

Kidney Exchange: an Operations Perspective

Itai Ashlagi and Alvin E. Roth

November 2020

Abstract

Many patients in need of a kidney transplant have a willing but incompatible (or poorly matched) living donor. Kidney exchange programs arrange exchanges among such patient-donor pairs, in cycles and chains of exchange, so each patient receives a compatible kidney. Kidney exchange has become a standard form of transplantation in the United States and a few other countries, in large part because of continued attention to the operational details that arose as obstacles were overcome and new obstacles became relevant. We review some of the key operational issues in the design of successful kidney exchange programs. Kidney exchange has yet to reach its full potential, and the paper further describes some open questions that we hope will continue to attract attention from researchers interested in the operational aspects of dynamic exchange.

1 Introduction

Kidney failure is a leading cause of death around the world. The best treatment is transplantation, but no country is presently able to supply all the transplants required by its patient population. In the U.S. and many other countries, most transplants today come from deceased donors. While efforts are underway to increase the availability of deceased donor kidneys, this is a naturally limited source of transplants, because only a tiny fraction of deaths allow kidneys to be recovered for transplantation.¹ But healthy people have two kidneys, and can remain healthy with one, and so another source of kidneys for transplantation is from healthy living donors, who can give a kidney to save someone with kidney failure. This also has some natural barriers, because kidneys have to be well matched to the patient's immune system, and so not everyone who is healthy enough to donate a kidney can donate one to whom they wish.

It is also against the law almost everywhere in the world to pay a living donor to donate a kidney. (The single exception is the Islamic Republic of Iran, where there is a legal monetary

¹Loosely speaking, a potential donor must die in a hospital, on a ventilator, so that his/her organs continue to receive oxygen, and of course the cause of death and general health of the deceased person must be consistent with having healthy kidneys at the time of death. One focus of efforts in the U.S. to increase transplantation of organs from deceased donors is to make more effective use of organs from older donors.

market for kidneys (Akbarpour et al., 2020c), although there are also black markets around the world.)² Kidney exchange (KE, also called Kidney Paired donation, KPD) arose as a way of increasing the availability of transplants from compatible living donors without violating the ban on compensating donors.³

Suppose someone with kidney failure has a healthy potential donor who loves them and would like to give them a kidney, but can't, because the donor's kidney is incompatible with the patient. Two or more such incompatible patient-donor pairs might be able to exchange kidneys, so that each patient gets a kidney that is compatible with him/her, from another patient's donor. A handful of early exchanges were identified by inspection, and conducted within individual transplant centers, and the question arose how to coordinate these, on a large scale, in ways that would make it feasible to do exchanges among many patients and donors, who might often be at different hospitals.

This paper recounts that effort, which began very slowly, and required constant adaptation of the market design, involving issues that engaged economists, computer scientists, and operations researchers in support of surgeons, physicians, and transplant professionals of all sorts. Today kidney exchange has become a standard form of transplantation in the United States and a few other countries, in part because of continued attention to the operational details that arose as previous obstacles were overcome and new ones became relevant. But much more remains to be done, because the full potential of kidney exchange has yet to be reached, and there are still many more patients in need of transplants than can presently be saved.

2 Background

At the beginning of 2020, there are about 100,000 patients with end-stage renal disease (ESRD) on the deceased donor waitlist in the United States. While patients with ESRD may be kept alive by dialysis, a transplant leads to better life quality and longer life expectancy (Wolfe et al., 1999). Each such transplant also saves hundreds of thousands of dollars in

²There is a considerable literature on the desirability or undesirability of allowing donors to be compensated, but we will not discuss that here (but see <https://marketdesigner.blogspot.com/search/label/compensation%20for%20donors> for more on that and related issues).

³The idea of kidney exchange seems to have first been proposed by Rapaport (1986), but the first clinical kidney exchanges did not occur until the 1990's: See Wallis et al. (2011) for a concise history. Other surveys include Sönmez and Ünver (2013a), Glorie et al. (2014), Anderson et al. (2015a), Biró et al. (2019c). Kidney exchange started to develop at a time when economists and market designers had begun to pay increased attention to matching markets, which are markets in which participants care to whom they are matched, and in which prices don't do all the work of deciding who gets what. (Other matching markets that were the object of study and design include the market for new doctors, (cf. Roth (1984), Roth (1991), Roth and Peranson (1999)), and the design of school choice systems, (cf. Abdulkadiroglu and Sönmez (2003) and Abdulkadiroglu et al. (2005a), Abdulkadiroglu et al. (2005b), Abdulkadiroglu et al. (2009)). The designs of both medical labor market clearinghouses and school choice were built upon adaptations of the deferred acceptance algorithm of Gale and Shapley (1962). However the algorithms that power contemporary kidney exchange are very different, and are today mostly formulated as mixed-integer programs.

medical costs compared to dialysis.⁴

In the coming year (judging from the experience of 2019, and hoping that the Corona pandemic does not too long depress transplantation rates) about 16,000 U.S. patients will be transplanted using a cadaver organ. About 6,500 kidneys will be donated from live donors in the U.S.⁵ But also about 8,000 patients on the waitlist will either die or become too sick to be transplanted.

Well over 1,000 of the U.S, living donor kidney transplants in 2019 resulted from kidney exchange, some through exchanges carried out within a single transplant center, and many through kidney exchange platforms which organize these exchanges among multiple hospitals.⁶

In each case, organization of exchanges on a large scale involves the creation of a database of patient-donor pairs, and the use of software to determine which donors are compatible with which patients, and then to find optimum collections of exchanges, according to some well defined criteria.

Early ideas about how to organize such exchanges (Roth et al., 2004) grew out of the work on “top trading cycles” by Shapley and Scarf (1974), and followup work by Roth (1982) and Abdulkadiroğlu and Sönmez (1999). The cycles produced by top trading cycle algorithms could potentially involve exchanges among a cycle consisting of many patient-donor pairs, and this was beyond the operational capability of most transplant centers in 2004. So the first inter-hospital exchange in the U.S., the New England Program for Kidney Exchange (Roth et al., 2005b) was organized around an algorithm that considered only exchanges between two patient-donor pairs (Roth et al., 2005a).

Next we describe some basic forms of exchange and some medical facts that determine compatibility.

2.1 Types of exchanges: cycles and chains

Exchanges of kidneys typically take one of two forms. A **cycle** involves a set of incompatible patient-donor pairs. The patient of each pair receives a kidney from the donor of another pair. To avoid situations in which a patient-donor pair donates a kidney but fails to receive a kidney, cycles are almost always arranged in a simultaneous manner, i.e. all transplants are done at the same time. This often limits cycles to small exchanges involving just 2 or 3 pairs, because a simultaneous exchange among n pairs requires the simultaneous availability of $2n$ operating rooms and surgical teams, to handle all the simultaneous nephrectomies (kidney removals) and transplants.

⁴In 2014, Medicare paid about \$90,000 per year per dialysis patient, but only about \$30,000 per year per transplant patient. With a median waiting time on dialysis of 3.5 years before a transplant, a transplant saves about \$200,000 (United States Renal Data System, 2016). A much more detailed cost analysis is given by Held et al. (2016).

⁵See <https://optn.transplant.hrsa.gov/data/view-data-reports/national-data/>.

⁶In the U.S. there are three active inter-hospital kidney exchanges, the National Kidney Registry (NKR), the Alliance for Paired Kidney Donation (APKD), and the UNOS pilot program. There are active kidney exchanges in Canada, the U.K.(see Manlove and O’malley (2015)), the Netherlands (Keizer et al., 2005), and smaller kidney exchanges in Australia (Cantwell et al., 2015) and several European countries.

For this reason, the first exchanges were organized only between two pairs at a time (Roth et al., 2005a,b). This allows the use of some elegant graph theory (and Edmonds (1965) algorithm), but it limits the number of transplants that can be arranged to those that involve a "double coincidence of wants," i.e. those that involve two patient-donor pairs whose patients are each compatible with the other patient's donor.⁷ It didn't take long before the logistics of doing three-way as well as two-way exchanges was mastered (Saidman et al., 2006), but real progress depended on enlarging exchange beyond cycles.

The other form of exchange is a **chain**, which is initiated by a **non-directed donor** (NDD) who has no particular intended recipient.⁸ The NDD initiates a chain by donating to the patient of the first pair in the chain, whose donor donates to the patient of the second pair and so forth. Chains usually end with donation to a patient on the deceased donor waitlist who has no affiliated living donor who could continue the chain. An advantage of chains is that they can be arranged sequentially, and non-simultaneously with each pair receiving a kidney before they donate one (Roth et al., 2006; Rees et al., 2009). This means that, if a link is broken, no pair is left without a kidney to offer in a future exchange, since a broken link that leaves a pair without a scheduled transplant comes when they have not yet donated their donor's kidney.⁹ Consequently, not all the operating rooms and surgical teams have to be available simultaneously (since a potential broken link is much less costly than in a cycle), so chains most often yield more (sometimes many more) transplants than a cycle.

The adoption of non-simultaneous nondirected donor chains was one of the biggest operational changes that separates the current situation, in which kidney exchange has become a standard form of transplantation in the U.S., from its earliest days. The New England Program for Kidney Exchange (NEPKE) operated from 2004 until 2011, when it was merged into a pilot program run by UNOS.¹⁰ In its seven years of operation it was responsible for a number of innovations (see e.g. Saidman et al. (2006)). But it did not employ nonsimultaneous chains, and employed only simultaneous cycles and chains, each resulting in only two or three transplants. Partly as a consequence, it produced only 83 transplants in its seven years of operation.¹¹ However with the publication of Rees et al. (2009) reporting the very

⁷Because this simplifies the mathematics, there are theoretical papers today that continue to model kidney exchange as involving only pairwise exchanges (and insights from these papers may continue to be useful), but the practice of kidney exchange has moved on.

⁸Although deceased donors are almost always non-directed, chains are so far initiated by living nondirected donors, although that may change in the future.

⁹As non-simultaneous chains have become a standard form of kidney exchange transplants, it turns out that broken links due to a donor backing out after his/her intended recipient has received a transplant are rare (Cowan et al., 2017). (This is partly due to operational practices in which not every pair is nominated to be a non-simultaneous "bridge" pair, and partly it is a behavioral economics observation: we humans are substantially nicer than some economic models give us credit for.)

¹⁰UNOS, the United Network for Organ Sharing, is the federal contractor that controls the allocation of organs from deceased donors in the U.S.

¹¹see <https://web.archive.org/web/20110727115230/http://www.nepke.org/>, <https://web.archive.org/web/20110727120120/http://www.nepke.org/theprogram.htm>, and <https://web.archive.org/web/20110727115642/http://www.nepke.org/livingdonors.htm>

first non-simultaneous chain, and subsequently Ashlagi et al. (2011a,b) establishing their ability to substantially increase the rate of exchange transplants, non-simultaneous chains were adopted by all the major U.S. inter-hospital kidney exchanges.¹²

2.2 Medical compatibility, crossmatches and pair types

ABO compatibility. For a donor to be compatible with a patient, there are two constraints. First, the patient must be blood-type compatible (ABO-compatible). This means the patient cannot receive a kidney from a donor who has a blood antigen (A or B) that the patient does not have (see Figure 1). So an O donor (who has neither the A nor the B antigen) is ABO-compatible with any patient whereas an O patient can only receive a kidney from an O donor, since other blood-type donors have A, B or both.




Figure 1: ABO-compatibility structure. A directed arc from X to Y means that a donor with blood-type X is compatible with a recipient with blood-type Y.

Recent desensitization technology sometimes makes it possible to transplant a patient who is ABO incompatible (ABOi); for example in some cases in which the patient’s relevant blood antigen antibody concentrations (titers) are sufficiently low.¹³

¹²Perhaps the earliest suggestion that long non-simultaneous chains could be started by NDDs was in Roth et al. (2006), but it took another half a dozen years, operational innovations, and the gathering of evidence to overcome substantial opposition before they became the dominant form of kidney exchange in the U.S. (cf. Anderson et al. (2015a).) In many other countries, chains, and non-simultaneous chains, are only now being slowly adopted. Other kinds of non-simultaneous exchanges have been conducted and considered, see e.g. Butt et al. (2009); Ausubel and Morrill (2014); Kute et al. (2016); Akbarpour et al. (2020b,a).

¹³For example, blood type A actually comes in two types, A1 and A2, and it is sometimes possible to transplant an A2 organ into an O or B patient, an A1 organ into a B patient, a B organ into an A patient, and an A2B organ into an A or B patient. To get an idea of the issues involved, note that A2 blood types have less A1 antigen on their cell surface so it is possible to transplant A2 organs into a B or O recipient with a low anti-A titer. The titer measures the amount or concentration of antibody. For patients receiving A2 organs we determine the A1 titer using a method (DTT) that detects only the IgG (Immunoglobulin G) portion of the antibody. If the DTT anti-A titer is low (typically 4 or less) this will be considered a “compatible” transplant. For true ABO incompatible transplants we use two methods – again the DTT anti-A or B titer (depending on the incompatibility) and a method that uses an anti-human globulin molecule that measures both IgG and IgM (Immunoglobulin M) portions of the antibody. The titer considered to be acceptable for desensitization

Tissue type compatibility. In addition to blood-type antigens, the donor has human leukocyte antigens (HLA). Each donor and patient inherit antigens from their parents, and for each location on the relevant chromosome a patient has 1 or 2 antigens.¹⁴ For a kidney to be compatible, the patient cannot have antibodies to the donor’s HLA, because otherwise the patient’s immune system will immediately try to reject the organ. A patient is “tissue-type compatible” to the donor if she has no antibodies to the donor’s HLA. This type of compatibility is verified using a *crossmatch* test.

Patients need to take immunosuppressive medications after transplantation to prevent their immune system from attacking any foreign antigens, and it is desirable to do as little suppression of the patient’s immune system as possible. For this reason it is also preferred to transplant an organ from a donor who has some of the same HLA antigens as the patient. Due to the development of desensitization technologies it is sometimes possible to transplant an organ to which the patient has an antibody. However, typically this results in worse outcomes than receiving a compatible kidney. One of the virtues of kidney exchange is that it allows more patients to receive transplants of compatible kidneys. However when no compatible kidney is available, “desensitizing” the patient to allow an otherwise incompatible kidney to be transplanted generally has a better outcome than remaining on dialysis. Consequently, extremely highly sensitized patients sometimes are transplanted through kidney exchange with a kidney to which they have sufficiently few antibodies to allow desensitization.¹⁵

A common measure for how difficult it will be for a patient to find a compatible donor (among those who are blood type compatible) is the Panel Reactive Antibody (PRA), which captures the likelihood, that based on her antibodies the patient is tissue-type incompatible with a random donor in the population despite being ABO compatible.¹⁶ A patient with high PRA is referred to as “highly sensitized.”¹⁷

More about HLAs, antibodies and differences across KE programs and hospitals.

In the early days of kidney exchange, databases included antigens (alleles) for three locations (loci) on chromosome 6 called A,B, and DR. As crossmatch testing became more accurate and it was clear that more antigens play a role, KE programs (and hospitals) gradually added antigens for loci C, DQ and DP. This turns out to help predict both initial compatibility, and the quality of the match, as some HLA mismatches are more important than others

varies by transplant center but typically it is 64 or 128. Guidelines for ABOi transplants through UNOS can be found here: https://optn.transplant.hrsa.gov/media/2223/mac_pcproposal_201707.pdf. (See also Sönmez et al. (2018) and Salvadori and Tsalouchos (2020).)

¹⁴One if the same antigen is inherited from both parents.

¹⁵Desensitization works better for ABO incompatibility than for tissue type incompatibility.

¹⁶An estimate of the PRA is calculated using the frequency of antigens. The method and antigen frequencies can be found here <https://unos.org/news/optn-public-comment-proposal-offers-alternate-cpra-calculation-method/>.

¹⁷Fifteen years ago it was common to view a PRA above 80 as highly sensitized. But progress in matching patients in kidney exchange pools has changed this view, and today most KE programs consider only PRA above 95 and even 98 to be high.¹⁸ We note that the deceased donor allocation system in the US shifted in 2014 from assigning all patients with PRA above 80 the same priority points, to assigning priority points that strictly increase with PRA.

(Manski et al., 2019).

Technologies and habits create differences between KE programs and hospitals with respect to HLA typing. We describe a few notable examples. Not all KE programs record all 6 loci. Another difference is in the resolution typing of alleles. As an example consider allele A*02:01 which has locus A, allele group 02 (serotype) and all 4 digits 02:01 are the specific allele.¹⁹ Some KE programs or hospitals record only the serotype A*02 and others will record the specific allele 02:01.²⁰ Some differences are due to continuous discoveries of alleles and their structure. HLA DQ for instance has two chains, α and β , that are adjacent on the same chromosome. Some KE programs record these as two different antigens DQA and DQB while others will view them as one “combined” antigen.

Hospitals differ also with respect to antibody typing. First, some record only serotypes and some record the entire alleles. Another notable difference is how hospitals determine what is an antibody (for the purpose of a virtual crossmatch). An antibody has a strength and hospitals set thresholds for what is considered an antibody. A common measure for the strength of an antibody is the mean fluorescence intensity (MFI). Hospitals use different thresholds to determine antibodies, as there are currently no accurate prediction tools for crossmatch tests.²¹ Figure 2 presents a partial list of antibodies of a patient in a database that also lists weak antibodies. In this example B7, B27, B48, B55 have MFI much lower than 4000 and are very unlikely to act as antibodies and reject a kidney. A29, A43, and B8 all have MFI above 4000 and are likely to cause a rejection of an organ that has at least one of these HLAs. Often 2-digit typing is sufficient. But sometimes the patient and the donor have different alleles of the same group which can result in an apparent but false match between a donor’s antigens and a patient’s antibodies. The patient in this example is unlikely to reject an organ that does not have A*29:01 or A*29:02, even if it has A*29:03.²²

The differences described above have a direct impact on frictions in the matching process, and can slow down merging and collaboration attempts between hospitals and KE programs. In the early days of kidney exchange, when hospitals all conducted their own blood tests and the technology for testing and the language for describing the immunological data was not so detailed, hospitals sometimes faced incentive problems about listing an antibody that

¹⁹For more about nomenclature see <http://hla.alleles.org/nomenclature/naming.html>.

²⁰There are also some equivalences between different groups. HLA equivalence tables can be found here: <https://optn.transplant.hrsa.gov/governance/public-comment/update-hla-equivalency-tables/>. These table keep updating.

²¹An (ad hoc) threshold of 4000 MFI is very common for most types of antibodies. KE databases used by Methodist Specialty and Transplant Hospital at San Antonio (MSTH), Israel and several other KE programs (using software developed by Ross Anderson, Itai Ashlagi, Sukolsak Sakshuwong and Jon Silberholz) allow thresholds for virtual crossmatching to be set flexibly; so all antibodies can be listed together with their MFI levels, whether weak or strong.

²²We believe that there is a lot to gain from international exchanges due to population heterogeneity at the specific allele level. This is because patients develop antibodies to the antigens to which they become exposed (e.g. through blood transfusions, childbirth, previous transplants, etc.). But different populations have different distributions of HLA’s, and so a highly sensitized patient who is very unlikely to find a match at her home location because of her many antibodies, may be more likely to find a match from a population of donors some of whose HLA may be relatively rare in the patient’s home location, because she may not have developed antibodies to those.

Antibodies

Serotype	Alleles	MFI
A29	29:02	9711
A29	29:01	14155
A43	43:01	17171
B7	07:02	575
B8	08:01	4126
B27	27:08	2370
B42	42:01	512
B48	48:01	873
B55	55:01	435

Figure 2: Example for a list of antibodies of a patient together with their MFI level.

(they thought) had a low concentration, since listing the antibody would mean that they would not be offered kidneys that appeared incompatible, but might actually be compatible. On the other hand, suppressing the report of an antibody entirely led to the offer of many kidneys that were ultimately incompatible, so that the offers had to be rejected, which delayed transplants.²³

Powerful donors. Just as some patients are hard to match, some donors are powerful in the sense that they have an increased chance of being compatible with some highly sensitized patients. These are donors who have rare HLAs, or, better yet, donors who are homozygous in rare HLAs, which is to say they have fewer than the usual number of distinct HLAs because they have inherited some of the same HLAs from both parents. Thus a high PRA patient, who has many antibodies, may sometimes not have antibodies to the particular HLAs of a powerful donor.²⁴

Crossmatches. Given the data on patient and donor ABOs, the donor HLA, and the patient antibodies one can determine *virtual* matches. This is sometimes called a *virtual crossmatch*. To verify whether the patient will not reject the donor’s kidney, a physical *crossmatch* test is required prior to the transplant, involving blood samples from the prospective donor and patient. A positive (negative) crossmatch between a patient and a donor means

²³Notice how the choices available to participating transplant centers change with changes in the market design resulting from changes in the available testing and reporting technologies, and also with innovations like having a centralized blood testing facility (to which hospitals send blood samples, rather than reports about those samples).

²⁴The frequencies can be found here: <https://unos.org/news/optn-public-comment-proposal-offers-alternate-cpra-calculation-method/>.

that the patient will (not) reject the donor’s kidney.²⁵

Pair types. The compatibility structure helps to classify pairs based on how difficult they are to match. It will be convenient to refer to the blood types of the patient-donor pairs, so that e.g. an A-B pair is one in which the patient has blood type A and the donor has blood type B. An exchange pool is likely to have fewer A-O pairs than O-A pairs because A-O pairs are often compatible (since they are ABO-compatible), and so can choose a direct live-donor transplant from the donor to the patient rather than exchange with another pair. When there are few highly sensitized (high PRA) patients, blood type compatibility is of first order importance. So pairs which are ABO compatible (even if they are tissue type incompatible) are *over-demanded*. That is (A-O, AB-O, B-O, AB-B, AB-A) pairs are offering to exchange more widely acceptable kidneys than they are seeking. In contrast, ABO incompatible pairs (O-A, O-B, O-AB, A-AB, B-AB) are *under-demanded*: they are seeking a more highly demanded kidney than they are offering, and so it will be harder to match all such pairs at any given time.

More generally, we will refer to some pairs as *hard-to-match* if they are either *under-demanded* or the patient is highly sensitized. Other pairs can be thought of as *easy-to-match*. In general, hard-to-match pairs may accumulate in the pool, and easy-to-match pairs will be quickly matched. Consequently the most efficient use of easy to match pairs is to match them to hard-to-match pairs: it is seldom efficient to match two easy-to-match pairs to each other, given the abundance of hard-to-match pairs who may remain unmatched.

2.3 Common operations in kidney exchange platforms

KE programs that serve multiple hospitals share similar operations and dynamics, with exact details varying across platforms.²⁶

- *Submissions.* Hospitals submit medical data to the KE program about their patient-donor pairs and non-directed donors as they become available or, in some hospitals, after attempts to arrange an internal in-hospital kidney exchange for those pairs has failed. This latter practice selects easy-to-match pairs out of the exchange pool, by matching them to each other. This makes it even harder to match hard-to-match pairs. So attention has to be paid to incentives for hospitals to participate fully (Ashlagi and Roth, 2014; Agarwal et al., 2019).²⁷
- *Matching and match offers.* KE programs periodically identify a set of exchanges within the patient-donor network. This is usually done using a weighted integer optimization program with weights assigned to each potential transplant. This includes deciding how and when to end a chain. For each potential transplant in an exchange, the associated hospital is informed and is given details about the potential donor.

²⁵This terminology can be confusing: a positive crossmatch is bad news.

²⁶While there are substantial similarities internationally, we give here a possibly U.S.-centric view.

²⁷Some platforms only enroll pairs that have already passed all screening procedures, while others enroll pairs at earlier stages.

- *Offer reviews and transplants.* Hospitals confirm whether the offers (matched donors) are acceptable. If all offers in a given exchange are acceptable, crossmatches between the relevant blood samples are done. If these crossmatches are negative, the next step is transplantation (organs are typically shipped to the patient’s hospital). But if one of these stages fails, the donors and patients in the exchange remain in the exchange pool.²⁸

2.4 Some differences across platforms

While in general the logistics are similar it is worth pointing out some differences. See also Biró et al. (2019c) for a detailed description of logistics of KE programs in European countries.

- *Upper bounds of cycle and chain lengths.* Multi-hospital KE programs usually bound the cycle length to 3.²⁹ Chains are usually limited to size 3 or 4.³⁰ Exceptions include the APKD, NKR, Israel, and the Czech Republic, which impose no limit on the chain length. Numerous countries do not allow non-directed donation and therefore don’t have chains at all.³¹
- *Prioritization/weight in matching.* KE programs use different prioritization when identifying matches. In fact almost every KE program uses different weights for various criteria. These include weights for HLA mismatches between the patient and the donor, ABO mismatch, age of the patients, difference between the age of the donor and recipient, distance, and so forth. Many countries, such as the UK, use an hierarchical set of objectives (see (Biró et al., 2019a)).
- *Matching frequency.* In the US KE programs identify matches almost on a daily basis. In European countries, Australia, and Canada, matching frequencies vary with many KE programs matching periodically every 1-4 months. We elaborate on this issue in Section 3.4.
- *Positive crossmatches and refusals of match offers.* KE programs suffer from high frictions due to high rates of offer refusals and positive crossmatches. In multi-hospital US KE programs (APKD, NKR, UNOS) refusals of offers are very common and KE programs attempt to elicit more refined preferences and request physicians’ to pre-select

²⁸When one link of a cycle fails, none of the transplants in the cycle can be conducted. But when the exchange is a chain, transplants can be performed up until the first positive crossmatch, and other links can be sought to “repair” and extend the chain. (This is one of the reasons that chains are such a productive source of kidney exchange transplants.)

²⁹This includes the Alliance for Paired Kidney Donation (APKD), UNOS, and the National Kidney Registry (NKR), and national KE programs in Netherlands, Spain, Portugal, UK Australia, and France.

³⁰Programs differ by whether the chain can continue into the next match run or not. For example, in the UNOS program a segment of a chain is limited to 4 but can proceed to the next match run, whereas the UK KE program terminated chains after a single segment (Manlove and O’malley, 2015).

³¹Such countries include France, Belgium, Austria, Sweden, Switzerland, Poland and Greece.

acceptable matches (Fumo et al., 2015). To overcome the delays from positive cross-matches, rather than shipping blood samples between centers, some KE programs use a centralized blood lab that carries out crossmatches in-house.³² This has implications over time differences between match offers. We return to these frictions in Section 3.5.

3 Operational issues

This section briefly describes operational challenges facing kidney exchange platforms. In the Appendix, we provide some simulation models for further research in this area.

3.1 Eliciting preferences

One early problem facing all of the multi-hospital KE programs was that transplant centers were consistently rejecting potential donors for reasons that could in principle have been stated in advance. But it was difficult for surgeons to pre-emptively accept or reject donors for each potential patient, particularly from a large set of donors, from most of whom any particular patient would never be offered a match. The reason is that each donor is a big set of attributes that take some time to be evaluated. One way this problem was addressed was by introducing a threshold language, so that, for each patient at a given time, a surgeon could specify the maximum age, body mass index, blood pressure, etc. for which a donor would be considered for that patient (Fumo et al., 2015). Thresholds help, but are far from perfect: i.e. a proposed transplant that meets all the specified thresholds can still be rejected.³³ Increasing the accuracy of the virtual matches, including by enhancing the language in which they can be described while preserving enough simplicity so that they will be effectively used by busy doctors, is an ongoing part of the market design.³⁴

3.2 Matching in thick pools

Early research in kidney exchange considered static pools that were assumed to be sufficiently large so that tissue-type incompatibilities were of lesser importance. Allocations that maximize the number of transplants have a simple and intuitive structure in such (idealized) pools (Roth et al., 2007; Ashlagi and Roth, 2014), which essentially carries over to sufficiently thick dynamic pools (Ünver, 2010). The main idea behind the structure, which is briefly described below, provides some initial guidance for measuring efficiency of a platform (Agarwal et al., 2019).

In sufficiently thick pools, efficient allocations are effectively determined by blood types, and there is essentially no need for cycles with more than 3 pairs. For some intuition consider

³²Examples include the UK, Austria, Belgium, the Czech Republic, and the APKD.

³³E.g. a surgeon might specify that he would consider donors up to age 60, moderately overweight, and with slightly elevated blood pressure, but mean in practice that a 60 year old would have to be otherwise ideal, and someone with elevated blood pressure would have to be under 50.

³⁴Issues related to the desirability of using simple languages comes up in many areas of market design, see e.g. Milgrom (2010).

the sets of A-O and O-A patient-donor pairs. If the only types of pairs were A-O and O-A, which are over- and under-demanded, it would be possible to match every A-O pair with a different O-A pair in 2-way cycles. Such an allocation maximizes the number of transplants (Figure 3). Allocations that would match some A-O pairs with each other would generate fewer transplants. This intuition carries over for other types of pairs (see Roth et al. (2007) and Ashlagi and Roth (2014) for the structure of efficient matching in large static pools).




Figure 3: A-O and O-A patient-donor pairs. In large pools with A-O and O-A pairs, all A-O pairs (in which the patient is not too highly sensitized) can match in 2-way exchanges with O-A pairs.

However “sufficiently thick” is an idealization not yet met in practice anywhere in the world, so in practice tissue-type incompatibilities play a very large role. Even relatively large pools do not match all their (overdemanded) O donors to their (underdemanded) O patients. For example, at the National Kidney Registry (NKR), which the largest platform, for about 7% of the transplanted O organs, the recipient is a non-O patient (Agarwal et al., 2019). Despite the scarcity of O donors, this may be efficient when the patient who receives the blood type O kidney is so highly sensitized that there may be only one kidney in the pool which is compatible. To put it another way, because there is a large population of highly sensitized patients, there are many hard to match pairs, including those with blood type O donors. Hard-to-match pairs in which the patient is very highly sensitized are not in fact “overdemanded,” regardless of their blood types.³⁵

One strategy for matching hard-to-match pairs is to increase the number of easy-to-match pairs included in the exchange pool. One class of easy-to-match pairs that are often excluded are compatible pairs, i.e. pairs in which the patient could receive a transplant directly from his/her intended donor. These pairs may nevertheless benefit from participating in kidney exchange, e.g. by getting a better matched kidney, or one from a younger donor, while incidentally helping a hard to match pair to be matched. Compatible pairs help in two ways: 1. they help highly sensitized pairs by making the pool of easy-to-match pairs thicker, so that not only is there another donor from whom a hard-to-match pair could receive a kidney, but it is more likely that that pair can form part of a cycle or chain; 2. they increase the supply of O-donors (who are selected out of the exchange pool when compatible pairs

³⁵See Ashlagi et al. (2012a) for a theory of sparse pools based on random graphs, where chains play an important role.

don't participate).³⁶

3.3 Optimization

KE programs typically use optimization to find matches using cycles and chains within their pool. The static optimization problem is to maximize the (weighted) number of transplants using logistically feasible cycles and chains subject to the constraint that no pair or NDD is matched more than once. Weights capture weak priorities of the platform. This problem is NP-complete (Abraham et al., 2007) even without chains.³⁷ Following Roth et al. (2007) and Abraham et al. (2007), researchers have produced different formulations and algorithms to solve this problem. We discuss some ideas briefly.

Consider the (directed) compatibility graph, $G = (V, E)$. The set V can be partitioned into nodes P consisting of patient-donor pairs, and N , consisting of NDDs. An edge connects one node to another if the donor at the first node is compatible with the recipient at the second node. Let w_e be the weight on edge e and for each exchange C let w_C be the sum of the weights of edges in that exchange. Let \mathcal{C}_k be the set of cycles with at most k edges and \mathcal{Ch}_j the set of chains with at most j edges. (By adding a directed edge from each pair to each NDD, chains can be viewed as cycles that include an NDD.) Let $\mathcal{C}_k(v)$ and $\mathcal{Ch}_j(v)$ be these subsets of cycles and chains that contain node v , respectively. A simple formulation that allows cycles of length at most k and chains of length at most j is:

$$\begin{aligned} \max \quad & \sum_{C \in \mathcal{C}_k \cup \mathcal{Ch}_j} w_C z_C \\ \text{s.t.} \quad & \sum_{C \in \mathcal{C}_k(v)} z_C + \sum_{C \in \mathcal{Ch}_j(v)} z_C \leq 1 \quad v \in V, \\ & z_C \in \{0, 1\} \quad C \in \mathcal{C}_k \cup \mathcal{Ch}_j \end{aligned}$$

For small pools and short cycles and chains, an optimization solver will often manage to handle the above formulation. Other algorithms use formulations based on flows. For each $v \in V$, let $\delta^-(v)$ be the edges pointing to v and $\delta^+(v)$ be the edges outgoing from v . Similarly define $\delta^-(S)$ and $\delta^+(S)$ the edges that point in to and out of a set of nodes $S \subseteq V$. When pools are large the number of variables (chains and cycles) grows exponentially large.

One simple formulation is to maximize flow (without cycles and chain variables) under the constraints: (i) incoming flow equals outgoing flow, and (ii) each node has at most one unit of outgoing flow. One should further rule out exchanges of infeasible size. If the pool is large or dense there can be exponentially many of these constraints. Instead one can solve

³⁶See e.g. Roth et al. (2005b, 2008); Gentry et al. (2007) and Sönmez et al. (2020). The latter paper considers a model with a continuum of pairs, in which no patients are too highly sensitized to be matched, but in which blood type O patients still benefit from the inclusion of compatible pairs.

³⁷The problem can be solved in polynomial time in two cases: if all exchanges are limited to size two this becomes a maximum matching problem, and if chains and cycles are of an unbounded size this becomes a maximum flow problem.

this formulation using constraint generation techniques, re-solving the problem constraining only against infeasible exchanges that have been previously found.

Researchers developed and adopted a variety of formulations (including objectives and constraints on feasible exchanges) as well as methods (e.g., column generation and branch and price) to address this optimization problem (Biro et al., 2009; Glorie et al., 2012; Manlove and O’malley, 2015; Anderson et al., 2015b; Constantino et al., 2013; Alvelos et al., 2019; Dickerson et al., 2012c; Alvelos et al., 2019).³⁸

How important is the use of sophisticated optimization? Our experience (and that of the American inter-hospital exchanges) is that most instances can be solved with simple algorithms.³⁹ This is because pools at the steady-state are usually not too large and not too dense (hard-to-match pairs are the ones that usually accumulate and easy-to-match pairs match quickly). In fact, a large and dense steady-state pool probably signals flaws in the matching process. It is important, however, not to miss these very highly sensitized patients, who are typically matched through chains or cycles of size longer than 2.

3.4 Matching in a Dynamic Pool

Kidney exchange pools are dynamic with patient-donor pairs and NDDs arriving and matching over time. There is an inherent trade-off between identifying matches faster to reduce waiting times and waiting for more enrollments to create more match opportunities. KE programs face the decision of when to identify matches. Figure 4 illustrates the potential benefit from waiting.



Figure 4: Waiting can lead to more matches. Pairs a,b,c are existing in the pool and d arrives in the future. Matching b with c without waiting results in one instead of two exchanges.

In the US, platforms shifted gradually towards small batch sizes, which means matching very frequently. National programs including the APKD and the NKR identify exchanges on a daily basis, and the UNOS program identifies matches twice a week. One concern is that this behaviour is driven by competition across platforms as hospitals may enroll pairs in multiple platforms. But even the (exceptionally productive) single-hospital program at Methodist Specialty and Transplant Hospital at San Antonio (MSTH) runs matches on a daily basis. In other countries the situation is different. National programs in Canada, United Kingdom, Netherlands, Czech Republic, Australia, and other countries search for exchanges every 2 – 4 months (Ferrari et al., 2015; Malik and Cole, 2014; Johnson et al., 2008). We briefly discuss research on this front.

³⁸The open source KE software mentioned above uses formulations in Anderson et al. (2015b) and is provided at <https://web.stanford.edu/~iashlagi/>. One formulation is based on the travelling salesman problem and is able to handle long chains effectively.

³⁹Even searching by brute force often works.

A matching policy employed by a platform determines which exchanges to implement and when. Two simple and commonly used policies are *greedy* and *batching*. The greedy policy forms exchanges as soon as they become available. A batching policy identifies an optimal set of exchanges in the pool every fixed number of periods.

In a simulation study Ashlagi et al. (2018) measure the effect batching policies have on efficiency (measured by the fraction of matched pairs and waiting times). The study uses APKD and MSTH data, which have different pool compositions. Pairs sampled from the data arrive according to a Poisson process and depart according to an exponential random variable, unless they match earlier. Matches are done through 2 and 3-way cycles and chains with different lengths. The main finding is that very frequent matching results in almost no harm. Figure 5 plots the fraction of matched pairs and the average waiting time under different matching frequencies and different sets of weights using the APKD data.⁴⁰ Assigning high priority to highly sensitized patients increases their match rate, but at the expense of hard-to-match and under-demanded pairs.

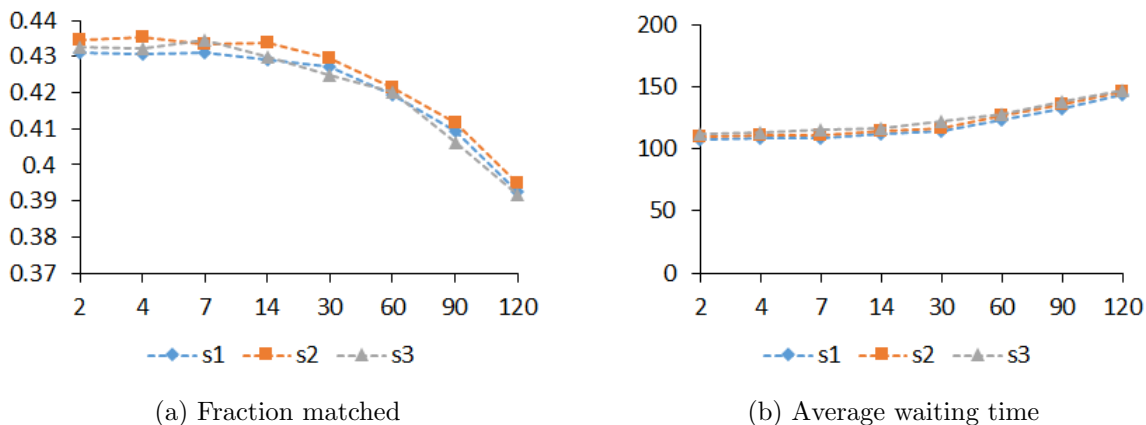


Figure 5: Batching policies using APKD data (from Ashlagi et al. (2018)). The x-axis represents the number of days between two match-runs. Each plot, S1-S3, stands for different prioritization/weights assigned to pairs based on the patients’ PRA.

The reason that the percentage matched goes down with the batch size is due to the departure rate. But even when the departure rate is low, Ashlagi et al. (2018) report that there is a very small benefit from waiting. Monteiro et al. (2020) find minor improvements from using batching when pairs either don’t depart or depart exponentially after the next batch.⁴¹

Figure 6 provide similar simulations results for numerous arrival rates using NKR data. While accumulating pairs does not increase the fraction matched, an increase in the arrival

⁴⁰The figure presented here assumes no frictions, but see Ashlagi et al. (2018) for similar findings from simulations that account also for delays due to blood shipping, refusal of offers and positive crossmatches.

⁴¹They further find interesting trade-offs between assigning or not assigning priority to pairs according to their waiting times.

rate does. We return to this insight in Section 4 when we discuss mergers and increasing the thickness of the KE pool.




Figure 6: Batching policies for different arrival rates using NKR data (from Agarwal et al. (2018)). Each plot corresponds to the batch size in days.

Why is this happening? Intuitively, under-demanded pairs as well as hard-to-match pairs with very highly sensitized patients accumulate in the pool. Suppose that now a potentially easy to match A-O patient-donor pair arrives to the pool. One possibility is that it can match immediately with one of the many O-A patient-donor pairs in the pool. If not, it is likely because the (newly arrived) A patient is very highly sensitized and part of a hard-to-match pair. This suggests that delaying other pairs from matching is unlikely to help this A-O pair in the near future. So when the departure rate is low, delaying easy-to-match pairs is unnecessary, because an easy-to-match pair is likely to match with one (of the many) hard-to-match pairs (and it would be inefficient to match together two easy to match pairs). When the departure rate is high, matching infrequently will result in many departures of easy-to-match pairs.

The class of batching policies is not exhaustive and the intuition above ignores some plausible scenarios. Suppose the donor of an easy-to-match pair is compatible with the patient of a hard-to-match pair but these pairs are not part of any exchange. It may be beneficial to wait for another pair to arrive that can close a 3-way cycle with both of these pairs. An open question is whether there are policies that can significantly outperform a greedy policy.⁴²

Theory. A growing literature helps in understanding matching policies on dynamic networked environments. Ünver (2010) characterizes the optimal policy in a dynamic kidney exchange model with linear waiting costs, in which compatibility is essentially determined only by blood types, i.e. in which highly sensitized pairs don't play any important role. A

⁴²We previously experimented with heuristic strategies, which haven't outperformed a greedy policy.

greedy algorithm is shown to be optimal for 2-way cycles and almost optimal for 2 and 3-way cycles.⁴³

One stream of papers considers models in which compatibilities are based on random graphs, modeling the sparsity due to sensitivity of patients. Several of these papers assume nodes (pairs) do not depart the pool unless matched (Anderson et al., 2017; Ashlagi et al., 2019b,c; Blum and Mansour, 2020). These papers find that greedy matching is optimal when minimizing average waiting times.⁴⁴ Other papers assume nodes depart the pool according to some hazard rate. Akbarpour et al. (2020d) seeks to maximize the number of matches and Ashlagi et al. (2019c) analyzes the trade-offs between waiting times and the number of matches. These papers also find greedy matching to be optimal in a large market.⁴⁵

More sophisticated policies have been considered by Dickerson et al. (2012b,a); Dickerson and Sandholm (2014). One of the key ideas is to take account of the potential of different kinds of patient-donor pairs in the compatibility graph (as shadow prices) to determine whether to match them or keep them in the pool.

What next? As kidney exchange grows, accounting for *match quality* (e.g., life years gained from transplant) may become increasingly attractive for platforms. Platforms do assign (different) weights to matches based on the characteristics of the donor and the patient. These weights are mostly determined in an ad-hoc manner within each KE program and are also largely fixed over time. More work is needed to explore the trade-offs between match qualities, number of matches, and keeping waiting times low.⁴⁶ Another theoretical avenue is to explore models that allow for more correlation in the compatibility graph than do the random-graph based models in the above papers. Finally, little is known about dynamic matching in the presence of match failures (described in more detail next).

3.5 Frictions

As already mentioned, doctors enter not only medical data but also preferences and thresholds for acceptable characteristics of a donor (e.g., age, creatinine level, BMI, weight). Despite these data, a significant fraction of proposed (computer-identified) matches fail to convert to transplants. For example in the UNOS KE program, prior to 2015, out of 2246 of

⁴³Using thresholds to facilitate 3-way cycles is beneficial, but generates relatively small improvements.

⁴⁴An elegant and simple analysis for policies based on chains is given by Blum and Mansour (2020). When 3-way cycles are also included, greedy is optimal among a large class of policies (Anderson et al., 2017).

⁴⁵The models assume different type structures and different limits regarding the growth of the market (e.g., whether the edge probability remains constant or decreases with the size of the pool).

⁴⁶Computational experiments that also incorporate compatible pairs have been done in Roth et al. (2004, 2005b) and Li et al. (2019). There are also a few theoretical studies that look at matchings of size two; Blanchet et al. (2020) characterizes, as a function of the quality distribution, the optimal amount of waiting in a stochastic model. Mertikopoulos et al. (2019) explore trade-offs between quality and number of matches. Ashlagi et al. (2020a) find that in a dynamic version of the classic assignment model (with known qualities) simple “queue-like” policies have almost no allocative efficiency loss. Kerimov et al. (2020) consider a dynamic model with deterministic match values and find that “small” batching is required to achieve low regret. For a dynamic model with no distributional information see Ashlagi et al. (2019a).

the potential matches identified (as part of exchanges), only 172 converted into transplants. This is due to the high rate of declined match offers as well as the high rate of positive cross-matches. We discuss these subtle issues briefly and point to relevant work and potential directions for research.

Match offers are presently declined at a rate of about 25-35% at the APKD and UNOS and around 20% at the NKR (Ashlagi et al., 2018; Agarwal et al., 2018; Hanto et al., 2008; Fumo et al., 2015). This suggests that the data and preference criteria are too coarse. Platforms now often ask physicians to pre-specify which donors (from a selected set of potential donors) would be acceptable to each of their patients. But in large pools this can still be a costly task. It is worth noting that the NKR (monetarily) penalizes a hospital that refuses a match offer it pre-selected to be acceptable. In contrast to the US, refusal rates are very low in countries with centralized KE programs (studies documenting the performance of such KE programs do not report such failures).⁴⁷

The chance that a virtual match results in a positive crossmatch is more than 35% at the APKD and UNOS for highly sensitized patients with PRA above 90 and around 10% on average for other patients. High numbers are also reported at the NKR and other countries such as Canada (Cole et al., 2015). These numbers were even higher when only a few HLA antigens were listed. As discussed in Section 2 hospitals determine in an ad-hoc manner the strength (MFI) threshold for listing an antibody.

How to reduce frictions? Numerous KE programs maintain a blood lab and do cross-matches in-house (examples include centralized KE programs in Europe, single center programs, and the APKD). This eliminates delays due to shipping blood samples after exchanges are identified.

It is common that a match between a patient and a donor is successful but the corresponding exchange is not (or only partially) implemented. KE programs attempt to use the knowledge from successful one-way matches to improve the information in the compatibility graph and increase the future match rate. See Fumo et al. (2015) for how APKD has handled some of these strategies. Some countries that find exchanges at a low frequency re-optimize over the compatibility graph after matches that result in a positive crossmatch. Matching infrequently without re-optimizing while having high match failures is likely to lead to a very low match rate (as this delays discovering the actual compatibility graph).⁴⁸ Simulations that consider this tension can be found in Santos et al. (2017) and Ashlagi et al. (2018).⁴⁹

Some studies consider the problem of maximizing the expected number of matches (Kliementova et al., 2016; Dickerson et al., 2019; Wang et al., 2017; Bidkhorji et al., 2020) or using robust optimization (McElfresh et al., 2019). Several papers consider identifying exchanges that are not necessarily disjoint, after which more than one crossmatch would be possible while having “backup” options (Blum et al., 2013; Bray et al., 2015; Glorie et al., 2015). Other papers ask which links in the compatibility graph one should test if pre-testing is

⁴⁷See, e.g., (Ferrari et al., 2015; Biró et al., 2019c; Kal-van et al., 2011).

⁴⁸This helps explaining the low match rate in UNOS.

⁴⁹See also Carvalho et al. (2020); Dickerson et al. (2020).

feasible (Assadi et al., 2019; Blum et al., 2020; McElfresh et al., 2020). These studies do not consider match rate over time.

An important challenge when attempting to reduce frictions is not harming highly sensitized patients, who are the most likely to have a positive crossmatch.⁵⁰ Adam Bingaman and Cathi Murphey who collaborate on kidney exchange at MSTH (a highly successful kidney exchange at a single transplant center) conduct many crossmatches for each highly sensitized patient. They feel that not having to coordinate with other hospitals gives them more freedom to conduct as many crossmatch tests as they think desirable for highly sensitized patients.⁵¹

4 Thickening the network

4.1 What does it mean to have a thicker pool?

The composition of the pool, arrival rates (enrollment), and departure rates (without being matched) are key factors for the potential number of transplants a KE program can facilitate. To illustrate this, Figure 7 plots simulation results using data from different KE programs between the years 2012-2014. Each plot measures the fraction of the pool matched for different rates of arrivals.⁵² During this time period, NKR, APKD and UNOS had very similar pool compositions,⁵³ but many more pairs and NDDs enrolled to the NKR.⁵⁴ Not surprisingly, countries with national KE programs also vary significantly with respect to the number of pairs registered and their pool compositions.⁵⁵

NDDs and easy-to-match pairs (including compatible pairs, who are invited to enroll in some kidney exchange pools to find a better matched kidney than that of their compatible intended donor) all positively contribute to the composition of the pool. These, together with a large arrival rate, will lower waiting times and increase the match rate for all pairs (especially hard-to-match pairs).⁵⁶

⁵⁰Indeed this is still poorly understood and just a few simulations of matching policies have been considered (Santos et al., 2017; Ashlagi et al., 2018).

⁵¹Cathi : “I don’t track how many crossmatches came out positive. A lot of the times I think it will be positive but I want to do it anyway”.

⁵²The experiments simulate operational details and use departure rate estimates. Pairs and NDDs are sampled with replacement from the corresponding data. The simulation generating the 1st quartile at the NKR uses pairs and NDDs from hospitals that enroll among the top quartile of hospitals with respect to the percentage of pairs they enroll to the NKR.

⁵³Roughly 13-15% over-demanded pairs and about 42-43% under-demanded pairs.

⁵⁴In 2014, 155, 244, and 243 pairs and 5, 4, and 51 NDDs enrolled to the APKD, UNOS, and NKR, respectively. The NKR and APKD now offer NDD’s a “voucher” intended to be used in the future by a family member, who would receive a chain-ending kidney (see e.g. Wall et al. (2017); Krawiec et al. (2017); Akbarpour et al. (2020b)).

⁵⁵For example, the Dutch program registered 655 pairs between 2004 and 2014, UK registered 991 pairs prior to 2014, Canada registered 491 pairs between November 2009 and December 2013 France had 78 pairs between 2013 and 2018 (Biró et al., 2019c; Ferrari et al., 2015).

⁵⁶That an increase in the match rate results in lower waiting times is intuitive and can formally be understood through Little’s law (Little and Graves, 2008), which states that the arrival rate multiplied by

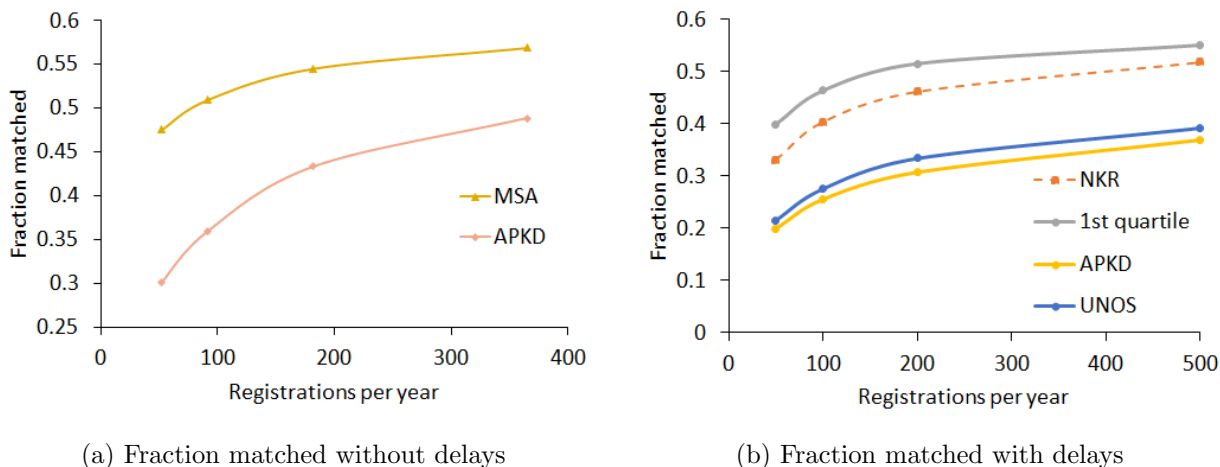


Figure 7: Pool composition and scale (adapted from Ashlagi et al. (2018) and Agarwal et al. (2018)). The left and right plots are generated without and with delays, respectively (delays due to receiving match offer responses and crossmatch tests).

The thickness of the pool depends not just on how many patient-donor pairs are available, but on how highly sensitized they are. Pools often contain a large fraction of highly sensitized patients, which means that the underlying compatibility graph is quite sparse. In such pools short cycles, especially between just 2 pairs, will be difficult to find. But chains initiated by NDDs will be very valuable, and can often be quite long and lead shorter waiting times.⁵⁷

To make this clear, consider a pool in which all patients have a PRA of 99%, (i.e. each patient has only a 1% chance of being compatible with an otherwise blood type compatible donor). In a large enough pool, the chance that one pair can find another to donate to isn't small, but the chance that the second pair can donate back to the first is still only 1%, so pairwise exchanges will be rare (This is why easy to match pairs are so valuable, since they make it possible to match hard to match pairs in short cycles containing the easy to match pair.). But if a NDD donates to a hard to match pair, the chance that pair can pass it forward to some other hard to match pair, who can continue to pass it forward is not so small, which is how long chains form.

Theoretical studies looked at the benefit of chains and cycles in dynamic settings, where the underlying model is based on random graphs (Anderson et al., 2015a; Ashlagi et al., 2019b; Blum and Mansour, 2020). These papers find that in sparse pools with many highly sensitized patients, chains result in significantly lower waiting times than relying only on cycles. The addition of easy-to-match pairs (or NDDs) reduce the need for long chains and make short cycles more common (Ashlagi et al., 2019b; Blum and Mansour, 2020).⁵⁸

the average waiting time equals the average number of pairs waiting.

⁵⁷See Ashlagi et al. (2012b); Ashlagi and Roth (2012, 2014); Ashlagi et al. (2013); Anderson et al. (2015b); Bingaman et al. (2018).

⁵⁸See also (Gentry et al., 2007; Sönmez and Ünver, 2014).

Another potential way to increase thickness is increasing the density of the compatibility graph, by increasing the heterogeneity of HLA and antibodies in the pool. For example highly sensitized patients in some region may be less sensitized to donors in other regions. We speculate that there is such potential in international exchange, especially due to heterogeneity in the specific allele level.

4.2 How and whether to merge

KE programs in countries like the US and Israel seek the participation of more hospitals, to merge patient-donor pools or collaborate with each other in other ways.⁵⁹ When measuring the fraction of the pool transplanted, the marginal benefit from increasing the pool with the same composition decreases (Figure 7).⁶⁰ While increasing the arrival rate beyond some efficient size will have a minor impact on the fraction of the pool transplanted, patients may still benefit from lower waiting times because this is similar to having a faster matching process. (Ashlagi et al., 2018).

When hospitals are not forced to participate, they can choose how to participate and which pairs or NDDs to enroll. Hospitals can choose to match some of their pairs internally (Sönmez and Ünver, 2013b; Ashlagi and Roth, 2014). This behavior is common in the US. Agarwal et al. (2019) find that more than 50% of the exchange transplants between 2012-2014 were done through internal exchanges. Hospitals select harder-to-match pairs on average to register with inter-hospital KE programs than they match internally (and internal matches are less efficient, e.g., more than 11% of the recipients of internally matched O transplanted kidneys have a non-O blood type, while in the large platforms this percentage is around 2.5%).

To engage participation of hospitals one needs to pay attention to the possibility of free-riding behavior. Several inter-hospital KE programs now use a point system to incentivize hospitals to enroll their NDDs, which create a big positive externality, since they can initiate large chains of transplants, that benefit patients of multiple hospitals. These points were used to decide in which hospital to terminate a chain.⁶¹ But easy-to-match pairs also have a positive benefit—the way to match hard to match pairs is to match them with easy to match pairs—while many hard-to-match pairs will remain unmatched, or will only be matched instead of another hard to match pair. A hospital that enrolls only under-demanded pairs does not help to generate extra transplants and in fact negatively affects other hospitals. Ashlagi and Roth (2014) advocated using a “frequent flyer” system to alleviate free-riding behavior. Agarwal et al. (2019) find that a simple point system that assigns for each type of pair and NDD its marginal benefit can improve transplant centers’ incentives to participate

⁵⁹The APKD, NKR and UNOS each serve more than 50 hospitals, which register their pairs and NDDs.

⁶⁰This fraction converges to the optimal fraction matched in a sufficiently large static pool of the same composition.

⁶¹The idea is to keep track of which hospitals have contributed NDD’s to the inter-hospital KE program (instead of having the NDD donate to someone on the originating hospital’s deceased donor waiting list). In this way, contributing hospitals can have some chain-ending donations directed back to their deceased donor waiting lists, so that they are ‘repaid’ for the NDD.

fully in an exchange platform (by registering easy to match as well as hard to match pairs).⁶² For example, in a thick pool, an additional over-demanded, easy to match A-O patient-donor pair would typically generate about two transplants, one for its own patient and one for an under-demanded O-A. The actual value of a particular pair, however, depends on the sparsity of the network and how sensitized the A patient is. In contrast, an additional under-demanded pair that is transplanted means that another under-demanded pair is not. (To carry forward the frequent flier analogy, enrolling easy to match pairs earns a transplant center frequent flier points, while getting a hard to match pair transplanted spends some of its points.)⁶³

The NKR is now using a mixture of incentive schemes for hospitals. First, it reduces registration fees for a hospital that commits to enroll all its pairs.⁶⁴ Second, inspired by Ashlagi and Roth (2014), it adopted a point system (called the Center Liquidity Contribution system) that assigns a value for each type of pair that is aligned with its' average marginal benefit. The points are utilized to break ties between "optimal solutions".

There are several theoretical and practical questions regarding point systems: When should values be updated? Should hospitals be penalized when they have a deficit? Should points be used only as tie-breakers? How should a point system behave when hospitals have different size and different compositions?

Single center KE programs. These incentive issues aren't present in the same way at single-center KE programs that achieve efficient scale. MSTH is the largest single center KE program in the U.S., with more than 500 transplants since inception in 2008.⁶⁵ Some features of the program include: prospective education, flexibility in assigning antibodies and many crossmatch tests⁶⁶, storage of blood samples, ABOi transplants, and involvement of compatible pairs (Bingaman et al., 2012).

Part of MSTH's success at recruiting compatible pairs to participate in kidney exchange has to do with the operational details of the way they introduce KE. While many centers invite pairs into KE only after they prove to be incompatible, MSTH reports that they indicate at the outset of compatibility testing that the object is to obtain the best transplant possible, and that this might involve KE even for a compatible pair, e.g. to obtain a kidney from a younger or better matched donor.

Compatible pairs are assigned a high priority in match-runs and are usually matched with young donors. An important consideration is the preferences of donors for when to donate, especially those of compatible pairs.⁶⁷ This success is likely to be hard for small

⁶²The paper offers a simulation framework to calculate these points, in a way that recalls the Shapley value (Shapley, 1953; Roth, 1988).

⁶³See also Hajaj et al. (2015) who study a credit system in dynamic kidney exchange.

⁶⁴The NKR is the only platform we are aware of that charges fees for different types of services.

⁶⁵More than UNOS and APKD and many countries. See here a report of MSTH from 2018 (Bingaman et al., 2018).

⁶⁶Including intensive testing of potential one-way matches especially with highly sensitized patients.

⁶⁷Cathi Murphey: "donors play a big role in the timing of exchanges."

hospitals to replicate.⁶⁸

Another remarkable single-center KE program is in India, at the Trivedi Institute of Transplantation Sciences in Ahmedabad, India (Kute et al., 2018). Kidney exchange is particularly important for India, because, in its absence, many patients are transplanted with incompatible kidneys, which requires them to be treated with large amounts of immunosuppression. In India this turns out to be quite dangerous, and many such patients succumb to opportunistic diseases. But kidney exchange allows patients to receive compatible kidneys, so that they require less immunosuppression and suffer less infection.⁶⁹

An interesting research direction is to understand better the trade-offs between single and multi-hospital KE programs.

Other collaborations. One way in which pools of patient-donor pairs available for kidney exchange can be merged is via international exchanges, across borders. Recent examples include a Czech-Austrian exchange (Böhmig et al., 2017), a chain between Israel and Czech Republic⁷⁰, the newly formed kidney exchange network of Scandiatransplant, which creates a single pool of patient-donor pairs across the Nordic countries, and efforts to merge pools among kidney exchange networks elsewhere in Europe (Biró et al., 2019c).

Several collaboration attempts between KE programs haven't been very fruitful (e.g., between NKR and APKD, and between Spain, Italy, and Portugal).⁷¹ The pattern is similar; the match-run was done on “left-over” pools after each KE program ran its own match. These pools consist mostly of under-demanded pairs and very sensitized patients.⁷²

Cross border collaborations may face free-riding like behavior. Some form of agreements or incentive schemes as well as transparency are likely to be necessary for these to succeed. In addition to the proposal by Agarwal et al. (2019), Biró et al. (2019b) offers to use matching mechanisms based on the Shapley value, and Klimentova et al. (2020) considers maximizing under fairness constraints, at the level of countries. Collaborations may sometimes be especially difficult when countries face different legal, regulatory, and operational restrictions governing the kinds of exchanges that are feasible.

4.3 Global kidney exchange

More patients can be helped if kidney exchange can include more countries. Indeed, kidney failure is a global problem, and transplantation is the treatment of choice, but there is a worldwide shortage of transplants. However this shortage has somewhat different causes

⁶⁸Indeed numerous hospitals attempted to follow the footsteps of MSTH but eventually chose to participate in a multi-hospital KE program.

⁶⁹While there are special difficulties associated with kidney transplantation in India, the available technology is exceedingly modern: see e.g. <https://marketdesigner.blogspot.com/2019/05/robot-assisted-kidney-transplantation.html>.

⁷⁰https://www.mzv.cz/telaviv/en/bilateral_relations/unprecedented_czech_israeli_cooperation_1.html

⁷¹The combined pool between the three countries included 113 pairs and resulted in one 3-way cycle <https://www.lamoncloa.gob.es/lang/en/gobierno/news/Paginas/2018/20180808transplant.aspx>.

⁷²See (Biró et al., 2019c; Mincu et al., 2020) for more forms of collaborations in European countries.

in the rich countries of the developed world than in the middle income countries of the developing world. In wealthy countries, a shortage of transplantable organs is the limiting factor, while in many middle income countries, financial barriers prevent transplantation for much of the population, despite the fact that there are transplant centers that can perform transplants for those who can afford them. Liyanage et al. (2015) estimate that 2-7 million people die every year worldwide due to inability to pay for dialysis or kidney transplantation.

Recent efforts to make kidney exchange truly global have included efforts to invite into American kidney exchange patient-donor pairs coming from countries in which kidney exchange is unavailable, or in which the national health insurance doesn't cover the full costs of transplantation for many citizens. One obstacle is that there will often be financial barriers that prevent the health insurance system of those countries from paying the costs incurred by their citizens in the U.S. However this needn't be an insuperable obstacle, because transplantation is much cheaper than dialysis, and so substantial savings accrue to American health care payers whenever an American patient is transplanted, sufficient to pay the costs of the foreign pair in the U.S., and after they return home (Krawiec and Rees, 2014; Rees et al., 2017a; Bozek et al., 2018). These pilot KE programs, still in their infancy as practical alternatives, offer the prospect of enabling patient-donor pairs from around the world to assist each other in receiving transplants.

5 Maintaining social support

Kidney exchange requires a lot of organization and logistics. Wouldn't it be simpler to buy kidneys from willing sellers, so that each patient in need could receive a compatible kidney directly from a living donor? It turns out that things are not so simple: there are laws almost everywhere in the world against paying donors for organs. In the U.S., the relevant legislation is the National Organ Transplant Act (NOTA), 42 U.S.C. 274e 1984, which states in part "it shall be unlawful for any person to knowingly acquire, receive or otherwise transfer any human organ for valuable consideration for use in human transplantation".

Transactions that some people would like to engage in but others think they should not be allowed to are called *repugnant transactions* (Roth, 2007).⁷³ Kidney exchange was organized with the intention of increasing transplantation without violating the law against compensating donors, or arousing the repugnance that motivates the law.

⁷³Of course some transactions may have negative externalities that directly harm third parties, and it is easy to understand the objections of those who are harmed. When we study repugnant transactions as such we normally concentrate on transactions in which the harms to those who object are difficult to measure—sometimes by specifying that a third party can't detect that the transaction has taken place unless someone tells them (Ambuehl et al., 2015). Thus e.g. same-sex marriage is a transaction that two people would like to engage in, by marrying each other, but others may think they shouldn't be allowed to, although you can't even tell if someone is married unless someone tells you (e.g. by wearing a wedding ring). Receiving payment for a kidney (as opposed to donating it altruistically) falls into this category of repugnant transactions. There are a wide variety of repugnant transactions that are banned in some places and legal in others, many involving payment for body parts (e.g. surrogacy is widely legal in the U.S. but widely illegal elsewhere, as is payment for blood plasma).

Note that the language of the 1984 NOTA bans "valuable consideration" for an organ for transplant, not just monetary compensation. So it wasn't completely clear that kidney exchange was fully legal in the U.S. But in 2007, some years after kidney exchanges started to be performed in the U.S., Congress amended the NOTA to say that kidney exchange does not involve valuable consideration of the kind considered by the Act, suggesting that kidney exchange does indeed avoid much the repugnance associated with kidney sales.⁷⁴

That is not to say that kidney exchange hasn't faced some opposition. In Germany, it is essentially illegal. (German law says that, barring the intervention of a judge, a patient can only receive a living donor transplant from a member of his or her immediate family.)

Global Kidney Exchange has also been viewed as repugnant in some quarters, (see e.g. Delmonico and Ascher (2017)) and replies by Rees et al. (2017b); Roth et al. (2017); Marino et al. (2017); Roth et al. (2020)). Positions against it have been taken by a number of European transplant societies (Europe, 2018). But it has received support from the American Society of Transplant Surgeons, from the European Society of Transplantation committee on Ethical, Legal and Psychological Aspects of Transplantation (Ambagtsheer et al., 2020) and from Italy at the World Health Organization⁷⁵. A very clear analysis of the ethical issues involved (concluding with very strong support of GKE) is given in the *Lancet* by a trio of moral philosophers, Minerva et al. (2019).⁷⁶ Perhaps most significant for the future of GKE is that it has been well received in countries whose patient-donor pairs have benefited from kidney exchange in the U.S.⁷⁷ We are thus optimistic that kidney exchange will continue to expand, so that more patients and donors from around the world can receive transplants of compatible kidneys.

6 Open questions and research directions

This section recapitulates a few of the research opportunities mentioned in previous sections.

1. Social support for transplantation innovation. Why does kidney exchange face big legal obstacles in Germany, but not in the U.S. or a number of other European countries? Why have some transplant professionals regarded Global Kidney Exchange with repugnance?

⁷⁴The amendment, known as the Charlie W. Norwood Living Organ Donation Act, Public Law 110-144, 110th Congress, 2007 passed without any dissenting votes in either the House or the Senate <https://www.govtrack.us/congress/bills/110/hr710>

⁷⁵see <https://asts.org/about-asts/position-statements#.Xm10nc5Kg2w>, <https://marketdesigner.blogspot.com/2017/12/more-on-endorsement-of-global-kidney.html>, and <https://marketdesigner.blogspot.com/2018/01/italy-recommends-global-kidney-exchange.html>

⁷⁶Minerva et al. (2019) summarize their conclusions by saying "Misguided objections should not be allowed to prevent the GKE from realising its potential to reduce suffering and save the lives of rich and poor patients alike." See also Singer et al. (2020).

⁷⁷see e.g. the cover story of the April 14, 2017 issue of *Newsweek En Espanol*, by Iván Carrillo, which covers GKE between the U.S. and Mexico, and speaks of it as building "UN PUENTE DE VIDA," a bridge of life. <https://marketdesigner.blogspot.com/2017/04/a-bridge-of-life-global-kidney-exchange.html>

Why is the U.S. progress on removing financial disincentives from living kidney donation so slow? One simple hypothesis is that laws facilitating or banning different markets reflect quite different popular opinion in different jurisdictions. But Roth and Wang (2020) find little support for this hypothesis: public opinion in Germany shows strong support for legalization of kidney exchange, for example.⁷⁸ Because innovations in transplantation require social and sometimes legislative support, a better understanding of the connection between those will help make further progress. (Section 5).

2. International exchanges. What are the benefits from international collaborations that stem from biological heterogeneity (and not simply due to larger scale)? How should the rules, incentives, and operational procedures be designed to promote fruitful, ongoing collaborations? (Section 4).

3. Initiating non-directed donor chains with a deceased donor kidney. Presently, the non-directed donor who initiates a chain is virtually always a living donor, but it has been proposed that such chains can be initiated with the kidney from an appropriate deceased donor (Melcher et al., 2016). (See also Furian et al. (2019), Furian et al. (2020), and Wang et al. (2020).) However regulatory barriers still need to be navigated before this becomes commonplace, and this will involve investigating and developing new operational procedures. The integration of the deceased donor waiting list raises several questions. Can one design a policy that does not harm patients on the deceased donor waiting list who have no intended living donor? For example initiating chains with O donors may harm O patients on the list. And when should a chain be terminated with a final donation back to someone on the deceased donor waiting list who does not have a willing living donor to continue the chain? Among other differences from conventional chains, non-directed living donors are known in advance and remain available for a while, but deceased donors often become available with little notice, at which point their organs must be transplanted immediately. So consideration will have to be given how to start such chains quickly, and continue them efficiently. (Section 2.1).

4.Improving on greedy matching. When exchange pools contain many hard to match pairs, greedy matching works well, because easy-to-match pairs and non-directed donors can often be efficiently matched when they arrive. How can greedy matching be improved? Suppose for example that a pair arrives that can donate a kidney to the patient in some very hard to match pair, but there isn't yet a cycle or chain that includes that transplant. Can efficiency be improved by keeping both pairs in the pool in the hope of facilitating such an exchange? (Section 3.4).

⁷⁸Roth and Wang (2020) study public attitudes towards global kidney exchange, commercial surrogacy, and prostitution in Germany, the Philippines, Spain and the U.S., where the laws vary widely. They find public support for legal surrogacy and global kidney exchange both where it is banned and where it is legal. On a different matter of repugnance related to kidney transplantation, see e.g. Leider and Roth (2010) and Elías et al. (2019) on public attitudes towards compensation for kidney donors.

5. Matching optimization. What should be optimized when choosing a set of transplants? Current optimization protocols embody (explicitly in match weights, and implicitly in match frequency) ad hoc trade-offs involving match qualities, number of matches, and waiting times. But match quality is multi-dimensional. There are immediate observables, such as number of HLA mismatches. But these are proxies for other objectives that need to be better estimated from data and effectively incorporated into matching objectives, such as expected graft life, and how sensitized the patient will be after the current graft fails and a new transplant is needed. Moreover, with heterogeneous match qualities (additional life years), greedy matching policies may not be efficient. We need a better understanding of dynamic matching when matches yield different values.⁷⁹ (Section 3.4).

6. Dynamics with frictions. Little is understood regarding optimal and simple dynamic matching policies in the presence of match failures (such as positive crossmatches). Fairness concerns as well as trade-offs between long and short term efficiency arise as highly sensitized patients are more likely to have a positive crossmatch (Section 3.5). To better match over time it is helpful to study when and why pairs leave the platform without a match (Section 3.4). Moreover, many donors often have certain periods in which they are available to donate.⁸⁰ We need further study of dynamic mechanisms to control the incentives for hospitals to participate fully in multi-hospital KE programs (Section 4).

7. Frequent flier point systems for contributions to the exchange pool. Systems for giving credit to hospitals for contributing nondirected donors and easy to match pairs to the pool (to be balanced against contributing hard to match pairs) need to be designed to accurately measure contributions, and to provide appropriate incentives. As behavior and pool composition change, the appropriate credits will change. And some hospitals may build up large credits and others large deficits, which will alter incentives. So the system must be designed to allow it to be dynamically maintained to retain its desirable properties. We therefore need to understand better dynamic mechanisms (without money) that achieve high efficiency. A natural approach is to prioritize matching hard-to-match pairs of hospitals that accumulated credits. If ties were common (the graph was dense) one can expect to sustain cooperation and maintain a stable system.⁸¹ Ashlagi and Kerimov (2019) find that even minimal ability to break ties between hospitals enables maintaining a stable credit system, which keeps credits from diverging.⁸² Some challenges may arise due to the different size and composition (of hard- and easy-to-match pairs) in hospitals. Moreover, when should the

⁷⁹See Ashlagi et al. (2019a); Blanchet et al. (2020); Mertikopoulos et al. (2019); Kerimov et al. (2020) for theoretical papers that consider centralized dynamic matching.

⁸⁰In this connection, there is often information both about when currently available donors will no longer be available, and when donors who are not currently available will become available.

⁸¹See for example Kash et al. (2015); Johnson et al. (2014); Hauser and Hopenhayn (2008); Möbius (2001) for literature on scrip systems and trading favors in dynamic games.

⁸²The main idea relies on the power of two choices in computer science (Mitzenmacher, 2001), which allows creating the desired drift in the credits of each hospital. This is encouraging due to the sparsity in kidney exchange pools.

value of credits be updated (noting that simple systems would be desired for transparency)? How should credits be utilized? Should hospitals be penalized if they have negative credits?⁸³ (Section 4.2)

8. Single center versus inter-hospital exchange. When might it be advantageous for a single transplant center with a sufficiently large exchange pool to operate independently of multi-hospital exchange? I.e. coordinating with other hospitals brings benefits in terms of thickness of the pool, but costs in coordinating crossmatch tests and other policies. And when it comes to innovation there are trade-offs between size and nimbleness. How should we assess these? How should the protocols of inter-hospital exchange networks be altered to make them more attractive to successful single center kidney exchange programs? (Section 4.2)

9. Eliciting and modeling preferences. What preferences should a KE program elicit, from patients, donors, surgeons, and transplant center directors? It may be useful to develop machine learning models to predict positive crossmatches and match refusals, and to refine optimization models that seek to maximize long-term consequences of transplants (such as graft survival and subsequent patient sensitivity), and to understand the trade-offs involved with waiting (while on dialysis) for a better match.⁸⁴ A good starting point for exploring this direction is the growing literature on waitlist mechanisms motivated by the allocation of cadaver organs (which also suffer a high refusal rate, and discarded organs), as well as waitlists for public housing.⁸⁵ While it is natural to explore patients' or surgeons preferences, we note that, in kidney exchange, donors' preferences can also play an important role. (Section 3.5.)

10. Understanding the bridge(s) between theory and practice. Kidney exchange has profited mightily from a literature ranging from stylized mathematical models published in analytical journals, to clinically oriented studies of designs adopted in practice and reported in medical journals. In between are theoretical papers that model important institutional details relevant to practice, and empirical papers that take careful account of clinical data. Theory and practice cross-pollinate each other, but there is room to learn how to do that better. In particular, tractable mathematical models that lead to theorems with potentially practical conclusions nevertheless often depend on simplifying assumptions that, on their face, ignore important facts on the ground.⁸⁶ Early, idealized models of kidney exchange (some written before we realized how idealized they were) sometimes looked at very

⁸³Kanoria and Qian (2019) identify settings, in which utilizing credits via a max-weight algorithm can increase efficiency.

⁸⁴Data from both deceased and live donations can assist with learning about and modeling refusals.

⁸⁵See, e.g., Su and Zenios (2006); Bloch and Cantala (2017); Leshno (2014); Baccara et al. (2020); Arnosti and Shi (2019); Ashlagi et al. (2020a,b) for theoretical studies and Agarwal et al. (2020a); Ata et al. (2020); Agarwal et al. (2020b) for empirical studies.

⁸⁶Market design theorists must of course be careful that the results they derive aren't driven primarily by those of their assumptions that are unrealistic.

large markets, because limit theorems or continuum assumptions can yield analytical results more easily than detailed finite models. But these models abstract away from perhaps the most important problem of kidney exchange, which is that there are some pairs that are very hard to match, because they have a very low probability of being able to accept any particular kidney. (When there are infinitely many donor kidneys, this problem vanishes by assumption.) But the analytical attractiveness of large market models ensures their continued supply, and we have learned about how to estimate some aspects of their connection to practical exchange by comparing the limit results to simulations.⁸⁷ Better theory about speeds of convergence and related matters will also help to bridge this kind of gap. Theory is also easier if we assume that all exchanges will be between only two pairs (because then we can tap into elegant nondirected graph theory). In the early days of kidney exchange this actually matched clinical practice, but kidney exchange today depends heavily on chains. Theoretical results that bridge this gap would be useful too, e.g. to help us understand which kinds of results from models of pairwise exchange can be expected to carry over to exchange that includes chains. More generally, models of varying degrees of abstraction sometimes give insights that could have been missed in the welter of important detail, but there is room for theory to guide modeling by directing our attention to critical details that cannot be ignored for at least some kinds of conclusions.

7 Concluding remarks

The design of kidney exchange networks has involved almost continuous modification, sometimes at high levels (e.g. what kinds of exchanges to perform) and sometimes at the most detailed of operational levels (e.g. how to communicate with hospitals).⁸⁸ This makes it quite different from the design experience of some other matching markets, like medical labor clearinghouses and school choice. In those applications, a given design could stay in place for a number of years, and be modified only occasionally. Perhaps the difference is that, hiring medical residents and fellows isn't the everyday work of hospitals, nor is matching students to schools the main work of schools. Instead, those markets only operate annually, and the rest of the time are not subjected to constant exploration by the participants. In contrast, kidney transplantation is the daily concern of transplant centers, and so every aspect of the kidney exchange market design is examined and tested on an almost daily basis. And, in the United States, the three principle inter-hospital exchanges compete with one another to a certain extent, including the fact that some pairs are registered in more than one exchange.

Nevertheless, the design of kidney exchanges shares a lot in common with other applications of market design (cf. Roth (2008, 2015, 2018)). The first task of a marketplace—in this case the inter-hospital kidney exchange network—is to make the market thick. For kidney exchange, this means assembling a sufficiently large, constantly renewed pool of patient-donor

⁸⁷Here too, care needs to be taken so that the simulations model relevant clinical detail.

⁸⁸The different levels of detail at which practical market design operates have generated a number of metaphors, from architecture (Wilson, 2002), to engineering (Roth, 2002), to plumbing (Duflo, 2017).

pairs.

A second task is to tame the congestion that can arise in arranging and conducting transplants in a large pool of pairs. At a high level, non-simultaneous chains removed the logistical bottleneck of having to assemble many surgical teams and operating rooms at once, and powerful integer programming heuristics allow the computationally complex computations needed to produce a match proposal to proceed quickly. At a more nitty-gritty level, arranging for reliable virtual crossmatches, for moving blood samples for physical cross matches, for transportation of organs, for communicating among transplant centers, are all operations that are designed more smoothly today, because of many incremental changes.

A third role of a well designed kidney exchange clearinghouse is to make all transactions safe and reliable. This involves high level issues like arranging standard financial charges (so that hospitals with different costs of nephrectomies can nevertheless exchange kidneys in a way that doesn't dismay their administrators (cf. Rees et al. (2012))), and issues closer to the viscera, like being able to rely on another hospital to capably perform a donor nephrectomy, and have the kidney shipped reliably to arrive predictably (with tracking in place in case of plane-schedule mishaps).

And a final, important task of market design is to find, and build, the social support needed for the market to thrive.

Looking back, kidney exchange has accomplished a lot, but not nearly enough. The number of people waiting for a kidney transplant is growing, despite the growth of exchange. But there is room for kidney exchange to continue to grow and to increase the availability of transplants further, by designing international kidney exchanges, by starting chains with deceased donor kidneys, and by introducing other market design innovations that have yet to be explored or even conceived.

References

- Atila Abdulkadiroğlu and Tayfun Sönmez. House allocation with existing tenants. *Journal of Economic Theory*, 88(2):233–260, 1999.
- Atila Abdulkadiroglu and Tayfun Sönmez. School choice: A mechanism design approach. *The American Economic Review*, 93(3):729–747, 2003.
- Atila Abdulkadiroğlu, Parag A Pathak, and Alvin E Roth. The new york city high school match. *American Economic Review*, 95(2):364–367, 2005a.
- Atila Abdulkadiroğlu, Parag A Pathak, Alvin E Roth, and Tayfun Sönmez. The boston public school match. *American Economic Review*, pages 368–371, 2005b.
- Atila Abdulkadiroğlu, Parag A Pathak, and Alvin E Roth. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the nyc high school match. *American Economic Review*, 99(5):1954–78, 2009.

- David J Abraham, Avrim Blum, and Tuomas Sandholm. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 295–304, 2007.
- Nikhil Agarwal, Itai Ashlagi, Eduardo Azevedo, Clayton Featherstone, and Ömer Karaduman. What matters for the productivity of kidney exchange? In *AEA Papers and Proceedings*, volume 108, pages 334–40, 2018.
- Nikhil Agarwal, Itai Ashlagi, Eduardo Azevedo, Clayton R Featherstone, and Ömer Karaduman. Market failure in kidney exchange. *American Economic Review*, 109(11):4026–70, 2019.
- Nikhil Agarwal, Itai Ashlagi, Michael A Rees, Paulo J Somaini, and Daniel C Waldinger. Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys. Forthcoming in *Econometrica*, 2020a.
- Nikhil Agarwal, Charles Hodgson, and Paulo Somaini. Choices and outcomes in assignment mechanisms: The allocation of deceased donor kidneys. 2020b.
- Mohammad Akbarpour, Julien Combe, Yinghua He, Victor Hiller, Robert Shimer, and Olivier Tercieux. Unpaired kidney exchange: Overcoming double coincidence of wants without money. 2020a. NBER Working Paper No. 27765, September.
- Mohammad Akbarpour, Julien Combe, Yinghua He, Victor Hiller, Robert Shimer, and Olivier Tercieux. Unpaired kidney exchange: Overcoming double coincidence of wants without money. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 465–466, 2020b.
- Mohammad Akbarpour, Farshad Fatemi, and Negar Matorian. The iranian market for kidneys. In *Working paper, Stanford GSB*, 2020c.
- Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. Thickness and information in dynamic matching markets. *Journal of Political Economy*, 128(3):000–000, 2020d.
- Filipe Alvelos, Xenia Klimentova, and Ana Viana. Maximizing the expected number of transplants in kidney exchange programs with branch-and-price. *Annals of Operations Research*, 272(1-2):429–444, 2019.
- Frederike Ambagtsheer, Bernadette Haase-Kromwijk, Frank JMF Dor, Greg Moorlock, Franco Citterio, Thierry Berney, and Emma K Massey. Global kidney exchange: Opportunity or exploitation? an elpat/esot appraisal. *Transplant International*, 33(9):989–998, 2020. URL <https://onlinelibrary.wiley.com/doi/full/10.1111/tri.13630>.
- Sandro Ambuehl, Muriel Niederle, and Alvin E Roth. More money, more problems? can high pay be coercive and repugnant? *American Economic Review*, 105(5):357–60, 2015.

- Ross Anderson, Itai Ashlagi, David Gamarnik, Michael Rees, Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Kidney exchange and the alliance for paired donation: Operations research changes the way kidneys are transplanted. *Interfaces*, 45(1):26–42, 2015a.
- Ross Anderson, Itai Ashlagi, David Gamarnik, and Alvin E Roth. Finding long chains in kidney exchange using the traveling salesman problem. *Proceedings of the National Academy of Sciences*, 112(3):663–668, 2015b.
- Ross Anderson, Itai Ashlagi, David Gamarnik, and Yash Kanoria. Efficient dynamic barter exchange. *Operations Research*, 65(6):1446–1459, 2017.
- Nick Arnosti and Peng Shi. How (not) to allocate affordable housing. In *AEA Papers and Proceedings*, volume 109, pages 204–08, 2019.
- I. Ashlagi, D. Gamarnik, Michael A. Rees, and A. E. Roth. The need for (long) NEAD chains. Working paper, 2012a.
- Itai Ashlagi and Süleyman Kerimov. Scrip systems with minimal availability. In *WINE*, page 335, 2019.
- Itai Ashlagi and Alvin E Roth. New challenges in multihospital kidney exchange. *American Economic Review*, 102(3):354–59, 2012.
- Itai Ashlagi and Alvin E Roth. Free riding and participation in large scale, multi-hospital kidney exchange. *Theoretical Economics*, 9(3):817–863, 2014.
- Itai Ashlagi, Duncan S Gilchrist, Alvin E Roth, and Michael A Rees. Nonsimultaneous chains and dominos in kidney-paired donation—revisited. *American Journal of transplantation*, 11(5):984–994, 2011a.
- Itai Ashlagi, Duncan S Gilchrist, Alvin E Roth, and Michael A Rees. Nead chains in transplantation. *American Journal of transplantation*, 11(12):2780–2781, 2011b.
- Itai Ashlagi, David Gamarnik, Michael A Rees, and Alvin E Roth. The need for (long) chains in kidney exchange. Technical report, National Bureau of Economic Research, 2012b.
- Itai Ashlagi, Patrick Jaillet, and Vahideh H Manshadi. Kidney exchange in dynamic sparse heterogeneous pools. In *In ACM Conference on Electronic Commerce (EC)*. Citeseer, 2013.
- Itai Ashlagi, Adam Bingaman, Maximilien Burq, Vahideh Manshadi, David Gamarnik, Cathi Murphey, Alvin E Roth, Marc L Melcher, and Michael A Rees. Effect of match-run frequencies on the number of transplants and waiting times in kidney exchange. *American Journal of Transplantation*, 18(5):1177–1186, 2018.
- Itai Ashlagi, Maximilien Burq, Chinmoy Dutta, Patrick Jaillet, Amin Saberi, and Chris Sholley. Edge weighted online windowed matching. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 729–742, 2019a.

- Itai Ashlagi, Maximilien Burq, Patrick Jaillet, and Vahideh Manshadi. On matching and thickness in heterogeneous dynamic markets. *Operations Research*, 67(4):927–949, 2019b.
- Itai Ashlagi, Afshin Nikzad, and Philipp Strack. Matching in dynamic imbalanced markets. *Available at SSRN 3251632*, 2019c.
- Itai Ashlagi, Jacob Leshno, Pengyu Qian, and Amin Saberi. Queue lengths as constantly adapting prices: Allocative efficiency under random dynamics. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 317–318, 2020a.
- Itai Ashlagi, Faidra Monachou, and Afshin Nikzad. Optimal dynamic allocation: Simplicity through information design. *Available at SSRN*, 2020b.
- Sepehr Assadi, Sanjeev Khanna, and Yang Li. The stochastic matching problem with (very) few queries. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–19, 2019.
- Baris Ata, John Friedewald, and A Cem Randa. Structural estimation of kidney transplant candidates’ quality of life scores: Improving national kidney allocation policy under endogenous patient choice and geographical sharing. *Structural Estimation of Kidney Transplant Candidates’ Quality of Life Scores: Improving National Kidney Allocation Policy Under Endogenous Patient Choice and Geographical Sharing (January 25, 2020)*, 2020.
- Lawrence M Ausubel and Thayer Morrill. Sequential kidney exchange. *American Economic Journal: Microeconomics*, 6(3):265–85, 2014.
- Mariagiovanna Baccara, SangMok Lee, and Leeat Yariv. Optimal dynamic matching. *Theoretical Economics*, 15(3):1221–1278, 2020.
- Hoda Bidkhorji, John P Dickerson, Duncan C McElfresh, and Ke Ren. Kidney exchange with inhomogeneous edge existence uncertainty. *arXiv preprint arXiv:2007.03191*, 2020.
- AA Bingaman, M Kapturczak, I Ashlagi, and CL Murphey. 400 kidney paired donor transplants at a single center; the methodist san antonio experience. In *American Journal of Transplantation*, volume 18, pages 420–421, 2018.
- AW Bingaman, FH Wright Jr, M Kapturczak, L Shen, S Vick, and CL Murphey. Single-center kidney paired donation: the methodist san antonio experience. *American Journal of Transplantation*, 12(8):2125–2132, 2012.
- Péter Biro, David F Manlove, and Romeo Rizzi. Maximum weight cycle packing in directed graphs, with application to kidney exchange programs. *Discrete Mathematics, Algorithms and Applications*, 1(04):499–517, 2009.
- Péter Biró, Bernadette Haase-Kromwijk, Tommy Andersson, Eyjólfur Ingi Ásgeirsson, Tatiana Baltsová, Ioannis Boletis, Catarina Bolotinha, Gregor Bond, Georg Böhmig, Lisa Burnapp, et al. Building kidney exchange programmes in europe—an overview of exchange practice and activities. *Transplantation*, 103(7):1514, 2019a.

- Péter Biró, Walter Kern, Dömötör Pálvölgyi, and Daniel Paulusma. Generalized matching games for international kidney exchange. 2019b.
- Péter Biró, Joris van de Klundert, David Manlove, William Pettersson, Tommy Andersson, Lisa Burnapp, Pavel Chromy, Pablo Delgado, Piotr Dworczak, Bernadette Haase, et al. Modelling and optimisation in european kidney exchange programmes. *European Journal of Operational Research*, 2019c.
- Jose H Blanchet, Martin I Reiman, Viragh Shah, and Lawrence M Wein. Asymptotically optimal control of a centralized dynamic matching market with general utilities. *arXiv preprint arXiv:2002.03205*, 2020.
- Francis Bloch and David Cantala. Dynamic assignment of objects to queuing agents. *American Economic Journal: Microeconomics*, 9(1):88–122, 2017.
- A. Blum and Y. Mansour. Kidney exchange and endless paths: on the optimal use of an altruistic donor. Working paper, 2020.
- Avrim Blum, Anupam Gupta, Ariel Procaccia, and Ankit Sharma. Harnessing the power of two crossmatches. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 123–140, 2013.
- Avrim Blum, John P Dickerson, Nika Haghtalab, Ariel D Procaccia, Tuomas Sandholm, and Ankit Sharma. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. *Operations research*, 2020.
- Georg A Böhmig, Jiří Fronek, Antonij Slavcev, Gottfried F Fischer, Gabriela Berlakovich, and Ondrej Viklicky. Czech-austrian kidney paired donation: first european cross-border living donor kidney exchange. *Transplant International*, 30(6):638–639, 2017.
- Danielle Bozek, Ty. B. Dunn, Christian S. Kuhr, Christopher L. Marsh, Jeffrey Rogers, Susan E. Rees, Laura Basagoitia, Robert J. Brunner, Alvin E. Roth, Obi Ekwenna, David E. Fumo, Kimberly D. Krawiec, Kopke Jonathan E., Jorge Sindhvani, Puneetand Ortiz, Miguel Tan, Siegfredo R. Paloyo, Jeffrey D. Punch, and Michael A. Rees. Complete chain of the first global kidney exchange transplant 4 and 3-yr follow-up. *European Urology Focus*, Forthcoming, 2018.
- M Bray, W Wang, PX-K Song, AB Leichtman, MA Rees, VB Ashby, R Eikstadt, A Goulding, and JD Kalbfleisch. Planning for uncertainty and fallbacks can increase the number of transplants in a kidney-paired donation program. *American Journal of Transplantation*, 15(10):2636–2645, 2015.
- Fauzia K Butt, HA Gritsch, P Schulam, GM Danovitch, A Wilkinson, J Del Pizzo, S Kapur, D Serur, S Katznelson, S Busque, et al. Asynchronous, out-of-sequence, transcontinental chain kidney transplantation: A novel concept. *American Journal of Transplantation*, 9(9):2180–2185, 2009.

- Linda Cantwell, Claudia Woodroffe, Rhonda Holdsworth, and Paolo Ferrari. Four years of experience with the australian kidney paired donation programme. *Nephrology*, 20(3):124–131, 2015.
- Margarida Carvalho, Xenia Klimentova, Kristiaan Glorie, Ana Viana, and Miguel Constantino. Robust models for the kidney exchange problem. *INFORMS Journal on Computing*, 2020.
- Edward H Cole, Peter Nickerson, Patricia Campbell, Kathy Yetzer, Nick Lahaie, Jeffery Zaltzman, and John S Gill. The canadian kidney paired donation program: a national program to increase living donor transplantation. *Transplantation*, 99(5):985–990, 2015.
- Miguel Constantino, Xenia Klimentova, Ana Viana, and Abdur Rais. New insights on integer-programming models for the kidney exchange problem. *European Journal of Operational Research*, 231(1):57–68, 2013.
- N Cowan, HA Gritsch, Nima Nassiri, Joseph Sinacore, and J Veale. Broken chains and renegeing: A review of 1748 kidney paired donation transplants. *American Journal of Transplantation*, 17(9):2451–2457, 2017.
- Francis L Delmonico and Nancy L Ascher. Opposition to irresponsible global kidney exchange. *American Journal of Transplantation*, 17(10):2745–2746, 2017.
- John Dickerson, Duncan McElfresh, Ke Ren, et al. Kidney exchange with inhomogeneous edge existence uncertainty. In *Conference on Uncertainty in Artificial Intelligence*, pages 161–170. PMLR, 2020.
- John P. Dickerson, Ariel D. Procaccia, and Tuomas Sandholm. Dynamic Matching via Weighted Myopia with Application to Kidney Exchange. *Proc of the 6th AAAI Conference on Artificial Intelligence*, pages 1340–1346, 2012a.
- John P Dickerson, Ariel D Procaccia, and Tuomas Sandholm. Dynamic matching via weighted myopia with application to kidney exchange. In *AAAI*, 2012b.
- John P Dickerson, Ariel D Procaccia, and Tuomas Sandholm. Optimizing kidney exchange with transplant chains: Theory and reality. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 711–718. International Foundation for Autonomous Agents and Multiagent Systems, 2012c.
- John P Dickerson, Ariel D Procaccia, and Tuomas Sandholm. Failure-aware kidney exchange. *Management Science*, 65(4):1768–1791, 2019.
- John Paul Dickerson and Tuomas Sandholm. Balancing efficiency and fairness in dynamic kidney exchange. In *AAAI Workshop: Modern Artificial Intelligence for Health Analytics*, 2014.

- Esther Dufo. Richard t. ely lecture: The economist as plumber. *American Economic Review*, 107(5):1–26, 2017.
- Jack Edmonds. Paths, Trees, and Flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.
- Julio J Elías, Nicola Lacetera, and Mario Macis. Paying for kidneys? a randomized survey and choice experiment. *American Economic Review*, 109(8):2855–88, 2019.
- Council of Europe. Statement on the global kidney exchange concept. In *European Committee on Organ Transplantation (CD-P-TO), European Directorate for the Quality of Medicines and HealthCare.*, 2018. URL https://search.coe.int/cm/Pages/result_details.aspx?ObjectID=09000016808aee9f.
- Paolo Ferrari, Willem Weimar, Rachel J Johnson, Wai H Lim, and Kathryn J Tinckam. Kidney paired donation: principles, protocols and programs. *Nephrology Dialysis Transplantation*, 30(8):1276–1285, 2015.
- DE Fumo, V Kapoor, LJ Reece, SM Stepkowski, JE Kopke, SE Rees, C Smith, AE Roth, AB Leichtman, and MA Rees. Historical matching strategies in kidney paired donation: the 7-year evolution of a web-based virtual matching system. *American Journal of Transplantation*, 15(10):2646–2654, 2015.
- Lucrezia Furian, Cristina Cornelio, Cristina Silvestre, Flavia Neri, Francesca Rossi, Paolo Rigotti, Emanuele Cozzi, and Antonio Nicolò. Deceased donor–initiated chains: First report of a successful deliberate case and its ethical implications. *Transplantation*, 103(10):2196–2200, 2019.
- Lucrezia Furian, Antonio Nicolò, Caterina Di Bella, Massimo Cardillo, Emanuele Cozzi, and Paolo Rigotti. Kidney exchange strategies: New aspects and applications with a focus on deceased-donor-initiated chains. *Transplant International*, 2020.
- David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- Sommer Elizabeth Gentry, DL Segev, Mary Simmerling, and RA Montgomery. Expanding kidney paired donation through participation by compatible pairs. *American Journal of Transplantation*, 7(10):2361–2370, 2007.
- Kristiaan Glorie, Albert Wagelmans, and Joris van de Klundert. Iterative branch-and-price for hierarchical multi-criteria kidney exchange. Technical report, 2012.
- Kristiaan Glorie, Bernadette Haase-Kromwijk, Joris van de Klundert, Albert Wagelmans, and Willem Weimar. Allocation and matching in kidney exchange programs. *Transplant International*, 27(4):333–343, 2014.

- Kristiaan Glorie, Margarida Carvalho, Miguel Constantino, Paul Bouman, and Ana Viana. Robust models for the kidney exchange problem. Technical report, Working paper, 2015.
- C Hajaj, J.P. Dickerson, A. Hassidim, and T. Sandholm. Strategy-proof and efficient kidney exchange using a credit mechanism. In *AAAI*, 2015.
- Ruthanne L Hanto, William Reitsma, and Francis L Delmonico. The development of a successful multiregional kidney paired donation program. *Transplantation*, 86(12):1744–1748, 2008.
- Christine Hauser and Hugo Hopenhayn. Trading favors: Optimal exchange and forgiveness. *Collegio Carlo Alberto Carlo Alberto Notebooks*, 88, 2008.
- Philip J Held, F McCormick, A Ojo, and John P Roberts. A cost-benefit analysis of government compensation of kidney donors. *American Journal of Transplantation*, 16(3):877–885, 2016.
- Kris Johnson, David Simchi-Levi, and Peng Sun. Analyzing scrip systems. *Operations Research*, 62(3):524–534, 2014.
- Rachel J Johnson, Joanne E Allen, Susan V Fuggle, J Andrew Bradley, Chris Rudge, et al. Early experience of paired living kidney donation in the united kingdom. *Transplantation*, 86(12):1672–1677, 2008.
- JA Gestel Kal-van, BJ Haase-Kromwijk, FH Claas, W Weimar, et al. Eight years of outcomes of the dutch living donor kidney exchange program. *Clinical transplants*, pages 287–290, 2011.
- Yash Kanoria and Pengyu Qian. Near optimal control of a ride-hailing platform via mirror backpressure. *arXiv*, pages arXiv–1903, 2019.
- Ian A Kash, Eric J Friedman, and Joseph Y Halpern. An equilibrium analysis of scrip systems. *ACM Transactions on Economics and Computation*, 3(3):13, 2015.
- KM Keizer, M De Klerk, BJJM Haase-Kromwijk, and W Weimar. The dutch algorithm for allocation in living donor kidney exchange. In *Transplantation proceedings*, volume 37, pages 589–591. Elsevier, 2005.
- Suleyman Kerimov, Itai Ashlagi, and Itai Gurvich. Dynamic matching in networks: on the optimality of static priorities. 2020.
- Xenia Klimentova, João Pedro Pedroso, and Ana Viana. Maximising expectation of the number of transplants in kidney exchange programmes. *Computers & Operations Research*, 73:1–11, 2016.
- Xenia Klimentova, Ana Viana, Joao Pedro Pedroso, and Nicolau Santos. Fairness models for multi-agent kidney exchange programmes. *Omega*, page 102333, 2020.

- Kimberly D. Krawiec and Michael A. Rees. Reverse transplant tourism. *77 Law & Contemporary Problems*, pages 145–173, 2014.
- Kimberly D Krawiec, Wenhao Liu, and Marc L Melcher. Contract development in a matching market: the case of kidney exchange. *Law & Contemp. Probs.*, 80:11, 2017.
- Vivek B Kute, Himanshu V Patel, Umesh T Varyani, Pankaj R Shah, Pranjal R Modi, Veena R Shah, Sayyed J Rizvi, Bipin C Pal, Priya S Shah, Pavan S Wakhare, et al. Six end-stage renal disease patients benefited from first non-simultaneous single center 6-way kidney exchange transplantation in india. *World journal of nephrology*, 5(6):531, 2016.
- Vivek B Kute, Narayan Prasad, Pankaj R Shah, and Pranjal R Modi. Kidney exchange transplantation current status, an update and future perspectives. *World journal of transplantation*, 8(3):52, 2018.
- Stephen Leider and Alvin E Roth. Kidneys for sale: Who disapproves, and why? *American Journal of transplantation*, 10(5):1221–1227, 2010.
- Jacob Leshno. Dynamic Matching in Overloaded Waiting Lists. Working paper, 2014.
- Zhuoshu Li, Kelsey Lieberman, William Macke, Sofia Carrillo, Chien-Ju Ho, Jason Wellen, and Sanmay Das. Incorporating compatible pairs in kidney exchange: A dynamic weighted matching model. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 349–367, 2019.
- John DC Little and Stephen C Graves. Little’s law. In *Building intuition*, pages 81–100. Springer, 2008.
- Thaminda Liyanage, Toshiharu Ninomiya, Vivekanand Jha, Bruce Neal, Halle Marie Patrice, Ikechi Okpechi, Ming-hui Zhao, Jicheng Lv, Amit X Garg, John Knight, et al. Worldwide access to treatment for end-stage kidney disease: a systematic review. *The Lancet*, 385(9981):1975–1982, 2015.
- Shafi Malik and Edward Cole. Foundations and principles of the canadian living donor paired exchange program. *Canadian journal of kidney health and disease*, 1(1):6, 2014.
- David F Manlove and Gregg O’malley. Paired and altruistic kidney donation in the uk: Algorithms and experimentation. *Journal of Experimental Algorithmics (JEA)*, 19:1–21, 2015.
- Charles F Manski, Anat R Tambur, and Michael Gmeiner. Predicting kidney transplant outcomes with partial knowledge of hla mismatch. *Proceedings of the National Academy of Sciences*, 116(41):20339–20345, 2019.
- Ignazio R Marino, Alvin E Roth, Michael A Rees, and Cataldo Doria. Open dialogue between professionals with different opinions builds the best policy. *American Journal of Transplantation*, 17(10):2749–2749, 2017.

- Duncan C McElfresh, Hoda Bidkhor, and John P Dickerson. Scalable robust kidney exchange. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1077–1084, 2019.
- Duncan C. McElfresh, Michael Curry, Tuomas Sandholm, and John P. Dickerson. Improving policy-constrained kidney exchange via pre-screening. In *Conference on Neural Information Processing Systems*, 2020.
- Marc L Melcher, John P Roberts, Alan B Leichtman, Alvin E Roth, and Michael A Rees. Utilization of deceased donor kidneys to initiate living donor chains. *American Journal of Transplantation*, 16(5):1367–1370, 2016.
- Panayotis Mertikopoulos, Heinrich H Nax, and Bary Pradelski. Quick or cheap? breaking points in dynamic markets. *Breaking points in dynamic markets (December 30, 2019)*, 2019.
- Paul Milgrom. Simplified mechanisms with an application to sponsored-search auctions. *Games and Economic Behavior*, 70(1):62–70, 2010.
- Radu-Stefan Mincu, Péter Biró, Márton Gyetvai, Alexandru Popa, and Utkarsh Verma. Ip solutions for international kidney exchange programmes. *Central European Journal of Operations Research*, pages 1–21, 2020.
- Francesca Minerva, Julian Savulescu, and Peter Singer. The ethics of the global kidney exchange programme. *The Lancet*, 394(10210):1775–1778, 2019.
- Michael Mitzenmacher. The power of two choices in randomized load balancing. *IEEE Transactions on Parallel and Distributed Systems*, 12(10):1094–1104, 2001.
- Markus Möbius. Trading favors. *Unpublished paper.[283]*, 2001.
- Tiago Monteiro, Xenia Klimentova, João Pedro Pedrosa, and Ana Viana. A comparison of matching algorithms for kidney exchange programs addressing waiting time. *Central European Journal of Operations Research*, pages 1–14, 2020.
- Felix T Rapaport. The case for a living emotionally related international kidney donor exchange registry. In *Transplantation proceedings*, volume 18(3), pages 5–9, 1986.
- Michael A Rees, Jonathan E Kopke, Ronald P Pelletier, Dorry L Segev, Matthew E Rutter, Alfredo J Fabrega, Jeffrey Rogers, Oleh G Pankewycz, Janet Hiller, Alvin E Roth, et al. A nonsimultaneous, extended, altruistic-donor chain. *New England Journal of Medicine*, 360(11):1096–1101, 2009.
- Michael A Rees, Mark A Schnitzler, EY Zavala, James A Cutler, Alvin E Roth, Frank D Irwin, Stephen W Crawford, and Alan B Leichtman. Call to develop a standard acquisition charge model for kidney paired donation. *American Journal of transplantation*, 12(6): 1392–1397, 2012.

- Michael A Rees, Ty B Dunn, Christian S Kuhr, Christopher L Marsh, Jeffrey Rogers, Susan E Rees, Alejandra Cicero, Laurie J Reece, Alvin E Roth, Obi Ekwenna, et al. Kidney exchange to overcome financial barriers to kidney transplantation. *American Journal of Transplantation*, 17(3):782–790, 2017a.
- Michael A Rees, Siegfredo R Paloyo, Alvin E Roth, Kimberly D Krawiec, Obi Ekwenna, Christopher L Marsh, AJ Wenig, and Ty B Dunn. Global kidney exchange: Financially incompatible pairs are not transplantable compatible pairs. *American journal of transplantation: official journal of the American Society of Transplantation and the American Society of Transplant Surgeons*, 17(10):2743, 2017b.
- Alvin E Roth. Incentive compatibility in a market with indivisible goods. *Economics letters*, 9(2):127–132, 1982.
- Alvin E Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of political Economy*, 92(6):991–1016, 1984.
- Alvin E Roth, editor. *The Shapley value: essays in honor of Lloyd S. Shapley*. Cambridge University Press, 1988.
- Alvin E Roth. A natural experiment in the organization of entry-level labor markets: regional markets for new physicians and surgeons in the united kingdom. *The American economic review*, pages 415–440, 1991.
- Alvin E Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, 2002.
- Alvin E Roth. Repugnance as a constraint on markets. *The Journal of Economic Perspectives*, 21(3):37–58, 2007.
- Alvin E Roth. What have we learned from market design? *The Economic Journal*, 118(527):285–310, 2008.
- Alvin E Roth. *Who gets what—and why: The new economics of matchmaking and market design*. Houghton Mifflin Harcourt, 2015.
- Alvin E Roth. Marketplaces, markets, and market design. *American Economic Review*, 108(7):1609–58, 2018.
- Alvin E Roth and Elliott Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American economic review*, 89(4):748–780, 1999.
- Alvin E Roth and Stephanie W Wang. Popular repugnance contrasts with legal bans on controversial markets. *Proceedings of the National Academy of Sciences*, 117(33):19792–19798, 2020.

- Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Kidney exchange. *The Quarterly journal of economics*, 119(2):457–488, 2004.
- Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Pairwise kidney exchange. *Journal of Economic theory*, 125(2):151–188, 2005a.
- Alvin E Roth, Tayfun Sönmez, et al. A kidney exchange clearinghouse in new england. *American Economic Review*, 95(2):376–380, 2005b.
- Alvin E Roth, Tayfun Sönmez, M Utku Ünver, Francis L Delmonico, and Susan L Saidman. Utilizing list exchange and nondirected donation through ‘chain’paired kidney donations. *American Journal of transplantation*, 6(11):2694–2705, 2006.
- Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. Efficient kidney exchange: Coincidence of wants in markets with compatibility-based preferences. *American Economic Review*, 97(3):828–851, 2007.
- Alvin E Roth, T Sonmez, and M Utku Unver. Kidney paired donation with compatible pairs. *Am J Transplant*, 8(2):463, 2008.
- Alvin E Roth, Kimberly D Krawiec, Siegfredo Paloyo, Obi Ekwenna, Christopher L Marsh, Alexandra J Wenig, Ty B Dunn, and Michael A Rees. People should not be banned from transplantation only because of their country of origin. *American Journal of Transplantation*, (10):2747–2748, 2017.
- Alvin E. Roth, Ignazio R. Marino, Obi Ekwenna, Ty B. Dunn, Siegfredo R. Paloyo, Miguel Tan, Ricardo Correa-Rotter, Christian S. Kuhr, Christopher L. Marsh, Jorge Ortiz, Giuliano Testa, Puneet Sindhvani, Dorry L. Segev, Jeffrey Rogers, Jeffrey D. Punch, Rachel C. Forbes, Michael A. Zimmerman, Matthew J. Ellis, Aparna Rege, Laura Basagoitia, Kimberly D. Krawiec, and Michael A. Rees. Global kidney exchange should expand wisely. *Transplant International*, 33(9):985–988, 2020. doi: 10.1111/tri.13656. URL <https://onlinelibrary.wiley.com/doi/full/10.1111/tri.13656>.
- Susan L Saidman, Alvin E Roth, Tayfun Sönmez, M Utku Ünver, and Francis L Delmonico. Increasing the opportunity of live kidney donation by matching for two-and three-way exchanges. *Transplantation*, 81(5):773–782, 2006.
- Maurizio Salvadori and Aris Tsalouchos. Current protocols and outcomes of aboincompatible kidney transplantation. *World Journal of Transplantation*, 10(7):191–205, 2020.
- Nicolau Santos, Paolo Tubertini, Ana Viana, and João Pedro Pedroso. Kidney exchange simulation and optimization. *Journal of the Operational Research Society*, 68(12):1521–1532, 2017.

- Dorry L Segev, Sommer E Gentry, Daniel S Warren, Brigitte Reeb, and Robert A Montgomery. Kidney paired donation and optimizing the use of live donor organs. *Jama*, 293(15):1883–1890, 2005.
- Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.
- Lloyd S Shapley. A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317, 1953.
- Peter Singer, Francesca Minerva, and Julian Savulescu. The global kidney exchange programme—authors’ reply. *The Lancet*, 395(10235):1485–1486, 2020.
- Tayfun Sönmez and M. Utku Ünver. Market design for kidney exchange. *The Handbook of Market Design*, 2013a.
- Tayfun Sönmez and M Utku Ünver. Market design for kidney exchange. *The Handbook of Market Design*, pages 93–137, 2013b.
- Tayfun Sönmez and M Utku Ünver. Altruistically unbalanced kidney exchange. *Journal of Economic Theory*, 152:105–129, 2014.
- Tayfun Sönmez, M Utku Ünver, and Özgür Yılmaz. How (not) to integrate blood subtyping technology to kidney exchange. *Journal of Economic Theory*, 176:193–231, 2018.
- Tayfun Sönmez, M Utku Ünver, and M Bumin Yenmez. Incentivized kidney exchange. *American Economic Review*, 110(7):2198–2224, 2020.
- X. Su and S. A. Zenios. Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model. *Management Science*, 52(11):1647–1660, 2006.
- United States Renal Data System. *2016 USRDS annual data report: Epidemiology of kidney disease in the United States*, volume 2. National Institutes of Health, National Institute of Diabetes and Digestive and Kidney Diseases, Bethesda, MD, 2016.
- M Utku Ünver. Dynamic kidney exchange. *The Review of Economic Studies*, 77(1):372–414, 2010.
- Anji E Wall, Jeffrey L Veale, and Marc L Melcher. Advanced donation programs and deceased donor-initiated chains—2 innovations in kidney paired donation. *Transplantation*, 101(12):2818–2824, 2017.
- C Bradley Wallis, Kannan P Samy, Alvin E Roth, and Michael A Rees. Kidney paired donation, 2011.

- W Wang, MA Rees, AB Leichtman, P X-K Song, M Bray, VB Ashby, T Shearon, A White-man, and JD Kalbfleisch. Deceased donors as non-directed donors in kidney paired donation. *American Journal of Transplantation*, 2020.
- Wen Wang, Mathieu Bray, Peter X-K Song, and John D Kalbfleisch. A look-ahead strategy for non-directed donors in kidney paired donation. *Statistics in biosciences*, 9(2):453–469, 2017.
- Robert Wilson. Architecture of power markets. *Econometrica*, 70(4):1299–1340, 2002.
- Robert A Wolfe, Valarie B Ashby, Edgar L Milford, Akinlolu O Ojo, Robert E Ettenger, Lawrence YC Agodoa, Philip J Held, and Friedrich K Port. Comparison of mortality in all patients on dialysis, patients on dialysis awaiting transplantation, and recipients of a first cadaveric transplant. *New England Journal of Medicine*, 341(23):1725–1730, 1999.
- Stefanos A Zenios, E Steve Woodle, and Lainie Friedman Ross. Primum non nocere: Avoiding harm to vulnerable wait list candidates in an indirect kidney exchange. *Transplantation*, 72(4):648–654, 2001.

A APPENDIX: Simulation models

Simulation models are important for testing different KE operational policies. We begin with simulating static pools and then discuss how to generate the dynamics in a KE program.

A.1 Generating compatibility graphs

The first ingredient is to be able to generate an exchange pool represented by a compatibility graph. A general framework for creating a pool is: (i) generate incompatible pairs and NDDs (nodes), and then (ii) generate compatibilities between patients and potential donors in the pool (directed edges).

With clinical data, generating the graph is straightforward. First sample pairs or NDDs from the data, and then create edges using virtual crossmatches using donors’ antigens and patients’ antibodies.

Without these data the following is useful for simulating whether a donor and a patient are compatible. Each patient is associated with a (virtual) PRA which is a number between 0 and 1. To replace the virtual crossmatch test between a donor and patient, we do a *PRA test*: draw a random number for the donor and if that number is higher than the PRA level, they passed the test successfully. The donor and the patient are *virtually compatible* if they are ABO compatible and the PRA test succeeded.

Sampling from pool compositions. A few papers provide statistics of pool compositions, from which one can sample pairs directly (and then generate edges after testing for virtual compatibility). Such a table is given in Glorie et al. (2012) for the Dutch data. Tables 1 and 2 provide such statistics of the historical APKD (2010-2019) and NKR (2010-2014) pools. To generate a pair first sample from a patient-donor ABO and then the PRA of the patient according to the marginal distribution.

Patient-Donor ABO	Percentage of pairs	Marginal frequencies (PRA intervals)						
		0-1	1-10	10-50	50-80	80-95	95-99	99-100
AB-AB	0.2	0	0	0	50	0	25	25
AB-B	0.4	0	0	0	16.7	16.7	0	66.7
AB-A	0.7	0	8.3	0	8.3	50	0	33.3
AB-O	0.6	10	0	20	10	0	20	40
B-AB	0.9	37.5	6.2	18.8	6.2	12.5	0	18.8
B-B	2.4	0	4.9	12.2	12.2	31.7	9.8	29.3
B-A	5.8	46.5	8.1	13.1	9.1	12.1	1	10.1
B-O	4.2	9.9	1.4	4.2	16.9	19.7	15.5	32.4
A-AB	1	41.2	5.9	5.9	11.8	17.6	0	17.6
A-B	3.6	30.6	9.7	6.5	14.5	9.7	1.6	27.4
A-A	9.7	4.2	1.8	16.9	19.3	18.1	10.8	28.9
A-O	8.8	12.7	4.7	9.3	19.3	15.3	18	20.7
O-AB	2.3	46.2	10.3	23.1	5.1	12.8	0	2.6
O-B	9.2	47.1	10.8	14	7.6	8.3	4.5	7.6
O-A	29.4	49.9	10	12.8	8.8	6.4	3.6	8.6
O-O	20.7	4.5	2.8	13.9	17.3	23.9	16.2	21.3

Table 1: APKD pool composition (2010-2019). PRA percentages are conditional on patient-donor ABO types.

Note: It is important to note that hospitals that participate at the APKD and NKR have conducted many internal exchanges (see Section 4.2). Without this behavior, these pools would likely contain a larger fraction of easy-to-match pairs.

Two (new) simulation models. We describe here two simulation models that are refinements of previously suggested models by Saidman et al. (2006) and Segev et al. (2005), who build on ideas from Zenios et al. (2001).

Model 1. *Generate incompatible (unrelated) pairs (p, d) as follows.*

1. *A patient p is generated by drawing a blood type and a PRA level using the general population distribution given in Table 3.*
2. *$k \geq 1$ incompatible donors are generated sequentially. For each donor, draw a blood*

Patient-Donor ABO	Percentage of pairs	Marginal frequencies (PRA intervals)						
		0-1	1-10	10-50	50-80	80-95	95-99	99-100
AB-AB	0.5	0	0	0	11.1	44.4	11.1	33.3
AB-B	0.3	20	0	0	20	20	40	0
AB-A	1	10.5	5.3	5.3	21.1	10.5	26.3	21.1
AB-O	0.7	7.7	0	7.7	30.8	15.4	7.7	30.8
B-AB	1.3	50	4.2	8.3	8.3	16.7	4.2	8.3
B-B	2.5	4.3	2.1	6.4	12.8	23.4	19.1	31.9
B-A	7.8	59.9	2.7	13.6	9.5	4.8	4.8	4.8
B-O	3.8	14.1	2.8	16.9	19.7	18.3	15.5	12.7
A-AB	1.9	50	5.6	13.9	8.3	11.1	5.6	5.6
A-B	5.3	52.5	2	10.1	14.1	11.1	2	8.1
A-A	9.7	7.7	1.6	15.3	21.9	18	14.2	21.3
A-O	8.5	6.9	1.9	13.1	15.6	22.5	13.1	26.9
O-AB	1.9	47.2	13.9	13.9	13.9	2.8	2.8	5.6
O-B	10	57.7	6.3	13.2	5.8	4.2	4.8	7.9
O-A	26.2	48.8	6.3	15.7	10.2	5.7	5.7	7.7
O-O	18.7	7.1	2	14.2	21.9	23.1	15.7	16

Table 2: NKR pool composition (2010-2014). PRA percentages are conditional on patient-donor ABO types.

type independently according to the ABO distribution in Table 3.⁸⁹ If a generated donor is compatible with patient (ABO and PRA test) , go back to step 1.

3. *If all donors are incompatible, pick one of the donors at random, call it d and (p, d) joins the pool.⁹⁰*

An adjustment to model 1 following (Saidman et al., 2006): Instead of step 2, draw for each donor also whether it is a spouse or a different unrelated donor (also non-biological) as well as the gender, according to the Table 3. In the case, in which the patient is a woman with PRA X , and the donor is her spouse, the PRA test is successful with probability $100 - 0.75(100 - X)$.⁹¹

The second model accounts for the relation between the patient and the donor (Zenios et al., 2001).

Model 2. *Generate incompatible (unrelated) pair (p, d) as follows.*

⁸⁹A few types of donors are omitted from the statistics (e.g., half siblings, life partner, twin.)

⁹⁰Sometimes a patient joins the pool with more than one incompatible donor. But often incompatible intended donors do not join the pool due to some failure in the work-up process (Zenios et al., 2001; Segev et al., 2005).

⁹¹This is because women are more likely to reject their husband's kidney than a random kidney, because of antibodies developed at childbirth to paternal antigens.

Characteristic	Frequency (%)
PRA range (average)	
0-1 (0)	62.56
1-50 (30)	16.48
50-80 (65)	6.9
80-95 (87)	5.06
95-99 (97)	2.74
99-100 (99.5)	6.26
ABO	
O	49.31
A	32.05
B	15.02
AB	3.62
Patient sex	
Male	61.53
Female	38.47
Donor relation	
Parent	10.56
Child	18.46
Sibling	23.02
Spouse	15.9
Other unrelated	32.06

Table 3: Patient and living donor distributions from OPTN data (2010-2019).

1. A patient p is generated by drawing two independent ABO proteins according to Table 4, which together determine p 's blood type.⁹²
2. The PRA of p is generated according to Table 3.
3. Generate 1 incompatible donor d and draw the relation of d to p according to Table 3. Generate the ABO of the donor according to Bayes Rule in case the donor is a relative and otherwise draw it independently. If d is compatible with the patient (ABO and PRA test), go back to step 1, and otherwise (p, d) joins the pool.

Comments and comparisons to previous models:

- One can adapt model 2 also to the ethnicity of the patient and donor based on Table 4.

⁹²For example, the chance that a random patient has blood type O is the chance she has no A or B proteins (with probability 70.22×70.22) and the chance that she has blood type A, means that she at least one A protein and no B). Using this distribution one can generate a random patient. Then, using Bayes rule, one can generate the ABO of potential relatives (including a spouse).

Population	O	A	B
General	70.22%	19.98%	9.98%
White	66.85%	25.43%	7.74%
Black	70.79%	16.02%	13.27%
Asian	62.52%	17.5%	20.10%
Hispanic	76.93%	16.96%	6.08%

Table 4: ABO protein distributions calculated from the ABO distributions in the OPTN data (O represents not A and not B).

- Both Saidman et al. (2006) and Segev et al. (2005) create compatibility graphs that are more densely connected than the ones generated by contemporary clinical data (Ashlagi and Roth, 2012). The ABO composition generated is also different than those at the APKD, NKR and Dutch pool (Glorie et al., 2012). One reason is due to modeling all patients with PRA 80-100 to have a simulated PRA of 90.⁹³ Instead Table 3 provides a much more refined PRA distribution.⁹⁴
- Model 1, even with the new PRA table (and its adjustment to capture spouses), still generates a pool with a highly skewed ABO distribution for values of $k \leq 4$. When k takes values between 5-10 the generated pool fits the data better (but this may vary across KE programs).
- Model 2 generates a compatibility graph that fits better contemporary multi-hospital exchange pools such as the APKD and NKR. We find that this modified model matches well the pool composition and also different connectivity measures (such as number of two-way cycles, one-way matches, size of maximum matching).

A.2 Simulating dynamics and operations of a pool

To simulate dynamics of a KE pool one needs to simulate (i) arrivals of pairs and NDDs (ii) the matching process including possible frictions, (iii) departures not due to matches, and (iv) the times during which donors are available.

Arrivals can naturally be simulated using a Poisson process, in which upon arrival of a new node one can draw the type of the node (what kind of pair or NDD). Departures are naturally simulated using estimated hazard rate (although more should be done to understand the departure process). The reasons for departure may matter for evaluating a policy.⁹⁵

⁹³For example the model by Saidman et al. (2006) generates a larger fraction of under-demanded pairs than in NKR and APKD pools (we observed 8% more O-A pairs and 6% more O-B pairs than in historical datasets from NKR and APKD).

⁹⁴See also a similar refined distribution in (Glorie et al., 2012).

⁹⁵Reasons can include too sick to be transplanted, or transplanted elsewhere (from a deceased or live donor).

In some situations, donors have a window, in which they know they will be available to donate. Incorporating these windows may have different implications for paired donors and non-directed donors.⁹⁶

There are a variety of possible match failures. Common to all KE programs are frictions due to crossmatches. Much more should be done to predict crossmatch failures. All studies we are aware of assumed that crossmatches fail independently for each donor and patient that are virtually compatible.⁹⁷ Typically patients with higher PRA have a higher likelihood to have a positive crossmatch with a donor. Statistics are given in numerous papers for a variety of KE programs (e.g., Ashlagi et al. (2011b); Wang et al. (2017); Agarwal et al. (2019); Biró et al. (2019a); Dickerson et al. (2019)).

Multi-hospital KE programs in the US typically make many match offers and a non-negligible percentage of offers are refused by hospitals. Statistics can be found in Ashlagi et al. (2011b); Wang et al. (2017); Agarwal et al. (2019), also assuming independent failures.⁹⁸

Other frictions include shipping delays of blood samples (when there is no central blood lab) and other delays due to communication. Statistics can be found in Agarwal et al. (2019); Ashlagi et al. (2011b).

Finally, for a detailed description of a simulation process that replicates common operations of a KE program (as described in Section 2.3, such the NKR, see Appendix C in Agarwal et al. (2019)).

⁹⁶Cathi Murphey: “I divide the donors into 2 categories- the ones that can go anytime and the other group that have a window. The ones that have a window are usually students or teachers or people who have jobs that have high and low peaks. Like landscapers. So they usually donate during spring break or the summer or over Christmas time. The other category are people that have jobs that they can take off of work at anytime or they don’t work.”

⁹⁷Whether using clinical or simulated data.

⁹⁸But it is very much possible that some hospitals refuse more offers than others.