

KIDNEY EXCHANGE*

ALVIN E. ROTH
TAYFUN SÖNMEZ
M. UTKU ÜNVER

Most transplanted kidneys are from cadavers, but there are also many transplants from live donors. Recently, there have started to be kidney *exchanges* involving two donor-patient pairs such that each donor cannot give a kidney to the intended recipient because of immunological incompatibility, but each patient can receive a kidney from the other donor. Exchanges are also made in which a donor-patient pair makes a donation to someone waiting for a cadaver kidney, in return for the patient in the pair receiving high priority for a compatible cadaver kidney when one becomes available. There are stringent legal/ethical constraints on how exchanges can be conducted. We explore how larger scale exchanges of these kinds can be arranged efficiently and incentive compatibly, within existing constraints. The problem resembles some of the “housing” problems studied in the mechanism design literature for indivisible goods, with the novel feature that while live donor kidneys can be assigned simultaneously, cadaver kidneys cannot. In addition to studying the theoretical properties of the proposed kidney exchange, we present simulation results suggesting that the welfare gains from larger scale exchange would be substantial, both in increased number of feasible live donation transplants, and in improved match quality of transplanted kidneys.

I. INTRODUCTION

Transplantation is the preferred treatment for the most serious forms of kidney disease. There are over 55,000 patients on the waiting list for cadaver kidneys in the United States, of whom almost 15,000 have been waiting more than three years. By way of comparison, in 2002 there were over 8,000 transplants of cadaver kidneys performed in the United States. In the same year,

* We thank Muriel Niederle, Erhan Sönmez, Stefanos Zenios, and the participants of the SITE 2003 Workshop on Matching Markets for their insightful comments and Hüseyin Çankaya for his help with the figures. After a first draft of this paper was completed, we also benefited from many helpful conversations with Dr. Frank Delmonico of the New England Organ Bank. Sönmez gratefully acknowledges the research support of KoçBank via the KoçBank scholar program and Turkish Academy of Sciences in the framework of the Young Scientist Award Program via grant TS/TÜBA-GEBİP/2002-1-19. We also received support from a National Science Foundation grant through the National Bureau of Economic Research. All errors are our own responsibility.

Email: aroth@hbs.edu; address: Harvard University, Department of Economics, Cambridge, MA 02138 and Harvard Business School, Boston, MA 02163.

Email: tsonmez@ku.edu.tr; address: Koç University, Department of Economics, College of Administrative Sciences and Economics, Rumeli Feneri Yolu, Sarıyer, İstanbul, 34450, Turkey.

Email: uunver@ku.edu.tr; address: Koç University, Department of Economics, College of Administrative Sciences and Economics, Rumeli Feneri Yolu, Sarıyer, İstanbul, 34450, Turkey.

© 2004 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology.

The Quarterly Journal of Economics, May 2004

TABLE I
U. S. KIDNEY TRANSPLANTS

Year	Cadaveric donors	Cadaveric transplants	Live donors	All wait-list patients	New wait-list additions
1992	4,276	7,202	2,535	22,063	15,224
1993	4,609	7,509	2,851	24,765	16,090
1994	4,797	7,638	3,009	27,258	16,538
1995	5,003	7,690	3,377	30,590	17,903
1996	5,038	7,726	3,649	34,000	18,328
1997	5,083	7,769	3,912	37,438	19,067
1998	5,339	8,017	4,361	40,931	20,191
1999	5,386	8,023	4,552	43,867	20,986
2000	5,490	8,089	5,324	47,596	22,269
2001	5,528	8,202	5,924	51,144	22,349
2002	5,630	8,534	6,233	54,844	23,494

The data for years 1992–2001 are constructed from the annual report of UNOS/OPTN, the data for 2002 are constructed from the national database of UNOS/OPTN. National database numbers are slightly higher than the annual report numbers due to continuous updating regarding previous years. Number of registrations may have multiple counts of patients since one patient may have registered in multiple centers for the wait-list.

about 3,400 patients died while on the waiting list, and another 900 became too ill to be eligible for transplantation. In addition to transplants of cadaver kidneys, in 2002 there were also somewhat over 6,000 transplants of kidneys from living donors, a number that has been increasing steadily from year to year. See Table I.

There is thus a considerable shortage of kidneys, compared with the demand. However, the substantial consensus in the medical community remains firmly opposed to allowing organs—even cadaveric organs—to be bought and sold, and this is a felony under the National Organ Transplant Act (NOTA) of 1984, and the Uniform Anatomical Gift Act of 1987.¹ The present paper considers ways to alleviate this shortage, and improve patient welfare, within the constraints of the current social and legal environment.

The National Organ Transplant Act of 1984 established the Organ Procurement and Transplantation Network (OPTN). Run

1. There is, however, a steady stream of literature both by doctors and by economists, considering how the shortage of organs might be alleviated by allowing their purchase and sale, and what effects this might have. See, e.g., Nelson et al. [1993] for an argument in favor of the status quo, and, e.g., Becker and Elias [2002] for an argument in favor of a market. Recent Congressional testimony endorsing the status quo but suggesting that empirical investigation of financial incentives might be in order can be found in Sade [2003].

by the United Network for Organ Sharing (UNOS), it has developed a centralized priority mechanism for the allocation of cadaveric kidneys.

Transplants from live donors generally have a higher chance of success than those from cadavers. The way such transplants are typically arranged is that a patient identifies a healthy willing donor (a spouse, for example) and, if the transplant is feasible on medical grounds, it is carried out. If the transplant from the willing donor is not feasible, the patient typically enters (or remains on) the queue for a cadaver kidney, while the donor returns home.

Recently, however, in a small number of cases, additional possibilities have been utilized when a transplant from a live donor and the intended recipient is infeasible. One of these, called a *paired exchange*, involves two patient-donor couples, for each of whom a transplant from donor to intended recipient is infeasible, but such that the patient in each couple could feasibly receive a transplant from the donor in the other couple [Rapaport 1986; Ross et al. 1997]. This pair of couples can then exchange donated kidneys. Compared with receiving cadaver kidneys at an unknown future time, this improves the welfare of the patients. In addition, it relieves the demand on the supply of cadaver kidneys, and thus potentially improves the welfare of those patients on the cadaver queue. A small number of these two-couple operations have been done, and the transplantation community has issued a consensus statement declaring them to be ethically acceptable [Abecassis et al. 2000].²

Another possibility is an *indirect exchange* (or *list exchange*) involving an exchange between one incompatible patient-donor couple, and the cadaver queue [Ross and Woodle 2000]. In this kind of exchange, the patient in the couple receives high priority on the cadaver queue, in return for the donation of his donor's kidney to someone on the queue.³ This improves the welfare of the patient in the couple, compared with having a long wait for a suitable cadaver kidney, and it benefits the recipient of the live

2. UNOS also published a legal opinion that such exchanges do not violate the NOTA (http://www.astso.org/ezefiles/UNOSSection_301_NOTA.pdf).

3. Priority on the cadaver queue is actually a bit complex, as queues are organized regionally, and consist of multiple queues, on which priority is determined by a scoring rule that gives points for how well matched the available kidney is to each patient, how long the patient has been waiting, etc. Giving high priority on the queue could be implemented by giving an appropriate number of points in the scoring rule.

kidney, and others on the queue who benefit from the increase in kidney supply due to an additional living donor. However, Ross and Woodle note that this may have a negative impact on type O patients already on the cadaver queue, an issue studied by Zenios, Woodle, and Ross [2001], to which we shall return.

In contrast to the system for cadaveric organs, and despite the growing interest in at least small-scale exchanges involving living donors, there is no national system, or even an organized registry at any level, for managing exchanges of kidneys from live donors. However, individual hospitals are beginning to think about larger scale living donor exchanges. As this paper was being written, the first three-couple kidney transplant exchange in the United States was reported at Johns Hopkins Comprehensive Transplant Center in Baltimore, among three couples for whom no two-couple transfer was feasible [Olson, August 2, 2003]. In the present paper we will consider how such a system of exchanges might be organized, from the point of view of achieving efficiency, and providing consistent incentives to patients, donors, and doctors, and what its welfare implications might be. We will see that the benefits of wider exchange accrue not only to the parties to the exchange. The resulting increase in live organ donation also benefits patients waiting for cadaver kidneys, including type O patients. The design we propose is partly inspired by the mechanism design literature on “house allocation,” and is intended to build on and complement the existing practices in kidney transplantation. In this respect and others it is in the modern tradition of engineering economics (see Roth [2002]) as applied to other problems of allocation, such as labor market clearinghouses (see Roth and Peranson [1999]), or auctions (see Milgrom [2004]), in which practical implementation often involves incremental change in existing practices.

II. BACKGROUND

II.A. Kidney Transplantation

End Stage Renal Disease (ESRD) is a fatal disease unless treated with dialysis or kidney transplantation. Transplantation is the preferred treatment. Two genetic characteristics play key roles in the feasibility and success of a kidney transplant. The first is the ABO blood-type: There are four blood types: A, B, AB, and O. Absent other complications, type O kidneys can be trans-

planted into any patient; type A or type B kidneys can be transplanted into same type or type AB patients; and type AB kidneys can only be transplanted into type AB patients. (So type O patients can only receive type O kidneys.) The second genetic characteristic is tissue type, also known as HLA type: HLA type is a combination of six proteins. As the HLA mismatch between the donor and the recipient increases, the likelihood of graft (i.e., transplanted organ) survival decreases [Opelz 1997]. HLA plays another key role in transplantation through the pretransplant "crossmatch" test. Prior to transplantation, the potential recipient is tested for the presence of preformed antibodies against HLA in the donor kidney. The presence of antibodies, called a *positive crossmatch*, effectively rules out transplantation.

When a cadaveric kidney becomes available for transplantation, the priority of each patient on the waiting list is determined by a point system based on factors including the blood type, HLA antigen-match, time spent on the waiting list, the region the kidney is harvested, etc. and the kidney is offered to the patient with the highest priority. If that patient declines, the kidney is offered to the patient with the next highest priority, and so on. Living donor kidney grafts have superior survival rates (and their availability can also avoid the long waiting time for a cadaver kidney). However, potential living donors can be eliminated from consideration due to incompatibility of the potential donor kidney with the intended recipient.

To minimize the elimination of physically eligible volunteer kidney donors on the basis of immunologic incompatibilities, Rapaport [1986] proposed the creation of a living donor pool for paired exchange. Ross et al. [1997] again proposed to increase the supply of living kidney donations by using kidneys from living incompatible donors through an exchange arrangement between two pairs. In 2000, UNOS initiated pilot testing of such programs.

Another exchange program is the *indirect exchange program* [Ross and Woodle 2000]: a potential donor who is incompatible with his intended recipient donates his kidney to the cadaveric waiting list, and his paired recipient will receive priority for the next compatible cadaveric kidney. There is widespread agreement in the transplantation community that indirect exchange can harm type O patients who have no living donors. First, they will be losing their priority to type O patients whose incompatible donors donate to the cadaveric pool, and second, very few type O living kidneys will be offered to their pool since a type O donor can

directly donate to his intended recipient unless there is a positive crossmatch. Despite this widespread concern, many transplant centers have also cautiously started pilot indirect exchange programs since 2000. For example, in UNOS Region 1 (New England), consisting of fourteen transplant centers and two Organ Procurement Organizations, four paired exchanges and seventeen indirect exchanges have been conducted from 2001 through 2003 (personal communication).

II.B. Mechanism Design

We will extend results in the mechanism design literature, which we first quickly review. Shapley and Scarf [1974] modeled a “housing market” consisting of n agents each of whom is endowed with an indivisible good, a “house.” Each agent has preferences over all the houses, and there is no money in the market, trade is feasible only in houses. They attribute to David Gale the “top trading cycle” algorithm that produces a house allocation in the core of the market. The algorithm works as follows: each agent points to her most preferred house (and each house points to its owner). There is at least one cycle in the resulting directed graph. In each such cycle, the corresponding trades are carried out, i.e., each agent in the cycle receives the house she is pointing to, and these agents and houses are removed from the market. The process continues (with each agent pointing to her most preferred house that remains on the market) until no agents remain, and the final allocation is the one in which each agent receives the house with which she left the market. When all preferences are strict, the procedure yields a unique outcome [Roth and Postlewaite 1977], and truthful preference revelation is a dominant strategy [Roth 1982].

Note that paired kidney exchanges similarly seek the gains from trade among patients with willing donors, but (with the recent Johns Hopkins three-pair exchange being a notable exception) mostly among just two pairs. In the kidney exchange to be considered below, if we consider exchange only among patients with donors, the properties of the housing market model essentially carry over unchanged, if we assume that donors’ preferences are aligned with those of their intended recipient. We will also assume that all surgeries in a given cycle are carried out simultaneously, which is the current practice, since a donor’s willingness to donate a kidney might change once her intended recipient received a transplant.

However, the kidney transplant environment consists not just of patients with donors, but also patients without donors, and cadaver kidneys not tied to any specific patient. Abdulkadiroğlu and Sönmez [1999] studied housing allocation on college campuses, which is in some respects similar: a set of rooms must be allocated to a set of students by a centralized housing office. Some of the students are existing tenants each of whom already occupies a room, and the rest of the students are newcomers. In addition to occupied rooms, there are vacant rooms. Existing tenants are entitled to keep their current rooms but may also apply for other rooms. Mechanisms used on a number of college campuses do not ensure the participation of existing tenants, and result in efficiency loss. This is the motivation for the generalization of the top trading cycles (TTC) mechanism proposed by Abdulkadiroğlu and Sönmez [1999], which they called *you request my house—I get your turn* (YRMH-IGYT): each student reports his strict preferences over all rooms, and an ordering of agents is randomly chosen. For any given preference list and ordering, the outcome is obtained as follows. (1) Assign the first student his top choice, the second student his top choice among the remaining rooms, and so on, until someone requests the room of an existing tenant who is yet to be served. (2) Whenever that happens, *modify the remainder of the ordering by moving the existing tenant to the beginning of the line*, and proceed with the procedure. (3) If at any point a cycle forms, it is formed exclusively by existing tenants, and each of them requests the room of the tenant who is next in the cycle. In such cases remove all students in the cycle by assigning them the rooms they request, and proceed with the procedure. The key innovation here is that an existing tenant whose current room is requested is upgraded to the first place in the line of agents remaining unassigned, before his room is allocated. As a result, the YRMH-IGYT mechanism assures every existing tenant a room that is no worse than his own. Therefore, existing tenants do not have any reason not to enter the market, and consequently the eventual allocation is Pareto efficient. Note that *the idea of upgrading an existing tenant whose current room is requested to the top of the line was also invented by the transplantation community in the form of an indirect exchange program*: when a potential donor donates his kidney to the highest priority patient on the waiting list, his intended recipient is upgraded to the top of the waiting list.

Note again that what prompted the introduction of simple

kidney exchange programs was the loss of volunteer kidney donors because of immunologic incompatibilities. Under these exchange programs, a potential donor who is incompatible with his intended recipient is given the incentive to go ahead with the donation, because his donation makes it possible for his intended recipient to receive a compatible kidney. Similarly, the potential efficiency loss in the campus housing problem is that some rooms might fail to be traded, even when welfare-enhancing trades are possible. The YRMH-IGYT is an attempt to address that problem in the housing context. We next consider how it must be adapted to kidney exchange.

III. KIDNEY EXCHANGE AND THE TOP TRADING CYCLES AND CHAINS (TTCC) MECHANISM

While there are clear similarities between house allocation and kidney exchange, there are also important differences. The counterpart of an existing tenant and his room is a donor-recipient pair, which we denote by (k_i, t_i) . We will often refer to donor k_i as kidney k_i , and recipient t_i as patient t_i . In the context of house allocation with existing tenants, there are also newcomers, none of whom owns a specific house, and vacant houses, none of which is owned by a specific student. The counterpart of newcomers are patients who have no living donors, and the counterpart of vacant houses are cadaveric kidneys that are not targeted for specific patients. This analogy reveals one important difference between the two models: in the house allocation model, the set of vacant houses is known. In the kidney exchange problem, it is not clear which cadaveric kidneys will be available, when they will be available, etc. *Therefore, while occupied houses and vacant houses are simultaneously allocated under the YRMH-IGYT mechanism, this is not possible in the context of kidney exchange.* Instead, patients with live donors who are not themselves allocated a live donor kidney will be assigned to the cadaver queue (with a priority reflecting whether their donor's kidney was donated to someone on the queue).

Let K denote the set of living donor kidneys at a particular time. While patients and their doctors may define their preferences over kidneys as they wish, here we consider, for specificity, the preferences that come from maximizing the probability of a successful transplant. Given any patient, part of K is outside the feasible set due to ABO blood-type incompatibility or a positive

crossmatch. Among feasible kidneys, HLA match [Opelz 1997], donor age, kidney size, etc. play a significant role in the graft survival. Therefore *patients have heterogeneous preferences over compatible kidneys*. In what follows, we will consider all preferences to be strict. *If only direct exchanges among donor-recipient pairs are considered, one can directly use Gale's Top Trading Cycles mechanism*. However, this will not allow for indirect exchanges. We will need to modify the model and the mechanism to allow for this possibility. Since the supply of specific cadaveric donor kidneys is not predictable, a patient who wishes to trade his donor's kidney in return for a priority in the cadaveric kidney waiting list is receiving a lottery. Taking this into consideration, the patient, doctor, and donor can decide whether this option is acceptable and, if so, where it ranks in the patient's preferences.

Given a patient t_i , let $K_i \subset K$ denote the set of living donor kidneys that are compatible with patient t_i . Let w denote the option of entering the waiting list with priority reflecting the donation of his donor's kidney k_i , and P_i denote his strict preferences over $K_i \cup \{k_i, w\}$. For our purposes the relevant part of P_i is the ranking up to kidney k_i or w , whichever ranks higher. If patient t_i ranks kidney k_i at the top of his preferences, that means he and his donor do not wish to participate in an exchange. If patient t_i ranks k_i on top of w , that means he and his donor do not consider exchanging kidney k_i with a priority in the cadaveric kidney waiting list.

We can now formalize a (static) *kidney exchange problem* consisting of a set of donor-recipient pairs $\{(k_1, t_1), \dots, (k_n, t_n)\}$, a set of compatible kidneys $K_i \subset K = \{k_1, \dots, k_n\}$ for each patient t_i , and a strict preference relation P_i over $K_i \cup \{k_i, w\}$ for each patient t_i . The outcome of a kidney exchange problem is a *matching* of kidneys/wait-list option to patients such that each patient t_i is either assigned a kidney in $K_i \cup \{k_i\}$ or the wait-list option w , and no kidney can be assigned to more than one patient although the wait-list option w can be assigned to several patients. A *kidney exchange mechanism* selects a matching for each kidney exchange problem. We are almost ready to introduce the *Top Trading Cycles and Chains* (TTCC) mechanism, a generalization of the TTC mechanism, for kidney exchange. First, we give a few definitions and observations to facilitate the description of the mechanism.

III.A. Cycles and w -Chains

The mechanism relies on an algorithm consisting of several rounds. In each round each patient t_i points either toward a kidney in $K_i \cup \{k_i\}$ or toward w , and each kidney k_i points to its paired recipient t_i .

A *cycle* is an ordered list of kidneys and patients $(k'_1, t'_1, k'_2, t'_2, \dots, k'_m, t'_m)$ such that kidney k'_1 points to patient t'_1 , patient t'_1 points to kidney k'_2 , \dots , kidney k'_m points to patient t'_m , and patient t'_m points to kidney k'_1 . Cycles larger than a single pair are associated with direct exchanges, very much like the paired-kidney-exchange programs, but may involve more than two pairs, so that patient t'_1 is assigned kidney k'_2 , patient t'_2 is assigned kidney k'_3 , \dots , patient t'_m is assigned kidney k'_1 . Note that each kidney or patient can be part of at most one cycle and thus *no two cycles intersect*.

A *w-chain* is an ordered list of kidneys and patients $(k'_1, t'_1, k'_2, t'_2, \dots, k'_m, t'_m)$ such that kidney k'_1 points to patient t'_1 , patient t'_1 points to kidney k'_2 , \dots , kidney k'_m points to patient t'_m , and patient t'_m points to w . We refer to the pair (k'_m, t'_m) whose patient receives a cadaver kidney in a w -chain as the *head* and the pair (k'_1, t'_1) whose donor donates to someone on the cadaver queue as the *tail* of the w -chain. W -chains are associated with indirect exchanges but unlike in a cycle, a kidney or a patient can be part of several w -chains. One practical possibility is choosing among w -chains with a well-defined *chain selection rule*, very much like the rules that establish priorities on the cadaveric waiting list. The current pilot indirect exchange programs in the United States choose the minimal w -chains, consisting of a single donor-recipient pair, but this may not be efficient. Selection of longer w -chains will benefit other patients as well, and therefore the choice of a chain selection rule has efficiency implications (see Theorem 1). Chain selection rules may also be used for specific policy objectives such as increasing the inflow of type O living donor kidneys to the cadaveric waiting list. Whenever w -chain $(k'_1, t'_1, \dots, k'_m, t'_m)$ is selected in the algorithm, patient t'_1 is assigned kidney k'_2 , patient t'_2 is assigned kidney k'_3 , \dots , patient t'_{m-1} is assigned kidney k'_m , patient t'_m receives high priority for the next compatible kidney in the cadaveric waiting list, and kidney k'_1 is offered either to the cadaveric waiting list or to another patient with a paired donor.

LEMMA 1. Consider a graph in which both the patient and the kidney of each pair are distinct nodes as is the wait-list option w . Suppose that each patient points either toward a kidney or w , and each kidney points to its paired recipient. Then either there exists a cycle, or each pair is the tail of some w -chain.

We can now introduce the TTCC mechanism. Because the exchange mechanism interacts with many parts of the kidney transplant environment, it will clarify the discussion to start by indicating which parts of the environment we take as fixed for our present purpose. First, we take the operation of the cadaver queue as fixed. The cadaver queue can be thought of as a stochastic arrival process of cadavers and patients, interacting with a scoring rule that determines which patients are offered which cadaver kidneys. We also take as fixed how patients whose donors donate a kidney to someone on the queue are given high priority on the queue, e.g., by being given points in the scoring rule.⁴ We also take as given the size of the live kidney exchange; i.e., the set of patient-donor pairs is taken to be fixed. In practice, the set of patient-donor pairs will grow as the geographic area served by the kidney exchange is increased, or as the time between exchanges is increased. A larger pool of possible exchanges will increase the potential efficiency gains that can be realized by exchange, but will also increase the size of the trading cycles/ w -chains that might be needed to achieve these efficiencies. We will keep track of both of these when we report simulations. Both the operation of the cadaver queue, and the frequency and scope of the kidney exchange will influence patients' "reservation utility," i.e., how they compare various opportunities for direct or indirect exchange to the option of not making any exchange now, but waiting for a future opportunity. Patients can express this reservation utility by where they rank their own donor in their preferences.

4. Depending on how this priority is given, patients may need to be aware of the current population of the queue to evaluate the desirability of the w option. Liran Einav and Muriel Niederle have shared with us interesting ideas about how to further model the desirability of the w option dynamically, taking into account that others may enter the queue with high priority, but we will not pursue this here.

III.B. The Exchange Mechanism

For the mechanism defined below, we assume that when one among multiple w-chains must be selected, a fixed chain selection rule is invoked. We will consider a number of such rules, and their implications for incentives, efficiency, and equity.

Throughout the procedure kidneys are assigned to patients through a series of exchanges. Some patients and their assigned kidneys will be immediately removed from the procedure, while others will remain with their assignments but they will assume a *passive* role. So at any point in the procedure, some agents may no longer be participants, some participants will be active, and the others passive.

For a given kidney exchange problem, the TTCC mechanism determines the exchanges as follows.

1. Initially, all kidneys are *available*, and all agents are active. At each stage of the procedure each remaining *active* patient t_i points to his most preferred remaining unassigned kidney or to the wait-list option w , whichever is more preferred, each remaining passive patient continues to point to his assignment, and each remaining kidney k_i points to its paired recipient t_i .
2. *By Lemma 1, there is either a cycle, or a w-chain, or both.*
 - (a) Proceed to Step 3 if there are no cycles. Otherwise, locate each cycle, and carry out the corresponding exchange (i.e., each patient in the cycle is assigned the kidney he is pointing to). Remove all patients in a cycle together with their assignments.
 - (b) Each remaining patient points to his top choice among remaining kidneys, and each kidney points to its paired recipient. Locate all cycles, carry out the corresponding exchanges, and remove them. Repeat until no cycle exists.
3. If there are no pairs left, we are done. Otherwise, by Lemma 1 each remaining pair is the tail of a w-chain. Select *only one* of the chains with the *chain selection rule*. The assignment is *final* for the patients in the selected w-chain. The chain selection rule also determines whether the selected w-chain is removed and the associated exchanges are all immediately assigned (including the kidney at the tail, which is designated to go to a patient on the cadaver queue), or if the selected w-chain is kept in the

procedure although each patient in it is passive henceforth.⁵

4. After a w-chain is selected, new cycles may form. Repeat Steps 2 and 3 with the remaining active patients and unassigned kidneys until no patient is left.

At the end of the procedure, each patient with a living donor is assigned a kidney (or a high priority place on the waiting list). However, that does not necessarily mean each of these patients receives a transplant. In particular, a minimal cycle (k_i, t_i) consisting of a single patient-donor pair may be a pair that was not offered a sufficiently desirable kidney in the current exchange, and chooses to wait in the hope of exchanging for a high quality living donor kidney the next time the exchange is run, after new donors have entered the system.

EXAMPLE 1. Consider a kidney exchange problem with 12 pairs $(k_1, t_1), \dots, (k_{12}, t_{12})$ with preferences as follows:

$t_1: k_9 \ k_{10} \ k_1$	$t_7: k_6 \ k_1 \ k_3 \ k_9 \ k_{10} \ k_1 \ w$
$t_2: k_{11} \ k_3 \ k_5 \ k_6 \ k_2$	$t_8: k_6 \ k_4 \ k_{11} \ k_2 \ k_3 \ k_8$
$t_3: k_2 \ k_4 \ k_5 \ k_6 \ k_7 \ k_8 \ w$	$t_9: k_3 \ k_{11} \ w$
$t_4: k_5 \ k_9 \ k_1 \ k_8 \ k_{10} \ k_3 \ w$	$t_{10}: k_{11} \ k_1 \ k_4 \ k_5 \ k_6 \ k_7 \ w$
$t_5: k_3 \ k_7 \ k_{11} \ k_4 \ k_5$	$t_{11}: k_3 \ k_6 \ k_5 \ k_{11}$
$t_6: k_3 \ k_5 \ k_8 \ k_6$	$t_{12}: k_{11} \ k_3 \ k_9 \ k_8 \ k_{10} \ k_{12} \ .$

Suppose that patients are ordered in a priority-list based on their indices starting with the patient with the smallest index. We use the following chain selection rule: choose the longest w-chain. In case the longest w-chain is not unique, choose the w-chain with the highest priority patient; if the highest priority patient is part of more than one, choose the w-chain with the second highest priority patient, and so on. Keep the selected w-chains until the termination.

The execution of the TTCC mechanism is given in Figures I–V.

5. The relevance of the last point is the following: whenever a w-chain $(k'_1, t'_1, \dots, k'_m, t'_m)$ is selected, even though the assignments of all patients in the w-chain are finalized, the kidney k'_1 at the tail of the w-chain can be utilized in two possible ways. It can immediately be offered to the waiting list (in which case the w-chain is removed), or it may be made available to the remaining patients as the process continues, and hence *the selected w-chain may possibly grow later on*, although the patients already in it are not affected.




FIGURE II
Example 1, Round 2

Upon removing cycle C_1 , a new cycle $C_2 = (k_7, t_7, k_6, t_6, k_5, t_5)$ forms. Remove it by assigning k_7 to t_5 , k_6 to t_7 , and k_5 to t_6 .

- d. (e.) Prioritize patient-donor pairs in a single list. Choose the w-chain starting with the highest priority pair, and remove it (keep it).

A w-chain that is formed at an interim step of the procedure may grow at subsequent steps unless it is removed; hence the immediate removal of w-chains has a potential efficiency cost. Therefore, the following “hybrid” of chain selection rules d and e may appeal to those who wish to moderate the efficiency loss while increasing the inflow of type O living kidneys to the cadaveric waiting list.

- f. Prioritize the patient-donor pairs so that pairs with type O donor have higher priorities than those who do not. Choose the w-chain starting with the highest priority pair; remove it in case the pair has a type O donor, but keep it otherwise.

III.C. Efficiency and Incentives

In what follows, we will speak of Pareto efficiency in terms of the agents in the kidney exchange problem, namely the paired




FIGURE III
Example 1, Round 3

No new cycle forms, and hence each kidney-patient pair starts a w-chain. The longest w-chains are $W_1 = (k_8, t_8, k_4, t_4, k_9, t_9)$ and $W_2 = (k_{10}, t_{10}, k_1, t_1, k_9, t_9)$. Since t_1 , the highest priority patient, is in W_2 but not in W_1 , choose and fix W_2 . Assign w to t_9 , k_9 to t_1 , and k_1 to t_{10} but do not remove them. Kidney k_{10} , the kidney at the tail of W_2 , remains available for the next round.

patients and donors who are available to participate in the kidney exchange. Given a kidney exchange problem, a matching is *Pareto efficient* if there is no other matching that is weakly preferred by all patients and donors and strictly preferred by at least one patient-donor pair. A kidney exchange mechanism is *efficient* if it always selects a Pareto efficient matching among the participants present at any given time.

THEOREM 1. Consider a chain selection rule such that any w-chain selected at a nonterminal round remains in the procedure, and thus the kidney at its tail remains available for the next round. The TTCC mechanism, implemented with any such chain selection rule, is efficient.

Chain selection rules that remove a selected w-chain before ter-




FIGURE IV

Example 1, Round 4

Upon fixing the w-chain W_2 , a new cycle $C_3 = (k_4, t_4, k_8, t_8)$ forms. Remove it by assigning k_4 to t_8 and k_8 to t_4 .

mination of the algorithm, on the contrary, may yield Pareto-inefficient outcomes.

Roth [1982] showed for the housing model that *truthful preference revelation is a dominant strategy* of the preference revelation game induced by the TTC mechanism, and hence an agent can *never* profit by misrepresenting his preferences. Recall that, *in the absence of indirect exchanges*, the static kidney exchange problem is a housing market, and therefore the Roth [1982] result immediately applies.⁶ When indirect exchanges are allowed,

6. That is, at any specific time, a patient cannot receive a more preferred kidney by misrepresenting his preferences, which include his option value for the possibility that his donor's kidney can be used in a future exchange. We emphasize that of course we speak of strategy-proofness in the limited strategy space, the space of stated preferences, we have modeled for the kidney exchange problem. There may remain strategic issues associated with other aspects of the organ transplant process, such as being registered at multiple transplant centers and hence appearing on multiple regional waiting lists.




FIGURE V

Example 1, Round 5

No new cycles form, and the pair (k_{12}, t_{12}) "joins" W_2 from its tail to form the longest w-chain $W_3 = (k_{12}, t_{12}, k_{10}, t_{10}, k_9, t_9, w)$. Fix W_3 , and assign k_{10} to t_{12} . Since no patient is left, w-chain W_3 is removed, and kidney k_{12} at its tail is offered to the highest priority patient at the cadaveric waiting list.

whether the TTCC mechanism is strategy-proof depends on the choice of the chain selection rule.

THEOREM 2. Consider the chain selection rules a, d, e, and f. The TTCC mechanism, implemented with any of these chain selection rules, is strategy-proof.

Among these four chain selection rules, the last two are especially appealing: Rule e yields an efficient and strategy-proof mechanism, whereas Rule f gives up efficiency in order to increase the inflow of type O kidneys to the cadaveric waiting list. On the negative side, strategy-proofness of TTCC is lost if one adopts a chain selection rule that chooses among the longest w-chains.

EXAMPLE 2. Consider the problem in Example 1, but suppose that patient t_4 misrepresents his preferences as $P'_4 =$

k_5, k_1, k_9, \dots improving the ranking of kidney k_1 . While Round 1 and Round 2 remain as in Example 1, Round 3 changes, and this time the longest w-chain at Round 3 is $W_4 = (k_8, t_8, k_4, t_4, k_1, t_1, k_9, t_9)$. Therefore, patient t_4 is assigned kidney k_1 instead of kidney k_8 , making his preference misrepresentation profitable.

IV. SIMULATIONS

The theoretical treatment of the TTCC mechanism makes clear that larger exchanges may yield welfare gains, but it gives no idea of their magnitude. The following simulations are meant as a first step in that direction, and as a “proof of concept” to demonstrate that the gains are potentially substantial. We use data where they exist, e.g., on the likelihood of mismatches and positive crossmatches (see Table II). Where no data exist—on the willingness of patients and donors to trade a live donation for priority on the cadaver queue—we do robustness checks by simulating a wide range of preferences.

IV.A. Patient and Donor Characteristics

In addition to characteristics reported in Table II, for the HLA characteristics of the population, we use the distribution reported in Zenios [1996] using the UNOS registration data for years between 1987 and 1991. We assume that all HLA proteins and blood type are independently distributed following Zenios. For simplicity, we consider unrelated donor-patient pairs. About 25.3 percent all living-donor transplants were in this category in 2001. We use UNOS data to find the conditional distribution of the age of a nonspousal unrelated donor given that he is an adult. We assume that HLA and blood-type characteristics of the donor have the same distribution as the patients’, the characteristics of a nonspousal unrelated donor are independently distributed with the patient, and the characteristics of a spouse are independently distributed with the patient except his or her age. We assume that the spouse age is the *same* as the patient age.

IV.B. Preference Construction

The preferences of patients are determined using the survival analysis of grafts reported in Mandal et al. [2003]. This analysis uses data obtained from first-time kidney-only transplants between 1995 and 1998 in the United States Renal Data

TABLE II
 AMERICAN CAUCASIAN PATIENT AND LIVING DONOR CHARACTERISTIC
 DISTRIBUTIONS USED IN SIMULATIONS

A. Patient ABO blood type	Frequency
O	45.6%
A	39.5%
B	11.1%
AB	3.8%
B. Patient gender	Frequency
Female	40.9%
Male	59.1%
C. Patient age	Frequency
<18	5.6%
18–34	13%
35–49	34.9%
50–64	38.9%
>64	7.6%
D. Unrelated living donors	Frequency
Spouse	53.5%
Other	46.5%
E. Living donor age	Frequency
<18	5.6%
18–34	13%
35–49	34.9%
50–64	38.9%
>64	7.6%
F. Positive crossmatch	Frequency
Female patient—husband	33.3%
Other	11.1%

The frequencies are obtained from the UNOS data for various years. Patients are the new wait-list additions recorded between January 1995 and April 2003, except the gender data. The gender and living donor data were recorded between 1992 and 2001. Based on UNOS/OPTN data and annual report as of 7/14/2003 retrieved from <http://www.optn.org>. Positive crossmatch probability is reported by Zenios, Woodle, and Ross [2001].

System (USRDS) database. We assume that the utility function of each patient depends on the number of HLA mismatches (x) and the donor age (y). In the “rational” preference construction, following Mandal et al., we assume that each patient younger than 60 has a utility function $u(x,y) = -0.514x - y/10$, and each patient 60 and older has a utility function $u(x,y) = -0.510x - y/10$. We also consider a “cautious” preference construction. Under cautious preferences, we assume that patient t_i prefers donor

$k_j \in K_i$ to his own donor k_i only if kidney k_i is not compatible with him, or if although kidney k_i is compatible with him, it has *more than* an equivalent of one additional HLA mismatch than kidney k_j has.

Under both preference scenarios, the wait-list option may or may not be considered acceptable by a patient. Since the expected quality of HLA match is very low when a patient is given priority in the waiting list and since the graft failure rates are significantly higher for cadaveric kidneys than living-donor kidneys, we assume that a patient considers the wait-list option acceptable *only if* his donor is not compatible with him. We also assume that the patients who consider this option acceptable prefer any compatible living-donor kidney to this option. Because there are no reliable data available on the rate of patients who consider this option acceptable and because it depends on how priority is given, we consider two treatments, in which 0 percent and 40 percent of the patients with incompatible donors prefer the wait-list option to their own donors.

IV.C. Simulated Mechanisms

We consider four exchange mechanisms to contrast with the *no-exchange* regime: (1) *paired-kidney-exchange mechanism*, (2) *TTC mechanism*, (3) *paired and indirect exchange mechanism*, and (4) *TTCC mechanism* with the efficient and strategy-proof chain selection rule e . In our simulations we randomly simulate a sample of n donor-patient pairs using the population characteristics explained above. Then, we determine the preferences of patients over kidneys in the sample: for each patient t_i , we first check whether a donor k_j is ABO-compatible. If k_j is ABO-compatible, then we check whether there is a positive crossmatch between t_i and k_j . If they test negative for crossmatch, then k_j is in the compatible donor set K_i of patient t_i . After finding the set of compatible kidneys for each patient, we obtain a preference ordering on this set, using the utility functions described above. We construct four sets of preferences for each patient using the rational or cautious preference construction and assuming that 0 percent or 40 percent of patients with incompatible donors consider the wait-list option acceptable. We simulate each of the five mechanisms under these four preference scenarios. We use a Monte-Carlo simulation size of 100 trials for three different population sizes of 30, 100, and 300 pairs.

IV.D. Discussion of Results

The simulation results suggest that substantial gains in the number and match quality of transplanted kidneys might result from adoption of the TTCC mechanism. We report the details of this analysis in three tables. The rows of the tables refer to different regimes, under different preference constructions, and different population sizes.

Table III reports the general patient statistics under each regime. The first column reports the *total live donor transplants* as percentage of the population size, which is the sum of the next two columns, *transplants from own compatible donor* and *transplants from trades*. The fourth column is the percentage of *patients upgraded to the top of the wait-list* through indirect exchanges. The fifth column reports the quality of matches in the live donor transplants: the lower the HLA mismatch is, the higher the odds of graft survival. Standard errors are reported below the estimates.

In Table IV we report the effect of each regime on the wait-list additions for each blood type. The columns are separated into two main groups. The first group reports the *net* percentage of patients sent to the top of the wait-list due to indirect exchanges (the percentages are taken with respect to all paired patients). This is a *net upgrade burden*, i.e., the difference between the patients added at the top of the list and the living-donor kidneys made available for the wait-list patients. The second group reports (again as a percentage of all paired patients) the rate of paired patients who nonetheless are sent to the cadaveric waiting list because the patient is not assigned a living-donor kidney.

In Table V, we report the sizes of cycles and w-chains under each mechanism. The columns are divided into two groups for cycles and w-chains. Each group reports the number, the average length, the average maximum (per group, over all 100 trials) length of cycles/w-chains, and the length of the *longest* cycle/w-chain encountered in all 100 trials. The lengths of cycles/w-chains are measured in pairs. Standard errors are reported below the estimates.

Next we highlight a number of these results.

1. A transition to the TTCC mechanism will significantly improve the utilization rate of potential unrelated living-donor kidneys: assuming a population size of 100 pairs, while approximately 55 percent of potential living-donor

TABLE III
NUMBER OF TRANSPLANTS AND QUALITY OF MATCH FOR N = 30, N = 100, AND N = 300

Pop. size	Pref.	Exchange regime	Total trans. %	Own donor trans. %	Trade %	Wait-list upgrade %	HLA mis.
n = 30	Wait-list 0%						
	All	None	54.83 (8.96)	54.83 (8.96)	0 (0)	0 (0)	4.79 (0.25)
	All	Paired	68.50 (9.90)	54.83 (8.96)	13.67 (9.40)	0 (0)	4.78 (0.24)
	Rational	TTC	82.47 (10.14)	23.03 (9.44)	59.43 (13.57)	0 (0)	4.16 (0.22)
	Cautious	TTC	81.07 (10.02)	34.17 (11.27)	46.90 (13.96)	0 (0)	4.29 (0.23)
	Wait-list 40%						
	All	Paired/ind.	68.50 (9.90)	54.83 (8.96)	13.67 (9.40)	13.20 (6.73)	4.78 (0.24)
	Rational	TTCC e	84.70 (8.49)	21.23 (9.60)	63.47 (12.39)	6.37 (4.88)	4.17 (0.22)
	Cautious	TTCC e	83.57 (8.53)	32.93 (10.98)	50.63 (12.54)	6.13 (4.39)	4.29 (0.22)
	n = 100	Wait-list 0%					
All		None	54.79 (4.48)	54.79 (4.48)	0 (0)	0 (0)	4.83 (0.14)
All		Paired	73.59 (4.97)	54.79 (4.48)	18.80 (3.81)	0 (0)	4.82 (0.11)
Rational		TTC	87.85 (4.54)	11.51 (3.44)	76.34 (5.45)	0 (0)	3.72 (0.10)
Cautious		TTC	87.23 (4.73)	24.01 (4.48)	63.22 (5.46)	0 (0)	3.86 (0.11)
Wait-list 40%							
All		Paired/ind.	73.59 (4.97)	54.79 (4.48)	18.80 (3.81)	10.24 (3.07)	4.82 (0.11)
Rational		TTCC e	89.44 (3.85)	10.29 (3.26)	79.15 (4.40)	3.96 (1.97)	3.71 (0.10)
Cautious		TTCC e	88.97 (4.17)	22.81 (4.83)	66.16 (4.79)	4.72 (2.60)	3.85 (0.11)
n = 300		Wait-list 0%					
	All	None	53.92 (2.82)	53.92 (2.82)	0 (0)	0 (0)	4.81 (0.08)
	All	Paired	75.03 (2.72)	53.92 (2.82)	21.11 (2.51)	0 (0)	4.81 (0.07)
	Rational	TTC	91.05 (3.35)	5.72 (1.28)	85.32 (3.61)	0 (0)	3.29 (0.06)
	Cautious	TTC	90.86 (3.31)	15.36 (2.20)	75.51 (4.07)	0 (0)	3.40 (0.06)
	Wait-list 40%						
	All	Paired/ind.	75.03 (2.72)	53.92 (2.82)	21.11 (2.51)	9.77 (1.73)	4.81 (0.07)
	Rational	TTCC e	92.29 (2.98)	5.00 (1.29)	87.29 (3.05)	3.02 (1.36)	3.29 (0.06)
	Cautious	TTCC e	92.17 (2.93)	14.42 (2.10)	77.75 (3.26)	3.19 (1.40)	3.39 (0.06)

TABLE IV
 ABO COMPOSITION OF NET WAIT-LIST UPGRADES AND OTHER WAIT-LIST PATIENTS AS PERCENTAGE OF N = 30, N = 100, AND N = 300

Pop. size	Pref.	Exchange regime	Net wait-list upgrades %					Other wait-list additions %					
			O	A	B	AB	Total	O	A	B	AB	Total	
n = 30	Wait-list 0%	All	0	0	0	0	0	28.83	10.13	5.77	0.43	45.17	
		Paired	0	0	0	0	0	24.50	4.07	2.80	0.13	31.50	
		TTC	0	0	0	0	0	12.47	3.17	1.77	0.13	17.53	
	Wait-list 40%	All	0	0	0	0	0	13.77	3.17	1.87	0.13	18.93	
		Paired/ind.	9.10	-6.33	-1.23	-1.53	0	14.67	2.20	1.37	0.07	18.30	
		TTC e	5.13	-2.47	-1.07	-1.60	0	7.13	1.17	0.63	0.00	8.93	
	n = 100	Wait-list 0%	All	0	0	0	0	0	27.64	11.18	5.86	0.53	45.21
			Paired	0	0	0	0	0	21.89	2.90	1.57	0.05	26.41
			TTC	0	0	0	0	0	9.30	1.90	0.91	0.04	12.15
Wait-list 40%		All	0	0	0	0	0	9.78	2.00	0.95	0.04	12.77	
		Paired/ind.	8.25	-5.27	-1.70	-1.28	0	13.42	1.62	1.08	0.05	16.17	
		TTC e	3.49	-1.17	-1.07	-1.25	0	5.71	0.46	0.42	0.01	6.60	
n = 300		Wait-list 0%	All	0	0	0	0	0	5.50	0.41	0.40	0.00	6.31
			Paired	0	0	0	0	0	28.32	11.42	5.79	0.55	46.08
			TTC	0	0	0	0	0	21.70	2.19	1.06	0.02	24.97
	Wait-list 40%	All	0	0	0	0	0	7.07	1.34	0.52	0.01	8.95	
		Paired/ind.	8.40	-5.65	-1.42	-1.32	0	13.21	1.36	0.63	0.01	15.21	
		TTC e	2.77	-0.87	-0.63	-1.27	0	4.28	0.23	0.18	0.00	4.69	
	n = 300	Wait-list 0%	All	0	0	0	0	0	4.28	0.21	0.15	0.00	4.64
			Paired	0	0	0	0	0	7.23	1.37	0.52	0.01	9.14
			TTC	0	0	0	0	0	7.23	1.37	0.52	0.01	9.14
Wait-list 40%		All	8.40	-5.65	-1.42	-1.32	0	13.21	1.36	0.63	0.01	15.21	
		Paired/ind.	2.77	-0.87	-0.63	-1.27	0	4.28	0.23	0.18	0.00	4.69	
		TTC e	2.94	-0.96	-0.63	-1.34	0	4.28	0.21	0.15	0.00	4.64	

Net wait-list upgrades is the difference between the number of upgraded patients of certain blood type and the number of live donor kidneys sent to the wait-list of the same blood type.

TABLE V
 PROPERTIES OF CYCLES AND W-CHAINS FOR N = 30, N = 100, AND N = 300

		Cycles				W-chains				
Pop. size	Pref.	Exchange regime	Length (as pairs)			Length (as pairs)			Longest	
			Number	Mean	Max.	Longest	Number	Mean		Max
n = 30	Wait-list 0%	None	16.45 (2.69)	1 (0)	1 (0)	1	0 (0)	—	—	
		Paired	18.50 (2.46)	1.11 (0.08)	1.90 (0.30)	2	0 (0)	—	—	
		TTC	14.65 (2.22)	2.95 (0.43)	5.65 (1.15)	9	0 (0)	—	—	
		TTC e	16.79 (2.71)	2.59 (0.41)	5.42 (1.17)	10	0 (0)	—	—	
	Wait-list 40%	Paired/ind.	18.50 (2.46)	1.11 (0.08)	1.90 (0.30)	2	3.96 (2.02)	1 (0)	1 (0)	
		TTC e	13.37 (2.50)	3.04 (0.47)	5.65 (1.15)	9	1.91 (1.46)	1.86 (0.83)	2.55 (1.35)	
		TTC e	15.74 (2.75)	2.62 (0.44)	5.42 (1.17)	10	1.84 (1.32)	1.81 (0.68)	2.26 (1.01)	
		Cautious								6
	n = 100	Wait-list 0%	None	54.79 (4.48)	1 (0)	1 (0)	1	0 (0)	—	—
			Paired	64.19 (4.33)	1.15 (0.03)	2.00 (0.00)	2	0 (0)	—	—
			TTC	36.53 (3.44)	4.22 (0.41)	10.14 (1.65)	15	0 (0)	—	—
			TTC e	45.53 (4.29)	3.51 (0.33)	9.82 (1.81)	17	0 (0)	—	—
Wait-list 40%		Paired/ind.	64.19 (4.33)	1.15 (0.03)	2.00 (0.00)	2	10.24 (3.07)	1 (0)	1 (0)	
		TTC e	32.91 (4.16)	4.41 (0.43)	10.14 (1.65)	15	3.96 (1.97)	2.40 (0.93)	3.96 (1.97)	
		TTC e	42.00 (5.43)	3.61 (0.39)	9.82 (1.81)	17	4.72 (2.60)	2.17 (0.65)	3.76 (1.64)	
		Cautious								10
n = 300		Wait-list 0%	None	161.76 (8.47)	1 (0)	1 (0)	1	0 (0)	—	—
			Paired	193.42 (7.42)	1.16 (0.02)	2.00 (0.00)	2	0 (0)	—	—
			TTC	79.54 (4.91)	5.98 (0.43)	16.84 (2.41)	26	0 (0)	—	—
			TTC e	102.14 (7.00)	4.89 (0.33)	16.00 (2.00)	22	0 (0)	—	—
	Wait-list 40%	Paired/ind.	193.42 (7.42)	1.16 (0.02)	2.00 (0.00)	2	29.30 (5.18)	1 (0)	1 (0)	
		TTC e	71.17 (5.88)	6.30 (0.46)	16.84 (2.41)	26	9.05 (4.07)	2.88 (0.78)	6.47 (2.40)	
		TTC e	93.76 (7.51)	5.05 (0.33)	16.00 (2.00)	22	9.57 (4.20)	2.65 (0.72)	5.99 (2.23)	
		Cautious								12

Cycles of one pair length with an incompatible donor-patient pair are not counted.

kidneys are utilized under the no-exchange regime, this rate increases to 73.5 percent under the paired-kidney-exchange, and to 88–89.5 percent under the TTCC mechanism. The efficiency gain gets larger as the population size grows further.

2. A transition to TTCC significantly decreases the HLA mismatch, especially for the large populations.
3. Under the TTCC mechanism, average/maximal sizes of cycles/w-chains increase as the population grows, although the increase is less than proportional.
4. Type O patients *without* living donors *benefit* from the TTCC mechanism. The TTCC mechanism significantly reduces the incidence of type O patients with potential donors who are forced to rely on the cadaveric waiting list because of an incompatibility. Consider a population size of 100 donor-recipient pairs (with the distribution of blood types, etc. of the UNOS data set). When no exchange is allowed, on average 27.6 type O patients who have a willing donor nevertheless join the waiting list for a cadaver kidney because they are incompatible (by blood type or positive crossmatch) with their donor. This rate reduces to 21.9 under the paired-kidney-exchange and further to 5.5–5.7 under the TTCC mechanism. That means out of 100 patients with living donors, 16.2–16.4 patients with O blood types drop out of competition for cadaver kidneys as a result of a change from paired-kidney-exchange to the TTCC mechanism. The corresponding cost of this change to type O patients with no living donors is that only 3.5–4.2 type O patients with living donors are moved to the head of the cadaver queue.⁷ So the reduction in demand for O type cadaver kidneys is much larger than the number of patients who are inserted at the head of the queue.⁸

7. This benefit/cost rate is significantly better than the benefit/cost rate that results from a change from paired-kidney-exchange to paired/indirect-kidney-exchange because, in an indirect exchange, only occasionally will a type O living donor kidney be sent to the waiting list to compensate for the type O patient who receives priority on the queue. In our simulations with 100 paired patients, a transition from paired-kidney-exchange to paired/indirect-kidney-exchange drops $21.89 - 13.42 = 8.47$ patients from the waiting list at a cost of moving 8.25 paired patients to the head of the cadaver queue.

8. We emphasize that the welfare implications of this for type O patients on the cadaver queue is much clearer when we think of patients who enter the queue after kidney exchange is already well established. There can still be important

More detail on the simulations can be found in Roth, Sönmez, and Ünver [2003].

V. CONCLUSION

The TTCC mechanism is motivated by the present small-scale pilot paired and indirect kidney exchange programs. We have concentrated here on its advantages over the simplest kinds of exchange. But, as we begin to talk with members of the transplantation community about first steps toward implementing such exchanges, it is clear that there are obstacles to be overcome. Some are the same obstacles that have kept existing exchange programs very small. Among these is the absence of registries of incompatible or poorly matched patient-donor pairs.⁹ Also, exchanges require coordination of multiple operating rooms and surgical teams, two for each patient-donor pair, so larger exchanges will require more coordination.

As a registry starts to be assembled, it may be that the main initial advantage will simply be to allow paired exchanges to be conducted more often. However, as the registry grows larger, and the practical difficulties are overcome, we have seen that there will be additional benefits to be reaped from more flexible forms of exchange that enlarge the set of possible exchanges. Compared with simple paired and indirect exchanges, the wider exchange implemented by the TTCC mechanism creates additional welfare gains in several ways. First, allowing longer cycles of exchange will allow some transplantations that could not be arranged with pairwise exchanges, and it will increase the scope for improving the quality of the resulting matches. And by allowing more live donations, it will reduce the competition for cadaver kidneys. Second, longer chains for combined indirect and paired exchange

welfare concerns about individual patients during the transition to an indirect exchange regime, since an O type patient who has already been on the waiting list for three years receives relatively less benefit from the reduction in new O type patients joining the list than does a new patient, but suffers increased waiting time when an O type patient is given higher priority on the queue. What the simulations suggest is that, once a kidney exchange regime is up and running, type O patients without a living donor who enter the queue for a cadaver kidney will be helped, not harmed, by the fact that exchanges are being made.

9. Such a registry would need to include tissue typing data, and present medical privacy laws mean that this will be difficult to collect *ex post*. We have begun exploratory discussions with our medical colleagues in New England about the possibility of developing a program to identify incompatible and poorly matched patient-donor pairs, and collect tissue-typing data from potential donors.

will allow an indirect exchange to benefit more than one patient-donor pair, and by doing so will also increase the number of live donations. And third, by increasing the number of O type patients who can receive live donations, and by managing the flow of kidneys and patients to the cadaver queue, this can be done in ways that help O type patients who have no live donor.¹⁰

In summary, the design of practical exchange mechanisms is the “engineering” part of economic theory, and must deal with constraints omitted from more abstract endeavors. The organization of a kidney exchange faces some of the most stringent constraints we have encountered, arising from social/legal/ethical concerns, as well as from the practical requirements of kidney transplantation and patient care. In the future, some of those constraints may be relaxed, e.g., through advances in dealing with immunological incompatibilities, or in using xenotransplants (animal organs) to relieve the organ shortage, or through some other way of radically increasing organ availability. In the meantime, increasing kidney exchange among willing donor-recipient pairs offers a way to benefit those pairs who are not well matched, and, by increasing live organ donation and reducing competition for cadaver organs, also benefiting patients who do not have live donors. This benefit is sufficiently widespread that it helps even the most vulnerable patients, the type O patients without a live donor.

APPENDIX: PROOFS

Proof of Lemma 1. Consider a graph where each patient points toward either a kidney or w , and each kidney points to its paired recipient. Suppose that there is no cycle. Consider an arbitrary pair (k_i, t_i) . Start with kidney k_i , and follow the path in the graph. Since there are no cycles, no kidney or patient can be encountered twice. Hence by the finiteness of pairs, the path will

10. To eliminate or reduce the adverse affect of indirect exchange programs on patients with no living donors, Zenios, Woodle, and Ross [2001] propose preferential selection of O blood-type paired donors for patients with multiple potential donors who wish to participate in indirect exchange programs. Their proposal is consistent with a direct extension of the TTCC mechanism to a model with multiple potential donors when the flexibility on chain selection is used to increase the inflow of O blood-type kidneys to the cadaveric waiting list. See Roth, Sönmez, and Ünver [2003] for more on exchange when patients may have multiple potential donors.

terminate at w . This is the w -chain initiated by pair (k_i, t_i) completing the proof. \diamond

Proof of Theorem 1. Let the TTCC mechanism be implemented with a chain selection rule such that any w -chain selected at a nonterminal round remains in the procedure and the kidney at its tail remains available for the next round. Any patient whose assignment is finalized in Round 1 has received his top choice and cannot be made better off. Any patient whose assignment is finalized in Round 2 has received his top choice among the kidneys not already assigned as part of an exchange (since chains are not removed, so the kidney at their tail remains available), and cannot be made better off without hurting a patient whose assignment was finalized in Round 1. Proceeding in a similar way, no patient can be made better off without hurting a patient whose assignment is finalized in an earlier round. Therefore, TTCC mechanism selects a Pareto-efficient matching at any given time provided that w -chains are removed at the termination. \diamond

We will prove Theorem 2 for the chain selection rule a and for a class of “priority chain selection rules” that covers rules d, e, and f. Under this class each ordering of patient-donor pairs together with a fixed pair defines a chain selection rule, and it is given as follows: order donor-patient pairs in a single priority list. Fix a pair (k_j, t_j) . Whenever a w -chain is to be selected, select the w -chain starting with the highest priority pair (k_i, t_i) , and remove the w -chain if the pair (k_i, t_i) has *strictly* higher priority than the fixed pair (k_j, t_j) , and keep it until termination otherwise.

First, we prove the following lemma which will be useful for the proof of Theorem 2.

LEMMA 2. Consider the TTCC mechanism implemented with a priority-based chain selection rule. Fix the stated preferences of all patients except patient t_i at P_{-i} . Suppose that in the algorithm the assignment of patient t_i is finalized at Round s under P_i and at Round s' under P'_i . Suppose that $s \leq s'$. Then the remaining active patients and unassigned kidneys at the beginning of Round s are the same, whether patient t_i announces P_i or P'_i .

Proof of Lemma 2. Patient t_i fails to participate in a cycle or a selected w -chain prior to Round s under either preference. Therefore, at any round prior to Round s not only the highest priority active patient is the same, whether patient t_i announces

P_i or P'_i , but also the same cycles/w-chains form, and in case there are no cycles, the same w-chain is selected, whether patient t_i announces P_i or P'_i . Hence the remaining active patients and unassigned kidneys at the beginning of Round s are the same, whether patient t_i announces P_i or P'_i . \diamond

Proof of Theorem 2. We first consider the chain selection rule a. Recall that for each patient t_i , the relevant part of preference P_i is the ranking up to k_i or w , whichever is more preferred. Given the preference profile $(P_i)_{i=1}^n$, construct a new preference profile $(P'_i)_{i=1}^n$ as follows: for each patient t_i with $k_i P_i w$, let $P'_i = P_i$. For each patient t_i with $w P_i k_i$, construct P'_i from P_i by swapping the ranking of k_i and w . Note that $k_i P'_i w$ for each patient t_i and because the relevant part of preferences are the more preferred of k_i and w , $\langle \{(k_i, t_i)\}_{i=1}^n, (P'_i)_{i=1}^n \rangle$, is a housing market. Let μ denote the outcome of the TTC mechanism for this housing market, and construct matching ν from matching μ as follows: if $P'_i \neq P_i$ and $\mu(t_i) = k_i$, then $\nu(t_i) = w$, otherwise, $\nu(t_i) = \mu(t_i)$. The key observation is that ν is the outcome of the TTCC mechanism when it is implemented with the minimal w-chain selecting chain selection rule. Therefore, by Roth [1982], a patient can never receive a more preferred kidney by a preference misrepresentation. He can receive the wait-list option w by a misrepresentation but cannot profit from it. That is because the TTCC mechanism never assigns a patient a kidney that is inferior to w . Hence TTCC is strategy-proof with this choice of chain selection rule.

Next consider any of the priority-based chain selection rules. Consider a patient t_i with true preferences P_i . Fix an announced preference profile P_{-i} for all other patients. We want to show that revealing his true preferences P_i is at least as good as announcing any other preferences P'_i under the TTCC mechanism. Let s and s' be the rounds at which patient t_i leaves the algorithm under P_i and P'_i , respectively. We have two cases to consider.

Case 1. $s < s'$. By Lemma 2 the same kidneys remain in the algorithm at the beginning of Round s whether patient t_i announces P_i or P'_i . Moreover, patient t_i is assigned his top choice remaining at Round s under P_i . Therefore, his assignment under P_i is at least as good as his assignment under P'_i .

Case 2. $s \geq s'$. After announcing P'_i , the assignment of patient t_i is finalized either by joining a cycle, or by joining a selected w-chain. We will consider the two cases separately.

Case 2a. The assignment of patient t_i is finalized by joining a cycle under P'_i .

Let $(k^1, t^1, k^2, \dots, k^r, t_i)$ be the cycle patient t_i joins, and thus k^1 be the kidney he is assigned under P'_i . Next suppose that he reveals his true preferences P_i . Consider Round s' . By Lemma 2 the same active patients and available kidneys remain at the beginning of this round whether patient t_i announces P'_i or P_i . Therefore, at Round s' , kidney k^1 points to patient t^1 , patient t^1 points to kidney k^2 , \dots , kidney k^r points to patient t_i . Moreover, they keep on doing so as long as patient t_i remains. Since patient t_i truthfully points to his best remaining choice at each round, he either receives a kidney better than kidney k^1 or eventually points to kidney k^1 , completes the formation of cycle $(k^1, t^1, k^2, \dots, k^r, t_i)$, and gets assigned kidney k^1 .

Case 2b. The assignment of patient t_i is finalized by joining a selected w-chain under P'_i .

Let $(k^1, t^1, k^2, \dots, k^r, t_i = t^r, k^{r+1}, \dots, k^{r+m}, t^{r+m})$ be the selected w-chain patient t_i joins, where $r \geq 1$ and $m \geq 0$, under P'_i . Therefore, under P'_i , patient t_i is assigned the kidney k^{r+1} if $m \geq 1$, and the wait-list option w if $m = 0$. Also note that, given the considered class of priority-based chain selection rules, pair (k^1, t^1) is the highest priority pair in Round s' . Next suppose that patient t_i reveals his true preferences P_i . Consider Round s' . By Lemma 2 the same active patients and available kidneys remain at the beginning of this round whether patient t_i announces P'_i or P_i . We will complete the proof by showing that, upon announcing his truthful preferences P_i , the assignment of patient t_i is finalized in Round s' and thus he is assigned his top choice available at the beginning of Round s' : recall that for this case there is no cycle in Round s' when patient t_i announces P'_i . Therefore, when he announces his true preferences P_i , either there is no cycle in Round s' or there is one cycle that includes him. If it is the latter, then his assignment is finalized in Round s' , and we are done. Otherwise, each pair initiates a w-chain by Lemma 1, and one of these w-chains has to be selected. By the choice of a priority-based chain selection rule, this will be the w-chain that starts with the highest priority pair (k^1, t^1) . But the path starting with kidney k^1 passes through patient t_i and therefore the selected

w-chain includes patient t_i . Hence in this case as well his assignment is finalized in Round s' completing the proof.

HARVARD UNIVERSITY
KOÇ UNIVERSITY
KOÇ UNIVERSITY

REFERENCES

- Abdulkadiroğlu, Atila, and Tayfon Sönmez, "House Allocation with Existing Tenants," *Journal of Economic Theory*, LXXXVIII (1999), 233–260.
- Abecassis, Michael, et al. *The Journal of the American Medical Association*, CCLXXXIV (2002), 2919–2926.
- Becker, Gary S., and Julio J. Elias, "Introducing Incentives in the Market for Living Organ Donations." Working paper, University of Chicago, 2002.
- Mandal, Aloke K., Jon J. Snyder, David T. Gilbertson, Allan J. Collins, and John R. Siliksen, "Does Cadaveric Donor Renal Transplantation Ever Provide Better Outcomes Than Live-Donor Renal Transplantation?" *Transplantation*, LXXV (2003), 494–500.
- Milgrom, Paul, *Putting Auction Theory to Work*, (Cambridge University Press, 2004).
- Nelson, Edward W., James E. Childress, Jennie Perryman, Victor Robards, Albert Rowan, Michael S. Seely, Sylvester Sterioff, and Mary Rouelli Swanson, "Financial Incentives for Organ Donation: A Report of the Payment Subcommittee," United Network for Organ Sharing Ethics Committee, 1993.
- Olson, Elizabeth, "A Tale of 6 Kidneys, or She Gave Him Her Love, He Gave Her His Kidney," *New York Times*, Saturday, August 2, A7, 2003.
- Opelz, Gerhard, "Impact of HLA Compatibility on Survival of Kidney Transplants from Unrelated Live Donors," *Transplantation*, LXIV (1997), 1473–1475.
- Rapaport, Felix T., "The Case for a Living Emotionally Related International Kidney Donor Exchange Registry," *Transplantation Proceedings*, XVIII (1986), 5–9.
- Ross, Lainie Friedman, David T. Rubin, Mark Siegler, Michelle A. Josephson, J. Richard Thistlethwaite, Jr., and E. Steve Woodle, "Ethics of a Paired-Kidney-Exchange Program," *The New England Journal of Medicine*, CCCXXXVI (1997), 1752–1755.
- Ross, Lainie Friedman, and E. Steve Woodle, "Ethical Issues in Increasing Living Kidney Donations by Expanding Kidney Paired Exchange Programs," *Transplantation*, LXIX (2002), 1539–1543.
- Roth, Alvin E., "Incentive Compatibility in a Market with Indivisibilities," *Economics Letters*, IX (1982), 127–132.
- , "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics," *Econometrica*, LXX (2002), 1341–1378.
- Roth, Alvin E., and Elliot Peranson, "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review*, LXXXIX (1999), 748–780.
- Roth, Alvin E., and Andrew Postlewaite, "Weak versus Strong Domination in a Market with Indivisible Goods," *Journal of Mathematical Economics*, IV (1977), 131–137.
- Roth, Alvin E., Tayfon Sönmez, and M. Utku Ünver, "Kidney Exchange," NBER Working Paper No. 10002, September, 2003.
- Sade, R. M., "Assessing Initiatives to Increase Organ Donations," Prepared Witness Testimony, The House Committee on Energy and Commerce, W. J. "Billy" Tauzin, Chairman, Subcommittee on Oversight and Investigations, June 3, 2003.
- Shapley, Lloyd, and Herbert Scarf, "On Cores and Indivisibility," *Journal of Mathematical Economics*, (1974), 23–28.
- Zenios, Stefanos, "Health Care Applications of Optimal Control Theory," Ph.D. Thesis, Massachusetts Institute of Technology, 1996.
- Zenios, Stefanos, E. Steve Woodle, and Lainie Friedman Ross, "Primum non Nocere: Avoiding Increased Waiting Times for Individual Racial and Blood-Type Subsets of Kidney Wait List Candidates in a Living Donor/Cadaveric Donor Exchange Program," *Transplantation*, LXVI (2001), 406–407.