

A dynamic model of barter exchange

Ross Anderson Itai Ashlagi* David Gamarnik Yash Kanoria†

May 19, 2014

Abstract

We consider the problem of efficient operation of a barter exchange platform for indivisible goods. We introduce a dynamic model of barter exchange where in each period one agent arrives with a single item she wants to exchange for a different item. We study a homogeneous and stochastic environment: an agent is interested in the item possessed by another agent with probability p , independently for all pairs of agents. We consider three settings with respect to the types of allowed exchanges: a) Only two-way cycles, in which two agents swap their items, b) Two or three-way cycles, c) (unbounded) chains initiated by altruistic donors who provide an item but expect nothing in return. The goal of the platform is to minimize the average waiting time of an agent.

Somewhat surprisingly, we find that in each of these settings, a policy that conducts exchanges in a greedy fashion is near optimal, among a large class of policies that includes batching policies. Further, we find that for small p , allowing three-cycles can greatly improve the waiting time over the two-cycles only setting, and the presence of altruistic donors can lead to a further large improvement in average waiting time. Specifically, we find that a greedy policy achieves an average waiting time of $\Theta(1/p^2)$ in setting a), $\Theta(1/p^{3/2})$ in setting b), and $\Theta(1/p)$ in setting c). Thus, a platform can achieve the smallest waiting times by using a greedy policy, and by facilitating three cycles and chains, if possible.

Our findings are consistent with empirical and computational observations which compare batching policies in the context of kidney exchange programs.

Keywords: barter, matching, market design, random graphs, dynamics, kidney exchange

1 Introduction

A marketplace for barter exchange provides opportunities for agents to exchange items directly, without monetary payments. There is a growing number of such marketplaces facilitating the exchange of a variety of items. We consider the problem faced by an exchange

*Ashlagi acknowledges the research support of the National Science Foundation grant SES-1254768.

†Corresponding author: ykanoria@columbia.edu

platform for indivisible items, that seeks to enable users to complete a desirable trade as early as possible. A major lever a platform operator has at her disposal is the *policy* employed in conducting exchanges, for example, greedy, batching, etc. We study which policy should be used in order to minimize the waiting time of users in the face of stochastic individual demands for items. We investigate this question under a variety of settings determined by the feasible *types* of exchanges, which are often driven by the technology adopted by the marketplace.

1.1 Background and motivation

A number of barter exchange platforms exist for swapping a variety of items. For instance, www.homeexchange.com and www.ReadItSwapIt.com¹ are decentralized marketplaces that enable pairwise swapping by mutual agreement of homes for vacation and books respectively. Finding a pairwise exchange is challenging because it requires two users to each possess an item that the other desires (Jevons, 1876). www.Swap.com, which allows for swapping (and selling) of preowned childrens items, books and DVDs, is a centralized platform². Centralization allows it to execute multi-way exchanges, where each agent in the cycle has an item that the next agent in the cycle desires, cf. Figure 1 below, to improve the likelihood of finding available exchanges.³

Kidney exchange clearinghouses have also noticed the limitations of pairwise exchanges (Roth et al., 2007).⁴ Kidney exchange can be a suitable solution when a healthy person (with two kidneys) wishes to donate one to a friend or family member but is biologically incompatible with the intended recipient.⁵ An incompatible donor-patient pair can exchange their donor’s kidney with one or more other such pairs in a cycle so that every patient receives a compatible kidney.⁶ Kidney exchange is a growing market and there are multiple clearinghouses such as the Alliance for Paired Donation (APD) and the National Kidney Registry (NKR) in the United States. Since it is desired that a donor gives her kidney no later than her associated patient receives a kidney, cyclic exchanges are conducted simultaneously and are therefore rarely longer than three.⁷ On the other hand, chains initiated by an altruistic donor (cf. Figure 1) can be organized non-simultaneously while satisfying this constraint (Roth et al. (2006)), allowing them to be unbounded in length.

Each of these markets evolve dynamically with agents arriving over time. At any point in

¹There are numerous other similar websites such as www.tradeaway.com, www.barterquest.com etc.

²Users send items to a central location where the market operator photographs, lists and stores them, and later ships them out.

³Quoting from the website: “...Netcyler has developed unique technology for enabling multi-party swapping of any type of items. Mathematically speaking the Netcyler Trade Ring technology can make swaps three orders of magnitude more likely compared to traditional one-to-one swaps...”

⁴In the early stages (approximately until 2005 in the US) only pairwise exchanges were conducted (Roth et al., 2005).

⁵In kidney exchange, an incompatible donor-patient pair can be thought of as a single agent.

⁶It is illegal to buy and sell organs in most countries and several other countries, making *exchange* a key avenue for many patients to obtain organs.

⁷A four-way cycle requires 8 simultaneous surgeries, 4 nephrectomies and 4 transplants.

time, there is a pool of agents waiting for an exchange, and possibly some feasible exchanges supported by the platform technology. In centralized platforms, the *policy* employed by the platform determines which feasible exchanges are executed, and when. We ask which policy centralized platforms should use, given the types of supported exchanges.⁸ One natural option is the *Greedy* policy, where the platform executes a feasible exchange as soon as the possibility emerges. Alternatively, the platform can adopt a *Batching* policy where it waits for a ‘batch’ of agents to accumulate, and then identifies a set of exchanges that maximizes the number of matched agents, with the remaining agents carried forward into the next batch. More complex policies are possible as well. Kidney exchange clearinghouses in the US have experimented with a variety of batch lengths in the past, and have subsequently shifted to very short batches resembling a greedy approach.⁹ When does the greedy policy perform well? What is the optimal policy, and how does it depend on the feasible types of exchanges? We address these questions in this paper.

We simultaneously also quantify the benefits resulting from facilitating multi-way exchanges, and chains, relative to allowing only pairwise exchanges, by comparing the performance of the optimal policy in each case. (In choosing its technology, a platform would compare such benefits with the cost of facilitating additional types of exchanges. We do not model this cost.)

1.2 Approach and contributions

We consider a stylized dynamic model of agent arrival and departure that allows us to study how agent outcomes depend on the policy and types of allowed exchanges.

In each period, an agent arrives with a single indivisible item that she wishes to exchange. Our model has a homogeneous and independent stochastic demand structure, in which every agent A is willing to exchange his item for any other agent B ’s item with probability p . Figure 1 shows a ‘compatibility graph’ representation of a snapshot of a barter exchange marketplace, which captures the interest of agents in the items of other agents. Exchanges are conducted via (directed) cycles or chains, and a policy determines which potential exchanges to conduct. The compatibility graph evolves in time as new agents arrive, and existing agents depart after completing exchanges.

Agents who arrive to a barter marketplace want to quickly complete desirable exchanges. Motivated by this, the performance metric we consider is the average waiting time of agents in steady state. (In various contexts, it may be appropriate to consider some other cost function that is non-linear in the waiting time. However, for this foundational work we choose the simple linear cost function.) Thus, the optimal policy is one that minimizes the average waiting time.

⁸This question is also relevant to decentralized platforms, who can nevertheless control aspects like visibility. See [Section 8](#) for further discussion.

⁹The APD moved from matching monthly to daily matching. The NKR now conducts daily matches after experimenting with longer batches. UNOS Kidney exchange program have moved from monthly matching to weekly matching.

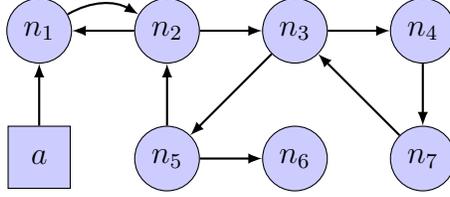


Figure 1: A compatibility graph representation of the potential trades in a market: Each circle node is an agent and the rectangle node, a is an “altruistic donor” who is willing to provide an item for free. A link from agent n_i to agent n_j means that n_j is willing to accept the good agent n_i has. This graph contains a 2-way cycle ($n_1 \rightarrow n_2 \rightarrow n_1$), 3-way cycles including ($n_2 \rightarrow n_3 \rightarrow n_5 \rightarrow n_2$) and multiple chains beginning from a including ($a \rightarrow n_1 \rightarrow n_2 \rightarrow n_3 \rightarrow n_4 \rightarrow n_7$).

We discuss three settings distinguished by the types of possible/allowed exchanges: (i) only 2-way cycles, (ii) 2 and 3-way cycles,¹⁰ and (iii) unbounded chains. In each setting, we seek to identify an approximately optimal policy. Further, we compare the best achievable expected waiting time across settings, with a view to quantifying the benefits of facilitating 3-way cycles and unbounded chains. Our key findings, informally stated, are as follows.

- In each of the three settings, the Greedy policy is approximately optimal in a large class of policies that includes all batching policies. Thus, in particular, batching does not provide any significant advantage.
- Under the Greedy policy, allowing three-way cycles leads to substantially smaller average waiting time than only allowing two-way cycles.
- Under the Greedy policy, the presence of altruistic donors who facilitate chain exchanges significantly reduces the average waiting time relative to the exchanges based on cycles of length two and three.

More precisely, we show that as $p \rightarrow 0$, the average waiting time under the Greedy policy scales as $\Theta(1/p^2)$ for the setting based on two-way cycles, as $\Theta(1/p^{3/2})$ for the setting based on two and three-way cycles, and as $\Theta(1/p)$ for the setting based on chains. Furthermore, for the first and third setting we show that a Greedy policy achieves the optimal scaling among essentially all possible policies¹¹. For the second setting with two and three-way cycles, we prove that greedy is scaling optimal in a broad class of policies called *monotone* policies (see the next section), which includes batching policies. We remark here that a small value of p is reasonable in many practical contexts, since agents are often interested in only a small fraction of the items offered by other agents.¹² Our results imply that, in each setting, for

¹⁰While our approach is potentially applicable to exchanges involving cycles longer than three, the technical details appear extremely challenging and we leave the question of sharp characterization of performance under four-way and longer exchanges for future research.

¹¹In fact, for the two-way cycles setting we show that even the constant factor is tight.

¹²Even kidney exchange clearinghouses observe a substantial fraction of highly sensitized patients that have probability 1 – 5% of matching.

all $p \in (0, 1)$, the waiting time under greedy is within a constant factor of the waiting time using any other batch size. Simulation experiments suggest that a batch of size of 1 is, in fact, truly optimal in each setting for any p .

Interestingly, our results are consistent with computational experiments in kidney exchange using clinical data (Ashlagi et al., 2013; Anderson et al., 2014), despite significant heterogeneity and other deviations from our model. These studies find that the benefit of batching relative to greedy, if any, is marginal, in line with our “greedy is approximately optimal among batching policies” finding. (This also matches practice.⁹) Ashlagi et al. (2013) and Dickerson et al. (2012) further demonstrate using computational experiments the significant benefits from using chains and 3-way cycles over 2-way cycles in dynamic settings. In practice, more than 85% of the transplants facilitated by multi-hospital kidney exchange clearinghouses are done through chains (Melcher et al., 2012). This is consistent with our theoretical findings regarding the benefits of 3-way cycles and chains in our model. We note that our results provide the first theoretical explanation of near optimality of greedy, whereas benefits from 3-cycles and chains have been previously found in related models (see below and Section 1.3).

Our model, while simplistic in its compatibility structure (which is described by a single parameter p), has several advantages. It notably avoids a “market size” parameter altogether (faster arrival of agents simply leads to an inverse rescaling of time), and further avoids a key drawback of previous models involving stochastic compatibilities (Ashlagi et al., 2012, 2013) that artificially require p to scale in a particular way with “market size”.¹³ Further, studying steady state behavior allows us to quantify performance exclusively in terms of waiting times. The alternative approach of studying a finite time horizon, as in Ashlagi et al. (2013), involves end-of-period effects that make it necessary to simultaneously consider both the waiting times and the number of matches, hindering performance comparisons.

These advantages of our model come at the cost of substantial new technical challenges. The most technical part of the paper involves obtaining the bounds of the form $\Theta(1/p^2)$, $\Theta(1/p^{3/2})$, and $\Theta(1/p)$. A key challenge we overcome is that the compatibility graph between currently waiting agents, conditional on running greedy so far, is not a directed Erdos-Renyi graph and has a complex distribution. It is sparser in terms of compatibilities in a very specific way: there are no possible exchanges, since the greedy policy would already have executed them. We develop methods to analytically control the graph with this distribution and the associated dynamical system. Another contribution is the technique we develop to prove lower bounds on average waiting times: this technique involves proof by contradiction, and is used in the case of three cycles and chains.

1.3 Related work

A number of previous theoretical studies in the kidney exchange literature have investigated the benefits from different types of exchanges: two-way, multi-way and chains. One stream

¹³These works have a market size parameter n , corresponding to the number of agents in a static or finite horizon setting, and require that $p = \Theta(1/n)$. Such a requirement is problematic, since p is typically exogenous.

of research studied large static marketplaces and counted the number of matches that can be achieved: [Roth et al. \(2007\)](#); [Ashlagi and Roth \(2011\)](#) found in a static large pool model with blood types that 3-way exchanges will increase the number of matches but there is no need for exchanges of size 4 or larger.

Recent works have studied a dynamic setting. [Ünver \(2010\)](#) has studied how to maximize the expected utility in a model in which there is no tissue-type incompatibility, i.e., compatibility is determined exclusively by blood groups, and is not stochastic. He shows that in this marketplace the optimal policy does not require cycles larger than 4. [Ashlagi et al. \(2013\)](#) studied a dynamic stochastic model in which nodes are either hard or easy to match (following empirical observations), abstracting away from blood types.¹⁴ Their model has a finite horizon, and different types of exchanges and policies are compared by counting number of matches. A main finding is that in such pools, 3-way cycles and long chains will result in more matches than pairwise exchanges. In each case, a batching policy maximizes number of matches. Now, it is not surprising that longer batches lead to more matches, however there is also the associated cost of making agents wait, and this cost is not captured in their setup. Another key weakness of their setup is that the probability to match a hard-to-match node is required to scale as the inverse of the pool size¹⁵. Such a requirement is needed to ensure that it is possible to match a constant fraction of the pool, with the constant factor being a function of the policy. Compatibility, however, is driven by factors independent of the pool size or arrival rate.

There is a large literature on trading in markets without monetary transfers in which agents are endowed with a single good, often referred to by “housing markets” ([Shapley and Scarf \(1974\)](#)). In a housing market one considers a group of agents, each of whom owns a house and has preferences over the set of houses. Shapley and Scarf, studied the core of the market, and described the well-known Top-Trading Cycles algorithm (attributed to David Gale) which finds an element in the core. This literature has grown to be very mature, analyzing core properties (see e.g. [Roth and Postlewaite \(1977\)](#)), incentives and design (see e.g. [Roth \(1982\)](#), [Abdulkadiroğlu and Sönmez \(1998\)](#) and [Pycia and Ünver \(2011\)](#)). Some applied studies are [Wang and Krishna \(2006\)](#) on timeshare exchanges (which allow people to trade a “week” of vacation they own) and [Dur and Ünver \(2012\)](#) on tuition exchange (which allows dependents of faculty members to attend other colleges for no tuition). This literature focuses on static markets. Little is known in dynamic matching environments, where agents arrive and trades occur over time.

Periodic matching has been studied in other (non-barter) markets as well. [Mendelson \(1982\)](#) analyzes a clearinghouse that periodically searches for outcomes in a dynamic market with sellers and buyers who arrive according to a stochastic process and trade indivisible goods. He studies the behavior of prices and quantities resulting from periodic trading. [Budish et al. \(2013\)](#) finds that some very small batching increases efficiency over continuous time trading in financial markets as firms compete over price rather than over speed. These works study (in)efficiencies resulting from prices, while our work focuses on the waiting times

¹⁴This model is a dynamic version of the model in [Ashlagi et al. \(2012\)](#).

¹⁵Here the pool size is the arrival rate times the time horizon.

in a homogeneous environment.

We now say a few words to situate our work from a technical perspective. Our model and analysis bring together the rich literature on (static) random graph models, e.g., see [Bollobás \(2001\)](#); [Janson et al. \(2000b\)](#) with the rich literature on queuing systems [Kleinrock \(1975\)](#); [Asmussen \(2003\)](#). In our model, the queue of waiting agents has a graph structure (i.e., the compatibility graph). Our stochastic model of compatibilities mirrors the canonical Erdos-Renyi model of a directed (static) random graph (but the dynamics make it much more complex). Comparing with common models of queuing systems, our system is peculiar in that the queuing system does not contain “servers” per se. Instead, the queue, in some sense, serves itself by executing exchanges that the compatibility graph allows. ([Gurvich and Ward \(2012\)](#) study optimal control in a related setting where jobs -analogous to our agents- can be served by matching with other compatible jobs. A key difference is that they consider a fixed number of job types, whereas in our setting compatibility is stochastically drawn for each pair of agents which leads to an unbounded number of possible agent types.) Nodes form cycles or chains with other nodes. Each time an exchange is executed, the corresponding agents/nodes leave the system. As a result, it turns out that for any reasonable policy the system is stable, irrespective of the rate of arrival of agents. If we speed up the arrival rate of agents, the entire system speeds up by the same factor, and waiting time reduces by the same factor. Thus, without loss of generality we consider an arrival rate of 1, with one agent arriving in each time slot.

In a concurrent work, [Akbarpour et al. \(2014\)](#) study a dynamic stochastic model with only two-way exchanges. Their model includes abandonments¹⁶, and the cost of waiting is suppressed (their results focus on discount factors close to 1) leading them to reach a conclusion that is the opposite of ours, namely that the policy should wait to thicken the market as opposed to being greedy. A limitation of their work is that they do not consider three-way cycles and chains (which together account for most kidney exchanges), instead focusing on the special and technically simplest case of two-way exchanges. Finally, their model also artificially requires the probability of compatibility to scale in a particular way with agent arrival rate; one consequence is their choice of scaling would allow agents to be matched almost instantaneously if three-cycles or chains are permitted.¹⁷

1.4 Notational conventions

We conclude with a summary of the mathematical notation used in the paper. Throughout, \mathbb{R} (\mathbb{R}_+) denotes the set of reals (nonnegative reals). We write that $f(p) = O(g(p))$ where $p \in (0, 1]$, if there exists $C < \infty$ such that $|f(p)| \leq Cg(p)$ for all $p \in (0, 1]$. We adapt the $\Theta(\cdot)$ and $\Omega(\cdot)$ notations analogously. We write that $f(p) = o(g(p))$ where $p \in (0, 1]$, if for any $C > 0$, there exists $p_0 > 0$ such that we have $|f(p)| \leq Cg(p)$ for all $p \leq p_0$. We adapt the $\omega(\cdot)$ notation analogously.

¹⁶They also give a truthful mechanism to elicit abandonment time.

¹⁷Undirected edges in their model are equivalent to a two-cycle of directed edges, and this occurs under the natural refinement where each directed edge is present independently with the same probability.

We let $\text{Bernoulli}(p)$, $\text{Geometric}(p)$, and $\text{Bin}(n, p)$, denote a Bernoulli variable with mean p , a geometric variable with mean $1/p$, and a Binomial random variable which is the sum of n independent identically distributed (iid) Bernoulli(p) random variables (r.v.s). We write $X \stackrel{d}{=} \mathcal{D}$ when the random variable X is distributed according to the distribution \mathcal{D} . We let $\text{ER}(n, p)$ be a *directed* Erdős Rényi random graph with n nodes where every two nodes form a directed edge with probability p , independently for all pairs. We let $\text{ER}(n, M)$ be the closely related *directed* Erdős Rényi random graph with n nodes and M directed edges, where the set of edges is selected uniformly at random among all subsets of exactly M directed edges. The two models are almost indistinguishable and, as is common in the literature on random graphs ([Janson et al., 2000a](#)), depending on the context we will use one model or the other. We let $\text{ER}(n_L, n_R, p)$ denote a bipartite *directed* Erdős Rényi random graph with two sides. This graph contains n_L nodes on the left, n_R nodes on the right, and a directed edge between every pair of nodes containing one node from each side is formed independently with probability p . Given a Markov chain $\{X_t\}$ defined on a state space \mathcal{X} and given a function $f: \mathcal{X} \rightarrow \mathbb{R}$, for $x \in \mathcal{X}$, we use the shorthand

$$\mathbb{E}_x[f(X_t)] \triangleq \mathbb{E}[f(X_t) \mid X_0 = x].$$

1.5 Organization

The rest of the paper is organized as follows: We describe our model formally in [Section 2](#) and state the main results of the paper in [Section 3](#). In [Section 4](#), we describe simulation results supporting our theoretical findings.

In [Section 5](#) we prove our main results for cycles of length two only, [Section 6](#) proves our results for two and three-cycles (technically the most challenging), and [Section 7](#) proves our results for chains. We conclude in [Section 8](#).

2 Model

We first state our model for settings with only cyclical exchanges and no chains. Later we augment it to accommodate altruistic/bridge donors and chains.

Consider the following model of a barter exchange where each agent arrives with an item that she wants to exchange for another item. In our simple binary model, each agent is (equally) interested in the items possessed by some of the other agents, and not interested in the items possessed by the rest.

Compatibility graph representation. The state of the system at any time can be represented by a directed graph where each agent is represented by a node, and a directed edge (i, j) exists if agent j wants the item of agent i . Let $\mathcal{G}(t) = (\mathcal{V}(t), \mathcal{E}(t))$ denote the directed graph of compatibilities observed before time t .

Dynamics. Initially the system may start in any state with a finite number of waiting

agents. We consider discrete times $t = 0, 1, 2, \dots$. At each time, one new agent arrives¹⁸. The new node representing this agent v has an incoming edge from each waiting agent who wants the item of v , and an outgoing edge to each waiting agent whose item v wants.

Stochastic compatibility model. The item of the new agent v is of interest to each of the waiting agents independently with probability p , and independently, the agent v is interested in the item of each waiting agent independently with probability p . Mathematically, there is a directed edge (in each direction) with probability p between the arriving node v and each other node that currently exists in the system, independently for all nodes and directions.

Allocation and policies. An *allocation* in a compatibility graph is a set of disjoint exchanges, namely a set of disjoint cycles and chains. We say that a node that is part of an allocation is *matched*. When an allocation consisting of cycles is executed, the compatibility graph is updated by eliminating the matched nodes and all their incident edges. Immediately after the arrival of a new node, the platform can choose to perform one or more exchanges, based on its chosen *policy*. Here, a policy is a mapping from the history of the system so far to an allocation. An exchange can happen via a cycle, where a *k-way cycle* is a directed cycle in the graph involving k nodes. It can also happen via a chain, which we define below.

Three types of settings (or technologies) are considered, differing by the exchanges permitted in an allocation. In the first two settings, allocations can output only cycles of length at most k , for $k = 2, 3$. These are termed *Cycle Removal*. In the third setting, called *Chain Removal*, allocations consist of only a single chain originating from a bridge node. We augment our model as follows for the chain removal setting.

Altruistic/bridge donors and chains. At the first time period, there is one altruistic donor present in the system, possibly along with other regular agents, and no further altruistic donors arrive to the system later. An altruistic donor is willing to give an item without getting anything in return. We represent an altruistic/bridge donor by a special *bridge* node.¹⁹ Bridge nodes can have only outgoing edges. For a new arrival v , there is an edge from a bridge node to v with probability p , independent of everything else. A *chain* is a directed path that begins with a bridge node. Once a chain is executed by the platform, the last node in a chain becomes a bridge node who can continue the chain in a later period. (All incoming edges to the last node in the chain are eliminated.) Notice that only one bridge donor remains in the system in the system at all times.

One natural policy that will play a key role in our results is the *greedy* policy. The greedy policy attempts to match the maximum number of nodes upon each arrival.

Definition 2.1. *The greedy policy for each of the settings is defined as follows:*

- *Cycle Removal: At the beginning of each time period the compatibility graph does not contain cycles with length at most k . Upon arrival of a new node, if a cycle with length at most k can be formed with the newly arrived node, it is removed, with a uniformly*

¹⁸One can instead consider a stochastic model of arrivals, e.g., Poisson arrivals in continuous time. In our setting, such stochasticity would leave the behavior of the model essentially unchanged, and indeed, each of our main results extend easily to the case of Poisson arrivals at rate 1.

¹⁹An example for a bridge node is a non-directed donor in kidney exchange programs.

random cycle being chosen if multiple cycles are formed. Clearly, at the beginning of the next time period the compatibility graph again does not contain any cycles with length at most k . The procedure is described on figure [Figure 3](#).

- *Chain Removal:* There is one bridge node in the system at the beginning of every time period. This bridge node does not have any incoming or outgoing edges. Upon the arrival of a new node at the beginning of a new time interval, the greedy policy identifies an allocation that includes the longest chain originating from the bridge node (breaking ties uniformly at random) and removes these nodes from the system and the last node in the chain becomes a bridge node. Note that such a chain has a positive length if and only if the bridge node has a directed edge from it to the new node. Observe that the new bridge node has in-degree and out-degree zero, so the process can repeat itself. This procedure is described on figure [Figure 2](#).

Under each of the settings, the system described above operated under the greedy policy is a Markov chain with a countably infinite number of states, each state corresponding to a compatibility graph, with a bridge node for the second setting, and no bridge nodes for the first setting. Further, this Markov chain is irreducible since an empty graph is reachable from any other state. This raises the question of whether this Markov chain is positive recurrent. If the answer is positive one can further study various performance measures.²⁰ The performance measure we focus on in this paper is the average (steady state) waiting time, which we define to be the average steady state time interval between the arrival of a node and the time when this node is removed²¹. We also consider policies other than the greedy policy, in general the class of policies under which the system is stationary/periodic and ergodic in the $t \rightarrow \infty$ limit. This includes²² the following class of policies that generalize Markov policies:

Definition 2.2. *We call a policy a periodic Markov policy if it employs τ homogenous first order Markov policies in round robin for some $\tau \in \mathbb{N}$.*

In other words, a periodic Markov policy implements a heterogeneous first order Markov chain, where the transition matrices repeat cyclically every τ rounds. Now suppose the resulting Markov chain is irreducible and periodic with period τ' . Without loss of generality, assume that τ is a multiple of τ' (if not, redefine τ as per $\tau \leftarrow \tau\tau'$). Now, clearly the subsequence of states starting with the state at time ℓ and then including states at time

²⁰The Markov chain turns out to be aperiodic for chain removal and also cycle removal, except for cycle removal with $k = 2$ where it is periodic with period 2. In any case, average (steady state) waiting time, cf. (1), is a natural metric for any periodicity.

²¹One may instead consider a cost function that is not linear in waiting time, depending on the intended application. For this first work on dynamic barter exchange, we focus on the simple metric of average waiting time. We remark that our Theorems 3.4 (on chains) and 3.3 (on three-cycles) are scaling results that also hold for any cost function that is bounded above and below by linear functions of waiting time. Similarly, Theorem 3.1 (on two-cycles) leads to a $\Theta(1/p^2)$ scaling result for any cost function of this type.

²²More precisely, positive recurrent periodic Markov policies (that stabilize the system) lead to a periodic and ergodic system. In any case we are not interested in policies that do not stabilize the system.

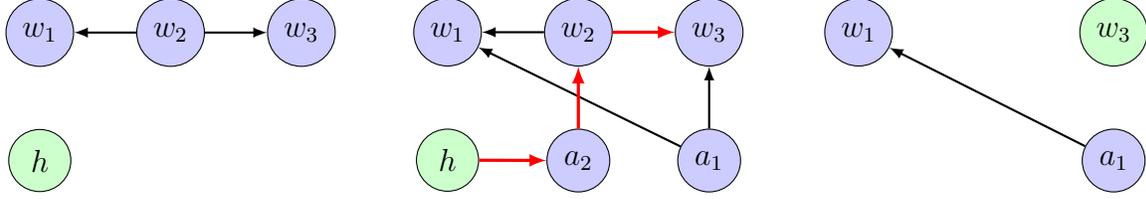


Figure 2: An illustration of chain matching under greedy. Initially, h is the head of the chain (the *bridge donor*), and nodes w_1 , w_2 , and w_3 are waiting to be matched, shown on the left. First, node a_1 arrives, and his good is acceptable by both w_1 and w_3 but no one has a good acceptable by a_1 . As h 's good is not acceptable by a_1 , it is not possible to move the chain. Then node a_2 arrives. His good is acceptable by w_2 and he is able to accept the good from h . The longest possible chain is shown in red in the center above. The chain is formed, h , a_2 , and w_2 are removed, and w_3 becomes the new head of the chain (bridge donor). Edges incident to the matched nodes are removed, as well as edges going in to w_3 . Note that in this case, the longest chain was not unique; w_1 could have been selected instead of w_3 .

intervals of τ , i.e., times $t = \ell, \ell + \tau, \ell + 2\tau, \dots$ forms an irreducible aperiodic first order Markov chain. If this ℓ -th ‘outer’ Markov chain is positive recurrent, we conclude that it converges to its unique steady state, leading to a periodic steady state for the original system with period τ . Define

$W_\ell \equiv$ Expected number of nodes in the system in the steady state
of the ℓ -th outer Markov chain.

Thus, W_ℓ is the expected number of nodes in the system at times that are $\ell \bmod \tau$ in steady state. Then we define the average waiting time for a periodic Markov policy as

$$W = (1/\tau) \sum_{\ell=0}^{\tau-1} W_\ell. \quad (1)$$

Note that this is the average number of nodes in the original system over a long horizon in steady state. Recalling Little’s law, this is hence identical to the average waiting time for agents who arrive to the system in steady state.

Remark 1. *We state our results formally for this broad class of periodic Markov policies, though our bounds extend also to other general policies that lead to a stationary/periodic and ergodic system in the $t \rightarrow \infty$ limit.*

3 Main results

We consider three different settings: a) two-way cycles only, b) two-way cycles and three-way cycles, and c) unbounded chains initiated by altruistic donors. In each setting we look for a policy that minimizes expected waiting time in steady state.

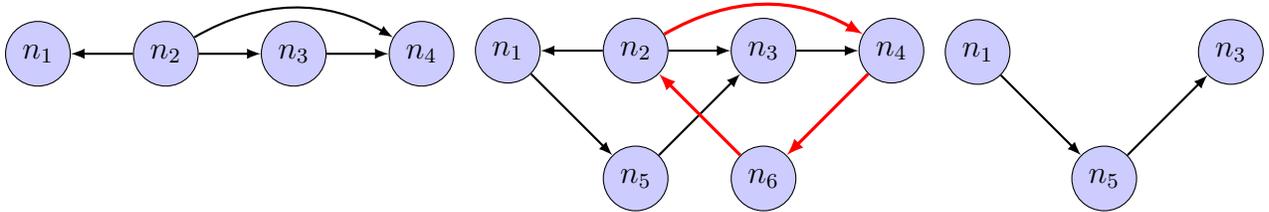


Figure 3: An illustration of cycle matching under the greedy policy, with a maximum cycle length of 3. Initially, nodes n_1 , n_2 , n_3 , and n_4 are all waiting, as shown on the left. Node n_5 arrives, but no directed cycles can be formed. Then n_6 arrives, forming the three cycle $n_6 \rightarrow n_2 \rightarrow n_4 \rightarrow n_6$. On the right, the three cycle is removed, along with the edges incident to any node in the three cycle. Note that when n_6 arrives, a six cycle is also formed, but under our assumptions, the maximum length cycle that can be removed is a three cycle.

Two-way cycles only. Our first result considers only 2-way cycles:

Theorem 3.1. *Under the Cycle Removal setting with $k = 2$, the greedy policy (cf. Definition 2.1) achieves an average waiting time of $\ln 2/p^2 + o(1/p^2)$. This is optimal, in the sense that for every periodic Markov policy, cf. Definition 2.2, the average waiting time is at least $\ln 2/(-\ln(1 - p^2)) = \ln 2/p^2 + o(1/p^2)$.*

The key fact leading to this theorem is that the prior probability of having a two-cycle between a given pair of nodes is p^2 , so an agent needs $\Theta(1/p^2)$ options in order to find another agent with whom a mutual swap is desirable. This result is technically the simplest to establish, but of equal interest in its implications. We prove Theorem 3.1 in Section 5.

Two-way cycles and three-way cycles. Our second result considers the case of cycle removals with $k = 3$. Our lower bound in this case applies to a specific class of policies which we now define.

Let \mathcal{G} denote the *global compatibility graph* that includes all nodes that ever arrive to the system, and directed edges representing compatibilities between them.

Definition 3.2. *A deterministic policy (under either Chain Removal or Cycle Removal) is said to be monotone if it satisfies the following property: Consider any pair of nodes (i, j) and an arbitrary global compatibility graph \mathcal{G} such that the edge (i, j) is present. Let $\bar{\mathcal{G}}$ be the graph obtained from \mathcal{G} when edge (i, j) is removed. Let T_i and T_j be the times of removal of nodes i and j respectively when the compatibility graph is \mathcal{G} and let $T_{ij} = \min(T_i, T_j)$. Then the policy must act in an identical fashion on $\bar{\mathcal{G}}$ and \mathcal{G} for all $t < T_{ij}$, i.e., the same cycles/chains are removed at the same times in each case, up to time T_{ij} . This property must hold for every pair of nodes (i, j) and every possible \mathcal{G} containing the edge (i, j) .*

A randomized policy is said to be monotone if it randomizes between deterministic monotone policies.

Remark 2. Consider the greedy policy for cycle removal defined above. It is easy to see that we can suitably couple the execution of greedy on different global compatibility graphs such that the resulting policy is monotone. The same applies to a batching policy which matches periodically (after arrival of x nodes), by finding a maximum packing of node disjoint cycles and removing them²³.

Note that the class of monotone policies includes a variety of policies in addition to simple batching policies. For instance, a policy that assigns weights to nodes and finds an allocation with maximum weight (instead of simply maximizing the number of nodes matched) is also monotone.

Theorem 3.3. Under the Cycle Removal setting with $k = 3$, the average waiting time under the greedy policy (cf. Definition 2.1) is $O(1/p^{3/2})$. Furthermore, there exists a constant $C < \infty$ such that, for any monotone policy that is periodic Markov (see Definitions 3.2 and 2.2), the average waiting time is at least $1/(Cp^{3/2})$.

Theorem 3.3 says that we can achieve a much smaller waiting time with $k = 3$, i.e., two and three-cycle removal, than the removal of two-cycles only (for small p). Further, for $k = 3$ greedy is again near optimal in the sense that no monotone policy can beat greedy by more than a constant factor. Theorem 3.3 is proved in Section 6. The proof overcomes a multitude of technical challenges arising from the complex distribution of the compatibility graph at a given time, and introduces several new ideas.

We remark that we could not think of any good candidate policy in our homogeneous model of compatibility that violates monotonicity but should do well on average waiting time. As such, we conjecture (but were unable to prove) that our lower bound on average waiting time applies to arbitrary and not just monotone policies.

The following fact may provide some intuition for the $\Theta(1/p^{3/2})$ scaling of average waiting time²⁴: In a static directed Erdős-Rényi graph with (small) edge probability p , one needs the number of nodes n to grow as $\Omega(1/p^{3/2})$ in order to, with high probability, cover a fixed fraction (e.g., 50%) of the nodes with node disjoint two and three cycles²⁵. Our rigorous analysis leading to Theorem 3.3 shows that this coarse calculation in fact leads to the correct scaling for average number of nodes in the dynamic system under the greedy policy, and that no monotone policy can do better.

Our result leaves open the case of larger cycles, i.e. $k > 3$, under the greedy, arbitrary monotone and arbitrary general policies. Based on intuition similar to the above, we conjecture that under the Cycle Removal setting with general k , the greedy policy achieves the average waiting time of $\Theta(p^{-\frac{k}{k-1}})$, and furthermore for every policy the average waiting time is lower bounded by $\Omega(p^{-\frac{k}{k-1}})$.

²³Note that such a policy is periodic Markov with a period equal to the batch size.

²⁴Recall that the average number of nodes is the same as the average waiting time, using Little's law.

²⁵The expected total number of three cycles is $n^3 p^3$ and the expected number of node disjoint three cycles is of the same order for $n^3 p^3 \lesssim n$. We need $n^3 p^3 \sim n$ in order to cover a given fraction of nodes with node disjoint three cycles, leading to $n \gtrsim 1/p^{3/2}$. For $n \sim 1/p^{3/2}$, the number of two-cycles is $n^2 p^2 \sim 1/p = o(n)$, i.e., very few nodes are part of two-cycles.

Unbounded chains initiated by altruistic donors. Our final result concerns the performance under the Chain Removal setting.

Theorem 3.4. *Under the Chain Removal setting, the greedy policy (cf. Definition 2.1) achieves an average waiting time of $O(1/p)$. Further, there exists a constant $C < \infty$ such that even if we allow removal of cycles of arbitrary length in addition to chains, for any periodic Markov policy, cf. Definition 2.2, the average waiting time is at least $1/(Cp)$.*

Thus, unbounded chains initiated by altruistic donors allow for a further large reduction in waiting time relative to the case of two-way and three-way cycles, for small p . In fact, as stated in the theorem, removal of cycles of arbitrary length (and chains), with any policy, cannot lead to better scaling of waiting time than that achieved with chains alone. Finally, greedy is near optimal among all periodic Markov policies for chain removal.²⁶

Theorem 3.4 involves a challenging technical proof presented in Section 7.

The intuition for the $\Theta(1/p)$ scaling of waiting time is somewhat involved: Since an agent finds the item of another agent acceptable with probability p , it is not hard to argue that no policy can sustain an expected waiting time that is $o(1/p)$; see our proof of the lower bound in Theorem 3.4 for a formalization of this intuition. On the other hand, under a greedy policy, the chain advances each time a new arrival can accept the item of the bridge donor, which occurs typically at $\Theta(1/p)$ intervals. One might hope that if there are many agents waiting, then typically, the next time there is an opportunity to advance the chain, we will be able to identify a long chain that will eliminate more agents than the number of agents that arrived since the last advancement. Our proof shows that this is indeed the case.

4 Computational experiments

We conducted simulation experiments which measure the average waiting times for nodes under Chain Removal and Cycle Removal with $k = 2$ and $k = 3$. For each of these matching technologies/settings, we simulated the performance of the batching policy with the batch size of x pairs, and compute the results for various values of x . For each scenario, we simulated a time horizon with 3500 arriving nodes, and measured the average number of nodes in the system after the the arrival of the 1000th node. (The first 1000 arrivals serve the role of a “warm-up” period.) 50 trials were conducted for each scenario simulated.

Figure 4(a) illustrates that when $p = 0.1$, the greedy policy, which corresponds to the batching policy with the batch size $x = 1$ performs the best among all batch sizes x . In addition, observe the significant difference between average waiting times corresponding to the Chain Removal setting on the one hand and the Cycle Removal setting with $k = 2$ on the other hand²⁷. Figures 4(b),(c) and (d) provide similar results for the cases $p = 0.08, 0.06$

²⁶One may ask what happens in the setting where chains, two-cycles and three-cycles are all allowed. We argue in Remark 3 that, for small p , this setting should be very similar to the setting with chains only.

²⁷We see that the difference between waiting times under chain removal and cycle removal with $k = 3$ is less pronounced. One reason for this could be that there are long intervals between consecutive times when a chain can be advanced, leading to a poor constant factor for chain removal. These intervals can be shortened by using non-maximal chains, and this may significantly improve the constant factor.

and 0.04.

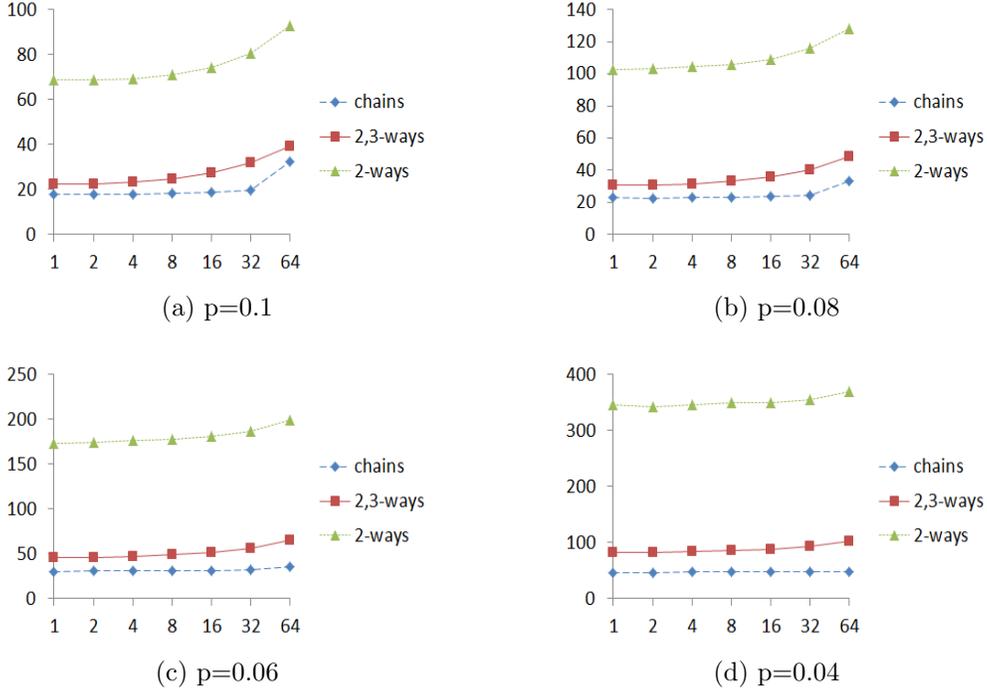


Figure 4: Average waiting time under the Chain Removal, Cycle Removal with $k = 2$, and Cycle Removal with $k = 3$, with batching sizes $x = 1, 2, 4, 8, 16, 32, 64$

5 Two-way Cycle Removal

In this section we consider Cycle Removal with $k = 2$. The greedy policy corresponding to the Cycle Removal setting when $k = 2$ is simple to characterize, since, as we show below, the underlying process behaves as a simple random walk. We observe that the random walk has a negative drift when $|\mathcal{V}(t)| \geq \log(2)/p^2$, and obtain a tight characterization of waiting time under greedy using a simple coupling argument. The key idea for the lower bound is that regardless of the implemented policy, the rate at which 2-cycles which will be eventually removed are formed must equal to the half of the rate at which new nodes arrive, which is equal to unity. Further, the probability that we do not form any cycles which will be eventually removed is lower bounded by the probability that we do not form any cycles at all. This probability depends only on the number of nodes in the system, the desired quantity.

Proof of Theorem 3.1. We first compute the expected steady state waiting time under the

greedy policy. Observe that for all $t \geq 0$,

$$|\mathcal{V}(t+1)| = \begin{cases} |\mathcal{V}(t)| + 1 & \text{with probability } (1-p^2)^{|\mathcal{V}(t)|}, \\ |\mathcal{V}(t)| - 1 & \text{with probability } 1 - (1-p^2)^{|\mathcal{V}(t)|}. \end{cases}$$

Let $\varepsilon > 0$ be arbitrary. If $|\mathcal{V}(t)| > (1+\varepsilon)\ln(2)/p^2$, then there exists a sufficiently small $p = p(\varepsilon)$ such that for all $p > p(\varepsilon)$

$$\mathbb{P}(|\mathcal{V}(t+1)| = |\mathcal{V}(t)| + 1) = (1-p^2)^{|\mathcal{V}(t)|} \leq \frac{1}{2^{1+\varepsilon}}.$$

Let $q = 1/2^{1+\varepsilon} < 1/2$, and let X_t be a sequence of i.i.d. random variables with distribution

$$X_t = \begin{cases} 1 & \text{with probability } q, \\ -1 & \text{with probability } 1-q. \end{cases}$$

Let $S_0 = 0$ and for $t \geq 1$, $S_{t+1} = (S_t + X_t)^+$, so S_t is a Birth-Death process. Letting $r = q/(1-q) < 1$, in steady state $\mathbb{P}(S_\infty = i) = r^i(1-r)$ for $i = 0, 1, \dots$, so

$$\mathbb{E}[S_\infty] = r/(1-r) = q/(1-2q) = \frac{1}{2^{1+\varepsilon} - 2}.$$

We can couple the random walk $|\mathcal{V}(t)|$ with S_t such that $|\mathcal{V}(t)| < (1+\varepsilon)\ln(2)/p^2 + S_t$ for all t . This yields

$$\mathbb{E}[|\mathcal{V}(\infty)|] \leq (1+\varepsilon)\frac{\ln(2)}{p^2} + \mathbb{E}[S_\infty] \leq (1+\varepsilon)\frac{\ln(2)}{p^2} + \frac{1}{2^{1+\varepsilon} - 2}.$$

Thus for every $\varepsilon > 0$, we have

$$\lim_{p \rightarrow 0} \frac{\mathbb{E}[|\mathcal{V}(\infty)|] - \ln(2)/p^2}{1/p^2} \leq \varepsilon \ln(2).$$

As ε was arbitrary, the result follows.

Now we establish the lower bound on $|\mathcal{V}(\infty)|$. Let v be a newly arriving node at time t , and \mathcal{W} be the nodes currently in system that are waiting to be matched. Let I be the indicator that at the arrival time of v (just before cycles are potentially deleted), no 2-cycles between v and any node in \mathcal{W} exist. Let \tilde{I} be the indicator that at the arrival time of v , no two cycles *that will be eventually removed* that are between v and any node in \mathcal{W} exist (in particular, \tilde{I} depends on the future). Thus $\tilde{I} \geq I$ a.s. Let \tilde{V}_t be the number of vertices in the system before time t such that the cycle which eventually removes them has not yet arrived. We let \tilde{V}_∞ be the distribution of \tilde{V}_t when the system begins in steady state. By stationarity

$$0 = \mathbb{E}[\tilde{V}_{t+1} - \tilde{V}_t] = \mathbb{E}_{\tilde{V}_\infty}[2\tilde{I} - 1],$$

giving $E[\tilde{I}] = 1/2$. Intuitively, in steady state, the expected change in the number of vertices not yet “matched” must be zero. Thus we obtain

$$\frac{1}{2} = \mathbb{E}[\tilde{I}] \geq \mathbb{E}[I] = \mathbb{E}[\mathbb{E}[I \mid |\mathcal{V}(\infty)|]] = \mathbb{E}[(1 - p^2)^{|\mathcal{V}(\infty)|}] \geq (1 - p^2)^{\mathbb{E}[|\mathcal{V}(\infty)|]},$$

by Jensen’s inequality. Taking logarithms on both sides and rearranging terms, we get

$$\mathbb{E}[|\mathcal{V}(\infty)|] \geq \frac{\log(1/2)}{\log(1 - p^2)} = \frac{\log(2)}{-\log(1 - p^2)}.$$

□

6 Three-way Cycle Removal

In this section we prove [Theorem 3.3](#). The proof is far more involved than for the case $k = 2$, especially the upper bound, and relies on delicate combinatorial analysis of 3-cycles random graph formed by nodes present in the system in steady state and those arriving over a certain time interval. We consider a time interval of the order $\Theta(1/p^{3/2})$ and assume that the system starts with at least order $\Theta(1/p^{3/2})$ nodes in the underlying graph. We establish a negative drift in the system and then, as in the case of Chain Removal mechanism, rely on the Lyapunov function technique in order to establish the required upper bound.

For the lower bound, we introduce a novel approach that allows us to prove a matching lower bound (up to constants) for monotone policies by contradiction. The rough idea is as follows: if the steady state expected waiting time is small (in this case smaller than $1/(Cp^{3/2})$ for appropriate C), then a typical new arrival sees a small number of nodes currently in the system, and so typically does not form a two or three-cycle with existing nodes or even the next few arrivals. Thus, the typical arrival typically has a long waiting time, which contradicts our initial assumption of a small expected waiting time.

Preliminaries

We first state a number of propositions and lemmas that will enable our proof of [Theorem 3.3](#). Proofs of these preliminaries are deferred to [Appendix A](#).

We begin by stating (without proof) the following version of the classical Chernoff bound (see, e.g. [Alon and Spencer 2008](#)).

Proposition 6.1 (Chernoff bound). *Let $X_i \in \{0, 1\}$ be independent with $\mathbb{P}(X_i = 1) = p_i$ for $1 \leq i \leq n$. Let $\mu = \sum_{i=1}^n p_i$.*

(i) *For any $\delta \in [0, 1]$ we have*

$$\mathbb{P}(|X - \mu| \geq \mu\delta) \leq 2 \exp\{-\delta^2\mu/3\} \tag{2}$$

(ii) *For any $R > 6\mu$ we have*

$$\mathbb{P}(X \geq R) \leq 2^{-R} \tag{3}$$

Next, we state a result which is based on the Lyapunov function technique. Given an irreducible aperiodic Markov chain $\{X_k\}$ on a countable statespace \mathcal{X} , suppose there exists a nonnegative function $V: \mathcal{X} \rightarrow \mathbb{R}_+$ which admits the following decomposition

$$V(X_k) = V(X_{k-1}) + A_k - D_k, \quad (4)$$

where the $A_k \geq 0$ is an i.i.d. sequence such that A_k is independent from state X_k , while the $D_k \geq 0$ may depend on X_k and A_k , but not on k directly. Specifically, D_k is a function of X_k and A_k . A_k and D_k are interpreted as the number of arrivals and departures in the time period of length k , respectively. Assume in addition that $\mathcal{B}(\alpha) \triangleq \{x \in \mathcal{X} \mid V(x) \leq \alpha\} \subset \mathcal{X}$ is finite for every α . Note that as $V(x) \geq 0$, we have that $D_k \leq V(x) + A_k$ a.s.

Proposition 6.2. *Suppose $\mathbb{E}[A_k^2]$ is finite and C_1 satisfies $\mathbb{E}[A_k^2] \leq C_1 \mathbb{E}[A_k]^2 < \infty$. Suppose there exists $\alpha, \lambda, C_2 > 0$ such that for every $x \notin \mathcal{B}(\alpha)$*

$$\mathbb{E} \left[A_k - \tilde{D}_k \mid X_k = x \right] \leq -\lambda \mathbb{E}[A_k], \quad (5)$$

where \tilde{D}_k is defined to be $\min\{D_k, C_2 A_k\}$. Then X_k is positive recurrent with the unique stationary distribution X_∞ and

$$\mathbb{E}[V(X_\infty)] \leq \max \left\{ \alpha, \frac{\max\{1, C_2 - 1\}^2 C_1 \mathbb{E}[A_k]}{\lambda} \right\} \left(2 + \frac{2}{\lambda} \right).$$

The reason for introducing a truncated downward jump process \tilde{D}_k as opposed to using just D_k is that in general the statement of the proposition is not true. Namely, there exists a process such the assumptions of the proposition above hold true when D_k replaces \tilde{D}_k in (5) and $\mathbb{E}[V(X_\infty)] = \infty$, as shown by [Example 1](#) in [Appendix A](#).

Next, we state a straightforward combinatorial bound: In a directed graph, a set \mathcal{M} of node disjoint three-cycles is said to be *maximal* if no three-cycle can be added to \mathcal{M} so that the set remains node disjoint.

Proposition 6.3. *Given an arbitrary directed graph G , let N be the number of three-cycles in a largest in cardinality set of node disjoint three-cycles in G . Then, any maximal set of node disjoint three-cycles consists of at least $N/3$ three-cycles.*

Finally, let \mathcal{G}_t denote the *global compatibility graph* that includes all nodes that ever arrive to the system up to time t , and directed edges representing compatibilities between them. Denote by \mathcal{W}_t the set of nodes out of $0, 1, \dots, t$ still present in the system at time t . The following is a key property of monotone policies:

Lemma 6.4. *Under any monotone policy, for every two nodes i, j arriving before time t (namely $i, j \leq t$) and every subset of nodes $\mathcal{W} \subset \{0, 1, \dots, t\}$ containing nodes i and j*

$$\mathbb{P}((i, j) \in \mathcal{G}_t \mid \mathcal{W}_t = \mathcal{W}) \leq p.$$

In words, pairs of nodes still present in the system at time t are no more likely to be connected at time t than at the time they arrive.

The following corollary follows immediately by linearity of expectations.

Corollary 6.5. *Let $W_t = |\mathcal{W}_t|$ and let E_t be the number of edges between nodes in \mathcal{W}_t . Then, under a monotone policy, $\mathbb{E}[E_t|W_t] \leq W_t(W_t - 1)p$.*

Proposition 6.2, Proposition 6.3 and Lemma 6.4 are proved in Appendix A.

Proof of Theorem 3.3

Proof of Theorem 3.3: the performance of the greedy policy. Suppose at time zero we observe $W \geq C^3/p^{3/2}$ nodes in the system with an arbitrary set of edges between them. Here C is a sufficiently large constant to be fixed later. Call this set of nodes \mathcal{W} . Consider the next $T = 1/(Cp^{3/2})$ arrivals, and call this set of nodes \mathcal{A} . Wlgl, label the times of these arrivals as $1, 2, \dots, T$, and use the label t for the node that arrives at time t . Let $\mathcal{A}_t \subseteq \{1, 2, \dots, t-1\}$ be the subset of nodes in \mathcal{A} that have arrived but have not been removed before time t . Similarly define \mathcal{W}_t to be the set of nodes from \mathcal{W} which are still in the system immediately before time t . Note that, in particular, $\mathcal{W}_1 = \mathcal{W}$.

Let N be the number of three cycles removed during the time period $[0, T]$, that include two nodes from \mathcal{W} . Let $\kappa = 1/C^2$ and consider the event

$$\mathcal{E}_1 \equiv \{|\mathcal{A}_{T+1}| - N \geq 2\kappa/p^{3/2}\} \quad (6)$$

Introduce the event

$$\mathcal{E}_2 \equiv \{\text{There exists a set of disjoint 2 and 3 cycles in } \mathcal{A} \\ \text{with cardinality at least } 3/(C^3p^{3/2}) \}. \quad (7)$$

First suppose that the event \mathcal{E}_1 does not occur. Then

$$|\mathcal{A}_{T+1}| \leq \frac{2}{C^2p^{3/2}} + N \leq \frac{1}{16Cp^{3/2}} + N, \quad (8)$$

for C sufficiently large. Also event \mathcal{E}_2 implies that (again for C sufficiently large) at most $9/(C^3p^{3/2}) \leq 1/(16Cp^{3/2})$ nodes in \mathcal{A} leave due to internal three-cycles or two cycles. Since $T = 1/(Cp^{3/2})$, then applying Eq. (8), at least $7/(8Cp^{3/2}) - N$ other nodes in \mathcal{A} also leave before $T + 1$. These other nodes belong to cycles of one of the following types:

- (i) A three cycle containing another node from \mathcal{A} and a node from \mathcal{W} .
- (ii) A two cycle with a node from \mathcal{W} .
- (iii) A three cycle containing two nodes from \mathcal{W} . There are exactly N nodes of this type.

Exactly N nodes in \mathcal{A} are removed due to cycles of type (iii) above, so we infer that at least $7/(8Cp^{3/2}) - 2N$ nodes in \mathcal{A} are removed due to cycles of type (i) or (ii) above, meaning that at least $(1/2)(7/(8Cp^{3/2}) - 2N)$ nodes in \mathcal{W} are removed as part of such cycles. Clearly, $2N$ nodes in \mathcal{W} are removed as part of cycles of type (iii).

It follows that

$$|\mathcal{W}_{T+1}| \leq |\mathcal{W}| - 2N - \frac{1}{2} \left(\frac{7}{8Cp^{3/2}} - 2N \right) \leq |\mathcal{W}| - \frac{7}{16Cp^{3/2}} - N. \quad (9)$$

Combining Eqs. (8) and (9), we deduce that

$$|\mathcal{W}_{T+1}| + |\mathcal{A}_{T+1}| \leq |\mathcal{W}| - \frac{3}{8Cp^{3/2}}. \quad (10)$$

We also have that the number of nodes in the system increases by at most $T = 1/(Cp^{3/2})$. We will show that $\mathbb{P}(\mathcal{E}_1 \cup \mathcal{E}_2) \leq \varepsilon = 1/9$. Before establishing this claim we show how this claim implies the result. We have

$$\mathbb{E}[|\mathcal{W}_{T+1}| + |\mathcal{A}_{T+1}| - |\mathcal{W}|] \leq \varepsilon T - (1 - \varepsilon) \frac{3}{8Cp^{3/2}} = -\frac{2}{9Cp^{3/2}}, \quad (11)$$

i.e., the number of nodes decrease by at least $2/(9Cp^{3/2})$ in expectation. We now apply [Proposition 6.2](#) to the embedded Markov chain observed at times which are multiples of T . Namely, let $T_i = i \cdot T$, and take $X_i = \mathcal{G}(T_i) = (\mathcal{V}(T_i), \mathcal{E}(T_i))$ and define $V(X_i) = |\mathcal{V}(T_i)|$. If we let \mathcal{D}_i be the set of nodes that are deleted in some cycle during the time interval $[T_i, T_{i+1})$, we obtain a decomposition

$$V(X_{i+1}) = |\mathcal{V}(T_{i+1})| = |\mathcal{V}(T_i)| + T - |\mathcal{D}_i| = V(X_i) + T - |\mathcal{D}_i|. \quad (12)$$

Since $T > 0$ is deterministic it is trivially independent from $\mathcal{G}(T_i)$. Thus the assumptions on decomposing V from (4) are satisfied. The assumption that $\{\mathcal{G} \mid V(\mathcal{G}) < n\}$ is finite for every n is satisfied as there are only finitely many graphs with n nodes. We take $\alpha = C^3/p^{3/2}$ making $\mathcal{B} = \{\mathcal{G} \mid |\mathcal{V}(\mathcal{G})| \leq C^3/p^{3/2}\}$. We can take $C_1 = 1$ as T is deterministic. We can take $C_2 = 3$, as trivially $|\mathcal{D}_i| \leq 3T$ since each newly arriving node can be in at most one three cycle (thus making $D_k = D_k$ in [Proposition 6.2](#)). Finally, we can take $\lambda = 2/9$, as by (11),

$$\mathbb{E}[T - |\mathcal{D}_i|] \leq -\frac{2}{9Cp^{3/2}} = -\frac{2}{9}\mathbb{E}[T].$$

Thus by applying [Proposition 6.2](#), we obtain that

$$\begin{aligned} \mathbb{E}[|\mathcal{V}(T_\infty)|] &\leq \max \left\{ \alpha, \frac{\max\{1, C_2 - 1\}^2 C_1 \mathbb{E}[A_k]}{\lambda} \right\} \left(2 + \frac{2}{\lambda} \right) \\ &= \max \left\{ \frac{C^3}{p^{3/2}}, \frac{4}{2/9} \frac{1}{Cp^{3/2}} \right\} \left(2 + \frac{2}{2/9} \right) \\ &= \frac{11C^3}{p^{3/2}}, \end{aligned}$$

for C sufficiently large. Finally, since the embedded chain is observed over deterministic time intervals, the bound above applies to the steady-state bound. We conclude

$$\mathbb{E}[|\mathcal{V}(\infty)|] \leq \frac{11C^3}{p^{3/2}}.$$

It remains to bound $\mathbb{P}(\mathcal{E}_1)$ and $\mathbb{P}(\mathcal{E}_2)$ to complete the proof. We do this below. We claim that $\mathbb{P}(\mathcal{E}_2) \leq \varepsilon/4$. We first show that there is likely to be a maximal set of node disjoint three-cycles in \mathcal{A} of size less than $2/(3C^3p^{3/2})$. This will imply, using [Proposition 6.3](#), that the maximum number of node disjoint three-cycles in \mathcal{A} is at most $2/(C^3p^{3/2})$. Reveal the graph on \mathcal{A} and simultaneously construct a maximal set of node disjoint three-cycles as follows: Reveal node 1. Then reveal node 2. Then reveal node 3 and whether it forms a three-cycle with the existing nodes. If it does remove this three-cycle. Continuing, at any stage t if a three-cycle is formed, choose uniformly at random such a three-cycle and remove it.

Since this process corresponds to a monotone policy (cf. [Definition 3.2](#)), then using [Corollary 6.5](#), the residual graph immediately before step t contains no more than $2\binom{t-1}{2}p$ edges in expectation, as the number of nodes is no more than $t-1$. It follows that the conditional probability of three-cycle formation at step t is no more than $\mathbb{E}[\text{Number of three-cycles formed}] = 2\binom{t-1}{2}p^3$. It follows that we can set up a coupling such that the total number of three-cycles removed (this is a maximal set of edge disjoint three-cycles resulting from our particular greedy policy) is no more than $Z = \sum_{t=1}^T X_t$ where $X_t \sim \text{Bernoulli}(2\binom{t-1}{2}p^3)$ are independent. Now $\mathbb{E}[Z] = 2\binom{T}{3}p^3 \leq 1/(3C^3p^{3/2})$. Using [Proposition 6.1](#) (i), we obtain that $\mathbb{P}(Z \geq 2/(3C^3p^{3/2})) < \varepsilon/8$, for large enough p , establishing the desired bound on the number of node disjoint three cycles. We have shown that the probability of having more than $2/(C^3p^{3/2})$ node disjoint three cycles in \mathcal{A} is less than $\varepsilon/8$.

Let Z' be the number of two cycles internal to \mathcal{A} . Then $Z' \sim \text{Bin}(\binom{T}{2}, p^2)$. Hence, $\mathbb{E}[Z'] \leq 1/(C^2p)$ and $\mathbb{P}(Z' \geq 1/(C^3p^{3/2})) \leq \varepsilon/8$ for sufficiently small p using [Proposition 6.1](#) (ii). It follows that the probability of having more than $1/(C^3p^{3/2})$ node disjoint two cycles in \mathcal{A} is less than $\varepsilon/8$. Now $\mathbb{P}(\mathcal{E}_2) \leq \varepsilon/4$ follows by union bound.

We now show $\mathbb{P}(\mathcal{E}_1) \leq 3\varepsilon/4$. To prove this, we find it convenient to define two additional events. Denote by $\mathcal{N}(\mathcal{S}_1, \mathcal{S}_2)$ the (directed) neighborhood of the nodes in \mathcal{S}_1 in the set of nodes \mathcal{S}_2 , i.e., $\mathcal{N}(\mathcal{S}_1, \mathcal{S}_2) = \{j \in \mathcal{S}_2 : \exists i \in \mathcal{S}_1 \text{ s.t. } (i, j) \in \mathcal{E}\}$. Abusing notation, we use $\mathcal{N}(i, \mathcal{S})$ to denote the neighborhood of node i in \mathcal{S} . Further, we find it convenient to define $\mathcal{B}_t = \mathcal{N}(t, \mathcal{A}_t)$. Define

$$\mathcal{E}_{3,t} \equiv \{|\mathcal{A}_t| \geq \kappa/p^{3/2}, \text{ and } |\mathcal{B}_t| < \kappa/(2p^{1/2})\}, \quad (13)$$

and $\mathcal{E}_3 = \cup_{0 \leq t \leq T} \mathcal{E}_{3,t}$. Define

$$\mathcal{E}_{4,t} \equiv \{|\mathcal{A}_t| \geq \kappa/p^{3/2}, \text{ and } |\mathcal{N}(\mathcal{B}_t, \mathcal{W}_t)| < C^3\kappa/(8p)\}, \quad (14)$$

and let $\mathcal{E}_4 = \cup_{0 \leq t \leq T} \mathcal{E}_{4,t}$. We make use of

$$\begin{aligned} \mathcal{E}_1 &\subseteq (\mathcal{E}_4^c \cap \mathcal{E}_1) \cup \mathcal{E}_4 \subseteq (\mathcal{E}_4^c \cap \mathcal{E}_1) \cup \mathcal{E}_3 \cup (\mathcal{E}_4 \cap \mathcal{E}_3^c) \\ \Rightarrow \mathbb{P}(\mathcal{E}_1) &\leq \mathbb{P}(\mathcal{E}_4^c \cap \mathcal{E}_1) + \mathbb{P}(\mathcal{E}_3) + \mathbb{P}(\mathcal{E}_4 \cap \mathcal{E}_3^c) \end{aligned}$$

Reveal the edges between t and \mathcal{A}_t when node t arrives. The existence of each edge is independent of the other edges and the current revealed graph. Thus we can bound

the probability of the event $\mathcal{E}_{3,t}$ using [Proposition 6.1](#) (i) by $2 \exp(-1/(12C^2p^{1/2}))$ for large enough C . It follows that for sufficiently small p , we have

$$\mathbb{P}(\mathcal{E}_3) \leq 2T \exp(-1/(12C^2p^{1/2})) \leq \varepsilon/4. \quad (15)$$

We now bound $\mathbb{P}(\mathcal{E}_4^c \cap \mathcal{E}_1)$. Let N_t be the number of three cycles removed before time t of type (iii) (recall that type (iii) three cycles include two nodes from \mathcal{W}). Define $Z_t \equiv |\mathcal{A}_t| - N_t$. Define

$$\mathcal{E}_{5,t} \equiv \{\text{Node } t \text{ is part of a three-cycle of type (i)}\}.$$

Note that

- If $\mathcal{E}_{5,t}$ then $|\mathcal{A}_{t+1}| = |\mathcal{A}_t| - 1$, $N_{t+1} = N_t$ if such a three cycle is removed and $|\mathcal{A}_{t+1}| = |\mathcal{A}_t|$, $N_{t+1} = N_t + 1$ if a three cycle of type (iii) is removed instead. In either case, we have $Z_{t+1} = Z_t - 1$.
- With probability one we have $|\mathcal{A}_{t+1}| \leq |\mathcal{A}_t| + 1$ and $N_{t+1} \geq N_t$. It follows that $Z_{t+1} \leq Z_t + 1$.

Now suppose $Z_t \geq \kappa/p^{3/2}$ and $\mathcal{E}_{4,t}^c$. Clearly $Z_t \geq \kappa/p^{3/2} \Rightarrow |\mathcal{A}_t| \geq \kappa/p^{3/2}$ and hence $\mathcal{E}_{4,t}^c \Rightarrow |\mathcal{N}(\mathcal{B}_t, \mathcal{W}_t)| \geq C^3 \kappa / (8p) = C / (8p)$. Revealing the edges between from $\mathcal{N}(\mathcal{B}_t, \mathcal{W}_t)$ to t , we see that

$$\mathbb{P}(\mathcal{E}_{5,t} | Z_t \geq \kappa/p^{3/2}, \mathcal{E}_{4,t}^c) \geq 1 - (1-p)^{C/(4p)} \geq 3/4, \quad (16)$$

for large enough C and small enough p , independent of everything so far. So, informally, if $\mathcal{E}_{4,t}^c$ then Z_t is a bounded above by a random walk with a downward drift whenever $Z_t \geq \kappa/p^{3/2}$. We now formalize this.

Define the random walk $(\tilde{Z}_t)_{t \geq 1}$ as follows: Let $\tilde{Z}_1 = 0$. Whenever $\tilde{Z}_t = 0$, we have $\tilde{Z}_{t+1} = 1$, else

$$\tilde{Z}_{t+1} = \begin{cases} \tilde{Z}_t + 1 & \text{w.p. } 1/4 \\ \tilde{Z}_t - 1 & \text{w.p. } 3/4 \end{cases} \quad (17)$$

So $(\tilde{Z}_t)_{t=1}^{T+1}$ is a downward biased random walk reflected upwards at 0.

Proposition 6.6. *There exists $C < \infty$ such that for any $T \in \mathbb{N}$ and $\nu > 0$, we have $\mathbb{P}(\tilde{Z}_{T+1} \geq \nu) \leq CT \exp(-\nu/C)$.*

The proof is omitted, as this is a standard result for random walks with a negative drift. Using [Proposition 6.6](#), we have that for sufficiently small p ,

$$\mathbb{P}(\tilde{Z}_{T+1} \geq \kappa/(2p^{3/2})) \leq \varepsilon/4.$$

Let τ be the first time at which event $\mathcal{E}_{4,t}$ occurs for $t \leq T$, and let $\tau = T + 1$ if \mathcal{E}_4 does not occur. We now show that the following claim holds:

Claim 6.7. We can couple Z_t and \tilde{Z}_t such that for all $t < \tau$, whenever $Z_t \geq \kappa/p^{3/2}$ we have $\tilde{Z}_{t+1} - \tilde{Z}_t \geq Z_{t+1} - Z_t$.

Proof of Claim. If $\mathcal{E}_{5,t}$ occurs, then (see above) we know that $Z_{t+1} = Z_t - 1$ and $\tilde{Z}_{t+1} - \tilde{Z}_t \geq -1$ holds by definition of \tilde{Z} . Hence, it is sufficient to ensure that $\tilde{Z}_{t+1} = \tilde{Z}_t + 1$ whenever $\mathcal{E}_{5,t}^c$ occurs. But this is easy to satisfy since Eq. (16) implies that

$$\mathbb{P}(\mathcal{E}_{5,t}^c | Z_t \geq \kappa/p^{3/2}, \mathcal{E}_{4,t}^c) \leq 1/4,$$

whereas $\mathbb{P}(\tilde{Z}_{t+1} = \tilde{Z}_t + 1) = 1/4$. This completes the proof of the claim. \square

The following claim is an immediate consequence:

Claim 6.8. We have $Z_t \leq \tilde{Z}_t + \lceil \kappa/p^{3/2} \rceil$ for all $t \leq \tau$.

Proof of Claim. The claim follows from Claim 6.7 and a simple induction argument. \square

It follows that

$$\begin{aligned} \mathbb{P}(Z_{T+1} \geq 2\kappa/p^{3/2}, \tau = T + 1) &\leq \mathbb{P}(\tilde{Z}_{T+1} \geq \kappa/p^{3/2}, \tau = T + 1) \\ &\leq \mathbb{P}(\tilde{Z}_{T+1} \geq \kappa/p^{3/2}) \\ &\leq \varepsilon/4. \end{aligned}$$

Thus we obtain

$$\mathbb{P}(\mathcal{E}_4^c \cap \mathcal{E}_1) \leq \varepsilon/4. \quad (18)$$

Finally, we bound $\mathbb{P}(\mathcal{E}_4 \cap \mathcal{E}_3^c)$. For any $\mathcal{S} \subseteq \mathcal{A}$, let $\mathcal{W}_{\sim \mathcal{S}} \subseteq \mathcal{W}$ be the set of waiting nodes that would have been removed before $T+1$ if (hypothetically) the nodes in \mathcal{S} had no incident edges in either direction, but we left all other compatibilities unchanged. Define the event \mathcal{E}_6 as follows: for all $\mathcal{S} \subseteq \mathcal{A}$ such that $|\mathcal{S}| = \kappa/(2p^{1/2})$, the bound

$$|\mathcal{N}(\mathcal{S}, \mathcal{W} \setminus \mathcal{W}_{\sim \mathcal{S}})| \geq C/(8p) \quad (19)$$

holds.

Claim 6.9. The event \mathcal{E}_6 occurs whp.

Before proving the claim, we show that it implies $\mathbb{P}(\mathcal{E}_4 \cap \mathcal{E}_3^c) \leq \varepsilon/4$. Suppose that \mathcal{E}_6 and \mathcal{E}_3^c occur. Consider any t such that $|\mathcal{A}_t| \geq \kappa/p^{3/2}$. Since \mathcal{E}_3^c , we have that $|\mathcal{B}_t| \geq \kappa/(2p^{1/2})$. Take any $\mathcal{S} \subseteq \mathcal{B}_t$ such that $|\mathcal{S}| = \kappa/(2p^{1/2})$. Notice that for our monotone greedy policy, cf. Remark 2, the set of waiting nodes that are removed before time t must be a subset of $\mathcal{W}_{\sim \mathcal{S}}$, i.e., we have $\mathcal{W}_t \supseteq \mathcal{W} \setminus \mathcal{W}_{\sim \mathcal{S}}$. Since \mathcal{E}_6 occurs, it follows that $|\mathcal{N}(\mathcal{S}, \mathcal{W}_t)| \geq C/(8p) \Rightarrow |\mathcal{N}(\mathcal{B}_t, \mathcal{W}_t)| \geq C/(8p)$. Thus we have \mathcal{E}_4^c . This argument just established that

$$\begin{aligned} \mathcal{E}_6 \cap \mathcal{E}_3^c &\subseteq \mathcal{E}_4^c \cap \mathcal{E}_3^c \\ \Rightarrow \mathcal{E}_6^c \cap \mathcal{E}_3^c &\supseteq \mathcal{E}_4 \cap \mathcal{E}_3^c. \end{aligned}$$

It follows that $\mathbb{P}(\mathcal{E}_4 \cap \mathcal{E}_3^c) \leq \mathbb{P}(\mathcal{E}_6^c \cap \mathcal{E}_3^c) \leq \mathbb{P}(\mathcal{E}_6^c) \leq \varepsilon/4$ using Claim 6.9, as required.

Proof of Claim 6.9. Consider any $\mathcal{S} \subseteq \mathcal{A}$ such that $|\mathcal{S}| = \kappa/(2p^{1/2})$. Clearly, since each node in \mathcal{A} can eliminate at most 2 nodes in \mathcal{W} , we have $|\mathcal{W}_{\sim\mathcal{S}}| \leq 2|\mathcal{A} \setminus \mathcal{S}| \leq 2|\mathcal{A}| = 2/(Cp^{3/2})$. It follows that $|\mathcal{W} \setminus \mathcal{W}_{\sim\mathcal{S}}| \geq C^3/p^{3/2} - 2/(Cp^{3/2}) \geq C^3/(2p^{3/2})$ for large enough C . Now notice that by definition $\mathcal{W}_{\sim\mathcal{S}}$ is a function only of the edges between nodes in $\mathcal{W} \cup (\mathcal{A} \setminus \mathcal{S})$, and is independent of the edges coming out of \mathcal{S} . Thus, for each node $i \in \mathcal{W} \setminus \mathcal{W}_{\sim\mathcal{S}}$ independently, we have that each node in \mathcal{S} has an edge to i independently w.p. p . We deduce $i \in \mathcal{N}(\mathcal{S}, \mathcal{W} \setminus \mathcal{W}_{\sim\mathcal{S}})$ w.p. $1 - (1-p)^{\kappa/(2p^{1/2})} \geq \kappa p^{1/2}/3$ for small enough p , i.i.d. for each $i \in \mathcal{W} \setminus \mathcal{W}_{\sim\mathcal{S}}$. It follows from Proposition 6.1 (i) that

$$|\mathcal{N}(\mathcal{S}, \mathcal{W} \setminus \mathcal{W}_{\sim\mathcal{S}})| < \frac{C^3}{2p^{3/2}} \cdot \frac{\kappa p^{1/2}}{3} \cdot \frac{3}{4} = \frac{\kappa C^3}{8p} = \frac{C}{8p}$$

occurs w.p. at most $2 \exp\{- (1/4)^2 \cdot C/(6p) \cdot (1/3)\} \leq \exp(-C/(300p))$ for small enough p . Now, the number of candidate subsets \mathcal{S} is $\binom{1/(Cp^{3/2})}{\kappa/p^{1/2}} \leq (1/(Cp^{3/2}))^{\kappa/p^{1/2}} \leq \exp(1/p^{\varepsilon+1/2})$ for small enough p . It follows from union bound that $|\mathcal{N}(\mathcal{S}, \mathcal{W} \setminus \mathcal{W}_{\sim\mathcal{S}})| < \frac{C}{8p}$ for one (or more) of these subsets \mathcal{S} with probability at most $\exp(-C/(300p)) \cdot \exp(1/p^{\varepsilon+1/2}) \leq \exp(-C/(400p)) \xrightarrow{p \rightarrow 0} 0$. Thus, whp, $|\mathcal{N}(\mathcal{S}, \mathcal{W} \setminus \mathcal{W}_{\sim\mathcal{S}})| < \frac{C}{8p}$ occurs for no candidate subset \mathcal{S} , i.e., event \mathcal{E}_6 occurs whp. \square

Proof of Theorem 3.3: lower bound for monotone policies. Denote by m the expected steady state number of nodes in the system, which by Little's Law equals the expected steady state waiting time. Suppose $m \leq 1/(Cp^{3/2})$, where C is any constant larger than 36. Fix a node i , and reveal the number of nodes W in the system when i arrives. Notice $W \leq 3m$ occurs with probability at least $1 - 1/3 = 2/3$ in steady state by Markov's inequality. Assume that $W \leq 3m$ holds. Let \mathcal{W} denote the nodes waiting in the system when i arrives, and let \mathcal{A} be the nodes that arrive in the next $3m$ time slots after node i arrives. Now, if node i leaves the system within $3m$ time slots of arriving, then i must form a two or three cycle with nodes in $\mathcal{A} \cup \mathcal{W}$. The probability of forming such a cycle is bounded above by

$$\mathbb{E}[\text{Number of two cycles between } i \text{ and } \mathcal{A} \cup \mathcal{W} | W] \tag{20}$$

$$+ \mathbb{E}[\text{Number of three cycles containing } i \text{ and two nodes from } \mathcal{A} \cup \mathcal{W} | W]. \tag{21}$$

Clearly,

$$\mathbb{E}[\text{Number of two cycles between } i \text{ and } \mathcal{A} \cup \mathcal{W} | W] \leq 6mp^2 \leq 1/C \tag{22}$$

for p sufficiently small. To bound the other term we notice that

$$\begin{aligned} & \mathbb{E}[\text{Number of three cycles containing } i \text{ and two nodes from } \mathcal{A} \cup \mathcal{W} | W] \\ &= p^2 \cdot \mathbb{E}[\text{Number of edges between nodes in } \mathcal{A} \cup \mathcal{W} | W] \end{aligned} \tag{23}$$

We use Corollary 6.5 to bound the expected number of edges between nodes in \mathcal{W} at the time when i arrives by $W(W-1)p$ and notice that other compatibilities (j_1, j_2) for $\{j_1, j_2\} \not\subseteq \mathcal{W}$ are present independently with probability p . Hence, we have

$$\mathbb{E}[\text{Number of edges between nodes in } \mathcal{A} \cup \mathcal{W} | W] \leq |\mathcal{W} \cup \mathcal{A}|(|\mathcal{W} \cup \mathcal{A}| - 1)p \leq 6m(6m - 1)p.$$

Using Eq. (23) we infer that

$$\begin{aligned} & \mathbb{E}[\text{Number of three cycles containing } i \text{ and two nodes from } \mathcal{A} \cup \mathcal{W} | W] \\ & \leq 6m(6m-1)p^3 \leq 36/C^2 \leq 1/C \end{aligned} \tag{24}$$

for $C > 36$. Using Eqs. (22) and (24) in (21), we deduce that the probability of node i being removed within $3m$ slots is no more than $2/C$.

Combining, the unconditional probability that node i stays in the system for more than $3m$ slots is at least $(2/3)(1-2/C) > 1/3$ for large enough C . This violates Markov inequality, implying that our assumption, $m \leq 1/(Cp^{3/2})$, was false. This establishes the stated lower bound. \square

The following conjecture results if we assume that n_t concentrates, and that the typical number of edges in a compatibility graph at time t with n_t nodes is close to what it would have been under an $\text{ER}(n_t, p)$ graph.

Conjecture 6.10. *For cycle removal with $k = 3$, the expected waiting time in steady state under a greedy policy scales as $\sqrt{\ln(3/2)}/p^{3/2} + o(1/p^{3/2})$, and no periodic Markov policy (including non-monotone policies) can achieve an expected waiting time that scales better than this.*

Here the constant $\sqrt{\ln(3/2)}$ results from requiring (under our assumptions) that a newly arrived node forms a triangle with probability $1/3$.

Our simulation results, cf. Figure 4, are consistent with this conjecture: the predicted expected waiting time for greedy from the leading term $\sqrt{\ln(3/2)}/p^{3/2}$ is $W = 80$ for $p = 0.04$, $W = 43$ for $p = 0.06$, $W = 28$ for $p = 0.08$ and $W = 20.1$ for $p = 0.1$. If proved, this conjecture would be refinement of Theorem 3.3. A proof would require a significantly more refined analysis for both the upper bound and the lower bound.

7 Chain Removal

In this section we prove Theorem 3.4. At any time there is one bridge donor in the system. Under a greedy policy, the chain advances when the newly arrived agent can accept the item of the bridge donor. The basic idea to establish that greedy achieves $O(1/p)$ waiting time is to show that when there are more than C/p waiting nodes just after we move the chain forward, then, on average, the next time the chain moves forward, it will remove more nodes than were added in the interim. This “negative drift” in number of nodes is crucial in establishing the bound (following which we again use Proposition 6.2 to infer a bound on the expected waiting time). The lower bound proof is based on the idea that the waiting time for a node must be at least the time for the node to get an in-degree of one. The lower bound is again proved by contradiction.

Preliminaries

We begin by stating a result on long chains in a static Erdős-Rényi random graph. The following result was first shown by [Ajtai et al. \(1981\)](#) and refined in a series of papers, see [Krivelevich et al. \(2012\)](#) for a historical account and the most tight result.

Proposition 7.1 ([Krivelevich et al. 2012](#)). *Fix any $\varepsilon > 0$ and any $\delta > 0$. There exist C and n_0 such that for all $c > C$ and all $n > n_0$ the following occurs: Consider an $\text{ER}(n, c/n)$ directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, and let D be the length of the longest directed cycle. We have*

$$\mathbb{P}(D > (1 - (2 + \delta)ce^{-c})n) > 1 - \varepsilon.$$

In words, (for large c) we have a cycle containing a large fraction of the nodes with high probability. From this, we can easily obtain a similar result about the longest path starting from a specific node.

Corollary 7.2. *Fix any $\varepsilon > 0$. There exist C and n_0 such that for all $c > C$ and all $n > n_0$ the following occurs: Consider a set \mathcal{V} of n vertices including a fixed vertex $v \in \mathcal{V}$, and draw an $\text{ER}(n, c/n)$ directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Let P_v denote the length of a longest path starting at v . Then*

$$\mathbb{P}(P_v < n(1 - \varepsilon)) \leq \varepsilon.$$

The proof is deferred to [Appendix B](#). We extend the result above to the case of bipartite random graphs.

Corollary 7.3. *Fix any $\kappa > 1$ and $\varepsilon > 0$. Then there exists $p_0 > 0$ and $C > 0$ such that the following holds: Consider any $c_L \in [1/\sqrt{\kappa}, \kappa]$, any $c_R > C$, and any $p < p_0$. Let \mathcal{L} be a set of c_L/p vertices and let \mathcal{R} be a set of c_R/p vertices. Fix a vertex $v \in \mathcal{L}$. Draw $\mathcal{G} = (\mathcal{L}, \mathcal{R}, \mathcal{E})$ as an $\text{ER}(c_L/p, c_R/p, p)$ bipartite random graph. We have*

$$\mathbb{P}\left(P_v < 2\frac{c_L}{p}(1 - \varepsilon)\right) \leq \varepsilon,$$

where again, P_v is the length of a longest path starting at v .

Again, the proof is in [Appendix B](#). The requirement $c_L \in [1/\sqrt{\kappa}, \kappa]$ here will correspond to p times the ‘typical’ interval between successive times when the chain advances under greedy. These intervals are distributed i.i.d. $\text{Geometric}(p)$, and hence typically lie in the range $[1/(p\sqrt{\kappa}), \kappa/p]$ for large κ , as stated in [Lemma 7.4](#) below. The $1/\sqrt{\kappa}$ term in the lower bound of this ‘typical’ range is a somewhat arbitrary choice we make that facilitates a proof of [Theorem 3.4](#) (a variety of other decreasing functions of κ would work as well).

Lemma 7.4. *There exist p_0 and κ_0 such that for all $p < p_0$ and all $\kappa > \kappa_0$, if $X \sim \text{Geometric}(p)$, then*

$$\mathbb{E}\left[X\mathbb{I}_{\{X < \frac{1}{p\sqrt{\kappa}} \text{ or } X > \frac{\kappa}{p}\}}\right] \leq \frac{2}{\kappa p}.$$

Again, the proof is in [Appendix B](#).

Proof of Theorem 3.4

We introduce the following notation. Let $\mathcal{G}(t) = (\mathcal{V}(t), \mathcal{E}(t), h(t))$ be the directed graph at time t describing the compatibility graph at time t . Here $h(t)$ is a special node not included in $\mathcal{V}(t)$ that is the head of the chain, which can only have out-going edges. We denote by $\mathcal{G}(\infty) = (\mathcal{V}(\infty), \mathcal{E}(\infty), h(\infty))$ the steady-state version of this graph (which exists as we show below).

According to the greedy policy, whenever $h(t)$ forms a directed edge to a newly arriving node, a largest possible chain starting from $h(t)$ is made. Thus before the new node arrives, $h(t+1)$ will always have an in degree and out degree of zero (as explained in Section 2), and we can only advance the chain when a newly arriving node has an in edge from $h(t)$. We refer to these periods between chain advancements as *intervals*. Let τ_i for $i = 1, 2, \dots$, denote the length of the i th interval, so that $\tau_i \sim \text{Geometric}(p)$. Let $T_0 = 0$ and $T_i = \sum_{j=1}^i \tau_j$ for $i = 1, 2, \dots$, be the time at the end of the i th interval. Additionally, let \mathcal{A}_i be the set of nodes that arrived during the i th interval $[T_{i-1}, T_i]$, so that $|\mathcal{A}_i| = \tau_i$, and let \mathcal{W}_i be the set of nodes that were “waiting” at the start of the i -th interval, namely at time T_{i-1} . Thus, right before the chain is advanced, every node in the graph is either in \mathcal{W}_i , \mathcal{A}_i or it is $h(t)$ itself.

The intuition for the upper bound on waiting time for the greedy policy in Theorem 3.4 is as follows. We will use the Lyapunov function argument (Proposition 6.2) to argue that for some C , if there are at least C/p nodes in the graph at the start of an interval $[T_i, T_i + \tau_{i+1}]$, then the number of vertices deleted in an interval is on average greater than the number of vertices that arrive in that interval, i.e. we have a negative drift on the number of vertices. We lower bound the number of nodes removed in the i th interval by the length of a longest path in the bipartite graph formed by putting nodes in \mathcal{A}_i (the newly arrived nodes) to the left part of the graph, and putting nodes \mathcal{W}_i (the nodes from the previous interval) to the right part of the graph, and maintaining only edges between these two parts (thus in particular preserving the bipartite structure). We bound the expected size of a longest path on this subgraph using Corollary 7.3. Observe, that the length of a longest path on our bipartite graph is at most $2|\mathcal{A}_i|$. This will enable us to truncate the downward jumps when applying Proposition 6.2.

Proof of Theorem 3.4: performance of the greedy policy. We apply Proposition 6.2, taking as our Markov chain $X_i = \mathcal{G}(T_i)$, and our Lyapunov function $V(\cdot)$ to be $V(\mathcal{G}(T_i)) \triangleq |\mathcal{V}(T_i)|$. For a constant $C > 0$ to be specified later, we let α from Proposition 6.2 be $\alpha = C/p$. Thus our finite set of exceptions is $\mathcal{B} = \{\mathcal{G} = (\mathcal{V}, \mathcal{E}, h) : |\mathcal{V}| \leq C/p\}$, the directed graphs with at most C/p nodes. Obviously our state space is countable and \mathcal{B} is finite. Let \mathcal{P}_i be the path of nodes that are removed from the graph in the i th interval. Thus

$$|\mathcal{V}(T_i)| = |\mathcal{V}(T_{i-1})| + |\mathcal{A}_i| - |\mathcal{P}_i|.$$

By taking $A_i = |\mathcal{A}_i| = \tau_i$ and $D_i = |\mathcal{P}_i|$, we have that $V(\cdot)$ satisfies the form of (4) and the independence assumptions on A_i and D_i . As $\tau_i \sim \text{Geometric}(p)$, $\mathbb{E}[|\mathcal{A}_i|^2] \leq 2/p^2 = 2\mathbb{E}[|\mathcal{A}_i|]^2$,

so we can take $C_1 = 2$. We set $C_2 = 2$, and so to apply [Proposition 6.2](#), we must find $\lambda > 0$ such that for every graph $\mathcal{G} \notin \mathcal{B}$,

$$\mathbb{E}_{\mathcal{G}} [|\mathcal{A}_i| - \min\{|\mathcal{P}_i|, 2|\mathcal{A}_i|\}] \leq -\lambda \mathbb{E}[|\mathcal{A}_i|], \quad (25)$$

where $\mathbb{E}_{\mathcal{G}}[\cdot]$ denote the expectation conditioned on the event $\mathcal{G}_{i-1} = \mathcal{G}$. We create an auxiliary bipartite graph $\mathcal{G}'_i = (\mathcal{A}_i, \mathcal{W}_i, \mathcal{E}'_i)$, where \mathcal{A}_i are the nodes on the left, \mathcal{W}_i are the nodes on the right, and \mathcal{E}'_i is the subset of $\mathcal{E}(T_i)$ consisting of edges (u, v) such that either $u \in \mathcal{A}_i, v \in \mathcal{W}_i$ or vice versa (thus ensuring \mathcal{G}'_i is bipartite). We let $v'_i \in \mathcal{A}_i$ be the node newly arrived at T_i that $h(T_i - 1)$ connected to. Finally, we let \mathcal{P}'_i be the longest path in \mathcal{G}'_i starting at v'_i . Trivially, $|\mathcal{P}'_i| \leq |\mathcal{P}_i|$. Observe that \mathcal{G}'_i is a $\text{ER}(|\mathcal{A}_i|, |\mathcal{W}_i|, p)$ bipartite random graph. Thus we apply [Corollary 7.3](#) to show $|\mathcal{P}'_i|$ is appropriately large with high probability. In particular, given arbitrary $\varepsilon > 0$ and $\kappa > \kappa_0 > 1$, where κ_0 is to be specified later, find C and p_0 according to [Corollary 7.3](#). Then $\mathcal{G}(T_{i-1}) \notin \mathcal{B}$ implies $|\mathcal{W}_i| = |\mathcal{V}(T_{i-1})| \geq C/p$. Then if $p < p_0$ and $a \in \left[\frac{1}{p\sqrt{\kappa}}, \frac{\kappa}{p}\right]$, then by [Corollary 7.3](#),

$$\mathbb{P}\left(|\mathcal{P}'_i| < 2|\mathcal{A}_i|(1 - \varepsilon) \mid |\mathcal{A}_i| = a\right) \leq \varepsilon. \quad (26)$$

We define the events E_i and F_i by

$$E_i = \left\{ \mathcal{A}_i \notin \left[\frac{1}{p\sqrt{\kappa}}, \frac{\kappa}{p} \right] \right\} \quad F_i = \{|\mathcal{P}'_i| < 2|\mathcal{A}_i|(1 - \varepsilon)\}.$$

We define $Z_i \triangleq 2|\mathcal{A}_i|(1 - \varepsilon)\mathbb{I}_{E_i^c \cap F_i^c}$. Thus $Z_i \leq |\mathcal{P}'_i| \leq |\mathcal{P}_i|$ from the definition of the event F_i , and $Z_i \leq 2|\mathcal{A}_i|$ by construction. We now use this to get an upper bound (25) as follows. First, we have:

$$\begin{aligned} \mathbb{E}_{\mathcal{G}} [|\mathcal{A}_i| - \min\{|\mathcal{P}_i|, 2|\mathcal{A}_i|\}] &\leq \mathbb{E}_{\mathcal{G}} [|\mathcal{A}_i| - Z_i] \\ &= \mathbb{E} [|\mathcal{A}_i|\mathbb{I}_{E_i}] + \mathbb{E} \left[|\mathcal{A}_i| - Z_i \mid E_i^c \right] \mathbb{P}(E_i^c), \end{aligned} \quad (27)$$

where in (27), we used that Z_i is zero on E_i . Now noting that for all $a \in \left[\frac{1}{p\sqrt{\kappa}}, \frac{\kappa}{p}\right]$, i.e. in the event E_i^c , we have

$$\mathbb{P}\left(Z_i = 0 \mid |\mathcal{A}_i| = a\right) = \mathbb{P}\left(F_i \mid |\mathcal{A}_i| = a\right) \leq \varepsilon,$$

by (26), and therefore

$$\mathbb{P}\left(Z_i = 2(1 - \varepsilon)|\mathcal{A}_i| \mid |\mathcal{A}_i| = a\right) = \mathbb{P}\left(F_i^c \mid |\mathcal{A}_i| = a\right) \geq 1 - \varepsilon.$$

as well. We now compute that

$$\begin{aligned}
& \mathbb{E} \left[|\mathcal{A}_i| - Z_i \mid E_i^c \right] \\
&= \sum_{a \in \left[\frac{1}{p\sqrt{\kappa}}, \frac{\kappa}{p} \right]} \mathbb{E} \left[|\mathcal{A}_i| - Z_i \mid |\mathcal{A}_i| = a \right] \mathbb{P} \left(|\mathcal{A}_i| = a \mid E_i^c \right) \\
&= \sum_{a \in \left[\frac{1}{p\sqrt{\kappa}}, \frac{\kappa}{p} \right]} \mathbb{E} \left[|\mathcal{A}_i| - Z_i \mid F_i \cap |\mathcal{A}_i| = a \right] \mathbb{P} \left(F_i \mid |\mathcal{A}_i| = a \right) \mathbb{P} \left(|\mathcal{A}_i| = a \mid E_i^c \right) \\
&\quad + \sum_{a \in \left[\frac{1}{p\sqrt{\kappa}}, \frac{\kappa}{p} \right]} \mathbb{E} \left[|\mathcal{A}_i| - Z_i \mid F_i^c \cap |\mathcal{A}_i| = a \right] \mathbb{P} \left(F_i^c \mid |\mathcal{A}_i| = a \right) \mathbb{P} \left(|\mathcal{A}_i| = a \mid E_i^c \right) \\
&\leq \sum_{a \in \left[\frac{1}{p\sqrt{\kappa}}, \frac{\kappa}{p} \right]} \mathbb{E} \left[|\mathcal{A}_i| \mid F_i \cap |\mathcal{A}_i| = a \right] \cdot \varepsilon \cdot \mathbb{P} \left(|\mathcal{A}_i| = a \mid E_i^c \right) \\
&\quad + \sum_{a \in \left[\frac{1}{p\sqrt{\kappa}}, \frac{\kappa}{p} \right]} \mathbb{E} \left[|\mathcal{A}_i| - 2(1 - \varepsilon)|\mathcal{A}_i| \mid F_i^c \cap |\mathcal{A}_i| = a \right] \cdot (1 - \varepsilon) \cdot \mathbb{P} \left(|\mathcal{A}_i| = a \mid E_i^c \right) \\
&\leq \varepsilon \mathbb{E} \left[|\mathcal{A}_i| \mid E_i^c \right] + (1 - \varepsilon) \mathbb{E} \left[(-1 + 2\varepsilon)|\mathcal{A}_i| \mid E_i^c \right] \\
&= (-1 + 4\varepsilon - 2\varepsilon^2) \mathbb{E} \left[|\mathcal{A}_i| \mid E_i^c \right] \\
&\leq (-1 + 4\varepsilon) \mathbb{E} \left[|\mathcal{A}_i| \mid E_i^c \right].
\end{aligned}$$

Now combining this with (27), we have

$$\begin{aligned}
\mathbb{E} [|\mathcal{A}_i| \mathbb{I}_{E_i}] + \mathbb{E} \left[|\mathcal{A}_i| - Z_i \mid E_i^c \right] \mathbb{P}(E_i^c) &\leq \mathbb{E} [|\mathcal{A}_i| \mathbb{I}_{E_i}] + (-1 + 4\varepsilon) \mathbb{E} \left[|\mathcal{A}_i| \mid E_i^c \right] \mathbb{P}(E_i^c) \\
&= \mathbb{E} [|\mathcal{A}_i| \mathbb{I}_{E_i}] + (-1 + 4\varepsilon) \mathbb{E} [|\mathcal{A}_i| \mathbb{I}_{E_i^c}] \mathbb{P}(E_i^c) \\
&\leq \frac{2}{\kappa p} + (-1 + 4\varepsilon) \left(\frac{1}{p} - \frac{2}{\kappa p} \right) \\
&\leq -\frac{1}{p} + \frac{4\varepsilon}{p} + \frac{4}{\kappa p}, \\
&= -\frac{1}{p} \left(1 - 4\varepsilon - \frac{4}{\kappa} \right),
\end{aligned} \tag{28}$$

where in (28) we used Lemma 7.4 twice. Now, we let $\delta \triangleq 4\varepsilon + 4/\kappa$, and observe that we can make δ arbitrarily and the inequality will still hold for sufficiently small p by our choice of ε and κ_0 . As we have

$$\mathbb{E}_{\mathcal{G}} [|\mathcal{A}_i| - \min\{|\mathcal{P}_i|, 2|\mathcal{A}_i|\}] \leq -\frac{1}{p}(1 - \delta) = -\mathbb{E}[|\mathcal{A}_i|](1 - \delta),$$

we can apply [Proposition 6.2](#) with $\lambda = (1 - \delta)$ to obtain that

$$\begin{aligned} \mathbb{E}[|\mathcal{V}(T_\infty)|] &\leq \max \left\{ \alpha, \frac{\max\{1, C_2 - 1\}^2 C_1 \mathbb{E}[A_k]}{\lambda} \right\} \left(2 + \frac{2}{\lambda} \right) \\ &= \max \left\{ \frac{C}{p}, \frac{2}{1 - \delta} \frac{1}{p} \right\} \left(2 + \frac{2}{1 - \delta} \right) \end{aligned}$$

Finally, recall that we are working with the “embedded Markov chain” as we are only observing the process at times T_i . We can relate the actual Markov chain to the embedded Markov chain as follows:

$$\mathbb{E}[|\mathcal{V}(\infty)|] = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=0}^t |\mathcal{V}(s)| \quad (29)$$

$$= \lim_{n \rightarrow \infty} \frac{1}{T_n} \sum_{s=0}^{T_n} |\mathcal{V}(s)| \quad (30)$$

$$\begin{aligned} &= \lim_{n \rightarrow \infty} \frac{n}{T_n} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \sum_{s=T_{i-1}}^{T_i} |\mathcal{V}(s)| \\ &= p \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \left(|\mathcal{V}(T_{i-1})|(T_i - T_{i-1}) + \frac{(T_i - T_{i-1})(T_i - T_{i-1} + 1)}{2} \right) \quad (31) \end{aligned}$$

$$\leq p \left(\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n |\mathcal{V}(T_{i-1})|(T_i - T_{i-1}) + \frac{1}{n} \sum_{i=1}^n (T_i - T_{i-1})^2 \right). \quad (32)$$

Here (29) follows from the positive recurrence of $\mathcal{G}(t)$. We have (30) as $a_n \rightarrow a$ implies that for every subsequence a_{n_i} , we have $a_{n_i} \rightarrow a$ as well, and using that as $T_n \rightarrow \infty$ a.s. We obtain the left term in (31) by observing that T_n is the sum of n independent $\text{Geometric}(p)$ random variables and then applying the SLLN. For the right term of (31), we simply use that $|\mathcal{V}(s+1)| = |\mathcal{V}(s)| + 1$ for $s \in [T_{i-1}, T_i - 1]$, and then the identity $\sum_{i=1}^n i = n(n+1)/2$.

We now considering each sum from (32) independently. For the first sum, observing that $|\mathcal{V}(T_{i-1})|(T_i - T_{i-1})$ is a function of our positive recurrent Markov chain $\mathcal{G}(T_i)$, we have that there exists a random variable X^* such that $|\mathcal{V}(T_{i-1})|(T_i - T_{i-1}) \Rightarrow X^*$ and the average value of $|\mathcal{V}(T_{i-1})|(T_i - T_{i-1})$ converges to $\mathbb{E}[X^*]$ a.s. The convergence in distribution $|\mathcal{V}(T_{i-1})|(T_i - T_{i-1}) \Rightarrow X^*$ implies the existence of $|\tilde{\mathcal{V}}(\tilde{T}_{i-1})|(\tilde{T}_i - \tilde{T}_{i-1})$ that converges to X^*

a.s. Putting these together, we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n |\mathcal{V}(T_{i-1})|(T_i - T_{i-1}) = \mathbb{E}[X^*] \quad (33)$$

$$\begin{aligned} &= \mathbb{E} \left[\lim_{i \rightarrow \infty} |\tilde{\mathcal{V}}(\tilde{T}_{i-1})|(\tilde{T}_i - \tilde{T}_{i-1}) \right] \\ &\leq \liminf_{i \rightarrow \infty} \mathbb{E} [|\mathcal{V}(T_{i-1})|(T_i - T_{i-1})] \end{aligned} \quad (34)$$

$$= \liminf_{i \rightarrow \infty} \mathbb{E} [|\mathcal{V}(T_{i-1})|] \mathbb{E}[T_i - T_{i-1}] \quad (35)$$

$$= \frac{1}{p} \mathbb{E}[|\mathcal{V}(T_\infty)|]. \quad (36)$$

Here we have (33) by the ergodic theorem for Markov chains, (34) by Fatou's lemma, (35) by the independence of $T_i - T_{i-1}$ from $\mathcal{V}(T_{i-1})$, and (36) by Theorem 2 from Tweedie (1983) (alternatively, (36) can be shown a little extra work using a simpler result from Holewijn and Hordijk (1975)).

For the second sum, as $T_i - T_{i-1} = \tau_i$ are i.i.d. Geometric(p), by the SLLN,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (T_i - T_{i-1})^2 = \mathbb{E}[\tau_1^2] = \frac{2-p}{p^2} \leq \frac{2}{p^2}$$

Thus

$$\mathbb{E}[|\mathcal{V}(\infty)|] \leq p \left(\frac{1}{p} \mathbb{E}[|\mathcal{V}(T_\infty)|] + \frac{2}{p^2} \right) = \mathbb{E}[|\mathcal{V}(T_\infty)|] + \frac{2}{p},$$

showing the result, as we have for the embedded process that $\mathbb{E}[|\mathcal{V}(T_\infty)|] = \Omega(1/p)$. \square

Finally, we mention that in moving from the “embedded Markov chain” back to the original Markov chain we make use of the fact that τ_i is light tailed in the sense that $\mathbb{E}[\tau_i^2] = O((\mathbb{E}[\tau_i])^2)$, to obtain a bound of $O(1/p)$ of the steady state expected number of nodes in the system.

The proof for the lower bound in Theorem 3.4 is based on the following key idea: the waiting time for a node must be at least the time for the node to get an in-degree of one. Using this, if the steady state average waiting time is $w = o(1/p)$, then by Little's Law when a typical vertex v arrives there are only $o(1/p)$ vertices in system, so v is likely not to have any in edges connecting with any of these existing nodes. After w steps, the number of newly arrived nodes is $w = o(1/p)$, so v is likely not to connect to any of these nodes either. Then the idea is to show that v will be in the system for greater than w steps with high enough probability (i.e. with probability at least $1/3$), contradicting that the expected waiting time v is w .

Proof of Theorem 3.4: lower bound. Let $C = 24$. We will show that the expected steady state waiting time w is at least $1/(Cp)$ for all p , giving the result. Assume for contradiction

that there exists p such that $w \leq 1/(Cp)$. By Little's law we have that $w = \mathbb{E}[|\mathcal{V}(\infty)|] \leq 1/(Cp)$ as well. Let i be a node entering at steady state, and let W_i be the waiting time of node i . Let \mathcal{W} be the set of nodes in the system when i arrives, so $\mathcal{W} \stackrel{d}{=} |\mathcal{V}(\infty)|$, and define the event $E_1 = \{|\mathcal{W}| \leq 3w\}$. By Markov's inequality, $\mathbb{P}(E_1) \geq 2/3$. Note that i cannot leave the system until it has an in degree of at least one. Let \mathcal{A} be the first $3w$ arrivals after i , and let the event E_2 be the event that either a node from \mathcal{W} or a node from \mathcal{A} has an edge pointing to i . We have

$$\mathbb{P}(E_2) = \mathbb{P}(\text{Bin}(|\mathcal{W}| + 3w, p) \geq 1),$$

making

$$\mathbb{P}(E_2 | E_1) \leq \mathbb{P}(\text{Bin}(6w, p) \geq 1) \leq \mathbb{P}(\text{Bin}(6/(Cp), p) \geq 1) \leq \frac{6}{C} = \frac{1}{4}$$

using the definition of E_1 , then that $w \leq 1/(Cp)$, then Markov's inequality, and finally that $C = 24$. Thus

$$w = \mathbb{E}[W_i] \geq 3w\mathbb{P}(E_2^c) \geq 3w\mathbb{P}(E_2^c|E_1)\mathbb{P}(E_1) \geq 3w(1 - 1/4)(2/3) = 3w/2 > w$$

providing the contradiction. □

Remark 3. Consider, instead, a setting where two and three-cycles can be removed in addition to chains. Theorem 3.4 tells us that under any policy, we still have a lower bound of $1/(Cp)$ on the expected waiting time (in fact, this holds for arbitrarily long cycles). It also tells us that under a greedy policy that executes only chains (ignoring opportunities to conduct two and three-cycles), the expected waiting time is $O(1/p)$. Further, it is not hard to see that this policy misses two and three cycle opportunities for $O(p)$ fraction of nodes. As such, we conjecture that a greedy policy that executes two and three-cycles in addition to chains will have almost identical performance to greedy with chains only, and in particular, the expected waiting time will still be $O(1/p)$.

8 Conclusion

Overcoming the rare coincidence of wants is a major obstacle in facilitating timely exchanges for agents in a barter marketplace. In this paper we studied how the policy adopted by the clearinghouse affect agents' waiting times in a thin marketplace. We investigated this question for a variety of settings determined by the feasible types of exchanges, which are largely driven by the technology adopted by the marketplace. We also studied how the feasible types of exchanges affect the waiting times.

We studied these questions in a dynamic model with a stylized homogenous stochastic demand structure. The market is represented by a compatibility graph: agents are represented by nodes, and each directed edge, which represents that the source agent has an item that is acceptable to the target agent, exists a priori with probability p . Exchanges take place in

the form of cycles and chains, where chains are initiated by an altruistic donor who is willing to give away his item without asking for anything in return. The key technical challenge we face is that in our dynamic setting, the compatibility graph between agents present at a particular time has a complicated distribution that depends on the feasible exchanges and the policy employed by the clearinghouse.

We analyzed the long run average time agents spend waiting to exchange their item, in a variety of settings depending on the feasible exchanges, 2-way cycles, 2 and 3-way cycles, or chains. Our main finding is that regardless of the setting, the greedy policy which attempts to match upon each arrival, is approximately optimal (minimizes average waiting time) among a large class of policies that includes batching policies. We also find that three-way cycles and chains lead to a large improvement in waiting times relative to two-cycles only. Although we do not model important details of kidney exchange clearinghouses, our findings are consistent with computational experiments and practice in that context.

We now discuss some additional implications of our work. First, consider a setting where there is competition between clearinghouses. In the presence of such competition, where the same agents may participate in multiple clearinghouses, there is an incentive for clearinghouses to complete exchanges as quickly as possible to avoid agents completing an exchange in a different clearinghouse. A priori one may worry that such an incentive may lead clearinghouses to hurt social welfare when they adopt a greedy-like policy. However, our work suggests that this is not the case, since greedy may be near optimal also from the users' perspective (at least in terms of average waiting times).

Second, though we motivated our work primarily in the context of centralized marketplaces, our results also have implications for decentralized marketplaces. One implication is that while organizing longer cycles and chains may require a centralized marketplace, our results imply that this may be an option worth considering (only two-way exchanges are typically possible in a decentralized marketplace). Another implication is more subtle: decentralized marketplaces can modulate the dynamic evolution of the marketplace via devices like visibility (hiding possible matches from market participants may allow simulation of policies different from greedy). Since greedy is near optimal, we conclude that the obvious approach of immediately informing participants of all available matches should typically work best.

Our work raises several further questions and we describe here a few of these. Allowing for heterogeneous agents or goods may lead to different qualitative results in some settings. For example, if Bob is a very difficult-to-please agent who is willing to accept only Alice's item but they are not both currently part of any feasible exchange, it may be beneficial to make Alice wait for some time in the hope of finding an exchange that can allow Bob to get Alice's item. Thus, when chains or cycles of more than two agents are permitted, some waiting may improve efficiency in the presence of heterogeneity (some evidence for this is given by [Ashlagi et al. \(2013\)](#)).

In kidney exchange, patient-donor pairs can roughly be classified as easy-to-match and hard-to-match.²⁸ Many hospitals internally match their easy-to-match pairs, and enrol their

²⁸Whether a pair is easy-to-match or not depends on the blood types, the donor's antigens and the patient's

harder-to-match pairs (Ashlagi and Roth, 2011) in centralized multi-hospital clearinghouses. An important question is how much waiting times of hard-to-match pairs will improve as the percentage of easy-to-match pairs grows.

Allowing for agents' abandonment and outside options are other issues worth exploring (note here that intuitively, our "greedy is near optimal" finding should be reinforced in a situation where agents abandon after an unknown, exponentially distributed time). Designing "good" mechanisms that make it incentive compatible for agents to participate in barter exchanges is another important direction (see a concurrent work by Akbarpour et al. (2014), cf. Section 1.3, who provide a truthful mechanism for agents to report their "deadlines" in a particular barter setting).

References

- ABDULKADIROĞLU, A. AND SÖNMEZ, T. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66, 3, 689–701.
- AJTAI, M., KOMLÓS, J., AND SZEMERÉDI, E. 1981. The longest path in a random graph. *Combinatorica* 1, 1, 1–12.
- AKBARPOUR, M., SHENGWU, L., AND GHARAN, S. 2014. Dynamic matching market design. Working paper.
- ALON, N. AND SPENCER, J. 2008. *The Probabilistic Method*. A John Wiley & Sons.
- ANDERSON, R. 2014. Stochastic models and data driven simulations for healthcare operations. Ph.D. thesis, Massachusetts Institute of Technology.
- ANDERSON, R., ASHLAGI, I., GAMARNIK, D., HIL, G., MELCHER, M., AND ROTH, A. 2014. The effect of match run frequencies and priorities for hard-to-match pairs on the number of transplants and waiting times in kidney paired donation.
- ASHLAGI, I., GAMARNIK, D., REES, M., AND ROTH, A. 2012. The need for (long) chains in kidney exchange. Tech. rep., National Bureau of Economic Research.
- ASHLAGI, I., JAILLET, P., AND MANSHADI, V. 2013. Kidney exchange in dynamic sparse heterogeneous pools. *arXiv preprint arXiv:1301.3509*.
- ASHLAGI, I. AND ROTH, A. E. 2011. Free riding and participation in large scale, multi-hospital kidney exchange. NBER paper no. w16720.
- ASMUSSEN, S. 2003. *Applied probability and queues*. Springer Verlag.
- BOLLOBÁS, B. 2001. *Random graphs*. Vol. 73. Cambridge university press.

antibodies.

- BUDISH, E., CRAMTON, P., AND SHIM, J. 2013. The high-frequency trading arms race: Frequent batch auctions as a market design response. Working paper.
- DICKERSON, J. P., PROCACCIA, A. D., AND SANDHOLM, T. 2012. Optimizing Kidney Exchange with Transplant Chains: Theory and Reality. *Proceedings of the eleventh international conference on autonomous agents and multiagent systems*.
- DUR, U. AND UNVER, U. 2012. Tuition exchange. Working paper.
- FOSTER, F. ET AL. 1953. On the stochastic matrices associated with certain queuing processes. *The Annals of Mathematical Statistics* 24, 3, 355–360.
- GAMARNIK, D. AND ZEEVI, A. 2006. [Validity of heavy traffic steady-state approximations in generalized Jackson networks](#). *The Annals of Applied Probability* 16, 1, 56–90.
- GLYNN, P. AND ZEEVI, A. 2006. [Bounding stationary expectations of Markov processes](#). In *Markov processes and related topics: A Festschrift for Thomas G. Kurtz. Selected papers of the conference, Madison, WI, USA, July*. 10–13.
- GURVICH, I. AND WARD, A. 2012. On the dynamic control of matching queues. *Preprint*.
- HOLEWIJN, P. AND HORDIJK, A. 1975. On the convergence of moments in stationary markov chains. *Stochastic Processes and Their Applications* 3, 1, 55–64.
- JANSON, S., LUCZAK, T., AND KOLCHIN, V. 2000a. *Random graphs*. Cambridge Univ Press.
- JANSON, S., LUCZAK, T., AND RUCINSKI, R. 2000b. *Random Graphs*. A John Wiley & Sons.
- JEVONS, W. S. 1876. *Money and the Mechanism of Exchange*. New York: D. Appleton and Company. <http://www.econlib.org/library/YPDBooks/Jevons/jvnMME.html>.
- KLEINROCK, L. 1975. *Theory, volume 1, Queueing systems*. Wiley-interscience.
- KRIVELEVICH, M., LUBETZKY, E., AND SUDAKOV, B. 2012. Longest cycles in sparse random digraphs. *Random Structures & Algorithms*.
- MELCHER, M., LEESER, D., GRITSCH, H., MILNER, J., KAPUR, S., BUSQUE, S., AND ET AL., J. R. 2012. Chain Transplantation: Initial Experience of a Large Multicenter Program. *American Journal of Transplantation* 12(9), 2429–2436.
- MENDELSON, H. 1982. Market Behavior in a Clearing House. *Econometrica* 50, 6, 1505–1524.
- PYCIA, M. AND UNVER, U. 2011. Incentive compatible allocation and exchange of discrete resources. Available at SSRN 1079505.

- ROTH, A. AND POSTLEWAITE, A. 1977. Weak versus strong dominance in a market with indivisible goods. *Journal of Mathematical Economics* 4, 131–137.
- ROTH, A. E. 1982. Incentive Compatibility in a Market with Indivisible Goods. *Economics Letters* 9, 127–132.
- ROTH, A. E., SÖNMEZ, T., AND ÜNVER, M. U. 2005. Pairwise kidney exchange. *Journal of Economic Theory* 125, 151–188.
- ROTH, A. E., SÖNMEZ, T., AND ÜNVER, M. U. 2007. Efficient kidney exchange: coincidence of wants in markets with compatibility-based preferences. *American Economic Review* 97, 828–851.
- ROTH, A. E., SÖNMEZ, T., ÜNVER, M. U., DELMONICO, F. L., AND SAIDMAN, S. L. 2006. Utilizing list exchange and nondirected donation through chain kidney paired donations. *American Journal of Transplantation* 6, 2694–2705.
- SAMUEL-CAHN, E. 1974. Asymptotic distributions for occupancy and waiting time problems with positive probability of falling through the cells. *The Annals of Probability* 2, 3, 515–521.
- SHAPLEY, L. AND SCARF, H. 1974. On the core and indivisibility. *Journal of Mathematical Economics* 1, 23–37.
- TWEEDIE, R. 1983. The existence of moments for stationary markov chains. *Journal of Applied Probability*, 191–196.
- ÜNVER, M. U. 2010. Dynamic Kidney Exchange. *Review of Economic Studies* 77(1), 372–414.
- WANG, Y. AND KRISHNA, A. 2006. Timeshare exchange mechanisms. *Management Science* 52, 8, 1223–1237.

A Proof of [Section 6](#) preliminaries

Proof of [Proposition 6.2](#)

[Proposition 6.2](#) is obtained as a corollary of [Proposition A.1](#) and [Proposition A.2](#) below.

Suppose $\{X_k\}$ is a discrete time irreducible Markov chain on a countable state space \mathcal{X} . First we give a condition for the positive recurrence of $\{X_k\}$ due to [Foster et al. \(1953\)](#), see [Asmussen \(2003\)](#) for a modern reference. We use \mathbb{E}_x to denote the expectation operator conditional on $X_0 = x$.

Proposition A.1 (Foster et al. 1953). *If there exists a function $V: \mathcal{X} \rightarrow \mathbb{R}$, $\gamma > 0$, and a finite set $B \subset \mathcal{X}$ such that for all $x \in B$,*

$$\mathbb{E}_x[V(X_1) - V(X_0)] < \infty, \quad (37)$$

and for all $x \in \mathcal{X} \setminus B$,

$$\mathbb{E}_x[V(X_1) - V(X_0)] \leq -\gamma,$$

then $\{X_k\}$ is positive recurrent.

Now, suppose that $\{X_k\}$ is positive recurrent, and let X_∞ denote the unique steady state distribution. We now give a bound on the first moment of $f(X_\infty)$ for any function f . The result below is from Anderson (2014) but is similar to Gamarnik and Zeevi (2006); Glynn and Zeevi (2006).

Proposition A.2 (Anderson 2014). *Suppose that X_t is positive recurrent and that there exist $\alpha, \beta, \gamma > 0$, a set $B \subset \mathcal{X}$ and functions $U: \mathcal{X} \rightarrow \mathbb{R}_+$ and $f: \mathcal{X} \rightarrow \mathbb{R}_+$ such that for $x \in \mathcal{X} \setminus B$,*

$$\mathbb{E}_x[U(X_1) - U(X_0)] \leq -\gamma f(x), \quad (38)$$

and for $x \in B$,

$$f(x) \leq \alpha, \quad (39)$$

$$\mathbb{E}_x[U(X_1) - U(X_0)] \leq \beta. \quad (40)$$

Then

$$\mathbb{E}[f(X_\infty)] \leq \alpha + \frac{\beta}{\gamma}.$$

Note that we need not assume that \mathcal{B} is bounded. Finally, we can prove the specialization of the above results as used in the paper.

Proof of Proposition 6.2. First, we apply Proposition A.1 to X_k using the same $V(x)$, \mathcal{B} , and $\gamma = \lambda \mathbb{E}[A_k]$ as in the statement of Proposition 6.2. For $x \notin \mathcal{B}$, we have

$$\mathbb{E}_x[V(X_1) - V(X_0)] = \mathbb{E}_x[A_0 - D_0] \leq \mathbb{E}_x[A_0 - \tilde{D}_0] \leq -\lambda \mathbb{E}_x[A_0] = -\gamma,$$

where in the inequalities we use $\tilde{D}_k \leq D_k$ and then (5). For all x , we have

$$\mathbb{E}_x[V(X_1) - V(X_0)] = \mathbb{E}_x[A_0 - D_0] \leq \mathbb{E}_x[A_0] < \infty.$$

Thus as \mathcal{B} is bounded, we can apply Proposition A.1 to obtain positive recurrence of X_k . Let X_∞ be the steady state version of the Markov chain X_k .

Next, we apply [Proposition A.2](#) taking $U(x) = V^2(x)$ and $f(x) = V(x)$. We let

$$\alpha' = \max \left\{ \alpha, \frac{\max\{1, C_2 - 1\}^2 C_1 \mathbb{E}[A_k]}{\lambda} \right\}, \quad (41)$$

thus making our set of exceptions from [Proposition A.2](#) $\mathcal{B}' = \{x \in \mathcal{X} \mid V(x) \leq \alpha'\}$. We have for $x \in \mathcal{B}'$ that

$$\mathbb{E}_x[U(X_1) - U(X_0)] = \mathbb{E}_x [(V(X_0) + A_0 - D_0)^2 - V(X_0)^2] \quad (42)$$

$$\leq \mathbb{E}_x [(V(X_0) + A_0)^2 - V(X_0)^2] \quad (43)$$

$$= 2V(x)\mathbb{E}[A_0] + \mathbb{E}[A_0^2] \quad (44)$$

$$\leq 2\alpha'\mathbb{E}[A_0] + C_1\mathbb{E}[A_0]^2 \quad (45)$$

$$\leq 2\alpha'\mathbb{E}[A_0] + \alpha'\lambda\mathbb{E}[A_0] \quad (46)$$

$$= \alpha'(2 + \lambda)\mathbb{E}[A_0] \quad (46)$$

$$\triangleq \beta' \quad (46)$$

where (42) follows from (4), (43) follows as $V(X_1) \geq 0$, and (44) follows from (41) and the definition of C_1 , and (45) follows again from (41) and as $\max\{1, C_2 - 1\}^2 \geq 1$ by definition. We have for $x \notin \mathcal{B}'$, that

$$\mathbb{E}_x[U(X_1) - U(X_0)] = \mathbb{E}_x [(V(X_0) + A_0 - D_0)^2 - V(X_0)^2] \quad (47)$$

$$\leq \mathbb{E}_x [(V(X_0) + A_0 - \tilde{D}_0)^2 - V(X_0)^2] \quad (48)$$

$$= 2V(x)\mathbb{E}_x[A_0 - \tilde{D}_0] + \mathbb{E}_x [(A_0 - \tilde{D}_0)^2] \quad (49)$$

$$\leq -2V(x)\lambda\mathbb{E}[A_0] + \mathbb{E} [\max\{1, C_2 - 1\}^2 A_0^2] \quad (49)$$

$$\leq -(V(x) + \alpha')\lambda\mathbb{E}_x[A_0] + \max\{1, C_2 - 1\}^2 C_1 \mathbb{E} [A_0]^2 \quad (50)$$

$$\leq -(V(x) + \alpha')\lambda\mathbb{E}_x[A_0] + \alpha'\lambda\mathbb{E} [A_0] \quad (51)$$

$$= -V(x)\lambda\mathbb{E}_x[A_0]$$

where (47) follows from (4), (48) follows as $V(X_1) \geq 0$ and $\tilde{D}_0 \leq D_0$, (49) follows from (5) and as $\tilde{D}_k \leq C_2 A_k$ a.s. implies that $|A_k - D_k| \leq \max\{1, C_2 - 1\} A_k$ a.s., (50) follows from (41) and the definition of C_1 , and finally (51) follows again from (41). Thus by taking $\gamma' = \lambda\mathbb{E}_x[A_0]$, we can now apply [Proposition A.2](#) with α' , β' and γ' to obtain that

$$\begin{aligned} \mathbb{E}[V(\infty)] &\leq \alpha' + \frac{\beta'}{\gamma'} = \alpha' + \frac{\alpha'(2 + \lambda)\mathbb{E}[A_0]}{\lambda\mathbb{E}[A_0]} = \alpha' \left(1 + \frac{2 + \lambda}{\lambda} \right) \\ &= \max \left\{ \alpha, \frac{\max\{1, C_2 - 1\}^2 C_2 \mathbb{E}[A_k]}{\lambda} \right\} \left(1 + \frac{2 + \lambda}{\lambda} \right) \end{aligned}$$

showing the result. □

Last, we give a quick counter example showing that without some assumptions beyond simply having negative drift, we may not even have a finite first moment.

Example 1. Consider the following random walk X_t on the nonnegative integers parametrized by some $\gamma \in (0, 1)$. From state 0, we always go up to state 1. For every other state $k = 1, 2, \dots$, with probability $(1 + \gamma)/(k + 1)$, we go to state 0, and with the remaining probability, $(k - \gamma)/(k + 1)$, we go up to state $k + 1$. This walk has the property that for all $k \geq 1$,

$$\mathbb{E}_k[X_1 - X_0] = (k + 1) \cdot \frac{k - \gamma}{k + 1} + 0 \cdot \frac{1 + \gamma}{k + 1} - k = -\gamma,$$

and thus is positive recurrent and has some stationary distribution $\pi_k = \mathbb{P}(X_\infty = k)$. However, we will show that $\mathbb{E}[X_\infty] = \infty$. A direct computation of the steady-state equations gives that $\pi_0 = \pi_1$, and for $n \geq 2$,

$$\pi_n = \frac{n - \gamma}{n + 1} \pi_{n-1} = \pi_0 \prod_{k=2}^n \frac{k - \gamma}{k + 1} = \frac{\pi_0}{\Gamma(2 - \gamma)} \frac{\Gamma(n + 1 - \gamma)}{\Gamma(n + 2)},$$

where Γ is the Gamma function. Using the identity

$$\lim_{n \rightarrow \infty} \frac{\Gamma(n + \alpha)}{\Gamma(n)n^\alpha} = 1$$

for all $\alpha \in \mathbb{R}$, we have that

$$\pi_n = \Theta\left(\frac{1}{n^{1+\gamma}}\right).$$

Thus there exists $c > 0$ and $\ell \in \mathbb{Z}_+$ such that

$$\mathbb{E}[X_\infty] = \sum_{n=0}^{\infty} n\pi_n \geq \sum_{n=\ell}^{\infty} \frac{c}{n^\gamma} = \infty,$$

showing the claim.

Proofs of Proposition 6.3 and Lemma 6.4

Proof of Proposition 6.3. We prove the result by contradiction. Assume that a maximal set of node disjoint three-cycles \mathcal{W} contained fewer than $N/3$ three-cycles. Then there must be a three-cycle X from a largest set of node disjoint three-cycles such that for every three-cycle $Y \in \mathcal{W}$, X and Y have no nodes in common. This yields a contradiction, as we could then add X to \mathcal{W} to make a larger set of node disjoint three-cycles, thus making \mathcal{W} not maximal. \square

Proof of Lemma 6.4. We assume that the removal policy is deterministic. The proof for the case of randomized policies follows immediately. Fix any two nodes i, j which arrive before time t (namely $i, j \leq t$). Given any directed graph \mathcal{G} on nodes $0, 1, \dots, t$ (that is nodes arriving up to time t) such that the edge (i, j) belongs to \mathcal{G} , denote by $\bar{\mathcal{G}}$ the same graph \mathcal{G} with edge (i, j) deleted. Let \mathcal{W} be any subset of nodes $0, 1, \dots, t$ containing i and j . Recall that we denote by \mathcal{G}_t the directed graph generated by nodes $0, 1, \dots, t$ and by \mathcal{W}_t the set of nodes observed at time t . Note that, since the policy is deterministic, graph \mathcal{G}_t uniquely determines the set of nodes \mathcal{W}_t .

We have

$$\mathbb{P}(\mathcal{W}_t = \mathcal{W}) = \sum_{\mathcal{G}} \mathbb{P}(\mathcal{G}) + \sum_{\bar{\mathcal{G}}} \mathbb{P}(\bar{\mathcal{G}}),$$

where the first sum is over graphs \mathcal{G} containing edge (i, j) such that the set of nodes observed at time t is \mathcal{W} when $\mathcal{G}_t = \mathcal{G}$, and the second sum is over graphs $\bar{\mathcal{G}}$ containing edge (i, j) , such that when $\mathcal{G}_t = \bar{\mathcal{G}}$, the set of nodes observed at time t is \mathcal{W} . Note, however that by our monotonicity assumption, if $\mathcal{G}_t = \mathcal{G}$ implies $\mathcal{W}_t = \mathcal{W}$, then $\mathcal{G}_t = \bar{\mathcal{G}}$ also implies $\mathcal{W}_t = \mathcal{W}$. Thus

$$\mathbb{P}(\mathcal{W}_t = \mathcal{W}) \geq \sum_{\mathcal{G}} (\mathbb{P}(\mathcal{G}) + \mathbb{P}(\bar{\mathcal{G}})),$$

where the sum is over graphs \mathcal{G} containing edge (i, j) such that $\mathcal{G}_t = \mathcal{G}$ implies $\mathcal{W}_t = \mathcal{W}$. At the same time note that $\mathbb{P}(\bar{\mathcal{G}}) = \mathbb{P}(\mathcal{G})(1 - p)/p$ since it corresponds to the same graph except edge (i, j) deleted. We obtain

$$\mathbb{P}(\mathcal{W}_t = \mathcal{W}) \geq \sum_{\mathcal{G}} \mathbb{P}(\mathcal{G})(1 + (1 - p)/p) = \sum_{\mathcal{G}} \mathbb{P}(\mathcal{G})/p.$$

We recognize the right-hand side as $\mathbb{P}(\mathcal{W}_t = \mathcal{W} | (i, j) \in \mathcal{G}_t)$. Now we obtain

$$\begin{aligned} \mathbb{P}((i, j) \in \mathcal{G}_t | \mathcal{W}_t = \mathcal{W}) &= \mathbb{P}(\mathcal{W}_t = \mathcal{W} | (i, j) \in \mathcal{G}_t) \mathbb{P}((i, j) \in \mathcal{G}_t) / \mathbb{P}(\mathcal{W}_t = \mathcal{W}) \\ &\leq \mathbb{P}((i, j) \in \mathcal{G}_t) \\ &\leq p, \end{aligned}$$

and the claim is established. □

B Proof of Section 7 preliminaries

We first prove Corollary 7.2. The proof follows relatively easily from Proposition 7.1. The idea is as follows. A sufficient condition to form a long chain from a node v is for v to be a member of the long cycle that will occur with high probability according to the proposition. Note that with constant probability e^{-c} , v will be isolated and thus not be part of the cycle, but we can make this probability small by taking C large.

Proof of Corollary 7.2. Given ε from the statement of Corollary 7.2, let \bar{C} and \bar{n}_0 be values guaranteed to exist from Proposition 7.1 applied when $\delta = 1$ and the probability of a long chain existing is at least $1 - \varepsilon/2$.

There exists C^* such that for all $c > C^*$, $3ce^{-c} < \varepsilon/2$ as the function $f(x) = xe^{-x}$ is strictly decreasing for $x > 1$. We claim that given our ε , Corollary 7.2 holds by taking $C = \max\{\bar{C}, C^*\}$ and $n_0 = \bar{n}_0$.

Given our $\text{ER}(n, c/n)$ graph where $n > n_0$ and $c > C$ and a fixed node v , let A be the event it contains a cycle of length at least $(1 - 3ce^{-c})n$, and let $B \subset A$ be the event that v is in the cycle. Observe that it suffices to prove that $P(B) > 1 - \varepsilon$ to show the result, as $3ce^{-c} < \varepsilon$ by our assumption that $c > C \geq C^*$ and the definition of C^* . Thus we compute that

$$\mathbb{P}(B) = P(B|A)\mathbb{P}(A) \geq (1 - \varepsilon/2)(1 - \varepsilon/2) \geq 1 - \varepsilon,$$

showing the result, where $\mathbb{P}(A) \geq 1 - \varepsilon/2$ follows from Proposition 7.1 and $P(B|A) \geq 1 - \varepsilon/2$ follows as the cycle is equally likely to pass through every node, so when the cycle hits $1 - \varepsilon/2$ fraction of the nodes, it has this probability of hitting v . \square

Next, we prove Corollary 7.3. The idea of the proof is as follows. First, we show with a simple calculation that a constant fraction of the nodes in \mathcal{R} will have both in and out degree one, as $p \rightarrow 0$. We consider paths which only use this subset of nodes from \mathcal{R} . Such a path is equivalent to a path in a modified graph on the set of nodes \mathcal{L} where there is an edge between two nodes u and v if and only if there is a path of length two between them via an intermediate node in \mathcal{R} which has in and out degree one. Such a graph behaves (approximately) as an Erdős-Rényi graph on the nodes of \mathcal{L} , with the number of edges proportionate to $|\mathcal{R}|$. Thus by ensuring that $|\mathcal{R}|$ is sufficiently large, we can apply Corollary 7.2 to obtain the result.

Proof of Corollary 7.3. Fix $\kappa > 1$ and $\varepsilon > 0$ from the statement of the corollary. For C and p_0 to be chosen later, let $c_L \in [1/\sqrt{\kappa}, \kappa]$, $c_R > C$, and $p < p_0$ be arbitrary. Given our graph $\mathcal{G} = (\mathcal{L}, \mathcal{R}, \mathcal{E})$ that is $\text{ER}(c_L/p, c_R/p, p)$, consider the subgraph $\mathcal{G}' = (\mathcal{L}, \mathcal{R}', \mathcal{E}')$ of \mathcal{G} where \mathcal{R}' is the set of vertices in \mathcal{R} with in degree one and out degree one in \mathcal{G} , and \mathcal{E}' are the edges in \mathcal{E} such that both endpoints are in \mathcal{G}' . From this graph, we create a new directed non-bipartite digraph $\mathcal{G}'' = (\mathcal{L}, \mathcal{E}'')$ where there is an edge from $u \in \mathcal{L}$ to $v \in \mathcal{L}$ iff there is at least one node $r \in \mathcal{R}'$ such that $(u, r) \in \mathcal{E}'$ and $(r, v) \in \mathcal{E}'$. Observe that a path of length k in \mathcal{G}'' gives a path of length $2k$ in \mathcal{G} by following the two edges in \mathcal{G}' for each edge in the path on \mathcal{G}'' , so it suffices to find a path of length $(1 - \varepsilon)c_L/p$ in \mathcal{G}'' .

For any vertex $r \in \mathcal{R}$, let I_r be the indicator variable that r has an in degree of one and an out degree of one. Note that these variables are independent. Further, we have

$$\mu(p) \triangleq \mathbb{P}(I_r = 1) = \mathbb{P}(\text{Bin}(|\mathcal{L}|, p) = 1)^2 = \left(\frac{c_L}{p} p(1-p)^{\frac{c_L}{p}-1} \right)^2 \rightarrow c_L^2 \exp(-2c_L),$$

as $p \rightarrow 0$. As each of the I_r are independent, we have that

$$|\mathcal{R}'| \stackrel{d}{=} \text{Bin} \left(\frac{c_R}{p}, \mu(p) \right).$$

Letting

$$A_1(\delta_1) = \left\{ (1 - \delta_1) \frac{c_R}{p} \mu(p) < |\mathcal{R}'| < (1 + \delta_1) \frac{c_R}{p} \mu(p) \right\}$$

we have by [Proposition 6.1](#) that for all p ,

$$\mathbb{P}(A_1(\delta_1)) \geq 1 - 2 \exp \left(-\delta_1^2 \frac{c_R}{p} \frac{\mu(p)}{3} \right).$$

We can view the edges of \mathcal{G}'' as being generated by the following process: for each $r \in \mathcal{R}'$, pick a source and then a destination uniformly at random from \mathcal{L} and add an edge from the source to the destination unless either:

- the source and destination are the same node,
- an edge between the source and destination already exists in the graph.

Thus $|\mathcal{E}''|$ is the number of non empty bins if we throw $|\mathcal{R}'|$ balls into $(c_L/p)^2$ bins and then throw out the c_L/p bins that correspond to self edges. (Alternatively, we can think of this process as throwing c_R/p balls, but each ball “falls through” only with probability $1 - \mu(p)$. This problem was studied extensively in [Samuel-Cahn \(1974\)](#), but here we need only a coarse analysis). Trivially, $|\mathcal{E}''| \leq |\mathcal{R}'|$. We now show that typically, the number of nonempty bins is almost equal to the number of balls thrown. For each $r \in \mathcal{R}'$, let X_r be the indicator that there is $\ell \in \mathcal{L}'$ such that $(\ell, r) \in \mathcal{E}'$ and $(r, \ell) \in \mathcal{E}'$. It is easy to see that the X_r are i.i.d. Bernoulli(p/c_L). For each $\{r, s\} \subset \mathcal{R}'$, let $Y_{\{rs\}}$ be the indicator that the nodes r and s are “colliding” on both their source and destination choices in \mathcal{L}' , i.e. there is $\ell, m \in \mathcal{L}'$, $\ell \neq m$, such that $(\ell, r), (\ell, s), (r, m), (s, m) \in \mathcal{E}'$. It is easy to see that $\mathbb{P}(Y_{\{rs\}} = 1) \leq p^2/c_L^2$ for each $\ell, m \in \mathcal{L}'$. We have

$$|\mathcal{E}''| \geq |\mathcal{R}'| - \sum_{r \in \mathcal{R}'} X_r - \sum_{\{r,s\} \subset \mathcal{R}'} Y_{\{rs\}}.$$

We compute that for any fixed \mathcal{R}'

$$\mathbb{E} \left[\sum_{r \in \mathcal{R}'} X_r + \sum_{\{r,s\} \subset \mathcal{R}'} Y_{\{rs\}} \right] \leq |\mathcal{R}'| \frac{p}{c_L} + \binom{|\mathcal{R}'|}{2} \frac{p^2}{c_L^2} \leq |\mathcal{R}'| \frac{p}{c_L} + \left(|\mathcal{R}'| \frac{p}{c_L} \right)^2$$

Letting

$$A_2(\delta_2) = \left\{ \sum_{r \in \mathcal{R}'} X_r + \sum_{\{r,s\} \subset \mathcal{R}'} Y_{rs} \leq \delta_2 |\mathcal{R}'| \right\},$$

we have that

$$\mathbb{P}(A_2(\delta_2)) \geq 1 - \frac{p}{\delta_2 c_L} - |\mathcal{R}'| \delta_2^{-1} \left(\frac{p}{c_L} \right)^2$$

Letting

$$B(\delta_1, \delta_2) = \left\{ (1 - \delta_1)(1 - \delta_2) \frac{c_R}{p} \mu(p) < |\mathcal{E}''| < (1 + \delta_1) \frac{c_R}{p} \mu(p) \right\},$$

we have that $B(\delta_1, \delta_2) \supset A_1(\delta_2) \cap A_2(\delta_2)$, and thus by taking complements and then applying the union bound,

$$\begin{aligned} \mathbb{P}(B(\delta_1, \delta_2)) &\geq 1 - \mathbb{P}(A_1(\delta_1)^c) - \mathbb{P}(A_2(\delta_2)^c) \\ &\geq 1 - 2 \exp\left(-\delta_1^2 \frac{c_R}{p} \mu(p)/3\right) - \frac{p}{\delta_2 c_L} - (1 + \delta_1) c_R \mu(p) \frac{p}{\delta_2^2 c_L^2}, \end{aligned}$$

thus giving us a high probability bound on the size of $|\mathcal{E}''|$ as $p \rightarrow 0$.

For our fixed ε , let \tilde{C} and \tilde{n}_0 be C and n_0 from [Corollary 7.2](#) such that for any $c > \tilde{C}$ and $n > \tilde{n}_0$, given a node in graph $\text{ER}(n, c/n)$, there exists a path with length at least $n(1 - \varepsilon/2)$ with probability at least $1 - \varepsilon/2$. We now specify p_0 from the corollary to be such that for all $c_L \in [1/\sqrt{\kappa}, \kappa]$, we have $c_L/p_0 > \tilde{n}_0$, i.e. $p_0 < 1/(n_0\sqrt{\kappa})$.

Let $\tilde{\mathcal{G}} = (\mathcal{L}, \tilde{\mathcal{E}})$ be an $\text{ER}(c_L/p, \tilde{C}p/c_L)$ directed random graph. We now couple \mathcal{G}'' (a directed $\text{ER}(n, M)$ graph, where M is random but independent of the edges selected) and $\tilde{\mathcal{G}}$ (a directed $\text{ER}(n, p)$ graph) in the standard way so that when $|\tilde{\mathcal{E}}| \leq |\mathcal{E}''|$, then $\tilde{\mathcal{E}} \subset \mathcal{E}''$ and when $|\mathcal{E}''| \leq |\tilde{\mathcal{E}}|$, then $\mathcal{E}'' \subset \tilde{\mathcal{E}}$. Thus if $\tilde{\mathcal{G}}$ has a long path and $|\tilde{\mathcal{E}}| < |\mathcal{E}''|$, then \mathcal{G}'' will have at least as long a path as well, as it will contain more edges on the same nodes. Let \tilde{P} be the length of a longest path starting at v in $\tilde{\mathcal{G}}$. Letting

$$A_3 = \left\{ \tilde{P} > \left(1 - \frac{\varepsilon}{2}\right) \frac{c_L}{p} \right\}$$

and recalling that $p_0 < 1/(n_0\sqrt{\kappa})$ implies that $c_L/p > \tilde{n}_0$, we have by [Proposition 7.1](#)

$$\mathbb{P}(A_3) \geq 1 - \frac{\varepsilon}{2}$$

We now need to show that \mathcal{G}'' will have more edges than $\tilde{\mathcal{G}}$ with high probability for all c_R sufficiently large (which we can control by choice of C from the statement of the corollary). We have $|\tilde{\mathcal{E}}| \sim \text{Bin}((c_L/p - 1)c_L/p, \tilde{C}p/c_L)$, thus by [Proposition 6.1](#), if

$$A_4(\delta_4) = \left\{ \tilde{C}(c_L/p - 1)(1 - \delta_4) < |\tilde{\mathcal{E}}| < \tilde{C}(c_L/p - 1)(1 + \delta_4) \right\}$$

then

$$\mathbb{P}(A_4(\delta_4)) \geq 1 - 2 \exp\left(-\delta_4^2 \tilde{C} \left(\frac{c_L}{p} - 1\right) / 3\right)$$

Now, for any fixed choice of $\delta_1, \delta_2, \delta_4$, there exists C sufficiently large such that if $c_R > C$ then for all $p < p_0$ and all $c_L \in [1/\sqrt{\kappa}, \kappa]$,

$$\tilde{C} \left(\frac{c_L}{p} - 1 \right) (1 + \delta_4) < (1 - \delta_1)(1 - \delta_2) \frac{c_R}{p} \mu(p)$$

(recall that $\mu(p)$ converges to a constant depending only c_L uniformly over $[1/\sqrt{\kappa}, \kappa]$ as $p \rightarrow 0$). For such c_R , we have

$$\{|\mathcal{E}''| > |\tilde{\mathcal{E}}|\} \subset B(\delta_1, \delta_2) \cap A_4(\delta_4),$$

as B makes $|\mathcal{E}''|$ big and A_4 ensures that $|\tilde{\mathcal{E}}|$ is small. Putting everything together, we have that

$$\left\{ P > 2 \frac{c_L}{p} (1 - \varepsilon) \right\} \supset B(\delta_1, \delta_2) \cap A_3 \cap A_4(\delta_4),$$

so by taking complements and then applying the union bound, we obtain

$$\mathbb{P} \left(P > 2 \frac{c_L}{p} (1 - \varepsilon) \right) \geq 1 - \mathbb{P}(B(\delta_1, \delta_2)^c) - \mathbb{P}(A_3^c) - \mathbb{P}(A_4(\delta_4)^c) = 1 - \frac{\varepsilon}{2} - O(p),$$

showing the result. □

Proof of Lemma 7.4. By the memoryless property of the geometric distribution, for all $t > 0$,

$$\mathbb{E}[X \mid X > t] = t + \mathbb{E}[X] = t + \frac{1}{p}. \quad (52)$$

Thus for all sufficiently large κ we have

$$\mathbb{E} \left[X \mathbb{I}_{X > \frac{\kappa}{p}} \right] = (1 - p)^{\frac{\kappa}{p}} \left(\frac{\kappa}{p} + \frac{1}{p} \right) \quad (53)$$

$$\leq e^{-\kappa} \frac{1 + \kappa}{p} \quad (54)$$

$$\leq \frac{1}{2\kappa p}, \quad (55)$$

where (53) follows from (52), (54) follows as $(1 - p)^{1/p} \leq e^{-1}$ for all p (take logarithms), and finally (55) holds provided $\kappa \geq \kappa_0$ for appropriately large κ_0 .

For the remaining term, we have that for all sufficiently large κ and sufficiently small p ,

$$\begin{aligned}\mathbb{E}\left[X\mathbb{I}_{X\leq\frac{1}{\sqrt{\kappa p}}}\right] &= \mathbb{E}[X] - \mathbb{E}\left[X\mathbb{I}_{X>\frac{1}{\sqrt{\kappa p}}}\right] \\ &= \frac{1}{p} - (1-p)^{\frac{1}{\sqrt{\kappa p}}}\left(\frac{1}{\sqrt{\kappa p}} + \frac{1}{p}\right)\end{aligned}\tag{56}$$

$$\leq \frac{1}{p} - \left(\frac{1-p}{e}\right)^{\frac{1}{\sqrt{\kappa}}}\left(\frac{1}{\sqrt{\kappa p}} + \frac{1}{p}\right)\tag{57}$$

$$\leq \frac{1}{p} - (1-p)^{\frac{1}{\sqrt{\kappa}}}\left(1 - \frac{1}{\sqrt{\kappa}}\right)\left(\frac{1}{\sqrt{\kappa p}} + \frac{1}{p}\right)\tag{58}$$

$$\begin{aligned}&= \frac{1}{p} - (1-p)^{\frac{1}{\sqrt{\kappa}}}\frac{1}{p} + (1-p)^{\frac{1}{\sqrt{\kappa}}}\frac{1}{\kappa p} \\ &\leq \frac{2}{\sqrt{\kappa}} + \frac{1}{\kappa p}\end{aligned}\tag{59}$$

$$\leq \frac{3}{2\kappa p}\tag{60}$$

where (56) follows from (52). To obtain (57), by Taylor's theorem, $(1-p)^{1/p} = e^{-1}(1-p/2) + o(p)$ as $p \rightarrow 0$, thus for sufficiently small p , we have $(1-p)^{1/p} \geq e^{-1}(1-p)$. In (58) we use that $e^{-x} \geq 1-x$, in (59), we use that for all x sufficiently small, $1 \geq (1-x)^n \geq 1-2xn$, and (60) follows by taking κ sufficiently large. Thus the result is shown. \square