Replication Archive for "An Empirical Framework for Sequential Assignment: The Allocation of Deceased Donor Kidneys"

Nikhil Agarwal, Itai Ashlagi, Michael Rees, Paulo Somaini, Daniel Waldinger.

Obtaining Original Data Files and Analysis Code

The data reported here have been supplied by UNOS as the contractor for the Organ Procurement and Transplantation Network (OPTN). The interpretation and reporting of these data are the responsibility of the author(s) and in no way should be seen as an official policy of or interpretation by the OPTN or the U.S. Government.

We will retain copies of the data until permitted by our Data Use Agreement with the Organ Procurement and Transplantation Network (OPTN). Further, we plan to send OPTN a copy of our replication archive if and when we are required to destroy our dataset. Researchers interested in using our dataset should directly contact OPTN to obtain permission: `https://optn.transplant.hrsa.gov/data/request-data/`

We are happy to provide copies of our data to researchers with permission and a data use agreement with the OPTN.

The analysis code is available on the following GitHub repository in addition to the compressed file accompanying this replication archive:

`https://bitbucket.org/nikhilagarwal12/kidneys-ecma`.

Data Description

Our data on patients, donors, transplants, and offers are based on information submitted to the Organ Procurement and Transplantation Network (OPTN) by its members. The main dataset on the waitlist is the Potential Transplant Recipient (PTR) dataset. It contains sequences of offers made to patients on the deceased donor kidney waitlist, their decisions, and their reasons for refusal. Detailed information on patient characteristics, donor characteristics, and transplant outcomes come from the Standard Transplantation Analysis and Research (STAR) dataset. UNOS also provided supplemental information for this study, including the ordering of distinct match runs conducted for the same deceased donor; the transplant centers of donors and patients in our dataset; and dates of birth for pediatric candidates, who joined the waitlist before turning 18 years of age.

The data contain unique identifiers that allow us to link the offer and acceptance data to patient and donor characteristics. Each deceased donor has a unique identifier. Similarly, each patient registration generates a unique patient waitlist identifier. Because patients may

move to different transplant centers or be registered in multiple centers simultaneously, some individual patients have multiple waitlist id's. Where appropriate, we de-duplicate offers so that each patient can receive at most one offer from each donor. The patient history file also contains a unique patient record identifier corresponding to a particular state of the patient on the waitlist, including the patient's CPRA, activity status, and pre-set screening criteria. Each offer in the PTR dataset contains the identifiers for the donor, the patient registration, and the patient history record that were used in the match run.

The PTR dataset contains offers made to patients on the deceased donor kidney waitlist that were not automatically rejected based on pre-specified criteria. Information include identifiers for the donor, patient, and patient history record that generated the offer; the order in which the offers were made; each patient's acceptance decision; and if the offer was not accepted, a reason for rejecting. Each offer record also contains certain characteristics of the match, including the number of tissue type mismatches.

The STAR dataset contains separate files on deceased donor characteristics, patient characteristics and transplant outcomes, and patient histories. The patient and donor characteristics from these tables are used to estimate our models of acceptance behavior, positive crossmatch probabilities, and patient departure rates. We also use these characteristics to replicate the mechanism and determine each patient's compatibility with and priority score for each deceased donor in our sample.

Sample Selection

This section explains the selection of patients, donors, and offers used in our structural model. We consider patients who were registered in NYRT and actively waiting for a deceased donor kidney between January 1st, 2010 and December 31st, 2013. Our donor sample includes all U.S. deceased kidney donors whose organs were allocated according to the standard mechanism. Our offer sample includes valid offers from all deceased donors to NYRT patients during this period recorded in the PTR data, as well as offers that were not made because of pre-specified screening criteria. Because the PTR data do not record offers that were not made due to screening criteria, we constructed our sample in two steps. First, we constructed the sample of patients and donors using the STAR and PTR datasets. Second, we ran a simulation of the kidney allocation mechanism to determine our sample of offers, including those "screened out" by pre-set criteria (and therefore not in the PTR data). The remainder of this section discusses the patient and donor samples, and the next section discusses our simulation of the allocation mechanism.

Because NYRT patients may be offered donors from across the U.S., our procedure first constructs a nationwide sample of deceased donors and patients. We also determine each donor's priority score cutoff at this stage. We then restrict the sample to NYRT and omit certain donors and patients who received non-standard treatment in the mechanism.

Our U.S. sample of deceased kidney donors comes from the intersection of donor identifiers in the PTR and STAR deceased donor files, excluding donors allocated using non-standard rules. Patients in our sample were active on the deceased donor kidney waitlist after 2010 and were not jointly registered for a pancreas transplant. Patient registration date and activity status are determined from the patient history file. We also exclude patients who departed the waitlist for reasons which indicate that they did not ultimately need a transplant. We exclude patients who were transplanted in another country, whose condition improved, or who could no longer be contacted. These departure reasons are recorded in the STAR patient and transplant outcome dataset. We then determine which offers in the PTR data were valid and could have been accepted by and transplanted into the patient; patients' acceptance decisions; and the resulting priority score cutoffs in each match run.

We first exclude PTR offers that are not valid. In certain cases, patients are bypassed when a donor is allocated to a specific recipient outside of the standard allocation rules. This can occur if the donor is an armed service member; if the donor specified a particular recipient (directed donation); if there is a medical emergency or expedited placement attempt; or if organ sharing among DSAs generates a "payback" in which one DSA allocates a kidney from another DSA as though it had been recovered in its own service area. There are also cases in which a patient is offered a tissue type incompatible donor, or a donor that did not meet the patient's pre-specified screening criteria. We identify these cases using a refusal reason code provided in the PTR data. In some cases, there is also text specifying specific circumstances justifying a rejection, which we parse to identify invalid offers in cases where the refusal code does not provide a specific reason. Finally, some offers are refused due to technological constraints if the patient needs a specific organ laterality or requires multiple simultaneous organ transplants. We do not consider these cases to be genuine refusals, and omit them from the offer dataset.

Next, we created an algorithm to de-duplicate offers and acceptances within and across match runs, and to determine the true priority score cutoff for each donor in each match run. For some donors, multiple match runs are conducted, and these match runs can include offers to overlapping sets of patients. A specific kidney (e.g. the left kidney) may also be accepted in multiple match runs. Finally, a patient can have multiple offers recorded from the same donor, even in the same match run. Our algorithm assumes that later match runs take precedence over earlier ones (using the match run numbers provided by OPTN), and that the last observed match run in which an organ is placed takes that organ out of circulation for subsequent match runs.

We then implement the sample restrictions for NYRT. We consider all patients who were registered in NYRT and had active status sometime between January 1st, 2010 and December 31st, 2013. Table 1 describes our specific sample restrictions for NYRT patients. The primary decision we made was to omit patients who received a transplant through non-standard allocation rules. This includes cases of medical urgency, an expedited placement attempt, a multi-organ transplant, or a military or directed donation. 362 patients were excluded

because they received deceased donor kidney transplants for these reasons – less than 4 percent of the NYRT patients who were actively waiting during our sample period – leaving 9,623 patients in our final NYRT sample.

Table 2 illustrates how our PTR data filters determines the sample of donors available to NYRT. We consider donors when a valid offer was made to at least one NYRT patient, where a valid offer is determined using the filters described above. Our final NYRT donor sample contains 5,642 donors.

The offers and patient acceptance decisions in the PTR data determine the priority score cutoff in each match run for each donor's available organs.

Replicating the Mechanism, Offer Dataset

Knowledge of the mechanism allows us to determine the set of offers that were declined through pre-set screening criteria, as well as the waiting time required for a particular patient to have access to a particular donor. These are essential for correctly modeling patient acceptance behavior and transplant opportunities under the current and counterfactual mechanisms. We wrote computer code to replicate the standard deceased donor kidney allocation rules in place between January 1st, 2010 and December 31st, 2013.

For each deceased donor and match run, the algorithm begins with all concurrent patient waitlist history records. It first determines which patients are incompatible with the donor due to their blood type and unacceptable human leukocyte antigen (HLA) antigens (if any). We use blood type and HLA equivalence tables followed by the OPTN, as well as the donor's HLA antigens and the current unacceptable antigens listed by each patient. Next, we check whether the donor met each patient's screening criteria. Finally, we determine the priority score of each patient given their CPRA, waiting time, geography, age, and number of HLA and DR mismatches with the donor. Given the priority score, we can calculate whether the patient was above the priority score cutoff for the donor. We can also determine the amount of additional waiting time (which may be infinite) after which the patient's priority score would exceed the donor's cutoff.

From the simulation, we obtain a set of offers predicted by our simulation of the mechanism. These are pairs of donors and patients where the patient met the priority score cutoff and was blood and tissue type compatible with the donor. Some of these offers met the patient's screening criteria, while others did not. Those that did should appear in the PTR data. This provides a check on the performance of our mechanism code. Table 4 tabulates offers appearing in our filtered PTR data and those predicted by simulation. The vast majority of offers in the PTR data (91.8%) were predicted by our simulation, and conversely, the vast majority of offers predicted by our simulation (93.4%) appear in the PTR data.

To estimate the patient acceptance model, we take as our offer sample the union of the PTR offer dataset and the set of offers that the simulation predicts were filtered due to the patient's

screening criteria but which would otherwise have appeared in the PTR data. In a final step, we de-duplicate offers at the patient level, since a patient registered at multiple centers will occasionally receive multiple offers from the same donor. Table 3 describes how we arrive at our final sample of offers from the output of the simulation. The final offer sample contains 2,713,043 offers. Offers that were screened out are interpreted as rejections since the patient deemed the donor's characteristics unacceptable.

To calculate patient value functions, we store all compatible patient and donor pairs, including patients who did not meet the donor's priority score cutoff.

Imputing Missing Donor DR Antigens

A donor's HLA antigens are needed to determine tissue type compatibility with transplant candidates as well as kidney points, which are in turn essential for replicating the mechanism. A limitation of our data is that we only observe a donor's DR antigens if one of their kidneys or pancreas was transplanted into a patient. In this case, they appear in the KIDPAN (patient/transplant) dataset. If no organs were placed, a donor's antigen information is recorded in the deceased donor file for kidney/pancreas donors. The deceased donor file lists the donor's HLA antigens at the A and B loci, but not at the DR locus.

We either obtain or impute a donor's missing DR antigens from two sources. First, some deceased donors had a liver, lung, or part of their intestine transplanted even though their kidneys and pancreas were not transplanted. The equivalent transplant files for these additional organs are part of the STAR dataset, and we take the donor's DR antigens directly from those files.

Second, for deceased donors who had no organs transplanted, we use the reported number of DR mismatches in the PTR offer dataset to impute the donor's DR antigens. Because we observe all patients' HLA antigens, the number of DR mismatches between a donor and patient is informative about the donor's antigens. For example, if a donor-patient pair has zero DR mismatches, the patient's tissue type limits the donor's antigens to a few possibilities.[1] A two DR mismatch pair also restricts the donor's DR antigens, though less so than a zero mismatch. Since deceased donors whose organs are not transplanted are usually offered to many patients, we can combine information across all offers to make an educated guess of the donor's DR antigens.

We use the following imputation algorithm. For each donor without DR antigen information, we take all of the donor's offers in the PTR data. Based on these offers, the recorded number of DR mismatches, and the patient's DR antigens and listed unacceptable antigens,

---

[1] In the zero DR mismatch case, the donor may not share the patient's exact DR antigens because of HLA equivalences. Some distinct HLA proteins are equivalent in the sense that a patient with one DR antigen may desensitize it to several DR antigens. UNOS publishes HLA equivalence tables for measuring HLA mismatches, and a separate table for equivalent unacceptable antigens. Furthermore, even ignoring equivalences, a zero DR mismatch donor could be homozygous at the DR locus.

we calculate a score for each DR antigen that the donor might have. We dock one point from a DR antigen's score for each PTR offer it contradicts in the following cases:

- The offer has zero DR mismatches, and the antigen is not equivalent to one of the patient's DR antigens

- The offer has two DR mismatches, and the antigen is equivalent to one of the patient's DR antigens

- The antigen was listed as unacceptable by the patient

For each donor, we take the two DR antigens with the highest scores. Ties are broken in favor of the antigens that appear most frequently among donors for whom DR antigens are recorded.

Table 1: Patient Sample Restrictions

	Number of Patients Registered
Kidney Candidates Registered in NYRT Between 2010 and 2013	14499
Excluding candidates who were not interested in a transplant	13950
Excluding inactive candidates	9985
Excluding candidates who received non-standard allocations	9623

Notes: candidates who were not interested in a transplant include patients who departed the waitlist because they refused transplantation, received a transplant in another country, could not be contacted, or had an improved condition. Inactive candidates are patients who registered on the waitlist but never changed their status to "active," and therefore never received kidney offers, during the sample period. Candidates receiving non-standard allocations include placements from military or directed donations, expedited placement attempts, and medical emergencies.

Table 2: Donor Sample Restrictions

	Number of Donors
All donors matched with NYRT patients	8181
Excluding non-standard allocations	8180
Excluding cases where patient did not meet donor's cutoff	6167
Excluding rejections based on pre-set criteria	6020
Excluding bypasses	5907
Excluding inactive and incompatible patients	5688
Excluding cases where surgeon or hospital was unavailable	5686
Excluding cases with no recorded patient response	5642

Notes: The donor sample is constructed using the PTR data. Non-standard allocations include placements from military or directed donations, expedited placement attempts, and medical emergencies. A patient met the donor's cutoff if one of the donor's available organs was accepted by that patient or a patient with lower priority. A donor was rejected by pre-set criteria if a patient who could have accepted the donor's organs automatically rejected the donor based on pre-set criteria. A bypass occurs when a patient would have received an offer under standard allocation rules, but did not because there was a non-standard allocation. Inactive patients were not actively waiting for a transplant; incompatible patients or could not accept the donor's organs due to immune system or size/laterality incompatibility. There was no recorded response if there was no yes/no answer or refusal code in the PTR data.

Table 3: Offer Sample Restrictions

	Number of Offers
All possible patient-donor pairs	15643712
Excluding blood and tissue-type incompatible pairs	8241934
Excluding offers below donor's cutoff	4333425
Excluding offers to patients not actively waiting	3158941
Excluding offers predicted to appear in PTR that did not	2724489
De-duplicating across match runs and patient listings	2713043

Notes: All possible patient-donor pairs include all donors in the NYRT sample and all NYRT patients registered at the time of the donor's match date. An offer is predicted to appear in PTR if the patient is medically compatible, actively waiting, above the donor's cutoff, and does not automatically reject the donor based on pre-set screening criteria.

Table 4: Fit of Mechanism Code

		Predicted by Simulation		
		No	Yes	Total
In PTR Data	No	1,476,556	103,046	1,579,602
	Yes	129,304	1,450,035	1,579,339
	Total	1,605,860	1,553,081	3,158,941