X

A non-asymptotic approach to analyzing kidney exchange graphs

YICHUAN DING, Sauder School of Business, University of British Columbia
DONGDONG GE and SIMAI HE, Shanghai University of Finance and Economics
CHRISTOPHER THOMAS RYAN, Booth School of Business, University of Chicago

We propose a novel methodology to study kidney exchange. Taking the random graph model of kidney exchange introduced in Ashlagi, Garmarnik, Rees and Roth's "The need for (long) chains in kidney exchange" (2012), we propose a non-asympotic approach to quantifying the effectiveness of transplant chains in reducing the number of unmatched highly sensitized patients. Our approach is based on a two phase random walk procedure where random walks are used to allocate chains, followed by allocation via matching in cycles. The benefit of random walks is that they preserve the probabilistic structure of residual graphs, greatly facilitating analysis. Our approach allows us to analytically show the benefit of chains, as compared to transplantation in two-way cycles only, in non-asymptotic (medium-sized) graphs. We also derive useful analytical bounds that illustrate the performance of our proposed allocation procedure and more general kidney allocation procedures. Our results complement previous findings on the benefits of chains that includes analytical results in large (limit) graphs and empirical results based on data from fielded kidney exchanges.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

Additional Key Words and Phrases: Kidney exchange, random graphs

1. INTRODUCTION

Allocating donated kidneys to deserving patients with end stage renal disease is an important challenge in today's health care system. Kidney exchanges are an integral part of the allocation system and pertain to the allocation of living donor kidneys. Exchanges consist of pools of patients each paired with a loved one willing to donate. A pair may be incompatible, due to differences in blood type or other tissue sensitivities. The goal of the exchange is to swap donors among incompatible pairs to allow for more transplants. Exchanging kidneys among incompatible patient-donor pairs creates *cycles* of donation within the exchange.

More possibilities for exchange occur in the presence of *altruistic donors* – individuals who are willing to donate their kidney to any patient in need. That is, the kidney of an altruistic donor is not directed to any particular patient. For this reason, altruistic donors are also called nondirected donors (NDDs). There are alternate uses for the donated kidney of an NDD. Until recently, the NDD kidneys were directed to the deceased donor wait-list managed nationally under the aegis of the United Network of Organ Sharing (UNOS). However, within kidney exchanges altruistic donors can initiate donation "chains". A *chain* starts with an altruistic donor offering a kidney to a compatible recipient. The paired donor of that recipient further donates his or her kidney to another compatible recipient, and so on. Since NDDs are not directed towards a particular recipient, a chain need not "cycle" back.

Allocating NDD kidneys among their alternate uses has sparked ethical and practical debate, including whether chains are needed at all ([Roth et al. 2007; Ünver 2010; Ashlagi et al. 2011; Gentry et al. 2009; Woodle et al. 2010]). Theoretical results show that under certain structures short cycles are sufficient, thus eliminating the need for chains [Roth et al. 2007; Ünver 2010]. On the other hand, empirical results and simulations consistently show that short cycles are important in practice. [Ashlagi et al. 2012] and [Dickerson et al. 2012b] resolve this discrepancy between theory and practice. Their analytical results reveal that the underlying sparseness of connections between patients and donors in the exchange is the main driver of the need for chains.

As the above demonstrates, analytical models of kidney exchange have been helpful in resolving debates and providing explanations of experimental data. As new empirical findings and practical issues have arisen, models have been adjusted to meet these challenges. The standard-bearer of analytical work has been random graph models. These models approximate exchanges by generating, in a probabilistic fashion, nodes and arcs that represent recipient-donor pairs and compatibilities, respectively. Recent papers use asymptotic analysis as their primary theoretical workhorse; that is, the examine kidney exchange graphs as the number of nodes tends to infinity. The development of [Ashlagi et al. 2012] described above is a primary example. Earlier theoretical findings (for instance [Roth et al. 2007]) were based on *dense* random graphs and were inadequate to explain empirical findings. By revising the standard model to include *sparse* random graphs, [Ashlagi et al. 2012] are able to theoretically justify the observed need for chains in fielded exchanges using asymptotic analysis. Asymptotic analysis has also been used to derive insights into incentives issues [Toulis and Parkes 2011], the effect of "failed" chains and cycles [Dickerson et al. 2013], myopic versus forward-looking considerations in dynamically allocating kidneys [Ashlagi et al. 2013; Dickerson et al. 2012a] and fairness issues [Dickerson et al. 2014].

However, asymptotic analysis has its limitations. Asymptotic results are best interpreted in the setting of "large" exchanges with many recipient-donor pairs, something not usually observed in practice [Melcher et al. 2012]. Researchers in the area of kidney exchange are well aware of this limitation. Indeed, [Ashlagi et al. 2012] states that analysis in large graphs and "medium" sized graphs should in fact be the target for analysis.

The goal of this paper is to develop a non-asymptotic methodology that applies to medium-sized exchanges. The core novelty of our methodology is to employ a random walk procedure with two distinct phases. The first phase is to allocate kidneys in chains via a memoryless random walk. After chains are removed, the second phase is to allocate via cycles. There are several analytical benefits to using random walks. First, we maintain an independence structure in the residual graph as we remove nodes incrementally. This provides useful formulas for the incremental probabilistic structure of the graph that facilitates the derivation of estimates and bounds.

Second, one of the challenges in analyzing kidney exchange graphs is balancing the efficiency of chains versus the efficiency of cycles. In our approach, the residual graph maintains its initial probabilistic *density*, unaffected by random walk realizations in the first stage. Since our procedure assigns nodes to chains randomly, it does not target high-degree nodes that would allow for longer chains at the cost of increased sparsity at the cycle-formation stage. The probabilistic structure of the residual graph, unaffected by the first stage, facilitates probabilistic comparison with the original graph. This idea is central to our analysis in Section 4.

Using the random walk procedure as a tool for analysis we achieve the following. We provide exact formulas and simple non-asymptotic analytical bounds for the tail probabilities and expectation of the random number of unmatched nodes after the termination of the first stage. Although non-asymptotic, these bounds can be used to recover asymptotic results (as demonstrated in Proposition 3.3). These bounds serve as inputs to further bound the expected number of unmatched nodes after both phases are implemented, assuming particular algorithms for assigning cycles in the second stage. These latter bounds allow us to assess the performance of our two stage procedure and quantify the benefits of chains in medium-sized graphs. Finally, we present results from numerical experiments that explore the tightness of our bounds and an investigation of the practical impact of assigning chains via random walk.

This paper is organized as follows. In Section 2 we introduce our random graph model of kidney exchange and propose our two-stage random walk procedure . Sec-

tion 3 provides analysis of the nature of the graph at the termination of the first stage Section 4 provides results on the nature of the graph after the termination of the second stage. Section 5 describes our numerical experiments. Section 6 concludes.

2. ANALYTICAL FRAMEWORK

2.1. Random graph model

We consider the random graph model of kidney exchange proposed in [Ashlagi et al. 2012]. Similar models are employed in [Ashlagi et al. 2013] and [Dickerson et al. 2013]. Careful justification of this model can be found in these cited papers.

The kidney exchange pool is modeled as a directed graph D which contains two types of nodes: patient-donor nodes and NDD nodes. A directed arc (u, v) connects nodes u and v if the patient of node v is compatible with the donor of node u. Following [Ashlagi et al. 2012], we suppress the issue of blood type matchings and focus instead on tissue-type matching of donors and patients. Arc (u, v) appears in the graph with a probability that depends only on the tissue-type characteristics of node v. Furthermore, patient-donor nodes are classified into two categories: *high-sensitization nodes* and *low-sensitization nodes*. For brevity we call high-sensitization nodes H-nodes and low-sensitization nodes L-nodes. Arc (u, v) appears in the graph with probability p_H (p_L) if v is an H-node (L-node), where u is an arbitrary node (not equal v) in the graph. Throughout we assume $p_H < p_L$. The assumption of two categories of patient-donor nodes is justified by empirical investigations found in [Ashlagi et al. 2013], where it is shown that the probability a patient is compatible with a randomly selected donor follows a bimodal distribution.

Our model considers the possibility that NDDs and bridge donors (donors freed to donate to extend the length of a chain) renege before the time of transplantation. This is captured by the probability r. That is, every time a chain is extended there is a probability r it terminates before the next link in the chain is transplanted.

We use the notation $D(h, \ell, t; p_H, p_L, r)$ to represent an exchange pool with h high sensitization nodes, ℓ low sensitization nodes, and t NDD donors, along with compatibility probabilities p_H and p_L and renege probability r. The proportion $\frac{\ell}{\ell+h}$ of low-sensitization nodes in the graph is denoted by λ. When certain parameters are understood as given we drop them in our notation. For instance, when the focus is on the size of the exchange with probabilities fixed, we will write $D(h, \ell, t)$ instead of $D(h, \ell, t; p_H, p_L, r)$.

A *clearing* of the kidney exchange graph is a collection of disjoint cycles and chains that represent the patients and donors involved in transplantation. Cycles and chains must be disjoint since each patient can receive at most one kidney and every donor can give at most one kidney. In practice, the manager of a kidney exchange clears at regular intervals (weekly, monthly or bimonthly) to balance the objectives of efficiency and fairness (see [Dickerson et al. 2012b] and [Melcher et al. 2012] for details). Our model is static and considers only a single decision period. However, the random walk technique can be adapted to dynamic matchings.

2.2. Random walk procedure

We propose the following two-phase clearing procedure for $D(h, \ell, t; p_H, p_L, r)$, illustrated in Figure 1.

Phase 1: While there exists at least one NDD, initiate a chain starting from an NDD. At each step, grow the chain by adding an H-node accessible from the last node of the chain (referred to as a tail node). If there is more than one accessible H-node, randomly select one with equal probability. If no H-nodes are accessible, terminate the chain and remove all selected nodes in the chain (including the initiating NDD donor).

Repeat until either all H-nodes have been removed or all NDDs have been consumed. Go to Phase 2.

Phase 2: Apply a cycle packing algorithm on the residual graph that remains at the termination of Phase 1.

A few remarks on the procedure are in order. First, chains in Phase 1 are executed within the subgraph of H-nodes and NDD-nodes. There are no L-nodes in the chains of our procedure. Second, Phase 2 does not specify a cycle packing algorithm. The random walk procedure supports many possible specifications of Phase 2. Our analysis in Section 4 provides theoretical bounds for the case where Phase 2 consists of bipartite matching between H- and L-nodes. Section 5 gives numerical results for when Phase 2 employs both 2- and 3-way cycles.

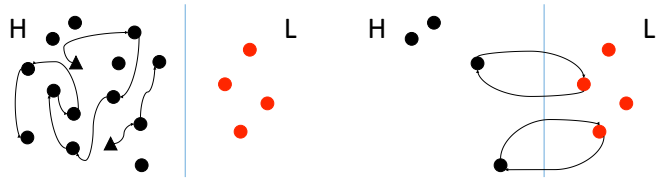

(a) Phase 1: Chain formation in H-nodes via random walk

(b) Phase 2: Cycle packing in the residual graph

Fig. 1. Random walk procedure. Black disks are H-nodes, triangles are NDDs and red disks are L-nodes.

The analytical power of the above procedure derives from the fact we are able to get upper bounds on the expected number of unmatched H-nodes after the termination of the algorithm and compare this to the expected number of unmatched when only cycles are permitted. There are several steps to this analysis. In Section 3 we analyze Phase 1, focusing on probabilistic statements about how many H-nodes have been transplanted. Section 4 explores what happens after Phase 2, leveraging results for Phase 1.

3. ANALYSIS OF PHASE 1

In this section, we define a two-dimensional-state stochastic process that tracks the progress of the random walks, both in terms of transplanting H-nodes and consuming NDD donors. The stochastic process is a Markovian pure death process with absorbing states. The state vector does not record the L-nodes as they are not involved in Phase 1. Counting arguments yield exact probabilities associated with the random number of H-modes left unmatched at the end of Phase 1. To yield more useful non-combinatorial bounds used in Section 4, we later define a potential function and construct martingales to get useful analytical estimates of the expected number of residual unmatched H-nodes.

Let $X(n)$ denote the number of unmatched H-nodes at the time when n nodes (either H-nodes or NDD donor nodes) have been removed from the original graph $D(h, \ell, t)$. Let $t(n)$ denote the number of remaining NDDs plus the one being used in the current chain at the time when n nodes (either H-nodes or NDD donor nodes) have been removed. The stochastic process is $\{(X(n), t(n))|n \geq 0\}$. Each increment of "time" n denotes the removal of a node from the graph. When either the donor reneges or there are no compatible donors an NDD node is removed and $t(n)$ is decremented by one. We call this a "failure". When a compatible match is found and a patient gets a transplant then an H-node is removed and $X(n)$ is decremented by one. We call this a "success". By definition, $X(0) = h$, and $t(0) = t$.

This process of node removal eventually terminates. There are two conditions for termination. The first is that all NDD donors have been consumed, corresponding to $t(n) = 0$. The second is that all H-nodes have been transplanted, corresponding to $X(n) = 0$. Thus, we can define the random stopping time

$$\tau_0 = \min\{n | t(n) = 0 \text{ or } X(n) = 0\}. \tag{1}$$

when Phase 1 terminates.

Observe that $\{(X(n), t(n)) | n \geq 0\}$ is a two-dimensional pure death process with absorbing states $\{(X, t) | X = 0 \text{ or } t = 0\}$. At each non-absorbing state, the transition probability is given by

$$(X(n+1), t(n+1)) = \begin{cases} (X(n) - 1, t(n)) & \text{w.p. } 1 - r_{X(n)} \\ (X(n), t(n) - 1) & \text{w.p. } r_{X(n)}, \end{cases}$$

where

$$r_i = r + (1 - r)(1 - p_H)^i$$

gives the probability that either the tail node reneges, or the tail node cannot find an accessible H-node. From this definition we see that $\{(X(n), t(n)) | n \geq 0\}$ is Markovian. Figure 2 provides a visual representation.

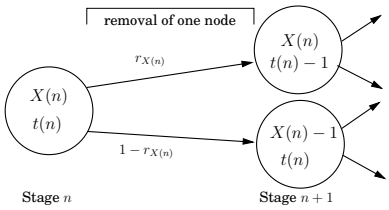

Fig. 2. The stochastic process $\{(X(n), t(n)) | n \geq 0\}$.

When the graph contains m H-nodes, the number of NDDs consumed reduce the number of unmatched H-nodes by one is a geometric random variables with success probability $1 - r_m$ and mean $\mu_m := \frac{r_m}{1 - r_m}$. For the ease of the subsequent analysis, we define the following potential function:

$$T(n) = \sum_{m=n+1}^{h} \mu_m. \tag{2}$$

The function T calculates the expected number of NDDs needed to reduce the number of H-nodes from h to m. In the special case of $r = 0$, $\mu_m = \frac{(1-p_H)^m}{1-(1-p_H)^m}$ and $\sum_{m=1}^{\infty} \mu_m < \infty$ and hence the revised potential function

$$T^0(n) := \sum_{m=n+1}^{\infty} \mu_m. \tag{3}$$

is well defined.

Observe that T is a strictly decreasing function on the discrete domain $0, 1, \ldots, h$. We extend T to be defined over the continuous domain $[0, h]$ via piecewise linear interpolation. This makes the inverse function T^{-1} well-defined on the range $[0, T(0)]$ of T where $T(0) < \infty$. For $x \geq T(0)$, we take $T^{-1}(x) = 0$, which will not modify the monotonicity of $T^{-1}(\cdot)$. Under this extension, both T and T^{-1} are convex functions because T has

increasing differences: $T(m) - T(m-1) = -\mu_m$ and μ_m is decreasing in m since r_m is decreasing in m. See Figure 3 for a visualization.

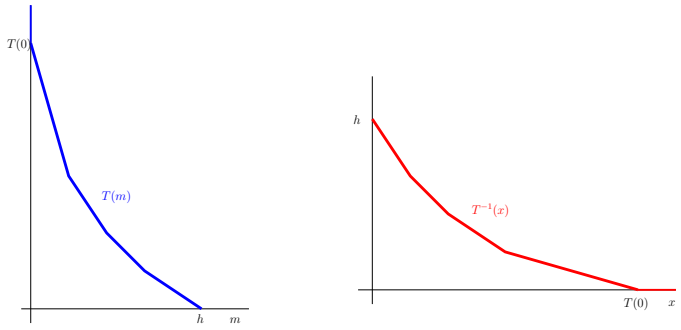

Fig. 3. The functions T and T^{-1}.

We seek distributional information on the random number $X(\tau_0)$ of unmatched H-node patients termination of Phase 1. To make the dependence on h and t explicit, let $Y_{h,t}$ denote the value of $X(\tau_0)$ when the initial graph is $D(h,\ell,t)$. We are interested in the following performance metrics: (a) the tail probability $\Pr(Y_{h,t} \le k)$ for a given non-negative integer k and (b) the expectation of $Y_{h,t}$.

THEOREM 3.1. *For a random kidney exchange graph* $D(h,\ell,t)$

(a) $\Pr(Y_{h,1} \le k) = \prod_{i=k+1}^{h}(1-r_i)$ *when* $t = 1$. *When* $t \ge 2$,

$$\Pr(Y_{h,t} \le k) = \prod_{i=k+1}^{h}(1-r_i) \sum_{k \le i_{t-1} \le \ldots \le i_1 \le h} \prod_{j=1}^{t-1} \xi_k(i_j) \tag{4}$$

where $\xi_k(i) = r_i$ *for* $i > k$, *and* $\xi(i) = 1$ *if* $i \le k$.

(b) Moreover,

$$\mathbb{E}(Y_{h,t}) = \sum_{k=0}^{h} \left(1 - \prod_{i=k+1}^{h}(1-r_i) \sum_{k \le i_{t-1} \le \ldots \le i_1 \le h} \prod_{j=1}^{t-1} \xi_k(i_j) \right).$$

PROOF. Note that $Y_{h,1}$ represents the number of unmatched H-nodes after matching with a single chain. The only way for $Y_{h,1} \le k$ is for there to be a string of consecutive "successes" in extending the chain to reduce the number of unmatched H-nodes from h to k. Again relying on independence, this happens with probability $\prod_{i=k+1}^{h}(1-r_i)$.

The event $\{Y_{h,t} \le k\}$ contains all scenarios where there are no more than t failures in the course of removing $h-k$ H-nodes. Without loss of generality, we assume there are $t' < t$ failures before $X(n)$ hits k. We let i_j denote the number of H-nodes remaining in the graph at the time of the j-th failure. For $j = 1, 2, \ldots, t'$, the failure rate $\xi_k(i_j)$ at i_j is r_{i_j}; whereas for $j = t'+1, \ldots, t$, we assign $\xi_k(i_j) = 1$ as these failures do not contribute to $Y_{h,t} \le k$. Thus, we derive the tail probability for $Y_{h,t}$ as

$$\Pr(Y_{h,t} \le k) = \prod_{i=k+1}^{h}(1-r_i) \sum_{k \le i_{t-1} \le \ldots \le i_1 \le h} \prod_{j=1}^{t-1} \xi_k(i_j).$$

The expression for $\mathbb{E}(Y_{h,t})$ directly follows from the equation $\mathbb{E}[X] = \sum_{k=0}^{\infty} \Pr(X > k)$ for nonnegative discrete random variables. □

The above combinatorial expressions for tail probabilities and expectation of $Y_{h,t}$ are precise but difficult to work with. The next result provides bounds that involve the potential function T and are more amenable to later analysis.

THEOREM 3.2. $\Pr(Y_{h,1} \leq k) \geq e^{T(h)-T(k)}$. *For $t \geq 2$,*

$$P(Y_{h,t} \leq k) \geq \exp\left(T(h) - T(k)\right) \frac{(1 + \sum_{i=k+1}^{H} r_i)^{t-1}}{(t-1)!}. \tag{5}$$

In the special case of $r = 0$,

$$P(Y_{h,t} \leq k) \geq \exp\left(-\frac{(1-p_H)^k}{p_H}\right) \frac{(1 + \sum_{i=k+1}^{H} r_i)^{t-1}}{(t-1)!}. \tag{6}$$

In the special case of $r = 0$ and $t = 1$, the above bound can be further strengthened to

$$P(Y_{h,1} \leq k) \geq 1 - \frac{1}{p_H}\left((1-p_H)^{k+1} - (1-p_H)^{h+1}\right). \tag{7}$$

PROOF. Let $a(x) = \frac{-\ln(1-x)}{x} = \sum_{j=1}^{+\infty} \frac{x^{j-1}}{j}$. Since a is an increasing function of x, $a(r_i)$ is a decreasing series in i. Notice that $\frac{1}{1+y} \geq e^{-y}$ for all $y \geq 0$, by taking $y = \frac{x}{1-x}$ we have $1 - x \geq e^{-\frac{x}{1-x}}$ for all $1 \geq x \geq 0$. Hence,

$$\begin{aligned} P(Y_{h,1} \leq k) = & \prod_{i=k+1}^{h}(1 - r_i) \\ \geq & \prod_{i=k+1}^{h} e^{-\mu_i} \\ = & e^{T(h)-T(k)}. \end{aligned}$$

For $t \geq 2$, note that the expansion of $(1 + \sum_{i=k+1}^{h} r_i)^{t-1}$ contains terms of the form $\prod_{j=1}^{t} \xi_k(i_j)$. The coefficient for each term is $(t'-1)!$, where t' is the number of $\xi_k(i_j)$ with $i_j > k$. Relaxing t' to t yields (5). When $r = 0$, we have

$$\begin{aligned} P(Y_{h,1} \leq k) &\geq \exp\left(-\sum_{i=k+1}^{h} \mu_i\right) \\ &= \exp\left(-\sum_{i=k+1}^{h} \frac{(1-p_H)^i}{1-(1-p_H)^i}\right) \\ &\geq \exp\left(-\sum_{i=k+1}^{h} \frac{(1-p_H)^i}{p_H}\right) \\ &\geq \exp\left(-\frac{(1-p_H)^k}{p_H}\right). \end{aligned}$$

We prove the strengthened bound (7) by backwards induction on k. When $k = h$, $\Pr(Y_{h,1} \leq h) = 1 = 1 - \frac{1}{p_H}\left((1-p_H)^{h+1} - (1-p_H)^{h+1}\right)$. So (7) holds in the base case of $k = h$. We now assume that (7) holds for $k+1$ and show it holds for k. We derive an

upper bound for $\Pr(Y_{h,1} \leq k)$ as follows:

$$\begin{aligned}
\Pr(Y_{h,1} \leq k) &= \Pr(Y_{h,1} \leq k) \Pr(Y_{h,1} \leq k | Y_{h,1} \leq k+1) \\
&= \left(1 - \frac{1}{p_H}((1-p_H)^{k+2} - (1-p_H)^{h+1})\right)\left(1 - (1-p_H)^{k+1}\right) \\
&\geq 1 - (1-p_H)^{k+1} - \frac{1}{p_H}(1-p_H)^{k+2} + \frac{1}{p_H}(1-p_H)^{h+1} \\
&= 1 - \frac{1}{p_H}(1-p_H)^{k+1} + \frac{1}{p_H}(1-p_H)^{h+1} \\
&= 1 - \frac{1}{p_H}\left((1-p_H)^{k+1} - (1-p_H)^{h+1}\right)
\end{aligned}$$

where the first equality follows from the random walk procedure, the second equality follows from the induction assumption on k, and the inequality holds by omitting the term $\frac{1}{p_H}((1-p_H)^{k+2} - (1-p_H)^{h+1})(1-p_H)^{k+1}$. This completes the induction. □

The above bounds are applicable to exchange graphs of arbitrary size. In particular, they can be leveraged in asymptotic settings to derive results similar to those in Ashlagi et al. [2012, 2013]. The following result demonstrates the approach.

PROPOSITION 3.3. *Suppose $r = 0$, h is in the order of $\frac{1}{p_H^{1+\epsilon}}$ for some $\epsilon > 0$, and $\lambda = \frac{\ell}{h+\ell}$ and p_L are both fixed positive constants. If $t \geq 1$, then with probability approaching one the exchange graph has a perfect clearing (that is, all nodes are transplanted) as $p_H \to 0$.*

PROOF. In (6) suppose $k = \frac{c}{p_H}\log(\frac{1}{p_H})$ for some constant $c > 1$. This yields $\Pr(Y_{h,1} \leq k) \geq \exp(-p_H^{c-1})$, which converges to 1 as $p_h \to 0$. Thus with probability approaching one, no matter how large the original graph is, only one NDD is sufficient to reduce the number of H-nodes to the order of $O\left(\frac{1}{p_H}\log(\frac{1}{p_H})\right)$. Since $t \geq 1$, with probability approaching one the number of H-nodes remaining unmatched after Phase 1, denoted by h', is in the order of $O(\frac{1}{p_H}\log(\frac{1}{p_H}))$. We claim that when $p_H \to 0$, with probability approaching one all of those $h'H$-nodes can be matched to L-nodes using two-way cycles.

To prove this claim, we construct an undirected bipartite graph $\tilde{G} = (V_H \cup V_L, \tilde{E})$ with partitioned node sets $V_H := \{\text{all remaining } h' \text{ unmatched } H\text{-nodes}\}$ and $V_L := \{\text{all } L\text{-nodes}\}$, and an undirected edge set $\tilde{E} = \{(v_H, v_L) \mid v_H \in V_H,\ v_L \in V_L,\ (v_H, v_L), (v_L, v_H) \in E\}$. Each edge occurs with probability of $p_H p_L$ – the probability to have a 2-way cycle between an H-node and an L-node. According to the Marriage Theorem [Hall 1935], if the H-nodes cannot be matched in $\tilde{G}$, then there exists a "bad" pair of subsets $A \subset V_H$ and $B \subset V_L$ with $a = |A| > b = |B|$ and the set B contains all vertices adjacent to vertices in A. Without loss of generality, we may assume that (A, B) is a minimal bad pair, which means that there is no bad pair (A', B') with $A' \cup B' \subset A \cup B$. When (A, B) is a minimal bad pair, we must have $b = a - 1$. The probability that any nodes outside B is not linked to any node inside A is given by $(1 - p_H p_L)^{(\ell-b)a}$. Since there at at most $C_{h'}^a C_\ell^b$ candidates for a minimal bad pair (A, B) of size a and b respectively, the probability that at least one minimal bad pair exists of

this size can be upper bounded by

$$\begin{aligned}\sum_{a=1}^{h'} C_{h'}^{a} C_{\ell}^{a-1}(1-p_H p_L)^{(\ell-a+1)a} &\le \sum_{a=1}^{h'} \frac{(h'\ell)^a}{(a!)^2}(1-p_H p_L)^{(\ell-a+1)a} \\ &\le \sum_{a=1}^{h'} \frac{1}{a!}\left[\ell h'(1-p_H p_L)^{\ell-a+1}\right]^a \\ &\le \exp\left(\ell h'(1-p_H p_L)^{\ell-h'+1}\right) - 1 \\ &\le \exp\left(\ell h' \exp(-p_H p_L(l-h'+1))\right) - 1 \\ &\le \ell h' \exp\left(-p_H p_L(\ell-h'+1)\right). \end{aligned} \tag{8}$$

Note that $\ell = O(\frac{1}{p_H^{1+\epsilon}}) >> h'$, thus $\ell - h' + 1 = O(\frac{1}{p_H^{1+\epsilon}})$ and $p_H p_L(\ell - h' + 1) = O(\frac{1}{p_H^{\epsilon}}) \ge 2\log(\frac{1}{p_H})$ when p_H is sufficiently small. Therefore, (8) is upper bounded by $\ell h' \exp\left(-2\log(\frac{1}{p_H})\right) = O(p_H^{1-\epsilon})$, implying that the probability of the occurrence of a bad pair converges to zero when $p_H \to 0$. Therefore, with probability approaching one, all the H-nodes can be matched by using H-L two-way cycles. After removal of all the H-nodes, the remaining subgraph contains L-nodes only. Each pair of L-nodes can be matched with a constant probability of p_L^2. Because the size of the remaining graph is $\ell - h' = O(\frac{1}{p_H^{1+\epsilon}}) \to \infty$ as $p_H \to \infty$, we know that a perfect matching exists by the well-known Erdő-Rényi theorem [Erdős and Rényi 1959]. □

Remark 3.4. When h is order $\frac{1}{p_H^{1+\epsilon}}$, Theorem 5.6(2) of Ashlagi et al. [2012] prove that all nodes can be matched using k-way cycles if $\lambda > \frac{1}{k}$, or using chains with length $\le m$ if $t \ge \frac{1-\lambda}{m}h$. Proposition 3.3 states that if we have a single chain of potentially infinite size (which is a weaker assumption than $\frac{1-\lambda}{m}h$ chains each with length $\le m$ with respect to the purpose of matching H-nodes) and a fixed proportion of L-nodes, we can transplant all nodes. This complements the result of [Ashlagi et al. 2012].

We now turn to deriving non-combinatorial lower and upper bounds on $\mathbb{E}[Y_{h,t}]$ for later analysis. Of course, one could combine Theorems 3.1 and 3.2 to achieve this, but a different method will yield cleaner bounds. The proof of the following result (in particular part (c)) is involved and can be found in Appendix A. It uses martingale theory and convexity properties of the potential function T and its inverse T^{-1}.

THEOREM 3.5. *The following conditions hold:*

(a) $\{T(X(n)) + t(n) | n \ge 0\}$ *is a martingale. As a consequence,*

$$\mathbb{E}[Y_{h,t}] \ge T^{-1}(T(h) + t). \tag{9}$$

(b) $\left\{\frac{X(n)+T^{-1}(T(X(n)+t(n)))}{2} | n \ge 0\right\}$ *is a super-martingale. As a consequence,*

$$\mathbb{E}[Y_{h,t}] \le \frac{1}{2}\left(T^{-1}(T(h)+t) + h\right). \tag{10}$$

(c) In the case of $r = 0$ *and* $p_H \le 0.1$*, we have the following strengthened upper bound,*

$$\mathbb{E}[Y_{h,t}] \le \frac{1}{p_H}\log\left(1 + \frac{1}{(T^0(n) + \frac{1}{4}t)p_H}\right) \tag{11}$$

where T^0 *was defined in* (3)*.*

Figure 4 plots the actual values of $\mathbb{E}[Y(h,t)]$ versus the bounds in Theorem 3.5. The lower bound from (a) is quite sharp. The upper bound from (b), is not so tight, but nonetheless helps us to understand the asymptotic behavior of $\mathbb{E}[Y(h,t)]$. In particular, when $r = 0$, the strengthened upper bound (11) indicates that the number of unmatched nodes after Phase 1 is upper bounded by $\frac{1}{p_H}\log(\frac{C(t,h)}{p_H})$ as long as $t \geq 1$. Also, additional NDDs do not help match many more patients since increasing t only lowers the constant $C(t,h)$, while the order remains $\frac{1}{p_H}\log(\frac{1}{p_H})$.

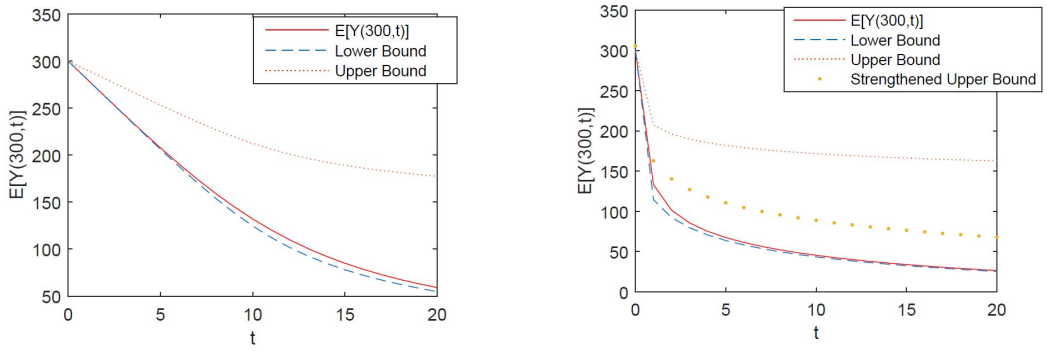

(a) $r = 0.05$, $p_H = 0.03$, $h = 300$ (b) $r = 0$, $p_H = 0.03$, $h = 300$

Fig. 4. $\mathbb{E}[Y(h,t)]$ and its lower bound (9), upper bound (10), and the strengthened upper bound (11) (applicable only to the $r = 0$ case).

4. ANALYSIS OF PHASE 2

The goal of this section is to provide analytical bounds on the number of unmatched H-nodes that are left after termination of both Phase 1 and Phase 2. This analysis proceeds by comparing against a benchmark, namely the number of unmatched H-nodes that remain if only Phase 2 was implemented from the beginning. In other words, we are interested in the *net* benefit of our procedure to reduce the number of unmatched H-nodes *beyond* what *could have* been transplanted via cycles alone. Our performance metric does not include the unmatched L-nodes as it is easier to clear all the L-nodes even using cycles only. Moreover, keeping L-nodes for a later clearing may even be preferred [Ashlagi et al. 2013]. As in the previous section, the underlying memorlyess property of the random walk procedure and arguments based on convexity will play a pivotal role here.

A first challenge is to understand the random structure of the residual graph $\mathcal{R}$ that remains at the termination of Phase 1. We show that this residual graph is again a graph of the form $D(h',\ell,t)$ where parameters ℓ, t, p_H, p_L and r fixed and $h' \leq h$. Throughout this section we suppose $h' > 0$ and all NDDs are consumed during Phase 1. The case where all H-nodes are matched before all NDDs are consumed is an uninteresting special case that is very unlikely to occur in practice and so does not warrant further analysis.

LEMMA 4.1. *Let $D = D(h,\ell,t)$ denote the initial random graph. Suppose during Phase 1, $h - h'$ H-nodes are transplanted and removed and let $\mathcal{R}(h')$ denote conditional residual graph; that is, the random graph resulting from D by removing exactly $h - h'$*

H-nodes in chains via Phase 1. Then the edge distribution in $\mathcal{R}(h')$ is identical to the random graph $D' = D(h', \ell, 0)$.

PROOF. Fix an ordering of the H-nodes in the starting graph D. There are $\binom{h-h'+t}{t}$ scenarios by which one starts with the graph D and ends with a residual graph that has h' remaining H-nodes and 0 NDD nodes. Indeed, we simply need to distribute t failures among $h - h' + t$ steps of removing either H-nodes or NDDs. Due to the memoryless property of random walk, each of these residual graphs is isomorphic as a random graph to D'. Now, consider the conditional random graph $\mathcal{R}(h')$. Since each of the scenarios is disjoint, we can conclude that the edge distribution in the remaining graph $\mathcal{R}(h')$ is identical to the random graph D'. □

Now, as a first step, we consider a simple instantiation of Phase 2. Construct an undirected bipartite graph with all the unmatched H-nodes on one side, and all L-nodes on the other side. An H-node is connected to an L-node if and only if two nodes form a two-way directed cycle in the original directed graph. The algorithm finds a maximal matching on this bipartite graph (for instance, using linear programming). This is admittedly not the optimal choice of algorithm for Phase 2, but it nonetheless forms a bedrock for analysis that we later extend.

Let $f(h)$ denote the expected number of H-nodes remaining when nodes are matched via Phase 2 from the outset (that is, Phase 1 is not implemented) on the original graph $D(h, \ell, t)$. In the definition of f, the number of L-nodes and the probabilities p_H, p_L and r are all fixed constants. Lemma 4.1 implies the following pivotal corollary.

COROLLARY 4.2. *The expected number of unmatched H-nodes after Phase 1 and Phase 2 is implemented on $D(h, \ell, t)$, given that Phase 1 eliminates $h - h'$ H-nodes, is $f(h')$.*

Thus, the value of implementing chains versus not implementing chains can be partially understood by comparing $\mathbb{E}[f(h')]$ to $f(h)$, where h' is now a random variable that represents the number of H-nodes remaining in the graph D after Phase 1 terminates; that is, $h' = Y_{h,t}$ where $Y_{h,t}$ is defined as in Section 3. The expectation in $\mathbb{E}[f(Y_{h,t})]$ is over the distribution of unmatched H-nodes after Phase 1. Whereas we do not have an explicit characterization of f (making a direct evaluation of $\mathbb{E}[f(Y_{h,t})]$ difficult), Section 3 does provide good estimates of $\mathbb{E}[Y_{h,t}]$. On order to leverage these estimates, we show (in Lemma 4.3 below) that the function f is convex. The proof of the lemma is in Appendix B. It involves the linear programming formulation of bipartite matching and submodularity arguments.

LEMMA 4.3. *The expected number $f(h)$ of unmatched H-nodes remaining after running the bipartite matching algorithm described above to $D(h, \ell, t)$ is convex in h.*

We leverage the convexity of f to bound the performance of the random walk procedure when Phase 2 implements bipartite matching.

THEOREM 4.4. *Let $\nu_C^*(h, t) := \mathbb{E}[f(Y_{h,t})]$ denote the expected number of unmatched H-nodes after completion of both Phase 1 and the bipartite matching algorithm in Phase 2 on $D(h, \ell, t)$ (when there is no ambiguity as to values of h and t we simply write $\nu_C^*(h, t) = \nu_C^*$). The following holds:*

$$\begin{aligned}\nu_C^* &\leq \frac{f(h)}{h}\mathbb{E}[Y_{h,t}] \\ &\leq \frac{1}{h}\left((h-\ell)^+ + \frac{(1-p_Lp_H)^{|h-\ell|+1} - (1-p_Lp_H)^{\max\{h,\ell\}+1}}{p_Lp_H}\right)\mathbb{E}[Y_{h,t}]. \qquad (12)\end{aligned}$$

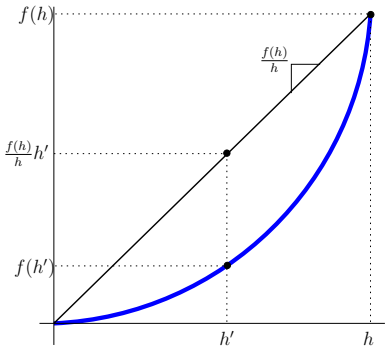

Fig. 5. Leveraging the convexity of f in the proof of Theorem 4.4.

PROOF. By Lemma 4.3, we know $f(h') \leq \frac{f(h)}{h}h'$ for any fixed $h' \leq h$ since f is convex. See Figure 5 and observe that $f(0) = 0$. Then by the monotonicity of expectation we have

$$\mathbb{E}[f(Y_{h,t})] \leq \mathbb{E}\left[\frac{f(h)}{h}Y_{h,t}\right] = \frac{f(h)}{h}\mathbb{E}[Y_{h,t}]. \tag{13}$$

This is the first inequality in (12).

The second inequality comes from bounding $f(h)$ from above. We derive the bound by analyzing the following naïve algorithm to match H-nodes. When $h > \ell$, sequence the L-nodes in an arbitrary order. For each L-node, attempt to match to an H-node using available edges (which correspond to 2-way cycles in the directed graph). If successful, remove the matched pair and proceed to the next L-node. The algorithm terminates when all the L-nodes have been matched. According to this algorithm, there are at least $h - i + 1$ unmatched H-nodes when the i-th L-node is next to be matched. Thus, the probability of matching the i-th L-node is at least $1 - (1 - p_H p_L)^{h-i+1}$. Summing up these probabilities, we derive a lower bound for the expected number of pairs

$$\sum_{i=1}^{\ell}\left(1 - (1 - p_H p_L)^{h-i+1}\right) = \ell - \frac{(1 - p_H p_L)^{h-\ell+1} - (1 - p_H p_L)^{h+1}}{p_H p_L}.$$

This leads to the upper bound on unmatched H-nodes

$$f(h) \leq h - \ell + \frac{(1 - p_H p_L)^{h-\ell+1} - (1 - p_H p_L)^{h+1}}{p_H p_L}. \tag{14}$$

When $h < \ell$, we propose a similar algorithm. This time we use two-way cycles to match the H-nodes sequentially. When the algorithm tries to match the i-th H-node, there are at least $\ell - i + 1$ unmatched L-nodes remaining, so the probability of matching the i-th H-node is at least $1 - (1 - p_H p_L)^{\ell-i+1}$. Using a similar logic as above, we derive the upper bound

$$f(h) \leq \frac{(1 - p_H p_L)^{\ell-h+1} - (1 - p_H p_L)^{\ell+1}}{p_H p_L}, \tag{15}$$

Together, (14) and (15) imply the following upper bound on $f(h)$ that applies to both cases ($h > \ell$ and $h < \ell$):

$$f(h) \;\leq\; (h - \ell)^+ + \frac{(1 - p_L p_H)^{|h-\ell|+1} - (1 - p_L p_H)^{\max\{h,\ell\}+1}}{p_L p_H}.$$

Plugging this upper bound into (13) yields the second inequality in (12). □

Of course, one may wonder how Phase 1 and Phase 2 as currently specified compare in performance to more sophisticated clearing algorithms. The next result shows how to bound the performance of an optimal 2-way cycle packing algorithm applied to the original graph in comparison to the performance using the random walk procedure with bipartite matching in Phase 2.

THEOREM 4.5. *Let ν_2^* denote the expected number of unmatched nodes when applying an optimal 2-way cycle packing algorithm to the original random graph $D(h,\ell,t)$. Recall, $\lambda = \frac{\ell}{h+\ell}$, then*

$$\frac{\nu_C^*}{\nu_2^*} \leq \frac{\mathbb{E}[Y_{h,t}]}{h} \frac{1-2\lambda}{1-2\lambda-(1-\lambda)^2 p_H^2(h+\ell)} \tag{16}$$

where v_C^* is defined as in the statement of Theorem 4.4 and assuming the denominator in the right-hand-side is positive.

PROOF. Suppose when running the optimal matching algorithm using 2-way cycles, the number of H-nodes that are matched by H-H cycles and H-L cycles are n_1 and n_2 respectively. The expected number of H-nodes matched in H-H-cycles is upper bounded by the expected total number of H-H cycles in the H-subgraph. Hence, $\mathbb{E}[n_1] \leq h(h-1)p_H^2 \leq p_H^2 h^2$. Clearly, $\mathbb{E}[n_2] \leq h - f(h)$. Therefore, $\mathbb{E}[n_1+n_2] \leq h - f(h) + p_H^2 h^2$. Thus, $\nu_2^* \geq f(h) - p_H^2 h^2$. By the definition of f(h), since there are at most $\frac{\lambda}{1-\lambda}h$ L-nodes, $f(h) \geq (1-\frac{\lambda}{1-\lambda})h$. This implies

$$\begin{aligned}
\frac{\nu_C^*}{\nu_2^*} &\leq \frac{\nu_C^*}{f(h)} \frac{f(h)}{\nu_2^*} \\
&\leq \frac{\mathbb{E}[Y_{h,t}]}{h} \frac{f(h)}{f(h)-(p_H^2 h)h} \\
&\leq \frac{\mathbb{E}[Y_{h,t}]}{h} \frac{(1-\frac{\lambda}{1-\lambda})h}{(1-\frac{\lambda}{1-\lambda}-p_H^2 h)h} \\
&= \frac{\mathbb{E}[Y_{h,t}]}{h} \frac{1-2\lambda}{1-2\lambda-(1-\lambda)^2 p_H^2(h+\ell)}
\end{aligned}$$

where the second inequality uses Theorem 4.4. □

Using bounds from Theorem 3.5 for $\mathbb{E}[Y_{h,t}]$ allows us to derive insights from Theorem 4.5 into the value of chains in reducing the number of unmatched H-nodes as compared to using cycles only. The ratio $\frac{v_C^*}{v_2^*}$ measures performance using chains as compared to using 2-cycles from the outset. The smaller is this ratio, the greater the marginal value of using chains. The bound (16) gives a guarantee on the expected marginal value of chains. Theorem 3.5 reveals that $\frac{\mathbb{E}Y_{h,t}}{h}$ even when p_H is small. In the case of $r=0$, even with one NDD, $\mathbb{E}Y_{h,t}$ can be upper bounded by $\frac{1}{p_H}\log(\frac{c}{p_H})$ no matter how large of h (see (11)). Therefore, the ratio $\frac{\nu_C^*}{\nu_2^*}$ is small as long as $1-2\lambda-(1-\lambda)p_H^2 h$ is not close to zero. Thus when both the proportion of L-nodes λ and the graph size $h+\ell$ are of small to moderate size, the benefits of using NDDs (chains) is most substantial.

4.1. Connections to integer programming formulations

In this subsection we further underscore the possibility that our bounds in Theorem 4.4 can be used to bound the performance of more sophisticated implementations of the

Phase 2 cycle packing algorithm. One example is the following integer linear programming (ILP) method for the clearing problem, formulated as:

$$\begin{aligned} \mu^*_{ILP} := \max & \sum_{c\in\mathcal{C}(M)} w_c x_c \\ \text{s.t.} & \sum_{\{c\in\mathcal{C}(M),\ v\in c\}} x_c \le 1,\ \forall \text{ vertices } v \\ & x_c \in \{0,1\} \end{aligned} \tag{17}$$

where $w(c)$ denotes the weight of cycle c, $\mathcal{C}(M)$ denotes a set of cycles with sizes no more than M and chains of arbitrary sizes. The ILP (17) was first proposed by Abraham et al. [2007] and improved upon in subsequent studies [Dickerson et al. 2012b, 2013]. This ILP serves as the basis of the allocation scheme currently used by UNOS in clearing its exchange.

One challenge of using this method is that the number of chains in $\mathcal{C}(M)$ is exponentially increasing with the graph size, so the ILP fails for graph of sizes > 200 [Dickerson et al. 2012b]. For this reason, the current UNOS solver uses a column generation method to strategically add potentially valuable chains and cycles in to the collection $\mathcal{C}(M)$ [Dickerson et al. 2013]. If Phase 1 of our procedure is used to generate chains in $\mathcal{C}(M)$, then whatever allocation obtained after Phase 1 and Phase 2 is a feasible solution to the ILP (17). Therefore, if the objective of the ILP is to minimize the unmatched number of H-nodes, then as long as $\mathcal{C}(M)$ contains the t random walks generated in Phase 1, it must do at least as well as our two-phase algorithm. Thus, Theorem 4.4 provides an analytical lower bound for optimal value of the ILP.

COROLLARY 4.6. *Let w_c denote the number of H-nodes contained in the cycle or chain c, and assume $\mathcal{C}(M)$ contains the t random walks that are generated during Phase 1. Then when cycles of at least length two are permitted in* (17) *(that is, $M \ge 2$) we have*

$$\begin{aligned} \mu^*_{ILP} &\ge H - \nu^*_C(T) \\ &\ge \frac{1}{hp_H}\left((h-\ell)^+ + \frac{(1-p_Lp_H)^{|h-\ell|+1} - (1-p_Lp_H)^{\max\{h,\ell\}+1}}{p_Lp_H}\right)\log(1+\frac{1}{(T^0(n)+\frac{1}{4}t)p_H}). \end{aligned}$$

Note that we have used the strengthened upper bound (11) for $\mathbb{E}[Y_{h,t}]$ since this is the $r = 0$ case. Maximizing the total number of H nodes is a reasonable strategy as the L-nodes are much easier to be matched.

5. NUMERICAL EXPERIMENTS

The random walk procedure presented in Section 2.2 is simple by design. It handles the issue of chains in as simple a way as possible (via random assignment) so that the probabilistic implications of allocating chains is minimal. The primary purpose of this was to aid in analysis. This section serves as a reality check. Numerical tests will give us some confidence in the strength of our analysis and how our procedure compares to other procedures implemented in the literature and in practice.

This numerical section has two related components. The first is to numerically test the strength of the bounds derived in Section 4. The second purpose of this numerical section is to provide some sense of the practical cost of using the random walk procedure.

Due to the specification of the random walk procedure and to fit our purposes, we conduct our numerical tests on simulated random graphs based on the model in Section 2.1 instead of the more complicated setting of a fully simulated kidney-exchange

Table I. Sharpness of bounds in Section 4.

No. nodes	No. NDDs	ν_C^* via simulation	Bound (12)	$\frac{\nu_C^*}{\nu_2^*}$ via simulation	Bound (16)
40	0	26.67	27.08		
40	1	25.53	25.79	0.88	0.95
40	5	20.88	21.94	0.61	0.69
60	0	37.09	37.76		
60	1	34.82	35.23	0.85	0.93
60	5	27.50	29.19	0.49	0.63
80	0	47.91	48.83		
80	1	43.20	45.68	0.75	0.94
80	5	33.41	36.66	0.42	0.60
100	0	56.75	58.01		
100	1	50.51	54.10	0.69	0.95
100	5	36.92	42.08	0.35	0.57

or using data based on fielded exchanges. A more thorough numerical investigation undertaken with more realistic data is left for future work.

In our first numerical experiment we compare the actual performance of the two-phase procedure with bipartite matching in Phase 2 on simulated random graphs with the theoretical upper bounds for ν_C^* derived in Theorem 4.4. We test random graphs with 40, 60, 80, and 100 nodes. In all those graphs, we set parameter values of $p_H = 0.05$, $p_L = 0.45$ and $\lambda = 0.27$. We tested three values for the number of NDD donors, $t = 0$, 1, and 5. This gives a total of 12 different random graph specifications.

For each random graph specification, we generate $1,000$ graphs and apply Phases 1 and 2 on the realized graphs and record the number of unmatched H-nodes at termination. This yields a total of $12,000$ observations. We average the number of unmatched H-nodes across all $1,000$ generated graphs in each specification and record those averages in the third column on Table I. We calculate the upper bounds using (12) and (10) and record these values in the fourth column of Table I. The simulated values are all within 10% of the upper bounds. If we compare across the different values for t, we find a significant reduction in the number of unmatched H-nodes when chains are allowed ($t > 0$). This demonstrates the marginal contribution by allowing chains in the matching. We repeated the experiment for different parameters values for p_H, p_L, λ and reach similar conclusions. The fifth and sixth columns of Table I provide a comparison of simulated values for the ratio $\frac{\nu_C^*}{\nu_2^*}$ and our bound in Theorem 4.5. We tried two different heuristics to proxy v_2^* and chose the minimal result. We do not conduct the tests when there were no NDDs, since we know that v_2^* is clearly superior to v_C^* when there are no cycles. The theoretical bound is reasonably close to simulated values, confirming the strength of our bounds.

In the second set of experiments we compare the performance of the random walk procedure to heuristic clearing algorithms that are commonly used. Our specific goal is to compare to the results given in Table 5 of [Ashlagi et al. 2012], which records the number of matched H-nodes after running their heuristic algorithm, and so we choose the same parameter settings. Their matching algorithm uses three-way cycles, and so to make the comparison fair, we also allow three-way cycle matching in the Phase 2 of the algorithm. The comparison is displayed in Table II, where the third column records the average number of unmatched H-nodes using Phase 1 and Phase 2 of our procedure across 1000 randomly generated graphs in each setting, and the right-most columns records the number of unmatched H-nodes inferred by Table 5 of [Ashlagi et al. 2012]. These results are encouraging. Although performance degrades as NDDs are added, the performance is comparable to that of [Ashlagi et al. 2012]. Of course, our goal is not to advocate that the random walk procedure be used in practice, instead our motivation is to show that a very simple allocation of chains in the

Table II. Performance of random walk procedure compared to heuristic in [Ashlagi et al. 2012].

No. nodes	No. NDDs	Random walk procedure	[Ashlagi et al. 2012] Table 5
40	0	23.46	23.81
40	1	22.20	22.80
40	5	18.52	19.97
60	0	28.68	30.79
60	1	26.47	27.21
60	5	22.91	21.98
80	0	33.54	36.59
80	1	31.17	29.68
80	5	24.17	21.87
100	0	36.11	40.67
100	1	31.70	29.91
100	5	24.46	19.13

first stage remains comparable to more elaborate methods. In personal communication (2014), John Dickerson confirmed that allocating cycles to NDD donors via random walk performs surprisingly well as compared the optimal ILP algorithms (within 5% of optimal). This illustrates that the analytical gain to using random walk does not overly compromise the practical impact of no longer choosing chain optimally. Further investigation into the use of random walks in practical implementations is left for future work.

6. DIRECTIONS FOR FUTURE WORK

In this paper we have developed a non-asymptotic approach to analyzing kidney exchange graphs that complements previous work that relies on asymptotic analysis. We demonstrate the power of this approach by providing analytical performance bounds on a random walk procedure for matching donors and recipients, and demonstrate how these bounds allow us to analytically show the benefit of chains in "medium-sized" (that is, non-limit) graphs.

We developed this approach in the stylized setting introduced by [Ashlagi et al. 2012] and one additional restrictions in our procedure to facilitate analysis. The chains our procedure produces consist entirely of H-nodes (initiated, of course, by an NDD donor). We did this to maintain the stochastic independence structure of residual graphs that was leveraged at several points in our proofs. Nonetheless, extending to "mixed chains" of both H- and L-nodes is approachable by adjusting our methodology and is a topic of further investigation.

Of course, there are several important assumptions inherent in the random graph model of [Ashlagi et al. 2012] itself. Discussion of the validity of these assumptions and the possibility of extension is well-documented (see, for instance [Ashlagi et al. 2012, 2013; Dickerson et al. 2012b, 2013]). A test bed for the power of our non-asymptotic approach is to systematically attempt to relax certain assumptions and see what tractability remains. Since our approach has important distinctions with the standard asymptotic methods, it is conceivable that we can achieve further generality in ways that are not amenable to other methods. Some promising avenues include embedding the random walk procedure in a dynamic setting of kidney exchange within an evolving patient and donor base. We feel the memoryless properties of random walks should prove useful in a dynamic setting.

Finally, although we introduced our procedure primarily for analytical investigation, we feel it is useful to comment on the practicality of the approach. Allocating organs on the basis of random walks may strike practitioners and patients alike as somewhat arbitrary and unfair. However, there is some flexibility in our procedure that could make it more palatable in practice. One can specify a *priority* by which to process the

nodes. This priority could increase the likelihood that a given patient (potentially a very deserving one) could get a kidney sooner than others. One way of doing this is to evaluate the potential of each vertex on the kidney exchange graph, as proposed in [Dickerson and Sandholm 2015]. Of course, a benefit of the random walk approach is its scalability, as opposed to complicated optimization algorithms proposed by other researchers. Investigation into how one might adapt our procedure in practice (or even decide if it has any practical merit) is also a potential avenue of future research.

Acknowledgements

We thank Tian Xie for his able technical assistance and John Dickerson for enlightening conversations. Part of this research was conducted at the Center for Management Science and Information Analytics at the Shanghai University of Finance and Economics. We thank the Center for its generosity in facilitating our collaboration. The research of the first author is partially supported by NSERC PGPIN 436156-13. The fourth author also thanks the Booth School of Business for its generous support.

REFERENCES

ABRAHAM, D., BLUM, A., AND SANDHOLM, T. 2007. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM conference on Electronic commerce*. 295–304.

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., GILCHRIST, D., ROTH, A., AND REES, M. 2011. Nonsimultaneous chains and dominos in kidney-paired donation–revisited. *American Journal of Transplantation 11,* 5, 984–994.

ASHLAGI, I., JAILLET, P., AND MANSHADI, V. 2013. Kidney exchange in dynamic sparse heterogenous pools. *arXiv preprint arXiv:1301.3509*.

DICKERSON, J., PROCACCIA, A., AND SANDHOLM, T. 2012a. Dynamic matching via weighted myopia with application to kidney exchange. In *Proceedings of the 2012 International Joint Conferences on Artificial Intelligence*.

DICKERSON, J., PROCACCIA, A., AND SANDHOLM, T. 2012b. Optimizing kidney exchange with transplant chains: Theory and reality. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. 711–718.

DICKERSON, J., PROCACCIA, A., AND SANDHOLM, T. 2013. Failure-aware kidney exchange. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. 323–340.

DICKERSON, J., PROCACCIA, A., AND SANDHOLM, T. 2014. Price of fairness in kidney exchange. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. 1013–1020.

DICKERSON, J. P. AND SANDHOLM, T. 2015. Futurematch: Combining human value judgments and machine learning to match in dynamic environments. In *AAAI Conference on Artificial Intelligence (AAAI)*.

ERDŐS, P. AND RÉNYI, A. 1959. On random graphs. *Publicationes Mathematicae Debrecen 6*, 290–297.

GENTRY, S., MONTGOMERY, R., SWIHART, B., AND SEGEV, D. 2009. The roles of dominos and nonsimultaneous chains in kidney paired donation. *American Journal of Transplantation 9,* 6, 1330–1336.

HALL, P. 1935. On representatives of subsets. *J. London Math. Soc 10,* 1, 26–30.

MELCHER, M., LEESER, D., GRITSCH, H., MILNER, J., KAPUR, S., BUSQUE, S., ROBERTS, J., KATZNELSON, S., BRY, W., YANG, H., LU, A., MULGAONKAR, S.,

DANOVITCH, G., HIL, Q., AND VEALE, J. 2012. Chain transplantation: initial experience of a large multicenter program. *American Journal of Transplantation 12,* 9, 2429–2436.

ROTH, A., SÖNMEZ, T., AND ÜNVER, M. 2007. Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *The American economic review*, 828–851.

TOPKIS, D. M. 1998. *Supermodularity and complementarity*. Princeton university press.

TOULIS, P. AND PARKES, D. 2011. A random graph model of kidney exchanges: Efficiency, individual-rationality and incentives. In *Proceedings of the 12th ACM conference on Electronic commerce*. 323–332.

ÜNVER, M. 2010. Dynamic kidney exchange. *The Review of Economic Studies 77,* 1, 372–414.

WOODLE, E., DALLER, J., AEDER, M., SHAPIRO, R., SANDHOLM, T., CASINGAL, V., GOLDFARB, D., LEWIS, R., GOEBEL, J., AND SIEGLER, M. 2010. Ethical considerations for participation of nondirected living donors in kidney exchange programs. *American Journal of Transplantation 10,* 6, 1460–1467.

A. PROOF OF THEOREM 3.5

PROOF. (a) At a given non-absorbing state $(X(n), t(n))$, compute $T(X(n+1))+t(n+1)$ by the transition probabilities as

$$\begin{aligned}\mathbb{E}[T(X(n+1))+t(n+1)] &= r_{X(n)}(T(X(n))+t(n)-1)+(1-r_{X(n)})(T(X(n)-1)+t(n)) \\ &= r_{X(n)}(T(X(n))+t(n)-1)+(1-r_{X(n)})\left(T(X(n))+\frac{r_{X(n)}}{1-r_{X(n)}}+t(n)\right) \\ &= T(X(n))+t(n).\end{aligned}$$

This invariant tells us along any random sample path, the expectation of $\mathbb{E}[T(X(n))+t(n)]$ is $T(h)+t$, its value in the initial state. Recall that τ_0 corresponds to the time at which either $X(n)$ hits zero or $t(n)$ hits zero, and so the optional stopping theorem for martingales

$$\begin{aligned}T(h)+t &= \Pr(t(\tau_0)=0)\mathbb{E}[T(Y_{h,t})+0]+(1-\Pr(t_{\tau_0}=0))\mathbb{E}[T(Y_{h,t})+t(\tau_0)] \\ &= \mathbb{E}[T(Y_{h,t})]+(1-\Pr(t_{\tau_0}=0))t(\tau_0) \\ &\geq \mathbb{E}[T(Y_{h,t})]\end{aligned}$$

which implies that $\mathbb{E}[T(Y_{h,t})] \leq T(h)+t$. By the convexity of $T(\cdot)$ and Jenson's inequality, $T(\mathbb{E}[Y_{h,t}]) \leq \mathbb{E}[T(Y_{h,t})] \leq T(h)+t$, which implies inequality (9) by the decreasing property of the inverse function T^{-1}.

(b) By the transition of the Markov chain, we have

$$\begin{aligned}&\mathbb{E}\left[\frac{X(n+1)+T^{-1}[T(X(n+1))+t(n+1)]}{2}\right] \\ =&(1-r_{X(n)})\frac{X(n)-1+T^{-1}[T(X(n)-1)+t(n)]}{2}+r_{X(n)}\frac{X(n)+T^{-1}[T(X(n))+(t(n)-1)]}{2} \\ \leq&\frac{X(n)+T^{-1}[T(X(n))+t(n)]}{2}-\frac{1-r_{X(n)}}{2}+r_{X(n)}\frac{T^{-1}[T(X(n))+(t(n)-1)]-T^{-1}[T(X(n))+t(n)]}{2} \\ \leq&\frac{X(n)+T^{-1}[T(X(n))+t(n)]}{2}-\frac{1-r_{X(n)}}{2}+\frac{r_{X(n)}}{2}\frac{1}{\mu_{X(n)}} \\ =&\frac{X(n)+T^{-1}[T(X(n))+t(n)]}{2}.\end{aligned}$$

where the last inequality follows since the slope of $T^{-1}(x)$ is upper bounded by $\frac{1}{\mu_{X(n)}}$ for $x > T(X(n))$.

Then the optional stopping theorem implies

$$\mathbb{E}\left[\frac{X(\tau_0)+T^{-1}(T(X(\tau_0))+t(\tau_0))}{2}\right] \leq \frac{h+T^{-1}(h+t(0))}{2}.$$

Note that when $t(\tau_0) > 0$ then $X(\tau_0)$ must be zero whence $T^{-1}(T(0)+t(0)) = 0$ (since $t(0) \geq 0$), so the left-hand-side always equals $\mathbb{E}\left[\frac{Y_{h,t}+T^{-1}(T(Y_{h,t})+0)}{2}\right] = \mathbb{E}[Y_{h,t}]$ as in the statement of the theorem. When $t(\tau_0) = 0$ this same property follows automatically. Therefore, we have proved (10).

(c) We prove (11) by induction on t and h. We first prove two base cases: (1) $h = 1$ and $t \geq 0$; (2) $h \geq 1$ and $t = 0$.

If $h = 1$, the unmatched number of H-nodes is either 1 (the H-node cannot be matched by any NDD) or 0 (the H-node can be matched by at least one NDD). Thus,

$\mathbb{E}[Y_{1,t}] = (1-p_H)^t$. Note that $T^0(1) = \sum_{k=2}^{\infty} \frac{(1-p_H)^k}{1-(1-p_H)^k} \leq \frac{1}{p_H} \sum_{k=1}^{\infty} (1-p_H)^k \leq \frac{1-p_H}{p_H^2}$. So we have

$$p_H \leq \frac{1}{1+T^0(1)p_H}. \tag{18}$$

Since $(1-p_H)^t \leq e^{-p_H t}$ for all $t \geq 0$ and $e^{p_H t} \geq 1 + p_H t$, we have

$$(1-p_H)^t \leq \frac{1}{1+tp_H}. \tag{19}$$

Multiplying (18) and (19) leads to $p_H(1-p_H)^t \leq \frac{1}{1+T^0(1)p_H} \frac{1}{1+tp_H} \leq \frac{1}{1+(T^0(1)+t)p_H}$, which further implies that $\exp(p_H(1-p_H)^t) \leq 1 + \frac{1}{(T^0(1)+t)p_H}$ by the inequality $\exp(x) \leq \frac{1}{1-x}$ for $x \in (0,1)$. Rearranging, we get

$$\mathbb{E}[Y_{1,t}] = (1-p_H)^t \leq \frac{1}{p_H} \log\left(1 + \frac{1}{(T^0(1)+t)p_H}\right) \leq \frac{1}{p_H} \log\left(1 + \frac{1}{(T^0(1)+\frac{1}{4}t)p_H}\right),$$

which is exactly inequality (11).

If $t = 0$, then $Y_{h,0} \equiv h$. So it suffices to prove that $h \leq \frac{1}{p_H} \log(1 + \frac{1}{T^0(h)p_H})$. To see this note

$$\begin{aligned} T^0(h) &:= \sum_{m=h+1} \frac{(1-p_H)^m}{1-(1-p_H)^m} \\ &\leq \frac{1}{1-(1-p_H)^{h+1}} \sum_{m=h+1}^{\infty} (1-p_H)^m \\ &= \frac{1}{p_H} \frac{(1-p_H)^{h+1}}{1-(1-p_H)^{h+1}} \\ &\leq \frac{1}{p_H} \frac{e^{-p_H(h+1)}}{1-e^{-p_H(h+1)}} \end{aligned} \tag{20}$$

where the last inequality follows from $1 - p_H \leq e^{-p_H}$ and $\frac{x}{1-x}$ is increasing in x over $(0,1)$. The above inequality implies that $h+1 \leq \frac{1}{p_H} \log(1 + \frac{1}{T^0(h)p_H})$, which verifies (11) in the $t = 0$ case.

To prove (11) for the general cases of $h \geq 2$ or $t \geq 1$, we need to use induction. The induction step actually reduces to verifying a single inequality, which is detailed below.

To simplify the notation, define functions $F(x) := \log(1 + \frac{1}{x})$ and $x(t,h) := (T^0(h) + \frac{1}{4}t)p_H$. Under this notation, proving (11) is equivalent to proving

$$\mathbb{E}[Y_{h,t}] \leq F(x(t,h)). \tag{21}$$

Suppose (11) holds for $\mathbb{E}[Y_{h-1,t}]$ and $\mathbb{E}[Y_{h,t-1}]$ with $h, t \geq 2$. Then the random walk procedure implies that

$$\begin{aligned} \mathbb{E}[Y_{h,t}] &= (1-(1-p_H)^h)\mathbb{E}[Y_{h-1,t}] + (1-p_H)^h \mathbb{E}[Y_{h,t-1}] \\ &\leq (1-(1-p_H)^h) \log\left(1 + \frac{1}{(T^0(h-1)+\frac{1}{4}t)p_H}\right) + (1-p_H)^h \log\left(1 + \frac{1}{(T^0(h)+\frac{1}{4}(t-1))p_H}\right) \\ &= (1-(1-p_H)^h) F(x(t,h-1)) + (1-p_H)^h F(x(t-1,h)). \end{aligned}$$

Therefore, in order to prove (21), it suffices to establish

$$\begin{aligned} 0 &\le F(x(t,h)) - \big((1-(1-p_H)^h)F(x(t,h-1)) + (1-p_H)^h F(x(t-1,h))\big) \\ &= (1-(1-p_H)^h)(F(x(t,h)) - F(x(t,h-1))) + (1-p_H)^h(F(x(t,h)) - F(x(t-1,h))) \end{aligned} \tag{22}$$

Since $F'(x) = -\frac{1}{x(x+1)} < 0$, using convexity of $F(\cdot)$, we have

$$\begin{aligned} (1-p_H)^h(F(x(t,h)) - F(x(t-1,h))) &\ge (1-p_H)^h F'(x(t,h-1))(x(t,h) - x(t-1,h)) \\ &= -\frac{(1-p_H)^h p_H}{4x(t-1,h)(x(t-1,h)+1)}, \end{aligned} \tag{23}$$

and

$$\begin{aligned} &(1-(1-p_H)^h)(F(x(t,h)) - F(x(t,h-1))) \\ \ge &(1-(1-p_H)^h)F'(x(t,h-1))(x(t,h) - x(t,h-1)) \\ = &-\frac{1-(1-p_H)^h}{x(t,h-1)(x(t,h-1)+1)} p_H(T^0(h) - T^0(h-1)) \\ = &\frac{(1-p_H)^h p_H}{x(t,h-1)(x(t,h-1)+1)} \end{aligned} \tag{24}$$

where the last equality follows from $(T^0(h) - T^0(h-1))(1-(1-p_H)^h) = -\mu_h(1-(1-p_H)^h) = -(1-p_H)^h$. Therefore, in order to prove inequality (22), it suffices to to verify that the RHS of (23) and (24) has a positive sum, which is equivalent to the following condition,

$$\frac{x(t,h-1)(x(t,h-1)+1)}{x(t-1,h)(x(t-1,h)+1)} \le 4. \tag{25}$$

Therefore, the induction step reduces to verifying (25). We call (25) the induction condition.

We next prove (11) by discussing various sub-cases of $t \ge 1$ or $h \ge 2$. In each case, we either directly prove (21), or verify the induction condition (25).

(1) $t = 1$, $h = 2$.

In this case, we prove (21) directly. Using (20), we deduce

$$\begin{aligned} T^0(h) &\le \frac{1}{p_H}\,\frac{(1-p_H)^{h+1}}{1-(1-p_H)^{h+1}} \\ &= \frac{1-p_H}{p_H}\,\frac{(1-p_H)^h}{1-(1-p_H)^{h+1}} \\ &\le \frac{1-p_H}{p_H}\,\frac{(1-p_H)^h}{1-(1-p_H)^h} \\ &= \frac{1-p_H}{p_H}\,\mu_h \end{aligned}$$

Plugging $h = 2$ into the above inequality yields

$$\begin{aligned} T^0(2)p_H + 1 &\le 1 + (1-p_H)\frac{(1-p_H)^2}{1-(1-p_H)^2} \\ &\le \frac{p_H(2-p_H) + (1-p_H)^3}{p_H(2-p_H)} \\ &= \frac{1-(1-p_H)^2 + (1-p_H)^3}{p_H(2-p_H)} \\ &= \frac{1-(1-p_H)^2 p_H}{p_H(2-p_H)} \\ &\le \frac{1-p_H/2}{p_H(2-p_H)} \\ &\le \frac{1-p_H/2}{p_H(2-p_H)} = \frac{1}{2p_H}. \end{aligned}$$

Thus,

$$\begin{aligned} (1 + T^0(2)p_H + \tfrac{1}{4}p_H)p_H(2-p_H^2)(1-p_H) &\le (\tfrac{1}{2p_H} + \tfrac{1}{4}p_H)2p_H(1-p_H) \\ &\le (1 + \tfrac{1}{2}p_H^2)(1-p_H) \\ &= 1 - p_H + \tfrac{1}{2}p_H^2 - \tfrac{1}{2}p_H^3 \\ &\le 1 \end{aligned} \tag{26}$$

where the last inequality follows from $p_H \le 0.1$. Using the concavity of $\log(x)$, we know that $\log(x+1) - \log(x) \ge \frac{1}{1+x}$. Thus, by replacing x with $T^0(2)p_H + \frac{1}{4}p_H$, we get

$$\frac{1}{1 + T^0(2)p_H + \frac{1}{4}p_H} \le \log\left(1 + \frac{1}{T^0(2)p_H + \frac{1}{4}p_H}\right). \tag{27}$$

According to the random walk procedure, $\mathbb{E}[Y_{2,1}] = 2r_2 + 1(1-r_2)r_1 + 0(1-r_2)(1-r_1) = (2-p_H^2)(1-p_H)$. Therefore, using (26) and (27),

$$\mathbb{E}[Y_{2,1}] = \frac{1}{(1 + T^0(2)p_H + \frac{1}{4}p_H)p_H}\left((1 + T^0(2)p_H + \tfrac{1}{4}p_H)p_H(2-p_H^2)(1-p_H)\right) \tag{28}$$

$$\le \frac{1}{p_H}\log\left(1 + \frac{1}{T^0(2)p_H + \frac{1}{4}p_H}\right) = F(x(1,2)). \tag{29}$$

(2) $t = 1$, $T^0(h) \le \frac{1}{e} - \frac{1}{4}$, $h \ge 3$.

We still prove (21) directly in this case. Let $n_0 := \lceil \frac{1}{p_H} \log(\frac{1}{p_H}) \rceil$. Using the upper bound on $\Pr(Y_{h,1} \leq k)$ given in (7), we have

$$\begin{aligned}
\mathbb{E}[Y_{h,1}] &= \sum_{k=0}^{h} (1 - \Pr(Y_{h,1} \leq k)) \\
&\leq n_0 - 1 + \frac{1}{p_H} \sum_{k=n^0-1}^{h} \left((1-p_H)^{k+1} - (1-p_H)^{h+1} \right) \\
&\leq n_0 - 1 + \frac{1}{p_H} \sum_{k=n^0-1}^{\infty} (1-p_H)^{k+1} - \frac{h - n_0 + 1}{p_H} (1-p_H)^{h+1} \\
&= n_0 - 1 + \frac{(1-p_H)^{n^0}}{p_H^2} - \frac{h - n_0 + 1}{p_H} (1-p_H)^{h+1} \\
&\leq n_0 - 1 + \frac{(1-p_H)^{n^0}}{p_H^2} \\
&\leq \tfrac{1}{p_H} \log(\tfrac{1}{p_H}) + \tfrac{1}{p_H} \qquad (30)
\end{aligned}$$

When $T^0(h) < \frac{1}{e} - \frac{1}{4}$, $\log\left(1 + \frac{1}{(T^0(h) + \frac{1}{4}) p_H}\right) > \log(\frac{e}{p_H}) = 1 + \log(\frac{1}{p_H})$. Thus, (30) implies

$$\mathbb{E}[Y_{h,1}] \leq \tfrac{1}{p} \left(\log\left(\tfrac{1}{p_H}\right) + 1 \right) < \frac{1}{p_H} \log\left(1 + \frac{1}{(T^0(h) + \frac{1}{4}) p_H}\right).$$

(3) $t = 1$, $T^0(h) \geq \frac{1}{e} - \frac{1}{4}$, $h \geq 3$.
We proceed by verifying the induction condition (25). First, we show that $\mu_h \leq \frac{11}{18} T^0(h)$ for $h \geq 3$. This follows from the following inequality that holds for all $m \geq k$

$$\begin{aligned}
\frac{\mu_{m+1}}{\mu_m} &= (1-p_H) \frac{1 - (1-p_H)^m}{1 - (1-p_H)^{m+1}} = (1-p_H) \left(1 - \frac{p_H (1-p_H)^m}{1 - (1-p_H)^{m+1}}\right) \\
&\geq (1-p_H) \left(1 - \frac{p_H (1-p_H)^m}{1 - (1-p_H)^m}\right) = (1-p_H)(1 - \mu_m p_H) \geq (1-p_H)(1 - \mu_k p_H).
\end{aligned}$$

Thus,

$$\begin{aligned}
T^0(h) = \sum_{k=h+1}^{\infty} \mu_k &\geq \left(\sum_{m=0}^{\infty} (1-p_H)^m (1 - \mu_h p_H)^m \right) \mu_h \\
&= \left(\frac{1}{1 - (1-p_H)(1 - \mu_h p_H)} - 1 \right) \mu_h \qquad (31)
\end{aligned}$$

Since $(1-p_H)^h \leq e^{-p_H h} \leq \frac{1}{1 + p_H h}$, we derive an upper bound for μ_h as $\mu_h = \frac{(1-p_H)^h}{1 - (1-p_H)^h} \leq \frac{1}{p_H h}$. Plugging this bound for μ_h into (31) leads to an lower bound for $T^0(h)$ as

$$\begin{aligned}
T^0(h) &\geq \left(\frac{1}{1 - (1-p_H)(1 - \frac{1}{h})} - 1 \right) \mu_h \\
&\geq \frac{(h-1)(1-p_H)}{1 + (h-1) p_H} \mu_h.
\end{aligned}$$

Since $p_H \leq 0.1$, when $h \geq 3$, the above inequality implies that $T^0(h) \geq \frac{0.9(h-1)}{1+0.1(h-1)}\mu_h \geq \frac{18}{11}\mu_h$. Using the inequality $\mu_h \leq \frac{11}{18}T^0(h)$, we upper bound the LHS of (25) as follows,

$$\begin{aligned}\frac{x(t,h-1)(x(t,h-1)+1)}{x(t-1,h)(x(t-1,h)+1)} &= \left(\frac{T^0(h)+\mu_h+\frac{1}{4}}{T^0(h)}\right)\left(\frac{1+(T^0(h)+\mu_h+\frac{1}{4})p_H}{1+T^0(h)p_H}\right)\\ &\leq \left(\frac{\frac{29}{18}T^0(h)+\frac{1}{4}}{T^0(h)}\right)\left(\frac{1+(\frac{29}{18}T^0(h)+\frac{1}{4})p_H}{1+T^0(h)p_H}\right)\\ &= \left(\frac{\frac{29}{18}T^0(h)+\frac{1}{4}}{T^0(h)}\right)\left(\frac{\frac{29}{18}T^0(h)+\frac{1}{4}}{T^0(h)}+\frac{1-\frac{\frac{29}{18}T^0(h)+\frac{1}{4}}{T^0(h)}}{1+T^0(h)p_H}\right).\end{aligned} \tag{32}$$

It is straightforward to check that the RHS in the second row in (32) is increasing in p_H. Thus, it suffices to prove that when $p_H = 0.1$, the RHS of (32) is upper bounded by 4, which is equivalent to showing

$$g(T^0(h)) := \left(\frac{29}{18}T^0(h)+\frac{1}{4}\right)\left(1+0.1\left(\frac{29}{18}T^0(h)+\frac{1}{4}\right)\right) - 4T^0(h)\left(1+0.1T^0(h)\right) \leq 0.$$

It is not difficult to verify that $g'(x) \leq 0$ for all $x \geq 0$, and $g(\frac{1}{e}-\frac{1}{4}) = -0.0178 \leq 0$. Since $T^0(h) \geq \frac{1}{e}-\frac{1}{4}$, we deduce that $g(T^0(h)) \leq 0$, which verifies (25).

(4) $t \geq 2$, $h \geq 2$

We first show that $T^0(h) \geq \mu_h$ for all $h \geq 2$. In the previous case, we already showed $T^0(h) \geq \frac{18}{11}\mu_h \geq \mu_h$ for all $h \geq 3$. It remains to show for $h = 2$; that is, $T^0(2) \geq \mu_2$. Note that $T^0(2) = \mu_3 + T^0(3) \geq \frac{29}{11}\mu_3$, and the ratio $\frac{\mu_2}{\mu_3} = \left(\frac{(1-p_H)^2}{1-(1-p_H)^2}\right)\left(\frac{(1-p_H)^3}{1-(1-p_H)^3}\right)^{-1}$ is upper bounded by 1.585 (the maximum is attained at $p_H = 0.1$). Therefore, $T^0(2) \geq \frac{29}{11}\mu_3 \geq \frac{29}{11}\frac{1}{1.585}\mu_2 \geq \mu_2$.

Given that $t \geq 2$ and $T^0(h) \geq \mu_h$ for all $h \geq 2$, we have the inequality

$$2x(t-1,h) = 2(T^0(h)+\tfrac{1}{4}(t-1))p_H \geq (T^0(h)+\mu_h+\tfrac{1}{4}t)p_H = x(t,h-1). \tag{33}$$

Using the above inequality, we further deduce that

$$2(x(t-1,h)+1) := 2((T^0(h)+\tfrac{1}{4}(t-1))p_H+1) \geq (T^0(h)+\mu_h+\tfrac{1}{4}t)p_H+1 = x(t,h-1)+1 \tag{34}$$

Multiplying (33) and (34) leads to (25).

According to the above, the induction step works for all combinations of (h,t), so (11) is proved. □

B. PROOF OF LEMMA 4.3

PROOF. Given any realization of a random graph $D = D(h,\ell,t)$. Let V_H and V_L denote the sets of H-nodes and L-nodes in D, respectively. For any set $V_S \subseteq V_H$, define a set function $F(\cdot) : 2^H \to \mathbb{Z}_+$ such that $F(V_S)$ denotes the number of unmatched H-nodes in set V_S after the bipartite matching algorithm in Phase 2 has been applied to the bipartite subgraph $D_S = V_S \cup V_L$. We next prove that $F(\cdot)$ is a supermodular set

function. To do that, observe that $F(V_S) = |V_S| - G(V_S)$ where

$$\begin{aligned} G(V_S) \; = \max \; & \sum_{i \in V_S} \sum_{j \in L} x_{ij} \\ \text{st.} \quad & \sum_{j \in L} x_{ij} \leq 1 \text{ for all } i \in V_S \\ & \sum_{i \in H} x_{ij} \leq 1 \text{ for all } j \in L \\ & x_{ij} \in \{0,1\} \text{ for all } (i,j) \in \tilde{E} \\ & x_{ij} = 0 \text{ for all } (i,j) \notin \tilde{E}. \end{aligned} \tag{35}$$

where $\tilde{E} := \{(i,j) | i \in V_H, j \in V_L, (i,j), (j,i) \in E\}$ denotes the edge set we define on the bipartite graph (node i and j can be matched using a two-way cycle if and only if the directed arcs (i,j) and (j,i) both lie in the edge set E of graph D). According to the linear integer programming formulation, $G(V_S)$ gives the maximum number of H-nodes being matched in the bipartite subgraph D_S. If we can show that $G(V_S)$ is submodular then since $|V_S|$ is a modular function this implies that $F(V_S)$ is supermodular.

To show that G is submodular we use an approach similar in spirit to Theorem 3.4.1 of [Topkis 1998]. Recall the well-known fact that (35) has no integrality gap compared to its linear relaxation:

$$\begin{aligned} G(V_S) \; = \max \; & \sum_{i \in V_S} \sum_{j \in L} x_{ij} \\ \text{st.} \quad & \sum_{j \in L} x_{ij} \leq 1 \text{ for all } i \in V_S \\ & \sum_{i \in H} x_{ij} \leq 1 \text{ for all } j \in L \\ & x_{ij} \geq 0 \text{ for all } (i,j) \in \tilde{E} \\ & x_{ij} = 0 \text{ for all } (i,j) \notin \tilde{E}. \end{aligned}$$

Taking the linear programming dual yields

$$\begin{aligned} G(V_S) \; = \min \; & \sum_{i \in V_S} y_i + \sum_{j \in L} z_j \\ \text{st.} \quad & y_i + z_j \geq 1 \text{ for all } (i,j) \in \tilde{E} \\ & y, z \geq 0. \end{aligned}$$

Replacing z by $-z$ gives

$$\begin{aligned} G(V_S) \; = \min \; & \sum_{i \in V_S} y_i - \sum_{j \in L} z_j \\ \text{st.} \quad & y_i - z_j \geq 1 \text{ for all } (i,j) \in \tilde{E} \\ & y \geq 0 \\ & z \leq 0. \end{aligned}$$

Note that the feasible region of this linear program is a lattice. This is useful to know because we can then apply the following lemma to deduce that $G(\cdot)$ is a submodular set function.

LEMMA B.1. *Let*

$$\begin{aligned} h(a) \ &= \min \ a^\top w \\ &\text{st.} \quad w \in W \subseteq \mathbb{R}^n \end{aligned} \tag{36}$$

where $W \subset \mathbb{R}^n$ lattice. Then h is submodular in a.

PROOF. Follows directly from Theorem 2.7.6 in [Topkis 1998]. □

We then leverage Lemma B.1 to prove submodularity of $G(\cdot)$. We define a function $\tilde{G}$ as

$$\begin{aligned} \tilde{G}(\chi_{V_S}) \ &= \min \ (\chi_{V_S})^\top y - (\chi_L)^\top z \\ &\text{st.} \quad y_i - z_j \geq 1 \text{ for all} (i,j) \in \tilde{E} \\ &\qquad y \geq 0 \\ &\qquad z \leq 0. \end{aligned}$$

Observe that $G(V_S) = \tilde{G}(\chi_{V_S})$, where χ_{V_S} is a the indicator function of V_S. Clearly, if $\tilde{G}$ is submodular then G is submodular. Since $\tilde{G}$ is in the form of h in the claim, we can conclude that $G(\cdot)$ is a submodular set function.

By applying Lemma B.1, we prove that $G(\cdot)$ is supermodular and therefore $F(\cdot)$ is also supermodular. The remaining task is to show that f is convex by leveraging the supermodularity of $F(\cdot)$. The connection between f and F is the following. As the function $f(h)$ represents the expected number of remaining H-nodes for the random graph $D(h, \ell, t)$, then $f(h)$ is the expectation over $F(H)$ over all realized sets H of size h. To show that f is discrete convex, the target is to show that f has increasing differences:

$$f(h) - f(h-1) \leq f(h+1) - f(h) \tag{37}$$

for every $h \geq 1$ (we require $h \geq 1$ since f can only take nonnegative arguments and still make sense as defined).

Now, assume that $|V_S| = h-1$ for some $h \geq 1$, and $i, j \notin V_S$, $i \neq j$. Then $|V_S \cup \{i\}| = h$, $|V_S \cup \{j\}| = h$, $|V_S \cup \{i,j\}| = h+1$. Then by the supermodularity of F we have

$$F(V_S \cup \{i\}) + F(V_S \cup \{j\}) \leq F(V_S \cup \{i,j\}) + F(V_S)$$

and putting this in terms of f (and using the fact that expectations is monotone and additive) yields:

$$f(h) + f(h) \leq f(h+1) - f(h-1).$$

A little rearranging yields (37).