Thickness and Information in Dynamic Matching Markets[*][†]

Mohammad Akbarpour[‡] Shengwu Li[§] Shayan Oveis Gharan[¶]

First Draft: February 2014 This Draft: January 2016

Abstract

We introduce a simple model of dynamic matching in networked markets, where agents arrive and depart stochastically, and the composition of the trade network depends endogenously on the matching algorithm. We show that if the planner can identify agents who are about to depart, then waiting to thicken the market is highly valuable, and if the planner cannot identify such agents, then matching agents greedily is close to optimal. We characterize the optimal waiting time (in a restricted class of mechanisms) as a function of waiting costs and network sparsity. The planner's decision problem in our model involves a combinatorially complex state space. However, we show that simple local algorithms that choose the right time to match agents, but do not exploit the global network structure, can perform close to complex optimal algorithms. Finally, we consider a setting where agents have private information about their departure times, and design a continuous-time dynamic mechanism to elicit this information.

Keywords: Market Design, Matching, Networks, Continuous-time Markov Chains, Mechanism Design

JEL Classification Numbers: D47, C78, C60

---

[*]This paper has been previously distributed with the title "Dynamic Matching Market Design."

[†]We thank Paul Milgrom and Alvin Roth for valuable guidance. We also thank Eric Budish, Ben Golub, Scott Kominers, and Alex Wolitzky for valuable comments. We are also grateful to Daron Acemoglu, Nikhil Agarwal, Nick Arnosti, Itai Ashlagi, Peter Biro, Timothy Bresnahan, Jeremy Bulow, Gabriel Carroll, Sylvain Chassang, Constantinos Daskalakis, Darrell Duffie, Alex Frankel, Matthew Jackson, Fuhito Kojima, Soohyung Lee, Jacob Leshno, Malwina Luczak, Joshua Mollner, Markus Mobius, Roger Myerson, Stephen Nei, Muriel Niederle, Afshin Nikzad, Michael Ostrovsky, Parag Pathak, Robert Shimer, Takuo Sugaya, and Bob Wilson for their valuable discussions, as well as several seminar participants for helpful suggestions. All errors remain our own.

[‡]Becker-Friedman Institute, University of Chicago. Email: akbarpour@uchicago.edu

[§]Department of Economics, Stanford University. Email: shengwu@stanford.edu

[¶]Department of Computer Science and Engineering, University of Washington. Supported by a Miller fellowship. Email: shayan@cs.washington.edu

Contents

1 Introduction

The theory of matching has guided the design of many markets, from school choice, to kidney exchange, to the allocation of medical residents. In a series of classic papers, economists have extensively characterized good matching algorithms for static settings.[1] In the canonical set-up, a social planner faces a set of agents who have preferences over partners, contracts, or combinations thereof. The planner's goal is to find a matching algorithm with desirable properties (*e.g.* stability, efficiency, or strategy-proofness). The algorithm is run, a match is made, and the problem ends.

Of course, many real-world matching problems are dynamic. In a dynamic matching environment, agents arrive gradually over time. A social planner continually observes the agents and their preferences, and chooses how to match agents. Matched agents leave the market. Unmatched agents either persist or depart. Thus, the planner's decision today affects the sets of agents and options tomorrow.

Some seasonal markets, such as school choice systems and the National Resident Matching Program, are well-described as static matching problems without intertemporal spillovers. However, some markets are better described as dynamic matching problems. Some examples include:

- **Kidney exchange:** In paired kidney exchanges, patient-donor pairs arrive over time. They stay in the market until either they are matched to a compatible pair, or their condition deteriorates so that they leave the market unmatched.
- **Markets with brokers:** Some markets, such as real estate, aircraft, and ship charters, involve intermediary brokers who receive requests to buy or sell particular items. A broker facilitates transactions between compatible buyers and sellers, but does not hold inventory. Agents may withdraw their request if they find an alternative transaction.
- **Allocation of workers to time-sensitive tasks:** Both within firms and online labor markets, such as Uber and oDesk, planners allocate workers to tasks that are profitable to undertake. Tasks arrive continuously, but may expire. Workers are suited to different tasks, but may cease to be available.

[1]See Gale and Shapley (1962); Crawford and Knoer (1981); Kelso Jr and Crawford (1982); Roth and Sotomayor (1992); Abdulkadiroglu and Sonmez (2003); Roth et al. (2004); Hatfield and Milgrom (2005); Abdulkadiroglu et al. (2005a,b); Roth et al. (2007); Hatfield and Kojima (2010); Hatfield et al (2013); Kominers and Sonmez (2013).

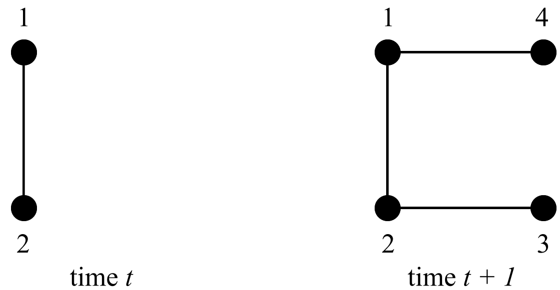

Figure 1: Waiting expands the set of options. If the planner matches agent 1 to agent 2 at time t, then the planner will be unable to match agents 3 and 4 at time $t+1$.

In dynamic settings, the planner must decide not only *which* agents to match, but also *when* to match them. In choosing when to match agents, the planner faces a trade-off between matching agents quickly and waiting to thicken the market. If the planner matches agents frequently, then matched agents will not have long to wait, but it will be less likely that any remaining agent has a potential match (a thin market). On the other hand, if the planner matches agents infrequently, then there will be more agents available, making it more likely that any given agent has a potential match (a thick market). For example, in Figure 1, where each node represents an agent and each link represents a potential 'match', if the planner matches agent 1 to agent 2 at time t, then the planner will be unable to match agents 3 and 4 at time $t+1$. By contrast, if the planner waits until $t+1$, he can match all four agents by matching 1 to 4 and 2 to 3.

The problem is more subtle if agents depart stochastically. Then, one drawback of waiting is that agents may depart. However, waiting might bring information about which agents will soon depart, enabling the planner to give priority to those agents. For example, in Figure 2, the planner learns at $t+1$ that agent 3 will imminently leave the market. If the planner matches agent 1 to agent 2 at time t, then he will be unable to react to this information at $t+1$.

This paper deals with identifying features of optimal matching algorithms in dynamic environments. The previous examples illustrate that static matching models do not capture important features of dynamic matching markets. In a static setting, the planner chooses the best algorithm for an exogenously given set of agents and their preferences. By contrast, in a dynamic setting, the set of agents and trade options at each point in time depend endogenously on the matching algorithm. Only in a dynamic framework can we study the trade-off between the option value of waiting and the potential costs of waiting.

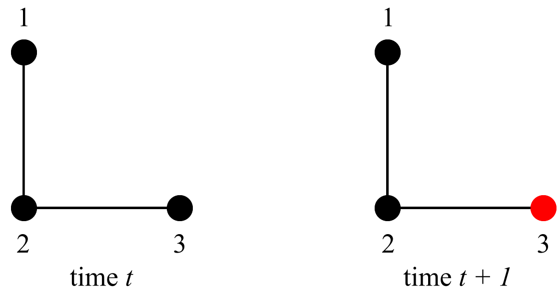

Figure 2: Waiting resolves uncertainties about agents' departure times. The planner will learn at $t+1$ that agent 3 will depart imminently, and waiting allows him to use this information.

The optimal timing policy in a dynamic matching problem is not obvious *a priori*. In practice, many paired kidney exchanges enact static matching algorithms ('match-runs') at fixed intervals.[2] Even then, matching intervals differ substantially between exchanges: The Alliance for Paired Donation conducts a match-run once a weekday, the United Network for Organ Sharing conducts a match-run once a week[3], the South Korean kidney exchange conducts a match-run once a month, and the Dutch kidney exchange conducts a match-run once a quarter (Akkina et al., 2011). This shows that policymakers select different timing policies when faced with seemingly similar dynamic matching problems. It is therefore useful to identify good timing policies, and to investigate how policy should depend on the underlying features of the problem.

In this paper, we create and analyze a stylized model of dynamic matching on networks, motivated mainly by the problem of kidney exchange. Agents arrive and depart stochastically. We use binary preferences, where a pairwise match is either acceptable or unacceptable, generated according to a known distribution. These preferences are persistent over time, and agents may discount the future. The set of agents (vertices) and the set of potential matches (edges) form a random graph. Agents do not observe the set of acceptable transactions, and are reliant upon the planner to match them to each other. We say that an agent *perishes* if she leaves the market unmatched.

The planner's problem is to design a matching algorithm; that is, at any point in time, to select a subset of acceptable transactions and broker those trades. The planner observes the current set of agents and acceptable transactions, but has only probabilistic knowledge

[2]In graph theory, a *matching* is a set of edges that have no nodes in common.

[3]See http://www.unos.org/docs/Update_MarchApril13.pdf

about the future. The planner may have knowledge about which agents' needs are *urgent*, in the sense that he may know which agents will perish imminently if not matched. The goal of the planner is to maximize the sum of the discounted utilities of all agents. In the important special case where the cost of waiting is zero, the planner's goal is equivalent to minimizing the proportion of agents who perish. We call this the *loss* of an algorithm.

We can interpret this model as a stylized representation of paired kidney exchange. Vertices in the graph represent patient-donor *pairs*, and there is an edge between two vertices if the donor in each pair is compatible with the patient in the other pair. In kidney exchanges with highly-sensitized patients, the feasibility of a match depends on both blood-type and tissue-type compatibility. Tissue-type, or human leukocyte antigens (HLA), is a combination of the level of six proteins, which implies a very large type space for patient-donor pairs. To make the model tractable, we make the simplifying assumption that any patient-donor pair is compatible with any other patient-donor pair with some independent probability. This interpretation abstracts from k-cycles, for $k > 2$.

What are the key features of the optimal matching algorithm? Since we explicitly model the network of potential matches, the resulting Markov Decision Problem is combinatorially complex. Thus, it is not feasible to compute the optimal solution with standard dynamic programming techniques. Instead, we employ a different approach: First, we derive bounds on optimum performance. Then, we formulate simple and tractable matching algorithms with different properties, and compare them to those bounds. If a simple algorithm is close to optimal, then that suggests that it has the essential features of the optimal matching algorithm.

The simple algorithms are as follows: The *Greedy* algorithm attempts to match each agent upon arrival; it treats each instant as a static matching problem without regard for the future.[4] The *Patient* algorithm attempts to match only urgent agents (potentially to a non-urgent partner). Both these algorithms are *local*, in the sense that they look only at the immediate neighbors of the agent they attempt to match rather than at the global graph structure[5]. We also study the *Patient*(α) algorithm, a family of algorithms that speeds up the trade frequency of the Patient algorithm. This algorithm attempts to match urgent cases, and additionally attempts to match each non-urgent case at some rate determined by

[4] The Greedy policy is inspired by some real-world markets, such as the Alliance for Paired Donation's policy of conducting a 'match-run' everyday. Our analysis of the Greedy Algorithm encompasses waiting list policies where brokers make transactions as soon as they are available, giving priority to agents who arrived earlier.

[5] Example 3.5 shows a case in which 'locality' of the Patient algorithm makes it suboptimal.

α[6].

We have three families of results. First, we analyze the performance of algorithms with different timing properties, in the benchmark setting where the planner can identify urgent cases. Second, we relax our informational assumption, and thereby establish the value of short-horizon information about urgent cases. Third, we exhibit a dynamic mechanism that truthfully elicits such information from agents. To state our results, we define the *density parameter* of the model, d. For each agent, d is the average number of new acceptable partners that arrive at the market before her needs are urgent.[7]

Our first family of results concerns the value of thickness in dynamic matching markets.

First, we establish that the loss of the Patient algorithm is exponentially smaller than the loss of the Greedy algorithm, where the exponent is in d. For example, in a market where $d = 8$, the loss of the Patient algorithm is no more than 15% of the loss of the Greedy algorithm. This entails that if waiting costs are negligible, for a wide range of parameters, the Patient algorithm substantially outperforms the Greedy algorithm. This shows that the upper bound on the value of waiting is large.

Thickness in matching markets is analogous to *liquidity* in financial markets. Many financial models imply that liquidity is desirable because it facilitates time-sensitive transactions. Our model shows that this intuition extends even to markets that do not have prices, where agents have heterogeneous preferences over trading partners. Since the Greedy algorithm conducts matches as soon as they become available, the market that remains is very thin (illiquid). Consequently, it is not feasible to match urgent cases. By contrast, the Patient algorithm maintains a thick (liquid) market. Thus, the planner can match urgent cases with high probability.

The second finding in our first family of results establishes that the loss of the Patient algorithm is "close to" the loss of the optimum algorithm. Recall that the Patient algorithm is local; it looks only at the immediate neighborhood of the agents it seeks to match. By

[6]More precisely, every non-urgent agent is treated as urgent when an exogenous "exponential clock" ticks and attempted to be matched either in that instance, or when she becomes truly urgent.

[7]Some approximate, back of the envelope calculations on the data of national kidney registry are helpful in having a reasonable estimate for the range of d. In the national kidney registry data NKR Quarterly Report (2013), the arrival rate of pairs (in 2013) was 467 pairs per year. Patients' average waiting time, if they are not matched, is nearly 3 years (Sonmez and Ünver (2015)). So $m \simeq 1400$. In addition, p is the probability that two random pairs are cross compatible, which is $\bar{p}^2$, where $\bar{p}$ is the probability that a random patient is compatible to a random donor. $\bar{p}$ is a function of the patients' ethnicity, age, sex, and many other factors, so it is hard to have an exact estimate, but given PRA and blood-type distribution data, $\bar{p}$ can be estimated to be a number between 0.05 and 0.1, implying $\bar{p}^2 = p \in [.0025, .01]$ and $d \in [3.5, 14]$. For the purpose of our results, it is worthwhile to remember the magnitude of d.

contrast, the optimum algorithm chooses the optimal time to match agents, as well as the optimal agents to match: It is global and potentially very complex, since the matchings it selects depend on the entire network structure. When we compare the performance of the (practically popular) Greedy algorithm to the optimum algorithm, we find that most of the gain is achieved merely by being patient and thickening the market, rather than optimizing over the network structure.[8]

Third, we find that it is possible to accelerate the Patient algorithm and still achieve exponentially small loss[9]. That is, we establish a bound for the tuning parameter α such that the Patient(α) algorithm has exponentially small loss. For instance, if agents stay in the market for 1 year on average before getting urgent, and if $d = 8$, then under the Patient(α) algorithm, the planner can promise to match agents in less than 4 months (in expectation) while the loss is at most 55% of the loss of the Greedy algorithm. Thus, even moderate degrees of waiting can substantially reduce the proportion of perished agents.

Fourth, we examine welfare under discounting. We show that for a range of discount rates, the Patient algorithm delivers higher welfare than the Greedy algorithm, and for a wider range of discount rates, there exists α such that the Patient(α) algorithm delivers higher welfare than the Greedy algorithm. Then, in order to capture the trade-off between the trade frequency and the thickness of the market, we solve for the optimal waiting time as a function of the market parameters. Our comparative statics show that the optimal waiting time is increasing in the sparsity of the graph.

Our second family of results relaxes the informational assumptions in the benchmark model. Suppose that the planner cannot identify urgent cases; i.e. the planner has no individual-specific information about departure times. We find that the loss of the optimum algorithm that lacks such information is exponentially more than the loss of the Patient algorithm, which naïvely exploits urgency information.

On the other hand, suppose that the planner has *more* than short-horizon information about agent departures. The planner may be able to forecast departures long in advance, or foresee how many new agents will arrive, or know that certain agents are more likely than others to have new acceptable transactions. We prove that *no* expansion of the planner's information allows him to achieve a better-than-exponential loss. Taken as a pair, these results suggest that short-horizon information about departure times is especially valuable to the planner. Lacking this information leads to large losses, and having more than this

[8] This result on the near-optimality of a "local" algorithm (Theorem 3.6) also contributes to the literature of local algorithms and local computation mechanism design in theoretical computer science (Suomela, 2013).

[9] As before, the exponent is in the average degree of agents.

information does not yield large gains.

In some settings, however, agents know when their cases are urgent, but the planner does not. For instance, doctors know whether their patients have urgent needs, but kidney exchange pools do not. Our final result concerns the incentive-compatible implementation of the Patient(α) algorithm.[10] Under private information, agents may have incentives to misreport their urgency so as to hasten their match or to increase their probability of getting matched. We show that if agents are not too impatient, a dynamic mechanism without transfers can elicit such information. The mechanism treats agents who report that their need is urgent, but persist, as though they had left the market. This means that as an agent, I trade off the possibility of a swifter match (by declaring that I am in urgent need now) with the option value of being matched to another agent before I truly become urgent. We prove that it is arbitrarily close to optimal for agents to report the truth in large markets.

The rest of the paper is organized as follows. Section 2 introduces our model and defines the objective. Section 3 presents our main contributions; we recommend that readers consult this section to see a formal statement of our results without getting into the details of the proofs. Section 4 models our algorithms as Markov Chains and bounds their mixing times. Section 5 goes through a deep analysis of the Greedy algorithm, the Patient algorithm, and the Patient(α) algorithm and bounds their performance. Section 6 considers the case where the urgency of an agent's need is private information, and exhibits a truthful direct revelation mechanism. Section 7 discusses key assumptions and suggests extensions. Section 8 concludes.

1.1 Related Work

There have been several studies on dynamic matching in the literatures of economics, computer science, and operations research. To the best of our knowledge, no prior work has examined dynamic matching on a general graph, where agents stochastically depart.

Kurino (2009) and Bloch and Houy (2012) study an overlapping generations model of the housing market. In their models, agents have deterministic arrivals and departures. In addition, the housing side of the market is infinitely durable and static, and houses do not have preferences over agents. In the same context, Leshno (2012) studies a one-sided dynamic housing allocation problem in which houses arrive stochastically over time, and the waiting list is overloaded, and characterizes desirable queuing policies while incentive

[10] Note that the Patient(α) algorithm contains the Patient algorithm as a special case.

constraints exist. In a subsequent independent paper, Baccara et al. (2015) study the problem of dynamic matching in a two-sided model with 'high' and 'low' types on each side; every agent on one side is compatible with every agent on the other side. In their setting, the optimal mechanism matches agents assortatively, while accumulating uneven pairs up to a threshold.

In the context of live-donor kidney exchanges, Ünver (2010) studies a model of dynamic kidney exchange in which agents have multiple types. In his model, agents never perish. Thus, one insight of his model is that waiting to thicken the market is not helpful when only bilateral exchanges are allowed. We show that this result does not continue to hold when agents depart stochastically. In a more recent study, Sonmez and Ünver (2015) introduce a continuum model of organ transplantation. Their goal is to study the effects of various transplantation technologies, including a new incentive scheme they introduce to increase the number of compatible pairs. In the Operations Research and Computer Science literatures, dynamic kidney matching has been extensively studied, see e.g., (Zenios, 2002; Su and Zenios, 2005; Awasthi and Sandholm, 2009; Dickerson et al., 2012). Ashlagi et al. (2013) construct a discrete-time finite-horizon model of dynamic kidney exchange. Unlike our model, agents who are in the pool neither perish, nor bear any waiting cost, and so they are not concerned with criticality time, or agents' incentives to misreport departure time's information. Finally, a follow-up paper uses our conceptual framework and techniques to model the competition of two platforms with Greedy and Patient algorithms (Das et al., 2015).

In an independent concurrent work, Anderson et al. (2014) analyze a model in which the main objective is to minimize the average waiting time, and agents never perish. They analyze both pairwise and 3-way cycles and show that when only pairwise exchanges are allowed, the Greedy algorithm is optimal in the class of 'periodic Markov policies', which is similar to Theorem 3.10 in this paper. Our paper, on top of that, shows that when agents' departure times are observable, then Greedy performs weakly, and the option value of waiting can be huge. In another independent concurrent study, Arnosti et al. (2014) model a two-sided dynamic matching market. Their main goal is to analyze congestion in decentralized dynamic markets, whereas we study the role of timing and information in dynamic markets from a central planning perspective.

Some recent papers study the problem of stability in dynamic matching markets. Livne and Du (2014); Kadam and Kotowski (2014); Doval (2015) all study the problem of dynamic stability in a two-period model, but with different environments and preference structures.

The literature on online advertising is also related to our work. In this setting, advertisements are static, but queries arrive adversarially or stochastically over time. Unlike our model, queries persist in the market for exactly one period. Karp et al. (1990) introduced the problem and designed a randomized matching algorithm. Subsequently, the problem has been considered under several arrival models with pre-specified budgets for the advertisers, (Mehta et al., 2007; Goel and Mehta, 2008; Feldman et al., 2009; Manshadi et al., 2012; Blum et al., 2014).

The problem of dynamic matching has been extensively studied in the literature of labor search and matching in labor markets. Shimer and Smith (2001) study a decentralized search market and discuss efficiency issues. This paper and its descendants are different from ours in at least two ways: First, rather than modeling market thickness via a fixed match-function, we explicitly account for the network structure that affects the planner's options. This allows market thickness to emerge as an endogenous phenomenon. In addition, in Shimer and Smith (2001), the benefit of waiting is in increasing the *match quality*, whereas in our model we show that even if you cannot increase match quality, waiting can still be beneficial because it increases the *number* of agents who get matched. Ebrahimy and Shimer (2010) study a decentralized version of the Greedy algorithm from a labor-search perspective.

In contrast to dynamic matching, there are numerous investigations of dynamic auctions and dynamic mechanism design. Budish et al. (2013) study the problem of timing and frequent batch auctions in the high frequency setting. Parkes and Singh (2003) generalize the VCG mechanism to a dynamic setting. Athey and Segal (2007) construct efficient and incentive-compatible dynamic mechanisms for private information settings. We refer interested readers to Parkes (2007) for a review of the dynamic mechanism design literature.

2 The Model

In this section, motivated by the problem of kidney exchange, we provide a stylized stochastic continuous-time model for a bilateral matching market that runs in the interval $[0, T]$. Agents arrive at the market at rate m according to a Poisson process. Hence, in any interval $[t, t+1]$, m new agents enter the market in expectation. Throughout the paper we assume $m \geq 1$. For $t \geq 0$, let A_t be the set of the agents in our market at time t, and let $Z_t := |A_t|$. We refer to A_t as the *pool* of the market. We start by describing the evolution of A_t as a function of $t \in [0, T]$. Since we are interested in the limit behavior of A_t, without loss of generality, we may assume $A_0 = \emptyset$. We use A_t^n to denote[11] the set of agents who enter the market at time t. Note that with probability 1, $|A_t^n| \leq 1$. Also, let $|A_{t_0,t_1}^n|$ denote the set of agents who enter the market in time interval $[t_0, t_1]$.

Each agent becomes *critical* according to an independent Poisson process with rate λ. This implies that, if an agent a enters the market at time t_0, then she becomes critical at some time $t_0 + X$ where X is an exponential random variable with parameter λ. Any critical agent leaves the market immediately; so the last point in time that an agent can get matched is the time that she gets critical. We say an agent a *perishes* if a leaves the market unmatched.[12]

We assume that an agent $a \in A_t$ leaves the market at time t, if any of the following three events occur at time t:

- a is matched with another agent $b \in A_t$,
- a becomes critical and gets matched
- a becomes critical and leaves the market unmatched, i.e., a perishes.

It is essential to note that the arrival of the criticality event with some Poisson rate is *not* equivalent to discounting with the same rate, because the criticality event might be observed by the planner and the planner can react to that information. Stochastic departure is a key feature of our model, for which we consider several informational structures: When the planner has perfect foresight; when the planner knows the set of critical agents at each point in time; and when the planner does not know anything about critical agents.

Say a enters the pool at time t_0 and becomes critical at time $t_0 + X$ where X is an exponential random variable with parameter λ. By the above discussion, for any matching algorithm, a leaves the pool at some time t_1 where $t_0 \leq t_1 \leq t_0 + X$ (note a may leave sooner than $t_0 + X$ if she gets matched before becoming critical). The *sojourn* of a is the length of the interval that a is in the pool, i.e., $s(a) := t_1 - t_0$.

We use A_t^c to denote the set of agents that are critical at time t.[13] Also, note that for

[11] As a notational guidance, we use subscripts to refer to a point in time or a time interval, while superscripts n, c refer to *new* agents and *critical* agents, respectively.

[12] We intend this as a term of art. In the case of kidney exchange, perishing can be interpreted as a patient's medical condition deteriorating in such a way as to make transplants infeasible. In the case of Uber, perishing can be interpreted as the point where a potential customer chooses an alternative mode of transportation, or when a driver disconnects from the platform.

[13] In our proofs, we use the fact that $A_t^c \subseteq \cup_{0 \leq \tau \leq t} A_\tau$. In the example of the text, we have $a \in A_{t_0+X}^c$. Note that even if agent a is matched before getting critical (i.e., $t_1 < t_0 + X$), we still have that $a \in A_{t_0+X}^c$. Hence, A_t^c is not necessarily a subset of A_t since it may have agents who are already matched and left the pool. This generalized definition of A_t^c is helpful in our proofs.

any $t \geq 0$, with probability 1, $|A_t^c| \leq 1$.

For any pair of agents, the probability that a bilateral transaction between them is *acceptable* is p, where $0 \leq p \leq 1$, and these probabilities are independent. Let $d = mp$ be the *density parameter* of the model. Note that d is the number of expected acceptable transactions of agents if the planner does not match anyone.[14] In the paper, we use this definition and replace p with d/m. In proofs, we may employ p and $q = 1 - d/m$ to simplify notation. For any $t \geq 0$, let $E_t \subseteq A_t \times A_t$ be the set of acceptable bilateral transactions between the agents in the market (the set of *edges*) at time t, and let $G_t = (A_t, E_t)$ be the *exchange possibilities graph* at time t. Note that if $a, b \in A_t$ and $a, b \in A_{t'}$, then $(a, b) \in E_t$ if and only if $(a, b) \in E_{t'}$, *i.e.* the acceptable bilateral transactions are persistent throughout the process. For an agent $a \in A_t$ we use $N_t(a) \subseteq A_t$ to denote the set of neighbors of a in G_t. It follows that, if the planner does not match any agents, then for any fixed $t \geq 0$, G_t is distributed as an Erdös-Réyni graph with parameter d/m and d is the average degree[15] of agents (Erdös and Rényi, 1960).

Let $A = \cup_{t \leq T} A_t^n$, let $E \subseteq A \times A$ be the set of acceptable transactions between agents in A, and let $G = (A, E)$[16]. Observe that any realization of the above stochastic process is uniquely defined given A_t^n, A_t^c for all $t \geq 0$ and the set of acceptable transactions, E. A vector (m, d, λ) represents a *dynamic matching market*. Without loss of generality, we can scale time so that $\lambda = 1$ (by normalizing m and d). Therefore, throughout the paper, we assume $\lambda = 1$, unless otherwise specified[17].

Online Matching Algorithms. A set of edges $M_t \subseteq E_t$ is a *matching* if no two edges share the same endpoints. An *online matching algorithm*, at any time $t \geq 0$, selects a (possibly empty) matching, M_t, in the current acceptable transactions graph G_t, and the endpoints of the edges in M_t leave the market immediately. We assume that any online matching algorithm at any time t_0 only knows the current graph G_t for $t \leq t_0$ and does not know anything about $G_{t'}$ for $t' > t_0$. In the benchmark case that we consider, the online algorithm can depend on the set of critical agents at time t; nonetheless, we will extend several of our theorems to the case where the online algorithm does not have this knowledge.

[14]In principle, for a given market, one can measure m (the arrival rate of agents) and p (the probability that two random agents have an acceptable trade) and calculate d.

[15]In an undirected graph, *degree* of of a vertex is equal to the total number of edges connected to that vertex.

[16]Note that $E \supseteq \cup_{t \leq T} E_t$, and the two sets are not typically equal, since two agents may find it acceptable to transact, even though they are not in the pool at the same time because one of them was matched earlier.

[17]See Proposition 5.12 for details of why this is without loss of generality.

As will become clear, this knowledge has a significant impact on the performance of any online algorithm.

We emphasize that the random sets A_t (the set of agents in the pool at time t), E_t (the set of acceptable transactions at time t), $N_t(a)$ (the set of an agent a's neighbors), and the random variable Z_t (pool size at time t) are all functions of the underlying matching algorithm. We abuse notation and do not include the name of the algorithm when we analyze these variables.

The Goal. The goal of the planner is to design an online matching algorithm that maximizes social welfare, i.e., the sum of the utility of all agents in the market. Let $\mathrm{ALG}(T)$ be the set of matched agents by time T,

$$\mathrm{ALG}(T) := \{a \in A : a \text{ is matched by ALG by time } T\}.$$

We may drop the T in the notation $\mathrm{ALG}(T)$ if it is clear from context.

An agent receives zero utility if she leaves the market unmatched. If she is matched, she receives a utility of 1 discounted at rate r. More formally, if $s(a)$ is the sojourn of agent a, then we define the utility of agent a as follows:

$$u(a) := \begin{cases} e^{-rs(a)} & \text{if } a \text{ is matched} \\ 0 & \text{otherwise.} \end{cases}$$

We define the social welfare of an online algorithm to be the expected sum of the utility of all agents in the interval $[0, T]$, divided by a normalization factor:

$$\mathbf{W}(\mathrm{ALG}) := \mathbb{E}\left[\frac{1}{mT}\sum_{a\in A} u(a)\right]$$

In this model, the passage of time brings two costs. First, agents discount the future, so later matches are less valuable than earlier matches. Second, agents may perish, in which case they cannot be matched at all. In a stylized way, these stochastic departures represent medical emergencies such as vascular access failure (Roy-Chaudhury et al., 2006), which may be difficult to predict far ahead of time (Polkinghorne and Kerr, 2002).

It is instructive to consider the special case where $r = 0$, i.e., the cost of waiting is negligible compared to the cost of leaving the market unmatched. In this case, the goal of the planner is to match the maximum number of agents, or equivalently to minimize the number of perished agents. The *loss* of an online algorithm ALG is defined as the ratio of the *expected*[18] number of perished agents to the expected size of A,

$$\mathbf{L}(\mathrm{ALG}) := \frac{\mathbb{E}\left[|A - \mathrm{ALG}(T) - A_T|\right]}{\mathbb{E}\left[|A|\right]} = \frac{\mathbb{E}\left[|A - \mathrm{ALG}(T) - A_T|\right]}{mT}.$$

When we assume $r = 0$, we will use the $\mathbf{L}$ notation for the planner's loss function. When we consider $r > 0$, we will use the $\mathbf{W}$ notation for social welfare.

Each of the above optimization problems can be modeled as a Markov Decision Problem (MDP)[19] that is defined as follows. The state space is the set of pairs (H, B) where H is any undirected graph of any size, and if the algorithm knows the set of critical agents, B is a set of at most one vertex of H representing the corresponding critical agent. The action space for a given state is the set of matchings on the graph H. Under this conception, an algorithm designer wants to minimize the loss or maximize the social welfare over a time period T.

Although this MDP has infinite number of states, with small error one can reduce the state space to graphs of size at most $O(m)$.[20] Even in that case, this MDP has an exponential number of states in m, since there are at least $2^{\binom{m}{2}}/m!$ distinct graphs of size m[21], so for even moderately large markets[22], we cannot apply standard dynamic programming techniques to find the optimum online matching algorithm.

Optimum Solutions. In many parts of this paper we compare the performance of an online algorithm to the performance of an optimal *omniscient* algorithm. Unlike any online algorithm, the omniscient algorithm has full information about the future, i.e., it knows the full realization of the graph G.[23] Therefore, it can return the maximum matching in this graph as its output, and thus minimize the fraction of perished agents. Let $\mathrm{OMN}(T)$ be the set of matched agents in the *maximum* matching of G. The loss function under the

[18] It is a modeling choice to use expected value as the objective function. One may also be interested in objective functions that depend on the *variance* of the performance, as well as the expected value. As we later show, the performance of our algorithms are highly concentrated around their expected value, which guarantees that the variance is very small in most of the cases.

[19] We recommend Bertsekas (2000) for background on Markov Decision Processes.

[20] This claim will be clarified later in the paper, when we prove our concentration bounds. The bounds show that the probability of having a pool size larger than $O(m)$ is exponentially small in m.

[21] This lower bound is derived as follows: When there are m agents, there are $\binom{m}{2}$ possible edges, each of which may be present or absent. Some of these graphs may have the same structure but different agent indices. A conservative lower bound is to divide by all possible re-labellings of the agents ($m!$).

[22] For instance, for $m = 30$, there are more than 10^{98} states in the approximated MDP.

[23] In computer science, these are equivalently called *offline* algorithms.

omnsicient algorithm at time T is

$$\mathbf{L}(\mathrm{OMN}) := \frac{\mathbb{E}\left[|A - \mathrm{OMN}(T) - A_T|\right]}{mT}$$

Observe that for any online algorithm, ALG, and any realization of the probability space, we have $|\mathrm{ALG}(T)| \leq |\mathrm{OMN}(T)|$.[24]

It is also instructive to study the optimum online algorithm, an online algorithm that has access to *unlimited* computational power. By definition, an optimum online algorithm can solve the exponential-sized state space Markov Decision Problem and return the best policy function from states to matchings. We first consider OPT^c, the algorithm that knows the set of critical agents at time t (with associated loss $\mathbf{L}(\mathrm{OPT}^c)$). We then relax this assumption and consider OPT, the algorithm that does not know these sets (with associated loss $\mathbf{L}(\mathrm{OPT})$).

Let ALG^c be the loss under any online algorithm that knows the set of critical agents at time t. It follows that

$$\mathbf{L}(\mathrm{ALG}^c) \geq \mathbf{L}(\mathrm{OPT}^c) \geq \mathbf{L}(\mathrm{OMN}).$$

Similarly, let ALG be the loss under any online algorithm that does not know the set of critical agents at time t. It follows that[25]

$$\mathbf{L}(\mathrm{ALG}) \geq \mathbf{L}(\mathrm{OPT}) \geq \mathbf{L}(\mathrm{OPT}^c) \geq \mathbf{L}(\mathrm{OMN}).$$

3 Our Contributions

In this section, we present our main contributions for the limit cases, discuss intuitions, and provide overviews of the proofs. The rest of the paper includes detailed analysis of the model, exact statement of the results with no limits, and full proofs.

In our model, solving for the optimal matching algorithm is computationally complex. Nevertheless, we can characterize key features of the optimal solution by employing some algorithmic techniques. The first simple observation is that we are not fully agnostic about the optimal algorithm. In particular, we know that when waiting cost is negligible (i.e. $r = 0$), OPT^c has two essential characteristics:

[24] This follows from a straightforward revealed-preference argument: For any realization, the optimum offline policy has the information to replicate any given online policy, so it must do weakly better.

[25] Note that $|\mathrm{ALG}|$ and $|\mathrm{OPT}|$ are generally incomparable, and depending on the realization of G we may even have $|\mathrm{ALG}| > |\mathrm{OPT}|$.

i) A pair of agents a, b get matched in OPT^c only if one of them is critical. Since $r = 0$, if a, b can be matched and neither of them is critical, then we are weakly better off if we wait and match them later.

ii) If an agent a is critical at time t and $N_t(a) \neq \emptyset$ then OPT^c matches a. This property is a corollary of the simple fact that matching a critical agent does not increase the number of perished agents in any online algorithm.

OPT^c is patient: It waits until some agent gets critical and if an agent is critical and has some acceptable partner, then OPT^c matches that agent. But the choice of match partner depends on the entire network structure, which is what makes the problem combinatorially complex. Our goal here is to detach these two effects: How much is achieved merely by being patient? And how much more is achieved by optimizing over the network structure?

To do this, we start by designing a matching algorithm (the Greedy algorithm), which mimics 'match-as-you-go' algorithms used in many real marketplaces[26]. It delivers maximal matchings at any point in time, without regard for the future.

Definition 3.1 (Greedy Algorithm:)**.** *If any new agent a enters the market at time t, then match her with an arbitrary agent in $N_t(a)$ whenever $N_t(a) \neq \emptyset$. We use* $\mathbf{L}(\text{Greedy})$ *and* $\mathbf{W}(\text{Greedy})$ *to denote the loss and the social welfare under this algorithm, respectively.*

Note that since $|A_t^n| \leq 1$ almost surely, we do not need to consider the case where more than one agent enters the market at any point in time. Observe that the graph G_t in the Greedy algorithm is almost always an empty graph. Hence, the Greedy algorithm cannot use any information about the set of critical agents.

To detach the value of waiting from the value of optimizing over the network structure, we design a second algorithm which chooses the optimal *time* to match agents, but ignores the network structure.

Definition 3.2 (Patient Algorithm)**.** *If an agent a becomes critical at time t, then match her uniformly at random with an agent in $N_t(a)$ whenever $N_t(a) \neq \emptyset$. We use* $\mathbf{L}(\text{Patient})$ *and* $\mathbf{W}(\text{Patient})$ *to denote the loss and the social welfare under this algorithm, respectively.*

To run the Patient algorithm, we need access to the set of critical agents at time t. We do not intend the timing assumptions about critical agents to be interpreted literally. An agent's point of perishing represents the point at which it ceases to be socially valuable to match

[26]For instance, the Alliance for Paired Donation conducts a 'match-run' everyday.

that agent. Letting the planner observe the set of critical agents is a modeling convention that represents high-accuracy short-horizon information about agents' departures.[27] Note that the Patient algorithm exploits *only* short-horizon information about urgent cases, as compared to the Omniscient algorithm which has full information about the future. We discuss implications of relaxing our informational assumptions in Subsection 3.3.

We also design a class of algorithms that interpolates between the Greedy and the Patient algorithms. The idea of this algorithm is to assign independent exponential clocks with rates $1/\alpha$ where $\alpha \in [0, \infty)$ to each agent a. If agent a's exponential clock ticks, the market-maker attempts to match her. If she has no neighbors, then she remains in the pool until she gets critical, where the market-maker attempts to match her again.

A technical difficulty with the above matching algorithm is that it is not memoryless; since when an agent gets critical and has no neighbors, she remains in the pool. Therefore, instead of the above algorithm, we study a slightly different matching algorithm (with a worse loss).

Definition 3.3 (The Patient(α) algorithm)**.** *Assign independent exponential clocks*[28] *with rate* $1/\alpha$ *where* $\alpha \in [0, \infty)$ *to each agent* a*. If agent* a*'s exponential clock ticks or if an agent* a *becomes critical at time* t*, match her uniformly at random with an agent in* $N_t(a)$ *whenever* $N_t(a) \neq \emptyset$*. In both cases, if* $N_t(a) = \emptyset$*, treat that agent as if she has perished; i.e., never match her again. We use* $\mathbf{L}(\text{Patient}(\alpha))$ *and* $\mathbf{W}(\text{Patient}(\alpha))$ *to denote the loss and the social welfare under this algorithm, respectively.*

It is easy to see that an upper bound on the loss of the Patient(α) algorithm is an upper bound on the loss of our desired interpolating algorithm. Under this algorithm each agent's exponential clock ticks at rate $\frac{1}{\alpha}$, so we search their neighbors for a potential match at rate $\bar{\alpha} := 1 + \frac{1}{\alpha}$. We refer to $\bar{\alpha}$ as the *trade frequency*. Note that the trade frequency is a decreasing function of α.

[27] An example of such information is the Model for End-Stage Liver Disease (MELD) score, which accurately predicts 3-month mortality among patients with chronic liver disease. The US Organ Procurement and Transplantation Network gives priority to individuals with a higher MELD score, following a broad medical consensus that liver donor allocation should be based on urgency of need and not substantially on waiting time. (Wiesner et al., 2003) Another example of such information is when a donor needs to donate his/her kidney in a certain time interval.

[28] The use of exponential clocks is a modeling convention that enables us to reduce waiting times while retaining analytically tractable Markov properties.

Stylized Facts and Limit Results. Section 4 and Section 5 study the performance of each algorithm as a function of m, T, and d, without taking limits. However, compatibility graphs for kidney exchange tend to be both large ($m > 500$) and sparse ($p < .02$, $d < 15$).[29] Consequently, we now state simplified results for the case of large markets with sparse graphs, in the steady state: $m \to \infty$, d is held constant, and $T \to \infty$. Clearly, this implies that $\frac{d}{m} = p \to 0$, which should not be taken literally. This method eliminates nuisance terms that are small given typical parameter values for kidney exchange, and is a standard way to state results for large but sparse graphs (Erdös and Rényi, 1960).

The results that follow do not depend sensitively on large values of m and T. Simulations in Appendix H indicate that the key comparisons hold for small values of m. Moreover, the algorithms we examine converge rapidly to the stationary distribution (Theorem 4.2). See Theorem 5.1, Theorem 5.2, and Theorem F.1 for the dependence of our results on T.

3.1 Timing and Thickness in Dynamic Matching Markets

Does timing substantially affect the performance of dynamic matching algorithms? Our first result establishes that varying the timing properties of simple algorithms has large effects on their performance. In particular, we show that the number of perished agents under the Patient algorithm is exponentially (in d) smaller than the number of perished agents under the Greedy algorithm. This shows that gains from waiting can be substantial.

Theorem 3.4. *For (constant) $d \geq 2$, as $T, m \to \infty$,*

$$\begin{aligned} \mathbf{L}(\text{Greedy}) &\geq \frac{1}{2d+1} \\ \mathbf{L}(\text{Patient}) &\leq \frac{1}{2} \cdot e^{-d/2} \end{aligned}$$

As a result,

$$\mathbf{L}(\text{Patient}) \leq (d + \frac{1}{2}) \cdot e^{-d/2} \cdot \mathbf{L}(\text{Greedy})$$

We already knew that the Patient algorithm outperforms the Greedy algorithm. What this theorem shows is that, in a world with negligible waiting costs, the Patient algorithm *exponentially* outperforms the Greedy algorithm. The intuition behind this finding is that, under the Greedy algorithm, there are no acceptable transactions among the set of agents in the market (the market is *thin*) and so all critical agents perish. On the contrary, under the

[29] See footnote 7 for a calculation using US data.

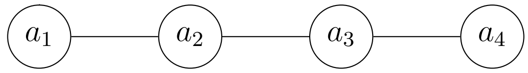

Figure 3: If a_2 gets critical in the above graph, it is strictly better to match him to a_1 as opposed to a_3. The Patient algorithm, however, chooses either of a_1 or a_3 with equal probability.

Patient algorithm, pool is always an Erdös-Réyni random graph (see Proposition 4.1), and so the market is thicker. This market thickness helps the planner to react to critical cases.

Theorem 3.4 provides an *upper bound* for the value of waiting: We shut down both channels by which waiting can be costly (negligible discounting, while the planner observes critical agents) and show that in this world, the *option value* of waiting is large. We will soon see that waiting is not necessarily valuable if either of those two channels is open.

The next question is, are the gains from market thickness large compared to the total gains from optimizing over the network structure and choosing the right agents to match? First, in the following example, we show that the Patient algorithm is naïve in choosing the right agent to match to a critical agent, because it ignores the global network structure.

Example 3.5. *Let G_t be the graph shown in Figure 3, and let $a_2 \in A_t^c$, i.e., a_2 is critical at time t. Observe that it is strictly better to match a_2 to a_1 as opposed to a_3. Nevertheless, since the Patient algorithm makes decisions that depend only on the immediate neighbors of the agent it is trying to match, it cannot differentiate between a_1 and a_3 and will choose either of them with equal probability.*

The next theorem shows that no algorithm achieves better than exponentially small loss. Furthermore, the gains from the right timing decision (moving from the Greedy algorithm to the Patient algorithm) are larger than the remaining gains from optimization (moving from the Patient algorithm to the optimum algorithm).

Theorem 3.6. *For (constant) $d \geq 2$, as $T, m \to \infty$,*

$$\frac{e^{-d}}{d+1} \quad \leq \quad \mathbf{L}(\mathrm{OPT}^c) \leq \mathbf{L}(\mathrm{Patient}) \leq \frac{1}{2} \cdot e^{-d/2}.$$

This constitutes an answer to the "when to match versus whom to match" question. Recall that OPT^c is the globally optimal solution. In many settings, optimal solutions may be computationally demanding and difficult to implement. Thus, this result suggests that, under some conditions, it will often be more worthwhile for policymakers to find ways to thicken the market, rather than to seek potentially complicated optimal policies.

It is worth emphasizing that this result (as well as Theorem 3.10) proves that *local* algorithms are close-to-optimal and since in our model agents are *ex ante* homogeneous, this shows that "whom to match" is not as important as "when to match". In settings where agents have multiple types, however, the decision of "whom to match" can be an important one even when it is local. For instance, suppose a critical agent has two neighbors, one who is *hard-to-match* and one who is *easy-to-match*. Then, *ceteris paribus*, the optimal policy should match the critical agent to the hard-to-match neighbor and breaking the ties in favor of hard-to-match agents reduces the loss.

The planner may not be willing to implement the Patient algorithm for various reasons. For instance, the cost of waiting is usually not zero; in other words, agents prefer to be matched earlier (we discuss this cost in detail in Subsection 3.2). In addition, the planner may be in competition with other exchange platforms and may be able to attract more agents by advertising reasonably short waiting times.[30] Hence, we study the performance of the Patient(α) algorithm, which introduces a way to speed up the Patient algorithm. The next result shows that when α is not 'too small' (i.e., the exponential clocks of the agents do not tick at a very fast rate), then Patient(α) algorithm still (strongly) outperforms the Greedy algorithm. In other words, even waiting for a moderate time can substantially reduce perishings.

Theorem 3.7. *Let $\bar{\alpha} := 1/\alpha + 1$. For (constant) $d \geq 2$ and $\alpha \geq 0$, as $T, m \to \infty$,*

$$\mathbf{L}(\text{Patient}(\alpha)) \leq (d+1) \cdot e^{-d/2\bar{\alpha}} \cdot \mathbf{L}(\text{Greedy})$$

A numerical example clarifies the significance of this result. Consider the case of a kidney exchange market, where 1000 new pairs arrive to the market every year, their average sojourn is 1 year and they can exchange kidneys with a random pair with probability $\frac{1}{100}$; that is, $d = 10$. The above result for the Patient(α) algorithm suggests that the market-maker can promise to match pairs in less than 4 months (in expectation) while the fraction of perished pairs is at most 37% of the Greedy algorithm. Note that if we employ the Patient algorithm, the fraction of perished pairs will be at most 7% of the Greedy algorithm.

We now present the ideas behind the proofs of the theorems presented in this section.

[30]We do not explicitly model platform competition since it is beyond the scope of this paper. One follow-up paper uses our conceptual framework and techniques to model platform competition (see Das et al. (2015)). In that model agents and platforms are not "strategic" (in its game-theoretic sense). The analysis of the Patient(α), nevertheless, can pave the way for studying competition in strategic settings, which is an area of future research.

The details of the proofs are discuss in the rest of the paper.

Proof Overview. [Theorem 3.4, Theorem 3.6, and Theorem 3.7] We first sketch the proof of Theorem 3.4. We show that for large enough values of T and m, (i) $\mathbf{L}(\text{Patient}) \leq e^{-d/2}/2$ and (ii) $\mathbf{L}(\text{OPT}) \geq 1/(2d+1)$. By the fact that $\mathbf{L}(\text{Greedy}) \geq \mathbf{L}(\text{OPT})$ (since the Greedy algorithm does not use information about critical agents), Theorem 3.4 follows immediately. The key idea in proving both parts is to carefully study the distribution of the pool size, Z_t, under any of these algorithms.

For (i), we show that pool size under the Patient algorithm is a Markov chain, it has a unique stationary distribution and it mixes rapidly to the stationary distribution (see Theorem 4.2). This implies that for t larger than mixing time, Z_t is essentially distributed as the stationary distribution of the Markov Chain. We show that under the stationary distribution, with high probability, $Z_t \in [m/2, m]$. Therefore, any critical agent has no acceptable transactions with probability at most $(1 - d/m)^{m/2} \leq e^{-d/2}$. This proves (i) (see Subsection 5.2 for the exact analysis of the Patient Markov chain).

To show that $\mathbf{L}(\text{OPT}^c) \geq \frac{e^{-d}}{d+1}$, we prove a stronger result: we show that $\mathbf{L}(\text{OMN}) \geq \frac{e^{-d}}{d+1}$. Note that $\mathbf{L}(\text{OMN})$ is a lower bound for the loss of any algorithm. To see how we bound $\mathbf{L}(\text{OPT})$ and $\mathbf{L}(\text{OMN})$, see the proof overview of Theorem 3.10, and for a detailed proof, see Theorem B.1 and Theorem B.2.

We now sketch the idea of the proof of Theorem 3.7. By the additivity of the Poisson process, the loss of Patient(α) algorithm in a $(m, d, 1)$ matching market is equal to the loss of the Patient algorithm in a $(m, d, \bar{\alpha})$ matching market, where $\bar{\alpha} = 1/\alpha + 1$.

The next step is to show that a matching market $(m, d, \bar{\alpha})$ is *equivalent* to a matching market $(m/\bar{\alpha}, d/\bar{\alpha}, 1)$ in the sense that any quantity in these two markets is the same up to a time scale (see Definition 5.11). By this fact, the loss of the Patient algorithm on a $(m, d, \bar{\alpha})$ matching market at time T is equal to the loss of Patient algorithm on a $(m/\bar{\alpha}, d/\bar{\alpha}, 1)$ market at time $\bar{\alpha}T$. But, we have already upper bounded the latter in Theorem 3.4.

□

One alternative interpretation of the above results is that information (i.e. knowledge of the set of critical agents) is valuable, rather than that waiting is valuable. This is not our interpretation at this point, since the Greedy algorithm cannot improve its performance even if it has knowledge of the set of critical agents. The graph G_t is almost surely an empty graph, so there is no possibility of matching an urgent case in the Greedy algorithm. Because urgent cases depart imminently, maintaining market thickness at all times is highly valuable. The Patient algorithm strongly outperforms the Greedy algorithm because it *waits* long enough to create a thick market.

3.2 Welfare Under Discounting and Optimal Waiting Time

In this part we explictly account for the cost of waiting and study online algorithms that optimize social welfare. It is clear that if agents are very impatient (i.e., they have very high waiting costs), it is better to implement the Greedy algorithm. On the other hand, if agents are very patient (i.e., they have very low waiting costs), it is better to implement the Patient algorithm. Therefore, a natural welfare economics question is: For which values of r is the Patient algorithm (or the Patient(α) algorithm) socially preferred to the Greedy algorithm?

Our next result shows that for small enough r, there exists a value of α such that the Patient(α) algorithm is socially preferable to the Greedy algorithm.

Theorem 3.8. *For any $0 \leq r \leq \frac{1}{8\log(d)}$, there exists an $\alpha \geq 0$ such that as $m, T \to \infty$,*

$$\mathbf{W}(\text{Patient}(\alpha)) \geq \mathbf{W}(\text{Greedy}).$$

In particular, for $r \leq \frac{1}{2d}$ and $d \geq 5$, we have

$$\mathbf{W}(\text{Patient}) \geq \mathbf{W}(\text{Greedy}).$$

A numerical example illustrates these magnitudes. Consider a barter market, where 100 new traders arrive at the market every week, their average sojourn is one week, and there is a satisfactory trade between two random agents in the market with probability 0.05; that is, $d = 5$. Then our welfare analysis implies that if the cost associated with waiting for one week is less than 10% of the surplus from a typical trade, then the Patient(α) algorithm, for a tuned value of α, is socially preferred to the Greedy algorithm.

When agents discount the future, how should the planner trade off the frequency of transactions and the thickness of the market? To answer this, we characterize the optimal trade frequency under the Patient(α) algorithm. Recall that under this algorithm each agent's exponential clock ticks at rate $\frac{1}{\alpha}$, so we search their neighbors for a potential match at rate $\bar{\alpha} := 1 + \frac{1}{\alpha}$. The optimal $\bar{\alpha}$, the trade frequency, is stated in the following theorem.

Theorem 3.9. *Given (constant) d, r, as $m, T \to \infty$, there exists $\bar{d} \in [d/2, d]$ as a function*

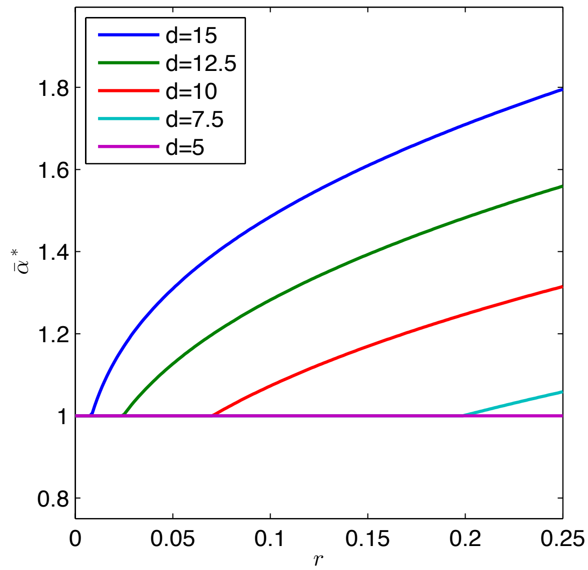

Figure 4: Optimal trade frequency ($\max\{\bar{\alpha}, 1\}$) as a function of the discount rate for different values of d. Our analysis shows that the optimal trade frequency is lower in sparser graphs.

of m, d[31] *such that the* Patient($\frac{1}{\max\{\bar{\alpha},1\}-1}$) *algorithm where $\bar{\alpha}$ is the solution of*

$$r - \left(r + \bar{d} + \frac{\bar{d}r}{\bar{\alpha}}\right)e^{-\frac{\bar{d}}{\bar{\alpha}}} = 0, \tag{3.1}$$

attains the largest welfare among all Patient(α) *algorithms. In particular, if $r < \bar{d}/4$, then $\bar{d}/\log(\bar{d}/r) \leq \bar{\alpha} \leq \bar{d}/\log(2\bar{d}/r)$.*

In addition, if $r < d/(2(e-1))$, then $\bar{\alpha}^*$ is weakly increasing in r and d.

Figure 4 illustrates $\max\{\bar{\alpha}, 1\}$ as a function of r. As one would expect, the optimal trade frequency is increasing in r. Moreover, Theorem 3.9 indicates that the optimal trade frequency is increasing in d. In Subsection 3.1, we showed that $\mathbf{L}$(Patient) is exponentially smaller in d than $\mathbf{L}$(Greedy). This may suggest that waiting is mostly valuable in dense graphs. By contrast, Theorem 3.9 shows that one should wait longer as the graph becomes sparser. Intuitively, an algorithm performs well if, whenever it searches neighbors of a critical agent for a potential match, it can find a match with very high probability. This probability is a function of both the pool size and d. When d is smaller, the pool size should be larger (i.e. the trade frequency should be lower) so that the probability of finding a match remains

[31] More precisely, $\bar{d} := xd/m$ where x is the solution of equation (5.12).

high. For larger values of d, on the other hand, a smaller pool size suffices.

Finally, we note that Patient (i.e. $\alpha = \infty$) has the *optimal* matching rate for a range of parameter values. To see why, suppose agents discount their welfare, but never perish. The planner would still wish to match them at some positive rate. For a range of parameters, this positive rate is less than 1. In a world where agents perish, $\bar{\alpha}$ is bounded below by 1, and we have a corner solution for such parameter values.

Proof Overview. [Theorem 3.8 and Theorem 3.9] We first show that for large values of m and T, (i) $\mathbf{W}(\text{Greedy}) \leq 1 - \frac{1}{2d+1}$, and (ii) $\mathbf{W}(\text{Patient}(\alpha)) \simeq \frac{2}{2-e^{-d/2\bar{\alpha}}+r/\bar{\alpha}}(1-e^{-d/2\bar{\alpha}})$, where $\bar{\alpha} = 1 + 1/\alpha$. The proof of (i) is very simple: $1/(2d+1)$ fraction of agents perish under the Greedy algorithm. Therefore, even if all of the matched agents receive a utility of 1, the social welfare is no more than $1 - 1/(2d+1)$.

The proof of (ii) is more involved and includes multiple approximation techniques. The idea is to define a random variable X_t representing the *potential* utility of the agents in the pool at time t, i.e., if all agents who are in the pool at time t get matched immediately, then they receive a total utility of X_t. We show that X_t can be approximated with a small error by studying the evolution of the system through a differential equation. Given X_t and the pool size at time t, Z_t, the expected utility of an agent that is matched at time t is X_t/Z_t. Using our concentration results on Z_t, we can then compute the expected utility of the agents that are matched in any interval $[t, t+dt]$. Integrating over all $t \in [0, T]$ proves the claim. (See Appendix F for an exact analysis of welfare under discounting for the Patient algorithm)

To prove Theorem 3.9, we characterize the unique global maximum of $\mathbf{W}(\text{Patient}(\alpha))$. The key point is that the optimum value of $\bar{\alpha}$ ($= 1 + 1/\alpha$) is less than 1 for a range of parameters. However, since $\alpha \geq 0$, we must have that $\bar{\alpha} \in [1, \infty)$. Therefore, whenever the solution of Equation 3.1 is less than 1, the optimal $\bar{\alpha}$ is 1 and we have a corner solution; i.e. setting $\alpha = \infty$ (running the Patient algorithm) is optimal. □

3.3 Value of Information and Incentive-Compatibility

Up to this point, we have assumed that the planner knows the set of critical agents; i.e. he has accurate short-horizon information about agent departures. We now relax this assumption in both directions.

First, we consider the case that the planner does not know the set of critical agents. That is, the planner's policy may depend on the graph G_t, but not the set of critical agents A_t^c. Recall that OPT is the optimum algorithm subject to these constraints. Second, we consider

OMN, the case under which the planner knows *everything* about the future realization of the market. Our main result in this section is stated below:

Theorem 3.10. *For (constant) $d \geq 2$, as $T, m \to \infty$,*

$$\frac{1}{2d+1} \leq \mathbf{L}(\mathrm{OPT}) \leq \mathbf{L}(\mathrm{Greedy}) \leq \frac{\log(2)}{d}$$

$$\frac{e^{-d}}{d+1} \leq \mathbf{L}(\mathrm{OMN}) \leq \mathbf{L}(\mathrm{Patient}) \leq \frac{1}{2} \cdot e^{-d/2}.$$

Proof Overview. We employ a novel trick to bound $\mathbf{L}(\mathrm{OPT})$, without knowing anything about the way OPT works. The idea is to provide lower bounds on the performance of *any* matching algorithm as a function of its expected pool size. Let ζ be the expected pool size of OPT. When the planner does not observe critical cases, all critical agents perish. Hence, in steady state, $\mathbf{L}(\mathrm{OPT}) \simeq \zeta/m$, which is the perishing rate divided by the arrival rate. Note that this is an increasing function of ζ, so from this perspective the planner prefers to decrease the pool size as much as possible.

Next, we count the number of agents who do not have any acceptable transactions during their sojourn. No matching algorithm can match these agents and so the fraction of those agents is a lower bound on the performance of *any* matching algorithm, including OPT. We do this in Subsection B.1 and show that $\mathbf{L}(\mathrm{OPT}) \geq \frac{1-\zeta(d/m+d^2/m^2)}{1+2d+d/m^2}$. From this perspective, the planner prefers to increase the pool size as much as possible. One can easily show that the optimal pool size (to minimize the loss) is $1/(2d+1)$. If $\zeta \leq 1/(2d+1)$, then the fraction of agents with no acceptable transactions is more than $1/(2d+1)$.

Providing a lower bound for the $\mathbf{L}(\mathrm{OMN})$ employs a similar idea. First, as a function of the pool size, we count the fraction of agents who come to market and have no acceptable transactions during their sojourn. This is a decreasing function of the pool size. But we know that the expected pool size can never be more than m, because that is the expected pool size when the planner does not match any agents. Hence, the fraction of agents with no acceptable transactions when the expected pool size is m is a lower bound on the loss of the OMN. (See Subsection B.2 for the details.)

□

Theorem 3.10 shows that the loss of OPT and Greedy are relatively close. This indicates that waiting and criticality information are complements: Waiting to thicken the market is substantially valuable only when the planner can identify urgent cases. Observe that OPT could in principle wait to thicken the market, but the gains from doing so (compared to running the Greedy algorithm) are not large. Under these new information assumptions, we once more find that local algorithms can perform close to computationally intensive global optima. Moreover, Theorem 3.10 together show that criticality information is valuable, since the loss of Patient, which naïvely exploits criticality information, is exponentially smaller than the loss of OPT, the optimal algorithm without this information.

What if the planner knows *more* than just the set of critical agents? For instance, the planner may have long-horizon forecasts of agent departure times, or the planner may know that certain agents are more likely to have acceptable transactions in future than other agents[32]. However, Theorem 3.10 shows that *no* expansion of the planner's information set yields a better-than-exponential loss. This is because $\mathbf{L}(\text{OMN})$ is the loss under a maximal expansion of the planner's information. Taken together, these results suggest that criticality information is particularly valuable.

However, in many settings, agents have privileged insight into their own departure times. In such cases, agents may have incentives to misreport whether they are critical, in order to increase their chance of getting matched or to decrease their waiting time. In kidney exchange, for instance, doctors (and hospitals) have relatively accurate information about the urgency of a patient-donor pair's need, but kidney exchange pools are separate entities and often do not have access to that information. We now exhibit a truthful mechanism *without* transfers that elicits such information from agents, and implements the Patient(α) algorithm.

We assume that agents are fully rational and know the underlying parameters, and that they believe that the pool is in the stationary distribution when they arrive, but they do not observe the actual realization of the stochastic process. That is, agents observe whether they are critical, but do not observe G_t, while the planner observes G_t but does not observe which agents are critical. Consequently, agents' strategies are independent of the realized sample path. Our results are sensitive to this assumption[33]; for instance, if the agent knew that she had a neighbor, or knew that the pool at that moment was very large, she would have an incentive under our mechanism to falsely report that she was critical.

The truthful mechanism, Patient-Mechanism(α), is described below.

[32] In our model, the number of acceptable transactions that a given agent will have with the next N agents to arrive is Bernoulli distributed. If the planner knows beforehand whether a given agent's realization is above or below the 50th percentile of this distribution, it is as though agents have different 'types'.

[33] This assumption is plausible in many settings; generally, centralized brokers know more about the current state of the market than individual traders. Indeed, frequently agents approach centralized brokers because they do not know who is available to trade with them.

Definition 3.11 (Patient-Mechanism(α)). *Assign independent exponential clocks with rate $1/\alpha$ to each agent a, where $\alpha \in [0, \infty)$. Ask agents to report when they get critical. If agent's exponential clock ticks or if she reports becoming critical, the market-maker attempts to match her to a random neighbor. If the agent has no neighbors, the market-maker treats her as if she has perished, i.e., she will never be matched again.*

Each agent a selects a mixed strategy by choosing a function $c_a(\cdot)$; at the interval $[t, t+dt]$ after her arrival, if she is not yet critical, she reports being critical with rate $c_a(t)dt$, and when she truly becomes critical she reports that immediately. Our main result in this section asserts that if agents are not too impatient, then the Patient-Mechanism(α) is incentive-compatible in the sense that the truthful strategy profile is a strong ϵ-Nash equilibrium.[34]

Theorem 3.12. *Suppose that the market is in the stationary distribution and*[35] $d = \text{polylog}(m)$. *Let $\bar{\alpha} = 1/\alpha + 1$ and $\beta = \bar{\alpha}(1 - d/m)^{m/\bar{\alpha}}$. Then, for $0 \leq r \leq \beta$, $c_a(t) = 0$ for all a, t (i.e., truthful strategy profile) is a strong ϵ-Nash equilibrium for* Patient-Mechanism(α), *where $\epsilon \to 0$ as $m \to \infty$.*

In particular, if $d \geq 2$ and $0 \leq r \leq e^{-d/2}$, the truthful strategy profile is a strong ϵ-Nash equilibrium for Patient-Mechanism(∞), *where $\epsilon \to 0$ as $m \to \infty$.*

Proof Overview. There is a hidden obstacle in proving that truthful reporting is incentive-compatible: Even if one assumes that the market is in a stationary distribution at the point an agent enters, the agent's beliefs about pool size may change as time passes. In particular, an agent makes inferences about the current distribution of pool size, *conditional* on not having been matched yet, and this conditional distribution is different from the stationary distribution. This makes it difficult to compute the payoffs from deviations from truthful reporting. We tackle this problem by using the concentration bounds from Proposition 5.9, and focusing on strong ϵ-Nash equilibrium, which allows small deviations from full optimality.

The intuition behind this proof is that an agent can be matched in one of two ways under Patient-Mechanism(∞): Either she becomes critical and has a neighbor, or one of her neighbors becomes critical and is matched to her. By symmetry, the chance of either happening is the same, because with probability 1 every matched pair consists of one critical agent and one non-critical agent. When an agent declares that she is critical, she is taking her chance that she has a neighbor in the pool right now. By contrast, if she waits, there is some

[34]Any strong ϵ-Nash equilibrium is an ϵ-Nash equilibrium. For a definition of strong ϵ-Nash equilibrium, see Definition 6.1.

[35]polylog(m) denotes any polynomial function of $\log(m)$. In particular, $d = \text{polylog}(m)$ if d is a constant independent of m.

probability that another agent will become critical and be matched to her. Consequently, for small r, agents will opt to wait. $\square$

The key insight of Theorem 3.12 is that remaining in the pool has a "continuation value": The agent, while not yet critical, may be matched to a critical agent. If agents are not too impatient, then the planner can induce truth-telling by using punishments that decrease this continuation value. Patient-Mechanism(α) sets this continuation value to zero, but in principle softer punishments could achieve the same goal. For instance, if there are multiple potential matches for a critical agent, the planner could break ties in favor of agents who have never mis-reported. However, such mechanisms can undermine the Erdös-Réyni property that makes the analysis tractable.

3.4 Technical Contributions

As alluded to above, most of our results follow from concentration results on the distribution of the pool size for each of the online algorithms that are stated in Proposition 5.5 and Proposition 5.9. In this last part we describe ideas behind these crucial results.

For analyzing many classes of stochastic processes one needs to prove concentration bounds on functions defined on the underlying process by means of central limit theorems, Chernoff bounds or Azuma-Hoeffding bounds. In our case many of these tools fail. This is because we are interested in proving that for any large time t, a given function is concentrated in an interval whose size depend only on d, m and not t. Since t can be significantly larger than d or m, a direct proof fails.

Instead, we observe that Z_t is a Markov Chain for a large class of online algorithms. Building on this observation, first we show that the underlying Markov Chain has a unique stationary distribution and it *mixes* rapidly. Then we use the stationary distribution of the Markov Chain to prove our concentration bounds.

However, that is not the end of the story. We do not have a closed form expression for the stationary distribution of the chain, because we are dealing with an infinite state space continuous time Markov Chain where the transition rates are complex functions of the states. Instead, we use the following trick. Suppose we want to prove that Z_t is contained in an interval $[k^* - f(m,d), k^* + f(m,d)]$ for some $k^* \in \mathbb{N}$ with high probability, where $f(m,d)$ is a function of m, d that does not depend on t. We consider a sequence of pairs of states $P_1 := (k^* - 1, k^* + 1), P_2 := (k^* - 2, k^* + 2)$, etc. We show that if the Markov Chain is at any of the states of P_i, it is more likely (by an additive function of m, d) that it jumps to a state of P_{i-1} as opposed to P_{i+1}. Using *balance equations* and simple algebraic manipulations, this implies that the probability of states in P_i *geometrically* decrease as i increases. In other words Z_t is concentrated in a small interval around k^*. We believe that this technique can be used in studying other complex stochastic processes.

4 Modeling an Online Algorithm as a Markov Chain

In this section, we establish that under both of the Patient and Greedy algorithms, the random processes Z_t are Markovian, have unique stationary distributions, and mix rapidly to the stationary distribution. We refer readers to the Appendix C for the necessary background on Markov Chains and mixing times.

First, we argue that the pool size Z_t is a Markov process under the Patient and Greedy algorithms. This follows from the following simple observation.

Proposition 4.1. *Under either of Greedy or Patient algorithms, for any $t \geq 0$, conditioned on Z_t, the distribution of G_t is uniquely defined. So, given Z_t, G_t is conditionally independent of $Z_{t'}$ for $t' < t$.*

Proof. Under the Greedy algorithm, at any time $t \geq 0$, $|E_t| = 0$. Therefore, conditioned on Z_t, G_t is an empty graph with $|Z_t|$ vertices.

For the Patient algorithm, note that the algorithm never looks at the edges between non-critical agents, so the algorithm is oblivious to these edges. It follows that under the Patient algorithm, for any $t \geq 0$, conditioned on Z_t, G_t is an Erdös-Rényi random graph with $|Z_t|$ vertices and parameter d/m. □

The following is the main theorem of this section.

Theorem 4.2. *For the Patient and Greedy algorithms and any $0 \leq t_0 < t_1$,*

$$\mathbb{P}\left[Z_{t_1} | Z_t \text{ for } 0 \leq t < t_1\right] = \mathbb{P}\left[Z_{t_1} | Z_t \text{ for } t_0 \leq t < t_1\right].$$

The corresponding Markov Chains have unique stationary distributions and mix in time $O(\log(m)\log(1/\epsilon))$ in total variation distance:

$$\tau_{\mathrm{mix}}(\epsilon) \leq O(\log(m)\log(1/\epsilon)).$$

The proof of the theorem can be found in the Appendix D.

This theorem is crucial in justifying our focus on long-run results in Section 3, since these Markov chains converge very rapidly (in $O(\log(m))$ time) to their stationary distributions.

5 Performance Analysis

In this section we upper bound $\mathbf{L}(\text{Greedy})$ and $\mathbf{L}(\text{Patient})$ as a function of d, and we upper bound $\mathbf{L}(\text{Patient}(\alpha))$ as a function of d and α.

We prove the following three theorems.[36]

Theorem 5.1. *For any $\epsilon \geq 0$ and $T > 0$,*

$$\mathbf{L}(\text{Greedy}) \leq \frac{\log(2)}{d} + \frac{\tau_{\text{mix}}(\epsilon)}{T} + 6\epsilon + O\Big(\frac{\log(m/d)}{\sqrt{dm}}\Big), \tag{5.1}$$

where $\tau_{\text{mix}}(\epsilon) \leq 2\log(m/d)\log(2/\epsilon)$.

Theorem 5.2. *For any $\epsilon > 0$ and $T > 0$,*

$$\mathbf{L}(\text{Patient}) \leq \max_{z\in[1/2,1]} \Big(z + \tilde{O}(1/\sqrt{m})\Big)e^{-zd} + \frac{\tau_{\text{mix}}(\epsilon)}{T} + \frac{\epsilon m}{d^2} + 2/m, \tag{5.2}$$

where $\tau_{\text{mix}}(\epsilon) \leq 8\log(m)\log(4/\epsilon)$.

Theorem 5.3. *Let $\bar{\alpha} := 1/\alpha + 1$. For any $\epsilon > 0$ and $T > 0$,*

$$\mathbf{L}(\text{Patient}(\alpha)) \leq \max_{z\in[1/2,1]} \Big(z + \tilde{O}(\sqrt{\bar{\alpha}/m})\Big)e^{-zd/\bar{\alpha}} + \frac{\tau_{\text{mix}}(\epsilon)}{\bar{\alpha}T} + \frac{\epsilon m\bar{\alpha}}{d^2} + 2\bar{\alpha}/m,$$

where $\tau_{\text{mix}}(\epsilon) \leq 8\log(m/\bar{\alpha})\log(4/\epsilon)$.

We will prove Theorem 5.1 in Subsection 5.1, Theorem 5.2 in Subsection 5.2 and Theorem 5.3 in Subsection 5.3. Note that the limit results of Section 3 are derived by taking limits from Equation 5.1 and Equation 5.2 (as $T, m \to \infty$).

5.1 Loss of the Greedy Algorithm

In this part we upper bound $\mathbf{L}(\text{Greedy})$. We crucially exploit the fact that Z_t is a Markov Chain and has a unique stationary distribution, $\pi : \mathbb{N} \to \mathbb{R}_+$. Our proof proceeds in three steps: First, we show that $\mathbf{L}(\text{Greedy})$ is bounded by a function of the expected pool size. Second, we show that the stationary distribution is highly concentrated around some point k^*, which we characterize. Third, we show that k^* is close to the expected pool size.

[36] We use the operators O and $\tilde{O}$ in the standard way. That is, $f(m) = O(g(m))$ iff there exists a positive real number N and a real number m_0 such that $|f(m)| \leq N|g(m)|$ for all $m \geq m_0$. $\tilde{O}$ is similar but ignores logarithmic factors, i.e. $f(m) = \tilde{O}(g(m))$ iff $f(m) = O(g(m)\log^k g(m))$ for some k.

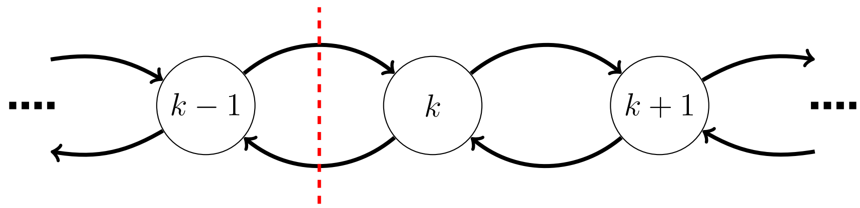

Figure 5: An illustration of the transition paths of the Z_t Markov Chain under the Greedy algorithm

Let $\zeta := \mathbb{E}_{Z\sim\mu}[Z]$ be the expected size of the pool under the stationary distribution of the Markov Chain on Z_t. First, observe that if the Markov Chain on Z_t is mixed, then the agents perish at the rate of ζ, as the pool is almost always an empty graph under the Greedy algorithm. Roughly speaking, if we run the Greedy algorithm for a sufficiently long time then Markov Chain on size of the pool mixes and we get $\mathbf{L}(\text{Greedy}) \simeq \frac{\zeta}{m}$. This observation is made rigorous in the following lemma. Note that as T and m grow, the first three terms become negligible.

Lemma 5.4. *For any $\epsilon > 0$, and $T > 0$,*

$$\mathbf{L}(\text{Greedy}) \leq \frac{\tau_{\text{mix}}(\epsilon)}{T} + 6\epsilon + \frac{1}{m}2^{-6m} + \frac{\mathbb{E}_{Z\sim\pi}[Z]}{m}.$$

The theorem is proved in the Appendix E.1.

The proof of the above lemma involves lots of algebra, but the intuition is as follows: The $\frac{\mathbb{E}_{Z\sim\pi}[Z]}{m}$ term is the loss *under* the stationary distribution. This is equal to $\mathbf{L}(\text{Greedy})$ with two approximations: First, it takes some time for the chain to transit to the stationary distribution. Second, even when the chain mixes, the distribution of the chain is not *exactly* equal to the stationary distribution. The $\frac{\tau_{\text{mix}}(\epsilon)}{T}$ term provides an upper bound for the loss associated with the first approximation, and the term $(6\epsilon + \frac{1}{m}2^{-6m})$ provides an upper bound for the loss associated with the second approximation.

Given Lemma 5.4 , in the rest of the proof we just need to get an upper bound for $\mathbb{E}_{Z\sim\pi}[Z]$. Unfortunately, we do not have any closed form expression of the stationary distribution, $\pi(\cdot)$. Instead, we use the balance equations of the Markov Chain defined on Z_t to characterize $\pi(\cdot)$ and upper bound $\mathbb{E}_{Z\sim\pi}[Z]$.

Let us rigorously define the transition probability operator of the Markov Chain on Z_t. For any pool size k, the Markov Chain transits only to the states $k+1$ or $k-1$. It transits to state $k+1$ if a new agent arrives and the market-maker cannot match her (i.e., the new agent does not have any edge to the agents currently in the pool) and the Markov Chain transits to the state $k-1$ if a new agent arrives and is matched or an agent currently in the pool gets critical. Thus, the transition rates $r_{k\to k+1}$ and $r_{k\to k-1}$ are defined as follows,

$$r_{k\to k+1} := m\Big(1-\frac{d}{m}\Big)^k \tag{5.3}$$

$$r_{k\to k-1} := k+m\Big(1-\Big(1-\frac{d}{m}\Big)^k\Big). \tag{5.4}$$

In the above equations we used the fact that agents arrive at rate m, they perish at rate 1 and the probability of an acceptable transaction between two agents is d/m.

Let us write down the balance equation for the above Markov Chain (see equation (C.3) for the full generality). Consider the cut separating the states $0,1,2,\dots,k-1$ from the rest (see Figure 5 for an illustration). It follows that,

$$\pi(k-1)r_{k-1\to k} = \pi(k)r_{k\to k-1}. \tag{5.5}$$

Now, we are ready to characterize the stationary distribution $\pi(\cdot)$. In the following proposition we show that there is a number $k^* \le \log(2)m/d$ such that under the stationary distribution, the size of the pool is highly concentrated in an interval of length $O(\sqrt{m/d})$ around k^*.[37]

Proposition 5.5. *There exists $m/(2d+1) \le k^* < \log(2)m/d$ such that for any $\sigma > 1$,*

$$\mathbb{P}_\pi\left[k^*-\sigma\sqrt{2m/d} \le Z \le k^*+\sigma\sqrt{2m/d}\right] \ge 1-O(\sqrt{m/d})e^{-\sigma^2}.$$

Proof. Let us define $f:\mathbb{R}\to\mathbb{R}$ as an interpolation of the difference of transition rates over the reals,

$$f(x) := m(1-d/m)^x - (x+m(1-(1-d/m)^x)).$$

In particular, observe that $f(k) = r_{k\to k+1} - r_{k\to k-1}$. The above function is a decreasing convex function over non-negative reals. We define k^* as the unique root of this function. Let $k^*_{\min} := m/(2d+1)$ and $k^*_{\max} := \log(2)m/d$. We show that $f(k^*_{\min}) \ge 0$ and $f(k^*_{\max}) \le 0$. This shows that $k^*_{\min} \le k^* < k^*_{\max}$.

$$\begin{aligned} f(k^*_{\min}) &\ge -k^*_{\min} - m + 2m(1-d/m)^{k^*_{\min}} \ge 2m\Big(1-\frac{k^*_{\min}d}{m}\Big) - k^*_{\min} - m = 0,\\ f(k^*_{\max}) &\le -k^*_{\max} - m + 2m(1-d/m)^{k^*_{\max}} \le -k^*_{\max} - m + 2me^{-(k^*_{\max})d/m} = -k^*_{\max} \le 0. \end{aligned}$$

[37]In this paper, $\log x$ refers to the natural log of x.

In the first inequality we used equation (A.4) from Appendix A.

It remains to show that π is highly concentrated around k^*. In the following lemma, we show that stationary probabilities decrease geometrically.

Lemma 5.6. *For any integer $k \geq k^*$*

$$\frac{\pi(k+1)}{\pi(k)} \leq e^{-(k-k^*)d/m}.$$

And, for any $k \leq k^*$, $\pi(k-1)/\pi(k) \leq e^{-(k^*-k+1)d/m}$.

This has been proved in Subsection E.2.

By repeated application of the above lemma, for any integer $k \geq k^*$, we get[38]

$$\pi(k) \leq \frac{\pi(k)}{\pi(\lceil k^* \rceil)} \leq \exp\Big(-\frac{d}{m}\sum_{i=\lceil k^* \rceil}^{k-1}(i-k^*)\Big) \leq \exp(-d(k-k^*-1)^2/2m). \tag{5.6}$$

We are almost done. For any $\sigma > 0$,

$$\begin{aligned}\sum_{k=k^*+1+\sigma\sqrt{2m/d}}^{\infty} \pi(k) \leq \sum_{k=k^*+1+\sigma\sqrt{2m/d}}^{\infty} e^{-d(k-k^*-1)^2/2m} &= \sum_{k=0}^{\infty} e^{-d(k+\sigma\sqrt{2m/d})^2/2m} \\ &\leq \frac{e^{-\sigma^2}}{\min\{1/2, \sigma\sqrt{d/2m}\}}\end{aligned}$$

The last inequality uses equation (A.1) from Appendix A. We can similarly upper bound $\sum_{k=0}^{k^*-\sigma\sqrt{2m/d}} \pi(k)$. $\square$

Proposition 5.5 shows that the probability that the size of the pool falls outside an interval of length $O(\sqrt{m/d})$ around k^* drops exponentially fast as the market size grows. We also remark that the upper bound on k^* becomes tight as d goes to infinity.

The following lemma exploits Proposition 5.5 to show that the expected value of the pool size under the stationary distribution is close to k^*.

Lemma 5.7. *For k^* as in Proposition 5.5 ,*

$$\mathbb{E}_{Z\sim\pi}[Z] \leq k^* + O(\sqrt{m/d}\log(m/d)).$$

[38] $\lceil k^* \rceil$ indicates the smallest integer larger than k^*.

This has been proved in Subsection E.3.

Now, Theorem 5.1. follows immediately by Lemma 5.4 and Lemma 5.7 because we have

$$\frac{\mathbb{E}_{Z\sim\pi}[Z]}{m} \leq \frac{1}{m}(k^* + O(\sqrt{m}\log m)) \leq \frac{\log(2)}{d} + o(1)$$

5.2 Loss of the Patient Algorithm

Let $\pi : \mathbb{N} \to \mathbb{R}_+$ be the unique stationary distribution of the Markov Chain on Z_t, and let $\zeta := \mathbb{E}_{Z\sim\pi}[Z]$ be the expected size of the pool under that distribution.

Once more our proof strategy proceeds in three steps. First, we show that $\mathbf{L}(\text{Patient})$ is bounded by a function of $\mathbb{E}_{Z\sim\pi}\left[Z(1-d/m)^{Z-1}\right]$. Second, we show that the stationary distribution of Z_t is highly concentrated around some point k^*. Third, we use this concentration result to produce an upper bound for $\mathbb{E}_{Z\sim\pi}\left[Z(1-d/m)^{Z-1}\right]$.

By Proposition 4.1, at any point in time G_t is an Erdös-Réyni random graph. Thus, once an agent becomes critical, he has no acceptable transactions with probability $(1-d/m)^{Z_t-1}$. Since each agent becomes critical with rate 1, if we run Patient for a sufficiently long time, then $\mathbf{L}(\text{Patient}) \approx \frac{\zeta}{m}(1-d/m)^{\zeta-1}$. The following lemma makes the above discussion rigorous.

Lemma 5.8. *For any $\epsilon > 0$ and $T > 0$,*

$$\mathbf{L}(\text{Patient}) \leq \frac{1}{m}\mathbb{E}_{Z\sim\pi}\left[Z(1-d/m)^{Z-1}\right] + \frac{\tau_{\text{mix}}(\epsilon)}{T} + \frac{\epsilon m}{d^2}.$$

Proof. See Appendix E.4. □

So in the rest of the proof we just need to lower bound $\mathbb{E}_{Z\sim\pi}\left[Z(1-d/m)^{Z-1}\right]$. As in the Greedy case, we do not have a closed form expression for the stationary distribution, $\pi(\cdot)$. Instead, we use the balance equations of the Markov Chain on Z_t to show that π is highly concentrated around a number k^* where $k^* \in [m/2, m]$.

Let us start by defining the transition probability operator of the Markov Chain on Z_t. For any pool size k, the Markov Chain transits only to states $k+1$, $k-1$, or $k-2$. The Markov Chain transits to state $k+1$ if a new agent arrives, to the state $k-1$ if an agent gets critical and the the planner cannot match him, and it transits to state $k-2$ if an agent gets critical and the planner matches him.

Remember that agents arrive with the rate m, they become critical with the rate of 1 and the probability of an acceptable transaction between two agents is d/m. Thus, the transition

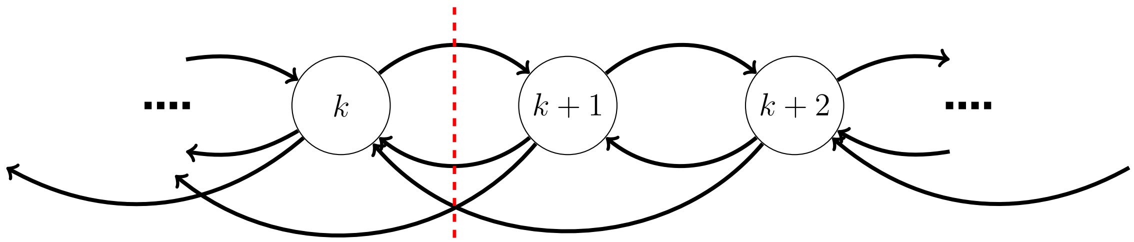

Figure 6: An illustration of the transition paths of the Z_t Markov Chain under the Patient Algorithm

rates $r_{k\to k+1}$, $r_{k\to k-1}$, and $r_{k\to k-2}$ are defined as follows,

$$
\begin{aligned}
r_{k\to k+1} &:= m && (5.7)\\
r_{k\to k-1} &:= k\Big(1-\frac{d}{m}\Big)^{k-1} && (5.8)\\
r_{k\to k-2} &:= k\Big(1-\Big(1-\frac{d}{m}\Big)^{k-1}\Big). && (5.9)
\end{aligned}
$$

Let us write down the balance equation for the above Markov Chain (see equation (C.3) for the full generality). Consider the cut separating the states $0, 1, 2, \ldots, k$ from the rest (see Figure 6 for an illustration). It follows that

$$\pi(k)r_{k\to k+1} = \pi(k+1)r_{k+1\to k} + \pi(k+1)r_{k+1\to k-1} + \pi(k+2)r_{k+2\to k} \tag{5.10}$$

Now we can characterize $\pi(\cdot)$. We show that under the stationary distribution, the size of the pool is highly concentrated around a number k^* where $k^* \in [m/2-2, m-1]$. Remember that under the Greedy algorithm, the concentration was around $k^* \in [\frac{m}{2d+1}, \frac{\log(2)m}{d}]$, whereas here it is at least $m/2$.

Proposition 5.9 (Patient Concentration)**.** *There exists a number $m/2-2 \le k^* \le m-1$ such that for any $\sigma \ge 1$,*

$$\mathbb{P}_\pi\Big[k^* - \sigma\sqrt{4m} \le Z\Big] \ge 1 - 2\sqrt{m}e^{-\sigma^2},\ \mathbb{P}\Big[Z \le k^* + \sigma\sqrt{4m}\Big] \ge 1 - 8\sqrt{m}e^{-\frac{\sigma^2\sqrt{m}}{2\sigma+\sqrt{m}}}.$$

Proof Overview. The proof idea is similar to Proposition 5.5. First, let us rewrite (5.10) by replacing transition probabilities from (5.7), (5.8), and (5.9):

$$m\pi(k) = (k+1)\pi(k+1) + (k+2)\Big(1-\Big(1-\frac{d}{m}\Big)^{k+1}\Big)\pi(k+2) \tag{5.11}$$

Let us define a continous $f : \mathbb{R} \to \mathbb{R}$ as follows,

$$f(x) := m - (x+1) - (x+2)(1 - (1 - d/m)^{x+1}). \tag{5.12}$$

It follows that

$$f(m-1) \leq 0, f(m/2-2) > 0,$$

so $f(.)$ has a root k^* such that $m/2 - 2 < k^* < m$. In the rest of the proof we show that the states that are far from k^* have very small probability in the stationary distribution, which completes the proof of Proposition 5.9. This part of the proof involves lost of algebra and is essentially very similar to the proof of the Proposition 5.5. We refer the interested reader to the Subsection E.5 for the complete proof of this last step. □

Since the stationary distribution of Z_t is highly concentrated around $k^* \in [m/2-2, m-1]$ by the above proposition, we derive the following upper bound for $\mathbb{E}_{Z\sim\pi}\left[Z(1-d/m)^Z\right]$, which is proved in the Appendix E.6.

Lemma 5.10. *For any $d \geq 0$ and sufficiently large m,*

$$\mathbb{E}_{Z\sim\pi}\left[Z(1-d/m)^Z\right] \leq \max_{z\in[m/2,m]} (z + \tilde{O}(\sqrt{m}))(1-d/m)^z + 2.$$

Now Theorem 5.2 follows immediately by combining Lemma 5.8 and Lemma 5.10.

5.3 Loss of the Patient(α) Algorithm

Our idea is to *slow down* the process and use Theorem 5.2 to analyze the Patient(α) algorithm. More precisely, instead of analyzing Patient(α) algorithm on a $(m, d, 1)$ market we analyze the Patient algorithm on a $(m/\bar{\alpha}, d/\bar{\alpha}, 1)$ market. First we need to prove a lemma on the equivalence of markets with different criticality rates.

Definition 5.11 (Market Equivalence)**.** *An* α-scaling *of a dynamic matching market (m, d, λ) is defined as follows. Given any realization of this market, i.e., given (A^c_t, A^n_t, E) for any $0 \leq t \leq \infty$, we construct another realization $(A^{c\prime}_t, A^{n\prime}_t, E')$ with $(A^{c\prime}_t, A^{n\prime}_t) = (A^c_{\alpha\cdot t}, A^n_{\alpha\cdot t})$ and the same set of acceptable transactions. We say two dynamic matching markets (m, d, λ) and (m', d', λ') are* equivalent *if one is an α-scaling of the other.*

It turns out that for any $\alpha \geq 0$, and any time t, any of the Greedy, Patient or Patient(α) algorithms (and in general any time-scale independent online algorithm) the set of matched agents at time t of a realization of a (m, d, λ) matching market is the same as the set of matched agents at time αt of an α-scaling of that realization. The following proposition makes this rigorous.

Proposition 5.12. *For any m, d, λ the $(m/\lambda, d/\lambda, 1)$ matching market is equivalent to the (m, d, λ) matching market.*[39]

Now, Theorem 5.3 follows simply by combining the above proposition and Theorem 5.2. First, by the additivity of the Poisson process, the loss of the Patient(α) algorithm in a $(m, d, 1)$ matching market is equal to the loss of the Patient algorithm in a $(m, d, \bar{\alpha})$ matching market, where $\bar{\alpha} = 1/\alpha + 1$. Second, by the above proposition, the loss of the Patient algorithm on a $(m, d, \bar{\alpha})$ matching market at time T is the same as the loss of this algorithm on a $(m/\bar{\alpha}, d/\bar{\alpha}, 1)$ market at time $\bar{\alpha}T$. The latter is upper-bounded in Theorem 5.2.

6 Incentive-Compatible Mechanisms

In this section we design a dynamic mechanism to elicit the departure times of agents. As alluded to in Subsection 3.3, we assume that agents only have statistical knowledge about the rest of the market: That is, each agent knows the market parameters $(m, d, 1)$, her own status (present, critical, perished), and the details of the dynamic mechanism that the market-maker is executing. Agents do not observe the graph G_t and their prior belief is the stationary distribution.

Each agent a chooses a *mixed strategy*, that is she reports getting critical at an infinitesimal time $[t, t + dt]$ with rate $c_a(t)dt$. In other words, each agent a has a clock that ticks with rate $c_a(t)$ at time t and she reports criticality when the clock ticks. We assume each agent's strategy function, $c_a(\cdot)$ is *well-behaved*, i.e., it is non-negative, continuously differentiable and continuously integrable. Note that since the agent can only observe the parameters of the market $c_a(\cdot)$ can depend on any parameter in our model but this function is constant in different sample paths of the stochastic process.

A *strategy profile* $\mathcal{C}$ is a vector of well-behaved functions for each agent in the market, that is, $\mathcal{C} = [c_a]_{a \in A}$. For an agent a and a strategy profile $\mathcal{C}$, let $\mathbb{E}[u_{\mathcal{C}}(a)]$ be the expected utility of a under the strategy profile $\mathcal{C}$. Note that for any $\mathcal{C}, a$, $0 \leq \mathbb{E}[u_{\mathcal{C}}(a)] \leq 1$. Given a strategy profile $\mathcal{C} = [c_a]_{a \in A}$, let $\mathcal{C} - c_a + \tilde{c}_a$ denote a strategy profile same as $\mathcal{C}$ but for agent a who is playing $\tilde{c}_a$ rather than c_a. The following definition introduces our solution concept.

[39]The proof is by inspection.

Definition 6.1. *A strategy profile $\mathcal{C}$ is a* strong ϵ-Nash equilibrium *if for any agent a and any well-behaved function $\tilde{c}_a(.)$,*

$$1 - \mathbb{E}\left[u_{\mathcal{C}}(a)\right] \leq (1+\epsilon)(1 - \mathbb{E}\left[u_{\mathcal{C}-c_a+\tilde{c}_a}\right]).$$

Note that the solution concept we are introducing here is slightly different from the usual definition of an ϵ-Nash equilibrium, where the condition is either $\mathbb{E}\left[u_{\mathcal{C}}(a)\right] \geq \mathbb{E}\left[u_{\mathcal{C}-c_a+\tilde{c}_a}\right] - \epsilon$, or $\mathbb{E}\left[u_{\mathcal{C}}(a)\right] \geq (1-\epsilon)\mathbb{E}\left[u_{\mathcal{C}-c_a+\tilde{c}_a}\right]$. The reason that we are using $1 - \mathbb{E}\left[u_{\mathcal{C}}(a)\right]$ as a measure of distance is because we know that under Patient(α) algorithm, $\mathbb{E}\left[u_{\mathcal{C}}(a)\right]$ is very close to 1, so $1 - \mathbb{E}\left[u_{\mathcal{C}}(a)\right]$ is a lower-order term. Thus, this definition restricts us to a stronger equilibrium concept, which requires us to show that in equilibrium agents can neither increase their utilities, *nor* the lower-order terms associated with their utilities by a factor of more than ϵ.

Throughout this section let $k^* \in [m/2-2, m-1]$ be the root of (5.12) as defined in Proposition 5.9, and let $\beta := (1-d/m)^{k^*}$. In this section we show that if r (the discount rate) is no more than β, then the strategy vector $c_a(t) = 0$ for all agents a and t is an ϵ-mixed strategy Nash equilibrium for ϵ very close to zero. In other words, if all other agents are truthful, an agent's utility from being truthful is almost as large as any other strategy.

Theorem 6.2. *If the market is at stationary and $r \leq \beta$, then $c_a(t) = 0$ for all a, t is a strong $O(d^4 \log^3(m)/\sqrt{m})$-Nash equilibrium for* Patient-Mechanism(∞).

By our market equivalence result (Proposition 5.12), Theorem 6.2 leads to the following corollary.

Corollary 6.3. *Let $\bar{\alpha} = 1/\alpha + 1$ and $\beta(\alpha) = \bar{\alpha}(1-d/m)^{m/\bar{\alpha}}$. If the market is at stationary and $r \leq \beta(\alpha)$, then $c_a(t) = 0$ for all a, t is a strong $O((d/\bar{\alpha})^4 \log^3(m/\bar{\alpha})/\sqrt{m/\bar{\alpha}})$-Nash equilibrium for* Patient-Mechanism(α).

The proof of the above theorem is involved but the basic idea is very easy. If an agent reports getting critical at the time of arrival she will receive a utility of $1-\beta$. On the other hand, if she is truthful (assuming $r = 0$) she will receive about $1 - \beta/2$. In the course of the proof we show that by choosing any strategy vector $c(\cdot)$ the expected utility of an agent interpolates between these two numbers, so it is maximized when she is truthful.

The precise proof of the theorem is based on Lemma 6.4. In this lemma, we upper-bound the the utility of an agent for any arbitrary strategy, given that all other agents are truthful.

Lemma 6.4. *Let Z_0 be in the stationary distribution. Suppose a enters the market at time 0. If $r < \beta$, and $10d^4 \log^3(m) \leq \sqrt{m}$, then for any well-behaved function $c(.)$,*

$$\mathbb{E}\left[u_c(a)\right] \leq \frac{2(1-\beta)}{2-\beta+r} + O\Big(d^4 \log^3(m)/\sqrt{m}\Big)\beta,$$

Proof. We present the sketch of the proof here. The full proof can be found in Appendix G.

For an agent a who arrives the market at time t_0, let $\mathbb{P}\left[a \in A_{t+t_0}\right]$ be the probability that agent a is in the pool at time $t+t_0$. Observe that an agent gets matched in one of the following two ways: First, a becomes critical in the interval $[t, t+\epsilon]$ with probability $\epsilon \cdot \mathbb{P}\left[a \in A_t\right](1+c(t))$ and if she is critical she is matched with probability $\mathbb{E}\left[(1-(1-d/m)^{Z_t-1}|a \in A_t\right]$. Second, a may also get matched (without being critical) in the interval $[t, t+\epsilon]$. Observe that if an agent $b \in A_t$ where $b \neq a$ becomes critical she will be matched with a with probability $(1-(1-d/m)^{Z_t-1})/(Z_t-1)$. Therefore, the probability that a is matched at $[t, t+\epsilon]$ without being critical is

$$\begin{aligned}\mathbb{P}\left[a \in A_t\right] \cdot \mathbb{E}\left[\epsilon \cdot (Z_t - 1)\frac{1-(1-d/m)^{Z_t-1}}{Z_t-1}|a \in A_t\right] \\ = \epsilon \cdot \mathbb{P}\left[a \in A_t\right]\mathbb{E}\left[1-(1-d/m)^{Z_t-1}|a \in A_t\right],\end{aligned}$$

and the probability of getting matched at $[t, t+\epsilon]$ is:

$$\epsilon(2+c(t))\mathbb{E}\left[1-(1-d/m)^{Z_t-1}|a \in A_t\right]\mathbb{P}\left[a \in A_t\right].$$

Based on this expression, for any strategy of agent a we have,

$$\mathbb{E}\left[u_c(a)\right] \leq \frac{\beta}{m} + \int_{t=0}^{t^*}(2+c(t))\mathbb{E}\left[1-(1-d/m)^{Z_t-1}|a \in A_t\right]\mathbb{P}\left[a \in A_t\right]e^{-rt}dt$$

where t^* is the moment where the expected utility that a receives in the interval $[t^*, \infty)$ is negligible, i.e., in the best case it is at most β/m.

In order to bound the expected utility, we need to bound $\mathbb{P}\left[a \in A_{t+t_0}\right]$. We do this by writing down the dynamical equation of $\mathbb{P}\left[a \in A_{t+t_0}\right]$ evolution, and solving the associated differential equation. In addition, we need to study $\mathbb{E}\left[1-(1-d/m)^{Z_t-1}|a \in A_t\right]$ to bound the utility expression. This is not easy in general; although the distribution of Z_t remains stationary, the distribution of Z_t conditioned on $a \in A_t$ can be a very different distribution. Therefore, we prove simple upper and lower bounds on $\mathbb{E}\left[1-(1-d/m)^{Z_t-1}|a \in A_t\right]$ using the concentration properties of Z_t. The details of all these calculations are presented in Appendix G, in which we finally obtain the following closed form upper-bound on the expected utility of a:

$$\mathbb{E}\left[u_c(a)\right] \le \frac{2d\sigma^5}{\sqrt{m}}\beta + \int_{t=0}^{\infty}(1-\beta)(2+c(t))\exp\Big(-\int_{\tau=0}^{t}(2+c(\tau)-\beta)d\tau\Big)e^{-rt}dt. \tag{6.1}$$

Finally, we show that the right hand side is maximized by letting $c(t) = 0$ for all t. Let $U_c(a)$ be the right hand side of the above equation. Let c be a function that maximizes $U_c(a)$ which is not equal to zero. Suppose $c(t) \neq 0$ for some $t \geq 0$. We define a function $\tilde{c} : \mathbb{R}_+ \to \mathbb{R}_+$ and we show that if $r < \beta$, then $U_{\tilde{c}}(a) > U_c(a)$. Let $\tilde{c}$ be the following function,

$$\tilde{c}(\tau) = \begin{cases} c(\tau) & \text{if } \tau < t, \\ 0 & \text{if } t \le \tau \le t+\epsilon, \\ c(\tau) + c(\tau-\epsilon) & \text{if } t+\epsilon \le \tau \le t+2\epsilon, \\ c(\tau) & \text{otherwise.} \end{cases}$$

In words, we push the mass of $c(.)$ in the interval $[t, t+\epsilon]$ to the right. We remark that the above function $\tilde{c}(.)$ is not necessarily continuous so we need to smooth it out. The latter can be done without introducing any errors and we do not describe the details here. In Appendix G, we show that $U_{\tilde{c}}(a) - U_c(a)$ is non-negative as long as $r \le \beta$, which means that the maximizer of $U_c(a)$ is the all zero function. Plugging in $c(t) = 0$ into (6.1) completes the proof of Lemma 6.4. □

The proof of Theorem 6.2 follows simply from the above analysis.

Proof of Theorem 6.2. All we need to do is to lower-bound the expected utility of an agent a if she is truthful. We omit the details as they are essentially similar. So, if all agents are truthful,

$$\mathbb{E}\left[u(a)\right] \ge \frac{2(1-\beta)}{2-\beta+r} - O\Big(\frac{d^4\log^3(m)}{\sqrt{m}}\Big)\beta.$$

This shows that the strategy vector corresponding to truthful agents is a strong $O(d^4\log^3(m)/\sqrt{m})$-Nash equilibrium. □

7 Assumptions and Extensions

In order to make this setting analytically tractable, we have made several important simplifying assumptions. Here we discuss how relaxing those assumptions would have implications for our results.

First, we assumed that agents are *ex ante* homogeneous: They have in expectation the same average degrees.[40] What would happen if the planner knew that certain agents currently in the pool were more likely to have edges with future agents? Clearly, algorithms that treat heterogeneous agents identically could be inefficient. However, it is an open question whether there are *local* algorithms, sensitive to individual heterogeneity, that are close to optimal.

Second, we assumed that agents' preferences are binary: All acceptable matches are equally good. We made this assumption so that waiting cannot raise welfare by improving the quality of matches. Our model shows that, even when this channel is not present, waiting can improve welfare by increasing the size of the matching. One could extend the model by permitting acceptable matches to vary in quality. We suspect that this would reinforce our existing results, since waiting to thicken the market could allow planners to make better matches, in addition to increasing the size of the matching.[41]

Third, we assumed that agents have the memoryless property; that is, they become critical at some constant Poisson rate. One might ask what would be different if the planner knew ahead of time which agents would be long-lived or short-lived. Our performance bounds on the Omniscient algorithm provide a partial answer to this question: Such information may be gainful, but a large proportion of the gains can be realized via the Patient algorithm, which uses only short-horizon information about agent's departure times.

In addition, our assumption on agents' departure processes can be enriched by assuming agents have a range of sequential *states*, while an independent process specifies transition rates from one state to the next, and agents who are at the "final" state have some exogenous criticality rate. The full analysis of optimal timing under discounting in such environment is a subject of further research. Nevertheless, our results suggest that for small waiting costs,

[40] Note, however, that agents are *ex post* heterogeneous as they have different positions in the trade network.

[41] To take one natural extension, suppose that the value of an acceptable match is v for both agents involved, where v is a random variable drawn iid across pairs of agents from some distribution $F(\cdot)$. Suppose that the Greedy and Patient algorithms are modified to select the highest match-value among the acceptable matches. Then the value to a matched agent under Greedy is (roughly) the highest among N draws from $F(\cdot)$, where N is distributed Binomial$(k^*_{\text{Greedy}}, \frac{d}{m})$. By contrast, the value to a matched agent under Patient is (roughly) the highest among N draws from $F(\cdot)$, where N is distributed Binomial$(k^*_{\text{Patient}}, \frac{d}{m})$. By our previous arguments, $k^*_{\text{Patient}} > k^*_{\text{Greedy}}$, so this strengthens our result.

if the planner observes critical agents, the Patient algorithm is close-to-optimal, and if the planner cannot observe the critical agents, waiting until agents transit to the final state (as there is no risk of agents perishing before that time) and then greedily matching those agents who have a risk of perishing is close to optimal.

Finally, our theoretical bounds on error terms, $O(\frac{1}{\sqrt{m}})$, are small only if m is relatively large. What happens if the market is small, e.g. if $m < 100$? To check the robustness of our results to the large market assumptions, we simulated our model for small markets. Our simulations (see Appendix H) suggest that our results also hold in small markets.

8 Conclusion

There are many real-world markets where the matching algorithm affects not only who gets matched today, but also what the composition of options will be tomorrow. Some examples are paired kidney exchanges, dating agencies, and labor markets such as Uber. In such markets, policymakers face a trade-off between the speed of transactions and the thickness of the market. It is natural to ask, "Should the matching algorithm wait to thicken the market? How much should it wait?"

Our analysis indicates that the answer depends on three factors: First, can the planner accurately identify urgent cases? Second, what is the discount rate? Third, how sparse is the graph of potential transactions? Waiting can yield large gains if the planner can identify urgent cases and the discount rate is not too high. This is because a thick market allows the planner to match urgent cases with high probability. If the planner cannot identify urgent cases or if the discount rate is high, then greedily and frequently matching agents is close to optimal. In general, the optimal waiting time increases if the planner can identify urgent cases, if the discount rate decreases, or if the graph becomes sparser.

The previous results show that it is valuable for the planner to identify urgent cases, such as by paying for predictive diagnostic testing or monitoring agents' outside options. When the urgency of individual cases is private information, we exhibit a mechanism without transfers that elicits such information from sufficiently patient agents.

A recurring theme of our results is that the welfare effects of timing are large compared to the total gains from optimization. In our setting, the optimal algorithm depends on a combinatorially complex state space. However, for a variety of parameters and informational assumptions, naïve local algorithms that choose the right time to match come close to optimal benchmarks that exploit the whole graph structure. This suggests that the dimension of time is a first-order concern in many matching markets, with welfare implications that static models do not capture.

Much remains to be done in the theory of dynamic matching. As market design expands its reach, re-engineering markets from the ground up, economists will increasingly have to answer questions about the timing and frequency of transactions. Many dynamic matching markets have important features (outlined above) that we have not modeled explicitly. We offer this paper as a step towards systematically understanding matching problems that take place across time.

References

Abdulkadiroglu, Atila, Parag A. Pathak, and Alvin E. Roth. (2005a). The New York City high school match. *American Economic Review*, 364-367. 1

Abdulkadiroglu, A., Pathak, P. A., Roth, A. E., and Sonmez, T. (2005b). The Boston public school match. *American Economic Review,* 368-371. 1

Abdulkadiroglu, Atila, Parag A. Pathak, and Alvin E. Roth. (2005). School choice: A mechanism design approach. *A*merican Economic Review, 93.3: 729-747. 1

Akkina, S. K., Muster, H., Steffens, E., Kim, S. J., Kasiske, B. L., and Israni, A. K. (2011). Donor exchange programs in kidney transplantation: rationale and operational details from the north central donor exchange cooperative. *American Journal of Kidney Diseases*, 57(1):152–158. 3

Anderson, R., Ashlagi, I., Gamarnik, D., and Kanoria, Y. (2014). A dynamic model of barter exchange. Technical report, Working paper, 2013. 8

Arnosti, N., Johari, R., and Kanoria, Y. (2014). Managing congestion in decentralized matching markets. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 451–451. ACM. 8

Ashlagi, I., Jaillet, P., and Manshadi, V. H. (2013). Kidney exchange in dynamic sparse heterogenous pools. In *EC*, pages 25–26. 8

Athey, S. and Segal, I. (2007). Designing efficient mechanisms for dynamic bilateral trading games. *American Economic Review*, 97(2):131–136. 9

Awasthi, P. and Sandholm, T. (2009). Online stochastic optimization in the large: Application to kidney exchange. In *IJCAI*, volume 9, pages 405–411. 8

Baccara, Mariagiovanna, SangMok Lee, and Leeat Yariv. (2015) Optimal Dynamic Matching. *Working Paper* 8

Bertsekas, D. P. (2000). *Dynamic Programming and Optimal Control*. Athena Scientific, 2nd edition. 13

Bloch, F. and Houy, N. (2012). Optimal assignment of durable objects to successive agents. *Economic Theory*, 51(1):13–33. 7

Blum, A., Haghtalab, N., Procaccia, A. D., and Sharma, A. (2014). Ignorance is almost bliss: Near-optimal stochastic matching with few queries. *arXiv preprint arXiv:1407.4094*. 9

Budish, Eric B., Peter Cramton, and John J. Shim. (2013). The high-frequency trading arms race: Frequent batch auctions as a market design response. *Fama-Miller Working Paper*, 14-03. 9

Crawford, V. P. and Knoer, E. M. (1981). Job matching with heterogeneous firms and workers. *Econometrica*, pages 437–450. 1

Das, S., Dickerson, J. P., Li, Z., and Sandholm, T. Competing Dynamic Matching Markets. In *Proceedings of the Conference on Auctions, Market Mechanisms and Their Applications (AMMA), 2015.* 8, 19

Dickerson, J. P., Procaccia, A. D., and Sandholm, T. (2012). Dynamic matching via weighted myopia with application to kidney exchange. In *AAAI*. 8

Doval, Laura. (2015). A theory of stability in dynamic matching markets. exchange. In mimeo. 8

Du, Songzi, and Yair Livne. (2014) Rigidity of transfers and unraveling in matching markets. Available at SSRN 2203910.

Ebrahimy, E. and Shimer, R. (2010). Stock–flow matching. *Journal of Economic Theory*, 145(4):1325–1353. 9

Erdös, P. and Rényi, A. (1960). On the evolution of random graphs. In *MATHEMATICAL INSTITUTE OF THE HUNGARIAN ACADEMY OF SCIENCES*, pages 17–61. 11, 17

Feldman, J., Mehta, A., Mirrokni, V., and Muthukrishnan, S. (2009). Online stochastic matching: Beating 1-1/e. In *FOCS*, pages 117–126. 9

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15. 1

Goel, G. and Mehta, A. (2008). Online budgeted matching in random input models with applications to adwords. In *SODA*, pages 982–991. 9

Grimmett, G. and Stirzaker, D. (1992). *Probability and random processes*. Oxford University Press, 2nd edition. 52

Hatfield, J. W. and Kojima, F. (2010). Substitutes and stability for matching with contracts. *Journal of Economic Theory*, 145(5):1704–1723. 1

Hatfield, J. W., Kominers, S. D., Nichifor, A., Ostrovsky, M., and Westkamp, A. (2013). Stability and competitive equilibrium in trading networks. Journal of Political Economy, 121(5), 966-1005. 1

Hatfield, J. W. and Milgrom, P. R. (2005). Matching with contracts. *AER*, pages 913–935. 1

Kadam, Sangram V., and Maciej H. Kotowski. (2014) Multi-period matching. mimeo 8

Karp, R. M., Vazirani, U. V., and Vazirani, V. V. (1990). An optimal algorithm for on-line bipartite matching. In *STOC*, pages 352–358. 9

Kelso Jr, A. S. and Crawford, V. P. (1982). Job matching, coalition formation, and gross substitutes. *Econometrica*, pages 1483–1504. 1

Kominers, S. D., and Sonmez, T. (2013). Designing for diversity in matching. In *EC* (pp. 603-604). 1

National Kidney Registry Quarterly Report (2013). Dynamic matching in overloaded systems. `http://www.kidneyregistry.org/pages/p283/NKR_ReportQ4_2013.php`. 5

Kurino, M. (2009). House allocation with overlapping agents: A dynamic mechanism design approach. *Jena economic research papers*, (075). 7

Leshno, J. D. (2012). Dynamic matching in overloaded systems. `http://www.people.fas.harvard.edu/~jleshno/papers.html2`. 7

Levin, D. A., Peres, Y., and Wilmer, E. L. (2006). *Markov chains and mixing times*. American Mathematical Society. 50, 53

Du, S., and Livne, Y. (2014). Rigidity of transfers and unraveling in matching markets. *Available at SSRN 2203910*. 8

Manshadi, V. H., Oveis Gharan, S., and Saberi, A. (2012). Online stochastic matching: Online actions based on offline statistics. *MOR*, 37(4):559–573. 9

Mehta, A., Saberi, A., Vazirani, U., and Vazirani, V. (2007). Adwords and generalized online matching. *JACM*, 54(5):22. 9

Norris, J. (1998). *Markov Chains*. Number No. 2008 in Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press. 50, 51, 52

Parkes, D. C. (2007). Online mechanisms. In Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V., editors, *Algorithmic Game Theory*, pages 411–439. Cambridge University Press, Cambridge. 9

Parkes, D. C. and Singh, S. P. (2003). An mdp-based approach to online mechanism design. In *Advances in neural information processing systems*, page None. 9

Polkinghorne, K. R., and Kerr, P. G. (2002). Predicting vascular access failure: A collective review. Nephrology, 7(4), 170-176. 12

Roy-Chaudhury, P., Sukhatme, V. P., and Cheung, A. K. (2006). Hemodialysis vascular access dysfunction: a cellular and molecular viewpoint. *J*ournal of the American Society of Nephrology, 17(4), 1112-1127. 12

Roth, A. and Sotomayor, M. (1992). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Econometric Society Monographs. Cambridge University Press. 1

Roth, A. E., Sönmez, T., and Ünver, M. U. (2004). Kidney exchange. *The Quarterly Journal of Economics*, 119(2):457–488. 1

Roth, A. E., Sönmez, T., and Ünver, M. U. (2007). Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *AER*, pages 828–851. 1

Shimer, R. and Smith, L. (2001). Matching, search, and heterogeneity. *Advances in Macroeconomics*, 1(1). 9

Sonmez, T. and Ünver, M. U. (2015). Enhancing the Efficiency of and Equity in Transplant Organ Allocation via Incentivized Exchange *Working Paper*. 5, 8

Suomela, J. (2013). Survey of local algorithms. *ACM Computing Surveys (CSUR), 45(2), 24.* 6

Su, X. and Zenios, S. A. (2005). Patient choice in kidney allocation: A sequential stochastic assignment model. *OR*, 53(3):443–455. 8

Ünver, M. U. (2010). Dynamic kidney exchange. *The Review of Economic Studies*, 77(1):372–414. 8

Wiesner, R., Edwards, E., Freeman, R., Harper, A., Kim, R., Kamath, P., Kremers, W., Lake, J., Howard, T., Merion, R. M., Wolfe, R. A., and Krom, R. (2003). Model for end-stage liver disease (meld) and allocation of donor livers. *Gastroenterology*, 124(1):91 – 96. 16

Zenios, S. A. (2002). Optimal control of a paired-kidney exchange program. *MS*, 48(3):328–342. 8

A Auxiliary Inequalities

In this section we prove several inequalities that are used throughout the paper. For any $a, b \geq 0$,

$$\begin{aligned}\sum_{i=a}^{\infty} e^{-bi^2} = \sum_{i=0}^{\infty} e^{-b(i+a)^2} \leq \sum_{i=0}^{\infty} e^{-ba^2-2iab} &= e^{-ba^2}\sum_{i=0}^{\infty}(e^{-2ab})^i \\ &= \frac{e^{-ba^2}}{1-e^{-2ab}} \leq \frac{e^{-ba^2}}{\min\{ab, 1/2\}}. \end{aligned} \tag{A.1}$$

The last inequality can be proved as follows: If $2ab \leq 1$, then $e^{-2ab} \leq ab$, otherwise $e^{-2ab} \leq 1/2$.

For any $a, b \geq 0$,

$$\sum_{i=a}^{\infty}(i-1)e^{-bi^2} \leq \int_{a-1}^{\infty} xe^{-bx^2}dx = \frac{-1}{2b}e^{-bx^2}\,|_{a-1}^{\infty} = \frac{e^{-b(a-1)^2}}{2b}. \tag{A.2}$$

For any $a \geq 0$ and $0 \leq b \leq 1$,

$$\sum_{i=a}^{\infty} ie^{-bi} = e^{-ba}\sum_{i=0}^{\infty}(i+a)e^{-bi} = e^{-ba}\Big(\frac{a}{1-e^{-b}} + \frac{1}{(1-e^{-b})}\Big) \leq \frac{e^{-ba}(2ba+4)}{b^2}. \tag{A.3}$$

The Bernoulli inequality states that for any $x \leq 1$, and any $n \geq 1$,

$$(1-x)^n \geq 1 - xn. \tag{A.4}$$

Here, we prove for integer n. The above equation can be proved by a simple induction on n. It trivially holds for $n = 0$. Assuming it holds for n we can write,

$$(1-x)^{n+1} = (1-x)(1-x)^n \geq (1-x)(1-xn) = 1 - x(n+1) + x^2 n \geq 1 - x(n+1).$$

B Performance of the Optimum Algorithms

In this section we lower-bound the loss of the optimum solutions. In particular, we prove the following theorems.

Theorem B.1. *If $m > 10d$, then for any $T > 0$*

$$\mathbf{L}(\mathrm{OPT}) \geq \frac{1}{2d + 1 + d^2/m}.$$

Theorem B.2. *If $m > 10d$, then for any $T > 0$,*

$$\mathbf{L}(\mathrm{OMN}) \geq \frac{e^{-d-d^2/m}}{d + 1 + d^2/m}$$

Before proving the above theorems, it is useful to study the evolution of the system in the case of the *inactive* algorithm, i.e., where the online algorithm does nothing and no agents ever get matched. We later use this analysis in this section, as well as Section 4 and Section 5.

We adopt the notation $\tilde{A}_t$ and $\tilde{Z}_t$ to denote the agents in the pool and the pool size in this case. Observe that by definition for any matching algorithm and any realization of the process,

$$Z_t \leq \tilde{Z}_t. \tag{B.1}$$

Using the above equation, in the following fact we show that for any matching algorithm

$\mathbb{E}[Z_t] \leq m$.

Proposition B.3. *For any $t_0 \geq 0$,*

$$\mathbb{P}\left[\tilde{Z}_{t_0} = \ell\right] \leq \frac{m^\ell}{\ell!}.$$

Therefore, $\tilde{Z}_t$ is distributed as a Poisson random variable of rate $m(1-e^{-t_0})$, so

$$\mathbb{E}\left[\tilde{Z}_{t_0}\right] = (1-e^{-t_0})m.$$

Proof. Let K be a random variable indicating the number agents who enter the pool in the interval $[0, t_0]$. By Bayes rule,

$$\mathbb{P}\left[\tilde{Z}_{t_0} = \ell\right] = \sum_{k=0}^{\infty} \mathbb{P}\left[\tilde{Z}_{t_0} = \ell, K = k\right] = \sum_{k=0}^{\infty} \mathbb{P}\left[\tilde{Z}_{t_0} = \ell | K = k\right] \cdot \frac{(mt_0)^k e^{-mt_0}}{k!},$$

where the last equation follows by the fact that arrival rate of the agents is a Poisson random variable of rate m.

Now, conditioned on the event that an agent a arrives in the interval $[0, t_0]$, the probability that she is in the pool at time t_0 is at least,

$$\mathbb{P}[X_{a_i} = 1] = \int_{t=0}^{t_0} \frac{1}{t_0} \mathbb{P}[s(a_i) \geq t_0 - t]\, dt = \frac{1}{t_0} \int_{t=0}^{t_0} e^{t-t_0} dt = \frac{1-e^{-t_0}}{t_0}.$$

Therefore, conditioned on $K = k$, the distribution of the number of agents at time t_0 is a Binomial random variable $B(k, p)$, where $p := (1-e^{-t_0})/t_0$. Let $\mu = m(1-e^{-t_0})$, we have

$$\begin{aligned}
\mathbb{P}\left[\tilde{Z}_{t_0} = \ell\right] &= \sum_{k=\ell}^{\infty} \binom{k}{\ell} \cdot p^\ell \cdot (1-p)^{k-\ell} \frac{(mt_0)^k e^{-mt_0}}{k!} \\
&= \sum_{k=\ell}^{\infty} \frac{m^k e^{-mt_0}}{\ell!(k-\ell)!} (1-e^{-t_0})^\ell (t_0 - 1 + e^{-t_0})^{k-\ell} \\
&= \frac{m^\ell e^{-mt_0} \mu^\ell}{\ell!} \sum_{k=\ell}^{\infty} \frac{(mt_0 - \mu)^{k-\ell}}{(k-\ell)!} = \frac{\mu^\ell e^{-\mu}}{\ell!}.
\end{aligned}$$

□

B.1 Loss of OPT

In this section, we prove Theorem B.1. Let ζ be the expected pool size of the OPT,

$$\zeta := \mathbb{E}_{t\sim \mathrm{unif}[0,T]}\left[Z_t\right]$$

Since OPT does not know A_t^c, each critical agent perishes with probability 1. Therefore,

$$\mathbf{L}(\mathrm{OPT}) = \frac{1}{m\cdot T}\mathbb{E}\left[\int_{t=0}^{T} Z_t dt\right] = \frac{\zeta T}{mT} = \zeta/m. \tag{B.2}$$

To finish the proof we need to lower bound ζ by $m/(2d+1+d^2/m)$. We provide an indirect proof by showing a lower-bound on $\mathbf{L}(\mathrm{OPT})$ which in turn lower-bounds ζ.

The key idea is to lower-bound the probability that an agent does not have any acceptable transactions throughout her sojourn, and this directly gives a lower-bound on $\mathbf{L}(\mathrm{OPT})$ as those agents cannot be matched under any algorithm. Fix an agent $a \in A$. Say a enters the market at a time $t_0 \sim \mathrm{unif}[0,T]$, and $s(a) = t$, we can write

$$\mathbb{P}\left[N(a) = \emptyset\right] = \int_{t=0}^{\infty} \mathbb{P}\left[s(a) = t\right]\cdot \mathbb{E}\left[(1-d/m)^{|A_{t_0}|}\right]\cdot \mathbb{E}\left[(1-d/m)^{|A^n_{t_0,t+t_0}|}\right] dt \tag{B.3}$$

To see the above, note that a does not have any acceptable transactions, if she doesn't have any neighbors upon arrival, and none of the new agents that arrive during her sojourn are not connected to her. Using the Jensen's inequality, we have

$$\begin{aligned}\mathbb{P}\left[N(a) = \emptyset\right] &\geq \int_{t=0}^{\infty} e^{-t}\cdot (1-d/m)^{\mathbb{E}\left[Z_{t_0}\right]}\cdot (1-d/m)^{\mathbb{E}\left[|A^n_{t_0,t+t_0}|\right]} dt\\ &= \int_{t=0}^{\infty} e^{-t}\cdot (1-d/m)^{\zeta}\cdot (1-d/m)^{mt} dt\end{aligned}$$

The last equality follows by the fact that $\mathbb{E}\left[|A^n_{t_0,t+t_0}|\right] = mt$. Since $d/m < 1/10$, $1 - d/m \geq e^{-d/m-d^2/m^2}$,

$$\mathbf{L}(\mathrm{OPT}) \geq \mathbb{P}\left[N(a) = \emptyset\right] \geq e^{-\zeta(d/m+d^2/m^2)}\int_{t=0}^{\infty} e^{-t(1+d+d^2/m)} dt \geq \frac{1-\zeta(1+d/m)d/m}{1+d+d^2/m} \tag{B.4}$$

Putting (B.2) and (B.4) together, for $\beta := \zeta d/m$ we get

$$\mathbf{L}(\mathrm{OPT}) \geq \max\left\{\frac{1-\beta(1+d/m)}{1+d+d^2/m}, \frac{\beta}{d}\right\} \geq \frac{1}{2d+1+d^2/m}$$

where the last inequality follows by letting $\beta = \frac{d}{2d+1+d^2/m}$ be the minimizer of the middle expression.

B.2 Loss of OMN

In this section, we prove Theorem B.2. This demonstrates that no expansion of the planner's information can yield a faster-than-exponential decrease in losses.

The proof is very similar to Theorem B.1. Let ζ be the expected pool size of the OMN,

$$\zeta := \mathbb{E}_{t\sim \text{unif}[0,T]}\left[Z_t\right].$$

By (B.1) and Proposition B.3,

$$\zeta \leq \mathbb{E}_{t\sim \text{unif}[0,T]}\left[\tilde{Z}_t\right] \leq m.$$

Note that (B.2) does not hold in this case because the omniscient algorithm knows the set of critical agents at time t.

Now, fix an agent $a \in A$, and let us lower-bound the probability that $N(a) = \emptyset$. Say a enters the market at time $t_0 \sim \text{unif}[0,T]$ and $s(a) = t$, then

$$\begin{aligned}
\mathbb{P}\left[N(a) = \emptyset\right] &= \int_{t=0}^{\infty} \mathbb{P}\left[s(a) = t\right] \cdot \mathbb{E}\left[(1-d/m)^{Z_{t_0}}\right] \cdot \mathbb{E}\left[(1-d/m)^{|A^n_{t_0,t+t_0}|}\right] dt \\
&\geq \int_{t=0}^{\infty} e^{-t}(1-d/m)^{\zeta+mt} dt \geq \frac{e^{-\zeta(1+d/m)d/m}}{1+d+d^2/m} \geq \frac{e^{-d-d^2/m}}{1+d+d^2/m}.
\end{aligned}$$

where the first inequality uses the Jensen's inequality and the second inequality uses the fact that when $d/m < 1/10$, $1 - d/m \geq e^{-d/m - d^2/m^2}$.

C Markov Chains: Background

We establish that under both of the Patient and Greedy algorithms the random processes Z_t are Markovian, have unique stationary distributions, and mix rapidly to the stationary distribution. To do so, this section contains a brief overview on continuous time Markov Chains. We refer interested readers to Norris (1998); Levin et al. (2006) for detailed discussions.

Let Z_t be a continuous time Markov Chain on the non-negative integers ($\mathbb{N}$) that starts from state 0. For any two states $i, j \in \mathbb{N}$, we assume that the rate of going from i to j is

$r_{i \to j} \geq 0$. The rate matrix $Q \in \mathbb{N} \times \mathbb{N}$ is defined as follows,

$$Q(i,j) := \begin{cases} r_{i \to j} & \text{if } i \neq j, \\ \sum_{k \neq i} -r_{i \to k} & \text{otherwise.} \end{cases}$$

Note that, by definition, the sum of the entries in each row of Q is zero. It turns out that (see e.g., (Norris, 1998, Theorem 2.1.1)) the transition probability in t units of time is,

$$e^{tQ} = \sum_{i=0}^{\infty} \frac{t^i Q^i}{i!}.$$

Let $P_t := e^{tQ}$ be the transition probability matrix of the Markov Chain in t time units. It follows that,

$$\frac{d}{dt} P_t = P_t Q. \tag{C.1}$$

In particular, in any infinitesimal time step dt, the chain moves based on $Q \cdot dt$.

A Markov Chain is irreducible if for any pair of states $i, j \in \mathbb{N}$, j is reachable from i with a non-zero probability. Fix a state $i \geq 0$, and suppose that $Z_{t_0} = i$, and let T_1 be the first jump out of i (note that T_1 is distributed as an exponential random variable). State i is *positive recurrent* iff

$$\mathbb{E}\left[\inf\{t \geq T_1 : Z_t = i\} | Z_{t_0} = i\right] < \infty \tag{C.2}$$

The ergodic theorem (Norris, 1998, Theorem 3.8.1) entails that a continuous time Markov Chain has a unique stationary distribution if and only if it has a positive recurrent state.

Let $\pi : \mathbb{N} \to \mathbb{R}_+$ be the stationary distribution of a Markov chain. It follows by the definition that for any $t \geq 0$, $P_t = \pi P_t$. The *balance equations* of a Markov chain say that for any $S \subseteq \mathbb{N}$,

$$\sum_{i \in S, j \notin S} \pi(i) r_{i \to j} = \sum_{i \in S, j \notin S} \pi(j) r_{j \to i}. \tag{C.3}$$

Let $z_t(.)$ be the distribution of Z_t at time $t \geq 0$, i.e., $z_t(i) := \mathbb{P}\left[Z_t = i\right]$ for any integer $i \geq 0$. For any $\epsilon > 0$, we define the *mixing time* (in total variation distance) of this Markov Chain as follows,

$$\tau_{\text{mix}}(\epsilon) = \inf\left\{t : \|z_t - \pi\|_{\text{TV}} := \sum_{k=0}^{\infty} |\pi(k) - z_t(k)| \leq \epsilon\right\}. \tag{C.4}$$

D Proof of Theorem 4.2

D.1 Stationary Distributions: Existence and Uniqueness

In this part we show that the Markov Chain on Z_t has a unique stationary distribution under each of the Greedy and Patient algorithms. By Proposition 4.1, Z_t is a Markov chain on the non-negative integers ($\mathbb{N}$) that starts from state zero.

First, we show that the Markov Chain is irreducible. First note that every state $i > 0$ is reachable from state 0 with a non-zero probability. It is sufficient that i agents arrive at the market with no acceptable bilateral transactions. On the other hand, state 0 is reachable from any $i > 0$ with a non-zero probability. It is sufficient that all of the i agents in the pool become critical and no new agents arrive at the market. So Z_t is an *irreducible* Markov Chain.

Therefore, by the ergodic theorem it has a unique stationary distribution if and only if it has a positive recurrent state (Norris, 1998, Theorem 3.8.1). In the rest of the proof we show that state 0 is positive recurrent. By (B.1) $Z_t = 0$ if $\tilde{Z}_t = 0$. So, it is sufficient to show

$$\mathbb{E}\left[\inf\{t \geq T_1 : \tilde{Z}_t = 0\} | \tilde{Z}_{t_0} = 0\right] < \infty. \tag{D.1}$$

It follows that $\tilde{Z}_t$ is just a continuous time birth-death process on $\mathbb{N}$ with the following transition rates,

$$\tilde{r}_{k \to k+1} = m \text{ and } \tilde{r}_{k \to k-1} := k \tag{D.2}$$

It is well known (see e.g. (Grimmett and Stirzaker, 1992, p. 249-250)) that $\tilde{Z}_t$ has a stationary distribution if and only if

$$\sum_{k=1}^{\infty} \frac{\tilde{r}_{0\to 1}\tilde{r}_{1\to 2} \ldots \tilde{r}_{k-1\to k}}{\tilde{r}_{1\to 0} \ldots \tilde{r}_{k\to k-1}} < \infty.$$

Using (D.2) we have

$$\sum_{k=1}^{\infty} \frac{\tilde{r}_{0\to 1}\tilde{r}_{1\to 2} \ldots \tilde{r}_{k-1\to k}}{\tilde{r}_{1\to 0} \ldots \tilde{r}_{k\to k-1}} = \sum_{k=1}^{\infty} \frac{m^k}{k!} = e^m - 1 < \infty$$

Therefore, $\tilde{Z}_t$ has a stationary distribution. The ergodic theorem (Norris, 1998, Theorem 3.8.1) entails that every state in the support of the stationary distribution is positive recurrent. Thus, state 0 is positive recurrent under $\tilde{Z}_t$. This proves (D.1), so Z_t is an ergodic Markov Chain.

D.2 Upper bounding the Mixing Times

In this part we complete the proof of Theorem 4.2 and provide an upper bound the mixing of Markov Chain Z_t for the Greedy and Patient algorithms. Let $\pi(.)$ be the stationary distribution of the Markov Chain.

D.2.1 Mixing time of the Greedy Algorithm

We use the *coupling* technique (see (Levin et al., 2006, Chapter 5)) to get an upper bound for the mixing time of the Greedy algorithm. Suppose we have two Markov Chains Y_t, Z_t (with different starting distributions) each running the Greedy algorithm. We define a joint Markov Chain $(Y_t, Z_t)_{t=0}^{\infty}$ with the property that projecting on either of Y_t and Z_t we see the stochastic process of Greedy algorithm, and that they stay together at all times after their first simultaneous visit to a single state, i.e.,

$$\text{if } Y_{t_0} = Z_{t_0}, \text{ then } Y_t = Z_t \text{ for } t \geq t_0.$$

Next we define the joint chain. We define this chain such that for any $t \geq t_0$, $|Y_t - Z_t| \leq |Y_{t_0} - Z_{t_0}|$. Assume that $Y_{t_0} = y, Z_{t_0} = z$ at some time $t_0 \geq 0$, for $y, z \in \mathbb{N}$. Without loss of generality assume $y < z$ (note that if $y = z$ there is nothing to define). Consider any arbitrary labeling of the agents in the first pool with $a_1, \ldots, a_y$, and in the second pool with $b_1, \ldots, b_z$. Define $z + 1$ independent exponential clocks such that the first z clocks have rate 1, and the last one has rate m. If the i-th clock ticks for $1 \leq i \leq y$, then both of a_i and b_i become critical (recall that in the Greedy algorithm the critical agent leaves the market right away). If $y < i \leq z$, then b_i becomes critical, and if $i = z + 1$ new agents a_{y+1}, b_{z+1} arrive to the markets. In the latter case we need to draw edges between the new agents and those currently in the pool. We use z independent coins each with parameter d/m. We use the first y coins to decide simultaneously on the potential transactions (a_i, a_{y+1}) and (b_i, b_{z+1}) for $1 \leq i \leq y$, and the last $z - y$ coins for the rest. This implies that for any $1 \leq i \leq y$, (a_i, a_{y+1}) is an acceptable transaction iff (b_i, b_{z+1}) is acceptable. Observe that if a_{y+1} has at least one acceptable transaction then so has b_{z+1} but the converse does not necessarily hold.

It follows from the above construction that $|Y_t - Z_t|$ is a non-increasing function of t. Furthermore, this value decreases when either of the agents $b_{y+1}, \ldots, b_z$ become critical (we note that this value may also decrease when a new agent arrives but we do not exploit this situation here). Now suppose $|Y_0 - Z_0| = k$. It follows that the two chains arrive to the same state when all of the k agents that are not in common become critical. This has the same distribution as the maximum of k independent exponential random variables with rate 1. Let E_k be a random variable that is the maximum of k independent exponentials of rate 1. For any $t \geq 0$,

$$\mathbb{P}[Z_t \neq Y_t] \leq \mathbb{P}\left[E_{|Y_0 - Z_0|} \geq t\right] = 1 - (1 - e^{-t})^{|Y_0 - Z_0|}.$$

Now, we are ready to bound the mixing time of the Greedy algorithm. Let $z_t(.)$ be the distribution of the pool size at time t when there is no agent in the pool at time 0 and let $\pi(.)$ be the stationary distribution. Fix $0 < \epsilon < 1/4$, and let $\beta \geq 0$ be a parameter that we fix later. Let (Y_t, Z_t) be the joint Markov chain that we constructed above where Y_t is started at the stationary distribution and Z_t is started at state zero. Then,

$$\begin{aligned}
\|z_t - \pi\|_{\mathrm{TV}} \leq \mathbb{P}[Y_t \neq Z_t] &= \sum_{i=0}^{\infty} \pi(i) \mathbb{P}[Y_t \neq Z_t | Y_0 = i] \\
&\leq \sum_{i=0}^{\infty} \pi(i) \mathbb{P}[E_i \geq t] \\
&\leq \sum_{i=0}^{\beta m/d} (1 - (1 - e^{-t})^{\beta m/d}) + \sum_{i=\beta m/d}^{\infty} \pi(i) \leq \frac{\beta^2 m^2}{d^2} e^{-t} + 2e^{-m(\beta-1)^2/2d}
\end{aligned}$$

where the last inequality follows by equation (A.4) and Proposition 5.5. Letting $\beta = 1 + \sqrt{2\log(2/\epsilon)}$ and $t = 2\log(\beta m/d) \cdot \log(2/\epsilon)$ we get $\|z_t - \pi\|_{\mathrm{TV}} \leq \epsilon$, which proves the theorem.

D.2.2 Mixing time of the Patient Algorithm

It remains to bound the mixing time of the Patient algorithm. The construction of the joint Markov Chain is very similar to the above construction except some caveats. Again, suppose $Y_{t_0} = y$ and $Z_{t_0} = z$ for $y, z \in \mathbb{N}$ and $t_0 \geq 0$ and that $y < z$. Let $a_1, \ldots, a_y$ and $b_1, \ldots, b_z$ be a labeling of the agents. We consider two cases.

Case 1) $z > y + 1$. In this case the construction is essentially the same as the Greedy algorithm. The only difference is that we toss random coins to decide on acceptable bilateral transactions at the time that an agent becomes critical (and not at the time of arrival). It follows that when new agents arrive the size of each of the pools increase by 1 (so the difference remains unchanged). If any of the agents $b_{y+1}, \ldots, b_z$ become critical then the size of second pool decrease by 1 or 2 and so is the difference of the pool sizes.

Case 2) $z = y + 1$. In this case we define a slightly different coupling. This is because, for some parameters and starting values, the Markov chains may not visit the same state for a long time for the coupling defined in Case 1 . If $z \gg m/d$, then with a high probability any critical agent gets matched. Therefore, the magnitude of $|Z_t - Y_t|$ does not quickly decrease (for a concrete example, consider the case where $d = m$, $y = m/2$ and $z = m/2 + 1$). Therefore, in this case we change the coupling. We use $z + 2$ independent clocks where the first z are the same as before, i.e., they have rate 1 and when the i-th clock ticks b_i (and a_i if $i \leq y$) become critical. The last two clocks have rate m, when the $z + 1$-st clock ticks a new agent arrives to the first pool and when $z + 2$-nd one ticks a new agent arrives to the second pool.

Let $|Y_0 - Z_0| = k$. By the above construction $|Y_t - Z_t|$ is a decreasing function of t unless $|Y_t - Z_t| = 1$. In the latter case this difference goes to zero if a new agent arrives to the smaller pool and it increases if a new agent arrives to the bigger pool. Let τ be the first time t where $|Y_t - Z_t| = 1$. Similar to the Greedy algorithm, the event $|Y_t - Z_t| = 1$ occurs if the second to maximum of k independent exponential random variables with rate 1 is at most t. Therefore,

$$\mathbb{P}[\tau \leq t] \leq \mathbb{P}[E_k \leq t] \leq (1 - e^{-t})^k$$

Now, suppose $t \geq \tau$; we need to bound the time it takes to make the difference zero. First, note that after time τ the difference is never more than 2. Let X_t be the (continuous time) Markov Chain illustrated in Figure 7 and suppose $X_0 = 1$. Using $m \geq 1$, it is easy to see that if $X_t = 0$ for some $t \geq 0$, then $|Y_{t+\tau} - Z_{t+\tau}| = 0$ (but the converse is not necessarily true). It is a simple exercise that for $t \geq 8$,

$$\mathbb{P}[X_t \neq 0] = \sum_{k=0}^{\infty} \frac{e^{-t}t^k}{k!} 2^{-k/2} \leq \sum_{k=0}^{t/4} \frac{e^{-t}t^k}{k!} + 2^{-t/8} \leq 2^{-t/4} + 2^{-t/8}. \tag{D.3}$$

Now, we are ready to upper-bound the mixing time of the Patient algorithm. Let $z_t(.)$ be the distribution of the pool size at time t where there is no agent at time 0, and let $\pi(.)$ be the stationary distribution. Fix $\epsilon > 0$, and let $\beta \geq 2$ be a parameter that we fix later. Let (Y_t, Z_t) be the joint chain that we constructed above where Y_t is started at the stationary

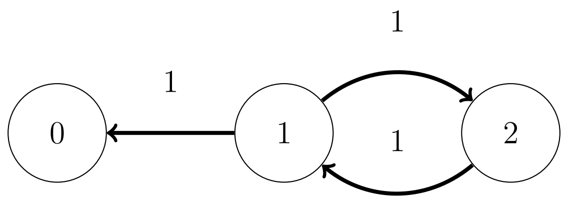

Figure 7: A three state Markov Chain used for analyzing the mixing time of the Patient algorithm.

distribution and Z_t is started at state zero.

$$
\begin{aligned}
\|z_t - \pi\|_{\mathrm{TV}} \le \mathbb{P}\left[Z_t \ne Y_t\right] &\le \mathbb{P}\left[\tau \le t/2\right] + \mathbb{P}\left[X_t \le t/2\right] \\
&\le \sum_{i=0}^{\infty} \pi(i)\mathbb{P}\left[\tau \le t/2 | Y_0 = i\right] + 2^{-t/8+1} \\
&\le 2^{-t/8+1} + \sum_{i=0}^{\infty} \pi(i)(1-(1-e^{-t/2})^i) \\
&\le 2^{-t/8+1} + \sum_{i=0}^{\beta m}(it/2) + \sum_{i=\beta m}^{\infty} \pi(i) \le 2^{-t/8+1} + \frac{\beta^2 m^2 t}{2} + 6e^{-(\beta-1)m/3}.
\end{aligned}
$$

where in the second to last equation we used equation (A.4) and in the last equation we used Proposition 5.9. Letting $\beta = 10$ and $t = 8\log(m)\log(4/\epsilon)$ implies that $\|z_t - \pi\|_{\mathrm{TV}} \le \epsilon$ which proves Theorem 4.2.

E Proofs from Section 5

E.1 Proof of Lemma 5.4

Proof. By Proposition B.3, $\mathbb{E}\left[Z_t\right] \le m$ for all t, so

$$
\begin{aligned}
\mathbf{L}(\mathrm{Greedy}) = \frac{1}{m \cdot T}\mathbb{E}\left[\int_{t=0}^{T} Z_t dt\right] &= \frac{1}{mT}\int_{t=0}^{T} \mathbb{E}\left[Z_t\right] dt \\
&\le \frac{1}{mT} m \cdot \tau_{\mathrm{mix}}(\epsilon) + \frac{1}{mT}\int_{t=\tau_{\mathrm{mix}}(\epsilon)}^{T} \mathbb{E}\left[Z_t\right] dt \qquad \text{(E.1)}
\end{aligned}
$$

where the second equality uses the linearity of expectation. Let $\tilde{Z}_t$ be the number of agents in the pool at time t when we do not match any pair of agents. By (B.1),

$$\mathbb{P}\left[Z_t \geq i\right] \leq \mathbb{P}\left[\tilde{Z}_t \geq i\right].$$

Therefore, for $t \geq \tau_{\mathrm{mix}}(\epsilon)$,

$$\begin{aligned}
\mathbb{E}\left[Z_t\right] = \sum_{i=1}^{\infty} \mathbb{P}\left[Z_t \geq i\right] &\leq \sum_{i=0}^{6m} \mathbb{P}\left[Z_t \geq i\right] + \sum_{i=6m+1}^{\infty} \mathbb{P}\left[\tilde{Z}_t \geq i\right] \\
&\leq \sum_{i=0}^{6m} (\mathbb{P}_{Z\sim\pi}\left[Z \geq i\right] + \epsilon) + \sum_{i=6m+1}^{\infty} \sum_{\ell=i}^{\infty} \frac{m^\ell}{\ell!} \\
&\leq \mathbb{E}_{Z\sim\pi}\left[Z\right] + \epsilon 6m + \sum_{i=6m+1}^{\infty} 2\frac{m^i}{i!} \\
&\leq \mathbb{E}_{Z\sim\pi}\left[Z\right] + \epsilon 6m + \frac{4m^{6m}}{(6m)!} \leq \mathbb{E}_{Z\sim\pi}\left[Z\right] + \epsilon 6m + 2^{-6m}. \quad \text{(E.2)}
\end{aligned}$$

where the second inequality uses $\mathbb{P}\left[\tilde{Z}_t = \ell\right] \leq m^\ell/\ell!$ of Proposition B.3 and the last inequality follows by the Stirling's approximation[42] of $(6m)!$. Putting (E.1) and (E.2) proves the lemma. $\square$

E.2 Proof of Lemma 5.6

Proof. For $k \geq k^*$, by (5.3), (5.4), (5.5),

$$\frac{\pi(k)}{\pi(k+1)} = \frac{(k+1) + m(1-(1-d/m)^{k+1})}{m(1-d/m)^k} = \frac{k - k^* + 1 - m(1-d/m)^{k+1} + 2m(1-d/m)^{k^*}}{m(1-d/m)^k}$$

where we used the definition of k^*. Therefore,

$$\frac{\pi(k)}{\pi(k+1)} \geq -(1-d/m) + \frac{2}{(1-d/m)^{k-k^*}} \geq \frac{1}{(1-d/m)^{k-k^*}} \geq e^{-(k^*-k)d/m}$$

[42]Stirling's approximation states that

$$n! \geq \sqrt{2\pi n}\left(\frac{n}{e}\right)^n.$$

where the last inequality uses $1-x \le e^{-x}$. Multiplying across the inequality yields the claim. Similarly, we can prove the second conclusion. For $k \le k^*$,

$$\begin{aligned}\frac{\pi(k-1)}{\pi(k)} &= \frac{k-k^*-m(1-d/m)^k+2m(1-d/m)^{k^*}}{m(1-d/m)^{k-1}}\\ &\le -(1-d/m)+2(1-d/m)^{k^*-k+1} \le (1-d/m)^{k^*-k+1} \le e^{-(k^*-k+1)d/m},\end{aligned}$$

where the second to last inequality uses $k \le k^*$. □

E.3 Proof of Lemma 5.7

Proof. Let Let $\Delta \ge 0$ be a parameter that we fix later. We have,

$$\mathbb{E}_{Z\sim\pi}[Z] \le k^* + \Delta + \sum_{i=k^*+\Delta+1}^{\infty} i\pi(i). \tag{E.3}$$

By equation (5.6),

$$\begin{aligned}\sum_{i=k^*+\Delta+1}^{\infty} i\pi(i) &= \sum_{i=\Delta+1}^{\infty} e^{-d(i-1)^2/2m}(i+k^*)\\ &= \sum_{i=\Delta}^{\infty} e^{-di^2/2m}(i-1) + \sum_{i=\Delta}^{\infty} e^{-di^2/2m}(k^*+2)\\ &\le \frac{e^{-d(\Delta-1)^2/2m}}{d/m} + (k^*+2)\frac{e^{-d\Delta^2/2m}}{\min\{1/2, d\Delta/2m\}},\end{aligned} \tag{E.4}$$

where in the last step we used equations (A.1) and (A.2). Letting $\Delta := 1+2\sqrt{m/d}\log(m/d)$ in the above equation, the right hand side is at most 1. The lemma follows from (E.3) and the above equation. □

E.4 Proof of Lemma 5.8

Proof. By linearity of expectation,

$$\mathbf{L}(\text{Patient}) = \frac{1}{m\cdot T}\mathbb{E}\left[\int_{t=0}^{T} Z_t(1-d/m)^{Z_t-1}dt\right] = \frac{1}{m\cdot T}\int_{t=0}^{T}\mathbb{E}\left[Z_t(1-d/m)^{Z_t-1}\right]dt.$$

Since for any $t \geq 0$, $\mathbb{E}\left[Z_t(1-d/m)^{Z_t-1}\right] \leq \mathbb{E}[Z_t] \leq \mathbb{E}\left[\tilde{Z}_t\right] \leq m$, we can write

$$\begin{aligned}
\mathbf{L}(\text{Patient}) &\leq \frac{\tau_{\text{mix}}(\epsilon)}{T} + \frac{1}{m \cdot T}\int_{t=\tau_{\text{mix}}(\epsilon)}^{T} \sum_{i=0}^{\infty}(\pi(i)+\epsilon) i(1-d/m)^{i-1} dt \\
&\leq \frac{\tau_{\text{mix}}(\epsilon)}{T} + \frac{\mathbb{E}_{Z\sim\pi}\left[Z(1-d/m)^{Z-1}\right]}{m} + \frac{\epsilon m}{d^2}
\end{aligned}$$

where the last inequality uses the identity $\sum_{i=0}^{\infty} i(1-d/m)^{i-1} = m^2/d^2$. □

E.5 Proof of Proposition 5.9

Let us first rewrite what we derived in the proof overview of this proposition in the main text. The balance equations of the Markov chain associated with the Patient algorithm can be written as follows by replacing transition probabilities from (5.7), (5.8), and (5.9) in (5.10):

$$m\pi(k) = (k+1)\pi(k+1) + (k+2)\Big(1-\Big(1-\frac{d}{m}\Big)^{k+1}\Big)\pi(k+2) \tag{E.5}$$

Now define a continous $f : \mathbb{R} \to \mathbb{R}$ as follows,

$$f(x) := m - (x+1) - (x+2)(1-(1-d/m)^{x+1}). \tag{E.6}$$

It follows that

$$f(m-1) \leq 0, f(m/2-2) > 0,$$

which means that $f(.)$ has a root k^* such that $m/2-2 < k^* < m$. In the rest of the proof we show that the states that are far from k^* have very small probability in the stationary distribution

In order to complete the proof of Proposition 5.9, we first prove the following useful lemma.

Lemma E.1. *For any integer $k \leq k^*$,*

$$\frac{\pi(k)}{\max\{\pi(k+1), \pi(k+2)\}} \leq e^{-(k^*-k)/m}.$$

Similarly, for any integer $k \geq k^*$, $\frac{\min\{\pi(k+1),\pi(k+2)\}}{\pi(k)} \leq e^{-(k-k^*)/(m+k-k^*)}$.

Proof. For $k \le k^*$, by equation (5.11),

$$\begin{aligned}
\frac{\pi(k)}{\max\{\pi(k+1), \pi(k+2)\}} &\le \frac{(k+1) + (k+2)(1-(1-d/m)^{k+1})}{m} \\
&\le \frac{(k-k^*) + (k^*+1) + (k^*+2)(1-(1-d/m)^{k^*+1})}{m} \\
&= 1 - \frac{k^*-k}{m} \le e^{-(k^*-k)/m},
\end{aligned}$$

where the last equality follows by the definition of k^* and the last inequality uses $1-x \le e^{-x}$. The second conclusion can be proved similarly. For $k \ge k^*$,

$$\begin{aligned}
\frac{\min\{\pi(k+1), \pi(k+2)\}}{\pi(k)} &\le \frac{m}{(k+1) + (k+2)(1-(1-d/m)^{k+1})} \\
&\le \frac{m}{(k-k^*) + (k^*+1) + (k^*+2)(1-(1-d/m)^{k^*+1})} \\
&= \frac{m}{m+k-k^*} = 1 - \frac{k-k^*}{m+k-k^*} \le e^{-(k-k^*)/(m+k-k^*)}.
\end{aligned}$$

where the equality follows by the definition of k^*. ☐

Now, we use the above claim to upper-bound $\pi(k)$ for values k that are far from k^*. First, fix $k \le k^*$. Let $n_0, n_1, \dots$ be sequence of integers defined as follows: $n_0 = k$, and $n_{i+1} := \arg\max\{\pi(n_i+1), \pi(n_i+2)\}$ for $i \ge 1$. It follows that,

$$\pi(k) \le \prod_{i: n_i \le k^*} \frac{\pi(n_i)}{\pi(n_{i+1})} \le \exp\Big(-\sum_{i: n_i \le k^*} \frac{k^* - n_i}{m}\Big) \le \exp\Big(-\sum_{i=0}^{(k^*-k)/2} \frac{2i}{m}\Big) \le e^{-(k^*-k)^2/4m}, \tag{E.7}$$

where the second to last inequality uses $|n_i - n_{i-1}| \le 2$.

Now, fix $k \ge k^* + 2$. In this case we construct the following sequence of integers, $n_0 = \lfloor k^* + 2 \rfloor$, and $n_{i+1} := \arg\min\{\pi(n_i+1), \pi(n_i+2)\}$ for $i \ge 1$. Let n_j be the largest number in the sequence that is at most k (observe that $n_j = k-1$ or $n_j = k$). We upper-bound $\pi(k)$ by upper-bounding $\pi(n_j)$,

$$\begin{aligned}
\pi(k) \le \frac{m \cdot \pi(n_j)}{k} \le 2 \prod_{i=0}^{j-1} \frac{\pi(n_i)}{\pi(n_{i+1})} \le 2\exp\Big(-\sum_{i=0}^{j-1} \frac{n_i - k^*}{m + n_i - k^*}\Big) \\
\le 2\exp\Big(-\sum_{i=0}^{(j-1)/2} \frac{2i}{m+k-k^*}\Big) \le 2\exp\Big(\frac{-(k-k^*-1)^2}{4(m+k-k^*)}\Big).
\end{aligned} \tag{E.8}$$

To see the first inequality note that if $n_j = k$, then there is nothing to show; otherwise we have $n_j = k-1$. In this case by equation (5.11), $m\pi(k-1) \geq k\pi(k)$. The second to last inequality uses the fact that $|n_i - n_{i+1}| \leq 2$.

We are almost done. The proposition follows from (E.8) and (E.7). First, for $\sigma \geq 1$, let $\Delta = \sigma\sqrt{4m}$, then by equation (A.1)

$$\sum_{i=0}^{k^*-\Delta} \pi(i) \leq \sum_{i=\Delta}^{\infty} e^{-i^2/4m} \leq \frac{e^{-\Delta^2/4m}}{\min\{1/2, \Delta/4m\}} \leq 2\sqrt{m}e^{-\sigma^2}.$$

Similarly,

$$\begin{aligned}\sum_{i=k^*+\Delta}^{\infty} \pi(i) \leq 2 \sum_{i=\Delta+1}^{\infty} e^{-(i-1)^2/4(i+m)} &\leq 2\sum_{i=\Delta}^{\infty} e^{-i/(4+\sqrt{4m}/\sigma)} \\ &\leq 2\frac{e^{-\Delta/(4+\sqrt{4m}/\sigma)}}{1-e^{-1/(4+\sqrt{4m})}} \leq 8\sqrt{m}e^{\frac{-\sigma^2\sqrt{m}}{2\sigma+\sqrt{m}}}\end{aligned}$$

This completes the proof of Proposition 5.9.

E.6 Proof of Lemma 5.10

Proof. Let $\Delta := 3\sqrt{m}\log(m)$, and let $\beta := \max_{z\in[m/2-\Delta, m+\Delta]} z(1-d/m)^z$.

$$\mathbb{E}_{Z\sim\pi}\left[Z(1-d/m)^Z\right] \leq \beta + \sum_{i=0}^{m/2-\Delta-1} \frac{m}{2}\pi(i)(1-d/m)^i + \sum_{i=m+\Delta}^{\infty} i\pi(i)(1-d/m)^m \quad \text{(E.9)}$$

We upper bound each of the terms in the right hand side separately. We start with upper bounding β. Let $\Delta' := 4(\log(2m)+1)\Delta$.

$$\begin{aligned}\beta &\leq \max_{z\in[m/2,m]} z(1-d/m)^z + m/2(1-d/m)^{m/2}((1-d/m)^{-\Delta}-1) + (1-d/m)^m\Delta \\ &\leq \max_{z\in[m/2,m]} (z+\Delta'+\Delta)(1-d/m)^z + 1. \end{aligned} \quad \text{(E.10)}$$

To see the last inequality we consider two cases. If $(1-d/m)^{-\Delta} \leq 1 + \Delta'/m$ then the inequality obviously holds. Otherwise, (assuming $\Delta' \leq m$),

$$(1-d/m)^{\Delta} \leq \frac{1}{1+\Delta'/m} \leq 1 - \Delta'/2m,$$

By the definition of β,

$$\beta \leq (m+\Delta)(1-d/m)^{m/2-\Delta} \leq 2m(1-\Delta'/2m)^{m/2\Delta-1} \leq 2me^{\Delta'/4\Delta-1} \leq 1.$$

It remains to upper bound the second and the third term in (E.9). We start with the second term. By Proposition 5.9,

$$\sum_{i=0}^{m/2-\Delta-1} \pi(i) \leq \frac{1}{m^{3/2}}. \tag{E.11}$$

where we used equation (A.1). On the other hand, by equation (E.8)

$$\sum_{i=m+\Delta}^{\infty} i\pi(i) \leq e^{-\Delta/(2+\sqrt{m})}\left(\frac{m}{1-e^{-1/(2+\sqrt{m})}} + \frac{2\Delta+4}{1/(2+\sqrt{m})^2}\right) \leq \frac{1}{\sqrt{m}}. \tag{E.12}$$

where we used equation (A.3).

The lemma follows from (E.9), (E.10), (E.11) and (E.12). $\square$

F Welfare under Discounting and Optimal Waiting Time

In this section, we provide closed form expressions for the welfare of the Greedy and Patient algorithms. We first state the theorems and then prove them in the following sections.

Theorem F.1. *There is a number $m/2-2 \leq k^* \leq m$ (as defined in Proposition 5.9) such that for any $T \geq 0$, $r \geq 0$ and $\epsilon < 1/2m^2$,*

$$\begin{aligned}
\mathbf{W}(\text{Patient}) &\geq \frac{T-T_0}{T}\Big(\frac{1-q^{k^*+\tilde{O}(\sqrt{m})}}{1+r/2-\frac{1}{2}q^{k^*-\tilde{O}(\sqrt{m})}} - \tilde{O}(m^{-3/2})\Big) \\
\mathbf{W}(\text{Patient}) &\leq \frac{2T_0}{T} + \frac{T-T_0}{T}\Big(\frac{1-q^{k^*-\tilde{O}(\sqrt{m})}}{1+r/2-\frac{1}{2}q^{k^*+\tilde{O}(\sqrt{m})}} + \tilde{O}(m^{-3/2})\Big)
\end{aligned}$$

where $T_0 = 16\log(m)\log(4/\epsilon)$. As a corollary, for any $\alpha \geq 0$, and $\bar{\alpha} = 1/\alpha + 1$,

$$\begin{aligned}
\mathbf{W}(\text{Patient}(\alpha)) &\geq \frac{T-T_0}{T}\Big(\frac{1-q^{k^*/\bar{\alpha}+\tilde{O}(\sqrt{m})}}{1+r/2\bar{\alpha}-\frac{1}{2}q^{k^*/\bar{\alpha}-\tilde{O}(\sqrt{m})}} - \tilde{O}(m^{-3/2})\Big) \\
\mathbf{W}(\text{Patient}(\alpha)) &\leq \frac{2T_0}{T} + \frac{T-T_0}{T}\Big(\frac{1-q^{k^*/\bar{\alpha}-\tilde{O}(\sqrt{m})}}{1+r/2\bar{\alpha}-\frac{1}{2}q^{k^*/\bar{\alpha}+\tilde{O}(\sqrt{m})}} + \tilde{O}(m^{-3/2})\Big)
\end{aligned}$$

Theorem F.2. *If $m > 10d$, for any $T \geq 0$,*

$$\mathbf{W}(\text{Greedy}) \leq 1 - \frac{1}{2d + 1 + d^2/m}.$$

F.1 Welfare of the Patient Algorithm

In this section, we prove Theorem F.1. Say an agent a is arrived at time $t_a(a)$. We let X_t be the sum of the potential utility of the agents in A_t:

$$X_t = \sum_{a \in A_t} e^{-r(t - t_a(a))},$$

i.e., if we match all of the agents currently in the pool immediately, the total utility that they receive is exactly X_t.

For $t_0, \epsilon > 0$, let $W_{t_0, t_0+\epsilon}$ be the expected total utility of the agents who are matched in the interval $[t_0, t_0 + \epsilon]$. By definition the social welfare of an online algorithm, we have:

$$\mathbf{W}(\text{Patient}) = \mathbb{E}\left[\frac{1}{T}\int_{t=0}^{T} W_{t,t+dt} dt\right] = \frac{1}{T}\int_{t=0}^{T} \mathbb{E}\left[W_{t,t+dt}\right] dt$$

All agents are equally likely to become critical at each moment. From the perspective of the planner, all agents are equally likely to be the neighbor of a critical agent. Hence, the *expected* utility of each of the agents who are matched at time t under the Patient algorithm is X_t/Z_t. Thus,

$$\mathbf{W}(\text{Patient}) = \frac{1}{mT}\int_{t=0}^{T} \mathbb{E}\left[2\frac{X_t}{Z_t} Z_t (1 - (1 - d/m)^{Z_t}) dt\right] = \frac{2}{mT}\int_{t=0}^{T} \mathbb{E}\left[X_t (1 - (1 - d/m)^{Z_t})\right] dt \tag{F.1}$$

First, we prove the following lemma.

Lemma F.3. *For any $\epsilon < 1/2m^2$ and $t \geq \tau_{\text{mix}}(\epsilon)$,*

$$\mathbb{E}\left[X_t\right]\left(1 - q^{k^* + \tilde{O}(\sqrt{m})}\right) - O(m^{-1/2}) \leq \mathbb{E}\left[X_t(1 - q^{Z_t})\right] \leq \mathbb{E}\left[X_t\right]\left(1 - q^{k^* - \tilde{O}(\sqrt{m})}\right) + O(m^{-1/2}).$$

Proof. Let $\Delta := 3\sqrt{m}\log(m)$. Let $\mathcal{E}$ be the event that $Z_t \in [k^* - \Delta, k^* + \Delta]$. First, we show the following inequality and then we upper-bound $\mathbb{E}\left[X_t|\overline{\mathcal{E}}\right]\mathbb{P}\left[\overline{\mathcal{E}}\right]$.

$$\mathbb{E}\left[X_t\right](1-q^{k^*+\Delta})-\mathbb{E}\left[X_t|\overline{\mathcal{E}}\right]\mathbb{P}\left[\overline{\mathcal{E}}\right]\leq\mathbb{E}\left[X_t(1-q^{Z_t})\right]\leq\mathbb{E}\left[X_t\right](1-q^{k^*-\Delta})+\mathbb{E}\left[X_t|\overline{\mathcal{E}}\right]\mathbb{P}\left[\overline{\mathcal{E}}\right] \tag{F.2}$$

We prove the right inequality and the left can be proved similarly.

By definition of expectation,

$$\begin{aligned}
\mathbb{E}\left[X_t(1-q^{Z_t})\right] &= \mathbb{E}\left[X_t(1-q^{Z_t})|\mathcal{E}\right]\cdot\mathbb{P}\left[\mathcal{E}\right]+\mathbb{E}\left[X_t(1-q^{Z_t})|\overline{\mathcal{E}}\right]\cdot\mathbb{P}\left[\overline{\mathcal{E}}\right] \\
&\leq \mathbb{E}\left[X_t|\mathcal{E}\right](1-q^{k^*-\Delta})+\mathbb{E}\left[X_t(1-q^{Z_t})|\overline{\mathcal{E}}\right]\cdot\mathbb{P}\left[\overline{\mathcal{E}}\right]
\end{aligned}$$

Now, for any random variable X and any event $\mathcal{E}$ we have $\mathbb{E}\left[X|\mathcal{E}\right]\cdot\mathbb{P}\left[\mathcal{E}\right]=\mathbb{E}\left[X\right]-\mathbb{E}\left[X|\overline{\mathcal{E}}\right]\cdot\mathbb{P}\left[\overline{\mathcal{E}}\right]$. Therefore,

$$\begin{aligned}
\mathbb{E}\left[X_t(1-q^{Z_t})\right] &\leq (1-q^{k^*-\Delta})(\mathbb{E}\left[X_t\right]-\mathbb{E}\left[X_t|\overline{\mathcal{E}}\right]\cdot\mathbb{P}\left[\overline{\mathcal{E}}\right])+\mathbb{E}\left[X_t(1-q^{Z_t})|\overline{\mathcal{E}}\right]\cdot\mathbb{P}\left[\overline{\mathcal{E}}\right] \\
&\leq \mathbb{E}\left[X_t\right](1-q^{k^*-\Delta})+\mathbb{E}\left[X_t|\overline{\mathcal{E}}\right]\mathbb{P}\left[\overline{\mathcal{E}}\right]
\end{aligned}$$

where we simply used the non-negativity of X_t and that $(1-q^{k^*-\Delta})\leq 1$. This proves the right inequality of (F.2). The left inequality can be proved similarly.

It remains to upper-bound $\mathbb{E}\left[X_t|\overline{\mathcal{E}}\right]\mathbb{P}\left[\overline{\mathcal{E}}\right]$. Let $\pi(.)$ be the stationary distribution of the Markov Chain Z_t. Since by definition of X_t, $X_t\leq Z_t$ with probability 1,

$$\begin{aligned}
\mathbb{E}\left[X_t|\overline{\mathcal{E}}\right]\mathbb{P}\left[\overline{\mathcal{E}}\right] &\leq \mathbb{E}\left[Z_t|\overline{\mathcal{E}}\right]\mathbb{P}\left[\overline{\mathcal{E}}\right] \\
&\leq \sum_{i=0}^{k^*-\Delta} i(\pi(i)+\epsilon)+\sum_{i=k^*+\Delta}^{6m} i(\pi(i)+\epsilon)+\sum_{i=6m+1}^{\infty} i\cdot\mathbb{P}\left[\tilde{Z}_t=i\right]
\end{aligned}$$

where the last term uses the fact that Z_t is at most the size of the pool of the inactive policy at time t, i.e., $\mathbb{P}\left[Z_t=i\right]\leq\mathbb{P}\left[\tilde{Z}_t=i\right]$ for all $i>0$. We bound the first term of RHS using Proposition 5.9, the second term using (E.12) and the last term using Proposition B.3.

$$\mathbb{E}\left[X_t|\overline{\mathcal{E}}\right]\mathbb{P}\left[\overline{\mathcal{E}}\right]\leq\frac{4}{\sqrt{m}}+6m\epsilon+\sum_{i=6m}^{\infty}\frac{m^i}{i!}\leq\frac{4}{\sqrt{m}}+\frac{3}{m}+2^{-6m}.$$

□

It remains to estimate $\mathbb{E}\left[X_t\right]$. This is done in the following lemma.

Lemma F.4. *For any $\epsilon < 1/2m^2$, $t_1 \geq 16\log(m)\log(4/\epsilon) \geq 2\tau_{\text{mix}}(\epsilon)$,*

$$\frac{m - O(m^{-1/2})}{2 + r - q^{k^* - \tilde{O}(\sqrt{m})}} \leq \mathbb{E}\left[X_{t_1}\right] \leq \frac{m + O(m^{-1/2})}{2 + r - q^{k^* + \tilde{O}(\sqrt{m})}}$$

Proof. Let $\eta > 0$ be very close to zero (eventually we let $\eta \to 0$). Since we have a $(m, d, 1)$ matching market, using equation (C.1) for any $t \geq 0$ we have,

$$\mathbb{E}\left[X_{t+\eta} | X_t, Z_t\right] = X_t(e^{-\eta r}) + m\eta - \eta Z_t\Big(\frac{X_t}{Z_t} q^{Z_t}\Big) - 2\eta Z_t\Big(\frac{X_t}{Z_t}(1 - q^{Z_t})\Big) \pm O(\eta^2)$$

The first term in the RHS follows from the exponential discount in the utility of the agents in the pool. The second term in the RHS stands for the new arrivals. The third term stands for the perished agents and the last term stands for the the matched agents. We use the notation $A = B \pm C$ to denote $B - C \leq A \leq B + C$.

We use $e^{-x} = 1 - x + O(x^2)$ and rearrange the equation to get,

$$\mathbb{E}\left[X_{t+\eta} | X_t, Z_t\right] = m\eta + X_t - \eta(1 + r)X_t - \eta X_t(1 - q^{Z_t}) \pm O(\eta^2).$$

Using Lemma F.3 for any $t \geq \tau_{\text{mix}}(\epsilon)$ we can estimate $\mathbb{E}\left[X_t(1 - q^{Z_t})\right]$. Taking expectation from both sides of the above inequality we get,

$$\frac{\mathbb{E}\left[X_{t+\epsilon}\right] - \mathbb{E}\left[X_t\right]}{\eta} = m - \mathbb{E}\left[X_t\right](2 + \delta - q^{k^* \pm \tilde{O}(\sqrt{m})}) \pm O(m^{-1/2}) - O(\eta)$$

Letting $\eta \to 0$, and solving the above differential equation from $\tau_{\text{mix}}(\epsilon)$ to t_1 we get

$$\mathbb{E}\left[X_{t_1}\right] = \frac{m \pm O(m^{-1/2})}{2 + \delta - q^{k^* \pm \tilde{O}(\sqrt{m})}} + C_1 \exp\Big(-(\delta + 2 - q^{k^* \pm \tilde{O}(\sqrt{m})})(t_1 - \tau_{\text{mix}}(\epsilon))\Big).$$

Now, for $t_1 = \tau_{\text{mix}}(\epsilon)$ we use the initial condition $\mathbb{E}\left[X_{\tau_{\text{mix}}(\epsilon)}\right] \leq \mathbb{E}\left[Z_{\tau_{\text{mix}}(\epsilon)}\right] \leq m$, and we can let $C_1 \leq m$. Finally, since $t_1 \geq 2\tau_{\text{mix}}(\epsilon)$ and $t_1/2 \geq 2\log(m)$ we can upper-bound the latter term with $O(m^{-1/2})$.

□

Let $T_0 = 16\log(m)\log(4/\epsilon)$. Since (for any matching algorithm) the sum of the utilities of the agents that leave the market before time T_0 is at most mT_0 in expectation, by the above two lemmas, we can write

$$\begin{aligned}
\mathbf{W}(\text{Patient}) &= \frac{2}{mT}\Big(mT_0 + \int_{T_0}^{T} \mathbb{E}\left[X_t(1-q^{Z_t})\right] dt\Big) \\
&\le \frac{2T_0}{T} + \frac{2}{mT}\int_{T_0}^{T}\Big(\frac{m(1-q^{k^*-\tilde{O}(\sqrt{m})})}{r+2-q^{k^*+\tilde{O}(\sqrt{m})}} + \tilde{O}(m^{-1/2})\Big)dt \\
&\le \frac{2T_0}{T} + \frac{T-T_0}{T}\Big(\frac{1-q^{k^*-\tilde{O}(\sqrt{m})}}{1+r/2-\frac{1}{2}q^{k^*+\tilde{O}(\sqrt{m})}} + \tilde{O}(m^{-3/2})\Big)
\end{aligned}$$

Similarly, since the sum of the utilities of the agents that leave the market by time T_0 is non-negative, we can show that

$$\mathbf{W}(\text{Patient}) \ge \frac{T-T_0}{T}\Big(\frac{1-q^{k^*+\tilde{O}(\sqrt{m})}}{1+r/2-\frac{1}{2}q^{k^*-\tilde{O}(\sqrt{m})}} - \tilde{O}(m^{-3/2})\Big)$$

F.2 Welfare of the Greedy Algorithm

Here, we upper-bound the welfare of the optimum online algorithm, OPT, and that immediately upper-bounds the welfare of the Greedy algorithm. Recall that by Theorem B.1, for any $T > 0$, $1/(2d+1+d^2/m)$ fraction of the agents perish in OPT. On the other hand, by the definition of utility, we receive a utility at most 1 from any matched agent. Therefore, even if all of the matched agents receive a utility of 1, (for any $r \ge 0$)

$$\mathbf{W}(\text{Greedy}) \le \mathbf{W}(\text{OPT}) \le 1 - \frac{1}{2d+1+d^2/m}.$$

F.3 Optimal α

One can show that as $m, T \to \infty$, then the closed form expression for the welfare of the $\lim_{m,T\to\infty} \mathbf{W}(\text{Patient}(\alpha)) \simeq \frac{2}{2-e^{-\bar{d}/\bar{\alpha}}+r/\bar{\alpha}}(1-e^{-\bar{d}/\bar{\alpha}})$, where we have used the fact that $k^* \in [m/2-2, m-1]$ with very high probability and for large values of m, $(1-d/m)^{k^*} \to e^{-dk^*/m}$. Hence, it is enough to find the value of $\bar{\alpha}$ that maximizes $\mathbf{W}(\text{Patient}(\alpha))$.

The first-order condition implies that the optimal value of $\bar{\alpha}$ solves $g(\bar{\alpha}) = 0$, where

$$g(\bar{\alpha}) = r - \Big(r + \bar{d} + \frac{\bar{d}r}{\bar{\alpha}}\Big)e^{-\frac{\bar{d}}{\bar{\alpha}}}.$$

To have an estimate of $\bar{\alpha}$, note that $g(.)$ is an increasing function of $\bar{\alpha}$ and $g(\bar{d}/\log(\bar{d}/r)) < 0$ and if $r < \bar{d}/4$, then $g(\bar{d}/\log(2\bar{d}/r)) > 0$. Thus, we must have that the solution of $g(\bar{\alpha}) = 0$

satisfies:

$$\bar{d}/\log(\bar{d}/r) \leq \bar{\alpha} \leq \bar{d}/\log(2\bar{d}/r)$$

F.4 Welfare Comparative Statics

Assume $r < \frac{d}{2(e-1)}$. We know that, as $T \to \infty$, $m \to \infty$, $\mathbf{W}(\text{Patient}(\alpha)) \simeq \frac{2}{2-e^{-d/2\bar{\alpha}}+r/\bar{\alpha}}(1-e^{-d/2\bar{\alpha}})$, where $\bar{\alpha} = 1 + 1/\alpha$. Note that $\bar{\alpha} \in [1, \infty)$.

Define

$$\begin{aligned}\tilde{\mathbf{W}}(\text{Patient}(\alpha)) &\equiv \frac{2}{2 - e^{-d/2\bar{\alpha}} + r/\bar{\alpha}}(1 - e^{-d/2\bar{\alpha}}) \\ &= \frac{2}{\frac{1+\frac{r}{\bar{\alpha}}}{1-e^{-\frac{d}{2\bar{\alpha}}}} + 1}.\end{aligned}$$

Maximizing $\tilde{\mathbf{W}}(\text{Patient}(\alpha))$ is equivalent to maximizing (with respect to $\bar{\alpha}$):

$$Q(\bar{\alpha}, d, r) \equiv \frac{1 - e^{-\frac{d}{2\bar{\alpha}}}}{1 + \frac{r}{\bar{\alpha}}}$$

$$\frac{\partial Q}{\partial \bar{\alpha}} = \frac{e^{-\frac{d}{2\bar{\alpha}}}(\bar{\alpha}(2r(e^{\frac{d}{2\bar{\alpha}}} - 1) - d) - dr)}{2\bar{\alpha}(\bar{\alpha} + r)^2}$$

Since $r < \frac{d}{2(e-1)}$, if $\bar{\alpha} \geq \frac{d}{2}$, then $\frac{\partial Q}{\partial \bar{\alpha}} < 0$; so the value of $\bar{\alpha}$ that maximizes Q is never above $\frac{d}{2}$. Thus, maximizing Q is equivalent to maximizing (with respect to $\bar{\alpha}$):

$$\hat{Q}(\bar{\alpha}, d, r) \equiv \begin{cases} Q(\bar{\alpha}, d, r) & \text{if } \bar{\alpha} < \frac{d}{2} \\ Q(\frac{d}{2}, d, r) & \text{otherwise} \end{cases}$$

Define $\bar{\alpha}^* \equiv \arg\max_{\bar{\alpha}} \hat{Q}(\bar{\alpha}, d, r) = \arg\max_{\bar{\alpha}} Q(\bar{\alpha}, d, r)$. Now note that

$$\frac{\partial Q}{\partial d} = \frac{e^{-\frac{d}{2\bar{\alpha}}}(\frac{1}{2\bar{\alpha}})}{1 + \frac{r}{\bar{\alpha}}} = \frac{e^{-\frac{d}{2\bar{\alpha}}}}{2(\bar{\alpha} + r)}$$

As a result,

$$\begin{aligned}\frac{\partial^2 Q}{\partial d \partial \bar{\alpha}} &= \frac{e^{-\frac{d}{2\bar{\alpha}}}}{2(\bar{\alpha}+r)^2}((\bar{\alpha}+r)\frac{d}{2\bar{\alpha}^2}-1) \\ &= \frac{e^{-\frac{d}{2\bar{\alpha}}}}{2(\bar{\alpha}+r)^2}((\bar{\alpha}+r)\frac{d}{2\bar{\alpha}^2}-1) \\ &= \frac{e^{-\frac{d}{2\bar{\alpha}}}}{2(\bar{\alpha}+r)^2}(\frac{d}{2\bar{\alpha}}+\frac{rd}{2\bar{\alpha}^2}-1)\end{aligned} \tag{F.3}$$

The right-hand side of Equation F.3 is positive if $\bar{\alpha} < d/2$. $\hat{Q}(\bar{\alpha}, d, r)$ is differentiable almost everywhere and continuous. Consequently, $\hat{Q}(\bar{\alpha}, d, r)$ has increasing differences in $(\bar{\alpha}, d)$. Thus, by Topkis' Theorem, $\bar{\alpha}^*$ is weakly increasing in d.

Similarly, maximizing $\tilde{\mathbf{W}}$ is equivalent to maximizing

$$\begin{aligned}R(\bar{\alpha}, d, r) &\equiv -\frac{1+\frac{r}{\bar{\alpha}}}{1-e^{-\frac{d}{2\bar{\alpha}}}} \\ \Rightarrow \frac{\partial R}{\partial r} &= -\frac{1}{\bar{\alpha}(1-e^{-\frac{d}{2\bar{\alpha}}})}\end{aligned}$$

As a result,

$$\begin{aligned}\frac{\partial^2 R}{\partial \bar{\alpha} \partial r} &= \frac{(1-e^{-\frac{d}{2\bar{\alpha}}}) - \frac{d}{2\bar{\alpha}}e^{-\frac{d}{2\bar{\alpha}}}}{\bar{\alpha}^2(1-e^{-\frac{d}{2\bar{\alpha}}})^2} \\ &= \frac{e^{-\frac{d}{2\bar{\alpha}}}(e^{\frac{d}{2\bar{\alpha}}}-1-\frac{d}{2\bar{\alpha}})}{\bar{\alpha}^2(1-e^{-\frac{d}{2\bar{\alpha}}})^2}\end{aligned} \tag{F.4}$$

Observe that $e^{\frac{d}{2\bar{\alpha}}} - 1 - \frac{d}{2\bar{\alpha}} = 0$ when $d = 0$. Moreover, $\frac{\partial}{\partial d}(e^{\frac{d}{2\bar{\alpha}}} - 1 - \frac{d}{2\bar{\alpha}} = 0) \geq 0$ for $d > 0$. Thus, $e^{\frac{d}{2\bar{\alpha}}} - 1 - \frac{d}{2\bar{\alpha}} \geq 0$ Consequently, $\frac{\partial^2 R}{\partial \bar{\alpha} \partial r} \geq 0$, which entails that $R(\bar{\alpha}, d, r)$ has increasing differences in $(\bar{\alpha}, r)$. Thus, by Topkis' theorem, $\bar{\alpha}^*$ is weakly increasing in r.

G Proofs from Section 6

G.1 Proof of Lemma 6.4

In this section, we present the full proof of Lemma 6.4. We prove the lemma by writing a closed form expression for the utility of a and then upper-bounding that expression.

In the following claim we study the probability a is matched in the interval $[t, t+\epsilon]$ and the probability that it leaves the market in that interval.

Claim G.1. *For any time $t \geq 0$, and $\epsilon > 0$,*

$$\mathbb{P}\left[a \in M_{t,t+\epsilon}\right] = \epsilon \cdot \mathbb{P}\left[a \in A_t\right](2 + c(t))\mathbb{E}\left[1 - (1 - d/m)^{Z_t}|a \in A_t\right] \pm O(\epsilon^2) \tag{G.1}$$

$$\mathbb{P}\left[a \notin A_{t+\epsilon}, a \in A_t\right] = \mathbb{P}\left[a \in A_t\right](1 - \epsilon(1 + c(t) + \mathbb{E}\left[1 - (1 - d/m)^{Z_t - 1}|a \in A_t\right]) \pm O(\epsilon^2)) \tag{G.2}$$

Proof. The claim follows from two simple observations. First, a becomes critical in the interval $[t, t+\epsilon]$ with probability $\epsilon \cdot \mathbb{P}\left[a \in A_t\right](1 + c(t))$ and if he is critical he is matched with probability $\mathbb{E}\left[(1 - (1 - d/m)^{Z_t - 1}|a \in A_t\right]$. Second, a may also get matched (without getting critical) in the interval $[t, t+\epsilon]$. Observe that if an agent $b \in A_t$ where $b \neq a$ gets critical she will be matched with a with probability $(1 - (1 - d/m)^{Z_t - 1})/(Z_t - 1),$. Therefore, the probability that a is matched at $[t, t+\epsilon]$ without getting critical is

$$\begin{aligned}\mathbb{P}\left[a \in A_t\right] \cdot \mathbb{E}&\left[\epsilon \cdot (Z_t - 1)\frac{1 - (1 - d/m)^{Z_t - 1}}{Z_t - 1}|a \in A_t\right] \\ &= \epsilon \cdot \mathbb{P}\left[a \in A_t\right]\mathbb{E}\left[1 - (1 - d/m)^{Z_t - 1}|a \in A_t\right]\end{aligned}$$

The claim follows from simple algebraic manipulations. □

We need to study the conditional expectation $\mathbb{E}\left[1 - (1 - d/m)^{Z_t - 1}|a \in A_t\right]$ to use the above claim. This is not easy in general; although the distribution of Z_t remains stationary, the distribution of Z_t conditioned on $a \in A_t$ can be a very different distribution. So, here we prove simple upper and lower bounds on $\mathbb{E}\left[1 - (1 - d/m)^{Z_t - 1}|a \in A_t\right]$ using the concentration properties of Z_t. By the assumption of the lemma Z_t is at stationary at any time $t \geq 0$. Let k^* be the number defined in Proposition 5.9, and $\beta = (1 - d/m)^{k^*}$. Also, let

$\sigma := \sqrt{6\log(8m/\beta)}$. By Proposition 5.9, for any $t \geq 0$,

$$\begin{aligned}
\mathbb{E}\left[1-(1-d/m)^{Z_t-1}|a\in A_t\right] &\leq \mathbb{E}\left[1-(1-d/m)^{Z_t-1}|Z_t < k^*+\sigma\sqrt{4m}, a\in A_t\right] \\
&\quad + \mathbb{P}\left[Z_t \geq k^*+\sigma\sqrt{4m}|a\in A_t\right] \\
&\leq 1-(1-d/m)^{k^*+\sigma\sqrt{4m}} + \frac{\mathbb{P}\left[Z_t\geq k^*+\sigma\sqrt{4m}\right]}{\mathbb{P}\left[a\in A_t\right]} \\
&\leq 1-\beta+\beta(1-(1-d/m)^{\sigma\sqrt{4m}}) + \frac{8\sqrt{m}e^{-\sigma^2/3}}{\mathbb{P}\left[a\in A_t\right]} \\
&\leq 1-\beta+\frac{2\sigma d\beta}{\sqrt{m}} + \frac{\beta}{m^2\cdot\mathbb{P}\left[a\in A_t\right]} \qquad \text{(G.3)}
\end{aligned}$$

In the last inequality we used (A.4) and the definition of σ. Similarly,

$$\begin{aligned}
\mathbb{E}\left[1-(1-d/m)^{Z_t-1}|a\in A_t\right] &\geq \mathbb{E}\left[1-(1-d/m)^{Z_t-1}|Z_t \geq k^*-\sigma\sqrt{4m}, a\in A_t\right] \\
&\quad \cdot \mathbb{P}\left[Z_t \geq k^*-\sigma\sqrt{4m}|a\in A_t\right] \\
&\geq (1-(1-d/m)^{k^*-\sigma\sqrt{4m}})\frac{\mathbb{P}\left[a\in A_t\right]-\mathbb{P}\left[Z_t< k^*-\sigma\sqrt{4m}\right]}{\mathbb{P}\left[a\in A_t\right]} \\
&\geq 1-\beta-\beta((1-d/m)^{-\sigma\sqrt{4m}}-1) - \frac{2\sqrt{m}e^{-\sigma^2}}{\mathbb{P}\left[a\in A_t\right]} \\
&\geq 1-\beta-\frac{4d\sigma\beta}{\sqrt{m}} - \frac{\beta^3}{m^3\cdot\mathbb{P}\left[a\in A_t\right]} \qquad \text{(G.4)}
\end{aligned}$$

where in the last inequality we used (A.4), the assumption that $2d\sigma \leq \sqrt{m}$ and the definition of σ.

Next, we write a closed form upper-bound for $\mathbb{P}\left[a\in A_t\right]$. Choose t^* such that $\int_{t=0}^{t^*}(2+c(t))dt = 2\log(m/\beta)$. Observe that $t^*\leq \log(m/\beta)\leq \sigma^2/6$. Since a leaves the market with rate at least $1+c(t)$ and at most $2+c(t)$, we can write

$$\frac{\beta^2}{m^2} = \exp\Big(-\int_{t=0}^{t^*}(2+c(t))dt\Big) \leq \mathbb{P}\left[a\in A_{t^*}\right] \leq \exp\Big(-\int_{t=0}^{t^*}(1+c(t))dt\Big) \leq \frac{\beta}{m} \qquad \text{(G.5)}$$

Intuitively, t^* is a moment where the expected utility of that a receives in the interval $[t^*,\infty)$ is negligible, i.e., in the best case it is at most β/m.

By Claim G.1 and (G.4), for any $t \le t^*$,

$$\begin{aligned}\frac{\mathbb{P}\left[a \in A_{t+\epsilon}\right] - \mathbb{P}\left[a \in A_t\right]}{\epsilon} &\le -\mathbb{P}\left[a \in A_t\right]\Big(2 + c(t) - \beta - \frac{4d\sigma\beta}{\sqrt{m}} - \frac{\beta^3}{m^3 \cdot \mathbb{P}\left[a \in A_t\right]} \pm O(\epsilon)\Big)\\ &\le -\mathbb{P}\left[a \in A_t\right]\Big(2 + c(t) - \beta - \frac{5d\sigma\beta}{\sqrt{m}} \pm O(\epsilon)\Big)\end{aligned}$$

where in the last inequality we used (G.5). Letting $\epsilon \to 0$, for $t \le t^*$, the above differential equation yields,

$$\mathbb{P}\left[a \in A_t\right] \le \exp\Big(-\int_{\tau=0}^{t}\Big(2+c(\tau)-\beta-\frac{5d\sigma\beta}{\sqrt{m}}\Big)d\tau\Big) \le \exp\Big(-\int_{\tau=0}^{t}(2+c(\tau)-\beta)d\tau\Big) + \frac{2d\sigma^3\beta}{\sqrt{m}}. \tag{G.6}$$

where in the last inequality we used $t^* \le \sigma^2/6$, $e^x \le 1+2x$ for $x \le 1$ and lemma's assumption $5d\sigma^2 \le \sqrt{m}$.

Now, we are ready to upper-bound the utility of a. By (G.5) the expected utility that a gains after t^* is no more than β/m. Therefore,

$$\begin{aligned}\mathbb{E}\left[u_c(a)\right] &\le \frac{\beta}{m} + \int_{t=0}^{t^*}(2+c(t))\mathbb{E}\left[1-(1-d/m)^{Z_t-1}|a \in A_t\right]\mathbb{P}\left[a \in A_t\right]e^{-rt}dt\\ &\le \frac{\beta}{m} + \int_{t=0}^{t^*}(2+c(t))((1-\beta)\mathbb{P}\left[a \in A_t\right] + \beta/\sqrt{m})e^{-rt}dt\\ &\le \frac{\beta}{m} + \int_{t=0}^{t^*}(2+c(t))\Big((1-\beta)\exp\Big(-\int_{\tau=0}^{t}(2+c(\tau)-\beta)d\tau\Big) + \frac{3d\sigma^3}{\sqrt{m}}\beta\Big)e^{-rt}dt\\ &\le \frac{2d\sigma^5}{\sqrt{m}}\beta + \int_{t=0}^{\infty}(1-\beta)(2+c(t))\exp\Big(-\int_{\tau=0}^{t}(2+c(\tau)-\beta)d\tau\Big)e^{-rt}dt.\end{aligned}$$

In the first inequality we used equation (G.3), in second inequality we used equation (G.6), and in the last inequality we use the definition of t^*. We have finally obtained a closed form upper-bound on the expected utility of a.

Let $U_c(a)$ be the right hand side of the above equation. Next, we show that $U_c(a)$ is maximized by letting $c(t) = 0$ for all t. This will complete the proof of Lemma 6.4. Let c be a function that maximizes $U_c(a)$ which is not equal to zero. Suppose $c(t) \neq 0$ for some $t \ge 0$. We define a function $\tilde{c} : \mathbb{R}_+ \to \mathbb{R}_+$ and we show that if $r < \beta$, then $U_{\tilde{c}}(a) > U_c(a)$.

Let $\tilde{c}$ be the following function,

$$\tilde{c}(\tau) = \begin{cases} c(\tau) & \text{if } \tau < t, \\ 0 & \text{if } t \leq \tau \leq t+\epsilon, \\ c(\tau) + c(\tau - \epsilon) & \text{if } t+\epsilon \leq \tau \leq t+2\epsilon, \\ c(\tau) & \text{otherwise.} \end{cases}$$

In words, we push the mass of $c(.)$ in the interval $[t, t+\epsilon]$ to the right. We remark that the above function $\tilde{c}(.)$ is not necessarily continuous so we need to smooth it out. The latter can be done without introducing any errors and we do not describe the details here. Let $S := \int_{\tau=0}^{t}(1 + c(t) + \beta)d\tau$. Assuming $\tilde{c}'(t) \ll 1/\epsilon$, we have

$$\begin{aligned} U_{\tilde{c}}(a) - U_c(a) &\geq -\epsilon \cdot c(t)(1-\beta)e^{-S}e^{-rt} + \epsilon \cdot c(t)(1-\beta)e^{-S-\epsilon(2-\beta)}e^{-r(t+\epsilon)} \\ &\quad + \epsilon(1-\beta)(2 + c(t+\epsilon))(e^{-S-\epsilon(2-\beta)}e^{-r(t+\epsilon)} - e^{-S-\epsilon(2+c(t)-\beta)}e^{-r(t+\epsilon)}) \\ &= -\epsilon^2 \cdot c(t)(1-\beta)e^{-S-rt}(2-\beta+r) + \epsilon^2(1-\beta)(2+c(t+\epsilon))e^{-S-rt}c(t) \\ &\geq \epsilon^2 \cdot (1-\beta)e^{-S-rt}c(t)(\beta - r). \end{aligned}$$

Since $r < \beta$ by the lemma's assumption, the maximizer of $U_c(a)$ is the all zero function. Therefore, for any well-behaved function $c(.)$,

$$\begin{aligned} \mathbb{E}\left[u_c(a)\right] &\leq \frac{2d\sigma^5}{\sqrt{m}}\beta + \int_{t=0}^{\infty} 2(1-\beta)\exp\Big(-\int_{\tau=0}^{t}(2-\beta)d\tau\Big)e^{-rt}dt \\ &\leq O\Big(\frac{d^4 \log^3(m)}{\sqrt{m}}\Big)\beta + \frac{2(1-\beta)}{2-\beta+r}. \end{aligned}$$

In the last inequality we used that $\sigma = O(\sqrt{\log(m/\beta)})$ and $\beta \leq e^{-d}$. This completes the proof of Lemma 6.4.

H Small Market Simulations

In Proposition 5.5 and Proposition 5.9, we prove that the Markov chains of the Greedy and Patient algorithms are highly concentrated in intervals of size $O(\sqrt{m/d})$ and $O(\sqrt{m})$, respectively. These intervals are plausible concentration bounds when m is relatively large. In fact, most of our theoretical results are interesting when markets are relatively large. Therefore, it is natural to ask: What if m is relatively small? And what if the d is not small

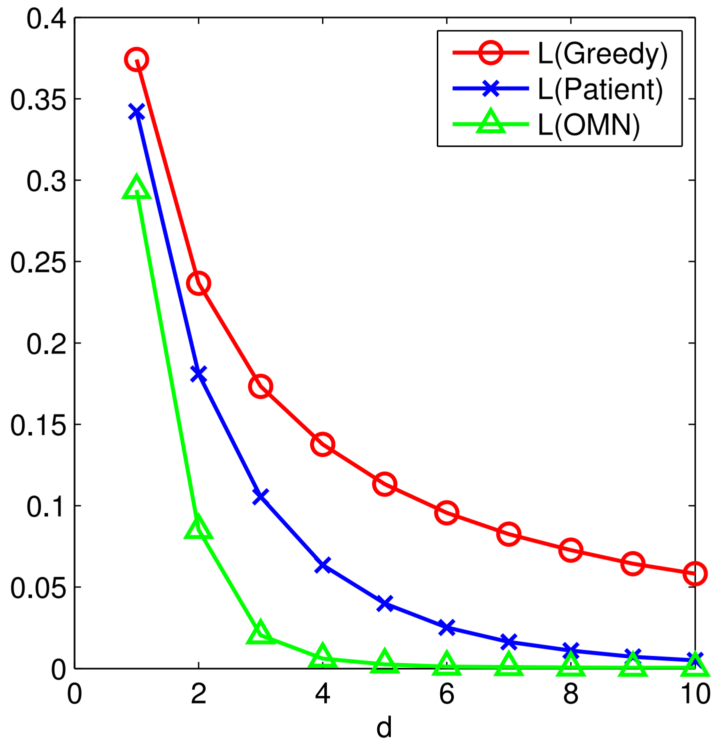

Figure 8: Simulated Losses for $m = 20$. For very small market sizes and even for relatively large values of d, the Patient algorithm outperforms the Greedy Algorithm.

relative to m?

Figure 8 depicts the simulation results of our model for small m and small T. We simulated the market for $m = 20$ and $T = 100$ periods, repeated this process for 500 iterations, and computed the average loss for the Greedy, Patient, and the Omniscient algorithms. As it is clear from the simulation results, the loss of the Patient algorithm is lower than the Greedy for any d, and in particular, when d increases, the Patient algorithm's performance gets closer and closer to the Omniscient algorithm, whereas the Greedy algorithm's loss remains far above both of them.