
MŰHELYTANULMÁNYOK

DISCUSSION PAPERS

MT-DP – 2019/15

Optimal Kidney Exchange with Immunosuppressants

HARIS AZIZ – ÁGNES CSEH

Discussion papers
MT-DP – 2019/15

Institute of Economics, Centre for Economic and Regional Studies,
Hungarian Academy of Sciences

KTI/IE Discussion Papers are circulated to promote discussion and provoke
comments.

Any references to discussion papers should clearly state that the paper is preliminary.
Materials published in this series may subject to further publication.

Optimal Kidney Exchange with Immunosuppressants

Authors:

Haris Aziz
senior lecturer
at the School of Computer Science and Engineering, UNSW
and the team leader
of the Algorithmic Decision Theory group
haris.aziz@unsw.edu.au

Ágnes Cseh
research fellow
Institute of Economics
Centre for Economic and Regional Studies
Hungarian Academy of Sciences
cseh.agnes@krtk.mta.hu

September 2019

Optimal Kidney Exchange with Immunosuppressants

Haris Aziz - Ágnes Cseh

Abstract

Potent immunosuppressant drugs suppress the body's ability to reject a transplanted organ up to the point that a transplant across blood- or tissue-type incompatibility becomes possible. In contrast to the standard kidney exchange problem, our setting also involves the decision about which patients receive from the limited supply of immunosuppressants that make them compatible with originally incompatible kidneys. We firstly present a general computational framework to model this problem. Our main contribution is a range of efficient algorithms that provide flexibility in terms of meeting meaningful objectives. We also show that these algorithms satisfy desirable axiomatic and strategic properties.

Keywords: kidney exchange, immunosuppressant drugs, matching, complexity

JEL classification: C63, C78

Optimális vesezsere immunszupresszáns gyógyszerek segítségével

Haris Aziz – Cseh Ágnes

Összefoglaló

Bizonyos erős immunszupresszáns gyógyszerek olyan mértékben képesek befolyásolni az emberi test idegen szervvel szemben tanúsított befogadókésztségét, hogy a vércsoport szerinti vagy szövettani inkompatibilias áthidalhatóvá válik. A szokványos vesezsereproblémával ellentétben modellünkben döntést kell hoznunk, hogy mely betegek kapják ezeket a korlátozott számban rendelkezésre álló, erős gyógyszereket. Először építünk egy általános, számítástudományi alapokon nyugvó modellt a problémához. Ezen a modellen később konstruálunk egy sokféleféppen módosítható, hatékony algoritmust, ami különböző célok elérését szolgálhatja. Azt is megmutatjuk, hogy ez az algoritmus vonzó aximomatikus és stratégiai tulajdonságokkal is rendelkezik.

Tárgyszavak: vesezsere, immunszupresszáns gyógyszerek, párosítások, komplexitás

JEL kódok: C63, C78

Optimal Kidney Exchange with Immunosuppressants

Haris Aziz* and Ágnes Cseh†

Abstract

Potent immunosuppressant drugs suppress the body’s ability to reject a transplanted organ up to the point that a transplant across blood- or tissue-type incompatibility becomes possible. In contrast to the standard kidney exchange problem, our setting also involves the decision about which patients receive from the limited supply of immunosuppressants that make them compatible with originally incompatible kidneys. We firstly present a general computational framework to model this problem. Our main contribution is a range of efficient algorithms that provide flexibility in terms of meeting meaningful objectives. We also show that these algorithms satisfy desirable axiomatic and strategic properties.

Introduction

The deployment of centralized matching algorithms for efficient exchange of donated kidneys is a major success story of market design (Biró et al. 2019). The standard model for kidney exchange involves information about patients’ compatibility with kidneys in the market. A patient can only be given a kidney that is compatible with the patient. The goal is to enable exchanges of kidneys via a centralized algorithm to satisfy the maximum number of patients.

We consider a new kidney exchange model which has an interesting feature that is informed by significant technological advances in organ transplant. The technology concerns immunosuppressants which if given to a patient can make her receptive to kidneys which she is not receptive to by default (Montgomery et al. 2011). We will refer to the model as *Kidney Exchange with Immunosuppressants (KEI)*. Immunosuppressants (abbreviated as suppressants from here onwards) have been successfully used in Japan and Korea and increasingly being considered in other countries (Chun, Heo, and Hong 2017), even though they are costly and may have side effects.

In this paper, the fundamental research problem that we explore is that of designing mechanisms for kidney ex-

change with immunosuppressants that satisfy desirable computational, incentive and monotonicity properties. A naive way of using suppressants is to clear the classic kidney exchange market without using them and then give suppressant to the patients who are left. However, there can be more efficient ways of giving suppressant to particular patients and then implementing exchanges of kidneys to facilitate as many transplants as possible. At first sight, the two-stage and connected process of using foresight to first giving suppressants to suitable patients and then finding a matching that satisfies suitable social objectives appears to be a complex problem. We design a flexible algorithmic approach for the problem.

Our contributions are as follows.

- We formalize a general model of kidney exchange with immunosuppressants that features compatible, half-compatible, and incompatible kidneys, and which allows for allocations as a result of multi-way exchanges. We then initiate a computational study of kidney exchange with immunosuppressants. Prior mechanism design work on the subject either only allows pairwise exchanges or focuses on a restricted model.
- Our central contribution is modeling important KEI problems in terms of graph algorithms involving minimum cost perfect matchings. Depending on how we set the edge costs in the graph problem, we can find in polynomial time, allocations corresponding to several important objectives. The objectives include maximizing the total number of transplants and given that maximizing the number of compatible transplants. Among the list of objectives captured by our algorithmic framework, we defer the choice of the exact objective to the policy-makers.
- We then focus on the problem where there is an upper bound on the number of suppressants that can be used. We present a polynomial-time algorithm for maximizing the number of transplants for a restricted model that we refer to as the silver bullet model. In the model, once a patient has been given a suppressant, then the patient can take any kidney. For our general model in which certain kidneys are inherently incompatible, we show that the problem of maximizing the number of transplants polynomial-time

*UNSW Sydney and Data61 (CSIRO), Australia, haris.aziz@unsw.edu.au

†Hungarian Academy of Sciences, Centre for Economic and Regional Studies, Institute of Economics, Hungary cseh.agnes@krtk.mta.hu

reduces to an interesting generalized matching problem whose complexity has been open for years.

- We also study axioms for our model. If patients prefer a transplant without potent suppressants and there is a central budget for these suppressants, then we show that strategyproofness, individual rationality, and a very weak form of efficiency called unanimity are incompatible. We also show that this impossibility disappears if we consider coarse preferences or do not impose hard upper bounds on the number of suppressants used.

Some of the techniques that we use such as to capture strong individual rationality, achieve strategyproofness, handle pairwise exchanges etc. are of independent interest and can be applied to a host of other problems in matching markets.

Related Work

Kidney exchange is one of the major research topics in matching market design (Sönmez and Ünver 2011; Abraham, Blum, and Sandholm 2007; Hatfield 2005; Roth, Sönmez, and Ünver 2005a; Biró, Manlove, and Rizzi 2009). In many of the papers, the algorithms only allow exchange cycles of limited size due to logistical and other constraints. In this paper, we allow exchange cycles of any size. Note that for any exchange cycle bounds that are three or more, even the kidney exchange problem in the traditional model without immunosuppressants is NP-hard (Abraham, Blum, and Sandholm 2007).

The use of suppressants to facilitate more efficient kidney exchange has been discussed in medical circles (see, e.g. Abramowicz et al. (2018)). The two market design papers directly relevant to our work are ones where kidney exchange with immunosuppressants has been mathematically modeled (Chun, Heo, and Hong 2017; Andersson and Kratz 2019). Chun, Heo, and Hong (2017) prove a couple of impossibility results as well as an exchange mechanism with some desirable monotonicity properties. They assume that once suppressants are administered to a patient, she can take a kidney from any donor. We consider a more general model in which half-compatibility is specific to particular patient-donor pairs.

Andersson and Kratz (2019) consider a more general model in which only certain donor-patient pairs can be made compatible after giving suppressants to the patient. They focus on *pairwise* kidney exchange (Roth, Sönmez, and Ünver 2005b) and provide some properties of half-compatibility priority matchings and present an efficient computational method for finding them. They also analyze the welfare implications of pairwise kidney exchange over the blood group barrier, both theoretically and via a simulation study. Our paper goes beyond the pairwise exchange model and presents results that can use longer cycles. We also show how our general results that do not impose restrictions on the exchange cycles can be adapted for the case of pairwise exchange.

Model and Concepts

A kidney exchange market is a tuple (P, D, C, H, I) where $P = \{p_1, p_2, \dots, p_n\}$ is a set of n patients (agents) and D is the set of donors. Some patients and donors come in pairs; others come single. If a patient p_i enters the market together with her donor, then the donor d_i will be equipped with the same index i .

Each patient p_i partitions the donors D into sets C_i , H_i , and I_i . The set C_i is the set of donors whom patient p_i is compatible with. The set H_i is the set of donors p_i is half-compatible with. Half-compatibility means that if a suppressant is given p_i , then p_i can accept a kidney from any donor in H_i . Donors in the set I_i are incompatible with patient p_i even if p_i is given a suppressant. These partitions at each patient form the collections of sets $C = (C_1, \dots, C_n)$, $H = (H_1, \dots, H_n)$, and $I = (I_1, \dots, I_n)$ in the input. For simplicity, we assume that compatible patient-donor pairs are not present, because their transplant is arranged outside of the system.

An *allocation* assigns each patient p_i at most one donor who is either in C_i or in H_i . Patients who are assigned a half-compatible donor receive suppressants.

We consider three models.

1. **BM (Baseline model)**: for each $p_i \in P$, $H_i = \emptyset$.
2. **SBM (Silver bullet model)**: for each $p_i \in P$, $I_i = \emptyset$.
3. **GM (General model)**.

The baseline model coincides with the traditional kidney exchange model in which immunosuppressants are not considered. SBM is the model in which we assume that if a patient is given a suppressant then she will be able to receive any kidney in the market. GM is the general model that also allows for some kidneys being inherently incompatible for a patient even if she has been given suppressants. Unless specified, we will focus on GM. In some cases, we will present some results that hold for the silver bullet model.

As far as a patient or the social designer is concerned, there are two types of preferences. We will treat matching with incompatible donors to be infeasible.

1. **Coarse preferences**: a patient is indifferent between a compatible donor and a half-compatible donor along with a suppressant, and prefers both options over no transplant at all.
2. **Refined preferences**: a patient prefers compatible donors over half-compatible donors along with a suppressant, which are preferred over no transplant at all.

Coarse preferences have the underlying assumption that a patient has no significant cost (in terms of money or side-effects) when receiving a half-compatible kidney.

Based on the preference relation one can define concepts such as Pareto optimality. Chun, Heo, and Hong (2017) considered SBM and coarse preferences. They also consider refined preferences when defining a monotonicity property. Andersson and Kratz (2019) considered GM and refined preferences.

A patient who is assigned a compatible donor or a half-compatible donor along with a suppressant is referred to as

satisfied. Our general goal is to satisfy the maximum number of patients while minimizing the need of suppressants. We will consider the following feasibility condition for all allocations, which captures a natural individual rationality requirement: either a patient donates her donor’s kidney to the market and gets a strict improvement or she and her donated kidney are not part of any allocation. We will refer to this condition as *strong individual rationality (S-IR)*. Notice that a patient who enters the market with a half-compatible donor improves her situation if she is assigned her own or another half-compatible donor along with a suppressant. Also, S-IR implies that a donor assigned to a patient will only donate a kidney if her patient also receives one. A weaker requirement is *individual rationality (IR)* whereby no patient who has a compatible kidney donor ends up with no transplant.

Example 1. Consider a kidney exchange problem in which there are three patients p_1, p_2, p_3 with corresponding donors d_1, d_2, d_3 . No patient’s donor has a kidney compatible with the patient. Patient p_1 finds the kidney of d_2 compatible, while $d_3 \in H_2$ and $d_1 \in H_3$. The problem is captured in Figure 1.

If immunosuppressants are not allowed, then no patient will be able to get a kidney without violating S-IR. Suppose the system has 2 suppressants available. In that case, one suppressant can be given p_2 and another to p_3 . Then p_1 can take a compatible kidney of d_2 , p_2 takes a half-compatible kidney of d_1 and p_3 takes a half-compatible kidney of d_3 .




Figure 1: A bipartite matching view of KEI. The dashed lines indicate half-compatible edges. Solid edges indicate compatibility edges. Dotted edges indicate a patient-donor pair.

A general graph theoretic approach

We construct a general bipartite matching based model capturing the most basic features of kidney exchange markets. It guarantees that each donor gives at most one kidney, each patient receives at most one kidney, and the donor in a pair is only part of the exchange if her patient receives a kidney. This framework gives us a set of feasible solutions for the problem. Then, by adding edge weights to the graph and finding a maximum weight matching, an optimal solution can be calculated. We specify a set of possible edge weights that can serve a large variety of goals of the decision maker, such as cost-efficiency or saving as many lives as possible.

Matching model

We build a bipartite graph to the instance (P, D, C, H, I) , see Figure 2. For convenience, we distinguish between patients with and without a related donor, who will form the sets P_2 and P_1 , respectively. Analogously, D_1 is the set of altruistic donors, while donors in D_2 enter the market along with their related patient in P_2 . The vertices of our graph are of three types.

- To each patient $p_i \in P$, we construct a patient vertex p_i .
- To each donor $d_i \in D$, we construct a donor vertex d_i .
- We construct a dummy donor vertex d_j for each patient $p_j \in P_1$.

If a donor-patient pair who applies for the scheme together, then the patient and donor and given the same index: we refer to them as p_i and d_i for some fixed i . Patients without a donor share the same index with their dummy donor vertex. Altruistic donors in the system do not have such a fixed-index pair.

The edges of the graph are as follows.

- Each patient $p_i \in P_2$ is connected to her related donor $d_i \in D_2$ via a *private edge*.
- Each patient $p_j \in P_1$ is connected to her dummy donor d_j via a private edge.
- A donor d_i has a *compatible* edge to a patient p_j , where i might be equal to j , if $d_i \in C_i$.
- A donor d_i has a *half-compatible* edge to a patient p_j , where i might be equal to j , if $d_i \in H_i$.

The three kinds of edges will play distinct roles when assigning weights to them. Private edges represent no transplant, while compatible and half-compatible edges stand for compatible and half-compatible transplants taking place. Notice that a patient and her donor forming a half-compatible pair are connected by two parallel edges, one private and one half-compatible.

Our goal is to calculate a P -complete matching in the constructed graph. By P -complete, we mean that each patient needs to be matched in the matching. A matching and the corresponding allocation are in trivial one-to-one correspondence with each other. A patient matched to her dummy donor or to her incompatible donor represents no transplant. The matching property ensures that each patient receives one kidney at most, and each donor also donates one kidney at most. Since a patient $p_i \in P_2$ is only connected to $d_i \cup C_i \cup H_i$, and we restrict our attention to P -complete matchings only, p_i is either satisfied or she remains unmatched, together with d_i . P -completeness ensures the following natural consequence of S-IR: either a patient uses her donor’s kidney and gets a strict improvement, or she and her donated kidney are not part of any allocation.

With coarse preferences assumed, there might be restrictions on the compatibility of a kidney given to a patient in a patient-donor pair who enter the market together. To ensure that a half-compatible pair only participates in a cross-donation or a cycle if and only if the patient receives a compatible kidney, we only need to delete the edges running from the patient to all half-compatible donors.




Figure 2: Example instance for our bipartite graph. Here, $P_1 = \{p_1\}$, and thus, d_1 is a dummy donor. The only altruistic donor is d_4 , forming set D_1 . Dotted edges are private, dashed edges mark half-compatible donations, and finally, solid edges mark compatible donations.

Objectives

We offer a variety of different weight functions defined on the edges of our graph. Each weight function serves a justifiable goal, as we argue later.

Table 1 summarizes the options for defining the weight function on each edge (p_i, d_j) , depending on the type of the edge. We assume N to be a sufficiently large integer. In an allocation, we denote the number of patients receiving a kidney from a compatible donor by CO , the number of patients receiving a kidney from a half-compatible donor by HC , while the total number of patients receiving a kidney by $TR = CO + HC$. Our objective functions are to be maximized in the lexicographic sense, e.g. $(TR, -HC)$ maximizes the number of transplants in total, and subject to this, it minimizes the number of half-compatible transplants.

Our objective functions can achieve the following.

- 1) The number of transplants is maximized if each pair chosen for surgery contributes weight 1, while no transplant adds no weight to the matching.
- 2) The number of transplants is maximized in the baseline model, if half-compatible donations are forbidden due to their infinitely large negative weight, and each compatible donation contributes weight 1.
- 3) If a compatible transplant carries a larger weight than the weight of all half-compatible transplants that can be carried out, then the main goal is to maximize the number of compatible donations. Since half-compatible donations do carry some small weight, their number will be maximized, but only subject to the first objective.
- 4) Since half-compatible donations now carry a small negative weight, they are only to be planned if they enable extra compatible transplants. However, any number of half-

compatible donations are welcome if they make only one more compatible donation happen, because we gain a lot in our objective function by adding N just one more time to it.

- 5) If compatible and half-compatible donations both carry a large weight, but the latter ones are somewhat less profitable, then the maximum number of donations will be calculated, and subject to this, as few half-compatible donations will be planned as possible.
- 6) The most general version is when we set an arbitrary weight to each transplant. This weight can express the expected utility in terms of life expectancy, risks, health-care savings, and it can differ for each pair. A max weight solution corresponds to a maximum utility allocation.

Our goal is to compute a maximum weight P -complete matching in the graph. A maximum weight matching can be computed in strongly polynomial time (Munkres 1957). To take care of P -completeness, or equivalently, maximum size, one only needs to apply the standard weight modification, in which each edge gets an additional uniform weight that is larger than the sum of the original weights. As a result, we obtain the following central result.

Theorem 1. *For each of the objectives 1) to 6), there exists a strongly polynomial-time algorithm to find an allocation achieving those objectives.*

Fix upper bound on HC

Immunosuppressants are highly useful to allow half-compatible kidneys to be allocated. However, they are not only extremely expensive but they also have undesirable side-effects. Given these issues, the market designer may wish to specify a fixed upper quota on $HC = h$, and wishes to maximize transplants subject to h . We show that even with an upper bound, we can solve the following central problem.

h -ALLKEI

Input: KEI instance $G = (N, O, e, A)$ and integer m .
 Question: Is there an allocation satisfying all the patients with at most h patients using immunos?

Theorem 2. *h -ALLKEI can be solved in polynomial time even for the general model.*

Proof. Construct the corresponding graph as defined in the previous section, but only keep compatible or half-compatible edges. We want to check whether there exists a P -complete matching M such that $|M \cap E'| \leq h$ and each patient gets a compatible or half-compatible allocation. This question can be answered by solving a simple weighted P -complete matching problem. In the graph with compatible or half-compatible edges only, we give each half-compatible edge weight 0 and each compatible edge weight 1. The weight of any P -complete matching M is $n - |M \cap E'|$. The maximum weight matching in this instance has weight at least $n - h$ if and only if there is an allocation using at most h immunosuppressants. \square

	Objective function	compatible	half-compatible	private
1	TR (coarse preferences)	1	1	0
2	CO (BM, where $HC = 0$)	1	$-\infty$	0
3	(CO, TR)	N	1	0
4	$(CO, -HC)$	N	-1	0
5	$(TR, -HC) \sim (TR, CO)$	N	$N - 1$	0
6	cost-optimal	compatible gain	half-compatible gain	0

Table 1: The variety of possible cost functions serving different goals.

A more general version of h -ALLKEI is the following problem.

h -MAXKEI

Input: KEI instance (P, D, C, H, I) and integers t and h .

Question: Is there an allocation giving at least t patients a compatible donor with at most h patients using immunos?

Our next result is that h -MAXKEI can be solved in polynomial time in the Silver Bullet Model.

Theorem 3. h -MAXKEI can be solved in polynomial time in the Silver Bullet Model.

Proof. If we have a model that excludes incompatibility, as the Silver Bullet Model, then a modification of the built graph solves this problem. The modification of the graph is as follows. The goal is to decompose each half-compatible edge into a set of paths, and then lead these paths through a gadget that will regulate the maximum number of used half-compatible edges through its size. First we add this gadget, which consists of $2h$ new vertices in sets A and B , and a set of h disjoint edges of weight 0 between them: $\{(a_1, b_1), (a_2, b_2), \dots, (a_h, b_h)\}$. Then we replace each edge (p_i, d_j) in the half-compatible class by a set of edges connecting p_i to each of a_1, a_2, \dots, a_h , and d_j to each of b_1, b_2, \dots, b_h . The weight on these edges are set to be half of the original weight of the replaced half-compatible edges. The rest of the graph remains unchanged. For an example, see Figure 3, which is a transformed version of our example instance from Figure 2, with $h = 2$.

Claim 1. A maximum weight $(P \cup A \cup B)$ -complete matching in the above-described graph corresponds to an allocation maximizing the weight subject to $HC \leq h$, if half-compatible donations are less desirable than compatible donations according to the weight function.

Proof. Due to the size of the gadget, no $(P \cup A \cup B)$ -complete matching allows more than h vertices in P to be matched along their edges to the gadget. Moreover, the number of vertices in P that are matched to a vertex in A equals the number of vertices in D that are matched to a vertex in B , because a $(P \cup A \cup B)$ -complete matching covers all vertices in the gadget. These vertices in P and D will be




Figure 3: Substituting all half-compatible edges by a gadget allowing at most $h = 2$ half-compatible donations.

the agents participating in half-compatible donations. Due to the assumptions of the Silver Bullet Model, any perfect matching on them is a set of executable transplants. The rest of the transplants are chosen based on the maximum weight matching criterion.

Notice that it is possible that a donor and a patient are connected via a 3-path through the gadget and via a direct compatible edge as well, but for all weight functions where half-compatible donations are less desirable than compatible donations (all our weight functions except for 1 and possibly 6), the path will carry the lower weight. Therefore, no $(P \cup A \cup B)$ -complete matching using such edges in the gadget can be of maximum weight. \square

One can calculate a maximum weight matching among all $(P \cup A \cup B)$ -complete matchings using the usual weight transformation of adding a large uniform weight to all edges incident to any vertex in $P \cup A \cup B$. This completes the proof. \square

This construction in the proof of Theorem 4 answers a question more general than h -MAXKEI. It actually decides

whether there is an allocation of weight at least t while using at most h suppressants.

Theorem 4. *In the silver bullet model, a minimum cost S-IR allocation can be computed in polynomial-time even if there is an upper bound on the number of suppressants that can be used.*

Next, we identify connections of h -MAXKEI with a special case of budgeted matching, a well-studied graph problem of unknown complexity.

UNIT-COST BUDGETED MATCHING

Input: Bipartite graph $G = (A \cup B, E)$, $E' \subseteq E$, edge weights, and an integer h .
 Question: Find a maximum weight matching M such that $|M \cap E'| \leq h$.

UNIT-COST BUDGETED MATCHING admits a PTAS (Berger et al. 2011; Mastrolilli and Stamoulis 2012). Berger et al. (2011) observe that for polynomial weights and costs (here we set the costs to be 1), BUDGETED MATCHING is very unlikely to be NP-hard, because it would imply $RP=NP$. However, after several decades, the problem of finding a deterministic algorithm to solve this problem is still open.

We now show how h -MAXKEI reduces to UNIT-COST BUDGETED MATCHING.

Lemma 1. *h -MAXKEI polynomial-time reduces to UNIT-COST BUDGETED MATCHING.*

Proof. We set E' to be the set of half-compatible edges, while G and the upper bound h are identical in the two problems. To make sure that the maximum weight matching in UNIT-COST BUDGETED MATCHING is P -complete, we modify the edge weights $w(e)$ from h -MAXKEI as follows.

$$w'(e) := w(e) + n$$

The weight of matching M is thus $w'(M) := w(M) + n \cdot |M|$. Since $w(e) \leq 1$, matchings of larger size are bound to have a larger weight as well, thus $w'(M)$ is maximized by a P -complete matching. Within the set of P -complete matchings, w' is maximized in the maximum weight matching according to w .

Notice that the weight transformation can be extended to other objective functions as well:

$$w'(e) := w(e) + n \cdot w_{\max}(e).$$

□

Regarding parametrized complexity, our trivial parameter is h , the number of suppressants available. If h is small, then one can try which h half-compatible edges are used, and then search for a maximum weight allocation in the rest of the instance built out of compatible and private edges only.

Axioms

We consider some natural axioms for allocation algorithms (mechanisms). A mechanism is *utilitarian (UTIL)* if it maximizes the total number of transplants (satisfied patients). A

mechanism is *Pareto optimal (PO)* if it returns a Pareto optimal allocation among the set of feasible outcomes. A mechanism is *strategyproof (SP)*, if no patient can misreport her compatibility relation to obtain a more preferred allocation. An extremely weak form of Pareto optimality is called *unanimity (UNAN)* that requires that a mechanism returns an allocation that gives each agent a fully compatible kidney whenever such an allocation exists. We discuss that our flexible objective-based algorithmic approach satisfies many of these properties.

Utilitarianism and Pareto optimality

Only objective functions 1, 2, and 5 reach UTIL, since they are the only ones that maximize the number of transplants in total.

PO (under coarse preferences) means that no inclusion-wise larger allocation exists. Only objective functions 1, 2, and 5 satisfy this.

PO (under refined preferences) means that no other allocation keeps the set of compatible donations, while extending the set of donations or upgrading a half-compatible donation to a compatible one. Only objective function 1 fails this: PO (under refined preferences) is defined for refined preferences, while objective function 1 is defined for coarse preferences.

UNAN is implied by both PO (under refined preferences) and PO (under coarse preferences), and thus all objective functions satisfy it.

Strategyproofness

Next, we consider strategic properties. We identify a fundamental incompatibility between strategyproofness and unanimity when agents have refined preferences and there is an upper bound on the number of suppressants.

Theorem 5. *Even for 3 patients, the combination SP (under-refined-preferences), UNAN, and IR are incompatible. Even for 3 patients, the combination SP (under-refined-preferences), UNAN, and S-IR are incompatible.*

Proof. Consider the problem instance depicted in Figure 4. Suppose there is exactly 1 suppressant available. Then PO requires that all three patients are matched: 1 gets d_2 , 2 gets d_3 and 3 gets d_1 . Suppose patient 2 misreports and reports d_3 as incompatible. In that case, IR and PO imply that patient 2 gets d_1 and patient p_1 gets d_2 , which is an improvement for p_2 with respect to refined preferences. □

Since UTIL implies PO (under coarse preferences), we also derive the following statements.

Corollary 1. *The combination SP (under-refined-preferences), PO (under refined preferences), and IR (or S-IR) are incompatible. The combination SP (under-refined-preferences), PO (under coarse preferences), IR (or S-IR) are incompatible.*

If there is no upper bound on the number of suppressants available, then PO, IR, and SP are compatible under refined or coarse preferences. This can be proved by reducing our problem to that housing markets with trichotomous




Figure 4: A KEI instance for the proof of Theorem 5. The dashed lines indicate half-compatible edges. Solid edges indicate compatibility edges. Dotted edges indicate a patient-donor pair.

preferences. Even for general weak preferences, suitable extensions of TTC achieve all the axioms (see, e.g. Saban and Sethuraman (2013)). The agents have trichotomous preferences: fully compatible kidneys are most preferred, half-compatible kidneys are second most preferred, and incompatible kidneys are least preferred.

Theorem 6. *If there is no upper bound on the number of suppressants available, then there exists an algorithm that is PO, IR, and SP.*

Our next insight is that as long as we consider strategyproofness with respect to coarse preferences, we can achieve all the objectives in conjunction with strategyproofness. The key is to overlay the priority mechanism over the algorithm that optimizes one of the global objectives while constrained to S-IR. The result follows from similar insights as that of Aziz (2019) and Hatfield (2005). The desirability of the overlay of the priority mechanism is that it can take as a black box any given algorithm that optimizes some criterion over the set of feasible allocations. In the overall algorithm, first, an optimal value W of a feasible matching is computed. Then, we go over an order of agents and check whether the agent in question can be satisfied in some optimal allocation which still satisfying all the preceding agents who could also be satisfied in some optimal allocation. It can be proved that an agent has no incentive to misreport her preferences in the Coarse preference model.

Theorem 7. *For all the objectives in the paper, there exists an algorithm that is SP (under coarse preferences) and meets the objective subject to S-IR and upper bound on the number of suppressants.*

Pairwise exchanges

In this section, we restrict our attention to pairwise exchanges. In many countries, the organ exchange programs only consider pairwise exchanges (Andersson and Kratz 2019). In this case, a simpler model suffices. The only possible transplants are cross-donations between two patient-donor pairs, and an altruistic donor donating a kidney to a patient, who might come with or without a donor.

We build a simple graph for these sorts of donations. First, we represent each patient-donor pair, each altruistic donor,

Algorithm 1 Priority Overlay

- 1 Construct graph G with appropriate edge weights.
- 2 Compute a maximum weight A -complete matching of G . Let the weight of the matching be W .
- 3 Take a permutation (ordering) π of the elements in N . Denote by $\pi(i)$ the i -th element in the ordering.
- 4 **for** $i=1$ to n **do**
- 5 **while** target=false **do**
- 6 Remove from G each of $p_{\pi(i)}$'s edges except the most preferred. Compute the maximum weight A -complete matching of G . If the weight is less than W or if there does not exist a A -complete matching, then put back the removed edges but remove most preferred edges in the graph. Otherwise set target to true.
- 7 **end while**
- 8 **end for**
- 9 Compute a maximum weight complete matching M of G . Consider the allocation x in which each patient $p_i \in P$ gets a donor that it is matched to in M . If $p_i \in P$ is matched to a dummy donor, its allocation is null.
- 10 **return** $(x(1), \dots, x(n))$

and each patient without a donor as a vertex. Then we connect two of these vertices if a cross-donation or an altruistic donation, as described above, is possible, because all donors are either compatible or half-compatible to the prospective recipient of the kidney. The edge carries the weight of the two transplants altogether, where transplant weights are calculated based on the relevant row of Table 1. We look for a not necessarily perfect maximum weight matching.

Conclusions

Immunosuppressants constitute one of the major developments in the transplant of organs. We have studied a general model of kidney exchange that also models the use of immunosuppressants. We envisage further work on both algorithmic as well as market design aspects of the problem. The most striking open question is undoubtedly to determine the complexity of MAXKEI in the general model. By solving that problem, one might even reach progress in budgeted matching, a long-standing open problem in matching theory. Finally, we note that some of our insights may be usable for other organ markets that involve lungs (Luo and Tang 2015), livers, or even a combination of these and other organs (Dickerson and Sandholm 2017).

References

- Abraham, D.; Blum, A.; and Sandholm, T. 2007. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM Conference on Electronic Commerce (ACM-EC)*, 295–304. ACM Press.
- Abramowicz, D.; Oberbauer, R.; Heemann, U.; Viklicky, O.; Peruzzi, L.; Mariat, C.; Crespo, M.; Budde, K.; and Oniscu, G. C. 2018. Recent advances in kidney transplantation: a

- viewpoint from the Descartes advisory board. *Nephrology Dialysis Transplantation* 33(10):1699–1707.
- Andersson, T., and Kratz, J. 2019. Pairwise kidney exchange over the blood group barrier. *The Review of Economic Studies*.
- Aziz, H. 2019. Strategyproof multi-item exchange under single-minded dichotomous preferences. *arXiv preprint arXiv:1905.10778*.
- Berger, A.; Bonifaci, V.; Grandoni, F.; and Schäfer, G. 2011. Budgeted matching and budgeted matroid intersection via the gasoline puzzle. *Mathematical Programming* 128(1-2):355–372.
- Biró, P.; Haase-Kromwijk, B.; Andersson, T.; Ásgeirsson, E. I.; Baltessová, T.; Boletis, I.; Bolotinha, C.; Bond, G.; Böhmig, G.; Burnapp, L.; et al. 2019. Building kidney exchange programmes in Europe: an overview of exchange practice and activities. *Transplantation* 103(7):1514–1522.
- Biró, P.; Manlove, D. F.; and Rizzi, R. 2009. Maximum weight cycle packing in optimal kidney exchange programs. *Discrete Mathematics, Algorithms and Applications* 1(4):499–517.
- Chun, Y.; Heo, E. J.; and Hong, S. 2017. Kidney exchange with immunosuppressants. Technical report, Vanderbilt University Department of Economics.
- Dickerson, J. P., and Sandholm, T. 2017. Multi-organ exchange. *Journal of Artificial Intelligence Research* 60:639–679.
- Hatfield, J. W. 2005. Pairwise kidney exchange: Comment. *Journal of Economic Theory* 125:189–193.
- Luo, S., and Tang, P. 2015. Mechanism design and implementation for lung exchange. In *IJCAI*, 209–215.
- Mastrolilli, M., and Stamoulis, G. 2012. Constrained matching problems in bipartite graphs. In *International Symposium on Combinatorial Optimization*, 344–355. Springer.
- Montgomery, R. A.; Lonze, B. E.; King, K. E.; Kraus, E. S.; Kucirka, L. M.; Locke, J. E.; Warren, D. S.; Simpkins, C. E.; Dagher, N. N.; Singer, A. L.; et al. 2011. Desensitization in hla-incompatible kidney recipients and survival. *New England Journal of Medicine* 365(4):318–326.
- Munkres, J. 1957. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics* 5(1):32–38.
- Roth, A. E.; Sönmez, T.; and Ünver, M. U. 2005a. A kidney exchange clearinghouse in New England. *American Economic Review* 95(2):376–380.
- Roth, A. E.; Sönmez, T.; and Ünver, M. U. 2005b. Pairwise kidney exchange. *Journal of Economic Theory* 125(2):151–188.
- Saban, D., and Sethuraman, J. 2013. House allocation with indifference: a generalization and a unified view. In *Proceedings of the 14th ACM Conference on Electronic Commerce (ACM-EC)*, 803–820. ACM Press.
- Sönmez, T., and Ünver, M. U. 2011. Matching, allocation, and exchange of discrete resources. In Benhabib, J.; Jackson, M. O.; and Bisin, A., eds., *Handbook of Social Economics*, volume 1. Elsevier. chapter 17, 781–852.