Strategyproofness in Kidney Exchange with Cancellations

Itai Feigenbaum*

Lehman College and the Graduate Center, City University of New York, Bronx, New York 10468, USA
itai.feigenbaum@lehman.cuny.edu

Abstract. Patients requiring kidney transplant may have proxy donors: people who want to donate a kidney to the patient, but cannot due to medical incompatibility. However, patients can swap proxy donors, so that each swapping patient ends up with a compatible donor. Some patients, called overloaded, have multiple proxy donors. A matching is a collection of planned transplants resulting from swaps: we consider both balanced matchings, which restrict overloaded patients to swapping just one of their proxy donors, and unbalanced matchings which do not. In practice, many planned transplants get canceled, and so we want matchings to maximize the expected number of actually executed transplants. Maximization of executed transplants introduces perverse incentives for overloaded patients, who can increase the probability they receive a kidney by hiding some of their proxy donors. We design the SuperGreedy Algorithm, which provably incentivizes patients to fully reveal their proxy donors. When cancellation probabilities are uniformly constant, we prove that SuperGreedy $O(1)$-approximates the maximum number of executed transplans; we also implement SuperGreedy and show via simulation that it performs well on realistic data.

Keywords: Kidney Exchange · Mechanism Design · Market Design.

1 Introduction

Kidney transplant is a treatment for renal disease. A living donor can donate one kidney without long-term harm. However, even if a patient finds a willing living donor, transplant might not be possible due to medical incompatibility. A *proxy donor* d of a patient p is a person who wishes to donate a kidney to p, but is incompatible with p; p is a *proxy patient* of d. While d is incompatible with p, d may be compatible with other patients, who in turn may have their own proxy donors. Kidney exchange programs (KEPs) allow patients to swap proxy donors, so that each swapping patient gets a compatible donor [3].

A *matching* is a plan for which transplants/swaps should take place. Some patients, called *overloaded*, have more than one proxy donor. Currently, KEPs require

* The author would like to thank John Dickerson, Loren Gragert and Navchetan Kaur for their insight into UNOS' practices. The data reported here have been supplied by the United Network for Organ Sharing as the contractor for the Organ Procurement and Transplantation Network. The interpretation and reporting of these data are the responsibility of the author and in no way should be seen as an official policy of or interpretation by the OPTN or the U.S. Government.

matchings to be *balanced*: if a patient receives a kidney through the program, only one of her proxy donors donate in return. However, we recently proposed utilization of *unbalanced* matchings, which allow several proxy donors to donate on behalf of a single patient, and can increase the number of transplants [8]. In this paper, we consider both balanced and unbalanced matchings.

Planned transplants often get canceled. For example, in the KEP organized by the United Network for Organ Sharing (UNOS), 93% of planned transplants get canceled [6]. A transplant can get canceled for reasons involving the patient or donor in the transplant, such as late discovery of incompatibility: we call these *direct cancellations*.[1] Another reason doesn't directly involve the patient or donor in a canceled transplant, but rather another transplant in the swap. Suppose a planned transplant to patient p gets canceled. Then, p's proxy donors no longer have an incentive to donate, so transplants involving them get canceled as well. Furthermore, if d is a proxy donor of p, and d was planned to donate to patient $\overline{p}$, then because the transplant from d to $\overline{p}$ gets canceled, so do transplants involving proxy donors of $\overline{p}$, and so on. Cancellations due to such propagation are called *indirect cancellations*. We would like to find a matching which induces a large expected number of *executed* transplants (transplants that actually occur, not just planned).

The matching algorithm OPT, which maximizes the expected number of executed transplants, introduces perverse incentives. Under OPT, an overloaded patient can increase the probability she receives a transplant by hiding some of her proxy donors from the KEP. In contrast, a *strategyproof* (SP) algorithm guarantees that each patient maximizes her probability of receiving a transplant by fully revealing her proxy donors, thus eliminating perverse incentives. Our goal is to design a SP matching algorithm which yields a large expected number of executed transplants.

Our main contribution is the *SuperGreedy* Algorithm, which we prove to be SP. We then consider *uniform* markets, where direct cancellation probabilities are constant across transplants; uniform markets have been studied by Dickerson et al. [6], and they provide some analytic and computational tractability to a highly complex problem. In such markets, we prove that SuperGreedy $O(1)$-approximates OPT. We also implement SuperGreedy for uniform markets, and show via simulation that it approximates OPT well in practice.

Ashlagi and Roth [3] provide a survey of kidney exchange. To the best of our knowledge, this paper is the first to consider strategic behavior of overloaded patients in kidney exchange with cancellations, and to consider unbalanced matchings with cancellations. The topic of balanced matching with cancellations was first studied by Dickerson et al. [6], leading to additional studies on the topic [4,5,10]. Strategic behavior of transplant centers in balanced matchings without cancellations has also

[1] Late discovery of incompatibility happens via a test called *crossmatch*, which is relatively expensive and requires mixing a patient's blood with a potential donor's blood, thus difficult to perform for all possible patients and donors in advance [4].

received significant attention [1,2,11]. Unbalanced matchings without cancellations are considered in our own working paper [8], and in a paper by Farina et al. [7].

The paper is organized as follows. In Section 2 we define the model and discuss some structural properties of matchings. In Section 3 we define SuperGreedy and show that it is SP (while OPT is not); we then provide a bound on SuperGreedy's approximation ratio in uniform markets. In Section 4 we show how to implement SuperGreedy and OPT in uniform markets via integer programming. In Section 5 we show that SuperGreedy approximates OPT well via simulation on realistic data in uniform markets. In Section 6, we conclude and discuss directions for future research. The appendix contains additional examples and results.

2 Model

For illustration of definitions, see Example 1 in this section and Example 2 in Appendix A. A *kidney exchange market* is a partially weighted bipartite digraph $G = (P, D, R, T, s)$, with vertices $P \cup D$, edges $R \cup T$ and function $s : T \in (0, 1)$. We drop G from the notation when it is clear from context. P and D represent patients and donors, respectively. R and T are sets of edges from P to D and from D to P respectively. $(p, d) \in R$ iff d is a proxy donor of p, and $(d, p) \in T$ iff d is deemed medically compatible with p based on initial tests (a later discovery of incompatibility between d and p is still possible). Each donor has at most one proxy patient. We denote the set of altruistic[2] donors—donors without a proxy patient—as $A \subseteq D$. For $d \in D - A$, we denote d's proxy patient as d^*. For each patient p, $p^* = \{d \in D : (p, d) \in R\}$ denotes the set of p's proxy donors; we assume $|p^*| \geq 1$. Note that d^* is a single patient for $d \in D - A$, while p^* is a set of donors for $p \in P$.

Let $M \subseteq T$. Patients and donors adjacent to an edge in M are (ex-ante) *matched* by M. A donor d s.t. either $d \in A$, or $d \in D - A$ and d^* is matched by M, is (ex-ante) *satisfied* by M. When we say matched or satisfied, we mean ex-ante unless otherwise specified. We define a *proxy donations cap* $\Lambda_{\max} \in \mathbb{Z}_{\geq 1}$. Informally, M is a *matching* if it describes a transplant plan (which might not be fully executed), where $(d, p) \in M$ means that d donates to p. Formally, M is a matching iff:

1. If $d \in D$ is matched by M, then d is satisfied by M.
2. Every donor/patient is involved in at most one transplant: for all $d \in D$, $|\{(d, p) \in M : p \in P\}| \leq 1$, and for all $p \in P$, $|\{(d, p) \in M : d \in D\}| \leq 1$.
3. No more than $\Lambda_{\max}$ proxy donors can donate on behalf of a single patient: $|\{(d, p') \in M : d \in p^*, p' \in P\}| \leq \Lambda_{\max}$ for all $p \in P$.

The *survival probability* $s(d, p)$ is the probability that a planned transplant from d to p is not directly canceled: for a matching M, the probability of direct cancellation of a transplant $(d, p) \in M$ is $1 - s(d, p)$, independently of all other cancellations in the matching. When s is a constant function—that is, when all transplants have the same

[2] Altruistic donors do not expect some patient to receive a kidney in return for their donation.

direct cancellation probability–we call the market *uniform*. An indirect cancellation of $(d,p) \in M$ occurs if $d \notin A$, and the incoming edge to d^* in M (the transplant involving d^*) is canceled, directly or indirectly; note that we allow a transplant to be simultaneously directly and indirectly canceled. A non-canceled transplant in M is *executed*. A patient who receives a kidney via an executed transplant in M is *ex-post matched* by M: let P_M be the set of patients who are ex-post matched.

In addition to transplants in M, M also induces donations to the waiting list of patients without proxy donors. We define the number of executed transplants induced by M to equal $|A| + \sum_{p \in P_M} \min\{\Lambda_{\max}, |p^*|\}$. This is equivalent to the assumption that all altruistic donors, and $\min\{\Lambda_{\max}, |p^*|\}$ of the proxy donors of any patient $p \in P_M$, execute a donation (either to P or to the waiting list). Let us justify this assumption. In practice, the waiting list's size is very large compared to the market.[3] Therefore, any donor has many compatible patients in the waiting list. For $d \in A$, if d is not matched by M (to a patient in P), we instead (informally) match d to a patient in the waiting list. Similarly, for any matched $p \in P$, if only $x < \min\{\Lambda_{\max}, |p^*|\}$ proxy donors from p^* are matched by M, we (informally) match $\min\{\Lambda_{\max}, |p^*|\} - x$ additional (arbitrary) donors from p^* to patients in the waiting list. Furthermore, when $d \in D$ is matched to a patient p (either in P or in the waiting list), and the donation from d to p gets directly but not indirectly canceled, we can immediately find a new compatible patient to d in the waiting list and redirect d's donation to that patient; we repeat this until a donation from d goes through. Therefore, if the donation from d to p doesn't get indirectly canceled, d will execute a donation with probability 1 (although perhaps not to p), and our assumption is justified. Our objective is to maximize the expected number of executed donations induced by M (to patients in P and to the waiting list), which we denote as $obj(M) = |A| + \mathbb{E}[\sum_{p \in P_M} \min\{\Lambda_{\max}, |p^*|\}]$. As the number of executed donations from altruistic donors is a constant $|A|$ independent of the matching, we are also interested in the *adjusted objective* $\overline{obj}(M) = obj(M) - |A|$.

2.1 Structure of Matchings

Each matching M can be uniquely extended to a *generalized matching* which includes the edges from patients to donors, namely $M \cup \{(p,d) \in R : \exists p' \in P, d' \in D \text{ s.t. } (d,p'), (d',p) \in M\}$. Consider the decomposition[4] of the generalized matching into weakly connected components; we abuse terminology and call it a decomposition of the matching. Each component is a graph we call *multi-arborescence*, or *marb*, defined next: see Example 1 below and Example 2 in Appendix A for illustration. An *arborescence* is a tree in which all edges point away from the root [9]. Informally, a *marb* consists of a unique directed cycle, with arborescences growing out of excess

[3] As of 5/30/2022, the number of patients in need of kidney transplant in the US is 89,950, the vast majority of whom do not have proxy donors (see https://optn.transplant.hrsa.gov/data/).

[4] The decomposition of a directed graph into weakly connected components is the decomposition of the associated undirected graph into connected components.

donors of overloaded patients in the cycle. Formally, a subgraph G' of a market G is a marb iff either it is an arborescense with an altruistic donor root or:

1. G' contains exactly one directed cycle $C = (P_C, D_C, T_C, R_C)$, called the *center*, where P_C and D_C are the vertices in the cycle from P and D respectively, while T_C and R_C are the edges in the cycle from T and R respectively; when convenient, we refer to T_C is the cycle.
2. Every vertex in G' has a unique directed trail (directed walk with distinct edges) from the center which includes all edges in the center; the trail of vertices in C is simply all the edges in C.

We consider arborescences with an altruistic donor root as marbs with an empty center; a single altruistic donor is itself a marb, and any altruistic donor unmatched by M is included in the decomposition as its own marb.

Given a matching M and a matched vertex v, $r(v, M)$ is the set of edges from T in the trail to v from the center of its marb; $|r(v, M)|$ is the number of edges in that trail, which we call the *distance* of v (from the center of its marb) in M. When the center is empty, that is, when the marb is an arborescence with an altruistic donor d as a root, $r(d, M) = \emptyset$ and $|r(d, M)| = 0$. For unmatched v, while $r(v, M)$ is undefined, we extend the notation and define $|r(v, M)| = \infty$. The *size* of a center $C = (P_C, D_C, T_C, R_C)$ is $|r(v, M)|$ for any $v \in P_C \cup D_C$, which equals $|P_C| = |D_C| = |T_C| = |R_C|$. A marb with center size i is called an i-marb ($i = 0$ for an empty center). Current practice caps cycle size for logistic reasons [3]: we therefore define a cap $C_{\max} \in \{0\} \cup \mathbb{Z}_{\geq 2}$ on the center size of each of the marbs used. When we refer to $M \subseteq T$ as a matching, we assume it respects $C_{\max}$.

For each edge $(d, p) \in M$, the probability $e(d, p, M)$ that (d, p) is executed is $\prod_{(d',p') \in r(p,M)} s(d', p')$; if $(d, p) \notin M$, then $e(d, p, M) = 0$. For $p \in P$, let $e(p, M)$ be the probability that p is ex-post matched. If p is ex-ante matched by M, $e(p, M) = \prod_{(d',p') \in r(p,M)} s(d', p')$; if p is not ex-ante matched by M, $e(p, M) = 0$. If C is a directed cycle, we define $e(C) = e(p, C)$ for (arbitrary) $p \in P_C$. Our objective function can be written as $obj(M) = |A| + \sum_{p \in P} e(p, M) \cdot \min\{\Lambda_{\max}, |p^*|\}$.

Example 1. Consider the market shown in Figure 1a. Assume that the survival probability of each transplant edge is α. Circular vertices are patients, and rectangular nodes are donors. Furthermore, assume $\Lambda_{\max} \geq 2$ and $C_{\max} \geq 3$. In Figure 1b we show an example of a generalized matching: the (non-generalized) matching M consists of the edges from donors to patients. The execution probability of each transplant in M is written next to it. M consists of three marbs. The first is a 0-marb which includes transplants (d_0, p_1) and (d_1, p_2). The second is a 2-marb: the center includes the transplants (d_3^2, p_9) and (d_9^2, p_3), and the marb also includes (d_3^1, p_4) and (d_9^1, p_{10}). The third is a 3-marb: the center includes transplants (d_5, p_6), (d_6^1, p_7), and (d_7^1, p_5), and the marb also includes (d_7^2, p_{11}) and (d_6^2, p_8). The 0-marb induces $1 + e(p_1, M) + e(p_2, M) = 1 + \alpha + \alpha^2$ donations in expectation (the 1 comes from the altruistic donor). The 2-marb induces $2e(p_3, M) + 2e(p_9, M) + e(p_4, M) +$

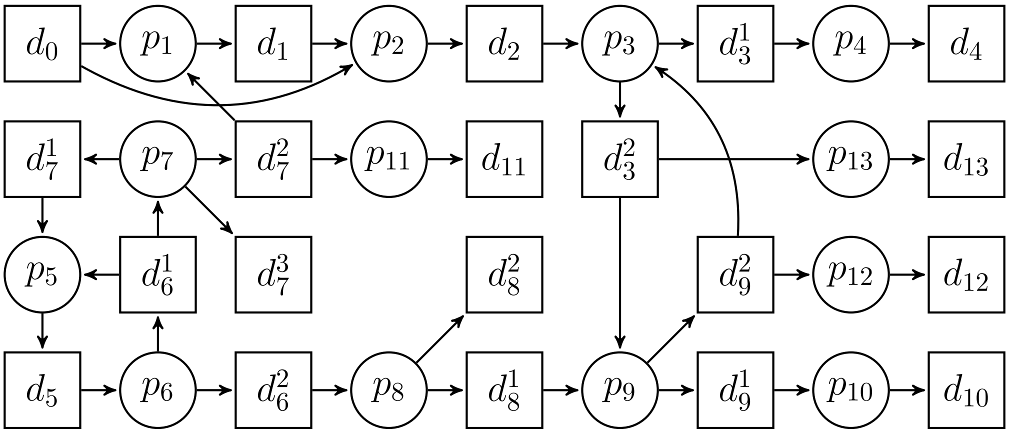

(a) Market for Example 1

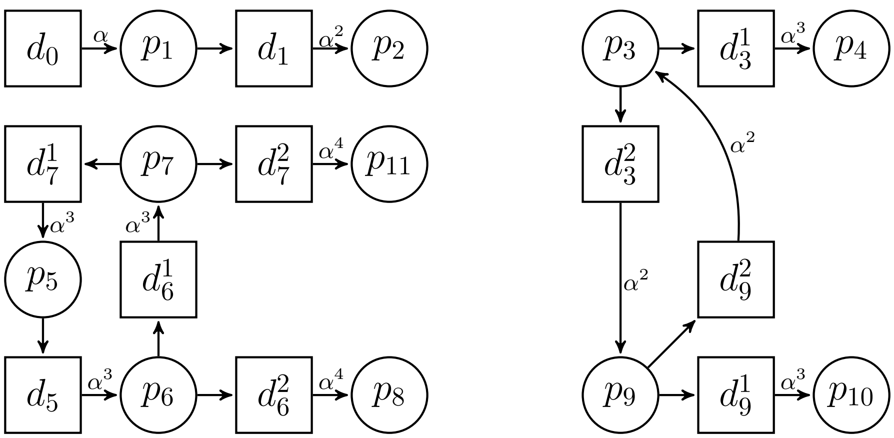

(b) Matching 1 for Example 1

Fig. 1: Market and Matching for Example 1.

$e(p_{10}, M) = 2\alpha^2 + 2\alpha^2 + \alpha^3 + \alpha^3$ donations in expectation (the coefficient is 2 for p_3 and p_9 because both have 2 proxy donors and $\Lambda_{\max} \geq 2$). The 3-marb induces $e(p_5, M) + 2e(p_6, M) + \min\{\Lambda_{\max}, 3\}e(p_7, M) + 2e(p_8, M) + e(p_{11}, M) = \alpha^3 + 2\alpha^3 + \min\{\Lambda_{\max}, 3\}\alpha^3 + 2\alpha^4 + \alpha^4$ donations in expectation. Overall, we get $1 + \alpha + 5\alpha^2 + (5 + \min\{\Lambda_{\max}, 3\})\alpha^3 + 3\alpha^4$ donations in expectation.

2.2 Mechanisms and Strategyproofness

A mechanism f is a function which maps a kidney exchange market G to a matching $f(G)$ on that market. We define OPT to be a mechanism that always chooses an optimal solution (matching M on G which maximizes $obj(M)$). Given a mechanism f and a market G, the *approximation ratio* of f on G is defined as $\frac{obj(OPT(G))}{obj(f(G))}$ (if

$obj(f(G)) = 0$, defined to be ∞ if $obj(OPT(G)) > 0$ and 1 if $obj(OPT(G)) = 0$). The *adjusted approximation ratio* is $\frac{\overline{obj}(OPT(G))}{\overline{obj}(f(G))}$. Since the adjusted approximation ratio is at least 1, and since the approximation ratio simply adds an identical constant $|A|$ to the numerator and the denominator, it follows that the adjusted approximation ratio upper bounds the approximation ratio.

Given a patient $p \in P$ and a proper subset of p's proxy donors $H \subset p^*$ s.t. $H \neq \emptyset$, let G_H be a market obtained from G by removing the donors in H and their adjacent edges. The mechanism is *strategyproof* (SP) if for every market G, every overloaded patient p in the market, and every proper subset $H \subset p^*$, $e(p, f(G_H)) - e(p, f(G)) \leq 0$. That is, p cannot benefit from reporting only $p^* - H$ as her proxy donors.

3 The SuperGreedy Mechanism

In this section, we introduce the SuperGreedy Mechanism, prove it is SP, and provide an upper bound on its adjusted approximation ratio in uniform markets. OPT is not SP (Appendix B, Theorems 3 and 4), so we look into a different class of mechanisms. A greedy mechanism operates as follows. Beginning with an empty matching M, in every iteration add a set of transplants $FLAT(\tilde{J})$ to M such that $M \cup FLAT(\tilde{J})$ is a matching, and such that all transplants in $FLAT(\tilde{J})$ have the same execution probability δ; the mechanism is greedy in that it chooses $FLAT(\tilde{J})$ which maximizes δ. For $M \cup FLAT(\tilde{J})$ to be a matching, $FLAT(\tilde{J})$ must consist of transplants that extend existing marbs in M and/or directed cycles (new marb centers) disjoint from M and from each other. Define the following notation/conventions, then see Algorithm 1 for pseudocode:

1. For $J \subseteq 2^T$ (where 2^T is the power set of T):
   (a) $FLAT(J) = \cup_{B \in J} B$: the set of edges in J.
   (b) $PATIENTS(J) = \{p \in P : \exists d \in D \text{ s.t. } (d, p) \in FLAT(J)\}$: the set of patients involved in a transplant in J.
   (c) $DONORS(J) = \{d \in D : \exists p \in P \text{ s.t. } (d, p) \in FLAT(J)\}$: the set of donors involved in a transplant in J.
   (d) $NEWSATISFIEDDONORS(J) = \{d \in D - A : d^* \in PATIENTS(J)\}$: the set of donors whose proxy patient is involved in a transplant in J.
2. For $Z \subseteq T$:
   (a) $WRAP(Z) = \{\{x\} : x \in Z\}$: "wrapping" the elements of B as subsets.
   (b) $DONORS(Z) = \{d \in D : \exists p \in P \text{ s.t. } (d, p) \in Z\}$: the set of donors involved in a transplant in Z.
3. For a set of cycles W, define $TRANSPLANTS(W) = \{T_C : C \in W\} \subseteq 2^T$.
4. For a set B and a function f, $\arg\max_{x \in B} f(x)$ is the set of all maximizers rather than an arbitrary one. Also, $\max_{x \in \emptyset} f(x) = -\infty$, and $\arg\max_{x \in \emptyset} f(x) = \emptyset$.
5. For $p \in P$, we call p^* the *proxy pool* of p. If $\Lambda_{\max}$ donors from a given proxy pool are matched by M, we say that the pool is *exhausted* in M.

Algorithm 1 Greedy Mechanism Pseudocode

Input: Market $G = (P, D, T, R, s)$, maximum cycle size $C_{\max}$, proxy donations cap $\Lambda_{\max}$

1: $M \leftarrow \emptyset$ ▷ Current matching
2: $EXH \leftarrow \emptyset$ ▷ Donors from exhausted proxy pools
3: $X \leftarrow A(G)$ ▷ Satisfied donors currently not donating to P
4: $Y \leftarrow P$ ▷ Currently unmatched patients
5: $E \leftarrow \{(d,p) \in T : d \in X, p \in Y\}$ ▷ Potential extensions to existing marbs
6: $C \leftarrow$ All directed cycles of size 2 through $C_{\max}$
7: **while** $E \neq \emptyset$ or $C \neq \emptyset$ **do**
8: $\quad \beta \leftarrow \max_{(d,p)\in E} e(p, M \cup (d,p))$ ▷ Max execution prob. of addition to existing marb
9: $\quad E' \leftarrow WRAP(\arg\max_{(d,p)\in E} e(p, M \cup (d,p))$ ▷ Max execution prob. extension transplants
10: $\quad \gamma \leftarrow \max_{\tilde{C}\in C} e(\tilde{C})$ ▷ Max execution prob. of a new cycle
11: $\quad C' \leftarrow TRANSPLANTS(\arg\max_{\tilde{C}\in C} e(\tilde{C}))$ ▷ Max. execution prob. cycles' transplants
12: $\quad$ **if** $\beta > \gamma$ **then**
13: $\quad\quad J \leftarrow E'$
14: $\quad$ **else if** $\beta < \gamma$ **then**
15: $\quad\quad J \leftarrow C'$
16: $\quad$ **else**
17: $\quad\quad J \leftarrow E' \cup C'$
18: $\quad$ **end if**
19: $\quad \tilde{J} \leftarrow$ Some $J' \subseteq J$ s.t. $M \cup FLAT(J')$ is a matching ▷ Tie-breaking rule dependent
20: $\quad M \leftarrow M \cup FLAT(\tilde{J})$
21: $\quad EXH \leftarrow \{d \in D : |(d^*)^* \cap DONORS(M)| = \Lambda_{\max}\}$
22: $\quad X \leftarrow X \cup NEWSATISFIEDDONORS(\tilde{J}) - DONORS(\tilde{J}) - EXH$
23: $\quad Y \leftarrow Y - PATIENTS(\tilde{J})$
24: $\quad E \leftarrow \{(d,p) \in T : d \in X, p \in Y\}$
25: $\quad C \leftarrow C - \{\tilde{C} \in C : P_{\tilde{C}} \cap PATIENTS(\tilde{J}) \neq \emptyset\}$ ▷ Removing cycles with matched patients
26: **end while**
27: **return** M

Different greedy mechanisms differ in tie-breaking rules between multiple options for $\tilde{J}$ (line 19). Not all natural tie-breaking rules yield SP mechanisms: see Theorem 5 in Appendix B. We define a tie-breaking rule that yields a SP greedy mechanism. Specifically, in every iteration, among the candidates for $\tilde{J}$, maximize the number of added donations: define $TB(J') = \sum_{p\in PATIENTS(J')} \min\{\Lambda_{\max}, |p^*|\}$ and choose $\tilde{J} \in \arg\max_{J'\subseteq J: M\cup FLAT(J') \text{ is a matching}} TB(J')$. Further tie-breaking is done according to a pre-fixed arbitrary order over 2^T, where we break in favor of $FLAT(\tilde{J})$ larger in the order.[5] We only require the pre-fixed orders we use to be *consistent*, meaning that when $T' \subseteq T$, the order for $2^{T'}$ agrees with the order for 2^T. This tie-breaking yields the *SuperGreedy* Mechanism $\mathbb{S}$: it is not only greedy w.r.t. maximizing δ, but also w.r.t. maximizing the number of added donations when tie-breaking. For illustration of SuperGreedy, see Example 3 in Appendix A. Next, we show that $\mathbb{S}$ is SP (we do not assume that the market is uniform for this result).

[5] If we instead tie-break uniformly at random among $\arg\max_{J'\subseteq J: M\cup FLAT(J') \text{ is a matching}} TB(J')$, our results continue to hold, assuming we extend our definition of SP to accommodate expected utilities under the additional randomization.

Theorem 1. $\mathbb{S}$ *is SP.*

Proof. Let $G = (P, D, R, T, s)$ be a market. Fix an overloaded patient $p \in P$ and a non-empty proper subset $H \subset p^*$. Let i be the iteration in which $\mathbb{S}$ adds a transplant involving p to the matching ($i = \infty$ if p is unmatched by $\mathbb{S}(G)$).

Subscript k indicates the value of a variable at the end of iteration k of $\mathbb{S}$ (when $k = 0$, it indicates the value of the variable before the first loop). The additional subscript H indicates that $\mathbb{S}$ runs on G_H, and its absence indicates that $\mathbb{S}$ runs on G. We also define T_H be the set of transplant edges in G_H, and $CYCLES(H)$ to be the set of cycles involving a donor from H in G.

Let $\delta = \min\{\beta, \gamma\}$, so in particular $\delta_k = \min\{\beta_k, \gamma_k\}$ and $\delta_{k,H} = \min\{\beta_{k,H}, \gamma_{k,H}\}$. It is easily seen that δ_k (and $\delta_{k,H}$) strictly decrease with k. Our proof structure is as follows. We will show that for $k < i$, $\mathbb{S}$ operates identically on G and G_H, meaning $M_{k,H} = M_k$: in particular, this implies that $\mathbb{S}$ doesn't match p before iteration i on G_H. Then, we will show that $\delta_{i,H} \leq \delta_i$. Since p is matched in iteration i on G, $e(p, \mathbb{S}(G)) = \delta_i$; since p is not matched before iteration i on G_H and $\delta_{k,H}$ decreases with k, $e(p, \mathbb{S}(G_H)) \leq \delta_{i,H} \leq \delta_i = e(p, \mathbb{S}(G))$, proving SP.

For iteration k, define B_k and $B_{k,H}$ to be the set of possibilities for J' on line 19 of the algorithm (on G and G_H respectively), namely:

$$B_k = \{J' \subseteq J_k \text{ s.t. } M_{k-1} \cup FLAT(J') \text{ is a matching}\}$$
$$B_{k,H} = \{J' \subseteq J_{k,H} \text{ s.t. } M_{k-1,H} \cup FLAT(J') \text{ is a matching}\}$$

or equivalently, we can define:

$$\begin{aligned} B_k = \{&J' \subseteq 2^{T - M_{k-1}} \text{ s.t. } M_{k-1} \cup FLAT(J') \text{ is a matching and} \\ &\forall p \in PATIENTS(J'), e(p, M_{k-1} \cup FLAT(J')) = \delta_k\} \\ B_{k,H} = \{&J' \subseteq 2^{T_H - M_{k-1,H}} \text{ s.t. } M_{k-1,H} \cup FLAT(J') \text{ is a matching and} \\ &\forall p \in PATIENTS(J'), e(p, M_{k-1,H} \cup FLAT(J')) = \delta_{k,H}\}. \end{aligned}$$

Let $k < i$, assume $M_{k-1,H} = M_{k-1}$, and define $M = M_{k-1,H} = M_{k-1}$. We claim that $\tilde{J}_{k,H} = \tilde{J}_k$: we will first show that $\tilde{J}_k$ is a valid candidate for $\tilde{J}_{k,H}$ on line 19 of $\mathbb{S}$ when run on G_H, and then show it is the candidate chosen based on the tie-breaking rule. Because we know that p is unmatched in M_k, and because $M_k = M \cup FLAT(\tilde{J}_k)$, we conclude that $\tilde{J}_k$ does not include any transplants involving H. Therefore, $M_k = M \cup FLAT(\tilde{J}_k)$ is a matching in G_H. Next, note:

1. In both markets, the transplant execution probabilities for M are the same, so in particular adding $FLAT(\tilde{J}_k)$ to M in G_H yields a per-transplant execution probability of δ_k for the transplants in $FLAT(\tilde{J}_k)$. δ_k is the maximum per-transplant execution probability of added transplants to M in G, and since G_H is a subgraph of G, it follows that the same is true on G_H, so $\delta_{k,H} = \delta_k$. Therefore,

$B_{k,H} \subseteq B_k$, and also since $\tilde{J}_k$ doesn't contain any transplants involving H, $\tilde{J}_k \in B_{k,H}$. Note that the latter means that $\tilde{J}_k$ is one of the candidates J' for $\tilde{J}_{k,H}$ on line 19 of the algorithm when run on G_H.

2. Let TB_W be the function TB when evaluated relatively to market W. For each $J' \subseteq 2^{T_H - M}$ where the subsets in J' are disjoint, if $p \notin PATIENTS(J')$, then $TB_{G_H}(J') = TB_G(J')$, and if $p \in PATIENTS(J')$, then $TB_{G_H}(J') \leq TB_G(J')$. Since $\tilde{J}_k \in \arg\max_{J' \in B_k} TB_G(J')$ and $p \notin PATIENTS(\tilde{J}_k)$, it follows that $\tilde{J}_k \in \arg\max_{J' \in B_k} TB_{G_H}(J')$. Since $B_{k,H} \subseteq B_k$, it follows that $\tilde{J}_k \in \arg\max_{J' \in B_{k,H}} TB_{G_H}(J')$. Consistency yields $\tilde{J}_{k,H} = \tilde{J}_k$. Therefore, $M_{k,H} = M_{k-1,H} \cup FLAT(\tilde{J}_{k,H}) = M_{k-1} \cup FLAT(\tilde{J}_{k-1}) = M_k$.

We claim that $\delta_{i,H} \leq \delta_i$. Let $M = M_{i-1,H} = M_{i-1}$. Because p is not yet matched by M, then $E_{i-1,H} = E_{i-1}$; also, $C_{i-1,H} = C_{i-1} - CYCLES(H)$. δ_i is the maximum per-edge execution probability for Z chosen from E_{i-1} and C_{i-1} s.t. $M \cup Z$ is a matching; $\delta_{i,H}$ is defined similarly, but with $C_{i-1,H} \subseteq C_{i-1}$ instead of C_{i-1}. As $\delta_{i,H}$ is maximized over a smaller selection than δ_i, $\delta_{i,H} \leq \delta_i$. □

In uniform markets with survival probability α, execution probability is determined by the distance of a vertex in the matching: if the distance is j, the execution probability is α^j. We can use this to get a simplified pseudocode (Algorithm 2) for $\mathbb{S}$ in uniform markets. Next, we prove a bound on the adjusted approximation ratio of $\mathbb{S}$ in uniform markets, under the assumption that the survival probability is less than $\min\{\frac{1}{\Lambda_{\max}}, \frac{1}{\sqrt{e}}\}$. We refer the reader to Theorem 6 in Appendix B for a slightly improved bound on the non-adjusted approximation ratio.

Algorithm 2 SuperGreedy for Uniform Markets

Input: Market $G = (P, D, T, R, s)$, maximum cycle size $C_{\max}$, proxy donations cap $\Lambda_{\max}$

$M \leftarrow \emptyset$ ▷ Current matching
$EXH \leftarrow \emptyset$ ▷ Donors from exhausted proxy pools
$X \leftarrow A(G)$ ▷ Satisfied donors currently not donating to P
$Y \leftarrow P$ ▷ Currently unmatched patients
$E \leftarrow \{(d,p) \in T : d \in X, p \in Y\}$ ▷ Potential extensions to existing marbs
for $\zeta = 1$ to $|D|$ **do** ▷ Distance of patients/donors to be added
  $E' \leftarrow WRAP(E)$ ▷ "Wrapping" edges as sets
  $C \leftarrow$ All directed cycles of size ζ involving only patients from Y
  $C' \leftarrow \{T_{\tilde{C}} : \tilde{C} \in C\}$ ▷ Size ζ cycles' transplants
  $J \leftarrow E' \cup C'$
  $\tilde{J} \leftarrow J'' \in \arg\max_{J' \subseteq J : M \cup FLAT(J') \text{ is a matching}} TB(J')$ ▷ Further tie-breaking is pre-fixed
  $M \leftarrow M \cup FLAT(\tilde{J})$
  $EXH \leftarrow \{d \in D : |(d^*)^* \cap DONORS(M)| = \Lambda_{\max}\}$
  $X \leftarrow X \cup NEWSATISFIEDDONORS(\tilde{J}) - DONORS(\tilde{J}) - EXH$
  $Y \leftarrow Y - PATIENTS(\tilde{J})$
  $E \leftarrow \{(d,p) \in T : d \in X, p \in Y\}$
end for
return M

Theorem 2. *Let $G = (P, D, T, R, s)$ be a uniform market. Assume that $s(d, p) = \alpha < \min\{\frac{1}{\Lambda_{\max}}, \frac{1}{\sqrt{e}}\}$ for all $(d, p) \in T$. The adjusted approximation ratio of $\mathbb{S}$ on G is upper bounded by $\frac{\Lambda_{\max}}{1-\Lambda_{\max}\alpha} \max\{(1 + 2\alpha - 2\alpha^2), (1 - \alpha)(1 + C_{\max}\alpha)\}$.*[6]

Proof. Throughout this proof, we refer to the implementation of $\mathbb{S}$ described in Algorithm 2. Our proof structure is as follows. First, we upper bound the number of donations induced by an i-marb. Then, we upper bound the number of i-marbs in $OPT(G)$ using the number of donations induced at distance up to i by $\mathbb{S}(G)$. Finally, we combine those bounds to bound the adjusted approximation ratio.

First, let us upper bound the expected number of donations induced by an i-marb. To do so, we assume that every patient in the marb is overloaded with at least $\Lambda_{\max}$ proxy donors, and that $\Lambda_{\max}$ of her proxy donors are matched; we call such marb *maximal*. Of course, this would mean that the marb matches an infinite number of patients/donors, but that is OK—we are calculating an upper bound.

We calculate the expected number of donations induced by a maximal i-marb. A maximal 0-marb has $\Lambda_{\max}^{j-1}$ patients at each distance $j \in \{1, 2, 3, \ldots\}$, each of which provides $\Lambda_{\max}$ proxy donors. Therefore, we get $\sum_{j=1}^{\infty} \Lambda_{\max}^{j-1}\Lambda_{\max}\alpha^j = \sum_{j=1}^{\infty} \Lambda_{\max}^{j}\alpha^j = \frac{\Lambda_{\max}\alpha}{1-\Lambda_{\max}\alpha}$ non-altruistic donations in expectation. For $i \geq 2$, a maximal i-marb has i patients matched at distance i, and $\Lambda_{\max}^{j-1}(\Lambda_{\max} - 1)i$ at distance $i + j$ for $j \in \{1, 2, 3, \ldots\}$.[7] Each matched patient provides $\Lambda_{\max}$ proxy donors, so we get $i\Lambda_{\max}$ donations induced at distance i, and $i(\Lambda_{\max} - 1)\Lambda_{\max}^{j}$ at distance $i + j$ for $j \in \{1, 2, 3, \ldots\}$. Overall, we get the expected number of donations:

$$i\Lambda_{\max}\alpha^i + i(\Lambda_{\max} - 1)\sum_{j=1}^{\infty} \Lambda_{\max}^{j}\alpha^{i+j} = i\Lambda_{\max}\alpha^i + i(\Lambda_{\max} - 1)\alpha^i\sum_{j=1}^{\infty} \Lambda_{\max}^{j}\alpha^{j} =$$

$$i\Lambda_{\max}\alpha^i + i(\Lambda_{\max} - 1)\alpha^i \cdot \frac{\Lambda_{\max}\alpha}{1 - \Lambda_{\max}\alpha} = i\alpha^i\Lambda_{\max}\frac{1-\alpha}{1 - \Lambda_{\max}\alpha}.$$

For $i \geq 1$, let $P_i^M = \{p \in P : |r(p, M)| = i\}$ be the set of patients matched at distance i by matching M. Let $\beta_i = \sum_{p \in P_i^{\mathbb{S}(G)}} \min\{\Lambda_{\max}, |p^*|\}$ be the number of induced donations at distance i in $\mathbb{S}(G)$ (including waiting list donations). Let N_0 be the number of 0-marbs in $OPT(G)$ with non-altruistic donations (have at least one edge included), and let N_i be the number of i-marbs in $OPT(G)$ for $i = 2, \ldots, C_{\max}$.

First, we bound N_0. Since every 0-marb inducing non-altruistic donations matches exactly one patient at distance 1, the number of such marbs in a matching equals the number of patients matched at distance 1 in that matching. Therefore, $N_0 = |P_1^{OPT(G)}|$. By Lemma 1 in Appendix B, $\mathbb{S}$ maximizes the number of patients matched at distance 1, so $|P_1^{OPT(G)}| \leq |P_1^{\mathbb{S}(G)}|$. Therefore, $N_0 \leq |P_1^{\mathbb{S}(G)}|$.[8]

[6] If $C_{\max} = 0$, then our analysis can actually yield an improved bound of $\frac{\Lambda_{\max}}{1-\Lambda_{\max}\alpha}$.

[7] The reason for the minus one is that each patient at level i "wastes" one donor on the cycle.

[8] We can also prove this theorem without Lemma 1, by relying on the fact that $|P_1^{OPT(G)}|$ is upper bounded by the number of donations induced at distance 1 in $OPT(G)$, which is upper bounded by β_1. However, that would complicate the proof of Theorem 6 in Appendix B.

Next, we bound N_i for $i \geq 2$. Let $\mathcal{C}_i$ be the set of centers of i-marbs in $OPT(G)$. We partition $\mathcal{C}_i$ into two sets. The first is the set of cycles in $\mathcal{C}_i$ which include no patients matched at distance less than i under $\mathbb{S}(G)$: $\mathcal{C}_i^a = \{C \in \mathcal{C}_i : \forall p \in P_C, |r(p, \mathbb{S}(G))| \geq i\}$. The second includes the remaining cycles in $\mathcal{C}_i$: $\mathcal{C}_i^u = \{C \in \mathcal{C}_i : \exists p \in P_C \text{ s.t. } |r(p, \mathbb{S}(G))| < i\}$. $\mathcal{C}_i^a$ includes the cycles from $OPT(G)$ that are available for $\mathbb{S}$ to add in iteration i, and $\mathcal{C}_i^u$ includes those that aren't.

All cycles in $\mathcal{C}_i^a$ are available for $\mathbb{S}$ in iteration i, and since they are all used in $OPT(G)$, they are disjoint. Thus, in iteration i, the collection of all cycles in $\mathcal{C}_i^a$ is an available choice in line 11 of Algorithm 2, which yields $i|\mathcal{C}_i^a|$ donations at distance i. Therefore, $\mathbb{S}$ induces at least $i|\mathcal{C}_i^a|$ donations at distance i. It follows that $i|\mathcal{C}_i^a| \leq \beta_i$.

Next, let $j \in \{2, \ldots, C_{\max}\}$. We bound $\sum_{i=2}^{j} |\mathcal{C}_i^u|$. The cycles in $\cup_{i=2}^{j} \mathcal{C}_i^u$ are disjoint, since they are all part of $OPT(G)$; also, trivially, $\mathcal{C}_z^u \cap \mathcal{C}_w^u = \emptyset$ for all $z \neq w$. By definition, for every cycle $C \in \mathcal{C}_i^u$, there exists a patient $p \in \cup_{k=1}^{i-1} P_k^{\mathbb{S}(G)}$ s.t. $p \in P_C$. Since all cycles in $\cup_{i=2}^{j} \mathcal{C}_i^u$ are disjoint, this defines a one-to-one mapping from $\cup_{i=2}^{j} \mathcal{C}_i^u$ to $\cup_{k=1}^{j-1} P_k^{\mathbb{S}(G)}$, which implies that $\sum_{i=2}^{j} |\mathcal{C}_i^u| \leq \sum_{k=1}^{j-1} |P_k^{\mathbb{S}(G)}|$.

Now let us bound the adjusted approximation ratio. When we replace the number of induced donations from each marb in $OPT(G)$ by the maximal marb upper bound we obtained, we get the following upper bound on $\overline{obj}(OPT(G))$:

$$\begin{aligned}
\overline{obj}(OPT(G)) &\leq \frac{\Lambda_{\max}\alpha}{1 - \Lambda_{\max}\alpha} N_0 + \sum_{i=2}^{C_{\max}} i\alpha^i \Lambda_{\max} \frac{1-\alpha}{1 - \Lambda_{\max}\alpha} N_i \\
&= \frac{\Lambda_{\max}\alpha}{1 - \Lambda_{\max}\alpha} N_0 + \sum_{i=2}^{C_{\max}} i\alpha^i \Lambda_{\max} \frac{1-\alpha}{1 - \Lambda_{\max}\alpha} (|\mathcal{C}_i^a| + |\mathcal{C}_i^u|)
\end{aligned}$$

We have shown $\sum_{i=2}^{j} |\mathcal{C}_i^u| \leq \sum_{k=1}^{j-1} |P_k^{\mathbb{S}(G)}|$ for all $j \in \{2, \ldots, C_{\max}\}$. In our bound, the coefficient of $|\mathcal{C}_i^u|$ is greater than the coefficient of $|\mathcal{C}_{i'}^u|$ for $2 \leq i < i'$: the derivative of $x\alpha^x$ w.r.t. x is $\alpha^x(x \log \alpha + 1)$, which is negative whenever $x \log \alpha + 1 < 0$, and since $\alpha < \frac{1}{\sqrt{e}}$, we get that for all $x \geq 2$, $x \log \alpha + 1 < 2 \log \frac{1}{\sqrt{e}} + 1 = 0$. A simple inductive argument shows that replacing $|\mathcal{C}_i^u|$ with $|P_{i-1}^{\mathbb{S}(G)}|$ for all $i \in \{2, \ldots, C_{\max}\}$ maximizes our bound. Since we also have $|\mathcal{C}_i^a| \leq \frac{1}{i}\beta_i$ and $N_0 \leq |P_1^{\mathbb{S}(G)}|$, we get:

$$\begin{aligned}
\overline{obj}(OPT(G)) &\leq \frac{\Lambda_{\max}\alpha}{1 - \Lambda_{\max}\alpha} |P_1^{\mathbb{S}(G)}| + \sum_{i=2}^{C_{\max}} i\alpha^i \Lambda_{\max} \frac{1-\alpha}{1 - \Lambda_{\max}\alpha} \left(\frac{\beta_i}{i} + |P_{i-1}^{\mathbb{S}(G)}|\right) \\
&= \frac{\Lambda_{\max}\alpha}{1 - \Lambda_{\max}\alpha} |P_1^{\mathbb{S}(G)}| + \sum_{i=2}^{C_{\max}} \alpha^i \Lambda_{\max} \frac{1-\alpha}{1 - \Lambda_{\max}\alpha} (\beta_i + i|P_{i-1}^{\mathbb{S}(G)}|).
\end{aligned}$$

Now, trivially, $|P_i^{\mathbb{S}(G)}| \leq \beta_i$, and we use that to get the bound:

$$\begin{aligned}\overline{obj}(OPT(G)) &\leq \frac{\Lambda_{\max}\alpha}{1-\Lambda_{\max}\alpha}\beta_1 + \sum_{i=2}^{C_{\max}} \alpha^i \Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha}(\beta_i + i\beta_{i-1}) \\ &= (\frac{\Lambda_{\max}\alpha}{1-\Lambda_{\max}\alpha} + 2\alpha^2\Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha})\beta_1 + \alpha^{C_{\max}}\Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha}\beta_{C_{\max}} \\ &+ \sum_{i=2}^{C_{\max}-1} \Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha}(\alpha^i + (i+1)\alpha^{i+1})\beta_i.\end{aligned}$$

Donations induced at distance at most $C_{\max}$ give $\overline{obj}(\mathbb{S}(G)) \geq \sum_{i=1}^{C_{\max}} \alpha^i\beta_i$. The adjusted approximation ratio is bounded from above by the ratio between our bounds for $\overline{obj}(OPT(G))$ and $\overline{obj}(\mathbb{S}(G))$. The coefficient ratio of β_i in this expression is:

1. For $i = 1$, $\frac{\Lambda_{\max}}{1-\Lambda_{\max}\alpha} + 2\alpha\Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha}$
2. For $1 < i < C_{\max}$, $\Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha}(1 + (i+1)\alpha)$
3. For $i = C_{\max}$, $\Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha}$

The numerator and the denominator are sums of positive terms, so the expression is bounded from above by the largest coefficient ratio, which is either the ratio for β_1 or for $\beta_{C_{\max}-1}$. Therefore, the adjusted approximation ratio is bounded by:

$$\max\{\frac{\Lambda_{\max}}{1-\Lambda_{\max}\alpha} + 2\alpha\Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha}, \Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha}(1 + C_{\max}\alpha)\},$$

and factoring out $\frac{\Lambda_{\max}}{1-\Lambda_{\max}\alpha}$ completes our proof. □

Our bound becomes tighter as α decreases. To get a sense of the bound, we note that UNOS' policy requires $C_{\max} = 3$, and that Dickerson et al. [6] estimate an upper bound of 0.3 on α. Using $C_{\max} = 3$ and $\alpha = 0.3$, for balanced matchings ($\Lambda_{\max} = 1$) the bound we get is approximately 2.03, and for $\Lambda_{\max} = 2$ the bound is approximately 7.1. We note that $\Lambda_{\max} = 2$ is pretty close to having unbalanced matchings without a proxy donations cap: based on the markets we used in our simulation in Section 5, on average approximately only 0.23% of all patients, and 4.14% of all overloaded patients have more than 2 proxy donors. So we have managed to get a constant bound, albeit with a fairly large constant for unbalanced matchings. Nevertheless, it is a worst-case bound, and furthermore, it is accomplished under very "aggressive" estimates even for a worst-case bound: our proof essentially assumes an infinite supply of overloaded patients for OPT, whereas in reality the number of patients is finite, and only a small fraction of them is overloaded. Also, as we mentioned, Dickerson et al.'s estimate of 0.3 is an upper bound on what α really is, and our bound improves as α gets smaller: for $\alpha = 0.2$, the bounds become 1.65 for $\Lambda_{\max} = 1$ and 4.4 for $\Lambda_{\max} = 2$, and for $\alpha = 0.1$ they become 1.31 for $\Lambda_{\max} = 1$ and 2.95 for $\Lambda_{\max} = 2$. Therefore, one could hope that on real data $\mathbb{S}$ performs better than our calculated bounds, and as we see in Section 5, it does indeed.

4 Implementation

In this section, we implement OPT and $\mathbb{S}$ for uniform markets with small values of $C_{\max}$ (in reality, $C_{\max} = 3$). Define a parameter $L_{\max}$ specifying the maximum distance allowed in a matching. Imposing[9] $L_{\max}$ in $\mathbb{S}$ is easily seen to preserve SP. The decrease in objective function value from when $L_{\max} = \infty$ to when $L_{\max} = l$ is bounded by $|D|\alpha^{l+1}$, which is negligible for even fairly small l. For each $a \in P \cup D$, let $z(a)$ be the set of donors/patients compatible with a. Let C_i be the set of all directed cycles of size i in the market, and for $p \in P$, define $c_i(p) \subseteq C_i$ be the set of directed cycles of size i involving patient p. Let $P = \{p_1, \ldots, p_{|P|}\}$ and $D = \{d_1, \ldots, d_{|D|}\}$. Our integer program (IP) uses the following binary variables: $p_{i,l}$ for each $p_i \in P$ and $l \in \{1, \ldots, L_{\max}\}$; $d_{i,j,l}$ for each $d_i \in D$, $p_j \in z(d_i)$, and either $l = 0$ if $d_i \in A$ or $l \in \{1, \ldots, L_{\max}\}$ if $d_i \notin A$; and $x_r \in \{0, 1\}$ for each directed cycle $r \in \cup_{i=2}^{C_{\max}} C_i$. $p_{i,l} = 1$ iff p_i is matched at distance l; $d_{i,j,l} = 1$ iff d_i is matched at distance l and donates to p_j; and $x_r = 1$ iff the matching uses cycle r. Consider the IP ($\sum_l$ means sum over l for which the variables are defined):

$$\underset{p,d}{\text{maximize}} \quad f(p) \tag{1a}$$

$$\text{subject to} \quad \sum_{d_j \in z(p_i)} d_{j,i,l-1} \geq p_{i,l} \qquad p_i \in P, l = 1 \text{ or } l > C_{\max}, \tag{1b}$$

$$\sum_{d_j \in z(p_i)} d_{j,i,l-1} + \sum_{r \in c_l(p_i)} x_r \geq p_{i,l} \qquad p_i \in P, 2 \leq l \leq C_{\max}, \tag{1c}$$

$$\sum_l \sum_{d_j \in z(p_i)} d_{j,i,l} \leq 1 \qquad p_i \in P, \tag{1d}$$

$$\sum_{p_j \in z(d_i)} d_{i,j,l} \leq p_{k,l} \qquad d_i \in D - A, p_k = d_i^*, l \geq 1, \tag{1e}$$

$$\sum_l p_{i,l} \leq 1 \qquad p_i \in P, \tag{1f}$$

$$\sum_l \sum_{p_j \in z(d_i)} d_{i,j,l} \leq 1 \qquad d_i \in D, \tag{1g}$$

$$\sum_{r \in \cup_{j \leq C_{\max}} c_j(p_i)} x_r \leq 1 \qquad p_i \in P, \tag{1h}$$

$$\sum_{(d_i, p_j) \in T_r} (d_{i,j,|T_r|} + p_{j,|T_r|}) \geq 2|T_r| x_r \qquad r \in \cup_{k=2}^{C_{\max}} C_k, \tag{1i}$$

$$\sum_{d_j \in p_i^*, p_{j'} \in z(d_j)} d_{j,j',l} \leq \Lambda_{\max} \quad p_i \in P, l \geq 1 \tag{1j}$$

The constraints capture exactly all possible matchings. Let us explain them:

[9] In Algorithm 1, stop when $\eta > \alpha^{L_{\max}}$; in Algorithm 2, run the loop up to $L_{\max}$ instead of $|D|$.

1b: When no size l cycles are allowed, either because they violate $C_{\max}$ or because $l = 1$ (there are no cycles of size 1), p_i is matched at distance l iff she receives a donation from a donor d_j matched at distance $l - 1$.
1c: When size l cycles are allowed, p_i is matched at distance l iff either she receives a donation from a donor d_j matched at distance $l - 1$, or she receives it via a cycle of size l.
1d: Each patient receives at most one donation.
1e: $d_i \in D - A$ can donate a kidney while matched at distance l only if d_i^* is matched at distance l.
1f: Each patient can only be matched at a single distance.
1g: Each donor can only be matched at a single distance, and each donor can donate to at most one patient.
1h: No patient can be used in more than one cycle.
1i: Recall that for a cycle r, T_r is the set of transplants in the cycle. For cycle r to be used, the donations in the cycle must occur at distance $|T_r|$—that is, all donations in the cycle must be used, and the donors/patients must be matched at distance $|T_r|$. The coefficient on the r.h.s. is $2|T_r|$ because that is the number of variables in the l.h.s. in total.
1j: No more than $\Lambda_{\max}$ proxy donors can donate on behalf of a single patient.

For $C_{\max} = 3$, all directed cycles of size up to $C_{\max}$ can be generated via brute force in $O(|D|)$ time (as there are $\binom{|D|}{i}$ such cycles for $i = 2, 3$). Generating the cycles took less than 3 hours per market in our simulation on a consumer PC.

Our IP can compute OPT and $\mathbb{S}$. $f(p) = \sum_{l=1}^{L_{\max}} \sum_{i=1}^{|P|} \alpha^l \min\{\Lambda_{\max}, |p_i^*|\} p_{i,l}$ yields OPT, as this is $\overline{obj}$.[10] For $\mathbb{S}$, the objective function involves very large coefficients, making feeding the IP directly into a solver impractical. The difficulty, surprisingly, is not in implementing the main part of the tie-breaking rule, but rather the secondary part, which breaks further ties according to some pre-specified ordering over 2^T. Nevertheless, **we can solve the IP iteratively**, if we use an appropriate ordering: we show the iterative solution method in Appendix C.

Let $\succ$ be some ordering on T. The extension of $\succ$ to an ordering $\succ^*$ over 2^T is done as follows: for any two distinct sets $B, B' \in 2^T$, $B \succ^* B'$ iff either $B' \subset B$, or $B' \not\subset B$, $B \not\subset B'$, and $\arg\max_{B-B'}(\succ) \succ \arg\max_{B'-B}(\succ)$.[11] This is equivalent to going through the edges in T in decreasing order of $\succ$, stopping at the first edge (d, p) that is in exactly one of the two sets, and choosing the set that contains it to be the larger one. Let $\mu(d_i, p_j) = k - 1$ iff (d_i, p_j) is the k-th largest edge in T according to $\succ$. We set a big-M value $\beta = |D| + 2^{|T|}$, and define

$$f(p) = \sum_{l=1}^{L_{\max}} \sum_{i=1}^{|P|} \beta^{2L_{\max}-2l+1} \min\{\Lambda_{\max}, |p_i^*|\} p_{i,l} + \sum_{l=1}^{L_{\max}} \sum_{(d_i,p_j)\in T} 2^{\mu(d_i,p_j)} \beta^{2L_{\max}-2l} d_{i,j,l}$$

[10] Since $\overline{obj}$ differs from obj by a constant, it does not matter for optimization.
[11] Since the element maximizing $\succ$ is unique, we slightly change our earlier convention and consider $\arg\max$ to be a specific element in this case instead of a set of elements.

Informally speaking, the first term is used for maximizing donations and the second term is used for tie-breaking. Our choice of β is large enough so that $\beta^{j+1} > \phi\beta^j$ for any number ϕ that will come up in our analysis. Therefore the algorithm maximizes the coefficients of β^j in decreasing order from $j = 2L_{\max} - 1$ to $j = 0$, meaning that for every j, we maximize the coefficients of β^j subject to maximizing the coefficients of β^{j+1}, subject to maximizing the coefficients of β^{j+2}, and so on until $\beta^{2L_{\max}-1}$. For odd powers $\beta^{2L_{\max}-1}, \beta^{2L_{\max}-3}, \ldots, \beta^1$ the coefficient of $\beta^{2L_{\max}-2l+1}$ is $\sum_{i=1}^{|P|} \min\{\Lambda_{\max}, |p_i^*|\} p_{i,l}$, which is the number of donations induced at distance l. For even powers $\beta^{2L_{\max}-2}, \beta^{2L_{\max}-4}, \ldots, \beta^0$, the coefficient of $\beta^{2L_{\max}-2l}$ is $\sum_{(d_i,p_j)\in T} 2^{\mu(d_i,p_j)} d_{i,j,l}$. Note that for all $w \in \mathbb{Z}_{\geq 1}$, $2^w = 1 + \sum_{k=0}^{w-1} 2^k$. Therefore, the coefficient of $\beta^{2L_{\max}-2l}$ breaks ties (between edges added at distance l) according to $\succ^*$.[12] Note that tie-breaking at distance l is favored over donation maximization at distances $l+1, \ldots, L_{\max}$, as $\mathbb{S}$ requires. Unfortuantely, a commercial solver fails to handle relatively small powers of β, since this method requires that $\beta > |D| + \sum_{i=1}^{|T|} 2^{i-1} = |D| + 2^{|T|} - 1$. Nevertheless, as mentioned, we show in Appendix C how to solve this IP iteratively for an appropriate choice of $\succ$.

5 Simulation

We simulate $\mathbb{S}$ and OPT on real markets from UNOS, under the assumption that the markets are uniform. Our data contains 519 real markets on which UNOS conducted matchings, from October 2010 to May 2019. For increased independence, we sampled 52 equidistant[13] markets. Basic statistics regarding our 52 markets are provided in Table 1. While the data does not contain comprehensive information about cancellations, we simulate each algorithm for constant survival probability of $\alpha = 0.1, 0.2, 0.3$. Dickerson et al. [6] estimate an upper bound of 0.3 on the survival probability. We set $L_{\max} = 12$: our largest market contains 301 donors, so the objective function value difference from $L_{\max} = \infty$ is bounded by $301 \cdot 0.3^{13} \leq 0.00005$.

Our simulation results[14] for the adjusted objective are shown in Table 2, and Figure 2 gives additional detail; we report similar results for the non-adjusted objective in Appendix D. When $\alpha = 0.3$, the average adjusted approximation ratio ranges from 1.142 for balanced matchings to 1.036 for unrestricted unbalanced matchings ($\Lambda_{\max} = 4$, which is unrestricted since no overloaded patient in the data has more than 4 proxy donors). Like our theoretical bounds, the simulated performance of $\mathbb{S}$ improves as α decreases. However, unlike our theoretical bounds, the simulated performance of $\mathbb{S}$ improves as $\Lambda_{\max}$ increases. The latter discrepancy is not surpris-

[12] Let $(d_{i'}, p_{j'}) \in T$, and let $B = \{(d,p) \in T : (d_{i'}, p_{j'}) \succ (d,p)\}$. Setting $d_{i',j',l} = 1$ and $d_{i'',j'',l} = 0$ for all $(d_{i''}, p_{j''}) \in B$ increases $\sum_{(d_i,p_j)\in T} 2^{\mu(d_i,p_j)} d_{i,j,l}$ strictly more than setting $d_{i',j',l} = 0$ and $d_{i'',j'',l} = 1$ for all $(d_{i''}, p_{j''}) \in B$.

[13] Distance defined as the number of additional match runs conducted in between the two runs.

[14] Note that columns 1-4 in Table 2 cannot be derived from dividing columns 9-12 by 5-8 respectively, as the adjusted approximation ratio is calculated before taking the mean.

patients #	donors #	overloaded %	altruistic %	compatible %
213.48 (40.28)	227.86 (43.33)	5.14% (0.98)	0.85% (0.9)	10.51% (1.27)

Table 1: Basic statistics of our sampled UNOS data. There are 52 markets in our data. We specify the mean (and standard deviation in parenthesis) of each characteristic over all markets. The following characteristics are considered, in order from left to right: number of patients; number of donors; percentage of patients who are overloaded; percentage of donors who are altruistic; percentage of donor-patient pairs who are compatible (out of all donor-patient pairs, proxies not excluded).

	adjusted approx. ratio				$\overline{obj}(\mathbb{S}(G))$				$\overline{obj}(OPT(G))$			
α \ $\Lambda_{\max}$	1	2	3	4	1	2	3	4	1	2	3	4
0.1	1.137	1.028	1.02	1.016	0.375	0.417	0.42	0.422	0.425	0.429	0.429	0.429
0.2	1.138	1.038	1.03	1.026	1.185	1.322	1.334	1.34	1.35	1.373	1.375	1.376
0.3	1.142	1.047	1.04	1.036	2.448	2.75	2.776	2.79	2.801	2.878	2.885	2.887

Table 2: Mean simulation results for $\mathbb{S}$. The mean is taken over all 52 markets.

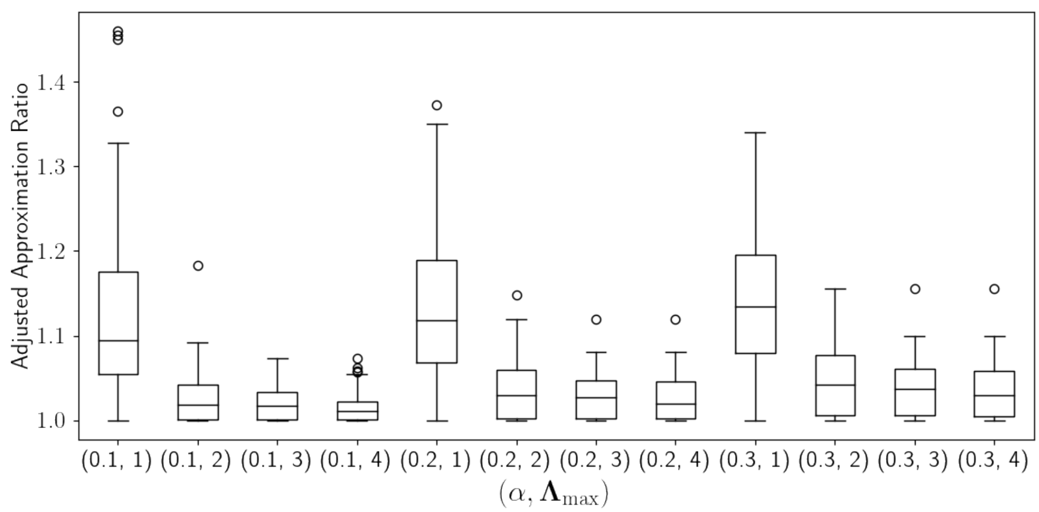

Fig. 2: Box plot for the adjusted approximation ratio of $\mathbb{S}$, showing quartiles q_1-q_3. Empty circles are outliers (points outside $[q_1 - 1.5(q_3 - q_1), q_3 + 1.5(q_3 - q_1)]$).

ing, as our theoretical bounds were derived using an infinite supply of overloaded patients with $\Lambda_{\max}$ proxy donors, but in the data there are few such patients.

6 Conclusion and Future Directions

In this paper, we have shown (Theorems 3 and 4 in Appendix B) that transplant cancellations can cause perverse incentives: overloaded patients might increase their chance of being matched by hiding some of their proxy donors. These incentives exist in both balanced and unbalanced matchings. We designed the SuperGreedy matching algorithm, which provably eliminates perverse incentives. We analytically bounded SuperGreedy's adjusted approximation ratio on uniform markets, and showed via simulation that SuperGreedy performs much better than those bounds in practice. In the process, we implemented SuperGreedy for uniform markets.

There are several questions stemming from our work. There is a significant gap between our analytic bounds and the real-world performance (in terms of expected donations) of SuperGreedy on uniform markets: it would be useful to get tighter analytic bounds to bridge this gap. Also, while SuperGreedy is SP in all markets, we do not have analytic or simulated results regarding its performance in non-uniform markets, which would be an important extension of our work. Finally, it would be interesting to extend our model to include additional edge weights representing the quality of the transplant (in terms of predicted graft survival).

References

1. Nikhil Agarwal, Itai Ashlagi, Eduardo Azevedo, Clayton R Featherstone, and Ömer Karaduman. Market failure in kidney exchange. *American Economic Review*, 109(11):4026–70, 2019.
2. Itai Ashlagi and Alvin E Roth. Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics*, 9(3):817–863, 2014.
3. Itai Ashlagi and Alvin E Roth. Kidney exchange: an operations perspective. 2020.
4. Avrim Blum, John P Dickerson, Nika Haghtalab, Ariel D Procaccia, Tuomas Sandholm, and Ankit Sharma. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. *Operations research*, 68(1):16–34, 2020.
5. John Dickerson, Duncan McElfresh, Ke Ren, et al. Kidney exchange with inhomogeneous edge existence uncertainty. In *Conference on Uncertainty in Artificial Intelligence*, pages 161–170. PMLR, 2020.
6. John P Dickerson, Ariel D Procaccia, and Tuomas Sandholm. Failure-aware kidney exchange. *Management Science*, 65(4):1768–1791, 2019.
7. Gabriele Farina, John P Dickerson, and Tuomas Sandholm. Operation frames and clubs in kidney exchange. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 199–205. AAAI Press, 2017.
8. Itai Feigenbaum. Unbalanced matchings in kidney exchange. *Working paper, SSRN 3885977.*
9. Jonathan L Gross, Jay Yellen, and Mark Anderson. *Graph theory and its applications.* Chapman and Hall/CRC, 2018.
10. Duncan McElfresh, Michael Curry, Tuomas Sandholm, and John Dickerson. Improving policy-constrained kidney exchange via pre-screening. *Advances in Neural Information Processing Systems*, 33, 2020.
11. Tayfun Sönmez and M Utku Ünver. Market design for kidney exchange. *The Handbook of Market Design*, 93137, 2013.

Appendix A Examples

In this appendix, we provide examples to illustrate some of our definitions.

Example 2. Consider the market shown in Figure 1a. Assume that the survival probability of each transplant edge is α. Circular vertices are patients, and rectangular nodes are donors. In Figures 3a, 3b and 3c we show examples of generalized matchings: the (non-generalized) matchings consist of just the edges from donors to patients. The execution probability of each transplant in the matching is shown next to it.

1. The graph in 3a is a (generalized) matching for all $\Lambda_{\max} \geq 1, C_{\max} \geq 3$. It consists of three marbs: the 0-marb which includes transplants (d_0, p_1) and (d_1, p_2), the 2-marb which includes transplants (d_3^2, p_9) and (d_9^2, p_3), and the 3-marb which includes transplants (d_5, p_6), (d_6^1, p_7) and (d_7^1, p_5). Both the 2-marb and the 3-marb consist of only their center and nothing else. The 0-marb induces $1 + \alpha + \alpha^2$ donations in expectation; the 2-marb induces $2 \min\{\Lambda_{\max}, 2\}\alpha^2$ since both patients have two proxy proxy donors; the 3-marb induces $(1 + \min\{\Lambda_{\max}, 2\} + \min\{\Lambda_{\max}, 3\})\alpha^3$ since p_5, p_6 and p_7 have 1, 2 and 3 proxy donors respectively. Overall, $1 + \alpha + (1 + 2\min\{\Lambda_{\max}, 2\})\alpha^2 + (1 + \min\{\Lambda_{\max}, 2\} + \min\{\Lambda_{\max}, 3\})\alpha^3$ donations are induced in expectation.
2. The graph in 3b is a (generalized) matching for all $\Lambda_{\max} \geq 2$, $C_{\max} \geq 2$. It consists of three marbs. The first is a 0-marb which includes transplants (d_0, p_1) and (d_1, p_2). The second is a 2-marb : the center includes transplants (d_5, p_6) and (d_6^1, p_5), and the marb also includes (d_6^2, p_8). The third is also a 2-marb: the center includes transplants (d_3^2, p_9) and (d_9^2, p_3), and the marb also includes (d_3^1, p_4) and (d_9^1, p_{10}). The first marb induces $1 + \alpha + \alpha^2$ donations in expectation; the second marb induces $\alpha^2 + 2\alpha^2 + 2\alpha^3$ donations (first, second and third terms correspond to p_5, p_6 and p_8 respectively); and the third marb induces $2 \cdot 2\alpha^2 + 2\alpha^3$ donations (the first term corresponds to p_3 and p_9, and the second one corresponds to p_4 and p_{10}). Overall, $1 + \alpha + 8\alpha^2 + 4\alpha^3$ donations are induced in expectation.
3. The graph in 3c is a (generalized) matching for all $\Lambda_{\max} \geq 2$, $C_{\max} \geq 3$. It consists of two marbs. The first is a 0-marb which includes transplants (d_0, p_2), (d_2, p_3), (d_3^1, p_4) and (d_3^2, p_{13}). The second is a 3-marb: the center includes transplants (d_5, p_6), (d_6^1, p_7), and (d_7^1, p_5), and the marb also includes (d_7^2, p_{11}), (d_6^2, p_8), (d_8^1, p_9), (d_9^1, p_{10}), and (d_9^2, p_{12}). The 0-marb induces $1+\alpha+2\alpha^2+\alpha^3+\alpha^3$ donations in expectation, while the 3-marb induces $\alpha^3+2\alpha^3+\min\{\Lambda_{\max}, 3\}\alpha^3+2\alpha^4+2\alpha^5+\alpha^6+\alpha^6$ donations. Overall, $1+\alpha+2\alpha^2+(5+\min\{\Lambda_{\max}, 3\})\alpha^3+2\alpha^4+2\alpha^5+2\alpha^6$ donations are induced in expectation.

Example 3. Consider the market shown in Figure 4. Assume $\Lambda_{\max} \geq 2$, and note that the largest number of proxy donors per patient in the market is 2, hence $\Lambda_{\max}$ plays no role. We will trace $\mathbb{S}$ for two examples: the first is when the market is

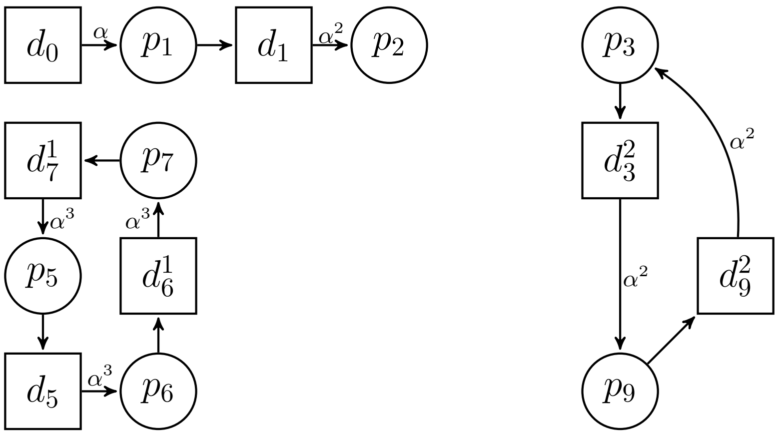

(a) Matching 2 for Example 2

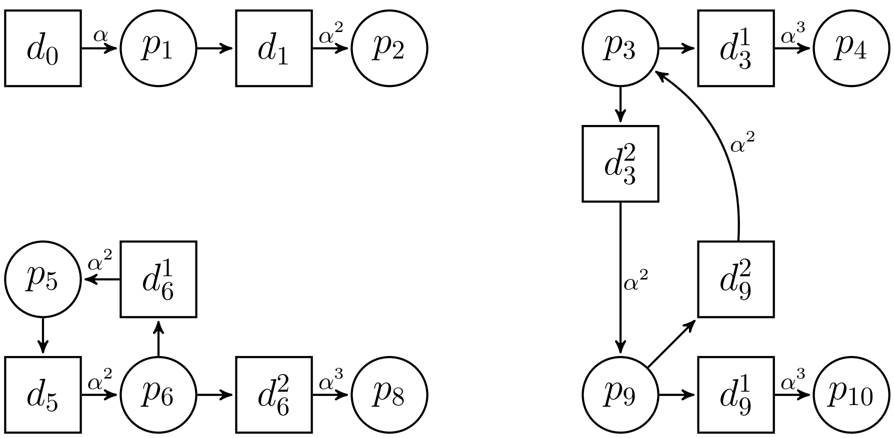

(b) Matching 3 for Example 2

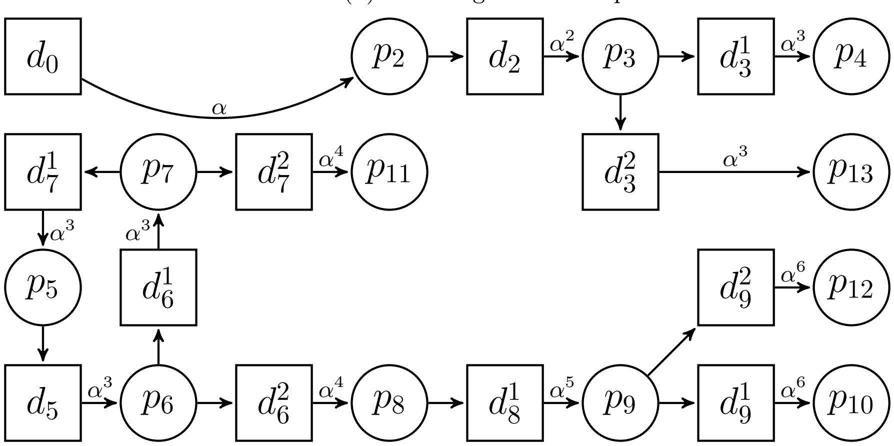

(c) Matching 4 for Example 2

Fig. 3: Matchings for Example 2.

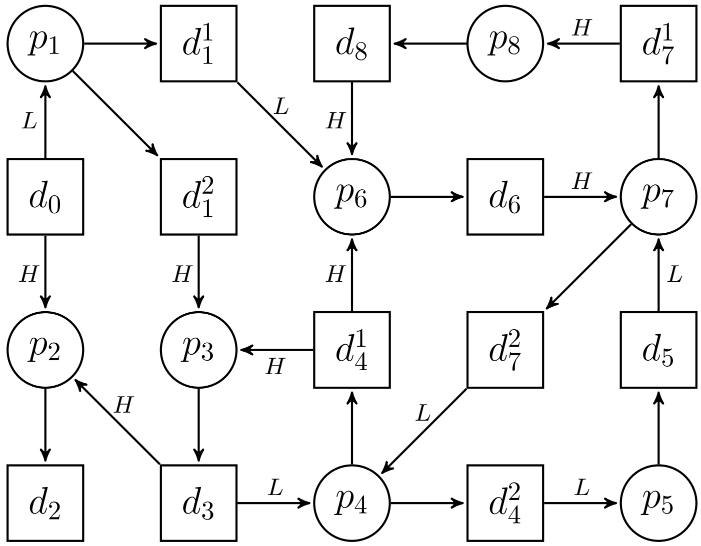

Fig. 4: Market for Theorem 3.

non-uniform, and the second when the market is uniform. In our first example, tie-breaking happens to not be relevant—that is, $\mathbb{S}$ and every other greedy mechanism will output the same matching. In the second one, tie-breaking does play a role. So first, assume that the survival probability of each edge is as specified next to the edge; furthermore, assume $0 < L < H < 1$ and $L < H^3$. We trace the operation of $\mathbb{S}$, beginning with an empty matching: the matching at the end of iteration i is denoted as M_i.

1. In the first iteration, the highest execution probability for a marb extension is H, obtained uniquely by adding (d_0, p_2). The highest execution probability for a cycle is H^3, obtained uniquely by the cycle $p_6 \to d_6 \to p_7 \to d_7^1 \to p_8 \to d_8 \to p_6$. Therefore, $M_1 = \{(d_0, p_2)\}$.
2. In the second iteration, there are no possible extensions to the 0-marb, and the highest execution probability cycle remains the same as in the first iteration. Therefore, we add a new marb center: the cycle $p_6 \to d_6 \to p_7 \to d_7^1 \to p_8 \to d_8 \to p_6$, namely $M_2 = M_1 \cup \{(d_6, p_7), (d_7^1, p_8), (d_8, p_6)\}$.
3. In the third iteration, the only possible marb extension is adding (d_7^2, p_4) to the 3-marb with execution probability H^3L. There is also one cycle left without any matched patients, namely $p_3 \to d_3 \to p_4 \to d_4^1 \to p_3$, with execution probability HL. As $HL > H^3L$, the cycle is added and we get $M_3 = M_2 \cup \{(d_3, p_4), (d_4^1, p_3)\}$.
4. In the fourth iteration, there are no more possible cycles to add. The only possible marb extension is adding (d_4^2, p_5) to the 2-marb with execution probability HL^2: therefore, $M_4 = M_3 \cup \{(d_4^2, p_5)\}$.

5. At the end of the fourth iteration, there is no way to extend any marb or add any cycle. Therefore the algorithm ends and the final ouptut is $M_4 = \{(d_0, p_2), (d_6, p_7), (d_7^1, p_8), (d_8, p_6), (d_3, p_4), (d_4^1, p_3), (d_4^2, p_5)\}$.

For the second example, ignore the survival probabilities in the figure and assume instead that the survival probability of each transplant edge is α. In that case:

1. In the first iteration, there are two transplants that can be added with execution probability α to the 0-marb containing d_0: (d_0, p_1) and (d_0, p_2). These are the only candidates, since there are no cycles with execution probability less than or equal to α. We cannot add both of these edges, since d_0 can only donate to one candidate. Since $|p_1^*| > |p_2^*|$, $\mathbb{S}$ adds (d_0, p_1) to the matching, and $M_1 = \{(d_0, p_1)\}$.
2. In the second iteration, the existing 0-marb can be extended by transplants (d_1^1, p_6) and (d_1^2, p_3), each with resulting execution probability α^2. The cycle $p_3 \to d_3 \to p_4 \to d_4^1 \to p_3$ can also be added with execution probability α^2. (d_1^1, p_6) does not overlap with the other extension edge or the cycle, and therefore will be added. On the other hand, adding (d_1^2, p_3) only matches p_3 and therefore adds $|p_3^*| = 1$ donations, while adding the cycle matches $|p_3^*| + |p_4^*| = 3$ donations, so we add the cycle. Overall, $M_2 = M_1 \cup \{(d_1^1, p_6), (d_3, p_4), (d_4^1, p_3)\}$.
3. In the third iteration, we can extend the 0-marb by adding (d_6, p_7) with execution probability α^3, and similarly we can extend the 2-marb by adding (d_4^2, p_5), and that's it. Since the two transplants are vertex disjoint, we can add both of them and get $M_3 = M_2 \cup \{(d_6, p_7), (d_4^2, p_5)\}$.
4. In the fourth iteration, the only extension possible is by adding (d_7^1, p_8) to the 0-marb: $M_4 = M_3 \cup \{(d_7^1, p_8)\}$.
5. At the end of the fourth iteration, there is no way to extend any marb or add any cycle. Therefore the algorithm ends and the final ouptut is $M_4 = \{(d_0, p_1), (d_1^1, p_6), (d_3, p_4), (d_4^1, p_3), (d_6, p_7), (d_4^2, p_5), (d_7^1, p_8)\}$.

Appendix B Additional Results

In this appendix, we provide additional results.

Theorem 3. *OPT is not SP.*

Proof. Consider market G in Figure 5, where the survival probability of each edge is written next to it. We assume that $\alpha >> \epsilon > 0$. $OPT(G) = \{(d_0, p_1), (d_1, p_2), (\tilde{d_0}, p_3)\}$, and also $OPT(G_{\{d_1\}}) = \{(\tilde{d_0}, p_1), (d_0, p_2)\}$. However, $e(p_1, OPT(G)) = \alpha$ while $e(p_1, OPT(G_{\{d_1\}})) = \alpha + \epsilon$, violating SP. □

Theorem 4. *OPT is not SP, even in uniform markets.*

Proof. We give two examples.

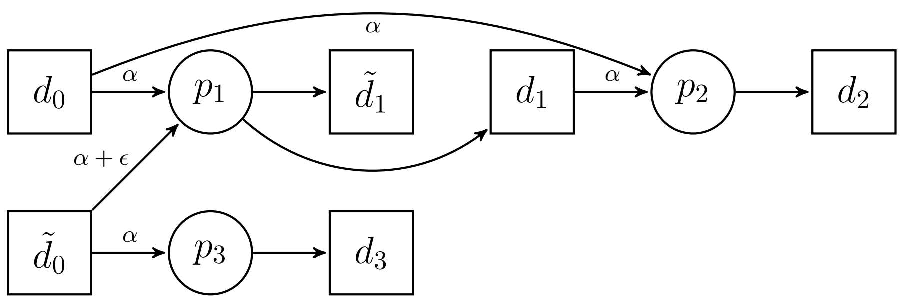

Fig. 5: Market for the proof of Theorem 3

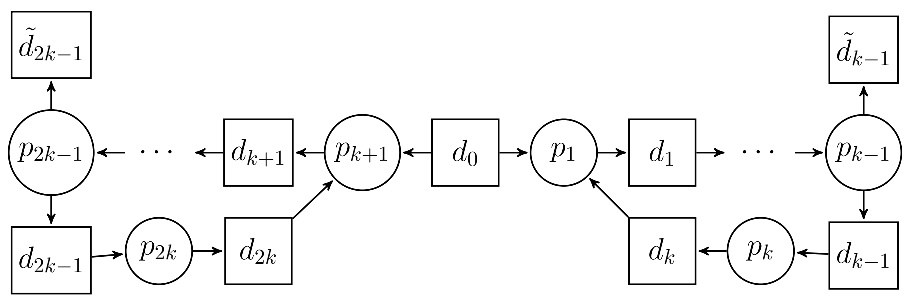

(a)

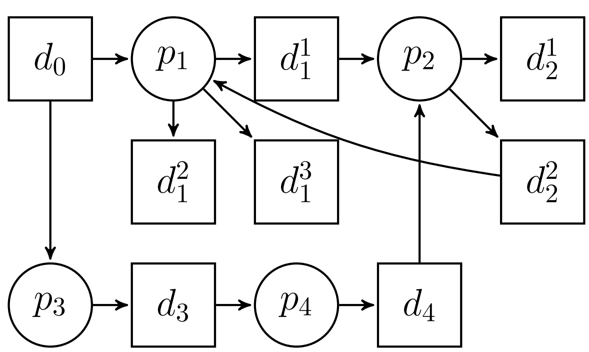

(b)

Fig. 6: Markets for the proof of Theorem 4

1. First, assume $C_{\max} \geq k$ for some $k \geq 2$. Consider a market $G = (P, D, R, T, s)$, shown on Figure 6a, where $s(d,p) = \alpha \in (\frac{1}{k}, 1)$ for all $(d,p) \in T$, and $P = \{p_1, \ldots, p_{2k}\}$. For $i \in \{k-1, 2k-1\}$, $p_i^* = \{d_i, \tilde{d}_i\}$, and for all other i, $p_i^* = \{d_i\}$; there is also one altruistic donor d_0. The edges in T consist of (d_i, p_{i+1}) for $i \in \{1, \ldots, k-1\} \cup \{k+1, \ldots, 2k-1\}$, (d_k, p_1), (d_{2k}, p_{k+1}), (d_0, p_1) and (d_0, p_{k+1}). Essentially, this market consists of two directed cycles $C_1 = p_1 \to d_1 \to p_2 \to d_2 \to \cdots \to d_k \to p_1$ and $C_2 = p_{k+1} \to d_{k+1} \to p_{k+2} \to d_{k+2} \to \cdots \to d_{2k} \to p_{k+1}$, an extra proxy donor for p_{k-1} and p_{2k-1}, and an altruistic donor compatible with p_1 and p_{k+1}.
   Assume WLOG that $OPT(G)$ consists of all edges in T except (d_0, p_1) and (d_{2k}, p_{k+1}): in other words, it consists of the transplants in C_1 as well as the chain (d_i, p_{i+1}) for $i \in \{k+1, \ldots, 2k-1\}$. Then, $e(p_{k-1}, OPT(G)) = \alpha^k$. On the other hand, consider $G_{\{d_{k-1}\}}$: in this market, if d_0 donates to p_{k+1}, the objective would be maximized by using the chain (d_i, p_{i+1}) for $i \in \{k+1, \ldots, 2k-1\}$ (and nothing else), which would yield $\alpha^{k-1}\mathbb{1}(\Lambda_{\max} \geq 2) + \sum_{i=0}^{k} \alpha^i$ donations in expectation ($\mathbb{1}(\Lambda_{\max} \geq 2)$ is an indicator variable, the first term comes from the extra donation from $\tilde{d}_{2k-1}$). On the other hand, if d_0 donates to p_1, then objective maximization is achieved by using C_2 and the chain (d_i, p_{i+1}) for $i \in \{1, \ldots, k-2\}$, which yields $\sum_{i=0}^{k-1} \alpha^i + (k + \mathbb{1}(\Lambda_{\max} \geq 2))\alpha^k$ expected transplants, where the first term comes from the chain, and the second term comes from C_2. Therefore, $OPT(G_{\{d_{k-1}\}})$ consists of the latter solution (C_2 and the chain (d_i, p_{i+1}) for $i \in \{1, \ldots, k-1\}$) if $\sum_{i=0}^{k-1} \alpha^i + (k + \mathbb{1}(\Lambda_{\max} \geq 2))\alpha^k > \alpha^{k-1}\mathbb{1}(\Lambda_{\max} \geq 2) + \sum_{i=0}^{k} \alpha^i$, or equivalently $(k - 1 + \mathbb{1}(\Lambda_{\max} \geq 2))\alpha^k > \alpha^{k-1}\mathbb{1}(\Lambda_{\max} \geq 2)$, which is satisfied when $\alpha > \frac{1}{k}$ if $\Lambda_{\max} \geq 2$ (and for all $\alpha \in (0, 1)$ if $\Lambda_{\max} = 1$). In that case, $e(p_{k-1}, OPT(G_{\{d_{k-1}\}})) = \alpha^{k-1} < \alpha^k = e(p_{k-1}, OPT(G))$, violating SP.
2. Assume instead that $C_{\max} = 1$, meaning that no cycles are allowed, and $\Lambda_{\max} \geq 3$. Consider a market $G = (P, D, R, T, s)$, shown on Figure 6b, where $s(d,p) = 1 - \epsilon$ for all $(d,p) \in T$ for some very small $\epsilon > 0$. For a small enough ϵ, $OPT(G) = \{(d_0, p_3), (d_3, p_4), (d_4, p_2), (d_2^2, p_1)\}$ while $OPT(G_{\{d_2^2\}}) = \{(d_0, p_1), (d_1^1, p_2)\}$. However, $e(p_2, OPT(G)) = (1-\epsilon)^3$ while $e(p_2, OPT(G_{\{d_2^2\}})) = (1-\epsilon)^2$, which violates SP. □

Theorem 5. *Consider the greedy mechanism $BADTB$ which tie-breaks so that the final objective function value is as large as possible (meaning, $obj(BADTB(G)) \geq obj(M)$ for all matchings M that can be obtained by greedy mechanisms on G). $BADTB$ is not SP.*

Proof. We give two examples of $BADTB$ failing SP. The key observation in both of them is that if an optimal solution can be built greedily, then $BADTB$ would output that optimal solution.

1. The proof of Theorem 4 on the market from Figure 6a works for $BADTB$ since in both G and $G_{d_{k-1}}$ the optimal solution can be generated greedily.

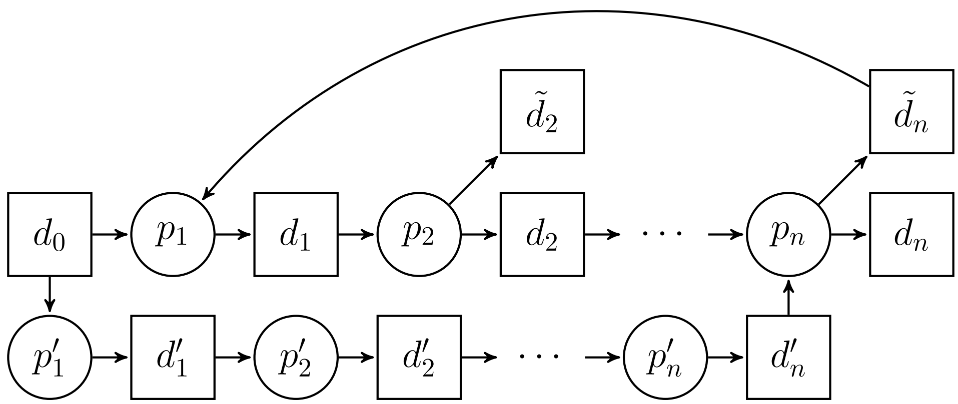

Fig. 7: Market for Theorem 5

2. Consider the market G shown in Figure 7, where for every $i = 1, \ldots, n$, there exist patients p_i and p'_i, as well as their respective proxy donors d_i and d'_i. In addition, there are three additional donors: an altruistic donor d_0, a proxy donor $\tilde{d}_2 \in p_2^*$ and a proxy donor $\tilde{d}_n \in p_n^*$. For every $1 \leq i \leq n-1$, d_i is compatible with p_{i+1} and d'_i is compatible with p'_{i+1}. In addition, d_0 is compatible with p_1 and p'_1, and $\tilde{d}_n$ is compatible with p_1. Assume $s(d,p) = 1-\epsilon$ for all donations in the market, where $\epsilon > 0$ is very small, and assume $n > C_{\max}$. Denote $BOTTOM = \{(d'_i, p'_{i+1}) : i = 1, \ldots, n-1\}$, $TOP = \{(d_i, p_{i+1}) : i = 1, \ldots, n-2\}$. Using the fact that the optimal solutions in both G and $G_{\{\tilde{d}_n\}}$ are unique and can be generated greedily, $BADTB(G) = \{(d_0, p'_1)\} \cup BOTTOM \cup \{(d'_n, p_n)\} \cup \{(\tilde{d}_n, p_1)\} \cup TOP$, and $BADTB(G_{\{\tilde{d}_n\}}) = \{(d_0, p_1)\} \cup TOP \cup \{(d_{n-1}, p_n\}$. However, $e(p_n, BADTB(G)) = (1-\epsilon)^{n+1}$ and $e(p_n, BADTB(G_{\{\tilde{d}_n\}})) = (1-\epsilon)^n$, violating SP. □

Lemma 1. *Let $G = (P, D, T, R, s)$ be a uniform market. Assume that $s(d,p) = \alpha < \min\{\frac{1}{\Lambda_{\max}}, \frac{1}{\sqrt{e}}\}$ for all $(d,p) \in T$. Let $P_1^M = \{p \in P : |r(p, M)| = 1\}$ be the set of patients matched at distance 1 for any matching M. Then, $|P_1^{\mathbb{S}(G)}| \geq |P_1^M|$ for all matchings M in G.*

Proof. For a matching M, define $M_A = \{(d,p) \in M : d \in A\}$. Let $\tilde{M}$ be a matching where $|P_1^{\tilde{M}}| > |P_1^{\mathbb{S}(G)}|$. Let $\{(d_1, p_1), \ldots, (d_k, p_k)\} = \mathbb{S}(G)_A$ (if the set is empty then $k = 0$). Then, define $M_0 = \tilde{M}_A$, and get M_i from M_{i-1} as follows:

1. If d_i and p_i are both unmatched by M_{i-1}, $M_i = M_{i-1} \cup \{(d_i, p_i)\}$.
2. If d_i and p_i are both matched by M_{i-1}, $M_i = M_{i-1}$.
3. If d_i is matched by M_{i-1} to $p \in P$ and p_i is unmatched by M_{i-1}, $M_i = M_{i-1} \cup \{(d_i, p_i)\} - \{(d_i, p)\}$.
4. If p_i is matched by M_{i-1} to $d \in A$ and d_i is unmatched by M_{i-1}, $M_i = M_{i-1} \cup \{(d_i, p_i)\} - \{(d, p_i)\}$.

When going from M_{i-1} to M_i, we remove at most one transplant involving an altruistic donor from M_{i-1}, and if we do, we add such a transplant (d_i, p_i) as well. Therefore, $|P_1^{M_i}| \geq |P_1^{M_{i-1}}|$. Thus, $|P_1^{M_k}| \geq |P_1^{M_0}| = |P_1^{\tilde{M}}| > |P_1^{\mathbb{S}(G)}|$. Furthermore, all patients matched at distance 1 by $\mathbb{S}(G)$ are also matched at distance 1 by M_k, so $P_1^{\mathbb{S}(G)} \subseteq P_1^{M_k}$, and since also $|P_1^{\mathbb{S}(G)}| < |P_1^{M_k}|$, then $P_1^{\mathbb{S}(G)} \subset P_1^{M_k}$. Therefore, the number of donations induced at distance 1 by $\mathbb{S}(G)$ is strictly less than the number of donations induced at distance 1 by M_k. But $\mathbb{S}(G)$ maximizes the number of donations induced at distance 1 by definition. Contradiction.

Theorem 6. *Let $G = (P, D, T, R, s)$ be a uniform market. Assume that $s(d, p) = \alpha < \min\{\frac{1}{\Lambda_{\max}}, \frac{1}{\sqrt{e}}\}$ for all $(d, p) \in T$. The approx. ratio of $\mathbb{S}$ on G is upper bounded by* $\max\{\frac{1}{1+\Lambda_{\max}\alpha} + \frac{\Lambda_{\max}^2\alpha}{1-\Lambda_{\max}^2\alpha^2} + 2\alpha^2\Lambda_{\max}^2\frac{1-\alpha}{1-\Lambda_{\max}^2\alpha^2}, \frac{\Lambda_{\max}}{1-\Lambda_{\max}\alpha}(1-\alpha)(1+C_{\max}\alpha)\}$.

Proof. We use the notation from the proof of Theorem 2. Let us set

$$B_1 = (\frac{\Lambda_{\max}\alpha}{1-\Lambda_{\max}\alpha} + 2\alpha^2\Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha})\beta_1 + \alpha^{C_{\max}}\Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha}\beta_{C_{\max}} + \sum_{i=2}^{C_{\max}-1}\Lambda_{\max}\frac{1-\alpha}{1-\Lambda_{\max}\alpha}(\alpha^i + (i+1)\alpha^{i+1})\beta_i$$

and

$$B_2 = \sum_{i=1}^{C_{\max}} \alpha^i\beta_i$$

We have established that $\overline{obj}(OPT(G)) \leq B_1$ and $\overline{obj}(\mathbb{S}(G)) \geq B_2$. Therefore, $obj(OPT(G)) \leq |A(G)| + B_1$ and $obj(\mathbb{S}(G)) \geq |A(G)| + B_2$, and we get the bound $\frac{obj(OPT(G))}{obj(\mathbb{S}(G))} \leq \frac{|A(G)|+B_1}{|A(G)|+B_2}$. Since $\frac{B_1}{B_2} \geq 1$ and both B_1 and B_2 are positive, it follows that $\frac{|A(G)|+B_1}{|A(G)|+B_2}$ decreases as $|A(G)|$ increases for $|A(G)| \geq 0$. Furthermore, we know that $|A(G)| \geq |P_1^{\mathbb{S}(G)}| \geq \frac{\beta_1}{\Lambda_{\max}}$: the fact that each patient in $P_1^{\mathbb{S}(G)}$ receives an altruistic donation justifies the first inequality, and the fact that each patient in $P_1^{\mathbb{S}(G)}$ induces at most $\Lambda_{\max}$ donations at distance 1 justifies the second inequality. Thus, we get $\frac{obj(OPT(G))}{obj(\mathbb{S}(G))} \leq \frac{\frac{\beta_1}{\Lambda_{\max}}+B_1}{\frac{\beta_1}{\Lambda_{\max}}+B_2}$. From this point, the analysis is identical to the non-adjusted case, with the coefficient of β_1 in both the numerator and the denominator increased by $\frac{1}{\Lambda_{\max}}$. □

Appendix C Solving SuperGreedy's Integer Program

While it is impractical to feed our IP for SuperGreedy into a solver directly, we can solve it iteratively. We will specify an ordering $\succ$ over T which will determine an ordering over 2^T as discussed in Section 4. Define $\succ$ as follows: for $(d_i, p_j), (d_{i'}, p_{j'}) \in T$, if either $j > j'$ or both $j = j'$ and $i > i'$, then $(d_i, p_j) \succ (d'_i, p'_j)$.

Informally, the iterative algorithm works as follows. In iteration l, we solve two IPs. First IP_l^1 determines which patients will receive a kidney at distance l, and then IP_l^2 breaks remaining ties. The advantage of this method is that we do not actually have to use large coefficients. IP_l^1 optimizes $\sum_{i=1}^{|P|} (\beta |p_i^*| + 2^i) p_{i,l}$: only the first power of β is involved, and it is actually sufficient to have $\beta = |D| + 2^{|P|}$, which is significantly smaller than $|D| + 2^{|T|}$. IP_l^2 only needs to consider a small subset of edges: edges from donors matched at distance $l-1$ who have yet to donate (in iteration $l-1$) to patients determined to match in IP_l^1, and edges in cycles including only patients determined to match in IP_l^1. Formally, we define IP_l^1 and IP_l^2, as well as their respective solutions $(p^{l,1}, d^{l,1})$ and $(p^{l,2}, d^{l,2})$ recursively. As the base case, $d_{i,j,0}^{0,2} = 1$ if d_i is altruistic, and all other variables in $d^{0,2}$ as well as all variables in $p^{0,2}$ are 0. We also define $DONORS_0 = A$. For iterations $l = 1, \ldots, L_{\max}$, do the following:

1. To get IP_l^1, modify our original IP as follows:
   (a) Set $p_{i,l'} = 0$ and $d_{j,k,l'} = 0$ for all $l' > l$ (and all i, j, k s.t. the variables exist).
   (b) Set $p_{i,l'} = 1$ when $p_{i,l'}^{l-1,2} = 1$ and set $d_{j,k,l'} = 1$ when $d_{j,k,l'}^{l-1,2} = 1$ for all $l' < l$ (and all i, j, k s.t. the variables exist).
   (c) Set the objective function to be $\sum_{i=1}^{|P|} (\beta |p_i^*| + 2^i) p_{i,l}$.
2. Solve IP_l^1 to get $(p^{l,1}, d^{l,1})$. Let the set of patients matched at distance l be $PATIENTS_l = \{p_i : p_{i,l}^{l,1} = 1\}$.
3. To get IP_l^2, modify our original IP as follows:
   (a) Set $p_{i,l'} = 0$ and $d_{j,k,l'} = 0$ for all $l' > l$ (and all i, j, k s.t. the variables exist).
   (b) Set $p_{i,l'} = 1$ when $p_{i,l'}^{l,1} = 1$ for all $l' \leq l$ and $d_{j,k,l'} = 1$ when $d_{j,k,l'}^{l-1,2} = 1$ for all $l' < l$ (and all i, j, k s.t. the variables exist).
   (c) Let $CYCLES_l$ be the set of all edges from T that are in cycles of size l that consist entirely of patients in $PATIENTS_l$. Let $EXTENSION_l = \{(d_i, p_j) \in T : d_i \in DONORS_{l-1}, p_j \in PATIENTS_l\}$. Define $B_l = CYCLES_l \cup EXTENSION_l$ and let $\mu_l(d, p) = k - 1$ if (d, p) is the k-th largest edge in B_l according to $\succ$.
   (d) Set the objective function to be $\sum_{i,j:(d_i,p_j) \in B_l} 2^{\mu_l(d_i,p_j)} d_{i,j,l}$
4. Solve IP_l^2 to get $(p^{l,2}, d^{l,2})$.
5. Set $DONORS_l = \cup_{i:p_i^l=1} p_i^* - \{d_i : \exists j \text{ s.t. } d_{i,j,l}^l = 1\}$ to be the set of all satisfied donors at distance l who have not donated to a patient at distance l (that is, who have not donated through a cycle).

Appendix D Additional Simulation Results for the Non-Adjusted Objective

This appendix includes the simulation results for the non-adjusted objective. We note that Dickerson et al. [6] also simulate OPT when $\Lambda_{\max} = 1$ and $\alpha = 0.3$, and report a significantly smaller number of transplants: this is to be expected, because they do not assume redirection of donors in indirectly-but-not-directly canceled transplants to the waiting list (in other words, their objective function is different than obj).

	approx. ratio				$obj(\mathbb{S}(G))$				$obj(OPT(G))$			
α \ $\Lambda_{\max}$	1	2	3	4	1	2	3	4	1	2	3	4
0.1	1.058	1.009	1.006	1.004	2.221	2.263	2.267	2.268	2.272	2.275	2.275	2.275
0.2	1.08	1.019	1.015	1.012	3.031	3.168	3.18	3.186	3.196	3.219	3.221	3.222
0.3	1.098	1.029	1.024	1.022	4.294	4.596	4.622	4.636	4.647	4.724	4.731	4.733

Table 3: Mean simulation results for $\mathbb{S}$. The mean is taken over all 52 markets.

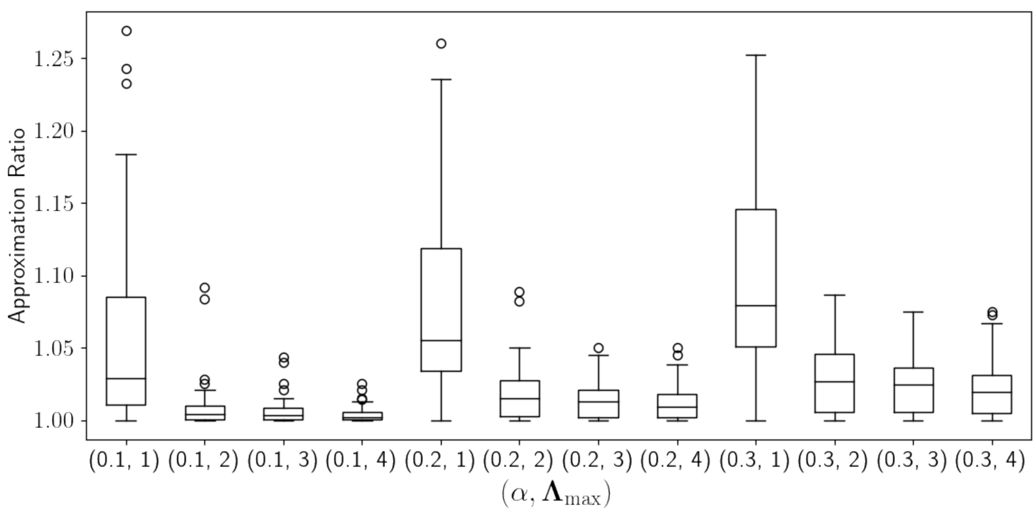

Fig. 8: Box plot for the approximation ratio of $\mathbb{S}$, showing quartiles q_1-q_3. Empty circles are outliers (points outside $[q_1 - 1.5(q_3 - q_1), q_3 + 1.5(q_3 - q_1)]$).