

Superiority of Instantaneous Decisions in Thin Dynamic Matching Markets

Johannes Bäumlér, Martin Bullinger, Stefan Kober, Donghao Zhu
Technical University of Munich

{johannes.baeumler, martin.bullinger, stefan.kober, donghao.zhu}@tum.de

Abstract

We study a dynamic matching procedure where homogeneous agents arrive at random according to a Poisson process and form edges at random yielding a *sparse* market. Agents leave according to a certain departure distribution and may leave early by forming a pair with a compatible agent. The primary objective is to maximize the number of matched agents. Our main result is to show that a mild condition on the departure distribution suffices to get almost optimal performance of instantaneous matching, despite operating in a *thin* market. We are thus the first to provide a natural condition under which instantaneous decisions are superior in a market that is both sparse and thin. This result is surprising because similar results in the previous literature are based on market *thickness*. In addition, instantaneous matching performs well with respect to further objectives such as minimizing waiting times and avoiding the risk of market congestion. We develop new techniques for proving our results going beyond commonly adopted methods for Markov processes.

1 Introduction

Matching problems are the basis of one of the most vibrant research areas at the intersection of economics, operations research, and computer science with an abundance of applications encompassing labor markets, school choice, child adoption, dating platforms, allocating advertisements, ride-sharing, or kidney exchange. In reality, many matching scenarios have a dynamic flavor in the sense that agents arrive and get matched over time. We contribute to the research on decision making in dynamic matching markets.

The central concern in matching markets is to find a matching of high quality. This leads to the algorithmic challenge of deciding when and how to match agents. In particular, this means a careful timing of performing matches. A repeatedly discussed and observed economic paradigm is that it can be beneficial to defer decisions (such as pairing two agents to form a match) to gain longterm benefits from improved later decisions based on an increased range of possibilities or larger knowledge. In other words, delaying decisions *thickens* a market, and especially in the research on matching markets, there are a lot of theoretical results promoting market thickness [see, e.g., Akbarpour et al., 2020, Baccara et al., 2020, Loertscher et al., 2022, Emek et al., 2016, Shimer and Smith, 2001, Leshno, 2021, and our extensive discussion in the related work]. Often, this means that agents have to wait in the market for some time to make optimal decisions. Still, under certain conditions, it is possible that instantaneous matching algorithms have a good performance [Ünver, 2010, Ashlagi et al., 2019b, 2022]. The reason for this is that in these models, a thick market arises even under quick matching decisions, and therefore, there is no reason to ‘wait and thicken the market’.

The implementation of kidney exchange markets promotes two ways of thickening markets: combining many small markets into a big nationwide or even transnational market and thickening the market by deferring matching decisions. This has led to the creation of kidney exchange

programs like the United Network for Organ Sharing (UNOS) in the USA or the Scandia Transplant Kidney Exchange Program (STEP) encompassing Denmark, Finland, Iceland, Norway, Sweden, and Estonia. A fundamental result by Akbarpour et al. [2020] suggests that, even under a large arrival rate of agents, it is necessary to defer matching decisions to guarantee a near-optimal performance. Our main result shows that this finding is sensitive to the distribution of the departure time of agents. In particular, if agents are guaranteed to stay in the market for any given (uniform) minimum amount of time, then there is no need to wait with matching decisions and instantaneous matching is superior.

Our result is remarkable because we operate in a *sparse* market (with few compatibilities) and obtain a *thin* market (with few agents present compared to the arrival rate) because of instantaneous matching. This is a contrast to all of the previous results cited above and allows us to combine a small loss with other desirable properties.

First, it is disputable whether market thickness is always desirable due to the risk of congestion. In fact, the analysis of real-life data indicates that market thickness may lead to bad outcomes with respect to the obtained size [Li and Netessine, 2020] and quality [Fong, 2020] of matchings. Another type of congestion leading to market inefficiencies is caused by an abundance of matching possibilities [Roth, 2018]. Since we are operating on a sparse market, we also avoid this type of congestion.

Second, while Akbarpour et al. [2020] show that the information gain of the exact *realizations* of departures has high value (if the departure times are exponentially distributed), instantaneous matching does not need such information. Hence, we can additionally avoid both cost and ethical concerns of the collection of departure information [Reese et al., 2015].

Third, apart from avoiding the danger of congestion, instantaneous matching has the advantage that agents do not sojourn in the market unnecessarily. The total waiting time of the agents is the integral of the pool size over time. Hence, we derive a short total waiting time of the agents by proving that the market stays small for most of the time. This stands in contrast to a trade-off between matching quality and waiting time observed in dense markets [Mertikopoulos et al., 2020]. In other words, transitioning to the goal of minimizing loss in a *sparse* market can circumvent this conflict of aims.

At this point, we briefly discuss why we can achieve our result and why it fails in other settings like the original model by Akbarpour et al. [2020]. The reason is that agents' matching opportunities are *sufficiently evenly* distributed. Then, for agents in the market, a guaranteed sojourn time means that they cannot be completely neglected within matching decisions when new agents arrive. Note that such a condition is usually fulfilled in realistic scenarios. Incentives that aim at a minimum sojourn time, like an entering fee or a cancellation fee to exit, make it unreasonable for agents to leave before there was a chance to match them. By contrast, if agents leave the market according to an exponential distribution, then there is a rather high probability that they are neglected because they sojourn in the market for a very short time.

We complement our theoretical results by a series of simulations. Our theoretical findings consider the limit case for growing arrival rates. However, the error terms vanish at rapid rates and the simulations also confirm our results for a rather small arrival rate indicating their robustness. Moreover, we test further specifications of the greedy algorithm. In both of our algorithms, the decision to form a matching is done by selecting a partner uniformly at random among the possible matches. This already yields approximately optimal results as our lower bounds for the performance of any algorithm show. On the other hand, we perform simulations for the case of a tie-breaking in favor of partners that have the smallest remaining sojourn, finding that this still yields a small improvement. Interestingly, our simulations provide evidence that the loss of the greedy mechanism and the patient mechanism are identical under a constant departure time. In a concluding section, we shed light on theoretical reasons for this observation. There, we also discuss extensions of our work to the case of heterogeneous agents.

On the technical side, we develop novel techniques to obtain our results. While most of

the existing literature heavily relies on a steady state analysis of the pool size [in particular Akbarpour et al., 2020, Anderson et al., 2017, Ashlagi et al., 2022], this quantity is not Markovian for general departure distributions, and we do not make use of a stationary distribution inherent to our process. Instead, we perform a detour by a careful direct analysis obtaining uniform bounds over time for the pool size which in turn lead to bounds for the loss. Similarly, our analysis of the patient algorithm demands new techniques and our key idea is a case analysis for the near past of a fixed point in time, performing a comparison of the critical case with an urn model.

1.1 Modeling Matching Markets

The rich literature on matching markets comprises a large number of formal models. These differ in their assumptions on the essence of agents, the matching technology, and procedural specifications as well as regarding the measurement of the desirability of matchings. First, agents may be homogeneous or heterogeneous. The latter is often due to having certain applications in mind. For instance, agents can be partitioned into two classes like workers and firms in labor markets, children and schools in school choice, or children and potential adoptive parents in adoption markets. Moreover, agents might assume various additional properties such as being hard or easy to match, which play for instance an important role in kidney exchange (Ashlagi et al. [2019b], Ashlagi et al. [2022], see also Loertscher et al. [2022] for a more general trade market model with superior and inferior buyers and sellers). Second, the matching technology specifies the possible matchings. For instance, it might be prohibited that agents of the same class match (which is very reasonable for the above examples of *bipartite* instances), or the possible matches could be given endogenously by a compatibility relation (which might even be deduced from an exogenous factor like affiliation to some type). Third, procedural specifications concern the chronological process of the model. Agents might arrive stochastically, in fixed time steps, or even according to an adversary. Then, they stay in the market for some time after which they depart according to some procedure. Some stylistic models even assume an indefinite sojourn of agents until they are matched. Finally, we have to measure the quality of matchings. The two most common approaches considered in the literature are measuring the quality of a matching simply by its cardinality or by maximizing a more complex predefined objective function, which is often defined specifically for the given model and the application in mind. The former is appropriate in models where sparse markets evolve from a compatibility relationship of the agents or in settings such as kidney exchange where losing an agent is severe. The latter often requires a careful consideration of model parameters and may quickly lead to complex measures. Objective functions usually include the quality of matches, but can also consider loss suffered by spending time in the market.

In the present paper, we follow the work by Akbarpour et al. [2020] and Anderson et al. [2017] and consider a matching market with ex ante homogeneous agents, that is, agents are a priori identical but, given a snapshot of the current market, they differ in their departure information and potential matches. An overview of our procedural specifications is given in Figure 1. Agents arrive stochastically according to a Poisson process and stay for a maximum time according to a departure distribution. Our focus lies on the departure information of the agents, which we treat as a variable component of our model. Much of our analysis concerns predicting the behavior of the market given certain fairly general assumptions on the departure information. The matching technology is given by a compatibility relationship based on the Erdős-Rényi random graph model [Erdős and Rényi, 1960], and is determined by independent Bernoulli random variables deciding on the compatibility of arriving agents with existing unmatched agents. We express our results with respect to the arrival rate of compatible agents, that is, the expected number of compatible agents arriving in one time unit. Since we consider a growing arrival rate of agents, while fixing the compatibility, this leads to decreasing compatibility probabilities. Therefore, our considered markets are mostly sparse. As a consequence, we measure the performance of an




Figure 1: Illustration of the matching process. Agents arrive stochastically according to a Poisson process. During their sojourn in the matching market, they can be matched with contemporaneous agents and leave the market. They perish according to a predefined departure distribution. Specific matching policies are highlighted in blue. The greedy algorithm seeks to match an agent instantaneously at arrival, while the patient algorithm waits for the last moment to match.

algorithm by the cardinality of the matching it achieves, i.e., we seek to minimize the obtained *loss*, defined as the fraction of agents departing from the market unmatched.

1.2 Matching Algorithms

Following the approach and terminology by Akbarpour et al. [2020], we focus on two matching procedures assuming the global perspective of a social planner responsible for a centralized execution of the matchings. The first one, called *greedy algorithm*, corresponds to the instantaneous implementation of matchings, whenever agents arrive at the market. In other words, whenever an agent arrives at the market, she either forms an edge with some existing agent and leaves the market right away as part of a pair, or she joins a pool of agents of which no pair can be matched. Once added to the pool, an agent can solely wait to be matched with a newly arriving agent, or until she perishes.

The second matching procedure, called *patient algorithm*, assumes that the exact departure times of agents are known to the social planner. Hence, once the point of departure arrives, there is one last chance for the agent to be matched before perishing. Therefore, the patient algorithm waits to match an agent until the very last moment and only matches upon the danger of losing an agent. Consequently, the only possibility of loss is if an agent at the edge of departure maintains no relationships with other agents in the pool. The main result by Akbarpour et al. [2020] is that the patient algorithm outperforms the greedy algorithm in case of exponentially distributed departure times.

As discussed before, our main result shows an exponentially small loss of the greedy algorithm under certain assumptions on the distribution of the departure time of agents, while agents simultaneously stay in the market for a short time.

By contrast, the patient algorithm causes agents to wait long, because at least half of the agents have to wait for their maximum possible duration to stay. Still, our analysis of the patient algorithm complements the findings by Akbarpour et al. [2020] and shows an exponentially small loss for a constant departure time. Hence, if exact departure points are known and waiting time is secondary, then the patient algorithm also performs very well.

2 Related Literature

There exist several streams of literature dealing with dynamic models of matching markets. Closest to our work are the articles by Anderson et al. [2017] and Akbarpour et al. [2020]. We extend their model and results offering new insights when greedy matching performs well. Anderson et al. [2017] seek to model barter exchange (having in mind the important application of kidney exchange) and the model is essentially a directed version of the model by Akbarpour

et al. [2020]. Arriving agents form directed relationships/edges with probability p , however they arrive in fixed time intervals for an indefinite sojourn, and the matching technology even allows for matching a larger group of agents or chains originating from altruistic donors. The goal is to minimize total waiting time and the main result is that allowing more matching possibilities significantly reduces waiting time. If the matching technology is restricted to pairs, then their model closely resembles our model, where compatibility is interpreted as having both directed edges present, amounting to a compatibility probability of p^2 . The waiting time in their result matches the waiting time that we observe for the greedy algorithm. Therefore, we complement the result by Anderson et al. [2017] for different procedural specifications. Kakimura and Zhu [2021] consider a bipartite version of this model, where the agents of the two partition classes arrive with two different Poisson rates and they also consider the greedy and patient algorithm. Performing a Markov chain analysis, they confirm the superiority of the patient algorithm in a two-sided market, but they observe a similar performance of the two algorithms for a one-sided market where one class of the agents are merely objects.

Much of the literature on dynamic matching markets originates from research in theoretical computer science on the online formation of bipartite matchings [Karp et al., 1990]. There, the primary goal is to maximize the cardinality of the matching formed instantaneously, while procedural specifications are secondary and algorithms are measured by their performance against an adversarial arrival of agents. With the purpose of modeling ad allocation, the model was generalized by Mehta et al. [2007], and a version with stochastic arrival of agents was considered subsequently [Feldman et al., 2009, Manshadi et al., 2012]. An adversarial arrival of agents is also considered by Emek et al. [2016] who consider an objective function combining the compatibility quality of matched agents and their waiting time, and by Ashlagi et al. [2019a] who consider a model with arrivals and departures after constant time.

The paradigm of waiting to create further possibilities, i.e., the desire of market thickness, observed by Akbargpour et al. [2020] is ubiquitous in the literature. This phenomenon often occurs in markets with heterogeneous types and unbounded sojourn of agents, and we identify two main circumstances for its appearance. The first encompasses markets with the concurrent objectives to maximize matching quality while minimizing waiting time, and the second concerns settings, where agents are offered matching opportunities which they can accept or deny, similar to the secretary problem [Ferguson, 1989, Bearden, 2006].

Regarding the former, Baccara et al. [2020] model child adoption by a bipartite matching market with two agent types and identify thresholds to decide when to conduct low-quality matches. In a similar vein, Loertscher et al. [2022] consider a bipartite trade model with a superior and inferior type of buyers and sellers, and determine how many agents to keep in the market. Interestingly, they find that few stored agents lead to a good market performance, but this result has to be interpreted with respect to their market design which requires to match many agents early. Blanchet et al. [2022] consider a bipartite market with exponentially distributed departure times. Their focus is on the distribution of matching values and they identify certain threshold rules which achieve a desired market thickness. In models with adversarial arrival of agents, Emek et al. [2016], Collina et al. [2020], and Pavone et al. [2020] consider the fine-tuned timing of matching decisions.

In the second type of markets, the challenge is to balance information gains through rejecting matches with the danger of turning down a promising match. There, Leshno [2021] studies a queuing model applied to an assignment problem, and finds that it is beneficial to decline mismatches. This can have both positive aspects for the declining agent who might obtain a better object, and for other agents in the market who might get served earlier. Furthermore, a recent paper by Agarwal et al. [2021] considers a matching market with arrivals and departures in discrete time steps and they investigate cutoff strategies for accepting a match. Loosely related, Shimer and Smith [2001] consider a model for labor markets where heterogeneous agents incur a cost for searching a suitable match. They find that agents of high productivity should wait

to improve their matching quality. Also, note that the matching problem in the above paper by Pavone et al. [2020] is equivalent to an adversarial version of the secretary problem.

There is far less evidence for the antagonistic principle of taking instantaneous decisions, with the notable exception of the research on kidney exchange, originating from seminal contributions by Roth et al. [2004, 2005]. In this line of research, agents usually stay indefinitely in the market and the performance is measured by means of the total waiting time. Ünver [2010] anticipates the conclusion by Anderson et al. [2017], namely the optimality of greedy matching. In his model, the role model for Anderson et al. [2017], he considers arrivals by a Poisson process and an indefinite sojourn, but the model contains agents of homogeneous types and the waiting time objective is specifically designed for the given setting. Follow-up work by Ashlagi et al. [2019b] offers a closer look at the case of easy-to-match and hard-to-match agents and their prioritization in greedy-type algorithms. These theoretical findings are in accordance with simulations based on real-world data [Ashlagi et al., 2018]. Ashlagi et al. [2022] show the optimality of greedy matching in a similar setting. All of this work has in common that the reason for the optimality of greedy matching is market *thickness*. By contrast, our work presents the optimality of an instantaneous algorithm measured by the *loss* of agents in a sparse market that is *thin* (for most of the time). Very recently, Kerimov et al. [2022] consider a model with heterogeneous types whose compatibility is globally fixed. They show that, under certain conditions on the compatibility structure, greedy matching is optimal with respect to regret.

Finally, Aveklouris et al. [2021] consider a bipartite model with heterogeneous agents where the objective function is a tradeoff of matching value and waiting costs. Since the model is quite complex model, a Markovian analysis is not feasible. They circumvent this obstacle by finding optimal matching rates by solving an optimization problem, and identify an algorithm that mimics these rates. Interestingly, this paper is the only related paper we are aware of which investigates the influence of the departure distribution.

3 Model

In this section, we present our formal model. First, we introduce our model for a dynamic matching market, and then we discuss our objectives.

3.1 Matching Markets

Following the models by Anderson et al. [2017] and Akbarpour et al. [2020], we consider a continuous-time matching market in the time window $[0, T]$ for some maximum time $T > 0$. Agents arrive with Poisson rate m and enter the market. We label the agents v_1, v_2, \dots in the order of their arrival. Given times $s, t \in [0, T]$ with $s < t$, we denote by A_t the set of *agents arriving at time t* and by $A_{[s, t]}$ and $a_{[s, t]}$ the set and number of agents arriving in the time interval $[s, t]$, respectively. At her arrival, an agent enters the *pool*, i.e., the set of agents in the market. The pool at time t is a random variable Z_t , and we denote its size by z_t , i.e., $z_t = |Z_t|$. We assume that the market is initially empty, i.e., $Z_0 = \emptyset$.

Once an agent enters the pool, we decide on her *compatibility* with agents already in the pool by independent Bernoulli random variables with parameter p . Hence, compatible agents arrive at Poisson rate $d = p \cdot m$, and we call this quantity the *density parameter*. Formally, we model compatibility as follows. We assume that agent v_i carries an independent Bernoulli $(p)^{\otimes \mathbb{N}}$ distributed random variable $U(v_i)$. There, $\otimes \mathbb{N}$ denotes the product space. Let the pool before the arrival of v_i consist of the agents $\{v_{l_1}, \dots, v_{l_K}\}$, where K is the size of the pool just before the arrival of v_i , and $1 \leq l_1 < l_2 < \dots < l_K < i$ are all integer numbers. At arrival, the agent v_i forms an edge with v_{l_j} if and only if $U_j(v_i) = 1$, where $U_j(v_i)$ is the j -th component of $U(v_i)$. So in particular, only the first K many components of $U(v_i)$ are of interest.

While an agent v_i is in the pool, she can be matched with a compatible agent in the pool,

and they leave the pool together. Otherwise, the agent leaves the pool after some independent random *maximum sojourn time* X_i has passed, which is distributed according to some measure μ with support contained in $[0, +\infty]$. The maximum sojourn time is also called *departure time* (but some caution is needed here, as an agent can also ‘depart’ from the market due to being matched). The last point, where an agent joining at time t can be matched is at time $t + X_i$, and we call an agent *critical* at the moment when she reached her last time to be matched. The time an agent spends in the pool is called her *sojourn time* and we say that an agent *perishes* if she leaves the pool unmatched.

Agents are matched by certain *matching algorithms* (or *matching policies*), which decide when two agents present in the pool leave the market as a pair. In principle, multiple pairs could be matched at the same time (as it is often done in batching algorithms, for instance for kidney exchange), and no agent can be part of more than one pair. We make the assumption that agents can only leave the market when they are matched or when they perish because their maximum sojourn time is reached. This avoids the consideration of algorithms who just remove agents from the market unmatched, which yields small waiting times but violates the purpose of matching algorithms.

Two important matching policies are the greedy algorithm and patient algorithm [Akbarpour et al., 2020]. The *greedy algorithm* GDY is defined by the policy that every agent arriving at the pool is matched to a compatible agent in the pool uniformly at random and instantaneously, or joins the pool if there is no compatible agent. On the other hand, the *patient algorithm* PAT is defined by having agents wait until they reach their maximum sojourn time at which they match with a compatible agent uniformly at random, or perish if no such agent exists. More formally, if agent v_i arrives at time t , then she waits until time $t + X_i$, and may only leave the market during this time when she is matched with a compatible critical agent selecting her as a match. If this does not happen, she checks the pool Z_{t+X_i} for compatible agents once she is about to perish, and selects a match if possible. Note that, for every time t , it holds almost surely that at most one agent arrives or perishes at this time, and we can therefore disregard a case analysis for the possibility of multiple simultaneous matching events.

3.2 Objectives

Our main objective is to minimize the loss of agents. Therefore, given an algorithm ALG and an arrival rate m , we consider the number of agents matched until time T , as defined by $\text{ALG}(m, T) = \{v_i : v_i \text{ matched until time } T\}$. Then, the *loss* of ALG until time T is defined by

$$\mathbf{L}_{\text{ALG}}(m, T) = \frac{\mathbb{E}[|A_{[0, T]} - \text{ALG}(m, T) - Z_T|]}{mT}.$$

In this expression, the denominator is equal to the expected number of agents arriving in the time window $[0, T]$. Note that the loss of an algorithm also depends on the density parameter d and the distribution μ of the departure times. For the sake of notational simplicity, we exclude these from our notation, but add the departure distribution as a subscript of probabilities or expectations in danger of ambiguity. We are interested in the limiting behavior for m and T tending to infinity. Therefore, we define

$$\mathbf{L}_{\text{ALG}} = \limsup_{m, T \rightarrow \infty} \mathbf{L}_{\text{ALG}}(m, T).$$

Note that considering the limiting behavior is mostly cosmetic. Error terms depending on m and T decay quickly in our theoretical analysis, and our simulations show the robustness of our results for reasonably small arrival rates (see Figure 2).¹

¹For instance, the error terms in the analysis of the greedy algorithm in Lemma 5.2, Corollary 5.3, and Lemma 5.4 are m^{-9} .

Our next objective is the waiting time of agents, i.e., the time that agents spend in the pool. For a given algorithm ALG, the *total waiting time* until time T is defined by

$$\mathbf{W}_{\text{ALG}}(m, T) = \mathbb{E} \left[\sum_{i \geq 1} \int_0^T \mathbf{1}_{v_i \in Z_s} ds \right].$$

There, $\mathbf{1}$ denotes the indicator function. Also, as in the definition of the loss, we suppress the dependence from the distribution of departure times. In other words, $\mathbf{W}_{\text{ALG}}(m, T)$ is the sum over all agents of the time spent in the pool until time T . Note that the total waiting time implicitly depends on m because the arrival rate influences the set of agents that has arrived until time T .

Our third objective is market congestion. As we have discussed before, congestion can be caused both by a thick market as well as by a dense market. Since we consider a fixed density parameter for a growing arrival rate, our markets are thin by design. Hence, we define the *congestion* of an algorithm as the average expected pool size as

$$\mathbf{C}_{\text{ALG}}(m, T) = \frac{1}{T} \int_0^T \mathbb{E}[z_s] ds$$

Once again, the expectation depends on the distribution of the departure times. As we will see, there exists a clear relationship between waiting time and congestion.

4 Superiority of the Greedy Algorithm

In this section, we discuss our results for the greedy algorithm. All technical proofs can be found in Section 5.

4.1 Loss Minimization

We start with our main objective, namely minimizing loss. We will prove the following theorem.

Theorem 4.1. *Assume that the departure time is distributed according to a probability measure μ with $\mu([0, 1)) = 0$. For $d \geq 2$, it holds that*

$$\mathbf{L}_{\text{GDY}} \leq e^{-\frac{d}{2 \ln(2)}}.$$

Notably, the assumption in Theorem 4.1 on a minimum departure time of one time unit is somewhat arbitrary. Indeed, for obtaining an exponentially small loss in d , we merely have to assume a uniform lower bound on the minimum departure time of an agent. However, transitioning to more general departure times also affects the constant appearing in the exponent of the loss. By rescaling the departure times, we obtain a scaled variant of the previous theorem.

Corollary 4.2. *Let $\varepsilon > 0$ and assume that the departure time is distributed according to a probability measure μ with $\mu([0, \varepsilon)) = 0$. For $d \geq 2$, we have*

$$\mathbf{L}_{\text{GDY}} \leq e^{-\frac{\varepsilon d}{2 \ln(2)}}.$$

Our analysis of the greedy algorithm is close to optimal. Up to a constant in the exponent, it cannot guarantee significantly better outputs. We state our next theorem for the case of a constant departure time, and restrict attention to the case $\mu(\{1\}) = 1$.

Proposition 4.3. *Assume that the departure time is given by the degenerate measure μ with $\mu(\{1\}) = 1$. Then,*

$$\mathbf{L}_{GDY} \geq \frac{1}{2}e^{-d}.$$

On the other hand, the inferiority of the greedy algorithm for exponentially distributed departure times observed by Akbarpour et al. [2020] really depends on a non-negligible chance of staying in the market for a very small time. We generalize the respective result by Akbarpour et al. [2020].

Proposition 4.4. *Let $c > 0$, $\varepsilon_0 > 0$, and assume that the departure time is distributed according to a probability measure μ with $\mu([0, \varepsilon]) \geq c\varepsilon$ for all $\varepsilon \leq \varepsilon_0$. Then, for d large enough, it holds that*

$$\mathbf{L}_{GDY} \geq \frac{c}{6d}.$$

Indeed, the greedy algorithm is even close to optimal if compared to any other algorithm. Hence, even having knowledge of all future agents and their compatibilities cannot lead to a large improvement with respect to loss compared to the greedy algorithm.

Proposition 4.5. *Assume that the departure time is given by a distribution μ with $\int x\mu(dx) = E$, i.e., E is the expected maximum sojourn time of an agent. Then, for any algorithm ALG, it holds that*

$$\mathbf{L}_{ALG} \geq e^{-2Ed}.$$

In particular, if the departure time is given by the degenerate measure μ with $\mu(\{1\}) = 1$, then, for any algorithm ALG, it holds that

$$\mathbf{L}_{ALG} \geq e^{-2d}.$$

Hence, by Theorem 4.1, the loss generated by the greedy algorithm in this case can be at most a factor of $4 \ln(2) \approx 2.77$ in the exponent worse compared to the loss of an optimal algorithm.

4.2 Comparison to Patient Algorithm

The fundamental insight by Akbarpour et al. [2020] is that the patient algorithm outperforms the greedy algorithm in terms of loss if the departure times are exponentially distributed. Even though the greedy algorithm performs very well under other departure distributions, the patient algorithm still is an alternative worth considering if access to the departure information is not costly. Here, we show that the patient algorithm still yields exponentially small loss for the case of agents with a maximum sojourn time of exactly one time unit. The central idea of the proof is to use the intuition that the probability of being matched is high if the pool size is high. Therefore, the key step is to show that the pool size is moderately high with very high probability. To this end, we perform a case analysis of the near past of a fixed point in time, considering four time intervals of length $1/3$. The critical insight is that few agents arriving during the first time interval can be present at the beginning of the last time interval. To show this, we perform a comparison with an urn model.

Theorem 4.6. *Assume that the departure time is distributed according to the degenerate probability measure μ with $\mu(\{1\}) = 1$. Then, it holds that*

$$\mathbf{L}_{PAT} \leq e^{-\frac{d}{5}}.$$

4.3 Further Objectives

For the patient algorithm, at least half of the agents stay for their maximum possible sojourn. Hence, under unit waiting times and for arrivals until time $T - 1$, the expected waiting time is at least $1/2$, a fixed constant. Even if we consider an arbitrary distribution of the maximum sojourn time, we obtain a constant average waiting time. The idea of the proof of this statement follows a similar reasoning. We divide the set of agents in a large set of agents whose maximum sojourn time is high and a small set of agents with low maximum sojourn time. At least half the agents of the former type which do not match with an agent of the latter type have to wait a long time.

Proposition 4.7. *Let $h > 0$ and assume that the departure time is distributed according to a probability measure μ with $\mu([h, \infty]) > \frac{4}{5}$. Then, for sufficiently large m and T , it holds that*

$$\mathbf{W}_{PAT}(m, T) \geq \frac{mTh}{4}. \quad (1)$$

The situation is different for the greedy algorithm. Approximately half of the agents get matched at arrival, but the waiting time of agents added to the pool is unclear. We will show that, in the greedy algorithm under an arbitrary departure time, $\mathbf{W}_{GDY}(m, T) \in \Theta\left(\frac{mT}{d}\right)$, i.e., the average waiting time is of order $\Theta\left(\frac{1}{d}\right)$.

Interestingly, the expected total waiting time can be expressed by an integral over the expected pool size, which is equivalent to the congestion. This can be seen as a variant of Little's law in our context.

Proposition 4.8. *Assume that the departure time is distributed according to an arbitrary probability measure μ . Then, for any algorithm ALG , it holds that*

$$\mathbf{W}_{ALG}(m, T) = \int_0^T \mathbb{E}_\mu[z_s] ds = T \cdot \mathbf{C}_{ALG}(m, T).$$

Hence, algorithms perform well with respect to congestion if and only if they perform well with respect to total waiting time. We can use this relationship to give an upper bound on the total waiting time.

Proposition 4.9. *Assume that the departure time is distributed according to an arbitrary probability measure μ . Then, for large enough m , under the greedy algorithm, it holds that*

$$\mathbf{W}_{GDY}(m, T) \leq \frac{6mT}{5d}.$$

We also prove a lower bound on the total waiting time. However, this lower bound needs minor additional assumptions on the measure of the departure times. Even if an algorithm has access to the whole information of incoming agents, compatibilities, and maximum sojourn times, it has a waiting time that is only a constant factor apart from the waiting time of the greedy algorithm. It is important for this result that agents can only leave the market by forming a matching or by perishing. So in particular, the algorithm cannot remove unmatched agents from the market. Without this restriction, the waiting time could of course be arbitrary small, by the cost of having a high proportion of unmatched agents.

Proposition 4.10. *Assume that the departure time is distributed according to a probability measure μ with $\mu\left(\left[\frac{1}{2d}, \infty\right]\right) > \frac{9}{10}$. Further, assume that $T \geq \frac{3}{d}$. Then, under any algorithm ALG , it holds that*

$$\mathbf{W}_{ALG}(m, T) \geq \frac{mT}{20d}.$$

4.4 Market Thickness

As we have already mentioned in the introduction, a feature that makes Theorem 4.1 special is that we operate on a thin market. The previous section allows us to quantify more precisely what this means in our model. First, the proof of Proposition 4.9 shows that the market under the greedy algorithm is of size $\mathcal{O}(\frac{m}{d})$ at most times for every distribution of the departure times. By contrast, the market size under the patient algorithm is of order $\Theta(m)$ most of the time under exponentially distributed departure times [Akbarpour et al., 2020] and the same result also holds for any guaranteed sojourn time (because at least half of the agents have to stay until their maximum sojourn times). Hence, the greedy algorithm maintains a market size that is smaller by a factor of $\frac{1}{d}$ compared to the arrival rate. In fact, under the mild assumptions of Proposition 4.10, every matching algorithm has to maintain a market size of $\Omega(\frac{m}{d})$ most of the time (this follows from the same arguments as in the proof of Proposition 4.10).

5 Obtaining Loss Bounds for the Greedy Algorithm

In this section, we discuss key steps in the proof of the bounds of the greedy algorithm.

5.1 Outline of the Analysis

The main problem in the analysis of any other model than the one with exponentially distributed departure times is the lack of Markovianity of the pool size. Knowing at which exact time points the agents in the current pool arrived, already reveals some information at which time points in the future we expect more or less agents to get critical and perish. So the knowledge of the entire history $(z_t)_{t \leq t_0}$ typically reveals more information than the pool size z_{t_0} , which shows that $(z_t)_{t \geq 0}$ is not a Markov chain in general. Hence, we cannot simply analyze a stationary distribution in order to control the pool size, but have to perform a direct analysis of the pool size to be able to investigate the loss of the greedy algorithm.

Determining an upper bound for the pool size follows from the following idea: Assume that the pool has a size of Cm for some constant $C > \frac{\ln(2)}{d}$ and an agent arrives. Then, by Lemma C.1, there is a probability of

$$\left(1 - \frac{d}{m}\right)^{Cm} \leq e^{-Cd} < \frac{1}{2}$$

that the new agent joins the pool causing the pool size to increase by 1. On the other hand, there is a probability of strictly more than $\frac{1}{2}$ that the agent matches and thus the pool size decreases by 1. Consequently, for a sufficiently large pool size, the arrival of new agents decreases the pool size in average. Besides the arrival of new agents, existing agents in the pool may also perish and cause a further decrease of the pool size. This analysis shows that we do not expect the pool size to be much larger than $\frac{\ln(2)}{d}m$ for most of the time. This key step is shown in Corollary 5.3 and Lemma 5.4. An important technique for the former result is to compare an arbitrary departure time with the case when agents do not perish at all, i.e., with the case of an infinite departure time. For this, we will create a coupling of the respective markets.

Given a uniform bound on the pool size, we can also bound the loss under various assumptions on the departure time. For instance, assume a constant departure after 1 time unit. Further, suppose that a new agent, called v , enters the pool at time t without being matched right away and that the pool size is bounded from above by $\frac{\ln(2)}{d}m$ for all times in $[t, t+1]$. This agent will stay for (at most) one time unit and each agent that arrives during this time has a probability of approximately (up to constant multipliers) $\frac{d}{m}$ to match with v . Hence, the probability that v perishes without being matched by another arriving agent is bounded by an expression of the form $(1 - \Theta(\frac{d}{m}))^m$ which explains the exponentially small loss under a constant departure time. By generalizing these arguments, we provide lower and upper bounds on the loss for different

distributions of departure times. We find quite general conditions that distinguish cases when we are guaranteed either an exponentially small or an inverse linear loss.

5.2 Comparison of Departure Distributions

In this section, we compare the pool sizes under the Greedy algorithm for different distributions of departure times and under different initial conditions.

In the first lemma, we use a coupling method to compare the evolution of the pool size under any departure time with the case of an infinite departure time. Therefore, let μ_∞ be the degenerate measure with $\mu_\infty(\{+\infty\}) = 1$. Given a time t , denote by Z_t^∞ and z_t^∞ the pool and the pool size for departure times according to μ_∞ , respectively. Note that our coupling lemma leaves a lot of freedom for the starting pool, and does for instance not require empty pools in the beginning.

Lemma 5.1. *Let μ be any probability measure on $[0, +\infty]$. Consider the pools Z_t and Z_t^∞ at time t for the greedy algorithm with departure times distributed according to μ and μ_∞ , respectively, and let Z_0 and Z_0^∞ be two starting pools with $z_0 \leq z_0^\infty$. Then, there exists a coupling of Z_t and Z_t^∞ such that*

$$z_t \leq z_t^\infty + 1 \quad (2)$$

for all $t \geq 0$.

In particular, Lemma 5.1 implies that when running the greedy algorithm without perishing (infinite departure time) and one starts with two different initial pool sizes, the process that started with the lower pool size, can surpass the other one by at most 1.

A key feature of the measure μ_∞ is that the pool size actually is Markovian under this measure, as the evolution of the pool size depends only on the current size of the pool and the arriving agents. The transition rates $(r_{k \rightarrow k+1})_{k \in \mathbb{N}_{\geq 0}}$ and $(r_{k \rightarrow k-1})_{k \in \mathbb{N}_{\geq 1}}$ are given by

$$r_{k \rightarrow k+1} = m \left(1 - \frac{d}{m}\right)^k, \quad r_{k \rightarrow k-1} = m \left[1 - \left(1 - \frac{d}{m}\right)^k\right]. \quad (3)$$

Interestingly, we have that $r_{0 \rightarrow 1} = m$ and, for all $k \geq 1$, that $r_{k \rightarrow k+1} + r_{k \rightarrow k-1} = m$. Thus, if the pool size is larger than 0, then the process waits an exponential time with expectation $1/m$ and then jumps to $k+1$ or $k-1$ with corresponding probabilities. At 0, the process also waits for an exponential time with expectation $1/m$ and then jumps $+1$. In particular, the transition probabilities $(p(k, l))_{k, l \in \mathbb{N}_{\geq 0}}$ of the underlying discrete-time Markov chain are simply given by

$$p(k, k+1) = \left(1 - \frac{d}{m}\right)^k \quad \text{and} \quad p(k, k-1) = 1 - \left(1 - \frac{d}{m}\right)^k \quad (4)$$

for $k \in \mathbb{N}_{\geq 1}$, $p(0, 1) = 1$, and all other probabilities are 0. These transition probabilities give rise to an irreducible and reversible Markov chain (every Markov chain with state space $\mathbb{N}_{\geq 0}$ making jumps between nearest neighbors only is reversible). We claim that there exists a stationary distribution ρ for this Markov chain. Indeed, one can construct a stationary measure with the aid of the measure $\tilde{\rho}$ defined by $\tilde{\rho}(0) = 1$ and $\tilde{\rho}(k+1) = \rho(k) \frac{p(k, k+1)}{p(k+1, k)} = \prod_{i=0}^k \frac{p(i, i+1)}{p(i+1, i)}$. This is *not* a probability measure, yet. However, the measure $\tilde{\rho}$ can be normalized as $\sum_{k=0}^{\infty} \prod_{i=0}^k \frac{p(i, i+1)}{p(i+1, i)} < \infty$, which follows because $p(i, i+1) < 1/3$ for all large enough i .

The existence of a stationary distribution tells us that the pool size does not go off to infinity as time grows. The balance equation of the stationary distribution for $j \in \mathbb{N}_{\geq 0}$ reads

$$\rho(j)p(j, j+1) = \rho(j+1)p(j+1, j), \quad (5)$$

and we will give bounds on the stationary distribution in Lemma 5.2 below.

5.3 Pool Size Bounds for the Greedy Algorithm

In this section, we provide bounds for the pool size of the greedy algorithm. The key idea is to distinguish different ranges of the pool size constraint by carefully chosen constants. For this, we define the constants $C_1 = C_1(m)$, $C_2 = C_2(m)$, and $C_3 = C_3(m)$ by

$$C_1 = 1 + \frac{10}{\ln(2)\ln(m)}, \quad C_2 = C_1 + 2\frac{d\ln(m)^2}{m\ln(2)}, \quad \text{and} \quad C_3 = C_1 + 4\frac{d\ln(m)^2}{m\ln(2)}.$$

First, we show how to control the stationary measure under infinite departure times. As a consequence, we obtain bounds for the pool size under an arbitrary departure distribution.

Lemma 5.2. *Let ρ be the stationary measure of the pool size without perishing. Then,*

$$\rho\left(\left(\frac{C_1\ln(2)m}{d} + \frac{3}{2}\ln(m)^2, \infty\right) \cap \mathbb{N}\right) \leq m^{-9}$$

for m large enough.

Combining the bounds under infinite departures with our coupling result, we can bound the pool size of the greedy algorithm under an arbitrary departure time.

Corollary 5.3. *Assume that the departure time is distributed according to an arbitrary probability measure μ . Then, it holds for all times t and arrival rates m large enough that*

$$\mathbb{P}_\mu\left(z_t > \frac{C_2\ln(2)}{d}m\right) \leq m^{-9}.$$

Proof. Assume that Z_0 is the empty pool and Z_0^∞ has size distributed according to the stationary distribution ρ . Then $z_0 \leq z_0^\infty$, so we can couple the processes according to Lemma 5.1 such that $z_t \leq z_t^\infty + 1$. The pool size z_t^∞ is still distributed according to ρ by stationarity, so we get that for large enough m

$$\mathbb{P}_\mu\left(z_t > \frac{C_2\ln(2)}{d}m\right) \leq \mathbb{P}_{\mu_\infty}\left(z_t^\infty > \frac{C_1\ln(2)}{d}m + \frac{3}{2}\ln(m)^2\right) \leq m^{-9}.$$

There, the first inequality follows from Lemma 5.1 and the second inequality from Lemma 5.2. \square

Our second key insight is that if the pool size is sufficiently small, then the probability of reaching a much larger pool size within one time unit is very low.

Lemma 5.4. *Assume that the departure time is distributed according to an arbitrary probability measure μ . Further, assume that Z is a pool of size at most $C_2\frac{\ln(2)}{d}m$. Then, for large enough m , it holds that*

$$\mathbb{P}_\mu\left(z_t \leq C_3\frac{\ln(2)m}{d} \text{ for all } t \leq 1 \mid Z_0 = Z\right) \geq 1 - m^{-9}.$$

5.4 Loss Bounds for the Greedy Algorithm

Before we leverage the bounds on the pool size to obtain bounds on the loss, we provide a useful lemma that lets us express the loss as an integral over the probability of perishing. This lemma holds for arbitrary matching algorithms and will be applied again for the patient algorithm. Its proof follows by an application of Mecke's equation.

Lemma 5.5. *Let a matching algorithm ALG be given and assume that departure time is distributed according to an arbitrary probability μ . Then,*

$$\mathbf{L}_{ALG}(m, T) = \frac{1}{T} \int_0^T \mathbb{P}_\mu(\text{an agent arriving at time } t \text{ perishes before time } T) dt,$$

and

$$\mathbf{L}_{ALG} = \limsup_{m, T \rightarrow \infty} \frac{1}{T} \int_0^T \mathbb{P}_\mu(\text{an agent arriving at time } t \text{ perishes}) dt.$$

Finally, we can apply our insights to bound the loss of the greedy algorithm. We start with an assumption on the measure that particularly covers the case of constant unit waiting times. We want to remark that all obtained loss bounds are also true for ex ante heterogeneous agents which have departure times with respect to different probability measures, all satisfying the respective assumptions.

Theorem 4.1. *Assume that the departure time is distributed according to a probability measure μ with $\mu([0, 1)) = 0$. For $d \geq 2$, it holds that*

$$\mathbf{L}_{GDY} \leq e^{-\frac{d}{2 \ln(2)}}.$$

Proof. By Lemma 5.5, we have

$$\mathbf{L}_{GDY} = \limsup_{m, T \rightarrow \infty} \frac{1}{T} \int_0^T \mathbb{P}_\mu(\text{an agent arriving at time } t \text{ perishes}) dt. \quad (6)$$

Now we define a certain good event \mathcal{G}_t on which the probability that an agent perishes is relatively low. Let

$$\begin{aligned} \mathcal{G}_t = & \left\{ z_t \leq \frac{C_2 \ln(2)}{d} m \right\} \\ & \cap \left\{ z_s \leq \frac{C_3 \ln(2)}{d} m \forall s \in [t, t+1] \right\} \cap \left\{ a_{[t, t+1]} \geq m - 30 \ln(m) m^{1/2} \right\}. \end{aligned}$$

We first want to bound the probability of the complement \mathcal{G}_t^c . By a union bound, we get

$$\begin{aligned} \mathbb{P}_\mu(\mathcal{G}_t^c) \leq & \mathbb{P}_\mu \left(z_t > \frac{C_2 \ln(2)}{d} m \right) + \mathbb{P}_\mu \left(a_{[t, t+1]} < m - 30 \ln(m) m^{1/2} \right) \\ & + \mathbb{P}_\mu \left(\left\{ z_s \leq \frac{C_3 \ln(2)}{d} m \text{ for all } s \in [t, t+1] \right\}^c \cap \left\{ z_t \leq \frac{C_2 \ln(2)}{d} m \right\} \right). \end{aligned}$$

The probability of the first summand is bounded by m^{-9} by Corollary 5.3. For the second summand, note that $a_{[t, t+1]}$ is a Poisson random variable with parameter m and thus we can use Lemma C.2 obtaining

$$\begin{aligned} \mathbb{P}_\mu \left(a_{[t, t+1]} < m - 30 \ln(m) m^{1/2} \right) &= \mathbb{P}_\mu \left(a_{[t, t+1]} < m \left(1 - \frac{30 \ln(m)}{m^{1/2}} \right) \right) \\ &\leq e^{-\frac{m}{3} \left(\frac{30 \ln(m)}{m^{1/2}} \right)^2} \leq m^{-10}. \end{aligned}$$

There, the last two inequalities hold for m large enough. The last summand is bounded by

$$\mathbb{P}_\mu \left(z_s > \frac{C_3 \ln(2)}{d} m \text{ for some } s \in [t, t+1] \mid z_t \leq \frac{C_2 \ln(2)}{d} m \right) \leq m^{-9}$$

by Lemma 5.4. Together, this shows that $\mathbb{P}_\mu(\mathcal{G}_t^c) \leq m^{-8}$ for large enough m .

Now assume that the event \mathcal{G}_t holds and an agent a arrives at time t . We would like to derive the probability that a does neither get matched directly nor between time t and $t+1$. Whenever a new agent v arrives at time $s \in (t, t+1)$ and at a pool of size z_s , then the probability that v matches with a equals

$$\frac{1}{z_s} \left[1 - \left(1 - \frac{d}{m} \right)^{z_s} \right].$$

Indeed, there is a probability of $1 - \left(1 - \frac{d}{m} \right)^{z_s}$ that the agent a matches at arrival, and all vertices in the pool have the same probability of getting matched. The function $x \mapsto \frac{1}{x} \left[1 - \left(1 - \frac{d}{m} \right)^x \right]$ is decreasing for $x > 0$ which can be seen as follows. Write $b = 1 - \frac{d}{m}$. The derivative of this function is given by $\frac{b^x - x b^x \ln(b) - 1}{x^2}$. As $1 - x \ln(b) \leq e^{-x \ln(b)} = b^{-x}$ we already have that

$$\frac{b^x - x b^x \ln(b) - 1}{x^2} = \frac{b^x}{x^2} (1 - x \ln(b) - b^{-x}) \leq 0$$

which shows that the function is decreasing. So on the event \mathcal{G}_t we have

$$\frac{1}{z_s} \left[1 - \left(1 - \frac{d}{m} \right)^{z_s} \right] \geq \frac{d}{C_3 \ln(2)m} \left[1 - \left(1 - \frac{d}{m} \right)^{\frac{C_3 \ln(2)m}{d}} \right]$$

for all $s \in [t, t+1]$. As $C_3 > 1$ we have that

$$\left(1 - \frac{d}{m} \right)^{\frac{C_3 \ln(2)m}{d}} \leq \left(1 - \frac{d}{m} \right)^{\frac{\ln(2)m}{d}} \stackrel{\text{Lemma C.1}}{\leq} \frac{1}{2}$$

and thus the probability that the agent v matches with a is at least

$$\frac{1}{z_s} \left[1 - \left(1 - \frac{d}{m} \right)^{z_s} \right] \geq \frac{d}{C_3 \ln(2)m} \left[1 - \left(1 - \frac{d}{m} \right)^{\frac{C_3 \ln(2)m}{d}} \right] \geq \frac{d}{2C_3 \ln(2)m}.$$

Using $a_{[t, t+1]} \geq m - 30 \ln(m) m^{1/2}$ and the assumption on the support of μ , implying that an agent can only perish after a sojourn of at least one time unit, we can further bound the probability that the agent a perishes from above by

$$\left(1 - \frac{d}{2C_3 \ln(2)m} \right)^{m - 30 \ln(m) m^{1/2}} \xrightarrow{m \rightarrow \infty} e^{-\frac{d}{2 \ln(2)}}.$$

Consider the event $\mathcal{A}_t = \{\text{an agent arriving at time } t \text{ perishes}\}$. To extend the former inequality without the restriction of \mathcal{G}_t , we compute

$$\begin{aligned} \mathbb{P}_\mu(\mathcal{A}_t) &= \mathbb{P}_\mu(\mathcal{A}_t \mid \mathcal{G}_t) \mathbb{P}_\mu(\mathcal{G}_t) + \mathbb{P}_\mu(\mathcal{A}_t \mid \mathcal{G}_t^c) \mathbb{P}_\mu(\mathcal{G}_t^c) \\ &\leq \mathbb{P}_\mu(\mathcal{A}_t \mid \mathcal{G}_t) + \mathbb{P}_\mu(\mathcal{G}_t^c) \leq \mathbb{P}_\mu(\mathcal{A}_t \mid \mathcal{G}_t) + m^{-8} \xrightarrow{m \rightarrow \infty} e^{-\frac{d}{2 \ln(2)}}. \end{aligned}$$

Inserting this into (6) finishes the proof. \square

6 Simulations

In the previous sections, we have provided strong theoretical bounds on the loss of the greedy and patient algorithm, particularly for the case of constant departure times. These bounds are

guaranteed to hold for m and T tending to infinity and give an indication on the market behavior for fixed m or T . In order to complete the picture, we provide simulated results for different input values. Note that matching markets are often of large scale in reality, see, e.g., Ünver [2010] for a discussion of kidney exchange markets and Baccara et al. [2014] for a discussion of child adoption markets. Their data on the poolsize (greater than 50 000) is confirmed by recent numbers from the Organ Procurement and Transplantation Network and the US Census 2020, respectively. We investigate the market behavior with respect to key characteristics like the percentage of unmatched agents or the average waiting time of agents for realistically sized and randomly generated examples. In particular, this analysis gives insights on the significance of error terms in practical examples.

In order to simulate the market behavior, we start with an empty pool at time 0, i.e., no agents are present initially. Then, agents join the pool at Poisson rate m . This means in particular that the difference between the arrival time of two successively arriving agents is distributed with respect to an exponential distribution with expectation $\frac{1}{m}$, i.e., the arrival of the next agent can always be simulated by evaluating a random variable. Every agent is assigned a maximum sojourn time at her arrival to the market according to a predetermined distribution. During the simulation, a matching algorithm is employed and used to determine pairs of agents from the pool to be matched.

Our focus lies on the greedy algorithm with constant unit departure times, for which we evaluate the average proportion of unmatched agents and the waiting time. The results of our simulations are depicted in Figure 2. First, we consider the average proportion of unmatched agents for the two algorithms in Figure 2a. The simulations indicate an exponentially small loss with respect to d for both algorithms, which is in line with the bounds obtained in our theoretical findings. Interestingly, the average proportion of unmatched agents in both algorithms appears to be surprisingly similar, hinting at a possible connection between two basic algorithms using neither structural information of the compatibility relation nor timing information beyond arrivals and critical states of agents. In Figure 2b, we display a detailed analysis of the average proportion of unmatched agents of the greedy algorithm under small density parameters by means of box-and-whisker plots.² We can see that there is little spread, indicating the robustness of the theoretically proved statements for very sparse markets.


Instead of choosing a match uniformly at random from the set of compatible agents, it is also possible to use further structural or timing information in order to perform a selection. If the social planner receives additional information on the maximum remaining sojourn time of individual agents, one natural matching choice is to select the compatible agent with the least remaining sojourn as a partner. This may avoid the perishing of an agent with small remaining sojourn time. In fact, this algorithm performs close to our guaranteed lower bound for the loss of greedy algorithms as can be seen in Figure 2c. Note that this lower bound does not depend on the choice of the match, but even holds for an arbitrary matching policy.

Finally, we investigate the average waiting time of agents for the greedy algorithm in Figure 2d. Note that the upper bound on the waiting time in Proposition 4.9 already holds for relatively small arrival rates, approximately at $m \approx 125$. Our simulations show that the aggregated waiting time is consistently below the expected upper bound, again demonstrating the superiority of the greedy algorithm over the patient algorithm with respect to waiting time.


7 Discussion and Conclusion

We have analyzed instantaneous matching as captured by the greedy algorithm and matching at the exact departure times as captured by the patient algorithm. We found out that the


²The whiskers denote the minimum and maximum of the average proportion of unmatched agents, respectively, while the lower and upper bound of the box represent the first and third quartile. Finally, the middle line corresponds to the median.




(a) Logarithmic plot of average proportion of unmatched agents for GDY and PAT.



(b) Box-and-whisker plots describing average proportion of unmatched agents for GDY under small density parameter.



(c) Logarithmic plot of average proportion of unmatched agents for GDY under different match selection.



(d) Inverse linear plot of average waiting time for GDY.

Figure 2: Visualization of results of simulations. Each datapoint in the first, third, and fourth figure corresponds to 10 runs for $m = 1000$ and $T = 100$. Each box-and-whisker plot in the second figure represents 200 runs for $m = 500$ and $T = 100$. Note that the linear behavior in the logarithmic plots of the first and third figure mean that the decay is exponential. For reference, we add the theoretical upper bounds and general and specific lower bounds obtained in Section 4.

Table 1: Bounds for the loss of GDY and PAT under different departure times. The bounds for unit departure times are novel, the bounds for exponential departure times are due to Akbarpour et al. [2020, Theorem 4]. The loss of GDY is exponentially small for unit departure times.

ALG	unit departure time		exponential departure time	
	lower bound	upper bound	lower bound	upper bound
GDY	$\frac{1}{2}e^{-d}$	$e^{-\frac{d}{2\ln(2)}}$	$\frac{1}{2d+1}$	$\frac{\ln 2}{d}$
PAT	e^{-2d}	$e^{-\frac{d}{5}}$	$\frac{e^{-d}}{d+1}$	$\frac{1}{2}e^{-\frac{d}{2}}$

performance of the greedy algorithm strongly depends on the distribution of the departure times. In Table 1, we display loss bounds of the greedy and patient algorithm. We showcase our results for the important special case of unit waiting times, and compare our bounds with the bounds for exponentially distributed waiting times considered by Akbarpour et al. [2020]. In particular, if agents are guaranteed to stay in the market for a minimum amount of time, then the greedy algorithm achieves an exponentially small loss. This is a striking result because we operate in a thin market, whereas previous theoretical results proving the optimality of matching algorithms heavily relied on market thickness (see Section 4.4 for a discussion of market thickness in our model). We thus offer a new type of justification to promote greedy algorithms in dynamic matching markets.

To conclude the paper, we elaborate on two important aspects of our work that give prospect to further investigation. First, we consider the loss equivalence of the greedy and patient algorithm under unit waiting times as observed in our simulations. Second, we discuss possible extensions of our work to the case of heterogeneous agents.

7.1 Equivalence of Greedy and Patient Algorithm

While our theoretical results only show an exponential loss in both cases, Figure 2a suggests that the loss of the two algorithms is (approximately) identical. Indeed, there is some intuition why the loss should be approximately $\frac{1}{2}e^{-\frac{d}{2\ln(2)}}$ for both the greedy and the patient algorithm. In this section, we provide two different, non-rigorous arguments which lead to exactly the loss observed in our simulations in both cases. Unfortunately, these arguments are heavily based on steady-state analysis and are therefore hard to be made precise for *non*-Markovian processes. Identifying a relationship of the two algorithms that establishes an identical loss is an intriguing problem for further research.

Greedy algorithm We start by reasoning about the pool size similar to the proof sketch in Section 5.1. Let z_{Eq} denote the pool size at equilibrium. Assuming that the loss is negligible, half of the agents instantaneously match at arrival. The probability of not forming an edge at arrival is about

$$\left(1 - \frac{d}{m}\right)^{z_{Eq}} \approx e^{-d\frac{z_{Eq}}{m}} \stackrel{!}{=} \frac{1}{2}$$

which suggests that the pool in equilibrium should be of size $\frac{\ln(2)}{d}m$. If an agent v does form an edge at arrival, but joins the pool, approximately m agents will arrive during her maximal sojourn time and approximately $\frac{m}{2}$ of these agents will match at arrival. The probability that a specific agent, matching at arrival, matches with v equals $\frac{1}{z_{Eq}}$, by symmetry. So assuming that

z_{Eq} stays approximately constant over time, the probability that v perishes can be estimated as

$$\left(1 - \frac{1}{z_{Eq}}\right)^{\frac{m}{2}} \approx \left(1 - \frac{d}{m \ln(2)}\right)^{\frac{m}{2}} \approx e^{-\frac{d}{2 \ln(2)}}.$$

We conditioned on v not matching at arrival, but joining the pool instead. This happens for approximately half of all agents. Thus, the total loss equals approximately $\frac{1}{2}e^{-\frac{d}{2 \ln(2)}}$.

Patient algorithm Again, we first want to control the pool size. Assuming a negligible loss, for a fixed agent v , approximately $\frac{m}{2}$ agents get critical in the interval $[t_v, t_v + 1]$, where t_v is the arrival time of v . Under the assumption that the pool size stays constant over time, by symmetry, the probability that none of these agents matches with v should equal

$$\left(1 - \frac{1}{z_{Eq}}\right)^{\frac{m}{2}} \approx e^{-\frac{m}{2z_{Eq}}} \stackrel{!}{=} \frac{1}{2}.$$

Solving this yields $z_{Eq} \approx \frac{m}{2 \ln(2)}$. Now, the probability that an agent getting critical perishes equals approximately

$$\left(1 - \frac{d}{m}\right)^{z_{Eq}} \approx \left(1 - \frac{d}{m}\right)^{\frac{m}{2 \ln(2)}} \approx e^{-\frac{d}{2 \ln(2)}}.$$

Assuming the loss is small, approximately half of the agents get critical, i.e., they are not matched with an other agent before the end of their maximal sojourn time. Thus, the total loss among all agents equals approximately $\frac{1}{2}e^{-\frac{d}{2 \ln(2)}}$. Hence, both arguments lead to an approximately equal loss.

7.2 Heterogeneous Agents

Many applications of dynamic matching markets comprise *heterogenous* agents differing in their market behavior, for instance with respect to their urgency to be matched or with respect to their capability to form matches with other agents. In this section, we give a brief overview of how to extend our results to markets with heterogeneous agents.

First, it is possible to extend our results to a setting with heterogeneous agents regarding their urgency to be matched, that is, agents differing with respect to their maximum sojourn times. As long as all agents *individually* satisfy the conditions on the maximum sojourn time in our results, they still work. To be more explicit, consider the following scenario. There are two types of agents arriving to the market at random. The first type consists of *urgent* agents, which have a maximum sojourn time of 1. The second type of agents consists of *patient* agents having a maximal sojourn time of $C > 1$. Assume that urgent agents arrive at the market at Poisson rate m_1 , whereas patient agents independently from urgent agents arrive at Poisson rate m_2 . Denote $m = m_1 + m_2$ and assume that the probability that two agents are compatible is $p = \frac{d}{m}$, where d is a constant. Then, the consequence of Theorem 4.1 still holds when m tends to infinity. Indeed, the sum of two Poisson processes is again a Poisson process, with the sum of the previous rate functions as a new rate function. Hence, the model with urgent and patient agents has the same distribution as the original model with incoming rate $m = m_1 + m_2$, where the maximal sojourn time of an agent is 1 with probability $\frac{m_1}{m_1 + m_2}$ and C with probability $\frac{m_2}{m_1 + m_2}$.

Second, a common type of heterogeneity considered in the literature differentiates easy-to-match and hard-to-match agents [Ashlagi et al., 2019b, 2022], differing regarding compatibility. In our setting, this can be modeled by having agents arrive with Poisson rate m where membership to a class is determined by a random variable. Then, the compatibility of agents is decided

according to different density parameters d_{HH} , d_{EH} , and d_{EE} with $d_{HH} \leq d_{EH} \leq d_{EE}$ representing the density of edges between two hard-to-match, two different-type, and two easy-to-match agents, respectively. This leads to matching probabilities $p_{HH} = d_{HH}/m$, $p_{EH} = d_{EH}/m$, and $p_{EE} = d_{EE}/m$. It is possible to extend our results under certain assumptions on the relationship between the different density parameters. Still, a rigorous treatment of this model is beyond the scope of this paper, and we leave it as an intriguing direction for further research.

References

- I. Agarwal, R. Cole, and Y. Tao. Selecting a match: Exploration vs decision. Technical Report arXiv:2106.08033, arXiv.org, 2021.
- M. Akbargpour, S. Li, and S. O. Gharan. Thickness and information in dynamic matching markets. *Journal of Political Economy*, 128(3):783–815, 2020.
- R. Anderson, I. Ashlagi, D. Gamarnik, and Y. Kanoria. Efficient dynamic barter exchange. *Operations Research*, 65(6):1446–1459, 2017.
- I. Ashlagi, A. Bingaman, M. Burq, V. H. Manshadi, D. Gamarnik, C. Murphey, A. E. Roth, M. L. Melcher, and M. A. Rees. Effect of match-run frequencies on the number of transplants and waiting times in kidney exchange. *American Journal of Transplantation*, 18(5):1177–1186, 2018.
- I. Ashlagi, M. Burq, C. Dutta, P. Jaillet, A. Saberi, and C. Sholley. Edge weighted online windowed matching. In *Proceedings of the 20th ACM Conference on Economics and Computation (ACM-EC)*, pages 729–742, 2019a.
- I. Ashlagi, M. Burq, P. Jaillet, and V. H. Manshadi. On matching and thickness in heterogeneous dynamic markets. *Operations Research*, 67(4):927–949, 2019b.
- I. Ashlagi, N. Afshin, and P. Strack. Matching in dynamic imbalanced markets. *Review of Economic Studies*, 2022. Forthcoming.
- A. Aveklouris, L. DeValve, and A. R. Ward. Matching impatient and heterogeneous demand and supply. Technical Report arXiv:2102.02710, arXiv.org, 2021.
- M. Baccara, A. Collard-Wexler, L. Felli, and L. Yariv. Child-adoption matching: Preferences for gender and race. *American Economic Journal: Applied Economics*, 6(3):133–58, 2014.
- M. Baccara, S. Lee, and L. Yariv. Optimal dynamic matching. *Theoretical Economics*, 15(3):1221–1278, 2020.
- J. N. Bearden. A new secretary problem with rank-based selection and cardinal payoffs. *Journal of Mathematical Psychology*, 50(1):58–59, 2006.
- J. H. Blanchet, M. I. Reiman, V. Shah, L. M. Wein, and L. Wu. Asymptotically optimal control of a centralized dynamic matching market with general utilities. *Operations Research*, 2022. Forthcoming.
- N. Collina, N. Immorlica, K. Leyton-Brown, B. Lucier, and N. Newman. Dynamic weighted matching with heterogeneous arrival and departure rates. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*, pages 17–30, 2020.
- Y. Emek, S. Kutten, and R. Wattenhofer. Online matching: haste makes waste! In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 333–344, 2016.

- P. Erdős and A. Rényi. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, 5(1):17–61, 1960.
- J. Feldman, A. Mehta, V. Mirrokni, and S. Muthukrishnan. Online stochastic matching: Beating $1-1/e$. In *Proceedings of the 50th Symposium on Foundations of Computer Science (FOCS)*, pages 117–126, 2009.
- T. S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4(3):282–289, 1989.
- J. Fong. Search, selectivity, and market thickness in two-sided markets: Evidence from online dating. Technical Report 3458373, SSRN, 2020.
- T. Hagerup and C. Rüb. A guided tour of chernoff bounds. *Information Processing Letters*, 33(6):305–308, 1990.
- N. Kakimura and D. Zhu. Dynamic bipartite matching market with arrivals and departures. In *Proceedings of the 17th International Conference on Web and Internet Economics (WINE)*, page 544, 2021.
- R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 352–358, 1990.
- S. Kerimov, I. Ashlagi, and I. Gurvich. On the optimality of greedy policies in dynamic matching. In *Proceedings of the 23rd ACM Conference on Economics and Computation (ACM-EC)*, 2022. Forthcoming.
- A. Klenke. *Probability theory: a comprehensive course*. Springer Science & Business Media, 2013.
- G. Last and M. Penrose. *Lectures on the Poisson process*, volume 7. Cambridge University Press, 2017.
- J. D. Leshno. Dynamic matching in overloaded waiting lists. *Available at SSRN 2967011*, 2021. Working Paper.
- J. Li and S. Netessine. Higher market thickness reduces matching rate in online platforms: Evidence from a quasiexperiment. *Management Science*, 66(1):271–289, 2020.
- S. Loertscher, E. V. Muir, and P. G. Taylor. Optimal market thickness. *Journal of Economic Theory*, 200:105383, 2022.
- V. H. Manshadi, S. O. Gharan, and A. Saberi. Online stochastic matching: Online actions based on offline statistics. *Mathematics of Operations Research*, 37(4):559–573, 2012.
- A. Mehta, A. Saberi, U. V. Vazirani, and V. V. Vazirani. Adwords and generalized online matching. *Journal of the ACM*, 54(5):22–es, 2007.
- P. Mertikopoulos, H. H. Nax, and B. S. R. Pradelski. Quick or cheap? breaking points in dynamic markets. In *Proceedings of the 21st ACM Conference on Economics and Computation (ACM-EC)*, pages 877–878, 2020.
- M. Pavone, A. Saberi, M. Schiffer, and M. Tsao. Online hypergraph matching with delays. Technical Report arXiv:2009.12022, arXiv.org, 2020.
- P. P. Reese, N. Boudville, and A. X. Garg. Living kidney donation: outcomes, ethics, and uncertainty. *The Lancet*, 385(9981):2003–2013, 2015.

- A. E. Roth. Marketplaces, markets, and market design. *American Economic Review*, 108(7):1609–58, 2018.
- A. E. Roth, T. Sönmez, and M. U. Ünver. Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488, 2004.
- A. E. Roth, T. Sönmez, and M. U. Ünver. Pairwise kidney exchange. *Journal of Economic Theory*, 125(2):151–188, 2005.
- R. Shimer and L. Smith. Matching, search, and heterogeneity. *The BE Journal of Macroeconomics*, 1(1), 2001.
- M. U. Ünver. Dynamic kidney exchange. *The Review of Economic Studies*, 77(1):372–414, 2010.

A Analysis of the Greedy Algorithm

We start with the analysis of the greedy algorithm.

A.1 Loss Bounds for the Greedy Algorithm

First, we prove our coupling lemma.

Lemma 5.1. *Let μ be any probability measure on $[0, +\infty]$. Consider the pools Z_t and Z_t^∞ at time t for the greedy algorithm with departure times distributed according to μ and μ_∞ , respectively, and let Z_0 and Z_0^∞ be two starting pools with $z_0 \leq z_0^\infty$. Then, there exists a coupling of Z_t and Z_t^∞ such that*

$$z_t \leq z_t^\infty + 1 \quad (2)$$

for all $t \geq 0$.

Proof. We consider the procedure where we have two pools to which two identical copies of agents arrive according to Poisson rate m . The first pool acts according to the greedy algorithm under measure μ and initial pool Z_0 and the second one according to the greedy algorithm under measure μ_∞ and initial pool Z_0^∞ . Recall that, attached to each arriving agent v , there is a random vector $U(v) \in \{0, 1\}^\mathbb{N}$ storing the random compatibility with other agents, as described in Section 3. In particular, if the pool is of size K and an agent v arrives, then the agent directly forms an edge if and only if $U_j(v) = 1$ for some $j \in \{1, \dots, K\}$. If there are compatible agents in the pool, then the agent v selects a partner among them chosen uniformly at random.

We show the assertion (2) by a case distinction. The statement is clear for the case $t = 0$. For $t > 0$, the only interesting times are the ones where agents arrive, as agents can only perish from the first pool. Hence, z_t^∞ cannot get smaller without the arrival of a new agent.

If $z_t < z_t^\infty$ and an agent arrives at time t , then z_t can increase by at most 1 and z_t^∞ can decrease by at most 1, so (2) is still satisfied after the arrival of the agent. If $z_t = z_t^\infty$ and an agent arrives at time t , then the agent matches under the measure μ if and only if she matches under the measure μ_∞ , so (2) is also satisfied after the arrival of the agent. If $z_t = z_t^\infty + 1$ and an agent arrives, then if the agent does not match with an element in the first pool, she also does not match with an element in the second pool, so both pools increase in size. So it is not possible that z_t gets larger while z_t^∞ gets smaller at the same time which shows that (2) is always satisfied. \square

Now, we prove the two lemmas regarding the pool size.

Lemma 5.2. *Let ρ be the stationary measure of the pool size without perishing. Then,*

$$\rho\left(\left(\frac{C_1 \ln(2)m}{d} + \frac{3}{2} \ln(m)^2, \infty\right) \cap \mathbb{N}\right) \leq m^{-9}$$

for m large enough.

Proof. Let $j \geq \frac{C_1 \ln(2)m}{d}$. Then, we have

$$p(j, j+1) = \left(1 - \frac{d}{m}\right)^j \leq \left(1 - \frac{d}{m}\right)^{\frac{C_1 \ln(2)m}{d}} \leq \frac{1}{2} e^{-\frac{10}{\ln(m)}}.$$

On the other hand, we also have that $p(j, j-1) \geq 1/2$. Inserting this into (5), we get

$$\rho(j+1) = \rho(j) \frac{p(j, j+1)}{p(j+1, j)} \leq e^{-\frac{10}{\ln(m)}} \rho(j).$$

Thus, we get inductively

$$\rho\left(\left\lceil \frac{C_1 \ln(2)m}{d} \right\rceil + n\right) \leq e^{-n \frac{10}{\ln(m)}} \rho\left(\left\lceil \frac{C_1 \ln(2)m}{d} \right\rceil\right) \leq e^{-n \frac{10}{\ln(m)}}.$$

Summing this over different values of n , we get for m large enough

$$\begin{aligned} \sum_{k > \frac{C_1 \ln(2)m}{d} + \frac{3}{2} \ln(m)^2} \rho(k) &\leq \sum_{k=\lceil \ln(m)^2 \rceil}^{\infty} \rho\left(\left\lceil \frac{C_1 \ln(2)m}{d} \right\rceil + k\right) \leq \sum_{k=\lceil \ln(m)^2 \rceil}^{\infty} e^{-k \frac{10}{\ln(m)}} \\ &= e^{-\lceil \ln(m)^2 \rceil \frac{10}{\ln(m)}} \sum_{k=0}^{\infty} e^{-k \frac{10}{\ln(m)}} \\ &\leq m^{-10} \frac{1}{1 - e^{-\frac{10}{\ln(m)}}} \leq m^{-10} \frac{\ln(m)}{5} \leq m^{-9}. \end{aligned}$$

There, we used that $1 - e^{-\frac{10}{\ln(m)}} \geq \frac{5}{\ln(m)}$ for m large enough in the second to last step. \square

Lemma 5.4. *Assume that the departure time is distributed according to an arbitrary probability measure μ . Further, assume that Z is a pool of size at most $C_2 \frac{\ln(2)}{d} m$. Then, for large enough m , it holds that*

$$\mathbb{P}_{\mu}\left(z_t \leq C_3 \frac{\ln(2)m}{d} \text{ for all } t \leq 1 \mid Z_0 = Z\right) \geq 1 - m^{-9}.$$

Proof. We seek an upper bound for the probability that the pool size reaches $C_3 \frac{\ln(2)}{d} m$ when starting at a size of at most $C_2 \frac{\ln(2)}{d} m$. Let $M = \lceil C_2 \frac{\ln(2)}{d} m \rceil$. When the pool has a size of at least M and an agent arrives, then the probability that the agent joins the pool, i.e., that she is not matched directly, equals

$$\left(1 - \frac{d}{m}\right)^M \leq \left(1 - \frac{d}{m}\right)^{C_2 \frac{\ln(2)}{d} m} \leq e^{-C_2 \ln(2)} \leq \frac{1}{2} e^{-\frac{10}{\ln(m)}}. \quad (7)$$

There, we used Lemma C.1 in the second inequality. Whenever a new agent arrives, the pool size can increase or decrease by 1, where we have a tendency towards a decrease of the pool size. Thus, in particular, if the pool starts at size M , the probability that the pool size increases beyond $C_3 \frac{\ln(2)}{d} m$ before decreasing to $C_2 \frac{\ln(2)}{d} m$ is bounded by the probability of this event for the case of a random walk with drift to the left of at least $\frac{1}{2} e^{-\frac{10}{\ln(m)}}$. The euclidean distance between $C_2 \frac{\ln(2)}{d} m$ and $C_3 \frac{\ln(2)}{d} m$ is $2 \ln(m)^2$. Hence, there are more than $\ln(m)^2$ integer points between these values, for $m \geq 3$. We can apply Lemma C.3 with $\frac{1}{2} - \varepsilon = \frac{1}{2} e^{-\frac{10}{\ln(m)}}$ in order to bound the probability that the Markov chain with drift crosses such an interval. Thus, the probability of this event is bounded by

$$\left(2 \frac{1}{2} e^{-\frac{10}{\ln(m)}}\right)^{\ln(m)^2} = m^{-10}.$$

If the pool size starts at some value less than $\frac{C_2 \ln(2)}{d} m$ and then grows to some value greater than $\frac{C_3 \ln(2)}{d} m$ then at some time interval in between the pool size starts with size M and reaches size $\frac{C_3 \ln(2)}{d} m$ before going back to $M - 1$.

Assume that k agents arrive during the time interval $[0, 1]$, i.e., $a_{[0,1]} = k$. For $i \in \{1, \dots, k\}$, the probability that the pool size is M when the i -th agent arrives and then increases beyond $\frac{C_3 \ln(2)}{d} m$ before going to $M - 1$ is bounded by m^{-10} by the previous calculations. Thus, by a

union bound, the probability that the pool size increases beyond $\frac{C_3 \ln(2)}{d}m$ is bounded by km^{-10} . Conditioning on the number of arrivals in the time interval $[0, 1]$, we get

$$\begin{aligned} & \mathbb{P}_\mu \left(z_t > \frac{C_3 \ln(2)m}{d} \text{ for some } t \leq 1 \mid Z_0 = Z \right) \\ &= \sum_{k=1}^{\infty} \mathbb{P}_\mu \left(z_t > \frac{C_3 \ln(2)m}{d} \text{ for some } t \leq 1 \mid Z_0 = Z, a_{[0,1]=k} \right) \mathbb{P}(a_{[0,1]} = k) \\ &\leq \sum_{k=1}^{\infty} km^{-10} \mathbb{P}(a_{[0,1]} = k) = m^{-9} \end{aligned}$$

as $a_{[0,1]}$ is a Poisson random variable with expectation value m . \square

Next, we provide the technical lemma that allows us to compute the loss by integrating over the probability of perishing.

Lemma 5.5. *Let a matching algorithm ALG be given and assume that departure time is distributed according to an arbitrary probability μ . Then,*

$$\mathbf{L}_{ALG}(m, T) = \frac{1}{T} \int_0^T \mathbb{P}_\mu(\text{an agent arriving at time } t \text{ perishes before time } T) dt,$$

and

$$\mathbf{L}_{ALG} = \limsup_{m, T \rightarrow \infty} \frac{1}{T} \int_0^T \mathbb{P}_\mu(\text{an agent arriving at time } t \text{ perishes}) dt.$$

Proof. In the proof, we omit the subscript μ , because the distribution of the sojourn times are fixed. Recall that X_i denotes the maximum sojourn time of agent v_i . In Section 3 we described that every agent v carries a random variable $U(v)$ that determines with which agents in the pool she forms an edge at arrival. We also assume that attached to agent v there is a random variable $T(v)$, which has the distribution $\text{Uniform}[0, 1]^{\otimes \mathbb{N}}$. Whenever an agent v decides to match with another agent, but has several possibilities, it choose the agent v_i , where $T_i(v)$ is minimal. As the uniform distribution is absolutely continuous with respect to the Lebesgue measure, there is almost surely a unique minimizer. By symmetry, the agent that v matches with is chosen uniformly among all possibilities. So T helps with breaking ties, in case there are some. For the complete process we have four different sources of randomness: The incoming agents arrive random, according to a Poisson process, their departure times are random with distribution μ , the forming of edges which is determined by the vectors $U(v)$, and the random breaking of ties, if necessary, that is determined by the vectors $T(v)$. Let $Z = ((X_i, U(v_i), T(v_i)))_{i \geq 1}$ be the random variable containing all information about the random variables $X_i, U(v_i)$ and $T(v_i)$ for $i \geq 1$. Given Z , the evolution of an agent v_i depends only on the random variable $Y = (Y_t)_{t \in [0, T]}$ distributed according to a Poisson process with rate m , where Y_t is the number of agents arriving at time t . Thus, we can define the function

$$f_Z(t, Y) := \mathbf{1} \{ \text{an agent arrives at time } t \text{ and perishes before time } T \},$$

where $\mathbf{1}$ is the indicator function. The loss until time T can clearly be expressed as $\mathbf{L}_{ALG}(m, T) = \frac{1}{mT} \mathbb{E} \left[\int_0^T f_Z(t, Y) Y(dt) \right]$. Note that Z and Y are independent, so the process Y is still the same Poisson process, given the information contained in Z . Furthermore, the intensity measure of this Poisson process is simply m times the Lebesgue measure on $[0, T]$. Thus, we can apply the tower property in the first and fourth equality and Mecke's equation [see, e.g., Last and Penrose,

2017, Theorem 4.1] in the second equality to compute

$$\begin{aligned}
& \mathbb{E} \left[\int_0^T f_Z(t, Y) Y(dt) \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[\int_0^T f_Z(t, Y) Y(dt) \middle| \sigma(Z) \right] \right] = \mathbb{E} \left[\int_0^T \mathbb{E} \left[f_Z(t, Y + \delta_t) \middle| \sigma(Z) \right] m dt \right] \\
&= m \int_0^T \mathbb{E} \left[\mathbb{E} \left[f_Z(t, Y + \delta_t) \middle| \sigma(Z) \right] \right] dt = m \int_0^T \mathbb{E} [f_Z(t, Y + \delta_t)] dt \\
&= m \int_0^T \mathbb{P}(\text{an agent arriving at time } t \text{ perishes before time } T) dt.
\end{aligned}$$

There, $\sigma(Z)$ denotes the σ -algebra generated by Z , and δ_t is the Dirac delta function, i.e., $\delta_t(s) = \mathbf{1}_{\{s=t\}}$. Now, let us consider the second statement of the lemma. By the equality that we proved just before, we know that

$$\begin{aligned}
\mathbf{L}_{\text{ALG}} &= \limsup_{m, T \rightarrow \infty} \frac{1}{T} \int_0^T \mathbb{P}_\mu(\text{an agent arriving at time } t \text{ perishes before time } T) dt \\
&= \limsup_{m, T \rightarrow \infty} \int_0^1 \mathbb{P}_\mu(\text{an agent arriving at time } sT \text{ perishes before time } T) ds. \tag{8}
\end{aligned}$$

where the last equality holds by integration by substitution. For $T > 0$, we define the functions $\tilde{f}_T, f_T : [0, 1] \rightarrow [0, 1]$ by

$$\begin{aligned}
f_T(s) &:= \mathbb{P}_\mu(\text{an agent arriving at time } sT \text{ perishes before time } T), \\
\tilde{f}_T(s) &:= \mathbb{P}_\mu(\text{an agent arriving at time } sT \text{ perishes}).
\end{aligned}$$

We clearly have $f_t(s) \leq \tilde{f}_T(s)$ for all $s \in [0, 1]$, and for fixed $s < 1$ and $m > 0$ we also have

$$\begin{aligned}
0 \leq \tilde{f}_T(s) - f_T(s) &= \mathbb{P}_\mu(\text{an agent arriving at time } sT \text{ perishes after time } T) \\
&\leq \mathbb{P}_\mu(\text{the maximum sojourn of an agent is in } [(1-s)T, \infty)) \\
&= \mu([(1-s)T, \infty))
\end{aligned}$$

which converges to 0 as T tends to infinity. As the difference $f_T(s) - \tilde{f}_T(s)$ is bounded by 1 for all $s \in [0, 1]$ we get that

$$\lim_{T \rightarrow \infty} \int_0^1 f_T(s) - \tilde{f}_T(s) ds = 0$$

and as this convergence holds uniformly for all m , we get

$$\begin{aligned}
\mathbf{L}_{\text{ALG}} &= \limsup_{m, T \rightarrow \infty} \frac{1}{T} \int_0^T \mathbb{P}_\mu(\text{an agent arriving at time } t \text{ perishes before time } T) dt \\
&= \limsup_{m, T \rightarrow \infty} \frac{1}{T} \int_0^T \mathbb{P}_\mu(\text{an agent arriving at time } t \text{ perishes}) dt.
\end{aligned}$$

□

In the next lemma, we show some rather general scaling invariance for the connection between loss and distribution of the departure times. To state the result, we need to express the loss with respect to parameters that we usually suppress. Therefore, given a probability measure ρ and a density parameter κ , let $\mathbf{L}_{\text{ALG}}(\rho, \kappa)$ denote the loss of ALG with respect to departure times distributed according to ρ and density parameter κ .

Lemma A.1. *Let μ be a probability measure on $[0, +\infty]$, let $c > 0$, and let ν be the probability measure defined by $\nu([a, b]) = \mu([ca, cb])$ for all $a, b \in [0, +\infty]$. Then,*

$$\mathbf{L}_{GDY}(\nu, d) = \mathbf{L}_{GDY}(\mu, cd).$$

Proof. The claim follows by a change of time. If we consider the model with departure times distributed according to ν and speed up time by a factor of c , then agents arrive at Poisson rate cm and the departure time X of an agent satisfies

$$\mathbb{P}(X \leq a) = \nu([0, a/c]) = \mu([0, a]).$$

Thus, the departure times are distributed according to μ . However, two agents are still compatible with a probability of $\frac{d}{m} = \frac{cd}{cm}$. So, effectively we see a model with departure times distributed according to μ and density parameter cd . \square

Combining Theorem 4.1 and Lemma A.1, we obtain the scaled version of Theorem 4.1.

Corollary 4.2. *Let $\varepsilon > 0$ and assume that the departure time is distributed according to a probability measure μ with $\mu([0, \varepsilon)) = 0$. For $d \geq 2$, we have*

$$\mathbf{L}_{GDY} \leq e^{-\frac{\varepsilon d}{2 \ln(2)}}.$$

Now, we provide the proof for the auxiliary lemma used to find lower bounds for the loss of the greedy algorithm.

Lemma A.2. *Let $\varepsilon, \delta > 0$ and assume that the departure time is distributed according to a probability measure μ with $\mu([0, \varepsilon]) \geq \delta$. Then,*

$$\mathbf{L}_{GDY} \geq \frac{\delta}{2} e^{-\varepsilon d}.$$

Proof. Consider an agent v_i arriving at time t which is not matched upon arrival. Note that compatible agents come at Poisson rate $m \cdot \frac{d}{m} = d$. Recall that X_i is the maximum sojourn time of v_i . As the maximum sojourn time and the arrival of compatible agents are independent, we get

$$\begin{aligned} & \mathbb{P}_\mu(v \text{ perishes} \mid v \text{ not matched upon arrival}) \\ & \geq \mathbb{P}_\mu(\text{no compatible agent arrives in } [t, t + X_i] \mid v \text{ not matched upon arrival}) \\ & \geq \mathbb{P}_\mu(X_i \leq \varepsilon) \mathbb{P}_\mu(\text{no compatible agent arrives in } [t, t + \varepsilon]) \geq \delta e^{-\varepsilon d}. \end{aligned}$$

Additionally, $\mathbb{P}_\mu(v \text{ not matched upon arrival}) \geq \frac{1}{2}$, because at least half of the agents v are not matched upon arrival. Together,

$$\begin{aligned} \mathbb{P}_\mu(v \text{ perishes}) &= \mathbb{P}_\mu(v \text{ perishes} \mid v \text{ not matched upon arrival}) \\ & \cdot \mathbb{P}_\mu(v \text{ not matched upon arrival}) \geq \frac{1}{2} \delta e^{-\varepsilon d}. \end{aligned}$$

\square

We can now use Lemma A.2 to obtain our two propositions regarding lower bounds.

Proposition 4.3. *Assume that the departure time is given by the degenerate measure μ with $\mu(\{1\}) = 1$. Then,*

$$\mathbf{L}_{GDY} \geq \frac{1}{2} e^{-d}.$$

From the lemma, we can also deduce a lower bound for the loss of the greedy algorithm for probability measures with non-negligible probability for a quick departure. This proposition follows by taking d with $\frac{1}{d} \leq \varepsilon_0$ and applying Lemma A.2 for $\varepsilon = \frac{1}{d}$.

Proposition 4.4. *Let $c > 0$, $\varepsilon_0 > 0$, and assume that the departure time is distributed according to a probability measure μ with $\mu([0, \varepsilon]) \geq c\varepsilon$ for all $\varepsilon \leq \varepsilon_0$. Then, for d large enough, it holds that*

$$\mathbf{L}_{GDY} \geq \frac{c}{6d}.$$

Next, we consider the general lower bound for the loss of any algorithm.

Proposition 4.5. *Assume that the departure time is given by a distribution μ with $\int x\mu(dx) = E$, i.e., E is the expected maximum sojourn time of an agent. Then, for any algorithm ALG , it holds that*

$$\mathbf{L}_{ALG} \geq e^{-2Ed}.$$

Proof. We start with a technical observation. Let X be a non-negative random variable with distribution μ . It is a well-known fact that

$$\sum_{k=1}^{\infty} \mu([k, \infty)) = \sum_{k=1}^{\infty} \mathbb{P}(X \geq k) \leq \mathbb{E}[X].$$

Multiplying by a factor $\varepsilon > 0$ we get that

$$\sum_{k=1}^{\infty} \mu([k\varepsilon, \infty)) = \sum_{k=1}^{\infty} \mathbb{P}(X \geq k\varepsilon) = \sum_{k=1}^{\infty} \mathbb{P}\left(\frac{X}{\varepsilon} \geq k\right) \leq \mathbb{E}\left[\frac{X}{\varepsilon}\right] = \int \frac{x}{\varepsilon} \mu(dx). \quad (9)$$

We will use this inequality in the following.

We proceed with the main proof. For some time $t \in \mathbb{R}_{\geq 0}$, let V_t be the set of agents whose maximum sojourn time is reached after time t . So V_t is the set of all agents that are possibly in the pool at time t . We first prove the intermediary claim that

$$\mathbb{E}[|V_t|] \leq mE. \quad (10)$$

Fix $\varepsilon > 0$. For an agent v arriving at time $s \in [t - (k+1)\varepsilon, t - k\varepsilon]$, the probability of $v \in V_t$ is at most $\mu([k\varepsilon, \infty))$. The expected number of agents arriving in a time interval of length ε is εm .

Consequently,

$$\mathbb{E}[|V_t|] \leq \sum_{k=0}^{\infty} \varepsilon m \mu([k\varepsilon, \infty)) \leq \varepsilon m + \varepsilon m \sum_{k=1}^{\infty} \mu([k\varepsilon, \infty)) \stackrel{(9)}{\leq} \varepsilon m + \varepsilon m \int \frac{x}{\varepsilon} \mu(dx) \leq \varepsilon m + mE.$$

Hence, as $\varepsilon > 0$ was chosen arbitrarily, this already implies Equation (10).

Now, assume that an agent v arrives at time t . We want to estimate the probability that no agent in V_t is compatible with v . By the law of total probability we have

$$\begin{aligned} & \mathbb{P}(\text{No agent in } V_t \text{ is compatible with } v) \\ &= \sum_{k=0}^{\infty} \mathbb{P}(|V_t| = k) \mathbb{P}(\text{no agent in } V_t \text{ is compatible with } v \mid |V_t| = k) \\ &= \sum_{k=0}^{\infty} \mathbb{P}(|V_t| = k) \left(1 - \frac{d}{m}\right)^k = \mathbb{E}\left[\left(1 - \frac{d}{m}\right)^{|V_t|}\right] \geq \left(1 - \frac{d}{m}\right)^{\mathbb{E}[|V_t|]} \stackrel{(10)}{\geq} \left(1 - \frac{d}{m}\right)^{mE}. \end{aligned} \quad (11)$$

The first inequality holds by Jensen's inequality, as the function $x \mapsto (1 - \frac{d}{m})^x$ is convex. The second inequality holds because this function is also decreasing.

Next, we estimate the probability that no agent compatible with v arrives during the sojourn time of v . Let M_v denote the maximum sojourn time of v , i.e., the random variable M_v is distributed according to μ . The next idea is to condition on the event $\{M_v = s\}$ for some positive real s . Then, the law of total probability and Jensen's inequality imply that

$$\begin{aligned} & \mathbb{P}(\text{no agent arriving in } [t, t+s] \text{ is compatible with } v | M_v = s) \\ &= \sum_{k=0}^{\infty} \mathbb{P}(a_{[t, t+s]} = k | M_v = s) \mathbb{P}(\text{no agent arr. in } [t, t+s] \text{ is comp. with } v | M_v = s, a_{[t, t+s]} = k) \\ &= \sum_{k=0}^{\infty} \mathbb{P}(a_{[t, t+s]} = k) \left(1 - \frac{d}{m}\right)^k = \mathbb{E} \left[\left(1 - \frac{d}{m}\right)^{a_{[t, t+s]}} \right] \geq \left(1 - \frac{d}{m}\right)^{\mathbb{E}[a_{[t, t+s]}]} = \left(1 - \frac{d}{m}\right)^{ms} \end{aligned}$$

where we used Jensen's inequality as before. By the tower rule and the law of total probability we thus get

$$\begin{aligned} & \mathbb{P}(\text{no agent arriving in } [t, t+X] \text{ is compatible with } v) \\ &= \mathbb{E} [\mathbb{P}(\text{no agent arriving in } [t, t+X] \text{ is compatible with } v | X)] \\ &\geq \mathbb{E} \left[\left(1 - \frac{d}{m}\right)^{mX} \right] \geq \left(1 - \frac{d}{m}\right)^{\mathbb{E}[mX]} = \left(1 - \frac{d}{m}\right)^{mE}. \end{aligned} \tag{12}$$

In the last inequality, we have applied Jensen's inequality once again.

Some agent that arrives at time t , which is compatible with no agent in V_t , and for which no compatible agent arrives during her maximal sojourn time, cannot get matched under any algorithm. For some agent v arriving at time t , the compatibility to agents that arrived before time t and the compatibility to agents that arrive after time t is independent. So in particular, combining (11) with (12) gives that

$$\mathbb{P}(v \text{ perishes}) \geq \left(1 - \frac{d}{m}\right)^{2mE}$$

for any agent v . For $m \rightarrow \infty$, we obtain that $\mathbb{P}(v \text{ perishes}) \geq e^{-2dE}$, and as the time t at which v arrives was arbitrary we get that

$$\mathbf{L}_{\text{ALG}} \geq e^{-2dE}$$

for any algorithm ALG by Lemma 5.5. □

A.2 Further Objectives for the Greedy Algorithm

Finally, we consider the waiting time and congestion objectives.

First, we prove the lower bound for the loss of the patient algorithm.

Proposition 4.7. *Let $h > 0$ and assume that the departure time is distributed according to a probability measure μ with $\mu([h, \infty]) > \frac{4}{5}$. Then, for sufficiently large m and T , it holds that*

$$\mathbf{W}_{\text{PAT}}(m, T) \geq \frac{mTh}{4}. \tag{1}$$

Proof. Assume that the agents v_1, \dots, v_K arrive up to time $T - h$, and the agents u_1, \dots, u_L arrive in the time interval $(T - h, T]$. We will assume that $20L \leq K$, which will hold with high probability for T large enough.

We will first focus on the agents v_1, \dots, v_K . The expected number among them with a maximal sojourn time in $[h, \infty]$ is $\mu([h, \infty])K$. We call such an agent an agent with high maximal sojourn. By the strong law of large numbers there will be at least $\frac{4K}{5}$ such agents with high probability. We denote by $\mathcal{A}(m, T)$ the event that $20L \leq K$ and that at least $\frac{4K}{5}$ agents have a high maximal sojourn. Then $\limsup_{m, T \rightarrow \infty} \mathbb{P}(\mathcal{A}(m, T)) = 1$. Assume that there are at least $\frac{4K}{5}$ agents with high maximal sojourn. At most $\frac{K}{5} + L$ of them can leave the pool by matching with an agent that does not have a high maximal sojourn or arrived after time $T - h$. So at least $\frac{3K}{5} - L$ agents of v_1, \dots, v_K with high maximal sojourn need to stay in the market for a time of at least h , or they match with another agent that of v_1, \dots, v_K with high maximal sojourn. Since the latter case can only affect half of them, we obtain a total waiting time of

$$\frac{1}{2} \left(\frac{3K}{5} - L \right) h \geq 0.275Kh$$

where we used the assumption $20L \leq K$ in the last step. So in particular we have that

$$\mathbf{W}_{\text{PAT}}(m, T) \geq \mathbb{E}[0.275Kh | \mathcal{A}(m, T)] \mathbb{P}(\mathcal{A}(m, T)) \geq \mathbb{E}[0.275Kh \mathbf{1}_{\mathcal{A}(m, T)}] \mathbb{P}(\mathcal{A}(m, T)) \quad (13)$$

K is a random variable with $\text{Poisson}(m(T - h))$ distribution. So in particular its second moment $\mathbb{E}[K^2]$ is given by $(m(T - h))^2 + m(T - h)$. Thus we get by Cauchy-Schwarz that

$$\mathbb{E}[0.275Kh \mathbf{1}_{\mathcal{A}(m, T)^C}] \leq h \sqrt{\mathbb{E}[K^2]} \sqrt{\mathbb{P}(\mathcal{A}(m, T)^C)} \leq h \sqrt{(m(T - h))^2 + m(T - h)} \sqrt{\mathbb{P}(\mathcal{A}(m, T)^C)}$$

and inserting this into (13) we see that

$$\begin{aligned} \mathbf{W}_{\text{PAT}}(m, T) &\geq \mathbb{E}[0.275Kh \mathbf{1}_{\mathcal{A}(m, T)}] \mathbb{P}(\mathcal{A}(m, T)) \\ &= \left(\mathbb{E}[0.275Kh] - \mathbb{E}[0.275Kh \mathbf{1}_{\mathcal{A}(m, T)^C}] \right) \mathbb{P}(\mathcal{A}(m, T)) \\ &\geq 0.275m(T - h) \mathbb{P}(\mathcal{A}(m, T)) - h \sqrt{(m(T - h))^2 + m(T - h)} \sqrt{\mathbb{P}(\mathcal{A}(m, T)^C)} \geq \frac{mT}{4} \end{aligned}$$

where the last inequality holds for m, T large enough. \square

Next, we consider the integral expression of waiting time with respect to the expected pool size.

Proposition 4.8. *Assume that the departure time is distributed according to an arbitrary probability measure μ . Then, for any algorithm ALG , it holds that*

$$\mathbf{W}_{\text{ALG}}(m, T) = \int_0^T \mathbb{E}_\mu[z_s] ds = T \cdot \mathbf{C}_{\text{ALG}}(m, T).$$

Proof. Let V be the set of agents that arrive up to time T . By Fubini's theorem, we have

$$\sum_{v \in V} \int_0^T \mathbf{1}_{v \in Z_s} ds = \int_0^T \sum_{v \in V} \mathbf{1}_{v \in Z_s} ds = \int_0^T z_s ds.$$

Recall that $\mathbf{1}$ denotes the indicator function. Another application of Fubini's theorem yields that

$$\mathbf{W}_{\text{ALG}}(m, T) = \int_0^T \mathbb{E}_\mu[z_s] ds.$$

\square

We use this proposition to prove the upper bound on the total waiting time.

Proposition 4.9. *Assume that the departure time is distributed according to an arbitrary probability measure μ . Then, for large enough m , under the greedy algorithm, it holds that*

$$\mathbf{W}_{GDY}(m, T) \leq \frac{6mT}{5d}.$$

Proof. Using Proposition 4.8, we would like to get an upper bound on $\mathbb{E}_\mu[z_s]$. The key idea is to consider the Markov chain $(Z_t^\infty)_{t \geq 0}$ with infinite departure time started at stationarity and apply Lemma 5.1. By Lemma C.1, we have $p(k, k+1) = (1 - \frac{d}{m})^k \leq 1/3$ for $k \geq \frac{\ln(3)m}{d}$ for the transition probabilities with respect to μ_∞ . Thus, the stationary distribution satisfies $\rho(k+1) = \frac{p(k, k+1)}{p(k+1, k)} \rho(k) \leq \frac{1}{2} \rho(k)$ for these k . However, this already implies that when z_0^∞ is distributed according to the stationary distribution, then

$$\begin{aligned} \mathbb{E}_{\mu_\infty}[z_0^\infty] &= \sum_{k=1}^{\infty} \rho(k)k = \sum_{k=1}^{\lceil \frac{\ln(3)m}{d} \rceil + 10} \rho(k)k + \sum_{\lceil \frac{\ln(3)m}{d} \rceil + 11}^{\infty} \rho(k)k \\ &\leq \lceil \frac{\ln(3)m}{d} \rceil + 10 + \sum_{k=11}^{\infty} 2^{-k} \left(\lceil \frac{\ln(3)m}{d} \rceil + k \right) \\ &\leq \lceil \frac{\ln(3)m}{d} \rceil + 10 + 2^{-10} \lceil \frac{\ln(3)m}{d} \rceil + 1 \leq 1.01 \lceil \frac{\ln(3)m}{d} \rceil + 11. \end{aligned}$$

Thus, we can apply Lemma 5.1 to obtain

$$\mathbf{W}_{GDY}(m, T) = \int_0^T \mathbb{E}_\mu[z_s] ds \leq \int_0^T \mathbb{E}_{\mu_\infty}[z_s^\infty + 1] ds \leq T \left(1.01 \lceil \frac{\ln(3)m}{d} \rceil + 12 \right) \leq \frac{6mT}{5d},$$

where the last inequality holds for m large enough. \square

Next, we provide the proof for the lower bound on the waiting time in the greedy algorithm.

Proposition 4.10. *Assume that the departure time is distributed according to a probability measure μ with $\mu([\frac{1}{2d}, \infty]) > \frac{9}{10}$. Further, assume that $T \geq \frac{3}{d}$. Then, under any algorithm ALG, it holds that*

$$\mathbf{W}_{ALG}(m, T) \geq \frac{mT}{20d}.$$

Proof. We say that an agent arriving at time $t \in [\frac{1}{2d}, T - \frac{1}{2d}]$ is *blocked* if no compatible agent arrives in the interval $[t - \frac{1}{2d}, t + \frac{1}{2d}]$, and the maximum sojourn time of this agent is at least $\frac{1}{2d}$. So in particular the agent will still be in the market at time $t + \frac{1}{2d}$. For some agent arriving at time t , compatible agents arrive at rate d . So the probability that there arrives no compatible agent in the time interval $[t - \frac{1}{2d}, t + \frac{1}{2d}]$ is given by $\exp(-d \frac{1}{d}) = e^{-1}$.

As the arrival process of compatible agents is independent of the maximum sojourn time, we get that for an agent v arriving at time $t \in [\frac{1}{2d}, T - \frac{1}{2d}]$

$$\mathbb{P}(v \text{ is blocked}) = \mu\left(\left[\frac{1}{2d}, \infty\right]\right) e^{-1} \geq \frac{0.9}{e}$$

by our assumption on μ . Let $B(T)$ be the number of blocked agents arriving in the interval $[\frac{1}{2d}, T - \frac{1}{2d}]$. An application of Mecke's equation as in Lemma 5.5, yields that

$$\mathbb{E}[|B(T)|] \geq m \left(T - \frac{1}{d} \right) \frac{0.9}{e}.$$

In the following, we will argue that under *any* matching algorithm one, it holds that

$$\mathbf{W}_{\text{ALG}}(m, T) \geq \frac{\mathbb{E}[|B(T)|]}{4d}.$$

This will then imply the statement.

Agents can leave the market either alone or in pairs of two. Assume that a blocked agent v leaves the market alone, then it had a waiting time of at least $\frac{1}{2d}$. If a blocked agent v gets matched with an agent that arrived after v , then the blocked agent also has had to wait for a time of at least $\frac{1}{2d}$. If a blocked agent v gets matched with an agent that arrived before v , then the agent with whom v got matched needed to have a waiting time of at least $\frac{1}{2d}$, as no agent compatible with v arrived in the time window of length $\frac{1}{2d}$ before the arrival of v . This also leads to a higher total waiting time. However, we need to take care of the possibility of double counting here. If two blocked agents are blocked and compatible with each other, then we must not accumulate a contribution of $\frac{1}{2d}$ to the total waiting time for each of them but only for the pair. We can take care of this by a division by 2. Hence, by monotonicity of expectations, it holds that

$$\mathbf{W}_{\text{ALG}}(m, T) \geq \frac{1}{2d} \frac{\mathbb{E}[|B(T)|]}{2}.$$

Thus,

$$\mathbf{W}_{\text{ALG}}(m, T) \geq \frac{1}{4d} \mathbb{E}[|B(T)|] \geq \frac{9}{40ed} m \left(T - \frac{1}{d} \right) \geq \frac{mT}{20d}.$$

The last inequality holds, as $T \geq \frac{3}{d}$ which is equivalent to $T - \frac{1}{d} \geq \frac{2T}{3}$, and $\frac{9}{40e} \frac{2}{3} \geq \frac{1}{20}$. \square

B Analysis of the Patient Algorithm

In this section, we provide the proof for our theorem regarding the loss of the patient algorithm for the case of agents with a maximum sojourn time of exactly one time unit.

Theorem 4.6. *Assume that the departure time is distributed according to the degenerate probability measure μ with $\mu(\{1\}) = 1$. Then, it holds that*

$$\mathbf{L}_{\text{PAT}} \leq e^{-\frac{d}{5}}.$$

Proof. As μ is fixed throughout the proof, we omit it as subscript. By Lemma 5.5 and unit departure times, we have

$$\mathbf{L}_{\text{PAT}} = \limsup_{m, T \rightarrow \infty} \frac{1}{T} \int_1^T \mathbb{P}(\text{an agent arriving at time } t - 1 \text{ perishes}) dt.$$

Assume that the pool has a size of k at time t and an agent gets critical. Then, the probability that the agent perishes is given by $(1 - \frac{d}{m})^{k-1}$. Thus, in order to prove small loss, it suffices to show that the pool size is large with relatively high probability. Let $t \geq \frac{4}{3}$ be arbitrary. Our goal is to show that $z_t > \frac{m}{5}$ with high probability. For this, we condition on the pool size at time $t - \frac{1}{3}$. By the law of total probability, we have

$$\begin{aligned} \mathbb{P}\left(z_t \leq \frac{m}{5}\right) &= \mathbb{P}\left(z_t \leq \frac{m}{5} \mid z_{t-\frac{1}{3}} < \frac{m}{8}\right) \mathbb{P}\left(z_{t-\frac{1}{3}} < \frac{m}{8}\right) \\ &\quad + \mathbb{P}\left(z_t \leq \frac{m}{5} \mid z_{t-\frac{1}{3}} \geq \frac{m}{8}\right) \mathbb{P}\left(z_{t-\frac{1}{3}} \geq \frac{m}{8}\right) \end{aligned} \tag{14}$$

and thus it suffices to bound each of the two summands above.

We start with the first one. Each agent that is in the pool at time $t - \frac{1}{3}$ can take at most one additional agent out of the pool when she gets critical. Thus, conditioned on the event where $z_{t-\frac{1}{3}} < \frac{m}{8}$, it holds that $z_t \geq a_{(t-\frac{1}{3},t)} - z_{t-\frac{1}{3}} \geq a_{(t-\frac{1}{3},t)} - \frac{m}{8}$. By independence of $a_{(t-\frac{1}{3},t)}$ and $z_{t-\frac{1}{3}}$, this already implies that

$$\mathbb{P}\left(z_t \leq \frac{m}{5} \mid z_{t-\frac{1}{3}} < \frac{m}{8}\right) \leq \mathbb{P}\left(a_{(t-\frac{1}{3},t)} \leq \frac{m}{5} + \frac{m}{8} \mid z_{t-\frac{1}{3}} < \frac{m}{8}\right) = \mathbb{P}\left(a_{(t-\frac{1}{3},t)} \leq \frac{13}{40}m\right) \stackrel{\text{Lemma C.2}}{\leq} m^{-2}$$

for m large enough. This directly implies that

$$\mathbb{P}\left(z_t \leq \frac{m}{5} \mid z_{t-\frac{1}{3}} < \frac{m}{8}\right) \mathbb{P}\left(z_{t-\frac{1}{3}} < \frac{m}{8}\right) \leq m^{-2}$$

for m large enough. Now let us consider the bound for the second addend in (14). Here we have that

$$\mathbb{P}\left(z_t \leq \frac{m}{5} \mid z_{t-\frac{1}{3}} \geq \frac{m}{8}\right) \mathbb{P}\left(z_{t-\frac{1}{3}} \geq \frac{m}{8}\right) = \mathbb{P}\left(z_t \leq \frac{m}{5}, z_{t-\frac{1}{3}} \geq \frac{m}{8}\right). \quad (15)$$

We now condition on the pool size at time $t - \frac{1}{3}$ and on the number of arriving agents in certain intervals. This is necessary to avoid possible coherences between the evolution of the pool size and the arrival of new agents. For this, define the vector $a = (a_1, a_2, a_3) := (a_{[t-\frac{4}{3}, t-1]}, a_{[t-1, t-\frac{2}{3}]}, a_{[t-\frac{2}{3}, t-\frac{1}{3}]})$ to get

$$\begin{aligned} \mathbb{P}\left(z_t \leq \frac{m}{5}, z_{t-\frac{1}{3}} \geq \frac{m}{8}\right) &= \sum_{k_1, k_2, k_3=0}^{\infty} \sum_{l \geq \frac{m}{8}} \mathbb{P}\left(z_t \leq \frac{m}{5}, z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right) \\ &= \sum_{k_1, k_2, k_3=\lfloor m/3-m^{2/3} \rfloor}^{\lfloor m/3+m^{2/3} \rfloor} \sum_{l \geq \frac{m}{8}} \mathbb{P}\left(z_t \leq \frac{m}{5}, z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right) \\ &+ \sum_{\substack{(k_1, k_2, k_3) \in \mathbb{N}^3: \\ (k_1, k_2, k_3) \notin [\lfloor m/3-m^{2/3} \rfloor, \lfloor m/3+m^{2/3} \rfloor]^3}} \sum_{l \geq \frac{m}{8}} \mathbb{P}\left(z_t \leq \frac{m}{5}, z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right). \end{aligned} \quad (16)$$

The second summand in the above sum can be bounded by

$$\begin{aligned} &\sum_{\substack{(k_1, k_2, k_3) \in \mathbb{N}^3: \\ (k_1, k_2, k_3) \notin [\lfloor m/3-m^{2/3} \rfloor, \lfloor m/3+m^{2/3} \rfloor]^3}} \sum_{l \geq \frac{m}{8}} \mathbb{P}\left(z_t \leq \frac{m}{5}, z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right) \\ &\leq \sum_{\substack{(k_1, k_2, k_3) \in \mathbb{N}^3: \\ (k_1, k_2, k_3) \notin [\lfloor m/3-m^{2/3} \rfloor, \lfloor m/3+m^{2/3} \rfloor]^3}} \mathbb{P}(a = (k_1, k_2, k_3)) \\ &\leq \sum_{\substack{k \in \mathbb{N}: \\ k \notin [\lfloor m/3-m^{2/3} \rfloor, \lfloor m/3+m^{2/3} \rfloor]}} \mathbb{P}(a_1 = k) + \mathbb{P}(a_2 = k) + \mathbb{P}(a_3 = k) \\ &= 3 \sum_{\substack{k \in \mathbb{N}: \\ k \notin [\lfloor m/3-m^{2/3} \rfloor, \lfloor m/3+m^{2/3} \rfloor]}} \mathbb{P}(a_1 = k) \stackrel{\text{Lemma C.2}}{\leq} 6e^{-m^{1/3}} \leq m^{-2}. \end{aligned}$$

The equality follows because a_1, a_2, a_3 are identically distributed, and the last two inequalities hold for large enough m . Inserting this into (16), we see that

$$\mathbb{P}\left(z_t \leq \frac{m}{5}, z_{t-\frac{1}{3}} \geq \frac{m}{8}\right) \leq \sum_{\substack{k_1, k_2, k_3 \\ = \lfloor m/3 - m^{2/3} \rfloor}}^{\lfloor m/3 + m^{2/3} \rfloor} \sum_{l \geq \frac{m}{8}} \mathbb{P}\left(z_t \leq \frac{m}{5}, z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right) + m^{-2}$$

for all large enough m . Our next goal is to show that

$$\mathbb{P}\left(z_t \leq \frac{m}{5}, z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right) \leq m^{-2} \mathbb{P}\left(z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right), \quad (17)$$

as this already implies that

$$\begin{aligned} \mathbb{P}\left(z_t \leq \frac{m}{5}, z_{t-\frac{1}{3}} \geq \frac{m}{8}\right) &\leq \sum_{\substack{k_1, k_2, k_3 \\ = \lfloor m/3 - m^{2/3} \rfloor}}^{\lfloor m/3 + m^{2/3} \rfloor} \sum_{l \geq \frac{m}{8}} \mathbb{P}\left(z_t \leq \frac{m}{5}, z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right) + m^{-2} \\ &\leq \sum_{\substack{k_1, k_2, k_3 \\ = \lfloor m/3 - m^{2/3} \rfloor}}^{\lfloor m/3 + m^{2/3} \rfloor} \sum_{l \geq \frac{m}{8}} m^{-2} \mathbb{P}\left(z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right) + m^{-2} \leq 2m^{-2}. \end{aligned}$$

By the law of conditional probability, it suffices to show

$$\mathbb{P}\left(z_t \leq \frac{m}{5} \mid z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right) \leq m^{-2} \quad (18)$$

for all $k_1, k_2, k_3 \in [\lfloor m/3 - m^{2/3} \rfloor, \lfloor m/3 + m^{2/3} \rfloor]$ and $l \geq \frac{m}{8}$, because this already implies (17).

For $i \in \{1, 2, 3\}$, let $v_1^i, v_2^i, \dots, v_{k_i}^i$ be the k_i many agents that arrive during the time interval $[t - \frac{5-i}{3}, t - \frac{4-i}{3}]$. We define the random variable K_1 by

$$K_1 = \sum_{j=1}^{k_1} \mathbf{1}_{\{v_j^1 \text{ is in the pool at time } t - \frac{1}{3}\}},$$

where $\mathbf{1}$ denotes the indicator function. From our conditions, we already know that $K_1 \leq k_1$.

In the next step, we make use of the following observation: For each agent x that arrives during the interval $[t - \frac{4}{3}, t - 1)$ and each agent y that arrives during the time interval $[t - 1, t - \frac{1}{3})$, the chances of agent y of making it to the pool at time $t - \frac{1}{3}$ are at least as high as for agent x , as the agent x witnesses at least as many critical agents as y . Moreover, since $k_1, k_2, k_3 \in [\lfloor \frac{m}{3} - m^{2/3} \rfloor, \lfloor \frac{m}{3} + m^{2/3} \rfloor]$, it holds that $\frac{k_1}{k_1 + k_2 + k_3} \leq \frac{2}{5}$ for m large enough. We will assume that $\frac{k_1}{k_1 + k_2 + k_3} \leq \frac{2}{5}$ henceforth. Thus, we can infer that

$$\mathbb{E}\left[K_1 \mid z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right] \leq l \frac{k_1}{k_1 + k_2 + k_3} \leq \frac{2}{5}l.$$

We want to translate this bound on K_1 , that holds in expectation, to a bound that holds with high probability, i.e., our next goal is to show that $K_1 < l/2$ with high probability. The distribution of K_1 highly depends on the arrival/departure-process, as agents that arrive at a later point in time typically witness less critical agents up to time $t - \frac{1}{3}$. We expect K_1 to be maximal, when all agents getting critical in the time frame $[t - \frac{4}{3}, t - \frac{1}{3})$ get critical just before time step $t - \frac{1}{3}$, as all agents arriving in the time frame $[t - \frac{4}{3}, t - \frac{1}{3})$ have the same chance of making it to the pool $Z_{t-\frac{1}{3}}$ in this case. In all other cases, agents that arrive earlier always have a lower chance of making it to the pool. Thus, we can bound the probability of $K_1 > \frac{l}{2}$ with the outcome of a different experiment, namely an urn experiment, that we describe now.

Let $n = k_2 + k_3$ and consider an urn with $N = k_1 + k_2 + k_3 = k_1 + n$ balls in it, where k_1 of them are red and $k_2 + k_3$ of them are blue. We draw $l \leq N$ balls out of this urn without replacement. Let \tilde{K}_1 be the number of red balls taken. We defer the technical proof of the next steps to two lemmas shown after the proof of the theorem. By Lemma B.1, the random variable \tilde{K}_1 stochastically dominates K_1 , so we need to bound the probability that $\tilde{K}_1 \geq l/2$. By Lemma B.2, it holds that $\mathbb{P}\left(\tilde{K}_1 \geq \frac{l}{2}\right) \leq 2m0.98^{m/8}$. Consequently,

$$\mathbb{P}\left(K_1 \geq \frac{l}{2}\right) \leq \mathbb{P}\left(\tilde{K}_1 \geq \frac{l}{2}\right) \leq 2m0.98^{m/8} \leq 0.998^m \leq m^{-3},$$

where the last two inequalities hold for m large enough. If $K_1 \leq \frac{l}{2}$, then $z_t \geq z_{t-\frac{1}{3}} - 2K_1 + a_{(t-\frac{1}{3}, t)} \geq a_{(t-\frac{1}{3}, t)}$. So, we can conclude in particular that

$$\begin{aligned} & \mathbb{P}\left(z_t \leq \frac{m}{5} \mid z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right) \\ & \leq \mathbb{P}\left(K_1 > \frac{l}{2} \mid z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right) + \mathbb{P}\left(a_{(t-\frac{1}{3}, t)} < \frac{m}{5} \mid z_{t-\frac{1}{3}} = l, a = (k_1, k_2, k_3)\right) \\ & \stackrel{\text{Lemma C.2}}{\leq} m^{-3} + \mathbb{P}\left(a_{(t-\frac{1}{3}, t)} < \frac{m}{5}\right) \leq m^{-2} \end{aligned}$$

for m large enough, which finally shows (18). For the second inequality we used that $a_{(t-\frac{1}{3}, t)}$ is independent of $z_{t-\frac{1}{3}}$ and $a = (a_1, a_2, a_3)$. Together with (14), this already implies that

$$\mathbb{P}\left(z_t \leq \frac{m}{5}\right) \leq m^{-1}$$

for m large enough. With this result we get that

$$\begin{aligned} & \mathbb{P}(\text{an agent arriving at time } t-1 \text{ perishes}) \\ & \leq \mathbb{P}(\text{an agent getting critical at time } t \text{ perishes}) \\ & = \mathbb{P}\left(\text{an agent getting critical at time } t \text{ perishes} \mid z_t \geq \frac{m}{5}\right) \mathbb{P}\left(z_t \geq \frac{m}{5}\right) \\ & + \mathbb{P}\left(\text{an agent getting critical at time } t \text{ perishes} \mid z_t < \frac{m}{5}\right) \mathbb{P}\left(z_t < \frac{m}{5}\right) \\ & \leq \left(1 - \frac{d}{m}\right)^{\frac{m}{5}-1} + m^{-1} \xrightarrow{m \rightarrow \infty} e^{-\frac{d}{5}}. \end{aligned}$$

□

We conclude the proof of Theorem 4.6 by proving the two auxiliary lemmas.

Lemma B.1. *In the notation of the proof of Theorem 4.6, \tilde{K}_1 stochastically dominates K_1 .*

Proof. We first describe three different experiments: In the first one, we have a pool $Z_{t-\frac{4}{3}}$ of some size. We call the agents in the pool at time $t - \frac{4}{3}$ the uncolored *black agents*. These agents get critical exactly one time step after their arrival, so they will not be in the pool at time $t - \frac{1}{3}$ anymore. The k_1 many agents arriving in the interval $[t - \frac{4}{3}, t - 1)$ are called the *red agents* and the $n = k_2 + k_3$ many agents arriving in the interval $[t - 1, t - \frac{1}{3})$ are called the *blue agents*. Then agents can depart from the pool by perishing or by matching, so that l colored, i.e., red or blue, agents are in the pool at time $t - \frac{1}{3}$. We are interested in the number K_1 , which is the number of red agents in the pool at time $[t - \frac{1}{3}]$.

In the second experiment, we have an urn with k_1 red balls and n blue balls and we take $k_1 + n - l$ balls out of the urn, uniformly at random without replacement. We are interested in the number of red balls K'_1 that are left in the urn after this procedure.

In the third experiment, we have an urn with k_1 red balls and n blue balls and we take l balls out of the urn, uniformly at random without replacement. We are interested in the number of red balls \tilde{K}_1 that are taken out of the urn.

Leaving l balls in an urn or taking l balls out of an urn gives rise to the exact same distribution. Thus it suffices to show that K'_1 stochastically dominates K_1 . For this, we couple the urn to the pool in the following way: At time $t - \frac{4}{3}$, we start with the initial urn, i.e., the one with k_1 red balls and n blue balls in it. Let $(R_s)_{s \in [t - \frac{4}{3}, t - \frac{1}{3}]}$, respectively $(B_s)_{s \in [t - \frac{4}{3}, t - \frac{1}{3}]}$, be the number of red, respectively blue, agents in the pool at time s . Analogously, let $(R'_s)_{s \in [t - \frac{4}{3}, t - \frac{1}{3}]}$, respectively $(B'_s)_{s \in [t - \frac{4}{3}, t - \frac{1}{3}]}$, be the number of red, respectively blue, balls in the urn at time s . The underlying algorithm of the matching process in the first experiment does not distinguish between red and blue agents. So whenever an agent gets critical and matches with a colored agent, only the proportion of red/blue agents matters. So we can think of the process of the pool as follows: When an agent gets critical, say at time t_0 , and matches with a colored agent, there is an underlying random variable $U_{t_0} \sim \text{Uniform}[0, 1]$ that determines whether the agent matches with a blue or with a red agent. When $U_{t_0} \leq \frac{R_{t_0}}{R_{t_0} + B_{t_0}}$, then we choose uniformly at random one of the R_{t_0} many red agents. Otherwise we choose one of the B_{t_0} many blue agents uniformly at random. Whenever there is a matching with a colored agent in the pool, we also remove a ball from the urn. We pick a red ball if $U_{t_0} \leq \frac{R'_{t_0}}{R'_{t_0} + B'_{t_0}}$, and otherwise we pick a blue ball. As U_{t_0} is the uniform distribution, this also gives rise to the correct probabilities. Note that the proportion of red agents among the colored agents in the pool can change either when agents depart by matching or when a new agent arrives. Contrary, the proportion of red balls in the urn changes only at times when colored agents match in the pool. Let S_s , respectively S'_s be the number of red agents, respectively balls, that are removed from the pool, respectively urn, by time s . We claim that

$$S_s \geq S'_s \text{ for all } s \in \left[t - \frac{4}{3}, t - \frac{1}{3} \right]. \quad (19)$$

This is clear for $s < t - 1$, as S_s is simply the number of removed colored agents, as no blue agent entered the pool yet. Thus we can assume $s \geq t - 1$ from here on. Whenever $S_s > S'_s$ just before some time t_0 , at which an agent gets critical and matches with a colored agent, then $S_s \geq S'_s$ still holds after this match. If $S_s = S'_s$ and $s \geq t - 1$, then all red agents already arrived, which gives $R_s = k_1 - S_s = R'_s$. As the number of colored balls in the urn is at least as high as the number of colored agents in the pool, we have

$$\frac{R'_s}{R'_s + B'_s} \leq \frac{R_s}{R_s + B_s}.$$

Assume that $S_s = S'_s$ just before a match with a colored agent in the pool occurs. If a red ball is drawn from the urn, i.e., $U_s \leq \frac{R'_s}{R'_s + B'_s}$, then also a red agent leaves the pool in a matching. Thus $S_s \geq S'_s$ still holds directly after the match. So we also get that $S_{t - \frac{1}{3}} \geq S'_{t - \frac{1}{3}}$ which directly gives

$$K'_1 = R'_{t - \frac{1}{3}} = k_1 - S'_{t - \frac{1}{3}} \geq k_1 - S_{t - \frac{1}{3}} = R_{t - \frac{1}{3}}$$

and thus finishes the proof. □

Lemma B.2. *In the notation of the proof of Theorem 4.6, $\mathbb{P}\left(\tilde{K}_1 \geq \frac{l}{2}\right) \leq 2m 0.98^{m/8}$.*

Proof. Consider a set $U \subseteq \{1, \dots, l\}$. Then, the probability of drawing a red ball at the i -th draw if and only if $i \in U$ is given by

$$\prod_{i=0}^{|U|-1} \frac{k_1 - i}{N - i} \prod_{i=0}^{l-|U|-1} \frac{n - i}{N - |U| - i} = \frac{1}{\prod_{i=0}^{l-1} N - i} \prod_{i=0}^{|U|-1} (k_1 - i) \prod_{i=0}^{l-|U|-1} (n - i).$$

Among all $U \subseteq \{1, \dots, l\}$ with $|U| \geq \frac{l}{2}$, this is clearly maximized when $|U| = \lceil \frac{l}{2} \rceil$, as $n \geq k_1$. Let us assume that l is even from here on. The proof for l odd works completely similar, just with one additional term. So we have for all l even and all $U \subseteq \{1, \dots, l\}$ with $|U| \geq \frac{l}{2}$ the bound

$$\begin{aligned} & \frac{1}{\prod_{i=0}^{l-1} N - i} \prod_{i=0}^{|U|-1} (k_1 - i) \prod_{i=0}^{l-|U|-1} (n - i) \leq \frac{1}{\prod_{i=0}^{l-1} N - i} \prod_{i=0}^{l/2-1} (k_1 - i) \prod_{i=0}^{l/2-1} (n - i) \\ & = \prod_{i=0}^{l/2-1} \frac{(k_1 - i)(n - i)}{(N - 2i)(N - 2i - 1)} = \prod_{i=0}^{l/2-1} \frac{(k_1 - i)(n - i)}{(N - 2i)(N - 2i)} \prod_{i=0}^{l/2-1} \frac{N - 2i}{N - 2i - 1}. \end{aligned} \quad (20)$$

The second factor in (20) can be bounded by

$$\prod_{i=0}^{l/2-1} \frac{N - 2i}{N - 2i - 1} \leq \prod_{i=0}^{l/2-1} \frac{N - 2i + 1}{N - 2i - 1} = \frac{N + 1}{N - l + 1} \leq N + 1.$$

In order to bound the components of the first factor, first notice that

$$\frac{k_1 - i}{N - 2i} + \frac{n - i}{N - 2i} = \frac{N - 2i}{N - 2i} = 1,$$

so we multiply two numbers that sum to one. Furthermore, $\frac{k_1 - i}{N - 2i} \leq \frac{2N/5 - i}{N - 2i} = \frac{2}{5} \frac{N - 5i/2}{N - 2i} \leq \frac{2}{5}$. Recall that we assumed that $\frac{k_1}{N} \leq \frac{2}{5}$, which is used in the first inequality. The product of $z(1 - z)$ with $z \in [0, \frac{2}{5}]$ is maximized for $z = \frac{2}{5}$. Thus we already have

$$\prod_{i=0}^{l/2-1} \frac{(k_1 - i)(n - i)}{(N - 2i)(N - 2i)} \leq \left(\frac{2}{5} \cdot \frac{3}{5} \right)^{l/2} \leq 0.24^{l/2}.$$

Inserting this into (20) we get that for all $U \subseteq \{1, \dots, l\}$ with $|U| \geq l/2$

$$\mathbb{P}(\text{The } i\text{-th draw is red if and only if } i \in U) \leq (N + 1)0.24^{l/2}.$$

Recall that we have $l \geq m/8$ and, due $k_i \leq m/3 + m^{2/3}$ for $i \in \{1, 2, 3\}$, $N + 1 = k_1 + k_2 + k_3 + 1 \leq 2m$ for m large enough. By a union bound, we thus get

$$\begin{aligned} \mathbb{P}\left(\tilde{K}_1 \geq \frac{l}{2}\right) &= \sum_{\substack{U \subseteq \{1, \dots, l\}: \\ |U| \geq l/2}} \mathbb{P}(\text{The } i\text{-th draw is red if and only if } i \in U) \leq 2^l (N + 1) 0.24^{l/2} \\ &\leq 2m \left(2\sqrt{0.24}\right)^l \leq 2m 0.98^l \leq 2m 0.98^{m/8}. \end{aligned}$$

□

C Auxiliary Statements

In this section we give several auxiliary statements that we used above.

The first lemma is an auxiliary inequality that follows from the inequality $1 + x \leq e^x$ with $x = -\frac{c}{m}$.

Lemma C.1. *Let $c \geq 0$ and $m > c$. Then,*

$$\left(1 - \frac{c}{m}\right)^m \leq e^{-c}. \quad (21)$$

An application of Chernoff's inequality [see, e.g., Hagerup and Rüb, 1990] yields the next result.

Lemma C.2. *Let X be the sum of independent Bernoulli random variables or a Poisson random variable with expectation value μ and let $0 < \delta \leq 1$. Then,*

$$\mathbb{P}(X \geq (1 + \delta)\mu) \leq e^{-\mu\delta^2/3}, \quad (22)$$

$$\mathbb{P}(X \leq (1 - \delta)\mu) \leq e^{-\mu\delta^2/3}. \quad (23)$$

The next lemma is a typical application of Markov chains to analyze random walks with drift, and its proof uses a classical technique.

Lemma C.3. *Let $0 < \varepsilon < \frac{1}{2}$ and let $(X_n)_{n \in \mathbb{N}}$ be an irreducible Markov chain with state space \mathbb{N} , that makes jumps between nearest neighbours only and with transition probabilities satisfying $p(z, z+1) \geq \frac{1}{2} - \varepsilon$ for all $z \geq M$. When the Markov chain starts at $X_0 = M$, the probability of hitting some integer $N > M$ before going to $M - 1$ is bounded by $(1 - 2\varepsilon)^{N-M}$.*

Proof. We provide a proof for the case $M = 1$. The proof for general M works completely analogous.

We define the function

$$f : \mathbb{N} \rightarrow \mathbb{R}, \quad n \mapsto \sum_{i=0}^{n-1} \prod_{j=1}^i \frac{p(j, j-1)}{p(j, j+1)},$$

i.e.,

$$f(0) = 0, f(1) = 1, \text{ and } f(n+1) = f(n) + \prod_{j=1}^n \frac{p(j, j-1)}{p(j, j+1)} \text{ for } n \geq 1.$$

The function f is harmonic on $\mathbb{N}_{>0}$ for the transition probabilities $(p(k, l))_{k, l \in \mathbb{N}}$, as

$$\begin{aligned} & p(n, n+1)f(n+1) + p(n, n-1)f(n-1) \\ &= p(n, n+1) \left(f(n) + \prod_{j=1}^n \frac{p(j, j-1)}{p(j, j+1)} \right) + p(n, n-1) \left(f(n) - \prod_{j=1}^{n-1} \frac{p(j, j-1)}{p(j, j+1)} \right) \\ &= f(n) + p(n, n+1) \prod_{j=1}^n \frac{p(j, j-1)}{p(j, j+1)} - p(n, n-1) \prod_{j=1}^{n-1} \frac{p(j, j-1)}{p(j, j+1)} = f(n) \end{aligned}$$

for all $n \geq 1$. For an integer $t \in \mathbb{Z}$, define the stopping time $\tau_t = \min\{n : X_n = t\}$, and define $\tau = \tau_0 \wedge \tau_N = \min\{n : X_n \in \{0, N\}\}$. Assume that the Markov chain starts at $X_0 = 1$. Then the process $f(X_{n \wedge \tau})$ is a martingale with respect to the natural filtration. As the state space is finite, this martingale is also uniformly integrable. By the optional sampling theorem [Klenke, 2013, Theorem 10.21], we have

$$\begin{aligned} 1 &= \mathbb{E}[f(X_\tau) | X_0 = 1] = f(0) \cdot \mathbb{P}(X_\tau = 0 | X_0 = 1) + f(N) \cdot \mathbb{P}(X_\tau = N | X_0 = 1) \\ &= f(N) \cdot \mathbb{P}(\tau_N < \tau_0 | X_0 = 1). \end{aligned} \quad (24)$$

Next, we bound $f(N)$ by

$$\begin{aligned}
f(N) &= \sum_{i=0}^{n-1} \prod_{j=1}^i \frac{p(j, j-1)}{p(j, j+1)} \geq \sum_{i=0}^{n-1} \prod_{j=1}^i \frac{\frac{1}{2} + \varepsilon}{\frac{1}{2} - \varepsilon} = \sum_{i=0}^{n-1} \left(\frac{1+2\varepsilon}{1-2\varepsilon} \right)^i = \frac{\left(\frac{1+2\varepsilon}{1-2\varepsilon} \right)^N - 1}{\frac{1+2\varepsilon}{1-2\varepsilon} - 1} \\
&= \frac{\left(\frac{1+2\varepsilon}{1-2\varepsilon} \right)^N - 1}{\frac{4\varepsilon}{1-2\varepsilon}} = \frac{(1+2\varepsilon)^N - (1-2\varepsilon)^N}{4\varepsilon(1-2\varepsilon)^{N-1}} \geq \frac{4\varepsilon}{4\varepsilon(1-2\varepsilon)^{N-1}} = \frac{1}{(1-2\varepsilon)^{N-1}}.
\end{aligned}$$

Inserting this inequality into (24) and solving for $\mathbb{P}(\tau_N < \tau_0 | X_0 = 1)$ gives that

$$\mathbb{P}(\tau_N < \tau_0 | X_0 = 1) \leq (1-2\varepsilon)^{N-1}.$$

□