start at time zero with an empty pool (that is, $X_0 = Y_0 = X'_0 = Y'_0 = \varnothing$). We will define the coupling such that $X'_t \subseteq X_t$ and $Y_t \subseteq Y'_t$ would hold in every sample path of the coupling and for every time $t \geq 0$.

Definition of the coupling. Recall that the arrival times and criticality times of the agents in both processes $\mathcal{M}$ and $\mathcal{N}$ are determined by the same sequences of events, $\mathfrak{e}$, as described above. To define the coupling, we specify how, in the coupled process, $\mathcal{M}$ and $\mathcal{N}$ evolve as time moves forward. Consider the moment when the coupled process reaches time t_i corresponding to an event $e_i = (a, t_i, \zeta)$.

- **If a is an H agent and $\zeta = \mathsf{A}$:**
  - In the process $\mathcal{M}$:
    - (A1) If there exists an agent in the E pool that is compatible to a, then match a to such an agent, namely e; let $X_{t_i} = X_{t_{i-1}}$ and $Y_{t_i} = Y_{t_{i-1}} - \{e\}$.
    - (A2) Otherwise, a is added to the H pool; i.e., let $X_{t_i} = X_{t_{i-1}} \cup \{a\}$ and $Y_{t_i} = Y_{t_{i-1}}$.
  - In the process $\mathcal{N}$:
    - (B1) If in the process $\mathcal{M}$ agent a is matched to an agent e in the E pool—through item A1 above—then match a to e in $\mathcal{N}$ as well (this is possible because $Y_{t_{i-1}} \subseteq Y'_{t_{i-1}}$ holds by construction); in that case, let $X'_{t_i} = X'_{t_{i-1}} - \{e\}$.
    - (B2) Otherwise, if there exists an agent in the E pool that is compatible to a, then match a to such an agent, namely e'; let $X'_{t_i} = X'_{t_{i-1}}$ and $Y'_{t_i} = Y'_{t_{i-1}} - \{e'\}$. If there is no such agent e', then add a to the H pool, i.e., let $X'_{t_i} = X'_{t_{i-1}} \cup \{a\}$ and $Y_{t_i} = Y_{t_{i-1}}$.
- **If a is an E agent and $\zeta = \mathsf{A}$:**

– In the process $\mathcal{N}$:

(C1) If there exists an agent in the H pool that is compatible to a, then match a to such an agent, namely h; let $X'_{t_i} = X'_{t_{i-1}} - \{h\}$ and $Y'_{t_i} = Y'_{t_{i-1}}$.

(C2) Otherwise, a is added to the E pool; i.e., let $Y'_{t_i} = Y'_{t_{i-1}} \cup \{a\}$ and $X'_{t_i} = X'_{t_{i-1}}$.

– In the process $\mathcal{M}$:

(D1) If in the process $\mathcal{N}$ agent a is matched to an agent h in the H pool—through item C1 above—then match a to h in $\mathcal{M}$ as well (this is possible because $X'_{t_{i-1}} \subseteq X_{t_{i-1}}$ holds by construction); in that case, let $X_{t_i} = X_{t_{i-1}} - \{h\}$.

(D2) Otherwise, if there exists an agent in the H pool that is compatible to a, then match a to such an agent, namely h'; let $X_{t_i} = X_{t_{i-1}} - \{h'\}$ and $Y_{t_i} = Y_{t_{i-1}}$.

(D3) Otherwise (if neither D1 nor D2 hold), if there exists an agent in the E pool that is compatible to a, then match a to such an agent, namely e'; let $X_{t_i} = X_{t_{i-1}}$ and $Y_{t_i} = Y_{t_{i-1}} - \{e'\}$. If there exists no such agent, then a is added to the E pool; i.e., let $X_{t_i} = X_{t_{i-1}}$ and $Y_{t_i} = Y_{t_{i-1}} \cup \{a\}$.

- **If a is an H agent and $\zeta = \mathsf{C}$:** Remove a from the H pool in both $\mathcal{M}$ and $\mathcal{N}$ (if she has not been already removed). That is, let $X_{t_i} = X_{t_{i-1}} - \{a\}$ and $X'_{t_i} = X'_{t_{i-1}} - \{a\}$.

- **If a is an E agent and $\zeta = \mathsf{C}$:** Remove a from the E pool in both $\mathcal{M}$ and $\mathcal{N}$ (if she has not been already removed). That is, let $Y_{t_i} = Y_{t_{i-1}} - \{a\}$ and $Y'_{t_i} = Y'_{t_{i-1}} - \{a\}$.

We next show that the promised inclusion properties $X'_{t_i} \subseteq X_{t_i}$ and $Y_{t_i} \subseteq Y'_{t_i}$ hold for every integer $i \geq 0$. Recall that we defined $t_0 = 0$. The proof is by induction. The induction step is for the case of $i = 0$, and holds because $X_{t_0} = X'_{t_0} = Y_{t_0} = Y'_{t_0} = \varnothing$. For the induction hypothesis we assume that $X'_{t_{i-1}} \subseteq X_{t_{i-1}}$ and $Y_{t_{i-1}} \subseteq Y'_{t_{i-1}}$ hold for $i > 0$, and for the induction step we will show that $X'_{t_i} \subseteq X_{t_i}$ and $Y_{t_i} \subseteq Y'_{t_i}$ hold. To do this, consider the four cases that event e_i can have:

1. **e_i is the arrival of an H agent:** By the design of the coupling, if the arriving H agent is added to the pool in $\mathcal{N}$, then it is also added to the pool in $\mathcal{M}$. This, together with $X'_{t_{i-1}} \subseteq X_{t_{i-1}}$, implies that $X'_{t_i} \subseteq X_{t_i}$. On the other hand, if the arriving H agent is matched to an E agent in $\mathcal{N}$, then the H agent is matched to the same E agent in $\mathcal{M}$ if that E agent belongs to $Y_{t_{i-1}}$. This, together with $Y_{t_{i-1}} \subseteq Y'_{t_{i-1}}$, implies that $Y_{t_i} \subseteq Y'_{t_i}$.

2. **e_i is the arrival of an E agent:** By the design of the coupling, if the arriving E agent is added to the pool in $\mathcal{M}$, then it is also added to the pool in $\mathcal{N}$. This, together with $Y_{t_{i-1}} \subseteq Y'_{t_{i-1}}$, implies that $Y_{t_i} \subseteq Y'_{t_i}$. On the other hand, if the arriving E agent is matched to an H agent in $\mathcal{M}$, then the E agent is matched to the same H agent in $\mathcal{N}$ if that H agent belongs to $X'_{t_{i-1}}$. This, together with $X'_{t_{i-1}} \subseteq X_{t_{i-1}}$, implies that $X'_{t_i} \subseteq X_{t_i}$.

3. e_i **corresponds to the criticality of an H agent:** In this case, $X_{t_i} = X_{t_{i-1}} - \{e\}$ and $X'_{t_i} = X'_{t_{i-1}} - \{e\}$. This, together with $X'_{t_{i-1}} \subseteq X_{t_{i-1}}$, implies that $X'_{t_i} \subseteq X_{t_i}$. It also holds that $Y_{t_i} = Y_{t_{i-1}}$ and $Y'_{t_i} = Y'_{t_{i-1}}$. This, together with $Y_{t_{i-1}} \subseteq Y'_{t_{i-1}}$, implies that $Y_{t_i} \subseteq Y'_{t_i}$.

4. e_i **corresponds to the criticality of an E agent:** In this case, $Y_{t_i} = Y_{t_{i-1}} - \{e\}$ and $Y'_{t_i} = Y'_{t_{i-1}} - \{e\}$. This, together with $Y_{t_{i-1}} \subseteq Y'_{t_{i-1}}$, implies that $Y_{t_i} \subseteq Y'_{t_i}$. It also holds that $X_{t_i} = X_{t_{i-1}}$ and $X'_{t_i} = X'_{t_{i-1}}$. This, together with $X_{t_{i-1}} \subseteq X'_{t_{i-1}}$, implies that $X_{t_i} \subseteq X'_{t_i}$.

Thus, the we have completed the induction step in all of the four cases. This proves the promised inclusion property that $X'_{t_i} \subseteq X_{t_i}$ and $Y_{t_i} \subseteq Y'_{t_i}$ hold for all integers $i \geq 0$.

Next we note that for every $t \in [t_{i-1}, t_i)$, it holds that $X_t = X_{t_{i-1}}$ and $X'_t = X'_{t_{i-1}}$, as no agent arrives to or departs from the market in the time interval (t_{i-1}, t_i). Therefore, for every time $t \geq 0$, $X'_t \subseteq X_t$ holds in every sample path of the coupling. This implies that the number of H agents in $\mathcal{M}$ first-order stochastically dominates the number of H agents in $\mathcal{N}$ at the steady state.

i.3 Proof of Lemma 11

The proof works by coupling the Markov processes $\mathcal{M}$ and $\mathcal{M}_{\mathfrak{u}}$. To define the coupling, we need a few definitions.

An *event* is a tuple $e = (a, t, \zeta)$ such that a is an agent, t is a positive real number corresponding to the time of the event, and $\zeta \in \{\mathsf{A}, \mathsf{C}\}$ specifies whether the event corresponds to the arrival ($\zeta = \mathsf{A}$) or the criticality ($\zeta = \mathsf{C}$) of the agent a. Let $\mathfrak{e} = \langle e_1, e_2, \ldots \rangle$ denote a sequence of events defined as follows. For every agent that arrives to the market in the process $\mathcal{M}$, $\mathfrak{e}$ contains two events corresponding to the arrival time and the criticality time of a (regardless of whether agent a is matched or not). Hence, for every tuple (a, t, A) in the sequence $\mathfrak{e}$, there exists also a tuple (a, t', C) in $\mathfrak{e}$ such that $t' > t$. We emphasize that the sequence $\mathfrak{e}$ merely determines the arrival times and criticality times of the agents (and is not concerned with the matches made between the agents). We suppose that the events in $\mathfrak{e}$ are ordered according to their times. For every integer $i > 0$, let t_i denote the time of event e_i. Thus, $t_1 \leq t_2 \leq \ldots$. For notational convenience, let $t_0 = 0$.

In the coupling, the arrival times and criticality times of the agents in both processes $\mathcal{M}$ and $\mathcal{M}_{\mathfrak{u}}$ are defined by the sequence $\mathfrak{e}$. (We note that $\mathcal{M}_{\mathfrak{u}}$ ignores the events in $\mathfrak{e}$ that correspond to the criticality times of E agents, as E agents do not stay in the pool in $\mathcal{M}_{\mathfrak{u}}$.)

Let X_t and Y_t respectively denote the sets of agents in the H pool and in the E pool at time t in the process $\mathcal{M}$. Also, let X'_t denote the set of agents in the H pool at time t in the process $\mathcal{M}_{\mathfrak{u}}$. (We do not define a variable Y'_t, since the E pool is always empty in the process $\mathcal{M}_{\mathfrak{u}}$, where E agents depart the pool immediately after arrival whether they find a compatible match or not.)

In the coupling, both $\mathcal{M}$ and $\mathcal{M}_{\mathfrak{u}}$ start at time 0 with an empty pool (that is, $X_t = Y_t = X'_t = \varnothing$). We will define the coupling such that $X_t \subseteq X'_t$ would hold in every sample path of the coupling and for every time $t \geq 0$.

Definition of the coupling. Recall that the arrival times and criticality times of the agents in both processes $\mathcal{M}$ and $\mathcal{M}_\mathfrak{u}$ are determined by the same sequences of events, $\mathfrak{e}$, as described above. To define the coupling, we specify how, in the coupled process, $\mathcal{M}$ and $\mathcal{M}_\mathfrak{u}$ evolve as time moves forward. Consider the moment when the coupled process reaches time t_i corresponding to an event $e_i = (a, t_i, \zeta)$.

- **If a is an H agent and $\zeta = \mathsf{A}$:**
  - In the process $\mathcal{M}$, if there exists an agent in the E pool that is compatible to a, then match a to such an agent, namely a'; let $X_{t_i} = X_{t_{i-1}}$ and $Y_{t_i} = Y_{t_{i-1}} - \{a'\}$. Otherwise, a is added to the H pool; i.e., let $X_{t_i} = X_{t_{i-1}} \cup \{a\}$ and $Y_{t_i} = Y_{t_{i-1}}$.
  - In the process $\mathcal{M}_\mathfrak{u}$, let $X'_{t_i} = X'_{t_{i-1}} \cup \{a\}$.
- **If a is an E agent and $\zeta = \mathsf{A}$:**
  - In the process $\mathcal{M}$, if there exists no agent in the the H pool or in the E pool that is compatible to a, then a is added to the E pool; in that case, let $X_{t_i} = X_{t_{i-1}}$ and $Y_t = Y_{t_{i-1}} \cup \{a\}$. Otherwise:
    (i) If there exists an agent in the H pool that is compatible to a, then match a to such an agent, namely h; let $X_{t_i} = X_{t_{i-1}} - \{h\}$ and $Y_{t_i} = Y_{t_{i-1}}$.
    (ii) If there is no agent in the H pool that is compatible to a, then there must be an agent in the E pool that is compatible to a. Match a to such an agent, namely e; let $X_{t_i} = X_{t_{i-1}}$ and $Y_t = Y_{t_{i-1}} - \{e\}$.
  - In the process $\mathcal{M}_\mathfrak{u}$, if there exists no agent in the H pool that is compatible to a, then a departs the market immediately after arrival; in that case, $X'_{t_i} = X'_{t_{i-1}}$. Otherwise, there exists an agent in the H pool that is compatible to a. Then,
    (i') if in the process $\mathcal{M}$ agent a is matched to an agent h in the H pool—through item (i) above—then match a to h in $\mathcal{M}_\mathfrak{u}$ as well (this is possible because $X_{t_{i-1}} \subseteq X'_{t_{i-1}}$ holds by construction); in that case, let $X'_{t_i} = X'_{t_{i-1}} - \{h\}$;
    (ii') otherwise, there must be an agent in the H pool that is compatible to a. Match a to such an H agent, namely e'. Let $X'_{t_i} = X'_{t_{i-1}} - \{e'\}$.
- **If a is an H agent and $\zeta = \mathsf{C}$:** Remove a from the H pool in both $\mathcal{M}$ and $\mathcal{M}_\mathfrak{u}$ (if she has not been already removed). That is, let $X_{t_i} = X_{t_{i-1}} - \{a\}$, $X'_{t_i} = X'_{t_{i-1}} - \{a\}$, and $Y_{t_i} = Y_{t_{i-1}}$.
- **If a is an E agent and $\zeta = \mathsf{C}$:** Remove a from the E pool in $\mathcal{M}$ (if she has not been already removed). That is, let $Y_{t_i} = Y_{t_{i-1}} - \{a\}$. Also, let $X_{t_i} = X_{t_{i-1}}$ and $X'_{t_i} = X'_{t_{i-1}}$.

The promised inclusion property $X_{t_i} \subseteq X'_{t_i}$ holds by an inductive argument. The induction basis is $i = 0$, which holds by definition, since $X_{t_0} = X'_{t_0} = \varnothing$. The induction step shows that if

$X_{t_i} \subseteq X'_{t_i}$ holds, then $X_{t_{i+1}} \subseteq X'_{t_{i+1}}$ holds as well. This step is similar to the proof of the inclusion property for the coupling in Lemma 4.

Finally, we note that for every $t \in [t_{i-1}, t_i)$, it holds that $X_t = X_{t_{i-1}}$ and $X'_t = X'_{t_{i-1}}$, as no agent arrives to or departs from the market in the time interval (t_{i-1}, t_i). Hence, in every sample path of the coupling and for every time $t \geq 0$, it holds that $X_t \subseteq X'_t$. This implies that the number of H agents in $\mathcal{M}_{\mathrm{u}}$ stochastically dominates the number of H agents in $\mathcal{M}$.

i.4 Proof of Lemma 15

First, we need to define an additional notion for this proof. Recall the process of matching agents under the greedy policy. According to this process, when an agent a arrives to the market, the policy chooses a strict order over all H agents in the market; this order is chosen uniformly at random. Then, a is matched with the first compatible H agent according to the selected order.[39] If such an H agent does not exist, then then a strict order over all E agents in the market is selected uniformly at random, and a is matched with the first compatible E agent in that order if such an agent exists. When a finds a compatible (E or H) agent in this process, we say that a makes an *offer* to the compatible agent. In the greedy policy, an agent who receives an offer always *accepts* the offer, and hence, the match is always carried over. Throughout this proof, we assume that all offers are accepted unless we explicitly state otherwise.

We start the proof by showing that the waiting time of an H agent converges in distribution to $\mathsf{Exp}(1/d + 1/\lambda)$. To this end, fix an H agent, h, upon her arrival at time t_0. Let $\mathcal{Q}$ denote the stochastic process under the greedy policy starting from t_0 and ending when h leaves the market. We couple another process with $\mathcal{Q}$, namely $\mathcal{Q}'$. We define this coupling below. Loosely speaking, $\mathcal{Q}'$ is the same as the greedy matching process with the exception that the arrival of h is "ignored" in the sense that h does not interfere with the evolution of $\mathcal{Q}'$. Formally, the coupling is as follows.

- $\mathcal{Q}'$ runs from time t_0 to $t_0 + \log m$. Furthermore, we set the criticality clock of h to tick at time $t_0 + \log m$ in $\mathcal{Q}'$.

- If h finds a compatible match upon her arrival (at time t_0) in $\mathcal{Q}$, then we stop both $\mathcal{Q}, \mathcal{Q}'$. Otherwise, we let $\mathcal{Q}$ evolve according to the greedy process. The process $\mathcal{Q}'$ has a sample path identical to $\mathcal{Q}$ until one of the following disjoint events happens:

Event (i) In $\mathcal{Q}$, h receives an offer—and hence, is matched—before time $t_0 + \log m$. In this case, we stop $\mathcal{Q}$. In $\mathcal{Q}'$, h rejects the received offer as well as all offers she may receive in the future. Any E agent whose offer gets rejected by h will make an offer to the next compatible agent on her (random) list. $\mathcal{Q}'$ will continue evolving according to the greedy process with the exception that agent h is essentially ignored (and does not affect the evolution of the process).

[39] We note that if a is an H agent herself, then there will be no such compatible H agent, by definition.

Event (ii) The criticality clock of h ticks in $\mathcal{Q}$ before time $t_0 + \log m$ and h departs without being matched. In this case, we stop $\mathcal{Q}$ but continue to run $\mathcal{Q}'$. $\mathcal{Q}'$ continues to evolve according to the greedy process with the exception that h will ignore all offers she may receive in the future (i.e., h is essentially ignored and does not affect the evolution of $\mathcal{Q}'$).

For notational simplicity, let $t_0 = 0$ without loss of generality. Let $E_h(t)$ denote the event in which h departs in $\mathcal{Q}$ before time t (either matched or unmatched). Also, let $E'_h(t)$ denote the event in which h receives at least one offer in $\mathcal{Q}'$ before time t. For any constant $t > 0$,

$$\mathbb{P}\left[\overline{E_h(t)}\right] = e^{-t/d} \cdot \mathbb{P}\left[\overline{E'_h(t)}\right] \tag{45}$$

holds for sufficiently large m such that $\log m > t$. The factor $e^{-t/d}$ on the right-hand side accounts for the fact that h becomes critical in $\mathcal{Q}'$ only at time $\log m$, but she becomes critical in Q with rate $1/d$. Therefore, if we show that $\lim_{m\to\infty} \mathbb{P}\left[\overline{E'_h(t)}\right] = e^{-t/\lambda}$, then (45) implies that

$$\lim_{m\to\infty} \mathbb{P}\left[\overline{E_h(t)}\right] = e^{-t(1/d+1/\lambda)},$$

which proves the claim on the distribution of waiting time. Hence, it remains to prove the following.

Claim 10. *For any constant $t > 0$,*

$$\lim_{m\to\infty} \mathbb{P}\left[\overline{E'_h(t)}\right] = e^{-t/\lambda}.$$

Proof. First, we observe that the process $\mathcal{Q}'$ can be run from time $t_0 = 0$ to $\log m$ as follows: sample a state $(x, y) \sim \pi$. Then, let the stochastic process corresponding to the greedy policy start from the state (x, y) and run for $\log m$ units of time. We also consider the extra "dummy" agent h who is in the H pool but, as discussed earlier, rejects every offer made to her. We will show that the probability that h receives no offers in the period $[0, t]$, i.e., $\mathbb{P}\left[\overline{E'_h(t)}\right]$, approaches $e^{-t/\lambda}$ as m approaches infinity.

Let $\mathcal{F}_t$ be a filtration that contains the information about the number of E agents arriving until time t in $\mathcal{Q}'$, denoted by n_t, the arrival times of these agents, denoted by by $a_1, \ldots, a_{n_t}$, and the number of H agents in the pool at times $a_1, \ldots, a_{n_t}$, denoted by $x_1, \ldots, x_{n_t}$. Then, observe that

$$\mathbb{P}\left[\overline{E'_h(t)} \middle| \mathcal{F}_t\right] \geq \Pi_{i=1}^{n_t}\left(1 - \frac{1}{x_i}\right), \tag{46}$$

$$\mathbb{P}\left[\overline{E'_h(t)} \middle| \mathcal{F}_t\right] \leq \Pi_{i=1}^{n_t}\left(1 - \frac{1-(1-p)^{x_i}}{x_i}\right). \tag{47}$$

To show that (46) holds one can consider two events: (i) either the E agent arriving at time a_i is compatible with at least one existing H agent or (ii) she is not compatible with any existing H

agent. In case (ii) h will be matched with the agent arriving at time a_i with probably 0. In case (i) h will be matched with the arriving agent with probability $\frac{1}{x_i}$, as there are x_i waiting hard-to-match agents and each of them is matched with the same probability. Thus, an lower bound on the probability that h is not matched with the E agent arriving at time a_i is given by $1 - \frac{1}{x_i}$. This implies the lower bound we state in (46) on the probability that h is not matched before time t. Also, (47) holds because the chance that the E agent arriving at time a_i is compatible with at least one of the H agents is $1-(1-p)^{x_i}$, and conditional on this event, the chance that the E agent makes an offer to agent h is at least $\frac{1-(1-p)^{x_i}}{x_i}$.

Define $\underline{x} = \min\{x_1, \ldots, x_{n_t}\}$. In the rest of the proof, we will use probabilistic bounds to show that in almost all sample paths of $\mathcal{Q}'$, $\underline{x}$ is close to λm and n_t is close to tm. This will let us simplify (46) and (47) and prove the claim.

Define event G as

$$G: \ \lambda m - \sqrt{m}\log^2 m \leq \underline{x} \leq \lambda m + \sqrt{m}\log^2 m.$$

Next, we will show that G happens with high probability. To this end, let the event $\overline{G}$ be the complement of the event G. We will show that

$$\mathbb{P}\left[\,\overline{G}\,\right] \leq e^{-\Omega(\log^2 m)}. \tag{48}$$

First, we observe that n_t is a Poisson random variable with mean mt. Thus, by the concentration bound of Fact 3 for the Poisson distribution, for any constant $t > 0$ we have

$$\mathbb{P}\left[n_t \notin [mt - (mt)^{2/3}, mt + (mt)^{2/3}]\right] = e^{-\Omega(\sqrt{m})}. \tag{49}$$

Next, consider an agent who arrives at time a_i, for a positive integer $i \leq n_t$. Let x_i, y_i denote the sizes of the H and E pools, respectively. By the PASTA property, the (unconditional) distribution of (x_i, y_i) is identical to the steady-state distribution π. Therefore, by Theorem 3 and Theorem 4,

$$\mathbb{P}\left[x_i \notin (\lambda m - \sqrt{m}\log^2 m, \lambda m + \sqrt{m}\log^2 m)\right] = e^{-\Omega(\log^2 m)}. \tag{50}$$

Let X_i denote the event $x_i \notin (\lambda m - \sqrt{m}\log^2 m, \lambda m + \sqrt{m}\log^2 m)$. A union bound over the events $\{X_i\}_{i=1\ldots,n_t}$ implies that

$$\mathbb{P}\left[\,\overline{G}\,\right] \leq \mathbb{P}\left[\cup_{i=1}^{n_t} X_i\right]. \tag{51}$$

The fact that $t > 0$ is a constant independent of m, together with (49), (50) and (51), implies that

$$\mathbb{P}\left[\,\overline{G}\,\right] \leq e^{-\Omega(\sqrt{m})} + e^{-\Omega(\log^2 m)} \cdot (mt + (mt)^{2/3}).$$

The above bound means that $\mathbb{P}\left[\overline{G}\right] \leq e^{-\Omega(\log^2 m)}$. Hence, (48) holds.

To summarize, so far we have shown that, with high probability, n_t is close to mt and $\underline{x}$ is close to λm; this was done by establishing (48) and (49). We will use these concentration bounds together with (46) and (47) to conclude the proof.

First, we recall (46):

$$\mathbb{P}\left[\overline{E'_h(t)}\middle|\mathcal{F}_t\right] \geq \Pi_{i=1}^{n_t}(1 - \frac{1}{x_i}).$$

We now use (49) to bound n_t from above, and then (48) to bound $\underline{x} = \min\{x_1, \ldots, x_{n_t}\}$ from below. This, together with the above inequality, implies that

$$\begin{aligned}
\mathbb{P}\left[E'_h(t)\right] &\leq e^{-O(\log^2 m)} + e^{-O(\sqrt{m})} + 1 - (1 - \frac{1}{\lambda m - \sqrt{m}\log^2 m})^{mt+(mt)^{2/3}} \\
&\leq e^{-O(\log^2 m)} + e^{-O(\sqrt{m})} + 1 - \left(\exp\left(-\frac{1}{\lambda m - \sqrt{m}\log^2 m} \cdot \frac{1}{1 - \frac{1}{\lambda m - \sqrt{m}\log^2 m}}\right)\right)^{mt+(mt)^{2/3}} \qquad (52) \\
&= e^{-O(\log^2 m)} + e^{-O(\sqrt{m})} + 1 - \exp\left(-\frac{mt + (mt)^{2/3}}{\left(\lambda m - \sqrt{m}\log^2 m\right)\left(1 - \frac{1}{\lambda m - \sqrt{m}\log^2 m}\right)}\right)
\end{aligned}$$

where (52) uses the fact that $1 - z \geq e^{\frac{z}{1-z}}$ holds for all $z \in (0,1)$. The above inequality implies that $\lim_{m\to\infty} \mathbb{P}\left[E'_h(t)\right] \leq 1 - e^{-\frac{mt}{\lambda m}} = 1 - e^{-\frac{t}{\lambda}}$, which implies that

$$\lim_{m\to\infty} \mathbb{P}\left[\overline{E'_h(t)}\right] \geq e^{-\frac{t}{\lambda}}. \qquad (53)$$

To complete the proof of the lemma, it remains to show that $\lim_{m\to\infty} \mathbb{P}\left[\overline{E'_h(t)}\right] \leq e^{-\frac{t}{\lambda}}$. To this end, we first recall (47):

$$\mathbb{P}\left[\overline{E'_h(t)}\middle|\mathcal{F}_t\right] \leq \Pi_{i=1}^{n_t}(1 - \frac{1-(1-p)^{x_i}}{x_i}).$$

The above bound, together with the concentration bound (48) for $\min\{x_1, \ldots, x_{n_t}\}$ and the concentration bound (49) for n_t, implies that

$$\begin{aligned}
\mathbb{P}\left[E'_h(t)\right] &\geq e^{-O(\log^2 m)} + e^{-O(\sqrt{m})} + 1 - \left[\exp\left(\frac{1-(1-p)^{\lambda m + \sqrt{m}\log^2 m}}{\lambda m + \sqrt{m}\log^2 m}\right)\right]^{mt-(mt)^{2/3}} \\
&= e^{-O(\log^2 m)} + e^{-O(\sqrt{m})} + 1 - \exp\left(\frac{\left(1-(1-p)^{\lambda m + \sqrt{m}\log^2 m}\right)\left(mt - (mt)^{2/3}\right)}{\lambda m + \sqrt{m}\log^2 m}\right)
\end{aligned}$$

The above inequality implies that $\lim_{m\to\infty} \mathbb{P}\left[E'_h(t)\right] \geq 1 - e^{-\frac{mt}{\lambda m}} = 1 - e^{-\frac{t}{\lambda}}$, which implies that $\lim_{m\to\infty} \mathbb{P}\left[\overline{E'_h(t)}\right] \leq e^{-\frac{t}{\lambda}}$. This inequality together with (53) conclude the proof of the claim. □

ii Non-asymptotic Comparison of the Greedy and Batching Policies

In this section, we compare the batching and greedy policies in small markets. To this end, we consider a market with a fixed arrival rate m, and a batching policy with batch length T in this market. We then quantify an upper bound $\overline{T^*}$ such that the match rate under the batching policy is strictly smaller than under the greedy policy in this market if $T > \overline{T^*}$. We note that, as $\overline{T^*}$ is only an upper bound, $T < \overline{T^*}$ would not necessarily imply that the match rate is higher under the batching policy than under the greedy policy.

Recall from Definition 7 the definition of the Markov chain $\mathcal{M}$ (associated with the greedy policy) and its unique steady state distribution π; also, recall that π_x denotes $\sum_{y=0}^{\infty} \pi_{x,y}$ for every integer $x \geq 0$. Define $\Pi_x = \sum_{i=0}^{x} \pi_x$.

Proof of Proposition 8. We first prove the claim of the proposition for match rate. Observe that, under the greedy policy, the match rates of E and H agents are respectively at least z^* and $\frac{z^*}{1+\lambda}$.

Let $\overline{B}(T) = \frac{1-e^{-T/d}}{T/d}$. Recall that, by Lemma 20, this is an upper bound on the match rate of E agents under the batching policy.

Claim 11. *The function $\overline{B}(T)$ is strictly decreasing.*

Proof. We observe that

$$\frac{\partial \overline{B}(T)}{\partial T} = \frac{e^{-\frac{T}{d}}\left(d\left(1-e^{T/d}\right)+T\right)}{T^2} < \frac{e^{-\frac{T}{d}}\left(d\left(1-1-T/d\right)+T\right)}{T^2} = 0,$$

where the inequality follows from the fact that $e^{\alpha} > 1 + \alpha$ for every $\alpha > 0$. □

For every z, solving the equation $\frac{1-e^{-T/d}}{T/d} = z$ for T gives

$$T^* = \frac{z^* W\left(-\frac{e^{-1/z^*}}{z^*}\right) + 1}{z^*/d}.$$

The above equation, together with the fact that $\overline{B}(\cdot)$ is strictly decreasing (as shown by Claim 11), implies that the match rate of E agents under the batching policy is smaller than z^* if the batch length is larger than T^*. Recall that z^* is a lower-bound on the match rate of E agents under the greedy policy. Thus, the match rate of E agents under the batching policy is smaller than their match rate under the greedy policy if the batch length is larger than T^*.

To prove the claim for the match rate of H agents, we note that their match rate under the greedy policy is at least $\frac{z^*}{1+\lambda}$. Also, their match rate under the batching policy is at most $\frac{\overline{B}(T)}{1+\lambda}$, by Lemma 20. The two latter bounds are the same as their counterpart bounds for E agents divided by $1+\lambda$. Hence, the proof for H agents follows directly from the above proof for E agents.

It remains to prove the claim for waiting times. Lemma 18 shows that $w^G_\Theta(m) = d(1 - q^G_\Theta(m))$ and $w^B_\Theta(m) = d(1 - q^B_\Theta(m))$. Thus, if the batch length in the batching policy is larger than T^*, then $w^G_\Theta(m) < w^B_\Theta(m)$. □

Proposition 11. *Suppose there exist integers $j_1, j_2 \geq 0$ and positive reals $\overline{\Pi}_{j_1}, \overline{\Pi}_{j_2}$ such that $j_1 \geq j_2$, $\Pi_{j_1} \leq \overline{\Pi}_{j_1}$ and $\Pi_{j_2} \leq \overline{\Pi}_{j_2}$. Define*

$$\underline{z}^* = (1 - \overline{\Pi}_{j_1})\left(1 - (1-p)^{j_1}\right) + (\overline{\Pi}_{j_1} - \overline{\Pi}_{j_2})\left(1 - (1-p)^{j_2}\right).$$

Then, for every agent type (E or H), the match rate and waiting time of that type under the batching policy are respectively smaller and larger than under the greedy policy if $T > \overline{T^*}$, where

$$\overline{T^*} = \frac{\underline{z}^* W\left(-\frac{e^{-1/\underline{z}^*}}{\underline{z}^*}\right) + 1}{\underline{z}^*/d}$$

and $W(\cdot)$ is the Lambert W function.

Proof. We first prove the following claims.

Claim 12. *Recall from Proposition 8 that z^* is the probability that an E agent, upon her arrival, is matched to an H agent under the greedy policy. Then, $z^* \geq \underline{z}^*$.*

Proof. For every integer $j_1 \geq 0$,

$$\mathbb{P}_{(x,y)\sim\pi}\left[x > j_1\right] = 1 - \Pi_{j_1} \geq 1 - \overline{\Pi}_{j_1}.$$

Therefore, upon the arrival of an easy-to-match agent, there are at least j_1 hard-to-match agents in the pool with probability at least $1 - \overline{\Pi}_{j_1}$. Hence, the easy-to-match agent, upon her arrival, is compatible to at least one hard-to-match agent with probability at least $(1-\overline{\Pi}_{j_1})(1-(1-p)^{j_1})$. This also holds when j_1 is replaced with j_2. The two latter facts, together with $j_1 \geq j_2$ and $\overline{\Pi}_{j_1} \geq \overline{\Pi}_{j_2}$ imply that an easy-to-match agent, upon her arrival, is compatible to at least one hard-to-match agent with probability at least

$$(1 - \overline{\Pi}_{j_1})\left(1 - (1-p)^{j_1}\right) + (\overline{\Pi}_{j_1} - \overline{\Pi}_{j_2})\left(1 - (1-p)^{j_2}\right).$$

This proves the claim. □

Claim 13. *The function $h(z) = \frac{zW\left(-\frac{e^{-1/z}}{z}\right)+1}{z/d}$ is decreasing at z for every $z \in [0,1]$.*

Proof. To prove the claim it suffices to show that the function $W\left(\frac{-e^{-1/z}}{z}\right) + \frac{1}{z}$ is decreasing in z, and to prove this it suffices to show that $W\left(\frac{-e^{-1/z}}{z}\right)$ is decreasing in z. Define $\phi(z) = \frac{-e^{-1/z}}{z}$. We note that $\phi(z) \in [e^{-1}, 0)$ for $z \in (0, 1]$, and that W is well-defined and increasing over the interval $[e^{-1}, 0]$ (Corless et al., 1996). The fact that $\phi'(z) = \frac{e^{-1/z}(z-1)}{z^3} < 0$ for $z \in (0, 1]$ thus implies that $W(\phi(z))$ is decreasing in z for $z \in (0, 1]$. □

□

By Proposition 11, the greedy policy has a larger match rate and lower waiting time than the batching policy if $T > \overline{T^*}$. We emphasize that this condition is not necessary, i.e., the greedy policy may have a better performance that the batching policy even when $T < \overline{T^*}$ (and in fact this is what we observe in the simulations of Figure 5).

We next demonstrate how small $\overline{T^*}$ can be in practically relevant instances. To this end, we suppose that $\lambda = 1.33$, $p = 0.037$, $m \in [250, 450]$, and $d = 1$. For better exposition, we call each unit of time a *year*, and assume that a year is divided to 365 *days*. (See the calibration to the NKR data in Section 4.1.) Changing the unit of time to year allows us to assume $d = 1$, which simplifies notation.

In Online Appendix ii.1, we compute $\overline{T^*}$ for the above parametrization of our model using Proposition 11. The result is plotted in Figure 11. This is the same figure as Figure 4 except that the horizontal axis has a different scale: as noted earlier, here we have assumed that the unit of time is year rather than day. The step by step derivation of this plot is described in Section ii.1. We observe that, at $m = 250$, if the batching policy makes matches less frequently than every day, then it would have a lower match rate than the greedy policy. As mentioned earlier, our theoretical bound does not imply the other direction—it does not imply that making matches faster than every $\overline{T^*}$ units of time in the batching policy leads to a higher match rate than the greedy policy.

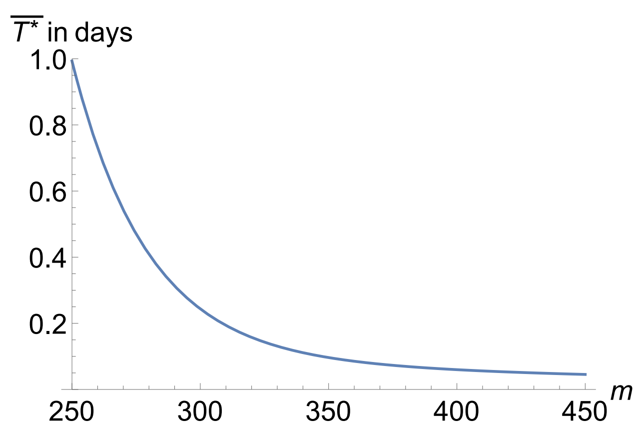

Figure 11: The vertical and horizontal axes respectively corresponds to m (arrival rate per year) and $\overline{T^*}$ (period length in days).

ii.1 The Step by Step Derivation of Figure 4

We will use Proposition 11 to derive a closed-form expression for $\overline{T^*}$ for the specified range of parameters, namely $\lambda = 1.33$, $p = 0.037$, $m \in [250, 450]$, and $d = 1$. To apply Proposition 11, we need to specify integers $j_1, j_2 \geq 0$ and positive reals $\overline{\Pi}_{j_1}, \overline{\Pi}_{j_2}$ such that $j_1 \geq j_2$ and

$$\Pi_{j_1} \leq \overline{\Pi}_{j_1}, \tag{54}$$

$$\Pi_{j_2} \leq \overline{\Pi}_{j_2}. \tag{55}$$

To this end, let $k_1 = 1.65 \log m$ and $k_2 = 2.46 \log m$. Define $j_1 = \lambda m - k_1\sqrt{m}$ and $j_2 = \lambda m - k_2\sqrt{m}$. Also, let $\overline{R} = 0.46$, and define

$$\overline{\Pi}_{j_1} = \frac{\overline{R}}{Q(k_1)} \cdot \frac{2(1 + \lambda + k_1/\sqrt{m})}{2 + \lambda}, \tag{56}$$

$$\overline{\Pi}_{j_2} = \frac{\overline{R}}{Q(k_2)} \cdot \frac{2(1 + \lambda + k_2/\sqrt{m})}{2 + \lambda}, \tag{57}$$

where we recall the definition of the function Q from Theorem 2:

$$Q(k) = e^k \left(\frac{1}{\sqrt{m}} \frac{k\sqrt{m} - 2m(1+\lambda)N_{k\sqrt{m}}}{2m(1+\lambda) + k\sqrt{m}} - \frac{1}{m} \right).$$

We will show that this choice of parameters j_1, j_2 and $\overline{\Pi}_{j_1}, \overline{\Pi}_{j_2}$ satisfies the conditions (54) and (55). This would imply that Proposition 11 applies, which then gives

$$\overline{T^*} = \frac{\underline{z}^* W\left(-\frac{e^{-1/\underline{z}^*}}{\underline{z}^*}\right) + 1}{\underline{z}^*/d},$$

where

$$\underline{z}^* = (1 - \overline{\Pi}_{j_1})\left(1 - (1-p)^{j_1}\right) + (\overline{\Pi}_{j_1} - \overline{\Pi}_{j_2})\left(1 - (1-p)^{j_2}\right).$$

Figure 4 plots the above expression for $\overline{T^*}$.

To complete the analysis, it remains to show that the choice of parameters j_1, j_2 and $\overline{\Pi}_{j_1}, \overline{\Pi}_{j_2}$ satisfies the conditions (54) and (55). This is done next.

ii.1.1 Conditions (54) and (55) are satisfied

Recall from Definition 7 that the Markov chain $\mathcal{M}$ models the stochastic process governing the greedy policy, and π denotes its steady state distribution. Also, recall from Section D.2.1 the simplified greedy process, its associated Markov chain $\mathcal{N}$, and its steady state distribution ρ. We will show that (54) and (55) hold using the following corollary of Theorem 2.

Corollary 4 (Corollary of Theorem 2)**.** *For every $k > 0$ define*

$$Q(k) = e^k \left(\frac{1}{\sqrt{m}} \frac{k\sqrt{m} - 2m(1+\lambda)(1-p)^{k\sqrt{m}}}{2m(1+\lambda) + k\sqrt{m}} - \frac{1}{m} \right),$$

$$R(k) = \sup_{0 \le z \le k\sqrt{m}} H(z),$$

$$\Phi(k) = \frac{R(k)}{Q(k)} \cdot \frac{2(1+\lambda+k/\sqrt{m})}{2+\lambda}.$$

Then, $\Pi_{\lambda m - k} \le \Phi(k)$ holds if $Q(k) > 0$.

Proof. If $x \le x - k\sqrt{m}$, then $(x, y) \notin B_k$ holds for every $y \ge 0$. Therefore,

$$\Pi_{\lambda m - k\sqrt{m}} = \mathbb{P}_{(x,y)\sim\pi}\left[x \le x - k\sqrt{m}\right] \le \mathbb{P}_{(x,y)\sim\rho}\left[(x,y) \notin B_k\right] \le \Phi(k),$$

where the first inequality holds by Lemma 4 and the second inequality holds by Theorem 2. This concludes the proof. □

Fact 7. *For every $m \in [250, 450]$, $Q(k_1), Q(k_2) > 0$.*

Proof. We demonstrate this by plotting $Q(k_1), Q(k_2)$ for $m \in [250, 450]$.

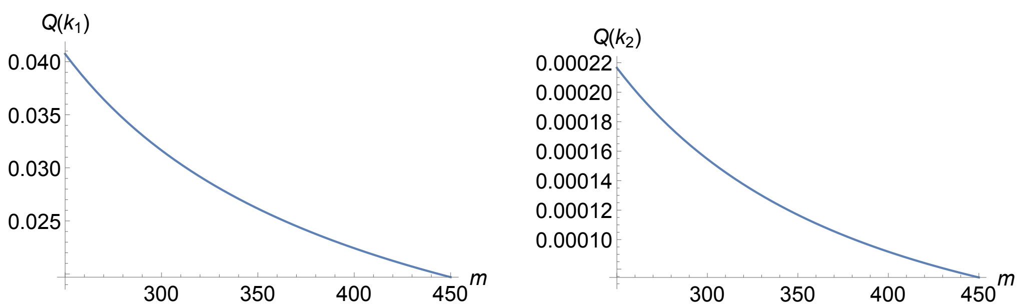

Figure 12: Recall that we set $k_1 = 1.65 \log m$ and $k_2 = 2.46 \log m$. We plot $Q(k_1), Q(k_2)$ using the expression given for the function Q in Corollary 4, while varying m.

□

Fact 8. *For every $m \in [250, 450]$, $R(k_1), R(k_2) < 0.46$.*

Proof. We first recall that, by Definition 12,

$$H(z) = -e^{\frac{z}{\sqrt{m}}} \left(\frac{1}{\sqrt{m}} \frac{z - 2m(1+\lambda)(1-p)^z}{2m(1+\lambda) + z} - \frac{1}{m} \right).$$

To prove the claim, it suffices to show that, for every $m \in [250, 450]$ and $z \in [0, k_2\sqrt{m}]$, $H(z) < 0.46$. We demonstrate this by plotting $H(z)$ for $m \in [250, 450]$ and $z \in [0, k_2\sqrt{450}]$. Figure 13 demonstrates that $H(z) < 0.46$ holds for $m \in [250, 450]$ and $z \in [0, k_2\sqrt{450}]$.

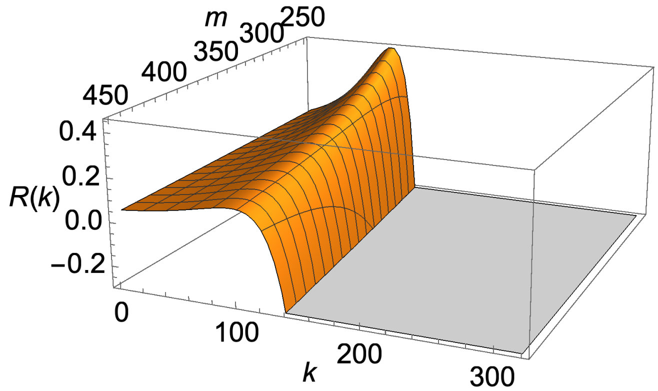

Figure 13: The range of the vertical axis is $[-0.3, 0.46]$. The plot is clipped at the gray area, as the value of the function R in the gray area falls below -0.3.

□

Let $\overline{R} = 0.46$. Corollary 4, together with Fact 7 and Fact 8, implies that

$$\Pi_{\lambda m - k_1\sqrt{m}} \leq \frac{\overline{R}}{Q(k_1)} \cdot \frac{2(1+\lambda+k_1/\sqrt{m})}{2+\lambda}, \tag{58}$$

$$\Pi_{\lambda m - k_2\sqrt{m}} \leq \frac{\overline{R}}{Q(k_2)} \cdot \frac{2(1+\lambda+k_2/\sqrt{m})}{2+\lambda}. \tag{59}$$

Recall that, by definition, $j_1 = \lambda m - k_1\sqrt{m}$, $j_2 = \lambda m - k_2\sqrt{m}$, and

$$\overline{\Pi}_{j_1} = \frac{\overline{R}}{Q(k_1)} \cdot \frac{2(1+\lambda+k_1/\sqrt{m})}{2+\lambda}, \qquad \overline{\Pi}_{j_2} = \frac{\overline{R}}{Q(k_2)} \cdot \frac{2(1+\lambda+k_2/\sqrt{m})}{2+\lambda}.$$

Hence, (58) and (59) are the same as (54) and (55), respectively. Therefore, (54) and (55) hold.

iii Proof of Proposition 6

The proof for the 'only if' part is relatively straightforward. Suppose $\lim_{m\to\infty} T_m \neq 0$. Then, there exists $\delta > 0$ and an increasing sequence of positive integers $m_1, m_2, \ldots$ such that $T_{m_i} > \delta$ for every $i > 0$. By Lemma 19, conditional on an agent arriving in an interval of length T_i, namely $(s, s+T_i)$, the agent does not become critical before time $s+T_i$ with probability $\frac{1-e^{-T_i/d}}{T_i/d}$. Thus, since $T_i > \delta$, the match rate of every agent type under the batching policy is at most $\frac{1-e^{-\delta/d}}{\delta/d}$ when the market size is m_i. On the other hand, the match rate of E agents under the greedy policy converges to 1 as the market grows large (Proposition 3). Therefore, the batching policy is not asymptotically optimal.

To prove the 'if' part of Proposition 6, we first prove the following theorems.

Theorem 5. *There exists a constant* $c_\lambda > 0$ *(independent of* m*) such that, at any time* $t \geq 0$*, the number of* H *agents in the pool is at least* $c_\lambda m$*, whp.*

Theorem 6. *If* $\lim_{m\to\infty} T_m = 0$*, then the match rate of* E *agents and the match rate of* H *agents respectively approach* 1 *and* $\frac{1}{1+\lambda}$ *as* m *approaches infinity.*

Theorem 7. *If* $\lim_{m\to\infty} T_m = 0$*, then the expected waiting time of* E *agents and the expected waiting time of* H *agents respectively approach* 0 *and* $\frac{\lambda d}{1+\lambda}$ *as* m *approaches infinity.*

iii.1 Preliminary Definitions and Lemmas

In the analysis, we suppose that time is indexed by non-negative real numbers. Consider a market of size m. A batching policy makes matches at times iT_m, for every positive integer i; these times are called *matching times*. For every $i \geq 0$ the interval $(iT_m, (i+1)T_m]$ is called *period* i.

The batching policy *executes* a matching at the *end* of every period i: at time $(i+1)T_m$, it finds the largest matching in the pool. If there are several such matchings, it selects the matching among them which has the maximum number of H agents. The policy then executes the selected matching (by performing the corresponding transplants) and removes the agents involved in that matching from the pool.

For any matching time t, A_t and B_t respectively denote the number of E and H agents in the pool after the execution of the matching at time t. For any time t that is not a matching time, we let A_t and B_t denote the number of E and H agents in the pool at time t, respectively.

We note that the sequence $\langle (A_{iT_m}, B_{iT_m}) \rangle_{i\geq 0}$ is a discrete time Markov chain with a state space $\mathbb{Z}_+ \times \mathbb{Z}_+$. Since this Markov chain is Ergodic, it has a steady-state distribution, which we denote by ρ. Since we are performing a steady-state analysis, we suppose that (A_0, B_0) is drawn from ρ. This assumption is without loss of generality by the Ergodic theorem for Markov chains.

iii.2 Proof of Theorem 5

Throughout the proof of this theorem, we scale (speed up) time with a factor of d, and suppose that the criticality rate d is normalized to 1. This is without loss of generality because scaling time only affects the speed, and not the distribution of the number of agents in the pool. At the same time, as d is a constant independent of m, the condition $\lim_{m\to\infty} T_m = 0$ would still hold after scaling time with a factor of d.

Lemma 24. *Let* $\delta > 0$ *be a constant and* $t \geq 0$ *be an arbitrary time. Then, whp,* $A_t < (1+\delta)m$*.*

Proof. Let $\mathcal{B}$ denote the stochastic process governing the batching policy. Consider a stochastic process $\mathcal{B}'$ which is the same as $\mathcal{B}$ with the difference that $\mathcal{B}'$ never executes a matching. Thus, $\mathcal{B}'$ has the same sequence of agents' arrivals as $\mathcal{B}$, and every agent in $\mathcal{B}'$ has the same criticality time as in $\mathcal{B}$. However, agents are not matched to each other in $\mathcal{B}'$. Let A'_t denote the number of

E agents at time t in process $\mathcal{B}'$. As A'_t is a pathwise upper bound on A_t, to prove the lemma, it suffices to show that $A'_t < (1+\delta)m$ holds whp.

The number of E agents in $\mathcal{B}'$ is a continuous-time Markov chain $\mathcal{M}'$. Since $\mathcal{M}'$ is Ergodic, it has a steady-state distribution, which we denote by ρ'. Also, let ρ'_i denote the steady-state probability assigned by ρ' to the state at which the number of E agents is i. The Markov chain $\mathcal{M}'$ may transit from each state i to a neighboring state $i-1$ (in case $i \geq 1$) or $i+1$. Let l_i and r_i respectively denote the transition rate from state i to states $i-1$ and $i+1$. Observe that $r_i = m$ holds for all $i \geq 0$, because the arrival rate of E agents is m. Also, $l_i = i$ holds for all $i \geq 0$ because $d = 1$. The balance equations for ρ' then imply that

$$\frac{\rho'_{i+1}}{\rho'_i} = \frac{r_i}{l_{i+1}} = \frac{m}{i+1}.$$

By the above equation, for $i \geq m + 2\sqrt{m}$ we have

$$\frac{\rho'_{i+1}}{\rho'_i} \leq \frac{m}{m+2\sqrt{m}} = 1 - \frac{2}{\sqrt{m}+2} \leq 1 - \frac{1}{\sqrt{m}} \leq e^{-\frac{1}{\sqrt{m}}},$$

where the second inequality holds for $m \geq 4$. Thus, for every integer $j \geq 0$ and $i = m + 2\sqrt{m} + j$,

$$\rho'_i \leq \frac{\rho'_i}{\rho'_m} \leq e^{-\frac{j}{\sqrt{m}}}.$$

By the above bound, for every integer $i \geq 0$ we have

$$\sum_{j=i}^{\infty} \rho'_{m+2\sqrt{m}+j} \leq \sum_{j=i}^{\infty} e^{-\frac{j}{\sqrt{m}}} = \frac{(e^{-1/\sqrt{m}})^i}{1 - e^{-1/\sqrt{m}}}.$$

For $i = \log m\sqrt{m}$, the right-hand side of the above equation goes to 0 as m goes to infinity. Therefore, $\mathbb{P}_{\rho'}\left[A'_t \geq m + (2 + \log m)\sqrt{m}\right]$ goes 0 as m goes to infinity. This proves the claim. □

Lemma 25. *Let $\epsilon > 0$ be a constant. Also, let $t = iT_m$ and $t' = jT_m$ be such that $0 \leq i < j$ and $t' - t = \epsilon$. Then, if $A_t \leq 2\epsilon(1+\lambda)m$, whp it holds that $A_{t'} \leq m\epsilon(4 + 2\lambda)$.*

Proof. The number of E agents who arrive to the pool in the period $[t, t']$ obeys the Poisson distribution with mean ϵm. This fact, together with the concentration bound of Fact 3 for the Poisson distribution, implies that the number of E agents who arrive to the pool in the period $[t, t']$ is at most $2\epsilon m$, whp. Therefore, $A'_t \leq A_t + 2\epsilon m$ holds whp. This concludes the proof. □

Lemma 26. *Let $\epsilon \in (0, \frac{\lambda/3}{3+2\lambda})$ be a constant. Also, let $t = iT_m$ and $t' = jT_m$ be such that $0 \leq i < j$ and $t' - t = \epsilon$. Then, if $A_t > 2\epsilon(1+\lambda)m$, whp it holds that $A_{t'} \leq A_t - \epsilon\lambda m/2$.*

Proof. For notational simplicity, let $\alpha = \epsilon\lambda/2$. Let X denote the set of all H agents who arrive at a point in $[t, t']$ and whose criticality time is larger than t'. Let Y_1 denote the set of all E agents who arrive at a point in $[t, t']$ and whose criticality time is larger than t'. Also, let Y_2 denote the set of all E agents who are present in the pool after the execution of the matching at time t whose criticality time is larger than t'. Let $Y = Y_1 \cup Y_2$.

Claim 14. *For every constant $\kappa \in (0, \epsilon]$, it holds whp that*

$$(e^{-\epsilon} A_t + (1 - e^{-\epsilon})m)(1 - \kappa) \leq |Y| \leq A_t + \epsilon(1 + \kappa)m$$

and

$$(1 - e^{-\epsilon})(1 - \kappa)(1 + \lambda)m \leq |X| \leq (1 - e^{-\epsilon})(1 + \kappa)(1 + \lambda)m.$$

Proof. The number of E agents that arrive in the interval $[t, t']$ and have a criticality time larger than t' is distributed according to the Poisson distribution with mean $(1 - e^{-\epsilon})m$, by Lemma 19. Therefore, $\mathbb{E}[|Y_1|] = (1 - e^{-\epsilon})m$. This, together with the concentration bound of Fact 3 for the Poisson distribution, implies that $|Y_1| \geq (1 - \kappa)(1 - e^{-\epsilon})m$ holds whp. On the other hand, $\mathbb{E}[|Y_2|] = e^{-\epsilon} A_t$, which holds because every agent that is in the pool after the execution of the matching at time t has a criticality time larger than t' with probability $e^{-\epsilon}$. Because the criticality clocks are independent, the Chernoff concentration bound (Fact 4) applies, and implies that $|Y_2| \geq (1 - \kappa)e^{-\epsilon} A_t$ holds whp. The lower concentration bounds that we developed on $|Y_1|, |Y_2|$ imply that $|Y| \geq (1 - \kappa)(m(1 - e^{-\epsilon}) + e^{-\epsilon} A_t)$ holds whp, which is the claimed lower bound for $|Y|$. To prove the upper bound on $|Y|$, we first observe that

$$|Y| \leq A_t + (1 + \kappa)(1 - e^{-\epsilon})m,$$

which holds by the same argument as above (but with the lower concentration bounds replaced with upper concentration bounds, i.e., the factor $1 - \kappa$ replaced with $1 + \kappa$). The latter inequality, together with the fact that $1 + z \leq e^z$ for $z \in \mathbb{R}$, implies that $|Y| \leq A_t + (1 + \kappa)\epsilon m$ holds whp.

To prove the bounds on $|X|$, first we observe that $\mathbb{E}[|X|] = (1 - e^{-\epsilon})(1 + \lambda)m$, by Lemma 19. Since $|X|$ has a Poisson distribution (by the independence of the criticality clocks of the agents), the concentration bound of Fact 3 for the Poisson distribution applies again, and immediately implies the claimed bounds for $|X|$. □

Claim 15. *$|Y| > |X|$ holds whp.*

Proof. By Claim 14, for every constant $\kappa \in (0, \epsilon]$, it holds whp that

$$|X| \leq (1 - e^{-\epsilon})(1 + \kappa)(1 + \lambda)m \leq \epsilon(1 + \kappa)(1 + \lambda)m.$$

Claim 14 also implies that, whp,

$$
\begin{aligned}
|Y| &\geq (e^{-\epsilon} A_t + (1 - e^{-\epsilon})m)(1 - \kappa) \\
&\geq \left((1 - \epsilon)A_t + (\epsilon - \epsilon^2)m\right)(1 - \kappa) \\
&\geq \left((1 - \epsilon)(2\epsilon(1 + \lambda)m) + (\epsilon - \epsilon^2)m\right)(1 - \kappa) \\
&\geq \left(2\epsilon(1 + \lambda)m - 3\epsilon^2(1 + \lambda)m\right)(1 - \kappa) \\
&\geq \frac{3}{2}\epsilon(1 + \lambda)m(1 - \kappa),
\end{aligned}
$$

where the second inequality holds because $1 - \epsilon \leq e^{-\epsilon} \leq 1 - \epsilon + \epsilon^2$, the third inequality uses the fact that $A_t > 2\epsilon(1 + \lambda)m$, and last inequality holds because $\epsilon < \frac{\lambda/3}{3+2\lambda} < \frac{1}{6}$.

Thus far, we have shown that $|X| \leq \epsilon(1+\kappa)(1+\lambda)m$ and $|Y| \geq \frac{3}{2}\epsilon(1+\lambda)m(1-\kappa)$ hold whp. To complete the proof of the claim, it suffices to show that the right-hand side of the former inequality is smaller than the right-hand side of the latter inequality, i.e., it suffices to show that

$$
\epsilon(1 + \kappa)(1 + \lambda)m < \frac{3}{2}\epsilon(1 + \lambda)m(1 - \kappa).
$$

The above inequality holds if and only if $\frac{1+\kappa}{1-\kappa} < \frac{3}{2}$. Recall that $\kappa \in (0, \epsilon)$; since $\epsilon < \frac{\lambda/3}{3+2\lambda} < \frac{1}{6}$, therefore $\frac{1+\kappa}{1-\kappa} < 1.4$. This completes the proof. □

Define the bipartite graph $G(X, Y)$ by connecting a node in X and a node in Y iff the agents corresponding to the nodes are compatible.

By Claim 15, $|Y| > |X|$ holds whp. If $A_{t'} > |Y| - |X| + \alpha m$, then at most $|X| - \alpha m$ hard-to-match agents are matched to easy-to-match agents when the matching is executed at time t'; this implies that $B_{t'} > |X| - (|X| - \alpha m) = \alpha m$. Therefore, if $A_{t'} > |Y| - |X| + \alpha m$, then there exists an $(\alpha m, |Y| - |X| + \alpha m)$-independent set in $G(X, Y)$. Such an independent set, however, does not exist whp, by Lemma 17 and Claim 14. Hence, whp, $A_{t'} > |Y| - |X| + \alpha m$ does not hold. Therefore, for any positive constant $\kappa < 1$, it holds whp that

$$
\begin{aligned}
A_{t'} &\leq |Y| - |X| + \alpha m \\
&\leq A_t + \epsilon(1 + \kappa)m - ((1 - e^{-\epsilon})(1 - \kappa)(1 + \lambda)m) + \epsilon m\frac{\lambda}{2} \qquad (60) \\
&\leq A_t + \epsilon m(1 + \kappa - (1 - \epsilon)(1 - \kappa)(1 + \lambda) + \frac{\lambda}{2}) \qquad (61) \\
&\leq A_t - \epsilon m\lambda/2, \qquad (62)
\end{aligned}
$$

where (60) holds by Claim 14, (61) holds because $e^{-\epsilon} \leq 1 - \epsilon + \epsilon^2$ for every $\epsilon \in (0, 1)$, and (62) holds by algebraic manipulation and the fact that $\kappa < \epsilon < 1$. The promised claim is proved. □

Lemma 27. *Let $\delta > 0$ be a constant such that $\delta \leq \frac{\lambda(1+\lambda)}{3+2\lambda}$, and $t > 0$. Then, whp, $A_t < \delta m$.*

Proof. Let $\epsilon = \frac{\delta}{4+2\lambda}$ and $\kappa = \lceil \frac{1+\delta}{\epsilon\lambda/2} \rceil$. Also, define $t_i = \epsilon i$, for $i = 0, \ldots, \kappa$. We recall that $A_{t_0} \leq (1+\delta)m$ holds whp, by Lemma 24. Also, Lemmas 25 and 26 apply because $\epsilon < \frac{\lambda/3}{3+2\lambda}$. These lemmas together imply that $A_{t_\kappa} \leq m\epsilon(4+2\lambda)$ holds whp, since $(1+\delta)m \leq \kappa\epsilon\lambda m/2$. The bound $A_{t_\kappa} \leq m\epsilon(4+2\lambda)$, together with $\epsilon = \frac{\delta}{4+2\lambda}$, implies that $A_{t_\kappa} \leq \delta m$, which completes the proof. □

Lemma 28. *Let $\epsilon \in (0, \frac{\lambda/3}{3+2\lambda}]$ be a constant and $t > 0$ be an arbitrary time. Then, whp, $B_t \geq \epsilon\lambda m/2$.*

Proof. Since $(A_0, B_0) \sim \rho$ by definition, then it suffices the prove the claim for the case of $t = \epsilon$. Let X denote the set of all H agents who are present in the pool at some point in the interval $[0, t]$ and whose criticality time is larger than t. By Claim 14, it holds whp that

$$|X| \geq (1 - e^{-\epsilon})(1-\epsilon)(1+\lambda)m. \tag{63}$$

Also, let Y denote the set of all E agents who are present in the pool at some point in the interval $[0, t]$. Let $\delta = \epsilon^2/2$ and $\kappa = \epsilon/2$. It holds whp that

$$|Y| \leq \delta m + \epsilon(1+\kappa)m \tag{64}$$

$$= \epsilon m(1 + \kappa + \epsilon/2) = \epsilon m(1+\epsilon), \tag{65}$$

where the first summand on the right-hand side of (64) corresponds to the E agents who are in the pool at time 0 and is due to Lemma 27, and the second summand corresponds to the E agents who arrive at some point in the interval $(0, t]$ and is due to Claim 14.

We now can write

$$B_t \geq |X| - |Y|$$

$$\geq ((1 - e^{-\epsilon})(1-\epsilon)(1+\lambda)m) - \epsilon(1+\epsilon)m \tag{66}$$

$$\geq \epsilon m((1-\epsilon)(1-\epsilon)(1+\lambda) - 1 - \epsilon) \tag{67}$$

$$= \epsilon m((1 + \epsilon^2 - 2\epsilon)(1+\lambda) - 1 - \epsilon)$$

$$\geq \epsilon m(\lambda - 3\epsilon - 2\epsilon\lambda)$$

$$\geq \epsilon m\lambda/2, \tag{68}$$

where (66) holds by (63) and (65), (67) holds because $e^{-\epsilon} \leq 1 - \epsilon + \epsilon^2$ for every $\epsilon \in (0, 1)$, and (68) holds because $\epsilon \leq \frac{\lambda/3}{3+2\lambda}$. □

iii.3 Proof of Theorem 6

Definition 16. *For an integer $i \geq 0$, an agent present in the pool at time $(i+1)T_m$ is called a* new *agent if she has arrived later than time iT_m.*

For the next lemma, we recall that in a graph G whose nodes correspond to E and H agents, an edge (u, v) is a *cross-edge* if u, v are agents of different types.

Lemma 29. *Let e_i denote the number of new E agents who are present in the pool at time $(i+1)T_m$. Then, the matching executed at the matching time $(i+1)T_m$ involves at least e_i cross-edges whp.*

Proof. We first construct a bipartite graph $G(X, Y)$, where Y is the set of new E agents who are present in the pool at time $(i+1)T_m$ and X is the set of all H agents present in the pool at time $(i+1)T_m$, before the matching is executed.

Claim 16. *Let c_λ be the constant given by Theorem 5. Then, whp, it holds that $|Y| < c_\lambda m \leq |X|$.*

Proof. The second inequality holds by Theorem 5. To prove that $|Y| < c_\lambda m$ holds whp, we note that $|Y|$ is distributed according to a Poisson distribution with mean $\eta_m = md(1 - e^{-T_m/d})$, by Lemma 19. Observe that

$$\eta_m = md(1 - e^{-T_m/d}) \leq md(1 - (1 - T_m/d)) = mT_m.$$

Hence, $\eta_m < mc_\lambda/2$ holds for sufficiently large m. The concentration bound of Fact 3 for the Poisson distribution then directly implies that $|Y| < mc_\lambda$ holds whp. □

The next claim concludes the proof.

Claim 17. *For any integer $i \geq 0$, the matching that is executed at time $(i+1)T_m$ involves at least $|Y|$ cross-edges, whp.*

Proof. By Theorem 5, $|X| \geq c_\lambda m$ holds whp. Also, by Claim 16, $|Y| < |X|$ holds whp. Thus, by Corollary 3, there exists a matching of size $|Y|$ in G whp. Therefore, by Claim 8, the matching that is executed at time $(i+1)T_m$ involves at least $|Y|$ cross-edges, whp. □

□

By Lemma 19, for every integer $i \geq 0$ we have that $\mathbb{E}[e_i] = mT_m \frac{d(1-e^{-T_m/d})}{T_m}$. Therefore, by Lemma 29, the match rate under the batching policy is at least

$$mT_m \frac{d(1 - e^{-T_m/d})}{T_m} \cdot \frac{1}{mT_m} = \frac{d(1 - e^{-T_m/d})}{T_m},$$

when the market size is m. This fact, together with $\lim_{T_m \to 0} \frac{d(1-e^{-T_m/d})}{T_m} = 1$, implies that the match rate of E agents approaches 1 as m approaches infinity. By the same argument, Lemma 29 also implies that the match rate of H agents approaches $\frac{1}{1+\lambda}$ as m approaches infinity.

iii.4 Proof of Theorem 7

Lemma 23 shows that, for type $\Theta \in \{E, H\}$, $w^B_\Theta(m) = d(1-q^B_\Theta(m))$. On the other hand, Theorem 6 shows that $\lim_{m\to\infty} q^B_E(m) = 1$ and $\lim_{m\to\infty} q^B_H(m) = \frac{1}{1+\lambda}$ hold when $\lim_{m\to\infty} T_m = 0$. Therefore,

$$\lim_{m\to\infty} w^B_E(m) = d(1-1) = 0,$$
$$\lim_{m\to\infty} w^B_H(m) = d(1 - \frac{1}{1+\lambda}) = \frac{d\lambda}{1+\lambda}$$

hold when $\lim_{m\to\infty} T_m = 0$. The proof is complete.

iv Analysis of the Patient Policy

In this section, we analyze the stochastic process corresponding to the patient policy and develop a concentration bound on the number of easy-to-match agents present in the pool at the steady state. Because one can renormalize the time scale and the arrival rates linearly with a factor of $1/d$, we suppose that $d = 1$ throughout this section. This is without generality, as speeding up or slowing down time does not change the steady state distribution of the number of easy-to-match (or hard-to-match) agents in the pool.

After some preliminary definitions, we present the core technical result in the analysis of the patient policy, Theorem 8, and prove it in Section iv.2. After that, we prove our results about the match rate and the distribution of waiting time under the patient policy in Section v.

iv.1 Definitions

For the analysis of the patient policy, we use a two-dimensional continuous-time Markov chain, $\mathcal{M}$, to model the dynamics. Let $V(\mathcal{M})$ denote the state space of $\mathcal{M}$. We represent each state by a pair (x, y) where x, y respectively denote the number of H agents and E agents. In other words,

$$V(\mathcal{M}) = \{(x, y) : \; x, y \in \mathbb{Z}_+\}.$$

We note that $\mathcal{M}$ is irreducible because there is only one communication class; i.e., it is possible to reach from any state to every other state. Also, $\mathcal{M}$ is positive recurrent by a proof identical to Lemma 2 (which shows that even if the policy makes no matches, the underlying stochastic process is positive recurrent; hence, so is $\mathcal{M}$). Therefore, $\mathcal{M}$ has a unique steady-state distribution, by the Ergodic theorem (Norris (1997), Theorem 3.5.3).

Definition 17. *Let π denote the stationary distribution of $\mathcal{M}$, with $\pi_{x,y}$ denoting the probability that π assigns to the state (x, y). Also, let $\pi_x = \sum_{y=0}^{\infty} \pi_{x,y}$.*

iv.2 The Analysis

The following theorem is the key to the analysis:

Theorem 8. $\mathbb{E}_{(x,y)\sim\pi}[y] = O(\sqrt{m})$.

Proof. First, we define another Markov process, namely $\mathcal{M}'$, which has the same state space as $\mathcal{M}$. The arrival processes and the criticality times for agents in $\mathcal{M}'$ are identical to $\mathcal{M}$. The only difference between $\mathcal{M}$ and $\mathcal{M}'$ is that, in $\mathcal{M}'$, no two E agents are considered to be compatible. We note that $\mathcal{M}'$ is irreducible because it has a single communication class. Also, $\mathcal{M}'$ is positive recurrent by a proof identical to Lemma 2 (which shows that even if the policy makes no matches, the underlying stochastic process is positive recurrent; hence, so is $\mathcal{M}'$). This implies that $\mathcal{M}'$ has a unique steady-state distribution due to the Ergodic theorem (Norris (1997), Theorem 3.5.3). Let π' denote the steady state distribution of $\mathcal{M}'$.

At any time, the size of the H pool in $\mathcal{M}$ is at least as large as the size of the H pool in $\mathcal{M}'$ in any sample path. (This follows from a similar argument we made earlier in Lemma 4.) Therefore, the steady state size of the H pool in $\mathcal{M}$ is at least as large as the steady state size of the H pool in $\mathcal{M}'$. Hence, the prove the claim, it suffices to show that

$$\mathbb{E}_{(x,y)\sim\pi'}[y] = O(\sqrt{m}). \tag{69}$$

Let $\widehat{\mathcal{M}'}$ denote the embedded Markov chain of $\mathcal{M}'$. To prove (69), we apply Proposition 9 on $\widehat{\mathcal{M}'}$ as follows. Let π'' denote the steady-state distribution of $\widehat{\mathcal{M}'}$. We will first show that

$$\mathbb{E}_{(x,y)\sim\pi''}[y] = O(\sqrt{m}). \tag{70}$$

To this end, let $x^* = (1+\lambda)m$. To apply Proposition 9 on $\widehat{\mathcal{M}'}$, we first define the functions $f, U : V(\widehat{\mathcal{M}'}) \to \mathbb{R}_+$ to be

$$f(x,y) = \frac{|x-x^*| + y}{\sqrt{m}},$$
$$U(x,y) = (x-x^*)^2 + y^2.$$

Let $\delta = \frac{112(1+\lambda)}{\min\{1,\lambda\}}$. Also, let $\tilde{y} = \max\{1, \log_{1-p}\frac{\lambda}{1+\lambda}\}$ and $k = \delta\tilde{y}$. Define

$$B = \{(x,y) \in V(\widehat{\mathcal{M}'}) : |x-x^*|, y \le k\sqrt{m}\}.$$

Definition 18. *Conditional on the Markov chain $\widehat{\mathcal{M}'}$ being at a state (x,y), let the random variable (x_1,y_1) denote the next state that the Markov chain moves to. Define*

$$\Delta(x,y) = \mathbb{E}\left[U(x_1,y_1)\right] - U(x,y).$$

We bound $\Delta(x, y)$ in a series of claims that follow next.

Claim 18. *The following holds for sufficiently large m: if $(x, y) \notin B$ and $x > x^*$, then*

$$\frac{\Delta(x, y)}{2} - 2 \leq -\frac{\lambda}{6(1+\lambda)} \cdot f(x, y).$$

Proof. All of the bounds that we write throughout this proof hold for sufficiently large m; i.e., there exists a constant m_0 such that the bounds hold for all $m > m_0$.

Let $z = x - x^*$. We observe that

$$\begin{aligned}\Delta(x, y) = (2y+1) \cdot \frac{m}{(2+\lambda)m + x + y} - (2y-1) \cdot \frac{y + x\bar{N}_y}{(2+\lambda)m + x + y} \\ - (2z-1) \cdot \frac{x + y\bar{N}_x}{(2+\lambda)m + x + y} + (2z+1) \cdot \frac{m(1+\lambda)}{(2+\lambda)m + x + y}.\end{aligned} \tag{71}$$

The above equality holds because the transition probability of $\widehat{\mathcal{M}'}$ from the state (x, y)

- to the state $(x, y+1)$ is $\frac{m}{(2+\lambda)m+x+y}$ (where the value of U increases by $2y+1$),
- to the state $(x+1, y)$ is $\frac{m(1+\lambda)}{(2+\lambda)m+x+y}$ (where the value of U increases by $2z+1$),
- to the state $(x, y-1)$ is $\frac{yN_x}{(2+\lambda)m+x+y}$ (where the value of U decreases by $2y-1$),
- to the state $(x-1, y-1)$ is $\frac{y\bar{N}_x + x\bar{N}_y}{(2+\lambda)m+x+y}$ (where the value of U decreases by $(2y-1)+(2z-1)$),
- to the state $(x-1, y)$ is $\frac{xN_y}{(2+\lambda)m+x+y}$ (where the value of U decreases by $2z-1$).

We note that the transition probabilities are 0 if $x-1$ or $y-1$ are negative. Multiplying the amount of change in U by the corresponding transition probability for each transition, summing up the resulting expressions, and factoring the terms $2y-1, 2y+1, 2z-1, 2z+1$ in the summation gives (71).

The proof considers two cases. Either $y \leq z$ or $y > z$.

Case 1. $y \leq z$**:** In this case, we must have $x \geq (1+\lambda)m + k\sqrt{m}$. Because otherwise $z < k\sqrt{m} \leq y$ holds, which would be a contradiction.

To simplify the inequality (71), we next observe that, for $y \geq \tilde{y}$,

$$-y\frac{y + x\bar{N}_y - m}{(2+\lambda)m + x + y} \leq -y\frac{x\bar{N}_y - m}{(2+\lambda)m + x + y} \leq 0,$$

where the latter inequality holds because $\tilde{y} \geq \log_{1-p} \frac{\lambda}{1+\lambda}$. On the other hand, for $y \leq \tilde{y}$,

$$-y\frac{y + x\bar{N}_y - m}{(2+\lambda)m + x + y} \leq -y\frac{-m}{(2+\lambda)m + x + y} \leq \tilde{y}\frac{m}{(2+\lambda)m + x + y} \leq \tilde{y}.$$

By the above two inequalities, for every $y \geq 0$ and $x \geq (1+\lambda)m$ it holds that

$$-y \frac{y + x\bar{N}_y - m}{(2+\lambda)m + x + y} \leq \tilde{y}. \tag{72}$$

We now use (71) to write

$$\begin{aligned}
\frac{\Delta(x,y)}{2} - 2 &= -y \cdot \frac{y + x\bar{N}_y - m}{(2+\lambda)m + x + y} - z \cdot \frac{x + y\bar{N}_x - m(1+\lambda)}{(2+\lambda)m + x + y} && (73)\\
&\leq \tilde{y} - z \cdot \frac{x - (1+\lambda)m}{(2+\lambda)m + x + z} = \tilde{y} - z \cdot \frac{x - (1+\lambda)m}{m + 2x} && (74)\\
&\leq \tilde{y} - z \cdot \frac{k\sqrt{m}}{(3+2\lambda)m + 2k\sqrt{m}} && (75)\\
&\leq \tilde{y} - z \cdot \frac{k}{3(1+\lambda)\sqrt{m}} && (76)\\
&\leq \tilde{y} - 2z \frac{\tilde{y}}{\sqrt{m}} && (77)\\
&\leq -z \frac{\tilde{y}}{\sqrt{m}} \leq -\frac{\tilde{y}}{2} \cdot f(x,y), && (78)
\end{aligned}$$

where (74) holds by (72) and the fact that $y \leq z$, (75) holds by Fact 6, (76) holds for sufficiently large m, (77) holds because $k \geq 6(1+\lambda)\tilde{y}$, the penultimate inequality holds because $z \geq \sqrt{m}$, and and the last inequality holds because $y \leq z$.

Case 2. $y > z$**:** From $y > z$ and $(x,y) \notin B$ it implies that $y \geq k\sqrt{m}$. We use (71) to write

$$\begin{aligned}
\frac{\Delta(x,y)}{2} - 2 &= -y \cdot \frac{y + x\bar{N}_y - m}{(2+\lambda)m + x + y} - z \cdot \frac{x + y\bar{N}_x - m(1+\lambda)}{(2+\lambda)m + x + y}\\
&\leq -y \cdot \frac{y + x\bar{N}_y - m}{(2+\lambda)m + x + y} && (79)\\
&= -y \cdot \frac{y + x\bar{N}_y - m}{(3+2\lambda)m + z + y}\\
&\leq -y \cdot \frac{z + x\bar{N}_y - m}{(3+2\lambda)m + 2z} && (80)\\
&\leq -y \cdot \frac{\lambda m}{3(1+\lambda)m} && (81)\\
&\leq -\frac{\lambda}{6(1+\lambda)} \cdot f(x,y) && (82)
\end{aligned}$$

where (79) holds because $x > (1+\lambda)m$, (80) holds for sufficiently large m by Fact 5, (81) holds because $x\bar{N}_y > (1+\lambda)m - \sqrt{m}$ for sufficiently large m, and the last inequality holds since $y > z$. Finally, $\tilde{y} \geq 1$, (78), and (82) together conclude the proof. □

Fact 9. *Let $z = x^* - x$. Then, when $x < x^*$,*

$$\frac{\Delta(x,y)-2}{2} = y \cdot \frac{m-(x\bar{N}_y + y)}{(2+\lambda)m+x+y} + z \cdot \frac{x+y\bar{N}_x-(1+\lambda)m}{(2+\lambda)m+x+y}. \tag{83}$$

Proof. We derive the following equality from the transition rates of $\widehat{\mathcal{M}'}$, in a way similar to (71).

$$\begin{aligned}\Delta(x,y) &= (2y+1)\cdot\frac{m}{(2+\lambda)m+x+y} - (2y-1)\cdot\frac{y+x\bar{N}_y}{(2+\lambda)m+x+y} \\ &\quad + (2z+1)\cdot\frac{x+y\bar{N}_x}{(2+\lambda)m+x+y} - (2z-1)\cdot\frac{m(1+\lambda)}{(2+\lambda)m+x+y}.\end{aligned}$$

Rearranging the terms in the above equation gives the bound claimed. ☐

Claim 19. *The following holds for sufficiently large m: if $x < x^* - k\sqrt{m}$ and $x+y < x^*$, then*

$$\frac{\Delta(x,y)-2}{2} \le -\min\left\{\frac{\tilde{y}}{2}, \frac{1}{1+\lambda}\right\}\cdot f(x,y).$$

Proof. Let $z = x^* - x$. We observe that $x+y < x^*$ implies that $y < z$. We consider 2 cases for the proof: $y \ge \tilde{y}$ or $y < \tilde{y}$.

Case 1: $y \ge \tilde{y}$**.** Recall that $z > y$. Hence, by (83),

$$\begin{aligned}\frac{\Delta(x,y)}{2} - 2 &\le y\cdot\frac{m-(x\bar{N}_y+y)}{(2+\lambda)m+x+y} + z\cdot\frac{x+y\bar{N}_x-(1+\lambda)m}{(2+\lambda)m+x+y} \\ &\le z\cdot\frac{m-(x\bar{N}_y+y)}{(2+\lambda)m+x+y} + z\cdot\frac{x+y\bar{N}_x-(1+\lambda)m}{(2+\lambda)m+x+y} \\ &= z\cdot\frac{-\lambda m + xN_y}{(2+\lambda)m+x+y} \\ &\le z\cdot\frac{-k\frac{\lambda}{1+\lambda}\sqrt{m}}{(3+2\lambda)m} \\ &= z\cdot\frac{-k\frac{\lambda}{1+\lambda}}{(3+2\lambda)\sqrt{m}} = \frac{-k\lambda}{2(3+2\lambda)(1+\lambda)}\cdot f(x,y) \le -\frac{1}{1+\lambda}\cdot f(x,y),\end{aligned}$$

where the first inequality holds by (83), the second inequality holds because $z > y$, the third inequality holds because $N_y \le (1-p)^{\tilde{y}} = \frac{\lambda}{1+\lambda}$, and the last inequality holds because $k \ge 6(1+\lambda)$.

Case 2: $y < \tilde{y}$**.** Using (83) we can write

$$\frac{\Delta(x,y)}{2} - 2 \le y \cdot \frac{m - (x\bar{N}_y + y)}{(2+\lambda)m + x + y} + z \cdot \frac{x + y\bar{N}_x - (1+\lambda)m}{(2+\lambda)m + x + y}$$
$$= \tilde{y}\frac{1}{2+\lambda} + z \cdot \frac{\tilde{y} - k\sqrt{m}}{(2+\lambda)m + x + y}$$
$$\le \tilde{y}\frac{1}{2+\lambda} - z \cdot \frac{(k-1)\sqrt{m}}{3(1+\lambda)m} \tag{84}$$
$$\le \tilde{y}\frac{1}{2+\lambda} - \frac{5\tilde{y}}{3}\frac{z}{\sqrt{m}} \tag{85}$$
$$= \tilde{y}\frac{1}{2+\lambda} - \tilde{y}\frac{2}{3}\frac{z}{\sqrt{m}} - \tilde{y}\frac{z}{\sqrt{m}} \le -\tilde{y}\frac{z}{\sqrt{m}} \tag{86}$$
$$\le -\frac{\tilde{y}}{2} \cdot f(x,y),$$

where (84) holds since $\tilde{y} < \sqrt{m}$ for sufficiently large m, (85) holds because $k \ge 6(1+\lambda)\tilde{y}$, the inequality in (86) holds because $z \ge k\sqrt{m} > \sqrt{m}$, and the last inequality holds because $y < z$.

Finally, we observe that the upper bounds that Case 1 and Case 2 provide on $\frac{\Delta(x,y)}{2} - 2$ together imply that $\frac{\Delta(x,y)}{2} - 2 \le -\min\{\frac{\tilde{y}}{2}, \frac{1}{1+\lambda}\} \cdot f(x,y)$. □

Claim 20. *The following holds for sufficiently large m: if $x < x^* - k\sqrt{m}$ and $x + y \ge x^*$ then*

$$\frac{\Delta(x,y)}{2} - 2 \le -\frac{\lambda}{21(1+\lambda)} \cdot f(x,y).$$

Proof. Let $z = x^* - x$. Observe that $x + y \ge x^*$ implies that $y \ge z \ge k\sqrt{m}$. We consider two cases: either $y \le 4(1+\lambda)m$ or not.

Case 1: $y \le 4(1+\lambda)m$**.** The inequality $y \ge z$, together with (83), implies that

$$\frac{\Delta(x,y)}{2} - 2 \le y \cdot \frac{m - (x\bar{N}_y + y)}{(2+\lambda)m + x + y} + y \cdot \frac{x + y - (1+\lambda)m}{(2+\lambda)m + x + y}$$
$$\le y \cdot \frac{-\lambda m + xN_y}{(2+\lambda)m + x + y}$$
$$\le y \cdot \frac{-2\lambda m/3}{(2+\lambda)m + x + y} \tag{87}$$
$$\le -\frac{\lambda m/3}{(2+\lambda)m + x + y} \cdot \frac{y}{y + z/\sqrt{m}} \cdot f(x,y) \tag{88}$$
$$= -\frac{\lambda m/3}{(2+\lambda)m + x + y} \cdot \frac{1}{1 + z/(y\sqrt{m})} \cdot f(x,y)$$
$$\le -\frac{\lambda}{21(1+\lambda)} \cdot f(x,y) \tag{89}$$

holds for sufficiently large m, where (87) holds because $xN_y \le \lambda m/3$ for sufficiently large m, (88)

holds because $f(x,y) \le y + \frac{z}{\sqrt{m}}$, and (89) holds since $(2+\lambda)m + x + y \le 7(1+\lambda)m$ and $z/y \le 1$.
Case 2: $y > 4(1+\lambda)m$**.** By (83),

$$
\begin{aligned}
\frac{\Delta(x,y)}{2} - 2 &\le y \cdot \frac{m - (x\bar{N}_y + y)}{(2+\lambda)m + x + y} + z \cdot \frac{x + y - m(1+\lambda)}{(2+\lambda)m + x + y} \\
&\le -y \cdot \frac{x\bar{N}_y - m + y}{(2+\lambda)m + x + y} + z \\
&\le -y \cdot \frac{x\bar{N}_y - m + 4m(1+\lambda)}{(2+\lambda)m + x + 4m(1+\lambda)} + z \qquad (90) \\
&\le -y \cdot \frac{(3+4\lambda)m}{(7+6\lambda)m} + z \qquad (91) \\
&\le -y\frac{3}{7} + z \le -\frac{3}{14} \le -\frac{3}{28} \cdot f(x,y), \qquad (92)
\end{aligned}
$$

where (90) holds by $y > 4m(1+\lambda)$ and Fact 5, (91) holds since $x \le (1+\lambda)m$, and (92) holds since $y \ge z$. Finally, (89) and (92), together with $\frac{3}{28} > \frac{\lambda}{21(1+\lambda)}$, conclude the proof. □

Claim 21. *When $(x,y) \notin B$ and $x^* - k\sqrt{m} \le x < x^*$, then*

$$\frac{\Delta(x,y)}{2} - 2 \le -\frac{\lambda}{28(1+\lambda)} \cdot f(x,y).$$

Proof. Let $z = x - x^*$. Observe that $z \le y$ holds because $z \le k\sqrt{m} \le y$. We consider two cases: either $y < 4(1+\lambda)m$ or not.
Case 1: $y < 4(1+\lambda)m$**.** By the inequalities (83) and $y \ge z$,

$$
\begin{aligned}
\frac{\Delta(x,y)}{2} - 2 &\le y \cdot \frac{m - (x\bar{N}_y + y)}{(2+\lambda)m + x + y} + y \cdot \frac{x + y - (1+\lambda)m}{(2+\lambda)m + x + y} \\
&\le y \cdot \frac{-\lambda m + xN_y}{(2+\lambda)m + x + y} \\
&\le -y \cdot \frac{\lambda m/2}{(2+\lambda)m + x + y} \qquad (93) \\
&\le -\frac{\lambda m/2}{(2+\lambda)m + x + y} \cdot \frac{y}{y + z/\sqrt{m}} \cdot f(x,y) \qquad (94) \\
&= -\frac{\lambda m/2}{(2+\lambda)m + x + y} \cdot \frac{1}{1 + z/(y\sqrt{m})} \cdot f(x,y) \\
&\le -\frac{\lambda}{14(1+\lambda)} \cdot \frac{1}{2} \cdot f(x,y) \qquad (95)
\end{aligned}
$$

holds for sufficiently large m, where (93) holds because $xN_y < \lambda m/2$ for sufficiently large m, (94) holds since $f(x,y) > y + \frac{z}{\sqrt{m}}$, and (95) holds since $z/y \le 1$.

Case 2: $y \geq 4(1+\lambda)m$**.** Inequalities $x + y \geq 4(1+\lambda)m$ and (83) imply that

$$
\begin{aligned}
\frac{\Delta(x,y)}{2} - 2 &\leq y \cdot \frac{m - (x\bar{N}_y + y)}{(2+\lambda)m + x + y} + z \\
&\leq y \cdot \frac{m - (x\bar{N}_y + 4(1+\lambda)m)}{(2+\lambda)m + x + 4(1+\lambda)m} + z \qquad (96) \\
&\leq -y \cdot \frac{(3+3\lambda)m}{(7+6\lambda)m} + z \\
&\leq -y\frac{3}{7} + z \leq -y\frac{3}{14} \leq -\frac{3}{28} \cdot f(x,y) \qquad (97)
\end{aligned}
$$

holds for sufficiently large m, where (96) holds by Fact 5, and (97) holds because $y \geq 4z$.

Finally, (95) and (97), together with $\frac{3}{28} > \frac{\lambda}{28(1+\lambda)}$, imply that $\frac{\Delta(x,y)}{2} - 2 \leq -\frac{\lambda}{28(1+\lambda)} f(x,y)$. This completes the proof of the claim. □

Claim 22. *The following holds for sufficiently large m: if $(x,y) \notin B$, then*

$$\frac{\Delta(x,y)}{2} - 2 \leq -\frac{\min\{1,\lambda\}}{28(1+\lambda)} \cdot f(x,y).$$

Proof. Each of the Claims 18, 19, 20, and 21 provides an upper bound on $\frac{\Delta(x,y)}{2} - 2$ for when (x,y) belongs to a particular subset of the states $V(\widehat{\mathcal{M}'})$. Since these subsets cover $V(\widehat{\mathcal{M}'})\backslash B$, then taking the minimum of the right-hand sides of the upper bounds provided by these claims yields a valid upper bound for $\frac{\Delta(x,y)}{2} - 2$ when $(x,y) \in V(\widehat{\mathcal{M}'})\backslash B$. Therefore, when $(x,y) \notin B$,

$$\frac{\Delta(x,y)}{2} - 2 \leq -\min\left\{\frac{\lambda}{6(1+\lambda)}, \frac{\tilde{y}}{2}, \frac{1}{1+\lambda}, \frac{\lambda}{21(1+\lambda)}, \frac{\lambda}{28(1+\lambda)}\right\} \cdot f(x,y).$$

Since $\tilde{y} > 1$, then

$$\min\left\{\frac{\lambda}{6(1+\lambda)}, \frac{\tilde{y}}{2}, \frac{1}{1+\lambda}, \frac{\lambda}{21(1+\lambda)}, \frac{\lambda}{28(1+\lambda)}\right\} \geq \frac{\min\{1,\lambda\}}{28(1+\lambda)},$$

which concludes the proof. □

Claim 23. *The following holds for sufficiently large m: For $(x,y) \in B$, $\Delta(x,y) \leq 2(\tilde{y} + k^2) + 2$.*

Proof. We consider three cases: either $x > x^*$, $x < x^*$, or $x = x^*$.

Case 1: $x > x^*$**.** We recall (72) and (73), which hold when $x > x^*$, and respectively state that

$$
\begin{aligned}
-y\frac{y + x\bar{N}_y - m}{(2+\lambda)m + x + y} &\leq \tilde{y}, \\
\frac{\Delta(x,y)}{2} - 2 &\leq -y\frac{y + x\bar{N}_y - m}{(2+\lambda)m + x + y} - z\frac{x + y\bar{N}_x - m(1+\lambda)}{(2+\lambda)m + x + y}.
\end{aligned}
$$

We next note that, since $x > m(1+\lambda)$,

$$-z \cdot \frac{x + y\bar{N}_x - m(1+\lambda)}{(2+\lambda)m + x + y} \leq 0.$$

This inequality and (72) provide upper bounds for the summands on the right-hand sides of (73). Adding up these upper bounds implies that $\frac{\Delta(x,y)}{2} - 2 \leq \tilde{y}$.

Case 2: $x < x^*$. First, we observe that, for sufficiently large m,

$$-y\frac{y + x\bar{N}_y - m}{(2+\lambda)m + x + y} \leq \tilde{y}. \tag{98}$$

This holds because, if $y \geq 2\tilde{y}$, then $x\bar{N}_y - m \geq 0$ holds for sufficiently large m, and thus the left-hand side is nonpositive. Otherwise, if $y < 2\tilde{y}$, the left-hand side is at most $\tilde{y}$ since $\frac{m}{(2+\lambda)m+x+y} \leq \frac{1}{2}$.

We recall that (83) holds when $x < x^*$. The inequality (83), together with (98), implies that

$$\begin{aligned}\frac{\Delta(x,y)}{2} - 2 &\leq \tilde{y} + z \cdot \frac{x + y\bar{N}_x - (1+\lambda)m}{(2+\lambda)m + x + y} \\ &\leq \tilde{y} + z \cdot \frac{k\sqrt{m}}{(2+\lambda)m + x + y} = \tilde{y} + \frac{k^2 m}{(2+\lambda)m + x + y} \leq \tilde{y} + k^2,\end{aligned} \tag{99}$$

which means that $\Delta(x,y) \leq 2(\tilde{y} + k^2) + 2$.

Case 3: $x = x^*$. In this case, $z = 0$. We have

$$\begin{aligned}\Delta(x,y) &= (2y+1) \cdot \frac{m}{(2+\lambda)m + x + y} - (2y-1) \cdot \frac{x\bar{N}_y + y}{(2+\lambda)m + x + y} \\ &+ (2z+1) \cdot \frac{x + y\bar{N}_x}{(2+\lambda)m + x + y} + (2z+1) \cdot \frac{m(1+\lambda)}{(2+\lambda)m + x + y} \\ &\leq 2y \cdot \frac{m - x\bar{N}_y - y}{(2+\lambda)m + x + y} + 2 \leq 2\tilde{y} + 2,\end{aligned}$$

where the equality holds due to the transition rates of $\widehat{\mathcal{M}'}$ (which were stated in the proof of Claim 18), and the last inequality is due to (72), which holds when $x \geq x^*$.

The upper bounds provided in the three cases above for $\Delta(x,y)$ imply that $\Delta(x,y) \leq 2(\tilde{y}+k^2)+2$ for every $(x,y) \in B$. $\square$

Let $g = \frac{\min\{1,\lambda\}}{28(1+\lambda)}$. Then, when m is sufficiently large, for every $(x,y) \notin B$,

$$\Delta(x,y) \leq -2gf(x,y) + 4 \leq -(2g - \frac{4}{\delta}) \cdot f(x,y) = -gf(x,y), \tag{100}$$

where the first inequality is by Claim 22, the second inequality holds since $f(x,y) \geq k \geq \delta$ for $(x,y) \notin B$, and the equality holds because $\delta = 4/g$, by the definition of δ.

We next observe that, for all $(x, y) \in B$, $f(x, y) \leq 2\delta$ holds by the definition of f. This inequality, together with (100) and Claim 23 provide all the components required for applying Proposition 9 on $\widehat{\mathcal{M}'}$. Applying this proposition implies that for sufficiently large m,

$$\mathbb{E}_{(x,y)\sim\pi''}\left[f(x,y)\right] \leq 2\delta + \frac{2(\tilde{y}+k^2)+2}{g}, \tag{101}$$

where we recall that π'' denotes the steady state distribution of $\widehat{\mathcal{M}'}$. Let θ denote the right-hand side of the above inequality. Since $\delta, g, k, \tilde{y}$ are constants (independent of m), so is θ. Recall that $f(x, y) = \frac{|x-x^*|+y}{\sqrt{m}}$. Hence, by the above inequality,

$$\mathbb{E}_{(x,y)\sim\pi''}\left[|x - x^*|\right] \leq \theta\sqrt{m} = O(\sqrt{m}), \tag{102}$$

$$\mathbb{E}_{(x,y)\sim\pi''}\left[y\right] \leq \theta\sqrt{m} = O(\sqrt{m}). \tag{103}$$

Claim 24. *The following inequalities hold for sufficiently large m:*

$$\mathbb{P}_{(x,y)\sim\pi''}\left[|x - x^*| \geq \theta\sqrt{m}\log m\right] \leq \frac{1}{\log m}, \tag{104}$$

$$\mathbb{P}_{(x,y)\sim\pi''}\left[y \geq \theta\sqrt{m}\log m\right] \leq \frac{1}{\log m}. \tag{105}$$

Proof. The proof follows directly from applying the Markov inequality on (102) and (103). □

Finally, we will conclude the proof of the theorem by showing that (69) holds. This is done in the following claim:

Claim 25. $\mathbb{E}_{(x,y)\sim\pi'}\left[y\right] \leq O(\sqrt{m})$ *holds, where π' is the steady-state distribution of $\mathcal{M}'$.*

Proof. Let $V = \mathbb{Z}_+^2$. We recall that the steady-state distribution of $\widehat{\mathcal{M}'}$ is denoted by π'', and $\pi'_{x,y}, \pi''_{x,y}$ denote the probabilities assigned to a node $(x, y) \in V$ by π', π'', respectively.

Let $w_{x,y} = \frac{1}{m(2+\lambda)+x+y}$ for all $(x, y) \in V$, and $x^* = (1 + \lambda)m$. Recall that θ denotes the right-hand side of (101), and let

$$C = \{(x, y) \in V : y, |x - x^*| \leq \theta\sqrt{m}\log m\}.$$

Also, define $\epsilon = \sum_{(x,y)\notin C} \pi''_{x,y}$. By Claim 24, $\epsilon \leq \frac{2}{\log m}$ holds for sufficiently large m. Define

$$\underline{w} = \min\{w_{x,y} : (x, y) \in C\}.$$

Let $s'' = \sum_{(x,y)\in V} w_{x,y}\pi''_{x,y}$. By Fact 2, we have $\pi'_{x,y} = \frac{w_{x,y}\pi''_{x,y}}{s''}$. Using this equation we write

$$\begin{aligned}\pi'_{x,y} &= \frac{w_{x,y}\pi''_{x,y}}{\sum_{(i,j)\in C} w_{i,j}\pi''_{i,j} + \sum_{(i,j)\notin C} w_{i,j}\pi''_{i,j}}\\ &\le \frac{w_{x,y}\pi''_{x,y}}{\sum_{(i,j)\in C} w_{i,j}\pi''_{i,j}} \le \frac{w_{x,y}\pi''_{x,y}}{(1-\epsilon)\underline{w}},\end{aligned}$$

where the last inequality follows from the fact that $\sum_{(i,j)\notin C} \pi''_{i,j} = \epsilon$. By the above bound,

$$\begin{aligned}\sum_{(x,y)\in V} \pi'_{x,y} y &\le \sum_{(x,y)\in V} \frac{w_{x,y}}{(1-\epsilon)\underline{w}} \cdot \pi''_{x,y} y\\ &\le \frac{w_{0,0}}{(1-\epsilon)\underline{w}} \cdot \sum_{(x,y)\in V} \pi''_{x,y} y = \frac{w_{0,0}}{(1-\epsilon)\underline{w}} \cdot \mathbb{E}_{(x,y)\sim\pi''}[y], \end{aligned} \tag{106}$$

where the last inequality holds because $w_{x,y} \le w_{0,0}$ for all $(x,y) \in V$. Also, we observe that

$$\frac{w_{0,0}}{\underline{w}} = \frac{m(2+\lambda) + m(1+\lambda) + 2\theta\sqrt{m}\log m}{m(2+\lambda)} \le 2$$

holds for sufficiently large m, where the inequality holds since θ is a constant (independent of m).

Plugging the latter bound into (106) implies that

$$\mathbb{E}_{(x,y)\sim\pi'}[y] = \frac{2}{1-\epsilon} \cdot \mathbb{E}_{(x,y)\sim\pi''}[y] = O(\sqrt{m}),$$

where the last equality holds by the fact that $\epsilon \le \frac{2}{\log m}$ and (103). □

□

v Match Rate and Waiting Times under the Patient Policy

Throughout this section, we let $\mathcal{M}$ denote the Markov chain associated with the patient policy and π denote its steady-state distribution, as defined in Section iv.

v.1 Match Rates under the Patient Policy

Lemma 30. *Under the patient policy, the match rate of hard-to-match agents is* $\frac{1}{1+\lambda} - O(m^{-1/6})$ *and the match rate of easy-to-match agents is* $1 - O(m^{-1/6})$.

Proof. Let a denote an easy-to-match agent at the steady state. Also, let M_a^H denote the event in which a is matched to a hard-to-match agent. Hence, $\mathbb{P}_\pi\left[M_a^H\right]$ denotes the steady-state probability that an easy-to-match agent is matched to a hard-to-match agent. To prove the lemma, it suffices to show that

$$\mathbb{P}_\pi\left[M_a^H\right] = 1 - O(m^{-1/6}). \tag{107}$$

To this end, let M_a^E denote the event in which a is matched to an easy-to-match agent, and N_a denote the event in which a leaves the pool unmatched. Observe that

$$\mathbb{P}_\pi\left[\overline{M_a^H}\right] = \mathbb{P}_\pi\left[M_a^E\right] + \mathbb{P}_\pi\left[N_a\right]. \tag{108}$$

In the rest of the proof, we will provide an upper bound on the right-hand side of the above equation. Let the random variable T_a denote the amount of time that a is in the pool.

Claim 26. $\mathbb{P}_\pi\left[T_a > m^{-1/3}\right] \leq O(m^{-1/6})$.

Proof. Let the random variable y denote the number of easy-to-match agents in the pool. Hence, $\mathbb{E}_\pi\left[y\right]$ is the expected number of easy-to-match agents in the pool at the steady state. Therefore, $\mathbb{E}_\pi\left[T_a\right] = \mathbb{E}_\pi\left[y\right]/m$ holds by the Little's law. Thus, by Theorem 8, $\mathbb{E}_\pi\left[T_a\right] = O(m^{-1/2})$. Applying the Markov inequality then implies that

$$\mathbb{P}_\pi\left[T_a > m^{-1/3}\right] \leq \frac{O(m^{-1/2})}{m^{-1/3}} = O(m^{-1/6}).$$

□

Claim 27. $\mathbb{P}_\pi\left[M_a^E\right] \leq O(m^{-1/6})$.

Proof. Let M_a^D denote the event in which a is not matched until she becomes critical, but is matched then. Since under the patient policy all matches are made upon the criticality times of agents, then

$$\mathbb{P}_\pi\left[M_a^E\right] \leq 2\mathbb{P}_\pi\left[M_a^D\right]. \tag{109}$$

Let z_a denote the time that agent a enters the pool and z'_a denote the time that agent a becomes critical. (Hence, $z'_a - z_a$ has an exponential distribution with mean d.) Then,

$$\begin{aligned}
\mathbb{P}_\pi\left[M_a^D\right] &\leq \mathbb{P}_\pi\left[M_a^D \cap (z'_a - z_a < m^{-1/3})\right] + \mathbb{P}_\pi\left[M_a^D \cap (z'_a - z_a > m^{-1/3})\right] \\
&\leq \mathbb{P}_\pi\left[z'_a - z_a < m^{-1/3}\right] + \mathbb{P}_\pi\left[T_a > m^{-1/3}\right] \\
&\leq 1 - e^{-m^{-1/3}/d} + O(m^{-1/6}) \qquad (110) \\
&\leq m^{-1/3}/d + O(m^{-1/6}), \qquad (111)
\end{aligned}$$

where (110) follows from the bound $1 + \alpha \leq e^{\alpha}$ for $\alpha \in \mathbb{R}$ and from Claim 26. Finally, plugging the above bound into (109) implies that $\mathbb{P}_\pi\left[M_a^E\right] = O(m^{-1/6})$, which concludes the proof. □

Claim 28. $\mathbb{P}_\pi [N_a] \leq O(m^{-1/6})$.

Proof. We observe that

$$\begin{aligned}\mathbb{P}_\pi [N_a] &= \mathbb{P}_\pi \left[N_a \cap (T_a \leq m^{-1/3})\right] + \mathbb{P}_\pi \left[N_a \cap (T_a > m^{-1/3})\right] \\ &\leq 1 - e^{-m^{-1/3}/d} + \mathbb{P}_\pi \left[T_a > m^{-1/3}\right] \\ &\leq O(m^{-1/3}/d) + O(m^{-1/6}) = O(m^{-1/6}),\end{aligned}$$

where the last inequality follows from Claim 26. □

Claim 27 and Claim 28 provide upper bounds for the summands on the right-hand side of (108). These bounds imply that $\mathbb{P}_\pi \left[\overline{M_a^H}\right] \leq O(m^{-1/6})$. That is, (107) holds and the proof is complete. □

v.2 Distribution of Waiting Times

Lemma 31. *Under the patient policy, as m converges to infinity, the waiting time of an E agent converges in distribution to the degenerate distribution at 0.*

Proof. Consider an E agent, a, and let w_a denote the waiting time of a. Hence, $\mathbb{E}_\pi [w_a]$ denotes the expected waiting time of an E agent at the steady state. For any fixed constant $t > 0$, we will show that $\lim_{m\to\infty} \mathbb{P}_\pi [w_a > t] = 0$. This will prove the claim about waiting times.

Let the random variable y denote the number of E agents in the pool. By Little's law, $\mathbb{E}_\pi [w_a] = \mathbb{E}_\pi [y] / m$, and by Theorem 8, $\mathbb{E}_\pi [y] = O(\sqrt{m})$. Therefore, by Markov inequality, $\mathbb{P} [w_a > t] < O(t/\sqrt{m})$, which means that $\lim_{m\to\infty} \mathbb{P} [w_a > t] = 0$. This proves the claim. □

Lemma 32. *Under the patient policy, as m converges to infinity, the waiting time of an H agent converges in distribution to an exponential random variable with rate $1/d$.*

Proof. Consider an H agent, namely h, upon her arrival. Suppose that h has arrived at a time z_0. Let E_h denote the event in which agent h receives an offer before her criticality clock ticks. Thus, $\mathbb{P}_\pi [E_h]$ denotes the steady state probability of the event E_h happens.

Claim 29. *If $\lim_{m\to\infty} \mathbb{P}_\pi [E_h] = 0$, then the waiting time of an H agent converges in distribution to an exponential random variable with rate $1/d$.*

Proof. Let w denote the waiting time of h and let z denote the time that her criticality clock is set to tick (which will be set upon the arrival of h). For any $t \geq 0$ we have

$$\begin{aligned}\mathbb{P}_\pi [z - z_0 < t] \leq \mathbb{P}_\pi [w < t] &\leq \mathbb{P}_\pi [z - z_0 < t] + \mathbb{P}_\pi [(z - z_0 > t) \wedge E_h] \\ &\leq \mathbb{P}_\pi [z - z_0 < t] + \mathbb{P}_\pi [E_h] \,.\end{aligned}$$

The above inequality and $\lim_{m\to\infty} \mathbb{P}_\pi [E_h] = 0$ imply that

$$\lim_{m\to\infty} \mathbb{P}_\pi [w < t] = \mathbb{P}_\pi [z - z_0 < t],$$

which proves the claim since $z - z_0$ is distributed exponentially with mean $1/d$, by definition. □

Hence, to complete the proof it suffices to prove the following claim:

Claim 30. $\lim_{m\to\infty} \mathbb{P}_\pi [E_h] = 0$.

Proof. The proof uses the proof of Claim 27. Recall from there that, for an E agent a, we define the M_a^D to be the event in which a is matched when she becomes critical. Recall that by (111), $\mathbb{P}_\pi \left[M_a^D\right] \leq O(m^{-1/6})$. Since H agents are not compatible, then

$$(1+\lambda)m \cdot \mathbb{P}_\pi [E_h] \leq m \cdot \mathbb{P}_\pi \left[M_a^D\right],$$

which just says that the average number of H agents per unit of time who are matched before their criticality clock ticks is at most equal to the average number of E agents per unit of time who are matched when they become critical. The above bound and (111) imply that $\mathbb{P}_\pi [E_h] = O(m^{-1/6})$, which proves the claim. □

□

Proof of Proposition 7. The claim about the match rate was proved in Lemma 30, and the claims about the waiting times of E and H agents were proved in Lemma 31 and Lemma 32, respectively. □

vi Comparison to the literature

We first consider our main setup (Section 3) but with a single type, i.e., $\lambda = -1$, and run simulations to compare the greedy, batching, and patient policies. In these simulations we vary the arrival rate per day m, while fixing the criticality rate to $\frac{1}{d} = \frac{1}{360}$ and the probability of compatibility to $q = 0.063$ (consistent with the values of d and q in our calibration to the NKR data). The results are reported in Figure 14. We observe that the match rate under the greedy policy approaches the optimal match rate as the market grows large. On the other hand, the waiting time is lower under the greedy policy than under the other policies considered.

We next consider our main setup but allow the compatibility probabilities to vanish with the market size. Figure 15 compares greedy and patient policies in this case for $\lambda \in \{-1, 0, 1.33\}$ and $p = q \in \{\frac{1}{\sqrt{m}}, \frac{1}{m}\}$. The arrival rate per unit of time is m and the criticality rate is $\frac{1}{d} = 1$.[40] When the compatibility probability equals $\frac{1}{\sqrt{m}}$, we observe that the match rate under the greedy

[40] This normalization would correspond to the setting of Akbarpour et al. (2020) when $\lambda = -1$ and $q = \frac{1}{m}$.

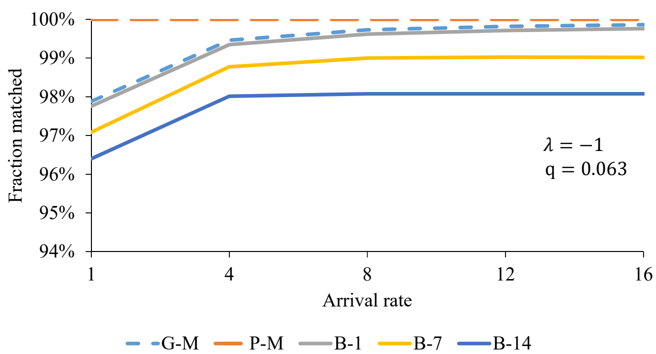

(a) Fraction matched

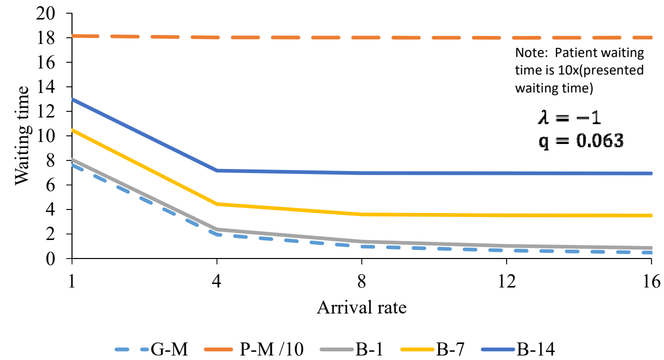

(b) Waiting time

Figure 14: Comparison of greedy (G-M), patient (P-M), and batching (B) policies for different arrival rates m in a single-type model. B-x denotes a batching policy that makes matches every x days.

policy converges to the match rate under the patient policy as the market grows large. When there are more hard-to-match agents than easy-to-match agents ($\lambda = 1.33$), the match rates are indistinguishable. On the other hand, the waiting time under the greedy policy is smaller than the waiting time under the patient policy in all cases.

When the compatibility probability is $\frac{1}{m}$, however, there seems to be a non-vanishing gap between the match rates of greedy and patient policies as well as between their waiting times: the patient policy has a higher match rate while the greedy policy has a lower waiting time.

vii Computing FWM for larger compatibility graphs

The FWP depicted in Figure 1 is 0.057 when the pool includes 4000 pairs (almost all pairs in the data set). One reason that the fraction of pairs without a compatible match does not equal zero is that the data from the four platforms used to generate Figure 1 consists of incompatible pairs that are selectively submitted by hospitals to these platforms. Indeed Agarwal et al. (2019) document that hospitals select to submit their "harder-to-match pairs" and patients of submitted pairs are, on average, much more sensitized than the patients of typical incompatible pairs Agarwal et al. (2019).

To explore the behavior of the FWP beyond our original pool size, we construct larger artificial pools from the same dataset by splitting patient-donor pairs and recreating new pairs while maintaining the blood-type composition of the original pool. Specifically, to generate a single pair, we (i) sample a patient x from all patients in the pool, and (ii) sample a donor from the pool among all donors that have the same blood type as the intended donor of x.

Figure 16 plots the FWP for different pool sizes. For a pool size of 4000 the FWP is about 0.06 and slightly larger than 0.057. For a pool size of around 15600, the FWP decreases to 0.0258.

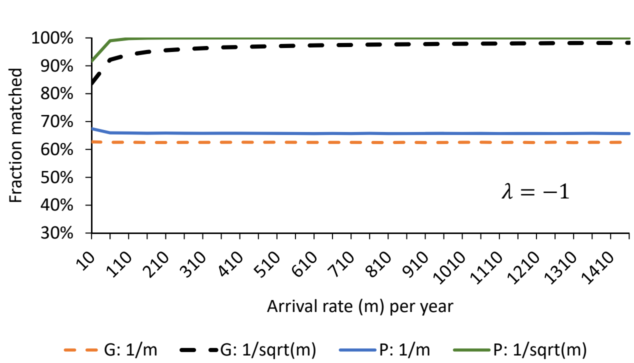

(a) Match rate

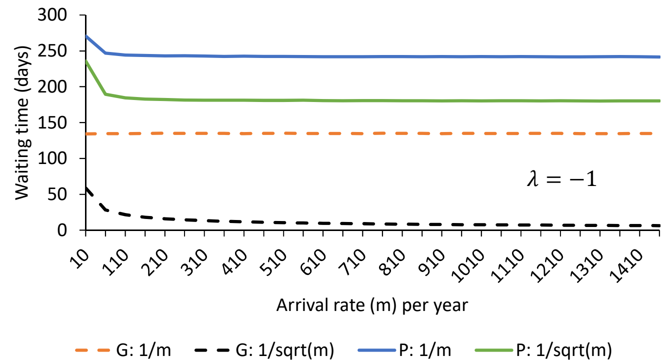

(b) Waiting time

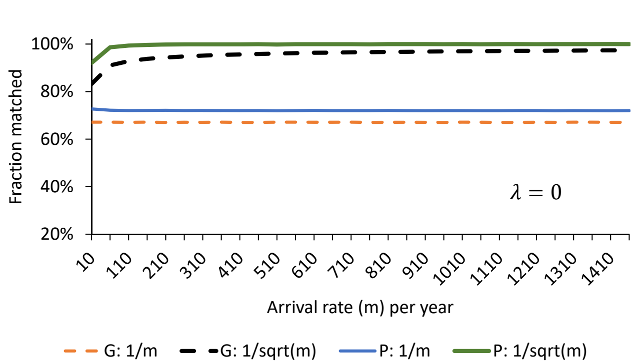

(c) Match rate

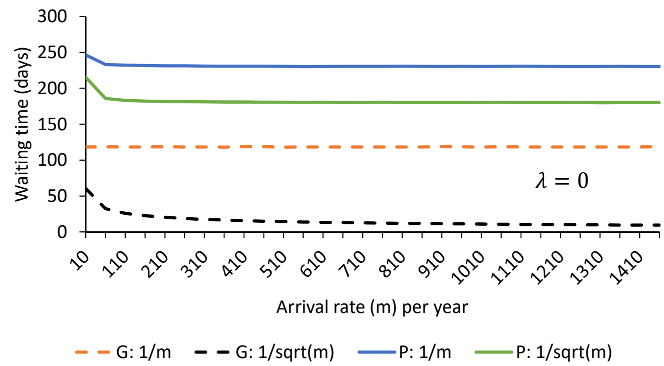

(d) Waiting time

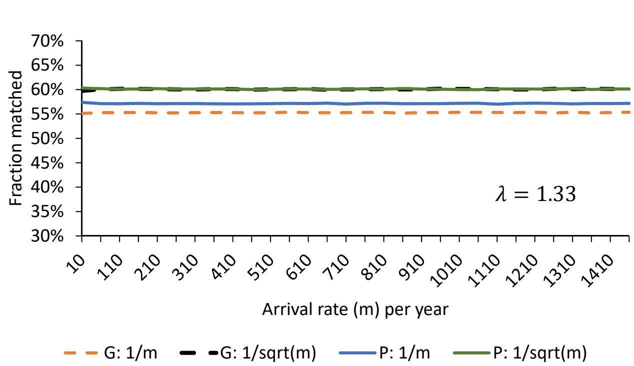

(e) Match rate

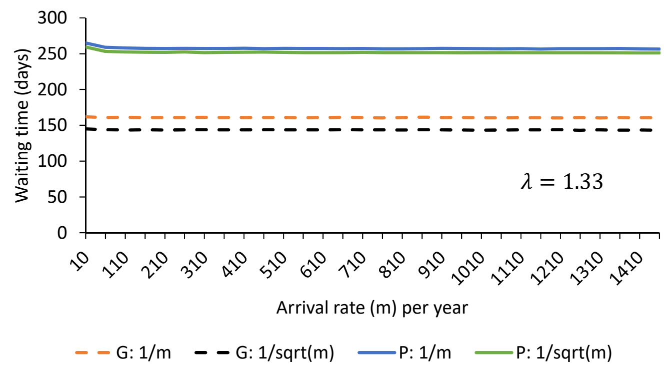

(f) Waiting time

Figure 15: Comparison of greedy (G-M) and patient (P-M) policies for different arrival rates m in a model with vanishing compatibility probabilities where $p \in \{\frac{1}{m}, \frac{1}{\sqrt{m}}\}$, $q = p$, and $\lambda \in \{-0.5, 0, 1.33\}$.

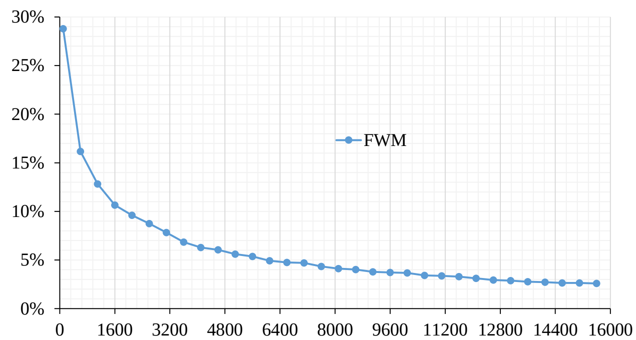

Figure 16: Average percentage of pairs without a compatible partner in a large artificially generated pool from the combined data sets of NKR, APD, UNOS and Methodist at San Antonio.