Published in final edited form as:

Stat Biosci. 2014 May 1; 6(1): 85–104. doi:10.1007/s12561-013-9082-0.

Optimal Decisions for Organ Exchanges in a Kidney Paired Donation Program

Yijiang Li

Department of Biostatistics, University of Michigan, Ann Arbor, MI 48109, USA

Peter X.-K. Song

Department of Biostatistics, University of Michigan, Ann Arbor, MI 48109, USA

Yan Zhou

Department of Biostatistics, University of Michigan, Ann Arbor, MI 48109, USA

Alan B. Leichtman

Department of Internal Medicine, University of Michigan, Ann Arbor, MI 48109, USA

Michael A. Rees

Department of Urology, University of Toledo Medical Center, Toledo, Ohio 43614, USA

John D. Kalbfleisch

Department of Biostatistics, University of Michigan, Ann Arbor, MI 48109, USA

Abstract

The traditional concept of barter exchange in economics has been extended in the modern era to the area of living-donor kidney transplantation, where one incompatible donor-candidate pair is matched to another pair with a complementary incompatibility, such that the donor from one pair gives an organ to a compatible candidate in the other pair and vice versa. Kidney paired donation (KPD) programs provide a unique and important platform for living incompatible donor-candidate pairs to exchange organs in order to achieve mutual benefit. In this paper, we propose novel organ allocation strategies to arrange kidney exchanges under uncertainties with advantages, including (i) allowance for a general utility-based evaluation of potential kidney transplants and an explicit consideration of stochastic features inherent in a KPD program; and (ii) exploitation of possible alternative exchanges when the originally planned allocation cannot be fully executed. This allocation strategy is implemented using an integer programming (IP) formulation, and its implication is assessed via a data-based simulation system by tracking an evolving KPD program over a series of match runs. Extensive simulation studies are provided to illustrate our proposed approach.

Keywords

Decision under uncertainty; Expected utility; Integer programming; Probabilistic modeling; Organ exchanges

1 Introduction

For patients with end-stage renal disease (ESRD), successful kidney transplantation provides a considerably better quality of life and improved survival as compared with dialysis (Evans et al, 1985; Russell et al, 1992; Wolfe et al, 1999). Kidney transplantation is also cost-effective as compared to continuing on dialysis (Laupacis et al, 1996). However, due to a growing demand and limited supply of deceased-donor kidneys, many patients who need a transplant must be placed on long waiting lists. According to the Organ Procurement and Transplantation Network (OPTN), as of December 2011, more than 90,000 kidney transplant candidates in the U.S. were on a waiting list. From 2008 to 2010, more than 33,000 candidates per year in the U.S. were added to the list, whereas only about 11,000 and 5,000 per year actually received a kidney transplant from a deceased donor and a living donor respectively; during that same period of time, about 4,700 patients died and another 1,900 became too sick to transplant each year while waiting for a kidney transplant. In response to this shortage of deceased-donor organs, candidates have increasingly undergone living-donor transplants. Moreover, living-donor transplants have the advantage of both short and long term higher graft survival rates than deceased-donor transplants (Terasaki et al, 1995; Hariharan et al, 2000).

A major barrier to living-donor kidney transplantation is the unfortunate fact that willing donors, even related living donors, are often incompatible with their intended recipients, because of ABO blood type incompatibility and/or antibodies against one or more of the donors' Human Leukocyte Antigens (HLA). With respect to blood type compatibility, A and B donors can donate to candidates of the same blood type or of type AB; AB donors can donate only to AB candidates; and O donors, known as universal donors, can donate to candidates of any blood types. The second form of incompatibility results from the presence of pre-existing antibodies in the blood of a candidate against the HLA antigens of a potential donor. Both forms of incompatibility can lead to rapid rejection of the transplanted organ and thus prohibit transplant.

Some of these incompatible donors, however, might be compatible with other candidates. In 1986, the idea of kidney paired donations (KPD) was set forth by Felix Rapaport. Rapaport (1986) and later Ross et al (1997) envisioned that two candidates with willing, but incompatible, donors could solve each other's problem by simultaneously exchanging the two donors' kidneys; see Figure 1-A for an example of two-way exchanges. More complex exchanges involving three or more pairs are also possible, as schematically illustrated in Figure 1-B.

While three-way or higher exchange cycles increase the chance of finding compatible matches, most individual transplant centers restrict exchanges to at most three ways mainly for two reasons. First, all operations on an exchange cycle must be performed simultaneously to avoid the risk that one of the donors may renege. This avoids a situation where a donor withdraws his or her commitment after the other donor has undergone nephrectomy and donated to the candidate associated with this reneging donor. This requirement of performing operations simultaneously creates substantial logistical difficulties of scheduling, for example, eight surgeons and eight operating rooms at the same

time for a four-way exchange. In addition, the greater the length of an exchange cycle, the less likely the potential transplants will actually occur since if any of the proposed transplants cannot proceed the whole cycle would collapse.

Despite these logistical difficulties, more and more KPD programs have recently been established, with the mission of promoting mutually beneficial organ exchanges among incompatible donor-candidate pairs. Regional programs in the U.S. include, for example, the New England Programs for Kidney Exchange, the Paired Kidney Exchange Program at Johns Hopkins Medical Center, the Alliance for Paired Donation (APD), the National Kidney Registry (NKR) Program, and the University of Michigan Paired Kidney Exchange Program. Internationally, KPD programs exist in the Netherlands (Keizer et al, 2005) and South Korea (Park et al, 1999). Additional developments include a recent OPTN approval of a national KPD pilot program.

Researchers from several disciplines, collaborating with transplant surgeons, have become heavily involved in the management of KPD programs. For instance, economists have modeled and analyzed kidney exchanges using game-theoretic approaches. Roth et al (2004) organized donor-candidate pairs as a “housing market”, a concept first proposed by Shapley and Scarf (1974), and produced an efficient organ exchange mechanism using Gale's Top Trading Cycles (TTC) algorithm. In the case where only two-way exchanges are considered, Roth et al (2005) and Segev et al (2005) applied a maximum cardinality matching algorithm (Edmonds, 1965) to select exchanges that allow the maximum number of transplants. More generally, organizing and optimizing kidney exchanges can be formulated as an integer programming (IP) problem. This novel IP-based formulation was first proposed by Roth et al (2007), where each potential transplant was assigned equal weight, resulting in an allocation strategy that enables the greatest number of transplants to be potentially implemented. Abraham et al (2007) adopted a more flexible weight assignment in this IP-based formulation and further developed an algorithm to reduce the computational complexity of managing large KPD programs. The efficiency of their proposed algorithm, however, relies on using equal edge weights and assuming only the blood type incompatibility.

This paper extends and improves upon the research described above. Our proposed kidney allocation strategy is innovative in several respects. First, it allows for a quality-oriented evaluation of a potential kidney transplant through medical-outcome-based utilities such as post-transplant graft and recipient survival. In addition, it explicitly takes into consideration the probability that a predicted compatible transplant will result in an actual transplant operation. Third, our approach allows for one or more contingency allocations should the originally planned exchanges fail to be executed. We also suggest a data-based simulation system for studying an evolving KPD program to evaluate different kidney allocation strategies, to compare their impact on performance outcomes, and to assess different choices of utility assignments. The knowledge learned from such simulation studies should provide invaluable guidance in implementing an actual clinical KPD allocation system.

The rest of this article is organized as follows: in Section 2, we introduce the representation and formulation of a KPD program and present a procedure for arranging kidney exchanges

according to the maximum utility cycle-based allocation. In Section 3, we explore some important issues and features in a KPD program that have not been addressed by previous studies and further propose a kidney allocation strategy based on the maximum expected-utility set-based allocation. In Section 4, we present a data-based simulation system for studying an evolving KPD program. Section 5 illustrates the application of the proposed simulation system and reports simulation results for comparing different allocation strategies. We conclude with some discussion in Section 6.

2 Problem formulation and the maximum utility cycle-based allocation

In this section, we first present a graph representation of a KPD program and then describe two IP formulations for organizing kidney exchanges. We also introduce a procedure for arranging organ exchanges according to the maximum utility cycle-based allocation.

2.1 Graph representation

We represent a KPD program as a *directed graph*, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the *vertex set*, $\mathcal{V} \equiv \mathcal{V}(\mathcal{G}) = \{1, 2, \dots, n\}$, is the set of n incompatible donor-candidate pairs, and the *edge set*, $\mathcal{E} \equiv \mathcal{E}(\mathcal{G})$ is a binary relation on $\mathcal{V}$, consisting of ordered pairs of vertices in $\mathcal{V}$. An edge from i to j , denoted as (i, j) , indicates that the donor of pair i is predicted to be compatible with the candidate of pair j . This predicted compatibility is based on a *virtual crossmatch*, which involves computer cross-checking for blood type compatibility as well as comparing candidate antibodies against donor HLA antigens. Before a predicted compatible transplant can be considered for an actual surgical operation, it must be confirmed by a labor intensive *laboratory crossmatch* to assure histocompatibility, which involves incubating the serum of a candidate with the white blood cells of a prospective donor. It is worth noting that prior research on KPD has not taken into consideration this uncertainty, and instead has proceeded as though a negative virtual crossmatch would guarantee, if selected, an actual transplant operation.

In such a directed graph, an *exchange cycle* of length k (or a k -way exchange cycle), $k \geq 2$, is defined as a sequence of distinct vertices in $\mathcal{V}$, denoted as $\langle c_1, c_2, \dots, c_k \rangle$, such that $(c_k, c_1) \in \mathcal{E}$ and $(c_{j-1}, c_j) \in \mathcal{E}$, $\forall j = 2, 3, \dots, k$. For an exchange cycle $C = \langle c_1, c_2, \dots, c_k \rangle$, we denote its *vertex set* as $\mathcal{V}(C) = \{c_1, c_2, \dots, c_k\}$ and its *edge set* as $\mathcal{E}(C) = \{(c_k, c_1), (c_{j-1}, c_j), j=2, 3, \dots, k\}$. In Figure 2 are shown one two-way exchange cycle, $\langle 2, 4 \rangle$, and two three-way exchange cycles, $\langle 1, 2, 3 \rangle$ and $\langle 2, 4, 5 \rangle$. Exchange cycles form a *disjoint* collection if their corresponding vertex sets are disjoint. A *cycle-based allocation* for a KPD program $\mathcal{G}$ is defined as a disjoint collection of exchange cycles, and further denoted as $\mathcal{C}(\mathcal{G})$. An alternative *set-based allocation*, denoted as $\mathcal{S}(\mathcal{G})$, will be introduced and discussed in Section 3.4.

For an edge $e \equiv (i, j) \in \mathcal{E}$, let u_e or u_{ij} denote a nonnegative *utility* associated with a kidney transplant from the donor in pair i to the candidate in pair j ; see more discussion on assigning edge utilities in Section 3.1. The utility of an exchange cycle, C , is the sum of its edge utilities, i.e. $U_C = \sum_{e \in \mathcal{E}(C)} u_e$, and the utility of a cycle-based allocation, $\mathcal{C}(\mathcal{G})$, is the sum of the utilities of its cycles, i.e. $\sum_{C \in \mathcal{C}(\mathcal{G})} U_C$. In this setting, the optimal cycle-based

allocation which we denote by $\mathcal{C}^*(\mathcal{G})$ is the one with the maximal utility. Roth et al (2007) proposed two IP formulations to determine $\mathcal{C}^*(\mathcal{G})$. One formulation encodes each exchange cycle as a decision variable, and the other one encodes each edge as a decision variable. We first look at the *cycle formulation* in Section 2.2. The *edge formulation* is presented in the Appendix.

2.2 Cycle formulation

Let $\mathcal{C}_k$ be the set of exchange cycles with lengths at most k . For each $i \in \mathcal{V}$, $\mathcal{C}_k(i)$ denotes the exchange cycles in $\mathcal{C}_k$ that involve pair i , i.e. $\mathcal{C}_k(i) = \{C \in \mathcal{C}_k : i \in \mathcal{V}(C)\}$. Define a *decision variable* Y_C for each cycle $C \in \mathcal{C}_k$, such that Y_C is 1 if organ exchanges indicated by C will be arranged, and $Y_C = 0$ otherwise. The problem of selecting $\mathcal{C}^*(\mathcal{G})$ can then be formulated as the following IP problem,

$$\begin{aligned} & \max_{\{Y_C\}_{C \in \mathcal{C}_k}} \sum Y_C U_C, \\ & \text{subject to } \sum_{C \in \mathcal{C}_k(i)} Y_C \leq 1, \forall i \in \mathcal{V}. \quad (1) \end{aligned}$$

Note that the constraint (1) codifies the fact that each pair cannot be allowed in more than one exchange cycle simultaneously. Each feasible solution to this IP problem corresponds to one $\mathcal{C}(\mathcal{G})$, and the optimal solution corresponds to $\mathcal{C}^*(\mathcal{G})$. As a special case in this IP problem, when the utility of an exchange cycle is defined as its length, i.e. $U_C = |\mathcal{V}(C)|$, the resulting objective is to maximize the total number of transplants. Most KPD studies have focused on this simplified utility assignment; see more discussion in Section 3.1.

Solving this IP optimization problem in general requires increased computation time as k grows, though in practice most KPD programs restrict k to be three or smaller due to logistical difficulties. Some particular values of k are worth mentioning. When k is equal to two, the optimization problem could be solved in polynomial time using a maximum weighted matching algorithm, which is an extended version of Edmonds' classical maximal cardinality matching algorithm. When k is equal to $|\mathcal{V}|$, i.e. no restriction is placed on the length of an exchange cycle, this problem can be efficiently solved by finding a maximum weight perfect matching arising from an edge formulation of the optimization problem (see the Appendix for a brief discussion). As noted in Roth et al (2007), this IP problem is NP-hard for k greater than two but less than $|\mathcal{V}|$, which poses associated computational challenges when the number of incompatible pairs is large.

2.3 The maximum utility cycle-based allocation

We summarize below a procedure for arranging kidney exchanges according to the maximum utility cycle-based allocation. Given a KPD program $\mathcal{G} = (\mathcal{V}, \mathcal{E})$,

- (i) Define $u: \mathcal{E} \rightarrow \mathbb{R}^+$, where u_e is the utility assigned to an edge $e \in \mathcal{E}$.
- (ii) Enumerate $\mathcal{C}_k$ and find U_C for each $C \in \mathcal{C}_k$.
- (iii) Determine $\mathcal{C}^*(\mathcal{G})$ by solving an IP problem.

- (iv) Arrange kidney exchanges according to $\mathcal{C}^*(\mathcal{G})$.

In (ii), we enumerate $\mathcal{C}_k$ by developing an algorithm based on depth-first search (DFS), which essentially prioritizes the direction of search always to offspring vertices and then to sibling vertices. The procedure outlined above is essentially the one proposed by Roth et al (2007); this approach is used in all the simulations of Section 5 in order to provide a benchmark that corresponds to the basic method underlying approaches currently in use.

In terms of the optimization problem, there is always advantage to increasing k since the possible exchanges for $k = 3$, for example, includes those exchanges for $k = 2$. To some degree, this illustrates the limitations of the above allocation strategy. In practice, large values of k do not lead to realistic or implementable solutions. For example, if $k = 5$, then the optimal solution $\mathcal{C}^*(\mathcal{G})$ obtained in step (iii) would typically include a cycle of five simultaneous transplants. Two problems arise: first, the resources to carry out these five simultaneous transplants would typically not be available since it would require 10 simultaneous operations. Second, however, there is a very high chance that such a cycle would break down due to the friction (i.e. the chance that planned transplants fail to proceed) in the system. Thus, even though $k = 4$ dominates $k = 3$, for example, the actual claimed utility based on an optimization with $k = 4$ could be less than that for $k = 3$, or similarly for $k = 3$ versus $k = 2$. Later in the paper, we take into account the uncertainty in completing identified transplants. In this case, there is still a theoretical advantage to increasing k , but the advantage become very slight for large k since we account for the likely failure of the chosen long cycle.

Based on the above considerations, we choose $k = 3$ as the typical constraint. On the one hand, allocations based on $k = 2$ may yield an insufficient number of options, whereas cycles of size 3 are logistically possible. As a result, it seems reasonable to consider two and three-way exchange cycles ($k = 3$) as the primary setup in this paper.

3 Optimal kidney allocation in a KPD program

In Section 2, we have introduced a graph representation of a KPD program and presented an allocation strategy based on $\mathcal{C}^*(\mathcal{G})$. In this section, we focus on issues in managing KPD programs that have not been previously addressed and further explore two alternative organ allocation strategies.

3.1 Utilities and uncertainties in a KPD program

Much of the prior work has focused on a simplified edge utility, namely $u_e = 1, \forall e \in \mathcal{E}$. As a result, the utility U_C of an exchange cycle, C , equals its length, and the objective function in the IP problem is the total number of arrangeable transplants. Some recent developments, however, have emerged in assigning rule-based utilities determined by pair attributes, such as degree of candidate sensitization or the age difference between the candidate and donor, so that they better evaluate potential transplants. For example, in an operational guideline recently posted by the U.S. national KPD pilot program, each potential transplant is initially assigned a base utility of 200 points. Extra points are added as bonuses to edges that, say, have zero antigen mismatches, or that involve a donor and a candidate in the same transplant

center; on the other hand, a certain number of points are deducted, for example, when a donor has one or more of the candidate's other antibody specificities.

For potential transplants in a KPD program, it is also desirable to examine outcome-based utilities derived from medical outcomes such as the expected graft or patient survival. To this end, we propose to associate with edge, (i, j) , a utility quantifying the medical outcome from a potential kidney transplant involving the donor in pair i and the candidate in pair j . This outcome could be, for example, graft survival, post-transplant recipient survival, or the incremental years of recipient life that would accrue with a kidney transplant as opposed to remaining on dialysis. By incorporating this outcome-based utility, we are able to evaluate and compare competing kidney allocations with a quality-oriented view, and therefore provide kidney transplant candidates with organs that are not only compatible, but that could potentially maximize opportunities to achieve good quality of life after transplants.

Prior research on KPD has implicitly assumed that a predicted compatible transplant, if attempted, would yield an actual transplant operation. In reality, predicted compatible transplants have to be confirmed by laboratory crossmatches, and hence may or may not lead to actual transplant operations. This uncertainty is a necessary ingredient since laboratory crossmatches cannot be undertaken on all possibly compatible donors and candidates due to labor and resource limitations. Further, even if the laboratory crossmatch is negative (non-reactive), an actual transplant operation may not occur due to other friction including, for example, pair preference, illness, and death. Throughout the rest of the paper, we use the term “*is viable*” to indicate that the edge actually results, if chosen, in a completed transplant operation. An exchange cycle is viable if each of its edges is viable.

Ignoring this uncertainty can result in a situation in which long exchange cycles are evaluated more favorably than short ones, despite the fact that longer cycles are much less likely to be implemented. To partially incorporate this uncertainty into the arrangement of kidney exchanges, the operational guidelines for the national KPD pilot program proposed a deduction of 30 points for a three-way exchange cycle, but not for an exchange cycle that is two-way.

To further address this stochastic feature inherent in a KPD program, we associate with each edge a probability corresponding to the chance of that edge being viable. This probability depends on factors such as the number of mismatched antigens, the donor-candidate travel distance, and age difference. For each $e \in \mathcal{E}$ let X_e be a Bernoulli random variable with $X_e = 1$ if e is viable, and $X_e = 0$ otherwise. By defining p_e as $P(X_e = 1)$, an exchange cycle, C , is then viable with probability equal to $P_C \equiv \prod_{e \in \mathcal{E}(C)} p_e$, under the assumption that edges in an exchange cycle have an independence relationship. More formally, this assumption is regarded as having a collection of independent random variables, $\{X_e, e \in \mathcal{E}(C)\}$. Though this independence relationship can be further relaxed, it is a reasonable assumption when the rate of pair withdrawal (due to factors such as pregnancy, illness, or death) is not high. In the rest of this paper, we proceed by assuming in general that $X_e, e \in \mathcal{E}$ are independent. It is worth pointing out that a set of common factors could affect both the edge probability and utility.

As we should discuss later in Section 3.3, incorporating these probabilities into the management of a KPD program also opens up the opportunity to identify and implement possible alternative fall-back options if the originally planned exchanges cannot be fully executed.

In this paper, we focus on the development of an allocation strategy for the management of a KPD program with certain desired properties that will be introduced in the remaining of this section. Our attention, however, is not on the estimation of utilities and probabilities associated with potential transplants, although Section 3.5 presents a discussion of some aspects of this.

3.2 The maximum expected-utility cycle-based allocation

One natural way to evaluate an exchange cycle is to use its *expected utility*, namely $EU_C \equiv U_C P_C$. Utilizing expected utilities in the management of a KPD program provides a way to incorporate uncertainty into the selection of kidney exchanges. For example, expected utility recognizes that a longer exchange cycle has a smaller chance of being viable, which counterbalances the fact that such a cycle might potentially contribute greater utilities and allow for more transplants.

The expected utility of a cycle-based allocation, $\mathcal{C}(\mathcal{G})$, is the sum of the expected utilities of its exchange cycles, i.e. $\sum_{C \in \mathcal{C}(\mathcal{G})} EU_C$. Among all cycle-based allocations, the one with the largest expected utility is the *maximum expected-utility cycle-based* allocation, which we denote as $\overline{\mathcal{C}}^*(\mathcal{G})$. The following procedure generates $\overline{\mathcal{C}}^*(\mathcal{G})$ in a KPD program $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and arranges exchanges accordingly.

- (i) Define $u: \mathcal{E} \rightarrow \mathbb{R}^+$, where u_e is the utility of an edge $e \in \mathcal{E}$ and $p: \mathcal{E} \rightarrow [0, 1]$, where p_e is the probability that e is viable.
- (ii) Enumerate $\mathcal{C}_k$ and calculate EU_C for each $C \in \mathcal{C}_k$.
- (iii) Find $\overline{\mathcal{C}}^*(\mathcal{G})$ by an IP-based approach as in Section 2, with U_C replaced by EU_C .
- (iv) Arrange kidney exchanges according to $\overline{\mathcal{C}}^*(\mathcal{G})$.

In (iii), the cycle formulation generalizes in a straightforward manner, but the edge formulation cannot be extended; this is because the edge formulation represents each edge as a separate decision variable and hence cannot describe EU_C . Notice that the above procedure and the one in Section 2.3 are both *fixed* in that they do not specify how to proceed if a chosen exchange cycle is not viable.

3.3 Contingency plans

Let us begin with a motivating example. In a small KPD program as represented in Figure 2, there are two three-way exchange cycles, $C_1 = \langle 1, 2, 3 \rangle$ and $C_2 = \langle 2, 4, 5 \rangle$, and one two-way exchange cycle, $C_3 = \langle 2, 4 \rangle$. We assume that $EU_{C_1} > EU_{C_2} > EU_{C_3}$, and thus conclude $\overline{\mathcal{C}}^*(\mathcal{G})$ is $\{C_1\}$. A further examination, however, reveals that pairs in C_3 are part of the pairs in C_2 , i.e. $\mathcal{V}(C_3) = \{2, 4\} \subset \mathcal{V}(C_2) = \{2, 4, 5\}$. This observation might suggest that, when

C_2 is selected but could not be completed because of problems in either the (4, 5) or (5, 2), the two-way exchange cycle C_3 could still be selected. Therefore, the contribution from this back-up exchange cycle C_3 would add some extra value to the exchange cycle C_2 . Does this extra value make $\{C_2\}$ a preferred allocation to $\{C_1\}$? How should we evaluate $\{C_2\}$ so as to correctly recognize $\{C_3\}$ as a possible back-up allocation? To address these questions in this specific example and address other related issues in general, we propose the following.

First, we give two definitions from graph theory: (i) a graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ is a *subgraph* of $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, if $\mathcal{V}' \subset \mathcal{V}$ and $\mathcal{E}' \subset \mathcal{E}$, and (ii) a graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ is a *induced subgraph* of $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, if $\mathcal{G}'$ is subgraph of $\mathcal{G}$ and in addition, $\mathcal{E}' = \{(u, v) \in \mathcal{E} : u, v \in \mathcal{V}'\}$. In the context of a KPD program, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, we then define an *exchange set* as an induced subgraph of $\mathcal{G}$, denoted by $S = (\mathcal{V}', \mathcal{E}')$, with the requirement that it admits at least one exchange cycle of length $|\mathcal{V}'|$, where $|\mathcal{V}'|$ is defined as the *size* of this exchange set. Figure 2 illustrates two exchange sets of size three, S_1 and S_2 , and one exchange set of size two, S_3 .

By definition, an exchange set would allow for one or more cycle-based allocations; if there is indeed more than one and not all of the exchange cycles in the first attempted allocation are viable, we might still have the option to select another allocation, which is called a *contingency*. Therefore, the expected utility generated from arranging exchanges in an exchange set depends on the order in which the possible cycle-based allocations are utilized. Such an ordering defines a more general procedure, which includes the two previously introduced procedures as special cases. For example, in an exchange set, $S = (\mathcal{V}', \mathcal{E}')$, if $\mathcal{C}^*(S)$ is adopted as top priority with no allocation assigned as a contingency, this in fact corresponds to the procedure presented in Section 2.3, and in consequence, the corresponding expected utility is $\sum_{C \in \mathcal{C}^*(S)} EU_C$. Similarly, the procedure described in Section 3.2 corresponds to selecting $\overline{\mathcal{C}}^*(S)$ as a first priority but again with no contingency at all, and hence generates an expected utility of $\sum_{C \in \overline{\mathcal{C}}^*(S)} EU_C$. According to what we have discussed, both of these two procedures are fixed in the sense that they each select an allocation only for $S = (\mathcal{V}', \mathcal{E}')$, but do not specify how to proceed on $(\mathcal{V}', \mathcal{E}' \setminus \mathcal{E}_f)$, where $\mathcal{E}_f \equiv \{e \in \mathcal{E}' : X_e = 0\}$ is observed when executing that chosen allocation. In contrast, the following “greedy” procedure is *sequential* and generates the largest expected utility.

- (i) Find $\mathcal{C}^*(S)$ in $S = (\mathcal{V}', \mathcal{E}')$ by applying the procedure presented in Section 2.3.
- (ii) If all cycles in $\mathcal{C}^*(S)$ are viable, finish with a claimed utility of $\sum_{C \in \mathcal{C}^*(S)} U_C$; if, however, certain edges, which we denote as $\mathcal{E}_f$, fail to be carried forward to actual operations, abort the original allocation and go back to (i) with $\mathcal{E}' \leftarrow \mathcal{E}' \setminus \mathcal{E}_f$.

Return to Figure 2, where both the three-way exchange cycle, $C_2 = \langle 2, 4, 5 \rangle$, and the two-way exchange cycle, $C_3 = \langle 2, 4 \rangle$, could be selected in the exchange set S_2 . If U_{C_2} is larger

than U_{C_3} , C_2 should be chosen as the first priority. When edge (4, 5) and/or edge (5, 2) in C_2 are not viable, C_3 could be tried as a contingency plan. On the other hand, if U_{C_2} is less than U_{C_3} , then the cycle C_3 , backed up by C_2 , should be selected as the first priority.

In light of this and similar examples, we quantify the potential value of selecting an exchange set S by applying the above greedy procedure for S . This potential value actually defines EU_S , the *expected utility* of S . Such an approach takes into consideration the full-potential contributions from all possible back-up allocations. Note that the expected utility of an exchange set can be exactly formulated as follows. For an exchange set, $S = (\mathcal{V}', \mathcal{E}')$, let $2^{\mathcal{E}'}$ denote the collection of all subsets of $\mathcal{E}'$, and for each $\tilde{\mathcal{E}} \in 2^{\mathcal{E}'}$, let

$$P(\tilde{\mathcal{E}}) \equiv \prod_{e \in \tilde{\mathcal{E}}} p_e \prod_{e \in (\mathcal{E}' \setminus \tilde{\mathcal{E}})} (1 - p_e)$$

which is the probability that the edges in $\tilde{\mathcal{E}}$ are viable whereas those in $\mathcal{E}' \setminus \tilde{\mathcal{E}}$ are not.

Consider $\tilde{S} \equiv (\mathcal{V}', \mathcal{E}')$, and let

$$U(\tilde{\mathcal{E}}) \equiv \sum_{C \in \mathcal{C}^*(\tilde{S})} U_C.$$

It follows that

$$EU_S = \sum_{\tilde{\mathcal{E}} \in 2^{\mathcal{E}'}} U(\tilde{\mathcal{E}}) P(\tilde{\mathcal{E}}).$$

In Figure 2, S_2 would be preferred to S_1 if $EU_{S_2} > EU_{S_1}$.

3.4 The maximum expected-utility set-based allocation

We have introduced a cycle-based allocation as a collection of disjoint exchange cycles. Among all such allocations, the procedure in Section 2.3 arranges kidney exchanges according to the allocation with the maximum utility, while the procedure in Section 3.2 arranges kidney exchanges according to the one with the maximum expected utility.

In this subsection, we define a set-based allocation as a collection of disjoint exchange sets, and further denote it as $\mathcal{S}(\mathcal{G})$ for a KPD program $\mathcal{G}$. Following the way in which the expected utility of an exchange set is defined in Section 3.3, the expected utility of a set-based allocation $\mathcal{S}(\mathcal{G})$ is, therefore, $\sum_{S \in \mathcal{S}(\mathcal{G})} EU_S$. Among all set-based allocations, the following procedure arranges kidney exchanges according to the maximum expected-utility set-based allocation, which we denote as $\overline{\mathcal{S}}^*(\mathcal{G})$.

- (i) Define $u: \mathcal{E} \rightarrow \mathbb{R}^+$ and $p: \mathcal{E} \rightarrow [0, 1]$ as in Section 3.2.
- (ii) Enumerate $\mathcal{S}_k$, the set of all exchange sets of size at most k , where $2 \leq k \leq |\mathcal{V}|$.

- (iii) For each $S \in \mathcal{S}_k$, calculate its expected utility as $EU_S = \sum_{\mathcal{E} \in 2^{\mathcal{E}'}} U(\tilde{\mathcal{E}}) P(\tilde{\mathcal{E}})$.
- (iv) Select $\overline{\mathcal{S}}^*(\mathcal{G})$ by forming an IP problem similar to the one discussed in Section 2.2.
- (v) Apply the aforementioned greedy procedure to each $S \in \overline{\mathcal{S}}^*(\mathcal{G})$.

Arranging kidney exchanges according to the above procedure allows for a more flexible utility-based evaluation of potential transplants, takes into consideration the uncertainties inherent in a KPD program, and provides fallback options when possible. Several remarks follow. In (ii), a reasonable k , say three or four, is required in practice due to logistical concerns as in Section 1. Enumerating $\mathcal{S}_k$ could be accomplished by a DFS-based algorithm similar to the one presented in Section 2.3. In (iii), calculating EU_S involves a summation over $|2^{\mathcal{E}'}|$ terms, which poses no computational difficulties in practice for small k .

3.5 Estimation of utilities and probabilities

So far in this paper, we have assumed that we are given a utility function, $u: \mathcal{E} \rightarrow \mathbb{R}^+$, and a probability function, $p: \mathcal{E} \rightarrow [0, 1]$. In practice, however, these utilities and probabilities are not available and have to be estimated.

In the literature, modeling of outcome-based utility has been considered in deceased-donor kidney transplants by Wolfe et al (2008), Schaubel et al (2006), and Schaubel et al (2009). Such models could be adopted relatively easily to living-donor kidney transplants. On the other hand, the model for probability can be established via a logistic modeling approach, based on clinical data from the Scientific Registry of Transplantation Recipients (SRTR), and multiple KPD programs including the Alliance for Paired Donation (APD) and the University of Michigan Transplant Center. In this logistic model, some primary predictors include, for example, the percentage of current panel reactive antibody (PRA) and the cross reactivity of antibody specificities.

Thus, in practice, all three procedures previously discussed for selecting the cycle-based or set-based allocations can be easily adopted by replacing the u and p with their estimated versions.

3.6 A KPD match run and an evolving KPD program

We consider a *match run* as a series of operations on a collection of incompatible pairs $\mathcal{V}$:

- (i) Form $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{E}$ is determined by checking virtual crossmatches on $\mathcal{V}$.
- (ii) Assign for each $e \in \mathcal{E}$ an estimated utility $\hat{u}_e$ and an estimated probability $\hat{p}_e$ according to the models discussed in Section 3.5.
- (iii) Arrange kidney exchanges according to $\mathcal{C}^*(\mathcal{G})$, $\overline{\mathcal{C}}^*(\mathcal{G})$, or $\overline{\mathcal{S}}^*(\mathcal{G})$.
- (iv) Recycle any donor-candidate pair that does proceed to an actual transplant operation back to the KPD pool awaiting future matches.

Every KPD program is constantly evolving in that successfully transplanted pairs leave and new incompatible pairs arrive over time. In addition, existing pairs in the pool could withdraw due to circumstances such as donor or candidate pregnancy, illness, or death. Such an evolving KPD program is managed by repeatedly executing match runs and updating the KPD pool on a regular basis over time. See Figure 3 for an illustration.

4 KPD simulation system

In Sections 2 and 3, we have presented three organ allocations for a KPD program $\mathcal{G}$, namely $\mathcal{C}^*(\mathcal{G})$, $\overline{\mathcal{C}}^*(\mathcal{G})$, and $\mathcal{T}^*(\mathcal{G})$. A comparison among them and other potential ones is certainly of great interest with respect to the practical management of a KPD program. Such a comparison, however, cannot be accomplished by traditional clinical trials due to the nature of kidney transplantation or any transplantation for that matter. Therefore, simulation plays an important role for achieving this purpose. In this section, key ingredients of such a simulation system will be discussed.

4.1 Generating incompatible pairs

To create an incompatible pair, we generate its donor and candidate separately according to their own population distributions. Candidates are *sampled at random with replacement* from databases of candidates presenting with a willing but incompatible donor. One such database is derived from the University of Michigan KPD program, in which candidates are described by variables of blood type, PRA, intensity of candidate antibody specificities, and so on. Currently, the UM database consists of 187 incompatible pairs, and additional databases from other KPD programs (as they become available to us through data-usage agreements) will be incorporated for wider variation in candidates. Donors, on the other hand, are generated by separately sampling their blood types and HLA haplotypes. Precisely, blood types are drawn from the U.S. population distribution: O, 44%; A, 42%; B, 10%; and AB, 4% (Stanford Blood Center, 2010). HLA haplotypes are sampled according to their frequencies in the U.S. population, which is derived from an extensive public database on potential bone marrow donors (Majers et al, 2007).

A simulated donor-candidate pair is regarded as an incompatible pair and hence included in the KPD pool if either their ABO blood types are mismatched, or the donor's HLA haplotypes contain any of the candidate's antibody specificities (with mean fluorescence intensity $> 5,000$), or both. In this paper, we do not consider KPD programs that include compatible pairs.

4.2 Simulating a match run and an evolving KPD program

We simulate a KPD match run on a collection of incompatible pairs, $\mathcal{V}$, by simulating each of its four steps, as discussed in Section 3.6.

- (i) Determine $\mathcal{E}$ by checking the simulated blood types and HLA haplotypes.
- (ii) Estimate both u_e and p_e based on generated pair characteristics.
- (iii) Apply $\mathcal{C}^*(\mathcal{G})$, $\overline{\mathcal{C}}^*(\mathcal{G})$, or $\mathcal{T}^*(\mathcal{G})$ to arrange kidney exchanges.

- (iv) Since a potential transplant may not be viable, we simulate such uncertainty via a Bernoulli trial with the probability of success equal to that edge probability; the realization of such a Bernoulli trial will indicate if a pair proceeds to an actual transplant and hence leaves the pool, or remains in the pool and awaits future matches.

To address the feature that a KPD program is evolving over time as discussed in Section 3.6, we first generate an initial KPD pool of N incompatible pairs as described in Section 4.1. Further, we assume that the arrival of new incompatible pairs follows a Poisson process with a rate λ . This rate may be governed by a log-linear model with covariates of age, candidate illness, and the relationship between donor and candidate, among others. Also, we assume that the withdrawal of existing pairs from the KPD pool follows another Poisson process with a rate μ ; and the log-linear model for μ could include other relevant covariates.

5 Simulation results

The proposed simulation system in Section 4 enables us to investigate various aspects associated with the management of a KPD program, among which comparing different allocation strategies is certainly of special interest. Simulation results including such comparisons are presented in this section. We implement the simulation system using the C++ programming language, and the related IP problems are solved by Gurobi Optimizer.

5.1 Problem complexity

The number of exchange cycles or sets can be enormous in a reasonably large KPD pool of, say, several hundred pairs, even though the size of exchange cycles or sets is restricted to at most three. Such complexity causes solving related IP problems computationally rather expensive. Table 1 summarizes the averaged numbers (over 200 rounds of simulations) of exchange cycles and sets of up to three pairs in pools of various sizes, 100, 200, 300, 400, and 500, where incompatible pairs are generated according to Section 4.1.

5.2 Simulation setup

We perform a total of 600 simulations, in which $\{p_e, e \in \mathcal{E}\}$ is generated from a uniform distribution, $U(0.1, 0.5)$; this distribution gives, for a virtual-crossmatch compatible transplant, an average successful rate of 30%, a reasonable number according to our experience at the University of Michigan and the Alliance for Paired Donation. In order to obtain a set of utilities for illustrative purpose, we generated utilities from two uniform distributions, $U(10, 20)$ and $U(10, 30)$; these are used to generate $\{u_e, e \in \mathcal{E}\}$. Each distribution of utilities is used for 200 simulations; edge utilities are fixed at 1 for the remaining 200 simulations. Notice that these simulation parameters are chosen in such ways to illustrate the application of our proposed simulation system; in practice, these utilities and probabilities should be estimated based on real data as we have discussed in Section 3.5.

We start each round of simulation by generating an initial pool of $N = 200$ incompatible pairs. Additionally generated pairs then enter the pool according to a Poisson process with $\lambda = 10$ pairs per month over a period of $m = 24$ months. For simplicity, we assume no existing pairs drop out of the pool.

A match run is executed at the end of each month on this evolving pool, starting with the initial 200 pairs. Pairs that arrive during the time of a match run will not participate in the current run, but wait for the next one. For each simulation, we make three copies of the evolving KPD program and execute each of the three match runs, which arrange kidney exchanges according to $\mathcal{C}^*(\mathcal{G})$, $\overline{\mathcal{C}}^*(\mathcal{G})$, and $\mathcal{S}^*(\mathcal{G})$, on each copy such that we can directly compare these allocation strategies. At the end of each match run, the KPD pool is updated with some pairs leaving or staying. We record several important measures needed for comparison, such as cumulative claimed utilities, cumulative number of transplants, and blood types of candidates receiving transplants.

The above description provides a simple simulation setup that enables us to examine and compare different allocation strategies. More realistic models will be further explored by incorporating more comprehensive data on pair characteristics and actual transplant operations, when such data become available to us via data-usage agreements.

5.3 Results

First, we report on the cumulative number of transplants over a period of 24 months across three allocation strategies, $\mathcal{C}^*(\mathcal{G})$, $\overline{\mathcal{C}}^*(\mathcal{G})$, and $\mathcal{S}^*(\mathcal{G})$, and under different models utilities, $U(1, 1)$, $U(10, 20)$, and $U(10, 30)$; see Figure 4, which unveils a consistent pattern regardless of utility models that $\mathcal{S}^*(\mathcal{G})$ results in the greatest number of transplants, whereas $\mathcal{C}^*(\mathcal{G})$ leads to the fewest number of transplants. Take Figure 4-C as an example. When the edge utility is fixed at 1, allocation strategy $\mathcal{S}^*(\mathcal{G})$ results in a median of 50 (over 200 simulations) transplants after match run 10 (at the end of month 9); in contrast, for exactly the same evolving KPD program, strategy $\mathcal{C}^*(\mathcal{G})$ only leads to a median of 34 transplants over the same period of time. Results from simulation also show that allocation strategy $\overline{\mathcal{C}}^*(\mathcal{G})$ allows for a much higher number of transplants than $\mathcal{C}^*(\mathcal{G})$ does, though it performs worse than $\mathcal{S}^*(\mathcal{G})$.

Figure 5 demonstrates that $\mathcal{S}^*(\mathcal{G})$ is advantageous over both $\mathcal{C}^*(\mathcal{G})$ and $\overline{\mathcal{C}}^*(\mathcal{G})$ in that $\mathcal{S}^*(\mathcal{G})$ on average achieves the largest cumulative claimed utility. Notice that when edge utilities are fixed at 1, the cumulative claimed utility is the same as the cumulative number of transplants; therefore, we only compare these allocation strategies under two other utility models, i.e. $U(10, 20)$ and $U(10, 30)$.

We can also examine via the proposed simulation system other aspects of a KPD program. For example, we are interested in exploring how the chance of having a transplant is associated with blood types. In practice, candidates with blood type O are usually in a disadvantageous position due to the limitation that they can only receive kidneys from blood type O donors, who, however, can donate to candidates of any blood type. This phenomenon is clearly observed in Figure 6, where about 60% of the incoming candidates are of blood type O, while only about 40% of the performed transplants involve a blood O candidate. As a consequence, candidates of the other blood types (A, B, and AB) are more represented among all candidates receiving a transplant. One possible solution to this difficulty is to

assign a bonus utility to a potential transplant that involves an O donor giving to an O candidate, or prohibit the option of an O donor being allocated to a non-O candidate.

6 Concluding remarks

In this paper, we have proposed a novel approach to organizing kidney exchanges in an evolving KPD program. Our approach identifies the maximum expected-utility set-based allocation that not only allows for a more general outcome-based evaluation of kidney transplants, but also takes into consideration stochastic features in managing a KPD program. Further, this approach exploits possible back-up exchanges when the originally planned allocation cannot be fully executed. Being able to provide such back-up exchanges is even more valuable considering that in practice utilities that relate to expected outcomes could usually decrease over time; it is generally better for a transplant to be performed sooner than later.

We have also suggested a data-based simulation system that enables us to examine various aspects of an evolving KPD program. This simulation system allows us to emulate genetic and demographic data from existing KPD programs, and to derive statistical models that mimic actual KPD programs. In particular, we have suggested (i) models for donor and candidate characteristics as well as their arrival in and withdrawal from a KPD pool; (ii) models for the estimation of the outcome-based utility of a potential transplant; and (iii) models for the prediction of the probability that a planned transplant would indeed occur. Utilizing such a simulation system, we are able to quantitatively compare different kidney allocation strategies; and results shed light on decision support applicable to actual KPD programs.

We have illustrated the proposed simulation system to compare several kidney allocation strategies. Results from simulation studies suggest advantages of adopting the maximum expected-utility set-based allocation over the other two allocation strategies. Such advantages are attributive to the introduction of an exchange set, in which uncertainties in a KPD system are properly incorporated and possible contingency allocations are allowed should the planned exchanges fail to be executed. Many issues of practical importance concern how allocation strategies should be implemented in a real KPD program, such as: what is the optimal timing for a match run; how long is the average waiting time for a pair with given properties, etc. Giving full answers to these questions requires more clinical data and realistic simulation models and is a direction we are pursuing.

Another possible future investigation is to consider exchanges initiated by an altruistic donor (Rees et al, 2009). An altruistic donor does not arrive with a designated candidate, but instead volunteers to donate a kidney to a pool of waiting candidates. Figure 7 gives an illustration of a chain of transplants initiated by an altruistic donor. Since transplants along the chain are not required to be performed simultaneously, a bridge donor (namely the donor whose incompatible candidate received a kidney but has yet to donate) at the end of the chain could later donate to a candidate who will arrive in the future, whose willing but incompatible donor then becomes a new bridge donor. In this respect, the chain is open-ended and increases the chance for a highly sensitized candidate to receive a compatible

kidney. The proposed work could be easily extended to include altruistic donors as participants in a KPD program. More specifically, a chain of kidney transplants can be viewed as a special exchange cycle, in which the bridge donor “donates” to a phantom candidate associated with the altruistic donor who first initiated that chain, and hence forms a hypothetical “edge” (denoted by dashed arrowed lines in Figure 7). Future research on KPD with altruistic donors is certainly of high importance and great interest; we will report those results in a future publication.

Acknowledgments

The authors thank the support from the Michigan Institute for Clinical and Health Research (MICHR), Michigan School of Public Health, Scientific Registry of Transplant Recipients, and National Institute of General Medical Sciences (NIGMS). The authors are also grateful for invaluable suggestions and comments from Drs. Yanhua Chen, David Abraham and Tuomas Sandholm. The authors would like to thank an anonymous referee whose comments helped to improve this manuscript. This work was supported in part by a grant from the National Institutes of Health (NIH) CTSA at the University of Michigan 2UL1TR000433-06 and by the NIH grant 1R01-DK093513.

Appendix - edge formulation

Let Y_{ij} be a decision variable for each $(i, j) \in \mathcal{E}$, such that Y_{ij} is 1 if edge (i, j) is chosen for a transplant and 0 otherwise. When no restriction is placed on the length of an exchange cycle, we solve the following IP problem:

$$\begin{aligned} & \max_{\{Y_{ij}\}} \quad \sum_{(i,j) \in \mathcal{E}} Y_{ij} u_{ij}, \\ & \text{subject to} \quad \sum_{j: (i,j) \in \mathcal{E}} Y_{ij} \leq 1, \forall i \in \mathcal{V}, \\ & \quad \quad \quad \sum_{j: (i,j) \in \mathcal{E}} Y_{ij} = \sum_{j: (j,i) \in \mathcal{E}} Y_{ji}, \forall i \in \mathcal{V}. \end{aligned}$$

This IP problem could be solved efficiently in polynomial time by finding a maximum weight perfect matching. If the exchange cycle length is restricted to be at most k , an additional set of constraints has to be added, i.e.

$$\sum_{e \in \mathcal{C}(\mathcal{C})} Y_e \leq k, \forall \mathcal{C} \in \mathcal{C}_k. \quad (2)$$

The number of additional constraints in (2), even when $k = 3$, is usually enormously large in a realistic KPD program with several hundred incompatible pairs, which makes it impossible to even store all the constraints in a typical IP solver such as CPLEX¹ or Gurobi².

Therefore, in this case, the cycle formulation presented in Section 2.2 is preferred.

References

- Abraham, DJ.; Blum, A.; Sandholm, T. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. Ec'07: Proceedings of The Eighth Annual Conference on Electronic Commerce; 2007. p. 295-304.
- Edmonds J. Paths, trees, and flowers. Canadian Journal of Mathematics. 1965; 17:449–467.

¹<http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>

²<http://www.gurobi.com/>

- Evans RW, Manninen DL, Garrison LP, Hart LG, Blagg CR, Gutman RA, Hull AR, Lowrie EG. The quality of life of patients with end-stage renal disease. *New England Journal of Medicine*. 1985; 312(9):553–559. [PubMed: 3918267]
- Hariharan S, Johnson CP, Bresnahan BA, Taranto SE, McIntosh MJ, Stablein D. Improved graft survival after renal transplantation in the united states, 1988 to 1996. *New England Journal of Medicine*. 2000; 342(9):605–612. [PubMed: 10699159]
- Keizer KM, de Klerk M, Haase-Kromwijk BJM, Weimar W. The dutch algorithm for allocation in living donor kidney exchange. *Transplantation Proceedings*. 2005; 37(2):589–591. [PubMed: 15848466]
- Laupacis A, Keown P, Pus N, Krueger H, Ferguson B, Wong C, Muirhead N. A study of the quality of life and cost-utility of renal transplantation. *Kidney International*. 1996; 50(1):235–242. [PubMed: 8807593]
- Maiers M, Gragert L, Klitz W. High-resolution hla alleles and haplotypes in the united states population. *Human Immunology*. 2007; 68(9):779–788. [PubMed: 17869653]
- Park K, Moon JI, Kim SI, Kim YS. Exchange donor program in kidney transplantation. *Transplantation*. 1999; 67(2):336–338. [PubMed: 10075605]
- Rapaport FT. The case for a living emotionally related international kidney donor exchange registry. *Transplantation Proceedings*. 1986; 18:5–9. [PubMed: 11649919]
- Rees MA, Kopke JE, Pelletier RP, Segev DL, Rutter ME, Fabrega AJ, Rogers J, Pankewycz OG, Hiller J, Roth AE, Sandholm T, Ünver MU, Montgomery RA. A nonsimultaneous, extended, altruistic-donor chain. *New England Journal of Medicine*. 2009; 360(11):1096–1101. [PubMed: 19279341]
- Ross LF, Rubin DT, Siegler M, Josephson MA, Thistlethwaite JR, Woodle ES. Ethics of a paired-kidney-exchange program. *New England Journal of Medicine*. 1997; 336(24):1752–1755. [PubMed: 9180096]
- Roth AE, Sönmez T, Ünver MU. Kidney exchange. *Quarterly Journal of Economics*. 2004; 119(2): 457–488.
- Roth AE, Sönmez T, Ünver MU. A kidney exchange clearinghouse in new england. *American Economic Review*. 2005; 95(2):376–380.
- Roth AE, Sönmez T, Ünver MU. Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *American Economic Review*. 2007; 97(3):828–851.
- Russell JD, Beecroft ML, Ludwin D, Churchill DN. The quality of life in renal transplantation-a prospective study. *Transplantation*. 1992; 54(4):656–660. [PubMed: 1412757]
- Schaubel DE, Wolfe RA, Port FK. A sequential stratification method for estimating the effect of a time-dependent experimental treatment in observational studies. *Biometrics*. 2006; 62(3):910–917. [PubMed: 16984335]
- Schaubel DE, Wolfe RA, Sima CS, Merion RM. Estimating the effect of a time-dependent treatment by levels of an internal time-dependent covariate: Application to the contrast between liver wait-list and posttransplant mortality. *Journal of the American Statistical Association*. 2009; 104(485): 49–59.
- Segev DL, Gentry SE, Warren DS, Reel B, Montgomery RA. Kidney paired donation and optimizing the use of live donor organs. *Journal of the American Medical Association*. 2005; 293(15):1883–1890. [PubMed: 15840863]
- Shapley L, Scarf H. On cores and indivisibility. *Journal of Mathematical Economics*. 1974; 1(1):23–37.
- Terasaki PI, Cecka JM, Gjertson DW, Takemoto S. High survival rates of kidney-transplants from spousal and living unrelated donors. *New England Journal of Medicine*. 1995; 333(6):333–336. [PubMed: 7609748]
- Wolfe RA, Ashby VB, Milford EL, Ojo AO, Ettenger RE, Agodoa LYC, Held PJ, Port FK. Comparison of mortality in all patients on dialysis, patients on dialysis awaiting transplantation, and recipients of a first cadaveric transplant. *New England Journal of Medicine*. 1999; 341(23): 1725–1730. [PubMed: 10580071]

Wolfe RA, McCullough KP, Schaubel DE, Kalbfleisch JD, Murray S, Stegall MD, Leichtman AB.
Calculating life years from transplant (lyft): Methods for kidney and kidney-pancreas candidates.
American Journal of Transplantation. 2008; 8(4p2):997–1011. [PubMed: 18336702]

Fig. 1.

(A): A two-way exchange cycle; (B): A three-way exchange cycle. The two diagrams in the bottom are the corresponding graphical representations of the top ones.

Fig. 2.

Two three-way exchange cycles, $C_1 = \langle 1, 2, 3 \rangle$ and $C_2 = \langle 2, 4, 5 \rangle$, and one two-way exchange cycle, $C_3 = \langle 2, 4 \rangle$; two exchange sets of size three, $S_1 = (\{1, 2, 3\}, \{(1, 2), (2, 3), (3, 1)\})$ and $S_2 = (\{2, 4, 5\}, \{(2, 4), (4, 5), (5, 2), (4, 2)\})$, and one exchange set of size two, $S_3 = (\{2, 4\}, \{(2, 4), (4, 2)\})$.

Fig. 3.

A flow diagram of an evolving KPD program

Fig. 4.

Cumulative number of transplants over 12 or 24 months across three allocation strategies, under three utility generating mechanisms (results are based on 200 simulations). Edge utilities are generated from $U(1, 1)$, $U(10, 20)$, and $U(10, 30)$; and edge probabilities are generated from $U(0.1, 0.5)$. The numbers at the top of each panel indicate match runs.

Fig. 5.

Cumulative claimed utilities of transplants over 24 months across three allocation strategies, under two utility generating mechanisms (results are based on 200 simulations). Edge utilities are generated from $U(10, 20)$, and $U(10, 30)$; and edge probabilities are generated from $U(0.1, 0.5)$. The numbers at the top of each panel indicates match runs.

Fig. 6.

Summaries of the percentage of candidates with certain blood type among all candidates joining the KPD pool and that percentage among all candidates who received a transplant (results are based on 200 simulations). Similar pattern are observed under three utility generating mechanisms, i.e. $U(1, 1)$, $U(10, 20)$, and $U(10, 30)$.

Fig. 7.

A chain of transplants initiated by an altruistic donor. The dashed arrow indicates a hypothetical “edge”, which represents a “donation” to a phantom candidate associated with the altruistic donor.

Table 1

The averaged numbers of exchange cycles or exchange sets up to three pairs in KPD pools of various sizes; standard deviations are given in the parentheses; the summary is calculated over 200 rounds of simulations.

pool size	exchange cycles (length 2 & 3)	exchange sets (size 2 & 3)
	mean (standard deviation)	mean (standard deviation)
100	388 (237)	383 (229)
200	2,659 (998)	2,630 (977)
300	9,413 (3,164)	9,305 (3,100)
400	21,076 (6,140)	20,829 (5,992)
500	40,290 (9,337)	39,815 (9,120)