

EQUILIBRIUM ALLOCATIONS UNDER ALTERNATIVE WAITLIST DESIGNS: EVIDENCE FROM DECEASED DONOR KIDNEYS

NIKHIL AGARWAL

Department of Economics, MIT and NBER

ITAI ASHLAGI

Management Science and Engineering, Stanford University

MICHAEL A. REES

Department of Urology, University of Toledo Medical Center

PAULO SOMAINI

Stanford Graduate School of Business and NBER

DANIEL WALDINGER

Department of Economics, New York University

Waitlists are often used to ration scarce resources, but the trade-offs in designing these mechanisms depend on agents' preferences. We study equilibrium allocations under alternative designs for the deceased donor kidney waitlist. We model the decision to accept an organ or wait for a preferable one as an optimal stopping problem and estimate preferences using administrative data from the New York City area. Our estimates show that while some kidney types are desirable for all patients, there is substantial match-specific heterogeneity in values. We then develop methods to evaluate alternative mechanisms, comparing their effects on patient welfare to an equivalent change in donor supply. Past reforms to the kidney waitlist primarily resulted in redistribution, with similar welfare and organ discard rates to the benchmark first-come, first-served mechanism. These mechanisms and other commonly studied theoretical benchmarks remain far from optimal. We design a mechanism that increases patient welfare by the equivalent of an 18.2% increase in donor supply.

KEYWORDS: Market design, organ allocation, dynamic mechanisms, dynamic discrete choice.

Nikhil Agarwal: agarwaln@mit.edu
Itai Ashlagi: iashlagi@stanford.edu
Michael A. Rees: michael.rees2@utoledo.edu
Paulo Somaini: soma@stanford.edu
Daniel Waldinger: dw120@nyu.edu

Subsumes “An Empirical Framework for Sequential Assignment: The Allocation of Deceased Donor Kidneys.” We are grateful to UNOS for facilitating access to the data. We thank Peter Arcidiacono, Lanier Benkard, Ben Handel, Liran Einav, Jose Fernandez, Gautam Gowrisankaran, Jacob Leshno, Anna Mikusheva, Michael Ostrovsky, Ariel Pakes, Parag Pathak, Peter Reiss, Andy Skrzypacz, Gabriel Weintraub, and Mike Whinston for helpful discussions. Lingjun Xia provided outstanding research assistance. The authors acknowledge support from the NSF (SES-1729090, SES-1254768), the NIH (R21-DK113626), the Sloan Foundation (FG-2019-11484), and the MIT SHASS Dean's Research Fund.

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.

1. INTRODUCTION

A KIDNEY TRANSPLANT IMPROVES THE QUALITY and length of a patient's life while saving hundreds of thousands of dollars in dialysis costs (Wolfe et al. (1999), Irwin, Bonagura, Crawford, and Foote (2012), Held, McCormick, Ojo, and Roberts (2016)). Deceased donor kidneys are allocated using a waitlist instead of traditional price-based market mechanisms. In 2018, fewer than one-fifth of the almost 100,000 patients on the kidney waitlist received a transplant.¹ At the same time, approximately one-fifth of kidneys recovered from deceased donors were discarded. Given the value and scarcity of these resources, it is essential to design mechanisms that are efficient, minimize waste, and achieve equitable outcomes. Similar design goals are important in the allocation of other deceased donor organs, public housing, long-term care, child care, and child adoption.

Previous research and policy guidance on waitlist design is based on restrictive assumptions about agents' preferences or behavior. The theoretical literature has found that even qualitative trade-offs are sensitive to whether objects are vertically or horizontally differentiated (cf. Su and Zenios (2004), Bloch and Cantala (2017), Leshno (2019), Agarwal, Ashlagi, Somaini, and Waldinger (2018)). Absent clear recommendations from theory, many organ allocation agencies use simulations to predict the effects of alternative allocation rules. The simulations, including those used to design organ allocation rules, do not allow agents' decisions to respond to changes in the system.²

This paper empirically evaluates equilibrium allocations under waitlist mechanisms that sequentially assign objects to forward-looking agents. We make several methodological and empirical contributions. First, we develop a method for estimating payoffs from various assignments and apply it to data from the deceased donor kidney waitlist in the New York City area. This step is based on an optimal stopping problem faced by a patient when they are offered a kidney. Second, we define a notion of steady state equilibrium that permits computationally tractable counterfactual analysis of a broad class of mechanisms. Finally, we use these techniques to compare alternative mechanisms in terms of efficiency, equity, and organ waste while accounting for changes in agent behavior. We evaluate equilibrium assignments under mechanisms used in practice, theoretical benchmarks, and welfare-maximizing mechanisms.

Our empirical application uses rich administrative data on the deceased donor kidney allocation system between 2010 and 2013 from the New York Organ Donor Network (NYRT), which serves the greater New York City area. The allocation mechanism offers each organ to patients on the waitlist in decreasing order of priority points, which are determined by a coarse set of patient and donor characteristics as well as the patient's waiting time. An offer can either be accepted or rejected. A patient who refuses an offer is not penalized and remains on the waitlist. We refer to the decision-maker as the patient although surgeons advise patients and often act on their behalf.³

The data suggest that although patients face extreme scarcity, they strategically wait for offers of desirable organs. While 1400 patients join the waitlist each year, fewer than 200 deceased donor kidneys are recovered in NYRT. These donors vary widely in quality;

¹Source: <https://optn.transplant.hrsa.gov/data/view-data-reports/national-data/>.

²For example, Kidney Pancreas Simulated Acceptance Module (KPSAM), used in the United States by the Organ Procurement and Transplantation Network kidney allocation committee to evaluate proposed mechanisms, assumes that acceptance decisions on the kidney waitlist do not depend on the mechanism. Similar methods are used by the organ allocation agencies in the U.K., Scandinavia, and France.

³Our data do not directly identify the decision maker. This approach is reasonable if each surgeon acts in the best interest of each patient when they are offered an organ.

some are accepted immediately while others are rejected by every patient and discarded. The chance of being offered desirable organs increases with waiting time, giving patients an incentive to reject offers and wait for a better kidney. Indeed, Agarwal et al. (2018) documented that patients who are less likely to receive offers in the future and, therefore, have lower option values of waiting, are more likely to accept an offer.

Motivated by these facts, we model an agent's decision to accept an offer as an optimal stopping problem. An agent accepts the current offer if the value of the object is higher than the expected value of continuing to wait. The distribution of future offers depends on the mechanism and the strategies of the other agents on the list. Our empirical strategy combines acceptance probabilities with detailed knowledge of the mechanism to recover the value of a transplant as a function of a rich set of patient and donor characteristics. The technique adapts methods for inverting conditional choice probabilities (Hotz and Miller (1993), Arcidiacono and Miller (2011), Arcidiacono, Bayer, Blevins, and Ellickson (2016)) to suit dynamic assignment mechanisms.

The estimated payoffs show that while some organs are systematically more desirable than others, there is substantial match-specific heterogeneity in values. For instance, organs from younger donors are preferred by all patients, but younger patients place a higher value on such organs. This and other sources of match-specific heterogeneity, such as immunological similarity, create scope for redesigning the allocation mechanism to improve match quality by incorporating detailed patient and donor characteristics into the priority system.

Next, we develop methods for computing equilibrium assignments under counterfactual mechanisms. This task requires us to solve two technical issues. First, we need to formulate a tractable notion of equilibrium. Computing counterfactuals is challenging because it involves solving a dynamic game with many players. To make progress, we define a notion of a steady state equilibrium in the spirit of previous approaches (e.g., Hopenhayn (1992), Krusell and Smith (1998), Weintraub, Benkard, and van Roy (2008), Fershtman and Pakes (2012)) and develop an algorithm for computing such an equilibrium. Second, we must ensure that assignments under counterfactuals of interest are identified. In dynamic models, counterfactual results may not be invariant to normalizing the payoff of a single action in all states because it restricts payoffs across states (see Aguirregabiria and Suzuki (2014), Kalouptsi, Scott, and Souza-Rodrigues (2020)). We formally show that normalizing the payoff from never receiving an assignment is appropriate for the mechanism design counterfactuals we consider if the value of declining all offers remains fixed. We argue that this assumption is reasonable in our context.

We then compare equilibrium assignments under alternative mechanisms, focusing on the trade-off between allocative efficiency and organ waste. These objectives have been emphasized in the theoretical literature (see Su and Zenios (2004), Bloch and Cantala (2017)) and in policy discussions (see OPTN (2011)). To summarize a change in the mechanism's expected welfare effect on a given patient, we compute the equivalent change in deceased donor supply (arrival rates) under the existing mechanism. We then aggregate these equivalent changes in donor supply across patients as a summary of the welfare effects. This aggregation implicitly makes interpersonal comparisons of utility across patients.

Previously used mechanisms and commonly studied theoretical benchmarks can yield either high average welfare or low discards, but not both. Reforms in 2014 prioritized healthier patients for the most desirable organs. We show that these changes primarily resulted in redistribution from older patients to younger patients, with little improvement in average patient welfare. In fact, both the pre- and post-2014 mechanisms yield welfare

and organ discard rates within 2.5% of the benchmark first-come, first-served mechanism (FCFS) studied in Bloch and Cantala (2017). FCFS induces agents to be selective because the expected quality of offers increases as time passes. Last-come, first-served (LCFS), which is theoretically studied in Su and Zenios (2004), makes agents less selective because they expect to receive lower-quality offers in the future if they decline an offer today. This mechanism dramatically reduces organ discard rates (25%), but at the cost of lowering welfare by 50% due to poor match quality.

In contrast, we use our estimates to identify mechanisms that significantly increase patient welfare and reduce discards. An approximation to the welfare-maximizing mechanism that tailors offer rates based on agent and object characteristics increases the average patient's welfare by the equivalent of an 18.2% increase in donor supply. This mechanism also reduces organ discard rates by more than 7%, and as a result, equilibrium queue lengths and waiting times are shorter than under the pre-2014 mechanism. A drawback is that some patient types may be harmed. To address this concern, we design a mechanism that increases welfare by an equivalent of 12.2% and reduces discards by 4.0% while ensuring that no patient type is substantially worse off.

Simultaneously considering heterogeneity and dynamic incentives is crucial for identifying these gains. The latter channel is ignored by the simulation models currently used by policymakers to predict the effects of various allocation systems (see SRTR (2015)). A comparison that ignores dynamic incentives and holds choice probabilities fixed predicts much smaller differences in match quality and discard rates between mechanisms, understating the effects of changing the mechanism. Moreover, a mechanism that naively prioritizes patients based on predicted transplant value only marginally improves upon the pre-2014 mechanism, increasing welfare and reducing discard rates by 3.8 and 0.2%, respectively. Taken together, our results highlight our empirical methods' ability to improve dynamic assignment mechanisms.

Related Literature

Zenios (2004) surveyed previous research on organ allocation mechanisms. Unlike our model, the empirical research in this area assumes that acceptance decisions do not depend on the mechanism. Most of the research within economics is on the much smaller living donor kidney exchange market (see, e.g., Roth, Sonmez, and Ünver (2004), Roth, Sönmez, and Ünver (2007), Agarwal, Ashlagi, Azevedo, Featherstone, and Karaduman (2019)).

The most closely related paper to ours is Zhang (2010), which also uses a dynamic model to argue that observational learning best fits acceptance patterns on the deceased donor kidney waitlist. The result is based on the fact that, conditional on a set of donor covariates included in the model, patients lower on the list are more likely to refuse an organ if patients higher on the list have refused it. Our approach abstracts away from learning but allows donor characteristics observed by agents but not included in the model to capture this correlation in acceptance behavior. We do this in accordance with our current institutional understanding and to simplify counterfactual analysis.⁴

⁴Zhang (2010) used data from 2002. The donor information available to patients was much more detailed and standardized during our sample period (2010–2013), reducing the scope for observational learning. Both observational learning and unobserved heterogeneity in organ quality result in correlated choices across agents. We leave developing methods for incorporating observational learning into our framework for future work.

The methods in this paper contribute to a growing literature on the empirical analysis of centralized assignment systems (see Agarwal and Somaini (2020), for a survey). Existing work has focused on static assignment mechanisms,⁵ where the theory is comparatively well-developed and empirical evidence suggests that differences among well-coordinated systems are small (Abdulkadiroglu, Agarwal, and Pathak (2017)). In contrast, theory suggests that the most desirable dynamic offer system depends on the nature of preferences.

Overview

Section 2 describes the institutions and data, and presents descriptive evidence. Sections 3 and 4 model the optimal stopping problem faced by each agent and detail our estimation methods. Section 5 describes our parameter estimates. Section 6 defines a steady state equilibrium, summarizes results on existence, and presents our approach to welfare comparisons. Section 7 compares predicted outcomes under alternative mechanisms. Section 8 concludes.

2. BACKGROUND, DATA, AND DESCRIPTIVE EVIDENCE

2.1. *Basics of Kidney Transplantation*

The two main treatments for kidney failure are dialysis and transplantation. Transplantation is usually preferable, improving health outcomes while saving an estimated \$270,000 or more over the life of a transplanted patient (Wolfe et al. (1999), Irwin et al. (2012), Held et al. (2016), USRDS (2018)). Medicare, which provides near-universal coverage for the disease, spent \$35.4 billion treating kidney failure in 2016, accounting for 1% of the federal budget (Chapter 9, USRDS (2018)).

The primary constraints on transplant feasibility are blood- and tissue-type compatibility. A patient's immune system will attack an incompatible organ, resulting in graft failure. Following transplantation, medications allow transplant physicians to limit new immune responses to foreign protein types, but preexisting immune responses lead to immediate loss of the transplanted organ if not avoided.⁶ For this reason, patients' transplantation possibilities differ based on their immune systems. Some patients react to most organs, even from donors with the same blood type, giving them fewer transplantation options. A patient's immune sensitization is commonly measured by Calculated Panel Reactive Antibodies (CPRA), which is the percentage of donors in a representative sample with whom they are tissue-type incompatible.

Transplant benefits vary considerably with patient and donor characteristics, even conditional on biological compatibility. There are several factors, including donor age and gender, health prior to death, circumstances of death, measures of kidney function, and

⁵To our knowledge, the only exceptions are Waldinger (2020) and Reeling and Verdier (2020). The optimal stopping rule we study differs from these models. It is most similar to Liu, Yang, and Wan (2019), which studies decisions in a ride-sharing platform.

⁶The immune system tags foreign objects (antigens) with antigen-specific antibodies so that white blood cells (leukocytes) and complement (a blood protein) can defend against them. Each donor has blood-type antigens and specific types of human leukocyte antigen proteins that are relevant for kidney transplantation. Some patients have preexisting antibodies to a subset of these antigens. A transplant recipient's immune system will immediately attack the donor kidney and reject the organ if the recipient has a preexisting antibody to any one of the donor antigens, making the donor incompatible (Danovitch (2009)). Following transplantation of a compatible kidney, the immune system will attack any donor antigen not present in the patient if such an attack is not attenuated by immunosuppressive medications.

kidney anatomy that are considered important determinants of kidney quality. In some cases, the donor may have an infectious disease that the patient is at risk of contracting if they proceed with the transplant. Size and weight match, tissue-protein similarity between the patient and donor, and many other factors influence specific organs' medical benefits for specific patients. We refer the reader to Danovitch (2009) for further details about kidney biology.

2.2. The Allocation of Deceased Donor Kidneys

The deceased donor kidney allocation process is regulated by U.S. federal law. It is designed and administered by the Organ Procurement and Transplantation Network (OPTN), which coordinates patient registration, organ recovery, and the allocation process using a system called UNet.

A patient with kidney disease may register on the waitlist when their kidney function falls below 20%. During registration, UNet collects detailed information about the patient's immunological profile, health conditions, and factors used to compute priority. Patients may also set criteria that automatically exclude offers from organs based on donor characteristics such as age, health, and kidney function.

The assignment of a potential donor's organs begins after brain death is declared or cardiac death is imminent and necessary consent has been obtained. Detailed information about the donor's medical history and the organ's condition is entered into UNet. This system screens out patients known to be incompatible with or to have excluded the donor; calculates the priority order for the remaining patients; transmits information about the donor to the transplant centers; and records their responses. In order to maintain organ viability, many potential recipients may be contacted at once. This process can take place while the donor is on life-support, before the organs are extracted. If several patients accept the donor's organ(s), the highest-priority patient(s) will receive the transplant(s).

When contacted, a potential recipient is first asked to indicate initial interest if further information about the organ is pending. Once a kidney has been recovered from the donor, those interested receive any new information discovered in the interim and must make a final decision without delay, usually within an hour. This decision, which is typically made in consultation with a transplant surgeon, commits the patient to accepting the organ pending the results of blood tests conducted using samples from the donor and patient.

The pre- and post-2014 priority systems are designed with both equity and efficiency concerns in mind. The priority system used in NYRT prior to 2014 first offers kidneys to patients with a perfect tissue-type match, then to patients from the local area in which the organs were recovered, then regionally, and finally nationally. Within each group, patients are prioritized based on tissue-type similarity, whether or not the patient is pediatric, patient CPRA, and waiting time (see OPTN (2014), for details). The new kidney allocation rules implemented in 2014 give the healthiest patients highest priority for the most desirable donors, increase priority for extremely hard-to-match (high-CPRA) patients, and reduce emphasis on waiting time. Israni et al. (2014) discusses this system and the rationale for the changes. Neither system uses past decisions to calculate priority and, unlike assignment systems for some other organs (e.g., livers), priority does not depend on patient urgency.

2.3. Data and Descriptive Evidence

This study uses data from the Organ Procurement and Transplantation Network (OPTN). The OPTN data system includes data on all donors, wait-listed candidates, and transplant recipients in the U.S., submitted by the members of the OPTN. The Health Resources and Services Administration (HRSA), U.S. Department of Health and Human Services provides oversight to the activities of the OPTN contractor.

The datasets we use are sourced from UNet and contain detailed patient and donor characteristics, the offers made, and accept/reject decisions. Our replication archive describes the data in greater detail. For tractability, we restrict attention to data on offers made to patients registered in NYRT between January 1, 2010, and December 31, 2013.⁷ NYRT has the most patients of any donor service area in the U.S. that used the standard allocation rules prior to 2014.

We now describe our sample of patients and donors and document choice patterns. A striking feature of the waitlist is that even though there is extreme scarcity, some donors are rejected by many patients. Choices suggest that large differences in donor quality, combined with substantial priority for waiting time, incentivize patients to reject low-quality donors and wait for more attractive offers.

Patients and Donors

Table I describes our patient sample. A total of 9623 patients were registered with NYRT and actively waiting at some point between 2010 and 2013. Panel A shows the state of the waitlist on January 1 of 2010 and 2013. Our dataset includes indicators of patient health (e.g., body mass index, age, total serum albumin) and medical history (e.g., diabetes, years on dialysis). The average waiting time of patients on the list is a little over 2 years and has increased over time. The average CPRA is about 12%, which indicates there is more than a one-in-ten chance that a patient is tissue-type incompatible with a randomly chosen donor. This mean is driven by a subset of patients who are incompatible with most donors.

Patients in NYRT face extreme scarcity, which results in long waiting times and many patients dying while waiting. Each year, fewer than 200 donors are recovered from the NYRT area (Table II, Panel A), while about 1400 patients join the NYRT waitlist (Table I, Panel B). This scarcity results in the average transplanted patient waiting for over 3 years before receiving a deceased donor. Patients who do not receive a transplant most often leave the list because they either die or become “too sick to transplant” (Table I, Panel B). A smaller number are fortunate enough to receive a living donor transplant, which is more likely for younger patients and often occurs within the first year on the waitlist. Finally, some patients leave for other or unknown reasons including a move outside the NYRT area.

Despite this scarcity, some kidneys are refused by many patients. Table II shows that across donors, the mean number of biologically compatible offers that meet the preset screening criteria is over 400, but the median is much lower, at 26. This skewed distribution arises because undesirable kidneys are rejected by many patients, while desirable kidneys are accepted after only a few offers. Over 20% of donors have at least one viable kidney discarded. Organs from these donors were refused by an average of almost 1500 patients.

⁷We end our sample in 2013 to rule out anticipatory effects and to avoid modeling transition dynamics as agents anticipate the new system introduced in December 2014.

TABLE I
PATIENT CHARACTERISTICS^a

	Patient Stocks, Arrivals, and Departures			
	Panel A: Patient Stocks January 1, 2010		January 1, 2013	
	Mean	S.D.	Mean	S.D.
Number of Patients	3901		4508	
Years on List	2.00	1.81	2.38	1.95
Years on Dialysis	3.24	3.63	3.24	3.63
Prior Transplant	14.8%	35.5%	14.0%	34.7%
Current Age	53.7	13.4	54.1	13.4
Calculated Panel Reactive Antibodies (CPRA)	9.8%	26.0%	14.4%	30.2%
Body Mass Index (BMI) at Arrival	27.8	5.9	27.9	5.8
Total Serum Albumin	4.0	0.6	4.0	0.6
Diabetic Patient	39.5%	48.9%	40.0%	49.0%
On Dialysis at Arrival	76.2%	42.6%	69.7%	46.0%

	Panel B: Patient Arrivals and Departures			
	Year 2010	Year 2011	Year 2012	Year 2013
# Patients Arriving	1379	1523	1503	1317
# Patients Departing	1093	1221	1304	1199
# Patients Received Deceased Donor Transplant	456	489	580	424
# Patients Received Living Donor Transplant	222	258	237	286
# Patients Died or Too Sick to Transplant	230	254	301	297
# Patients Departed for Other Reason	185	220	186	186

^a9623 patients were active on the NYRT waiting list at some time between January 1st, 2010 and December 31st, 2013. Panel A contains statistics for patients registered in NYRT on January 1st of each calendar year. Panel B contains statistics for patients who joined the NYRT waiting list (arrivals) and who were removed from the waiting list (departures) during each calendar year. Panel C classifies departures by reason. "Departed for Other Reason" includes transfers to non-NYRT transplant centers and miscellaneous departure reasons. Patients who received transplants at a non-NYRT center are included in the Received Deceased Donor Transplant and Received Live Donor Transplant categories.

Our observable donor covariates correlate with number of offers and discards in expected ways. Panel A of Table II summarizes certain organ characteristics by the allocation outcome for kidneys recovered in NYRT. Donors whose kidney(s) was/were discarded have higher creatinine levels (indicating lower kidney function) and are older, less likely to die of head trauma, more likely to be diabetic or hypertensive, and more likely to have donated after cardiac death.

NYRT patients are also offered donors from other parts of the country. Panel B shows that a total of 1410 donors were offered to patients registered with NYRT in the average year. Because most of these donors are offered to NYRT patients after being refused by many patients outside NYRT, these donors tend to have undesirable organs with many offers and high discard rates.

Waitlist, Offers, and Acceptance Rates

A patient's position on the list improves with waiting time, and the mean waiting time falls moving down the list of patients offered a given donor. The average waiting time among patients in the top five positions exceeds 4 years, while this average is approximately 3.5 years for patients in the one hundredth position. Nonetheless, the system is not well approximated by a first-come, first-served queue: points and priorities depend strongly on characteristics other than waiting time. We calculated the fraction of times

TABLE II
DONOR CHARACTERISTICS^a

	All		Yes		No	
	Mean	S.D.	Mean	S.D.	Mean	S.D.
			Panel A: Donors Recovered in NYRT; By Number of Organs Allocated			
Number of Donors Per Year	183.5		44.75		138.75	
Median Number of Offers per Donor	26.0		479.0		14.0	
Number of Offers per Donor	445.8	1442.6	1445.0	2620.9	123.5	346.1
Number of Kidneys Transplanted per Donor	1.52	0.80	0.35	0.48	1.90	0.43
Donor Age	43.8	18.0	56.0	14.2	39.9	17.3
Cause of Death – Head Trauma	25.2%	43.4%	11.7%	32.3%	29.5%	45.7%
Cause of Death – Stroke	43.9%	49.7%	60.9%	48.9%	38.4%	48.7%
Diabetic Donor	14.4%	35.2%	25.1%	43.5%	11.0%	31.3%
Hypertensive Donor	38.6%	48.7%	60.9%	48.9%	31.4%	46.4%
Expanded Criteria Donor (ECD)	31.1%	46.3%	58.1%	49.5%	22.3%	41.7%
Non-Heart Beating Donor (DCD)	9.7%	29.6%	12.3%	32.9%	8.8%	28.4%
Donor Creatinine	1.3	1.5	1.5	1.2	1.3	1.6
			Panel B: All Donors, By Number of Organs Allocated			
Number of Donors Per Year	1410.5		896		514.5	
Median Number of Offers per Donor	726.0		1183.5		230.5	
Number of Offers per Donor	1620.7	2564.3	2180.6	2965.0	645.6	1105.6
Number of Kidneys Transplanted per Donor	0.72	0.89	0.20	0.40	1.61	0.78

^a Panel A consists of all deceased kidney donors (734) recovered in NYRT and offered to NYRT patients between January 1st, 2010 and December 31st, 2013. Panel B includes all donors (5642) offered to NYRT patients during the same period, including donors recovered outside NYRT. Offers exclude cases in which the donor did not meet the patient's predetermined criteria for acceptable donors, or in which the patient was bypassed by the waitlist system due to operational considerations that did not involve an active choice by the patient or her surgeon.

that two patients who are offered the same donor are ordered identically on the list for the next donor they are both offered. This fraction is 81.5% (details available upon request). It would be 100% in a first-come, first-served system.

Because lower-quality kidneys are offered to a large number of potential recipients, patients receive many offers and the mean acceptance rate is low. Table III describes patterns of offer and acceptance rates.⁸ Panel A considers all feasible offers, including biologically compatible offers that did not meet the patient's preset criteria. A typical patient receives over 200 offers per year, but only 0.14% of offers are accepted. When interpreting these low acceptance rates, it is important to remember that undesirable organs are offered to a very large number of patients.

Offers from desirable donors are rare and much more likely to be accepted. Kidneys recovered in NYRT are accepted five times more often than those recovered outside NYRT, and 10.8% of perfect tissue-type match offers are accepted. Panels B further restricts to offers to patients in the first ten positions that met the patients' screening criteria. These offers are likely more attractive since the organs have not been refused by many patients. The typical patient can expect to receive less than one such offer each year and accepts one in four. Older patients are less likely to receive one of these desirable offers, in part because they are more likely to depart the waitlist before receiving a transplant. This fact puts them at a disadvantage in a mechanism that prioritizes waiting time.

Taken together, these statistics suggest that the supply of desirable donors in NYRT is scarce and it is necessary to wait for several years to gain priority to access these donors. This creates a strong incentive for patients to wait.

Motivating Evidence for Waitlist Redesign

Table IV describes transplants that occur. Pediatric patients receive high priority and are very likely to be transplanted, either with a deceased donor kidney or through a living donor. Transplant rates are lower among adults, and there is no significant gradient in transplant probability with age. Panel A describes deceased donors by age for transplanted patients. Pediatric patients almost exclusively receive transplants from young donors. In contrast, while there is some assortative matching by age among adults, many patients above age 65 receive kidneys from young and middle-aged donors.

The 2014 reforms were in part motivated by concerns that these observations reflect mismatch, and that healthier (typically younger) patients should receive the organs predicted to last longest (Israni et al. (2014)). However, it may be optimal to match some younger donors to older patients if there is considerable match-specific heterogeneity, or if other donor characteristics strongly predict organ quality. Our structural model will allow us to quantify the extent to which a welfare-maximizing mechanism should match donors and patients based on age rather than other factors, taking heterogeneity and dynamic incentives into account.

Another patient characteristic that will be important for the improvements we identify is whether the patient had already begun dialysis when they joined the waitlist. Some patients with kidney disease who qualify for the waitlist still have marginal kidney function and can avoid dialysis for some time. These patients are relatively healthy compared to patients already on dialysis at registration. The last two columns of Table IV show that while patients not yet on dialysis at registration represent 32% of our patient sample, they receive only 27% of deceased donor transplants.

⁸We exclude offers that did not result in assignment due to logistical reasons such as surgeon unavailability, OPO operational considerations, or special donor-specific considerations such as expedited organ placement or directed donation.

TABLE III
 RATES OF RECEIVING AND ACCEPTING OFFERS^a

	Number of Patients	Offer & Acceptance Rates					
		All Donors		NYRT Donors		Perfect Tissue Type Match	
		Annual Rate	% Accepted	Annual Rate	% Accepted	Annual Rate	% Accepted
All	9623	213.2	0.14%	37.0	0.74%	0.095	10.8%
On Dialysis at Registration	6513	211.5	0.15%	38.6	0.77%	0.090	11.4%
Not on Dialysis at Registration	3110	216.6	0.12%	33.6	0.66%	0.103	9.8%
Age 0-49 at Registration	3921	211.8	0.16%	38.3	0.81%	0.087	14.7%
Age 50+ at Registration	5702	214.1	0.13%	36.1	0.68%	0.100	8.6%
		Panel A: All Offers					
		Panel B: Offers Within the First 10 Positions that Met Screening Criteria					
All	9623	0.8	24.53%	0.8	25.81%	0.021	48.6%
On Dialysis at Registration	6513	0.8	24.70%	0.8	25.93%	0.017	48.5%
Not on Dialysis at Registration	3110	0.8	23.99%	0.8	25.40%	0.029	48.9%
Age 0-49 at Registration	3921	1.3	25.37%	1.3	26.23%	0.020	56.5%
Age 50+ at Registration	5702	0.5	23.34%	0.4	25.17%	0.022	42.7%

^aThere were 2,713,172 offers made to NYRT patients between January 1st, 2010 and December 31st, 2013. Panel B restricts to the first 10 NYRT patients in each donor's offer sequence. An offer Met Screening Criteria if the offer satisfied a patient's predetermined criteria for acceptable donors. "Annual Rate" columns report annual offer rates computed by patient and then averaged across patients.

TABLE IV
TRANSPLANT RATES BY AGE AND DIALYSIS STATUS^a

	Patient Age					On Dialysis at Registration	
	0–17	18–34	35–49	50–64	65+	Yes	No
Patient Population	2.1%	11.0%	27.3%	41.5%	18.1%	67.8%	32.2%
Transplanted Patients	4.6%	9.7%	26.7%	41.3%	17.6%	73.3%	26.7%
Panel A: Donor Age and Quality							
Donor Age 0–17	25.6%	14.7%	7.5%	4.5%	4.7%	6.8%	8.7%
Donor Age 18–35	72.2%	36.8%	26.3%	16.4%	12.5%	22.3%	24.6%
Donor Age 35–49	2.2%	30.5%	34.2%	28.7%	16.6%	27.6%	25.2%
Donor Age 50–64	0.0%	16.8%	29.9%	42.6%	53.4%	37.3%	34.6%
Donor Age ≥ 65	0.0%	1.1%	2.1%	7.9%	12.8%	5.9%	6.9%

^aPercentages in the first two rows represent the share in the patient population with a specific characteristics. Percentages in the second row restrict to patients that received a deceased donor transplant. Each column in Panel A describes the share of transplanted patients within the group indicated in the column that received a transplant from a given donor type.

Evidence of Response to Dynamic Incentives

A central assumption in our framework is that agents are forward-looking and respond to dynamic incentives. One implication of this assumption is that patients for whom the option value of waiting is lower should be less selective. Agarwal et al. (2018) presented descriptive evidence consistent with dynamic incentives using data from all areas of the United States and find that highly sensitized patients who are immunologically compatible with fewer donors—and who therefore expect to receive fewer offers in the future—are more likely than less sensitized patients to accept similar offers. We replicated this research strategy in the NYRT sample and found qualitatively similar patterns (see Section D.2 in the Supplemental Material (Agarwal, Ashlagi, Rees, Somaini, and Waldinger, 2021)). We refer the reader to this paper for a more detailed discussion of the empirical strategy and potential confounds.

3. A MODEL OF DECISIONS ON A WAITLIST

We use the following notation throughout the paper. Objects are indexed by j , agents by i . Let x_i denote the vector of observed characteristics of agent i ; z_j the observed characteristics of object j that are included in the model; η_j a donor quality index constructed from characteristics that are observed by agents but not included in the empirical specifications;⁹ and t indicates the amount of time the agent has been waiting on the list. The model does not include unobserved agent heterogeneity. We discuss this restriction in Section 4. Finally, let $c_{ij} = 1$ if j is compatible with i , and 0 otherwise. Incompatibility can arise for biological reasons in organ allocation but may arise due to other restrictions (e.g., legal) in other contexts.

This section begins by defining a class of sequential assignment mechanisms, primitives governing agents' decisions, and the agent arrival and departure processes. We then provide assumptions on agents' payoffs and beliefs which lead to a tractable optimal stopping problem from the agent's perspective.

⁹We include this term for two reasons. First, there are hundreds of donor characteristics in the dataset, far too many to include in the empirical specifications. Second, surgeons have information about the organ that are not included in the dataset.

3.1. Mechanisms and Primitives

3.1.1. Mechanisms

We consider sequential assignment mechanisms that use a priority score. The mechanism allocates each object as follows:

- Step 1 (Ordering): The priority score $s_{ijt} \equiv s(t; x_i, z_j)$ is calculated for all agents on the waitlist. Ties in the score, if any, are broken using a known tie-breaking rule.
- Step 2 (Offers): Each agent may decide to accept or reject the object, with acceptance denoted by $a_{ij} = 1$. The mechanism may solicit decisions from multiple agents simultaneously but may not skip any agents in the priority order. No offers are made to agents known to be incompatible with the object.
- Step 3 (Assignment): The objects are assigned to the agents with the highest q_j priorities for whom $a_{ij} = 1$, where q_j is the number of copies of the object available.
- Step 4 (Arrivals and Departures): An agent is removed from the waitlist once an object has been assigned to them. Other agents may join or leave the list.

Within the set of offer-based waitlist mechanisms, the primary restriction in our formulation is that an agent's priority does not depend on either the other agents in the market or their past actions. All deceased-donor organ allocation mechanisms and proposals considered prior to the 2014 reform were offer mechanisms based on priority scores which shared these properties.¹⁰ Our empirical exercises required us to develop computer code to calculate the order in which any organ would be offered using administrative data on patient and donor characteristics. We verified our code's output using the observed sequence of offers.

A complication in our setting is that organs must be allocated within a certain time frame. Several logistical constraints limit the number of offers that can be made. We treat the maximum number of offers that can be made for each object as exogenous because we do not have data on the factors that determine this limit. Our qualitative conclusions are robust to removing this limit.

3.1.2. Payoffs

There are three types of primitive payoffs in the model. The first is the expected net present value to agent i of being assigned object j after waiting t periods, $\Gamma_{ij}(t)$. In our application, $\Gamma_{ij}(t)$ captures a patient's value of various kidney transplants. The second is the expected net present value from departing without an assignment, $D_i(t)$. In our application, departures occur due to living donor transplants, death, or transfers to other listing centers (see Table I). The term $D_i(t)$ incorporates all of these reasons.¹¹ Finally, agents receive an expected flow payoff while waiting on the list, $d_i(t)$. In our application, $d_i(t)$ is the flow payoff from living without a functioning kidney.

Two implicit assumptions about the payoffs in our model are worth noting. First, we abstract away from costs of considering an offer. These costs are likely small relative to the value of transplants. Second, we assume that agents only value their own outcomes and not those of others. This may be violated if surgeons value the outcomes of other

¹⁰Based on an examination of committee reports and public comments downloaded from <https://optn.transplant.hrsa.gov/members/committees/kidney-committee/>.

¹¹We can represent the value of a departure as a weighted average over the value of the various events, that is, $D_i(t) = \sum_k p_{ik}(t)D_{ik}(t)$ where k denotes the type of departure (e.g., obtaining a live donor, death, etc.) and $p_{ik}(t)$ is the probability of each type of departure conditional on a departure occurring. The formulation only requires that each source satisfy Assumption 1(iii) below.

patients, especially those they are treating. NYRT has a total of 10 transplant hospitals staffed with many more kidney transplant surgeons, limiting common agency problems that surgeons might face. We assume that surgeons act in the best interest of each of their patients while the patient considers an offer.

We make the following assumptions about these payoffs.

ASSUMPTION 1: (i) *The expected net present value of an assignment is additively separable in a payoff shock ε_{ijt} :*

$$\Gamma_{ij}(t) \equiv \Gamma(t, x_i, z_j, \eta_j) + \varepsilon_{ijt}. \quad (1)$$

- (ii) *The random variables ε_{ijt} are independent and identically distributed (iid) with a known, nonatomic distribution with cumulative distribution function (cdf) G .*
 (iii) *The expected flow payoffs from waiting, $d_i(t)$, and the expected payoff from departing without an assignment, $D_i(t)$, are exogenous and depend only on (x_i, t) .*

The separability and independence restrictions on ε_{ijt} imposed in Assumptions 1(i) and 1(ii) are common in the dynamic discrete choice literature (see Rust (1987), e.g.). They enable an estimation approach based on the inversion technique of Hotz and Miller (1993). Our empirical specifications assume that G is the cdf of the standard normal, Φ . An important restriction in Assumption 1 is that it excludes persistent agent-level unobserved heterogeneity. Instead, our specifications include a rich set of patient and donor characteristics. Section 4 discusses challenges with relaxing this assumption and the advantages of our approach.

3.1.3. Arrivals and Departures

Time is continuous. Objects and agents arrive at Poisson rates λ and γ , respectively. The characteristics x of each arriving agent are iid. Similarly, each object's characteristics (z, η) are drawn iid from the cdf F upon arrival. We assume that each object must be assigned before the next object is offered. The Poisson arrival process and continuous time together imply that simultaneous arrivals are zero probability events.¹² Agents may depart the list due to or prior to assignment.

We make the following assumption on the arrival and departure processes.

ASSUMPTION 2: (i) *Departures prior to assignment and arrivals are governed by independent and exogenous Poisson processes.*

- (ii) *The departure rate of agent i at time t is given by $\delta_i(t) \equiv \delta(t; x_i)$. Further, each agent has a terminal date $T_i < \infty$ at which departure occurs with probability 1.*

In our application, we assume that patients die on or before their 100th birthday. T_i therefore corresponds to the waiting time for a patient on the day they turn 100 years of age.¹³

The primary economic restriction in our application is that departures prior to assignment and arrivals do not depend on the design of the deceased donor kidney waitlist.

¹²The continuous time formulation is consistent with the real world and computationally convenient. A discrete-time model would have to define a period length and specify how decisions are made when multiple offers arrive in one period. A time interval short enough so that only one offer can arrive at a time adds notational and computational complexity without substantive differences.

¹³It is straightforward to allow for agents that could remain on the list forever, $T_i = \infty$, as long as the value function for each patient approaches a constant.

Table I shows that the most common reason for departure without a deceased donor transplant is death or becoming “too sick to transplant.” It seems safe to assume that these events are not responsive to the design of the waitlist. The second most common reason is receiving a living donor transplant. Departures due to this reason are exogenous if the design of the waitlist does not affect the probability of finding a compatible living donor, and if patients always prefer a living donor to staying on the deceased donor waitlist. These conditions are plausible in our setting because living donors produce much better post-transplant outcomes.¹⁴

Similarly, we assume that agent arrivals do not depend on the design of the waitlist. During our sample period, patients could register as soon as they started dialysis or their kidney function was sufficiently low (below a glomerular filtration rate of 20 mL per minute). Therefore, it is in a patient’s interest to join the waitlist as soon as possible.

3.2. Individual Agent’s Problem

Agents on the waitlist who receive an offer of an object must decide whether to accept it or reject it and wait for future offers. This results in an optimal stopping problem from the perspective of the agent (Pakes (1986), Rust (1987)). This section starts by describing the offers an agent can expect to receive, which allows us to set up the value function. Finally, we introduce a normalization that simplifies the value function and is without loss of generality for our counterfactual exercises.

3.2.1. Offer Probabilities and Beliefs

To make an informed decision, an agent must form beliefs about the organs they can obtain in the future if they decline the current offer. Let s_j^* be the priority score of the pivotal agent offered object j . Agent i is assigned a compatible object j if they accept it and their score s_{ijt} is at least s_j^* . Therefore, it is sufficient for an agent to form beliefs over the probability distribution of s_j^* . In the kidney allocation context, s_j^* is the lowest priority score among the transplanted patients if none of donor j ’s kidneys are discarded. If at least one kidney is discarded, then s_j^* is the lowest priority score amongst the patients offered the kidney. The score s_j^* is random because it depends on the decisions of all agents offered object j , their compatibility, and the maximum number of offers that can be made.

Beliefs about the distribution of s_j^* depend on organ quality and the agent’s information about the competitive environment. In principle, the agent’s information set could include their previous offers and the other agents currently on the waitlist. However, there are several reasons, discussed below, why beliefs are unlikely to be sensitive to such detailed information. We therefore make the following assumption.

ASSUMPTION 3: *Each agent i believes that the probability that an object with characteristics (z_j, η_j) is compatible and will be available to them if it arrives when the agent has waited t periods is*

$$\pi(t; z_j, \eta_j, x_i) = H(s_{ijt}; z_j, \eta_j) \times \mathbb{P}(c_{ij} = 1 | z_j, x_i),$$

where $s_{ijt} = s(t; z_j, x_i)$ and $H(\cdot; z_j, \eta_j)$ is the cdf of the cutoff s_j^* given (z_j, η_j) .

¹⁴Hart et al. (2017) compared chances of graft failure 10 years after transplantation for living and deceased donor organs. This statistic for adults transplanted with a deceased donor kidney in 2005 is 52.8%, whereas it is only 37.3% for those who received a kidney from a living donor.

This assumption embeds three key restrictions. First, it assumes that beliefs are not sensitive to short-term variation in the set of other agents currently on the waitlist. This concern is limited because a surgeon cannot directly obtain information about patients they are not treating, and NYRT has ten transplant hospitals and many surgeons. Second, it abstracts away from inference about the likelihood of future offers based on past offers. The set and order of patients on the waitlist varies significantly across donors due to patient-donor specific compatibility and priority, limiting the ability to predict future offers based on recent experience. Furthermore, our data do not indicate that recent offers are predictive of future cutoffs,¹⁵ or that a patient's recent offer history predicts acceptance.¹⁶ Third, it assumes that the probability an organ is compatible depends only on observables and is independent of the cutoff. This restriction is appropriate in our context because surgeons list the blood and tissue types that are known to be incompatible with the patient.

Assumption 3 is a reasonable approximation if beliefs are based on average offer probabilities observed by a surgeon over a long period of time. This simplification is similar in spirit to other equilibrium concepts that condition on steady state or long-run quantities in order to make the analysis of dynamic games tractable (Hopenhayn (1992), Krusell and Smith (1998), Weintraub, Benkard, and van Roy (2008), Fershtman and Pakes (2012)). The assumption eases analysis relative to more general forms of beliefs because it avoids conditioning on the sequence of observed offers in the past.¹⁷ Despite this simplification, our state space continues to be quite rich because it contains all characteristics that influence payoffs, priorities, or compatibility.

3.2.2. Value Functions

We assume that agents make optimal accept/reject decisions by comparing the net present value of an object j to the value of waiting. Agent i decides to remain on the list instead of accepting object j if the payoff from an assignment $\Gamma_{ij}(t)$ is less than the value of continuing to wait conditional on the agent's type x_i and current waiting time t , denoted $V_i(t) \equiv V(t; x_i)$. The Hamilton–Jacobi–Bellman differential equation defining the value of waiting at time t is:

$$(\rho + \delta_i(t))V_i(t) = d_i(t) + \delta_i(t)D_i(t) + \lambda \int \pi_{ij}(t) \int \max\{0, \Gamma_{ij}(t) - V_i(t)\} dG dF + \dot{V}_i(t), \quad (2)$$

where ρ is the discount rate, the inner integral computes the expectation over the idiosyncratic payoff shocks ε_{ijt} in equation (1), and, with a slight abuse of notation, $\pi_{ij}(t) = \pi(t; x_i, z_j, \eta_j)$ defined in Assumption 3.

¹⁵We cannot reject zero autocorrelation in priority score cutoffs s_j^* across organs ordered by the date on which they arrived, even within donor subgroups. One would not expect this to be the case if the offers a patient observed contained information about the likelihood of receiving future offers. Details available on request.

¹⁶We tested whether, controlling for characteristics of both the patient and the current offer, recent offers predict acceptance behavior. Our estimates are precise and suggest that recent offer rates are not predictive of current acceptance behavior. This result is robust to various measures of recent offer rates and versions that only consider offer rates from desirable donors. Details available on request.

¹⁷Formally, beliefs would be conditioned on a richer information set, denoted $\mathcal{F}_{i,t}$, by replacing $H(s_{ijt}; z_j, \eta_j)$ with $H(s_{ijt}; \mathcal{F}_{i,t}, z_j, \eta_j) = \mathbb{P}(s_j^* < s_{ijt} | \mathcal{F}_{i,t}, z_j, \eta_j)$. In our setting, agents only know the past offers that they have received.

This expression can be derived by considering an agent's value of waiting at time t for an infinitesimal duration Δt . If no object arrives, the agent incurs flow payoffs from dialysis $d_i(t)\Delta t$ and may depart exogenously with probability $\delta_i(t)\Delta t$, receiving payoff $D_i(t)$. An object arrives with probability $\lambda\Delta t$, its characteristics drawn from the cdf F . The integral calculates the expected increment in the agent's value function due to each arrival. Specifically, the agent receives an offer for this object with probability $\pi_{ij}(t)$ and accepts it if $\Gamma_{ij}(t) > V_i(t)$, yielding an incremental value of $\int \max\{0, \Gamma_{ij}(t) - V_i(t)\} dG$. In the limit as $\Delta t \rightarrow 0$, the probability that both departures and object arrivals occur within the interval Δt tends to zero, yielding the differential equation above.¹⁸

Due to the option value of potential future offers, a patient may refuse an offer that they would accept if the alternative were never receiving a transplant. Moreover, certain marginal organs may be refused by all patients on the list and discarded even though some patients would benefit from the transplanted organ relative to no assignment.

3.2.3. Normalization and Simplifying the Value Function

A typical dataset from a sequential assignment mechanism only contains information about accept/reject decisions. As is well understood, data on actions alone do not identify all primitives of a dynamic discrete choice model, and the payoff from one action must be normalized in each state (Magnac and Thesmar (2002)). However, Aguirregabiria and Suzuki (2014) and Kalouptsi, Scott, and Souza-Rodrigues (2020) point out that such normalizations may affect certain counterfactual analyses because they arbitrarily restrict payoffs from specific actions across states. If this is the case, the counterfactuals are not identified. This fact poses a potentially serious barrier to answering questions that depend on primitives that are not identified from choice data.

Fortunately, our model admits a valid normalization such that counterfactuals involving changes in the mechanism are identified. Consider the inclusive value of an agent who declines every waitlist offer, $O_i(t)$.¹⁹ Such an agent receives the flow payoff of waiting and the value of departure when it occurs. The value at time t is defined by the differential equation

$$(\rho + \delta_i(t))O_i(t) = d_i(t) + \delta_i(t)D_i(t) + \dot{O}_i(t)$$

and the terminal condition $O_i(T_i) = D_i(T_i)$. Assumptions 1(iii) and 2(i) together imply that $O_i(t)$ does not depend on the mechanism. As we now show, it suffices to measure $V_i(t)$ and $\Gamma_{ij}(t)$ relative to $O_i(t)$ in order to analyze optimal decisions and changes in welfare induced by offer probabilities under current and alternative mechanisms:

¹⁸The discretized version of the equation defining the value of waiting at time t is

$$V_i(t) = \frac{1}{1 + \rho\Delta t} \left[d_i(t)\Delta t + \delta_i(t)\Delta t D_i(t) + \lambda\Delta t \int \pi_{ij}(t) \int \max\{V_i(t + \Delta t), \Gamma_{ij}(t)\} dG dF \right. \\ \left. + (1 - (\delta_i(t) + \lambda_i(t))\Delta t)V_i(t + \Delta t) + o(\Delta t) \right],$$

where $\lambda_i(t) = \lambda \int \pi_{ij}(t) dF$ is the rate at which agent i expects to receive an offer at time t . The leading fraction represents discounting due to time preferences. The first three terms inside the brackets are described in the text. The remainder term includes the payoff in the event that multiple donors or objects arrive, or that a donor arrives and the patient departs, within Δt . These events have probability of order $o(\Delta t)$. Therefore, the remainder is of order $o(\Delta t)$ as long as all expected payoffs are bounded. Taking the limit as $\Delta t \rightarrow 0$ under mild continuity conditions yields the differential equation above.

¹⁹It is the inclusive value of choosing the outside option defined as no assignment through the deceased donor waitlist, including the possibility of receiving a living donor transplant.

PROPOSITION 1: Let $V_i(t)$ be the solution to equation (2) with boundary condition $V_i(T_i) = O_i(T_i)$. Suppose Assumptions 1–3 hold. Then (i) $\tilde{V}_i(t) = V_i(t) - O_i(t)$ solves

$$(\rho + \delta_i(t))\tilde{V}_i(t) = \lambda \int \pi_{ij}(t) \int \max\{0, \tilde{\Gamma}_{ij}(t) - \tilde{V}_i(t)\} dG dF + \dot{\tilde{V}}_i(t)$$

with boundary condition $\tilde{V}_i(T_i) = 0$, where $\tilde{\Gamma}_{ij}(t) = \Gamma_{ij}(t) - O_i(t)$, and (ii) for any two functions π and π' , $V_i(t; \pi) - V_i(t; \pi') = \tilde{V}_i(t; \pi) - \tilde{V}_i(t; \pi')$ where the parameter π is introduced in the notation for clarity.

See Appendix A.1 for the proof. This result shows that optimal decisions and several counterfactual quantities are invariant to $O_i(t)$. Part (i) states that the value of waiting relative to either declining every waitlist offer or accepting any given offer does not depend on the function $O_i(t)$. Therefore, models with different functions $O_i(t)$ yield the same optimal decisions and are observationally equivalent. Part (ii) states that different functions $O_i(t)$ yield the same differences in the value function between any two offer probabilities and therefore under any two mechanisms that produce these offers.

We therefore normalize $O_i(t)$ to zero at all $t \in [0, T_i]$. The rest of the paper interprets $\Gamma_{ij}(t)$ and $V_i(t)$ as values relative to declining every waitlist offer. This normalization implies that $d_i(t) + \delta_i(t)D_i(t) = 0$ for all t and that $D_i(T_i) = 0$. Hence, a benefit is that we no longer need to estimate the payoffs from remaining on the list or departing without an assignment because equation (2) reduces to

$$(\rho + \delta_i(t))V_i(t) = \lambda \int \pi_{ij}(t) \int \max\{0, \Gamma_{ij}(t) - V_i(t)\} dG dF + \dot{V}_i(t). \quad (3)$$

It will be useful to rewrite this differential equation as

$$\begin{aligned} V(t; x_i) &= \int_t^{T_i} \exp(-\rho(\tau - t)) p(\tau|t; x_i) \\ &\quad \times \left(\lambda \int \pi(\tau; x_i, z, \eta) \int \max\{0, \Gamma(\tau, x_i, z, \eta) + \varepsilon_{ij} \right. \\ &\quad \left. - V(\tau; x_i)\} dG dF \right) d\tau, \end{aligned} \quad (4)$$

where

$$p(\tau|t; x_i) \equiv \exp\left(-\int_t^\tau \delta(\tau'; x_i) d\tau'\right)$$

is the probability that agent i does not exogenously depart before τ conditional on being on the list at t . We have explicitly reintroduced agent and object characteristics into the notation because this equation will form the basis of our empirical strategy. This expression uses the boundary condition $\lim_{t \rightarrow T_i} V(t; x_i) = D_i(T_i) = O_i(T_i) = 0$ because the probability of receiving an offer after t vanishes as $t \rightarrow T_i$. A similar result holds for an alternative model that considers the limit case with $T_i = \infty$.

4. ESTIMATION

The key primitives needed to predict equilibrium allocations and welfare under alternative mechanisms are the transplant values, $\Gamma(t, x, z, \eta)$. The challenge for estimation is

that acceptance decisions in our data depend on both the value of the offered organ and the value of continuing to wait. This section begins by representing the value function and the payoff from assignment in terms of Conditional Choice Probabilities (CCPs) and then describes the estimation approach, presents the empirical specification, and discusses alternative methods.

4.1. A CCP Representation

We develop a CCP representation which will allow us to estimate our model without solving the integral equation that defines the value function (equation (4)). Let $P_{ijt} = P(t, x_i, z_j, \eta_j)$ be the conditional probability that agent i refuses an offer of kidney j given (t, x_i, z_j, η_j) . We will estimate the function $P(\cdot)$ in the next section. The CCP approach is based on two key relationships.

PROPOSITION 2: *Suppose Assumptions 1–3 hold. Then (i) the conditional choice probabilities are given by*

$$P_{ijt} = G(V(t; x_i) - \Gamma(t, x_i, z_j, \eta_j)), \quad (5)$$

and (ii) there exists a known function ψ such that

$$\psi(P_{ijt}) = \int \max\{0, \Gamma(t, x_i, z_j, \eta_j) - V(t; x_i) + \varepsilon_{ijt}\} dG, \quad (6)$$

where $P_{ijt} = P(t, x_i, z_j, \eta_j) \in (0, 1)$.

Part (i) follows from $P_{ijt} = \mathbb{P}(V(t; x_i) > \Gamma(t, x_i, z_j, \eta_j) + \varepsilon_{ijt} | t, x_i, z_j, \eta_j)$ and G being the CDF of ε_{ijt} . Equation (5) is analogous to the well-known result that CCPs are a known function of differences between choice-specific conditional value functions. Part (ii) follows from substituting $V(t; x_i) - \Gamma(t, x_i, z_j, \eta_j) = G^{-1}(P_{ijt})$ from equation (5) into equation (6), which yields the function $\psi(\cdot)$. The result, a special case of Proposition 1 in Hotz and Miller (1993), shows that the incremental value of an offer, defined as the difference between the value of continuing to wait and the expected value from the optimal choice, is a known function of CCPs. Since our application assumes that $G = \Phi$, the cdf of the standard normal, we obtain the simple closed-form expression $\psi(P) = \Phi'(\Phi^{-1}(P)) - (1 - P)\Phi^{-1}(P)$.

These two equations can be used to rewrite the primitive of interest, $\Gamma(t, x_i, z_j, \eta_j)$, by substituting part (ii) into the integral equation for $V(t; x_i)$ (equation (4)):

$$V(t; x_i) = \int_t^{T_i} \exp(-\rho(\tau - t)) p(\tau | t; x_i) \left(\lambda \int \pi(\tau; x_i, z, \eta) \psi(P_{ij\tau}) dF \right) d\tau. \quad (7)$$

Using this expression in conjunction with part (i), we obtain

$$\begin{aligned} \Gamma(t, x_i, z_j, \eta_j) &= -G^{-1}(P_{ijt}) \\ &+ \int_t^{T_i} \exp(-\rho(\tau - t)) p(\tau | t; x_i) \left(\lambda \int \pi(\tau; x_i, z, \eta) \psi(P_{ij\tau}) dF \right) d\tau. \end{aligned} \quad (8)$$

The next subsection shows how to estimate the primitives of interest, $\Gamma(\cdot)$ and the distribution of η_j , using the expressions above.

4.2. Estimation Procedure

We estimate the model in four steps. First, we estimate $p(\tau|t; x_i)$ and λ using observed patient departures and donor arrivals, respectively. Second, we estimate conditional choice probabilities from patient accept/reject decisions. Third, we estimate the integral in equation (7) using the empirical distribution of donor types and cutoff scores to estimate F and π . In the final step, we recover transplant values $\Gamma(t, x, z, \eta)$ by solving for each patient's value function at each date using equation (8).

As is well known, time preferences are not identified from observed choices alone in dynamic discrete choice models (Magnac and Thesmar (2002)). We therefore set the discount rate ρ to a fixed value of 5% per year. Our results are robust to using an annual discount rate of 10%. This robustness in part arises because agents also devalue future offers due to exogenous departures. Since we estimate the departure rate $\delta_i(t; x)$ to be 16% per year for the average patient, this term dominates ρ in our value function (see equation (3)).

Step 1: Estimating Rates of Patient Departure and Donor Arrival

Our dataset contains information on how long each patient is on the waitlist and whether departure occurred due to a transplant. Time on the list is a censored measure of the departure time without a deceased donor transplant if the patient is either transplanted with a deceased donor or still on the list at the end of the sample period. Despite censoring, these durations can be used to estimate departure rates independently of pay-offs because Assumption 2 implies that, conditional on patient characteristics, departure from the list prior to assignment is exogenous.

We use maximum likelihood to estimate a censored Gompertz proportional hazards model in which the rate of departure takes the form

$$\delta(t; x_i) = \delta_1 \exp(\delta_2 t) \exp(x_i \beta), \quad (9)$$

where $\delta_1 \exp(\delta_2 t)$ is the baseline hazard function and x_i are observed patient covariates. This parametric form yields a simple expression for the survival function $p(\tau|t; x_i) \equiv \exp(-\int_t^\tau \delta(\tau'; x_i) d\tau')$, which eases computation of equation (7). It turns out that the estimates from this model yield a survival curve similar to those from the semiparametric Cox proportional hazards model.

Since donor arrivals are exogeneous, we estimate λ as the empirical mean arrival rate, which is the maximum likelihood estimator for a Poisson process.

Step 2: Estimating the CCPs, $P(t, x_i, z_j, \eta_j)$

Because the state space is multidimensional and includes continuous variables, a frequency estimator for the CCPs is infeasible. We therefore model the accept/reject decision using a flexible basis and a link function (see Arcidiacono and Ellickson (2011)).

Assume that $\Gamma(\cdot)$ is additively separable in η_j and that $\eta_j \sim N(0, \sigma_\eta^2)$, with a variance to be estimated. Then, equation (5) implies that the CCPs are a function of the difference $\Delta(x_i, z_j, t) - \eta_j \equiv V(t; x_i) - \Gamma(x_i, z_j, \eta_j, t)$. We estimate the CCPs using a flexible basis for $\Delta(\cdot)$ and the probit link function. Specifically, we set

$$P(t, x_i, z_j, \eta_j) = \Phi(\chi(x_i, z_j, t)\theta - \eta_j),$$

where $\chi(\cdot)$ is a flexible set of functions with interactions between its arguments, θ is an unknown parameter to be estimated, and $\Delta(\cdot) = \chi(\cdot)\theta$.²⁰ Our specification of $\chi(\cdot)$ includes dummies in x_i and z_j for categorical variables, piecewise linear splines for their continuous elements and for t , and interactions among several categorical variables and splines.

Standard arguments imply the identification of these parameters. The parameter θ is identified by the relationship between the covariates and the probability of acceptance. The variance, σ_η^2 , of the donor-specific unobservable is identified because many donors have two offered kidneys. If σ_η^2 is large, then conditional on the observables x_i , z_j , and t , early acceptance of a donor's first kidney predicts that the second acceptance should soon follow. In contrast, if σ_η^2 is small, then the position of the first acceptance should have little information about the second. The intuition is similar for the identification of measurement error models (see Kotlarski's theorem in Rao (1992), Hu and Schennach (2008)).

We estimate the parameters (θ, σ_η) using a Gibbs' sampler (McCulloch and Rossi (1994), Gelman, Carlin, Stern, and Rubin (2014)), which is asymptotically equivalent to a maximum likelihood estimator (see van der Vaart (2000, Theorem 10.1 (Bernstein–von Mises))).²¹ This procedure is more computationally tractable than maximum likelihood for models with unobserved heterogeneity because computing the likelihood function given data on decisions and the observables (x_i, z_j, t) requires integrating over the distribution of η_j . The Gibbs' sampler uses data augmentation and Markov Chain Monte Carlo sampling for this task.

Step 3: Estimating $V(t; x_i)$ Using Offer Arrivals and CCPs

Next, we use equation (7) to calculate $V(x_i; t)$ by numerically integrating the incremental value generated by offers that a patient can expect to receive in the future. This step requires us to integrate over future points in time, τ , and future donor types (z, η) . We have estimated the terms λ and $p(\tau|t; x_i)$ in Step 1 and have set ρ as discussed earlier.

The only remaining term in equation (7) is the expected value from the arrival of a donor at time τ :

$$W(x_i, \tau; \theta_0) = \int \pi(\tau; x_i, z_j, \eta_j) \psi(P_{ij\tau}) dF.$$

Estimating this value is challenging because we observe at most one offer at time τ , and only for patients who are on the list. Moreover, whether an offer is made to patient i depends on the arriving donor's type, the patient's priority score, and other patients' decisions. Instead of estimating this high-dimensional object using offers received by similar patients, we directly use the properties of the mechanism and Assumption 3 to make progress.

²⁰This formulation is an example of an arbitrary link function L^{-1} that need not equal the inverse CDF of ε_{ij} , Φ^{-1} . In the general case, $V(t; x_i) - \Gamma(x_i, z_j, \eta_j, t) = G^{-1}(L(\chi(x_i, z_j, t)\theta - \eta_j))$, where the argument of $G^{-1}(\cdot)$ equals $P(t, x_i, z_j, \eta_j)$. The remaining estimation steps can proceed exactly as described below, but with minor modifications to the expressions. We set $L^{-1} = G^{-1}$ to simplify computation.

²¹We obtain draws of the parameters θ and σ_η from a sequence of conditional posterior distributions using a Markov chain given dispersed priors and an initial set of parameters $(\theta^0, \sigma_\eta^0)$. The invariant distribution of the Markov chain is the posterior given the prior and the data. Details on the implementation, including burn-in procedures and convergence diagnostics, are in Appendix A.2.

Recall that agent i receives an offer for object j if their known priority score $s_{ij\tau}$ is at least s_j^* . Knowledge of the mechanism allows us to calculate $s_{ij\tau}$. Therefore, it is sufficient to characterize the joint distribution of the characteristics of the arriving donors (z_j, η_j) and their cutoff scores s_j^* . Formally, Assumption 3 implies that we can rewrite $\pi_{ij}(t)$ to obtain

$$\begin{aligned} W(x_i, \tau; \theta_0) &= \int \int \mathbb{P}(c_{ij} = 1 | z_j, x_i) 1\{s_{ij\tau} \geq s_j^*\} \psi(P_{ij\tau}) dH(s_j^*; z, \eta_j) dF \\ &\equiv \mathbb{E}[\mathbb{P}(c_{ij} = 1 | z_j, x_i) 1\{s_{ij\tau} \geq s_j^*\} \psi(P_{ij\tau}) | x_i, \tau]. \end{aligned}$$

We estimate this expectation using the observed cutoffs and characteristics of all donors in our data because their distributions are time-invariant. Specifically, we estimate $W(x_i, \tau; \theta_0)$ using a sample analog that replaces $\mathbb{P}(c_{ij} = 1 | z_j, x_i)$ and $\psi(P_{ij\tau})$ with estimated quantities:

$$\hat{W}(x_i, \tau; \hat{\theta}) = \frac{1}{J} \sum_{j=1}^J \hat{P}(c_{ij} = 1) 1\{s_{ij\tau} \geq s_j^*\} \psi(\hat{P}_{ij\tau}), \quad (10)$$

where j indexes a donor in our sample. The first term, $\hat{P}(c_{ij} = 1)$, is an estimate of $\mathbb{P}(c_{ij} = 1 | z_j, x_i)$ constructed using information on donor proteins and patient immune system characteristics.²² The third term replaces $\psi(P_{ij\tau})$ with the estimate $\hat{P}_{ij\tau} = \Phi(\chi(x_i, z_j, \tau)\hat{\theta} + \eta_j)$, where $\hat{\theta}$ is the estimated parameter and η_j is a draw from its distribution. Because the distribution of s_j^* and η_j are dependent, we draw from the conditional distribution of η_j given the observed accept/reject decisions of all patients offered donor j .²³

Therefore, our estimate $\hat{V}(t; x_i)$ is obtained by numerically integrating

$$\exp(-\rho(\tau - t)) \hat{p}(\tau | t; x_i) \hat{W}(x_i, \tau; \hat{\theta})$$

over the range $\tau \in [t, T_j]$. Details on this procedure, including the number of points used, are provided in Appendix A.2. Substituting $\hat{W}(x_i, t; \hat{\theta})$ for $W(x_i, t; \theta_0)$ is appropriate because the former is a \sqrt{J} -consistent estimator of the latter for each (x_i, t) under standard regularity conditions assuming that the dependence of future offers on current arrivals diminishes with the time horizon of the future offer.²⁴

Step 4: Estimating Γ

Once $\hat{V}(t; x_i)$ has been estimated, we recover $\Gamma(\cdot)$ using equation (8):

$$\hat{\Gamma}(t, x_i, z_j, \eta_j) = \hat{V}(t; x_i) - G^{-1}(\hat{P}_{ij\tau}) = \hat{V}(t; x_i) - (\chi(x_i, z_j, t)\hat{\theta} - \eta_j).$$

This quantity can be calculated for any value of (t, x_i, z_j, η_j) using the estimate $\hat{\theta}$.

²²We estimate $\mathbb{P}(c_{ij} = 1 | z_j, x_i)$ using the data on instances where a kidney was accepted but the transplant did not occur because the final crossmatch was positive.

²³Formally, for each j , we need to draw from a random variable with expectation $\mathbb{E}[1\{s_{ij\tau} \geq s_j^*\} \psi(P_{ij\tau})]$. For simplicity of exposition, we have dropped the term $\mathbb{P}(c_{ij} = 1 | z_j, x_i)$ and explicit conditioning on (x_i, z_j, τ) . Let \mathcal{I}_j denote all offers and accept/reject decisions for kidney j . By the law of iterated expectations, $\mathbb{E}[1\{s_{ij\tau} \geq s_j^*\} \psi(P_{ij\tau})] = \mathbb{E}[\mathbb{E}[\psi(P_{ij\tau}) | \mathcal{I}_j] 1\{s_{ij\tau} \geq s_j^*\}]$ because s_j^* is measurable with respect to \mathcal{I}_j and $s_{ij\tau}$ is a function of (x_i, z_j, τ) . Therefore, the random variable $1\{s_{ij\tau} \geq s_j^*\} \psi(P_{ij\tau})$, where $\psi(P_{ij\tau}) = \psi(G(\chi(x_i, z_j, \tau)\hat{\theta} + \eta_j))$ and η_j is drawn from the conditional distribution given \mathcal{I}_j , has the desired expectation.

²⁴Formal conditions and a proof are available upon request.

Discussion

It is worth discussing two decisions made in service of our goal of finding mechanisms that better target offers based on observable characteristics. First, our estimation approach is based on CCP methods instead of a full solution/nested fixed-point approach (Miller (1984), Wolpin (1984), Pakes (1986), Rust (1987)). The main advantage of a CCP approach is that it can accommodate rich observed heterogeneity while maintaining computational tractability. A full-solution approach would directly parametrize the object of interest $\Gamma(\cdot)$ in terms of parameters θ_{Γ} and compute the likelihood of accepting an offer using the implied value function. Although maximizing this likelihood yields an efficient estimator for θ_{Γ} , it is burdensome when the state space is large because the value function must be computed for many values of θ_{Γ} . Indeed, our state space includes the compatibility and priority score for each patient and donor using all the variables that enter the assignment mechanism. Even after the simplification afforded by Assumption 3, our preferred specification includes approximately 80 variables, many of which are continuous.

Second, our model abstracts away from patient-level unobserved heterogeneity. An extension that incorporates it would introduce three complications. First, both the departure rates $\delta_i(t)$ and choice probabilities would ideally depend on unobserved heterogeneity and would need to be jointly estimated. Second, we would need to solve the initial conditions problem for the agents already on the waitlist at the beginning of the sample. Finally, estimating dynamic models with both rich observed and unobserved heterogeneity is demanding of the data. We believe that abstracting away from patient unobserved heterogeneity still yields useful results because our dataset contains a rich set of patient characteristics. However, this restriction could bias our results if patients differ in ways not captured by the rich patient covariates in our data. For this reason, we explored a simple specification in which a limited form of unobserved heterogeneity was included in the choice model only. This specification, presented in Section D of the Supplemental Material, yields qualitatively similar counterfactual results.

5. PARAMETER ESTIMATES

Our estimated models include a rich set of patient and donor observed characteristics. The baseline characteristics, linear splines, and interactions were chosen by surveying the medical literature. Specifically, we use covariate and spline specifications from KPSAM, which was used by the kidney allocation committee to predict the outcomes of various allocation systems. We also include any covariates that were part of the survival models for kidney transplant patients used in Wolfe et al. (2008).

5.1. *Estimated CCPs and Departure Rates*

Conditional Choice Probabilities

We estimated three specifications for the conditional choice probability of accepting an offer. The first specification includes all of these baseline variables but not donor unobserved heterogeneity η or the state variable time t . The second specification adds donor unobserved heterogeneity, and the third specification adds waiting time interacted with a variety of characteristics. Table V presents select parameter estimates (see Table B.III in the Supplemental Material for the full specification).

The estimated coefficients are intuitive and robust in the last two specifications. Including donor unobserved heterogeneity influences the coefficients, especially those on donor

TABLE V
 CONDITIONAL CHOICE PROBABILITY OF ACCEPTANCE (SELECT COEFFICIENTS)^a

	Base Specification (1)		Unobserved Heterog. (2)		Waiting Time + UH (3)	
Calculated Panel Reactive Antibody (CPRA)	0.60	(0.05)	0.68	(0.06)	0.58	(0.09)
Donor Age < 18	0.27	(0.10)	-0.09	(0.19)	-0.04	(0.20)
Donor Age 18–35	0.59	(0.12)	-0.06	(0.19)	0.02	(0.19)
Donor Age 50+	-0.83	(0.16)	-0.77	(0.21)	-0.87	(0.22)
Expanded Criteria Donor (ECD)	-0.14	(0.02)	-0.53	(0.08)	-0.53	(0.10)
Donation from Cardiac Death (DCD)	-0.10	(0.02)	-0.51	(0.06)	-0.50	(0.09)
Perfect Tissue Type Match	2.33	(0.31)	2.92	(0.43)	2.89	(0.44)
Regional Offer	-1.38	(0.06)	-2.90	(0.19)	-2.92	(0.19)
National Offer	-1.54	(0.04)	-3.05	(0.12)	-3.11	(0.11)
Patient on Dialysis at Registration	-0.02	(0.02)	-0.10	(0.02)	-0.09	(0.02)
Log Waiting Time (years)					0.09	(0.06)
Log Waiting Time × 1{Over 1 Year}					-0.15	(0.07)
Log Waiting Time × 1{Over 2 Years}					-0.13	(0.12)
Log Waiting Time × 1{Over 3 Years}					0.30	(0.11)
Patient Age × 1{Donor Age < 18}	-0.01	(0.00)	0.00	(0.00)	0.00	(0.00)
Patient Age × 1{Donor Age 18–35}	-0.02	(0.00)	0.00	(0.01)	0.00	(0.01)
Patient Age × 1{Donor Age 50+}	0.02	(0.00)	0.02	(0.01)	0.02	(0.01)
Donor Unobservable Std. Dev.			1.02	(0.03)	1.04	(0.04)
Idiosyncratic Shock Std. Dev.	1.00		1.00		1.00	
Acceptance Rate	0.140%		0.140%		0.140%	
Number of Offers	2,713,043		2,713,043		2,713,043	

^aSelect parameter estimates for the conditional probability of accepting an offer. All variables except patient age, CPRA, and waiting time are indicators. For each specification, we simulate 200,000 draws from the Markov chain generated by the Gibbs' sampler, burn-in the first half of the chain, and keep every tenth draw. The point estimate is the sample mean and the standard error estimate is the standard deviation of these remaining draws.

characteristics, as can occur in a nonlinear model. In all specifications, offers from donors older than 50 years of age are less likely to be accepted than offers from 35- to 50-year-old donors. We also estimate significant patient-level and match-specific heterogeneity in acceptance rates. A perfect tissue type match is very likely to be accepted, much more so than a young donor. Regional and national offers are far less likely to be accepted. Adult patients of different ages are equally likely to accept a middle-aged donor, but older patients are more likely to accept a donor who is over 50 years old. This pattern is consistent with the idea that it is more important for younger patients to obtain kidneys that are likely to function for a long time.

Including donor unobserved heterogeneity and flexible waiting time controls significantly improves model fit. Specifications with donor unobserved heterogeneity best match the level of and sharp decline in the average acceptance rate as the offer number increases (see Figure 1(a)). Since undesirable organs are more likely to be offered to patients further down the list, specifications without donor unobserved heterogeneity over-predict the number of organs that are accepted in the early positions. Similarly, a flexible spline in waiting time best fits the average acceptance rates by years waited (see Figure 1(b)).

Finally, we assessed whether our specification is overfit by comparing the relative mean-squared error (RMSE) in our estimation sample and the first 6 months of 2014. As expected, the in-sample RMSE falls but the out-of-sample RMSE increases if the model is made richer (see Table B.IV in the Supplemental Material for details).

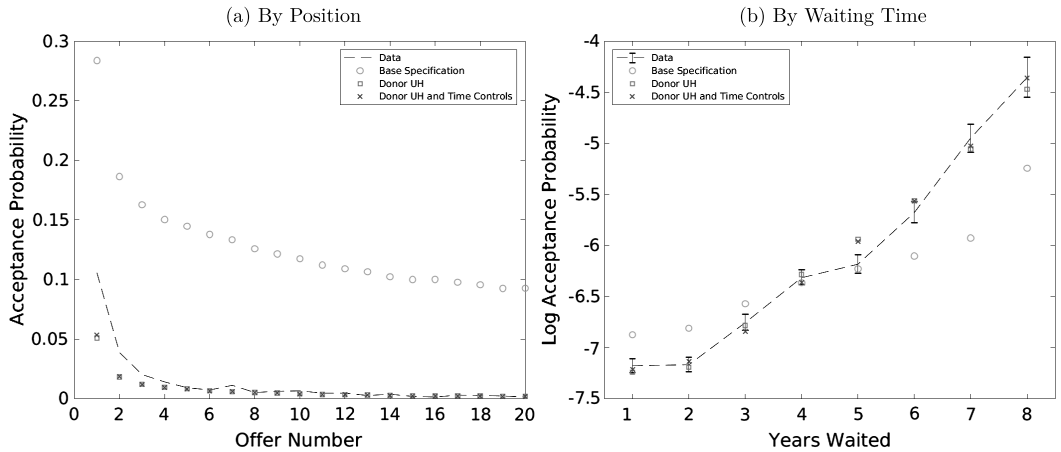


FIGURE 1.—Model fit.

This pattern suggests that our specification finds a good in-sample fit without overfitting.

Departure Rates

We estimated patient departure models under different parametric assumptions on the baseline hazard. These models include all patient-specific variables included in the CCP model.²⁵ Table B.II in the Supplemental Material presents these estimates. They reveal significant and robust heterogeneity in patient departure rates. For example, departure rates are higher for diabetic patients and for older adults. Patients with blood type A are also more likely to depart, potentially due to better chances of receiving a living donor transplant. Across specifications, we estimate an increasing baseline hazard of departure, consistent with patients becoming less healthy over time. The Gompertz hazard model results in a survival curve very similar to the Cox proportional hazards model.

For the reasons detailed above, we feel comfortable with the fit of the CCPs in the third model and the Gompertz hazard model. All results that follow use estimates from these specifications.

5.2. Estimated Value of Organ Offers

A common challenge when conducting welfare analysis in environments without money is the lack of a transferable numeraire good. We therefore introduce an interpretable measure for comparing values across agents in terms of an equivalent change in object arrival rates (supply) and use these units to describe our estimates. Our counterfactual results on patient welfare will also use these units.

Sequential assignment mechanisms can reassign offers from some agents to others. Motivated by this fact, consider the proportional increase in a patient's expected value of a

²⁵The hazard model includes CPRA as a linear term, but omits a dummy for CPRA > 0.8. Because priority is discontinuous in CPRA at 0.8, acceptance behavior can be discontinuous at this point. However, departures without a transplant should not be discontinuous.

one-time offer of object j at the time of registration:

$$EV_{ij} \equiv \frac{\int \max\{0, \Gamma(0, x_i, z_j, \eta_j) + \varepsilon_{ij} - V(0; x_i, \lambda)\} dG}{V(0; x_i, \lambda)},$$

where the dependence of V on the object arrival rate λ is reintroduced for clarity.

Instead of a one-time offer, the same change in value can be generated by an increase in the object arrival rate λ . Specifically, define λ_{ij} such that $V(0; x_i, \lambda_{ij}) - V(0; x_i, \lambda)$ equals the numerator of the expression above. We can rewrite EV_{ij} as

$$EV_{ij} = \frac{V(0; x_i, \lambda_{ij}) - V(0; x_i, \lambda)}{V(0; x_i, \lambda)} \approx \frac{\lambda_{ij} - \lambda}{\lambda},$$

where the approximation follows because $V(t; x_i, \lambda)$ is approximately linear in λ (see equation (4)). This approximation is appropriate for small changes in λ , so that offer probabilities and acceptance decisions do not change substantially.

Therefore, the value of a one-time offer as measured by EV_{ij} is equivalent to the value of an alternative policy that is able to marginally increase the supply of objects. Aggregating this measure yields a notion of social welfare in which the planner equally values the effect of a change in donor supply on all patients. This quantity is similar in spirit to Equivalent Variation at the time of registration. Note that by definition, EV_{ij} is invariant to the scale of utility units across agents.

Figure 2 describes our preferred estimates in these units. The plot shows how the value of an organ offer varies across specific patient and donor characteristics, holding all remaining characteristics fixed.

While all patients prefer younger donors, younger patients and patients not on dialysis at registration place a relatively higher value on younger donors. In contrast, older patients place a high value on offers from both young and old donors. These patterns are consistent with the differential life expectancy effects by age and health status of receiving a transplant from a higher-quality organ. Furthermore, because older patients receive fewer offers in the first ten positions (see Table III), each additional offer represents a greater equivalent increase in donor supply.




FIGURE 2.—Value of an organ offer. Notes: Patient and donor characteristics other than those explicitly varied are set to median values from the NYRT sample. Value functions are computed at registration.

Tissue type similarity is also valuable. Organs with no DR tissue protein mismatch are equivalent to a larger increase in overall donor supply than those with mismatches. The former are less likely to induce an adverse immune response post-transplant and, therefore, increase the expected life-years afforded by the transplant.

6. STEADY STATE EQUILIBRIA AND WELFARE COMPARISONS

6.1. *Equilibrium Concept*

We now define an equilibrium concept for counterfactual analysis. Our concept describes a large pool of agents waiting for offers from an allocation mechanism that uses a scoring rule which does not depend on an agent's past decisions. The endogenous objects are the equilibrium offer probabilities, the composition of agents on the waitlist, and optimal decision rules with their associated value functions.

Agents have types $x \in \mathcal{X}$, and objects have types $z \in \zeta$, where we henceforth include the unobserved donor characteristic η in z for notational simplicity. In what follows, we treat \mathcal{X} and ζ as finite sets. To simplify notation, albeit with a slight abuse, we replace subscripts that index individuals i and objects j with their types x and z , respectively. For instance, we write the value function as $V_x(t)$ instead of $V(t; x_i)$ and the scoring rule as $s_{xz}(t)$ instead of $s(t; x_i, z_j)$. The notation for other quantities such as π , Γ , and δ is adapted analogously.

Agents follow type-symmetric accept/reject strategies, $\sigma_x : \mathbb{R} \times \mathbb{R}_+ \rightarrow \{0, 1\}$, indexed by $x \in \mathcal{X}$. The first domain element is the payoff of being assigned a particular object, Γ , and the second element is time waited, $t \in \mathbb{R}_+$. We exclude strategies that depend on richer information because beliefs are restricted to satisfy Assumption 3.

We model the queue composition using a single steady state. Specifically, the queue composition is described by a probability density function, m , defined on the set $\mathcal{X} \times [0, T]$, where T is the maximum wait time. This density governs the distribution of agent types and their wait times. We write $m_x(t)$ to denote the density evaluated at (x, t) . The queue length is denoted by N .

DEFINITION 1: A **steady state equilibrium** consists of an accept/reject strategy σ^* , beliefs π^* , a queue size N^* , and a probability measure m^* such that the following conditions hold:

1. **Optimality:** For each agent of type $x \in \mathcal{X}$ and an offer with net present value Γ ,

$$\sigma_x^*(\Gamma, t) = 1\{\Gamma \geq V_x(t; \pi^*)\},$$

where $V_x(t; \pi^*)$ is the net present value to type x of declining the object and following the strategy σ_x^* after t given π^* .

2. **Consistent beliefs:** For each (t, x, z) , the beliefs $\pi^*(t; x, z)$ are consistent with equilibrium offer probabilities. In particular, for mechanisms that use the scoring rule s ,

$$\pi_{xz}^*(t) = H_z^*(s_{xz}(t)) \times \mathbb{P}(c = 1 | x, z),$$

where $H_z^*(s)$ is the probability that the object is available only to agents above the score s if N^* agents are drawn iid from m^* , and they follow strategy σ^* .²⁶

²⁶We do not restrict the queue length N^* to be an integer and draw the $(N^* + 1)$ st agent with probability $N^* - \lfloor N^* \rfloor$.

3. Steady state detailed balance condition: For each $x \in \mathcal{X}$, $m_x^*(t)$ and N^* satisfy

$$\dot{m}_x(t) = -m_x(t)\kappa_x(t) \quad \text{and} \quad m_x(0) = \frac{\gamma_x}{N^*},$$

where γ_x is the arrival rate of type x agents, $\kappa_x(t)$ is the equilibrium departure rate of a type x agent at waiting time t , and m^* is a density: $\sum_{x \in \mathcal{X}} \int_0^T m_x(\tau) d\tau = 1$.

The first condition states that each agent makes optimal decisions at each point in time given their beliefs. The value of declining an offer is given by the Hamilton–Jacobi–Bellman equation defined in Section 3.2. The second condition imposes that agents have correct beliefs about offer probabilities, which are determined by the steady state distribution of cutoff scores H_z^* that arises when agents use strategies σ^* and N^* agents are drawn from a distribution with density m^* . The final condition determines the distribution of waiting agents’ characteristics. The left-hand side is the derivative of the density of type x agents who have waited t years, and the right-hand side is the rate at which such an agent departs multiplied by the density of the type. Departures occur for both exogenous reasons and because agents are removed from the list once they are assigned; that is, $\kappa_x(t)$ is the sum of $\delta_x(t)$ and the rate at which an agent of type x is assigned at time t given the strategy σ^* and the offer probabilities π^* .

This equilibrium concept abstracts away from transitional dynamics in the size and composition of the queue. An alternative approach would be to assume that the queue length and composition follow a Markov process. However, this process is high-dimensional and would make the counterfactual exercises computationally intractable. We expect that our equilibrium concept will be a good approximation to the behavior of such a system. For example, part 3 sets N^* to be the expected queue length. A law of large numbers can be used to show that, in a long queue, the stationary distribution of the queue length concentrates mass on N^* (details available on request).

We compute steady state equilibria using an algorithm that iterates between computing the value function, optimal decisions, and the steady state waitlist composition. A detailed description of the procedure is provided in Section C.1 of the Supplemental Material. To keep the computational burden manageable, the results we present below are based on a type space given by a random sample of 300 patients and 500 donors drawn from our dataset and a discrete grid for time. Our results are not sensitive to varying these parameters.

We prove the existence of a steady state equilibrium for sequential assignment mechanisms that use a priority score in Theorem 1 of the Supplemental Material, Section E. The challenge in showing existence arises because the strategies, beliefs, and the queue composition are functions of time, which is a continuous variable. We use the Brower–Schauder–Tychonoff fixed-point theorem for general Banach spaces to prove existence. The primary assumptions are technical regularity conditions imposing bounds and Lipschitz continuity on primitive objects. The main substantive condition is that the set of scores used in the mechanism is finite. Our results do not rule out multiplicity of equilibria. However, we did not find multiple equilibria for the mechanisms and the main specification described below.²⁷

²⁷We checked for multiplicity by varying starting values, damping the updates in each iteration, and switching the sequence in which value functions and queue composition are updated.

6.2. Welfare Comparisons

Given a mechanism \mathcal{M} and a donor arrival rate λ , the **steady state value** of assignments to an agent of type x is

$$\bar{V}_x^{\mathcal{M}}(\lambda) = \frac{\gamma_x}{\rho} V_x^{\mathcal{M}}(0; \lambda) + \int_0^T N^* m_x^*(\tau) V_x^{\mathcal{M}}(\tau; \lambda) d\tau. \quad (11)$$

The first term represents the discounted value of all type x agents that arrive on the waitlist. The second term represents the value of all agents on the waitlist in steady state. Their sum equals the value of a social planner who considers the net present value of payoffs generated by assignments to type x agents for a system in steady state.²⁸

Let \mathcal{M}^0 be the baseline mechanism used during our sample period and λ^0 be the baseline donor arrival rate. For any mechanism \mathcal{M} , let the equivalent donor arrival rate for type x agents, $\lambda_x(\mathcal{M})$, be the solution to the equation $\bar{V}_x^{\mathcal{M}}(\lambda^0) = \bar{V}_x^{\mathcal{M}^0}(\lambda_x(\mathcal{M}))$. As discussed in Section 5.2, we can express a change in the value function for type x as an equivalent change in the donor arrival rate:

$$EV_x(\mathcal{M}) = \frac{\bar{V}_x^{\mathcal{M}}(\lambda^0) - \bar{V}_x^{\mathcal{M}^0}(\lambda^0)}{\bar{V}_x^{\mathcal{M}^0}(\lambda^0)} \approx \frac{\lambda_x(\mathcal{M}) - \lambda^0}{\lambda^0}.$$

This type-specific measure describes the welfare effects for each patient in terms of an alternative policy that keeps the mechanism fixed but is able to increase (or decrease) organ donation rates. Averaging $EV_x(\mathcal{M})$ across patients equally values an equivalent change in donor arrival rates for different patient types. Our results will also report effects on key subgroups so that the reader may consider alternative aggregates.

7. EVALUATING DESIGN TRADE-OFFS

The U.S. government charged the OPTN Kidney Transplantation Committee with reforming the system used prior to 2014. The stated goals were to find mechanisms that provided equitable outcomes for patients, efficiently allocated available organs, and minimized organ waste. Motivated by these goals, this section begins by comparing previously used mechanisms to benchmark mechanisms, before moving to optimal mechanisms and discussing the sources of gains we identify. Consistent with our sample definition, the results below change the mechanism for all donors offered to patients in NYRT.²⁹

The qualitative results described below are robust to four variations: first, to introducing a limited form of patient unobserved heterogeneity; second, to changing the annual discount factor to 10%; third, to varying the sample and number of patient and donor types; and fourth, to removing the limit on the number of offers that can be made (see the Supplemental Material, Section D). The last exercise suggests that being able to offer organs to more patients does not obviate the improvements we identify below.

²⁸Because this measure considers assignments in steady state, it ignores the payoffs to agents during the transition following a change in the mechanism. Analyzing this process requires modeling agent expectations and decisions during the transition, a task that is beyond the scope of this paper. The bias due to abstracting away from transition dynamics is small if the adjustment to the new steady state is quick.

²⁹Evaluating a nationwide change is burdensome due to the patchwork of variants on points used in approximately half of the states. To simplify this task, we keep the system used to prioritize patients from the rest of the United States fixed to the pre-2014 system. We also assume that acceptance decisions by patients from the rest of the U.S., which govern offers for nonlocal donors to patients in NYRT, remains fixed.

7.1. The 2014 Reform vs. Benchmark Mechanisms

This section compares the mechanisms used prior to 2014 to the redesigned mechanism and two benchmarks from the theoretical literature.

- **Post-2014:** The kidney allocation mechanism adopted in December 2014 awards greater priority to very high CPRA patients and prioritizes healthier patients for higher quality donors. Since high CPRA patients have few transplant opportunities, increasing their priority could yield more equitable outcomes. The second change offers high-quality kidneys to patients likely to benefit from them most.
- **First-Come, First-Served (FCFS):** This mechanism offers objects to agents in the order they joined the waitlist. In addition to being transparent and procedurally fair, FCFS has attractive efficiency properties: Bloch and Cantala (2017) show that it maximizes agent welfare when values for an object are drawn iid across agents. This result is driven by the fact that FCFS incentivizes agents to be selective and only accept objects with high match-specific values. We approximate FCFS by finely discretizing time on a grid t_0, t_1, \dots, t_L and set $s_{xz}(t) = l$ if $t \in [t_{l-1}, t_l)$.
- **Last-Come, First-Served (LCFS):** This mechanism offers objects to agents in reverse order of joining the waitlist. It provides strong incentives to accept an offer because agents are demoted if other agents arrive in the future. Lower selectivity can reduce discards if the number of offers that can be made is limited. Moreover, LCFS maximizes welfare if objects are vertically differentiated in quality and preferences are homogeneous because social welfare depends only on the fraction of objects allocated (Su and Zenios (2004)). We approximate LCFS by finely discretizing time on a grid t_0, t_1, \dots, t_L and set $s_{xz}(t) = L - l$ if $t \in [t_{l-1}, t_l)$.

Our analysis shows that these theoretical and practical recommendations can either increase welfare or reduce organ discard rates, but not both (Figure 3). Both the pre- and

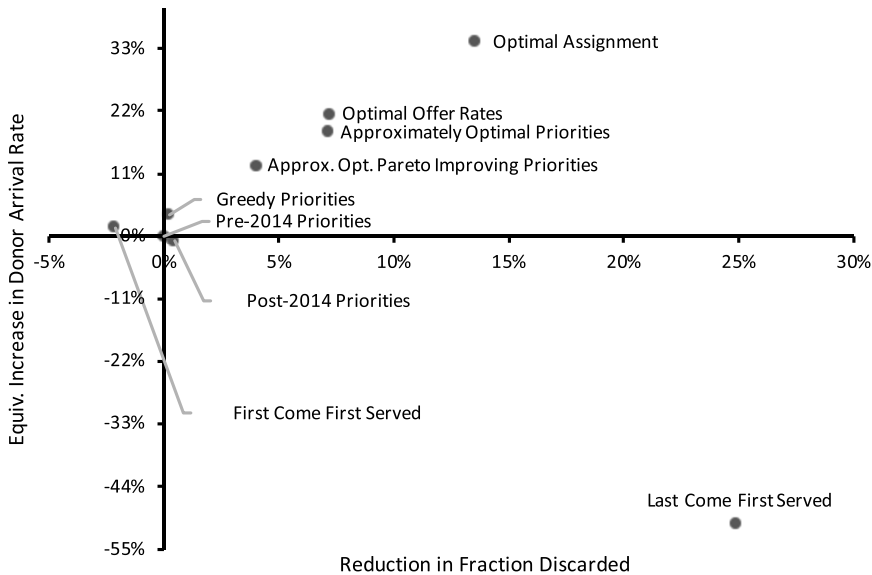


FIGURE 3.—Welfare and organ utilization under alternative mechanisms. Notes: Results are reported relative to pre-2014 priorities and based on each mechanism's steady state equilibrium. The reduction in fraction discarded is defined as the change in the number of kidneys rejected by all NYRT patients to whom the organ was offered.

post-2014 mechanisms are very similar to FCFS in terms of average patient welfare and organ discard rates. Although not reported, our results indicate that the 2014 reforms primarily resulted in redistribution toward younger and more highly sensitized patients. Unlike these three mechanisms, LCFS substantially reduces organ discard rates at the cost of lower welfare.³⁰

The welfare effects of these mechanisms are consistent with predicted effects on the waitlist and the types of donors transplanted. Table VI, Panel A shows that waiting times, queue lengths, and the quality of the average transplanted donor are similar across the first three mechanisms. In contrast, LCFS dramatically decreases discard rates and shortens the steady state queue length from 5113 to 2961. Lower discard rates come at a significant welfare cost because patients accept organs that are poorly matched to them and of low quality (i.e., from donors who are older and less likely to have died from head trauma). Indeed, decomposing the overall welfare changes into the portion predictable using our observable characteristics and the match-specific unobservable characteristics in our model shows that the primary differences are driven by observables.³¹

These empirical results are driven by the externality imposed by an agent's decision to reject an offer. When agents decline an object, they allow others to receive an assignment earlier, generating a positive externality. However, by refusing an object and remaining on the list, they also generate a negative externality as they take away future offers. Which externality dominates depends on preferences and the mechanism. When preferences are highly heterogeneous, the positive externality from being more selective dominates because agents accept only high-value offers while allowing other agents to receive an average-value offer sooner. Therefore, FCFS performs better than LCFS (Bloch and Cantala (2017)). This comparison is reversed when all agents value each object identically because selective agents pass on only lower-quality objects. Our empirical results indicate that horizontal preference heterogeneity is sufficiently strong that the positive externality from rejecting an offer outweighs the negative externality due to possible organ waste.

It is worth noting two details. First, our computed steady state queue length for the pre-2014 mechanism is 5113, which is a little longer than the queue length of 4508 on January 1, 2013 (Table I, panel A). This similarity is remarkable because this moment is not targeted by the CCP approach but nevertheless matches equilibrium computations from our model. It also suggests our sample is close to a steady state.³² Second, while our baseline results change the priority rules for NYRT patients for all donors, we also computed a version in which the change only affected donors recovered in NYRT. The quantitative magnitudes under this alternative are attenuated but directionally similar, with the largest changes for LCFS.

7.2. Optimal Mechanisms

This section designs mechanisms that aim to improve welfare subject to varying constraints. We first consider mechanisms that can choose whom to offer a kidney based on

³⁰This measure ignores transplantation's potential effects on healthcare costs that are not internalized by patients and doctors. A social planner who places weight on these costs may further favor mechanisms that reduce discards.

³¹The observable component includes the portion of welfare effects due to changes in total number of transplants, quality of the average transplanted organ, and the observable component of match-specific value. The unobservable part includes the change due to the term ε_{ijt} .

³²A slightly longer steady state queue should be expected as the NYRT kidney waiting list was growing during our sample period. However, panel B in Table I shows that the growth rate was declining.

TABLE VI
OUTCOMES IN VARIOUS MECHANISMS^a

	Waitlist			Transplanted Donors			EV_x Decomp.			
	EV_x (1)	Queue Length (2)	Reduction in Discard Rate (3)	Years on Waitlist (4)	Age (5)	Head Trauma (6)	Hypertensive (7)	Obs. (8)	Unobs. (9)	$\Delta V_x(0) > -5\%$ (10)
Pre-2014 Priorities	-	5113.2	-	2.73	44.9	16.3%	47.1%	-	-	-
Post-2014 Priorities	-0.8%	5042.5	0.5%	2.70	44.8	16.1%	47.1%	0.3%	-1.1%	97.0%
First-Come, First-Served	1.5%	5274.4	-2.1%	2.78	44.9	16.5%	47.2%	0.2%	1.3%	91.0%
Last-Come, First-Served	-50.6%	2961.4	24.8%	3.73	47.3	14.8%	51.6%	-43.7%	-6.8%	6.3%
Panel A: Steady State Equilibrium, Benchmark Mechanisms										
Optimal Offer Rates	21.4%	4576.2	7.2%	2.55	45.5	15.7%	48.4%	21.6%	-0.3%	68.0%
Approximately Optimal Priorities	18.2%	4555.0	7.2%	2.56	45.5	15.7%	48.4%	20.7%	-2.5%	67.3%
Approx. Opt. Pareto Improving Priorities	12.2%	4738.8	4.0%	2.62	45.3	15.9%	47.7%	16.2%	-4.0%	98.0%
Greedy Priorities	3.8%	5079.9	0.2%	2.78	44.7	16.3%	47.0%	8.9%	-5.1%	79.7%
Optimal Assignment	34.3%	3990.5	13.6%	2.38	45.5	15.6%	48.3%	-	-	94.0%
Panel B: Steady State Equilibrium, Alternative Systems										
Pre-2014 Priorities	-	5791.2	-	3.39	43.8	17.2%	44.6%	-	-	-
Post-2014 Priorities	-	5753.6	0.0%	3.37	43.8	17.2%	44.7%	-	-	-
First-Come, First-Served	-	5853.7	-0.1%	3.33	43.7	17.3%	44.5%	-	-	-
Last-Come, First-Served	-	5668.6	2.7%	5.99	44.1	16.6%	45.3%	-	-	-
Greedy Priorities	-	5718.5	1.0%	3.55	43.9	17.0%	44.9%	-	-	-

^aIn Panels A and B, Columns (1), (3), and (8)-(10) are reported relative to the steady state equilibrium under Pre-2014 Priorities. In Panel C, Reduction in Discard Rate is reported relative to outcomes under Pre-2014 Priorities assuming no equilibrium response. EV_x is defined as $EV_x(\mathcal{M})$ in the main text. Column (10) reports the fraction of patient types whose continuation values at registration fall by less than 5% relative to Pre-2014 priorities.

the characteristics observed in our data. However, the mechanism cannot dictate assignments or condition offers on past actions. Then, in the spirit of searching for a Pareto improvement, we constrain the mechanism to not make any type of patient significantly worse off than they were under the pre-2014 mechanism.

- **Optimal Offer Rates:** We solve for offer rates $\pi_{xz}(t)$ that maximize $\sum_x \frac{\bar{V}_x^\pi(\lambda_0)}{\bar{V}_x^{\mathcal{M}_0}(\lambda_0)}$, where, with a slight abuse of notation, $\bar{V}_x^\pi(\lambda_0)$ is the equilibrium steady state value for type x under offer rates π . The offer rates are subject to feasibility constraints: the steady state assignment rates implied by offer and acceptance rates must not exceed object arrival rates. For each donor type z ,

$$\sum_x \int_0^T N m_x(t) \pi_{xz}(t) \mathbb{P}(\Gamma_{xzt} + \varepsilon > V_x(t) | x, z) dt \leq q_z.$$

The left-hand side is the expected number of objects assigned, and the right-hand side is the number of objects available, where we have ignored the limit on the number of offers that can be made. Our calculations incorporate this limit (see the Supplemental Material, Section C.2). Because this constraint is placed only on expected quantities, the offer rates may not be implementable. The solution provides an upper bound on welfare under any offer mechanism that does not condition on past behavior.

- **Approximately Optimal Priorities:** We can use the solution to the problem above to implement similar equilibrium offer rates with a scoring rule. Specifically, we set $s_{xz}(t)$ to the values of $\pi_{xz}(t)$ that solve the problem above. The Kidney Transplantation Committee exclusively focused on mechanisms that use priority rules because they are transparent and simple to describe. We compute an equilibrium for these priorities using the algorithm outlined in the Online Appendix, Section C.1.
- **Approximately Optimal Pareto Improving Priorities:** We design a priority system that aims to increase average patient welfare without significantly hurting any patient type. We first solve for optimal offer rates as before but add a constraint that no agent type is worse off than under pre-2014 priorities at registration. We then solve for the equilibrium allocation setting $s_{xz}(t)$ equal to the solutions $\pi_{xz}(t)$ to this modified problem. While this mechanism may not yield a strict Pareto improvement, we expect fewer agents will be significantly worse off under these priorities.

We find that previously used mechanisms and both FCFS and LCFS are far from optimal; priorities designed using estimated preferences as inputs can substantially increase welfare while also reducing discard rates. Patient welfare under the approximately optimal priorities mechanism increases by 18.2%, and the discard rate falls by 7.2% (Table VI). Because patients are transplanted at higher rates, they spend less time waiting, and the queue length falls below 4600. This mechanism achieves most of the gains possible under any offer mechanism: optimal offer rates performs only marginally better, increasing patient welfare by 21.4%. While average gains are large, one drawback of approximately optimal priorities is that it makes some patients significantly worse off: nearly one-third of patient types are more than 5% worse off at registration than they were under the pre-2014 mechanism.

A significant fraction of these gains can be achieved while respecting stringent distributional constraints. The approximately optimal Pareto-improving priorities mechanism increases welfare by 12.2% and reduces discards by 4.0%. Recall that because this mechanism approximates the optimal Pareto improving offer rates with a scoring rule, the constraint that no type should be worse off may not be exactly satisfied. Fortunately, 98%

of patients are no more than 5% worse off at registration under our approximation than they were under pre-2014 priorities.

The mechanisms studied above are finely-tuned functions of patient and donor characteristics, which may reduce transparency and robustness to changes in the environment. To assess whether our results depend on very detailed priority scores, we coarsened the priorities into 20 and 50 equally sized bins. This change affected welfare and discard rates by less than five and one percentage points, respectively. While the coarsened priorities also depend on all characteristics in the model, this result points to the promise of finding less finely-tuned mechanisms that can produce large gains. An analysis that formalizes criteria such as transparency and simplicity and incorporates them into the design problem is left for future research.

Finally, the baseline and our alternative mechanisms open the possibility that agents may try to game the system by manipulating some characteristics. Addressing this issue requires analyzing the potential for manipulation and the availability of substitute observables. While this topic is beyond the scope of this paper, we note that all patient and donor characteristics we use to determine priority were used in either the pre-2014 mechanism, the post-2014 mechanism, or in a proposal considered by the Kidney Transplantation Committee.

7.3. *Incentives, Potential Improvements, and Sources of Gains*

7.3.1. *Importance of Incentives*

We now show two ways in which explicitly considering incentives in a waitlist offer mechanism is quantitatively important for finding the improvements described above. First, a naive mechanism that offers organs in order of predicted transplant value yields much smaller improvements. Consider the following mechanism:

- **Greedy Priorities:** Patients are prioritized in order of predicted transplant value, measured in donor supply units. Specifically, $s_{xz}(t)$ is set equal to $EV_{xz}(t)$. The UK deceased donor kidney allocation system now heavily prioritizes predicted transplant benefits (Watson, Johnson, and Mumford (2020)).

This mechanism only marginally improves patient welfare and discard rates relative to the pre-2014 mechanism (Figure 3). The contrast with approximately optimal priorities emphasizes the value of explicitly considering behavioral responses: incorporating rich preference heterogeneity alone is not sufficient.

Second, empirical approaches that ignore incentives mute the differences in predicted equilibrium allocations across mechanisms. Panel C in Table VI shows that predicted discard rates, queue lengths, and donor characteristics are similar across the priority mechanisms considered above when acceptance probabilities do not adjust to the new equilibrium. This finding suggests that the KPSAM module (see SRTR (2015)) used to advise the Kidney Transplantation Committee on allocation reform should be modified to incorporate incentives.

7.3.2. *Comparison With Maximum Possible Gains*

While the mechanisms described above improve both welfare and organ discards, it is unclear whether they achieve most of the possible gains. In particular, we have only considered indirect mechanisms that do not condition offers on past decisions. To analyze the potential scope for further improvements, we calculate an upper bound on possible gains if organs can be assigned by fiat under full information about preferences and rational expectations (but no foresight) about object arrival and agent departure processes:

- **Optimal Assignments:** We maximize steady state average welfare by choosing an assignment rule $a(\varepsilon; x, z, t) \in \{0, 1\}$ that depends on the agent and object types, time waited, and the preference shocks ε for all agents currently waiting. A feasibility constraint ensures that the assignment rate does not exceed the object arrival rate for any type z :

$$\sum_x \int_0^T N m_x(t) \mathbb{P}(c = 1, a = 1 | x, z, t) dt \leq q_z.$$

The term $\mathbb{P}(c = 1, a = 1 | x, z, t)$ is the probability that a type z object is compatible with and assigned to a randomly chosen type x agent who has waited for time t . In addition, we require that the waitlist be in steady state (Definition 1, part 3). A detailed description is provided in Section C.2.1 of the Supplemental Material.

The solution to this problem suggests that approximately optimal priorities achieves over half of the possible gains from redesigning the organ allocation system. Optimal assignments yield welfare that is higher than the pre-2014 outcome by an equivalent of a 34.3% increase in donor supply (Table VI). At the same time, discard rates for patients in NYRT would decrease by 13.6%. However, not all patients are better off even in this benchmark case. Moreover, the optimal assignments benchmark may be far from achievable by a mechanism because it assumes full information and the ability to dictate assignments.

7.3.3. Sources of Welfare Gains

The gains under approximately optimal priorities arise from offering organs to patients with high match-specific values for those organs. These patients are also more likely to accept the offers, reducing discard rates if the limit on the number of possible offers binds and increasing welfare if match quality is not sacrificed. We illustrate this intuition by describing assignments by patient age and dialysis status at the time of registration. As we discussed in Section 5, off-dialysis patients not only value each offer more but also are more likely to accept an offer. Similarly, older patients receive a higher value from each offer.

Approximately optimal priorities increases transplant rates for older patients and re-allocates desirable donors from patients on dialysis to patients off dialysis at registration (Table VII). For example, transplants from young NYRT donors to older off-dialysis patients rise more than four-fold compared to the baseline mechanism (Panel B). Meanwhile, older on-dialysis patients (Panel C) receive fewer kidneys each year from young healthy donors. The distributional constraints on the approximately optimal Pareto improving priorities result in less stark but similar patterns.

8. CONCLUSION

Empirical approaches to evaluating dynamic assignment mechanisms are particularly important because theory has not yet provided clear guidance on their optimal design. This paper provides a new equilibrium empirical framework for analyzing dynamic assignment systems that explicitly incorporates dynamic incentives and applies that framework to deceased donor kidney allocation. Organ allocation system reforms have been assisted by simulation models that do not allow accept/reject decisions to depend on agents' incentives.

We find that deceased donor kidney allocation mechanisms can be significantly improved. Previously used mechanisms and theoretical benchmarks can only reduce organ

TABLE VII
SOURCES OF WELFARE GAINS^a

	Welfare Change		EV_x Decomp.		Transplants by Donor Type				Young NYRT Offer	
	$\Delta V_x(0)$	EV_x	Obs.	Unobs.	Young NYRT	Old NYRT	Non-NYRT	Prob. (%)	Years Waited	
Pre-2014 Priorities (Baseline)	-	-	-	-	28.4	28.2	53.1	4.7	1.61	
Greedy Priorities	-0.7%	-1.2%	0.4%	-1.6%	13.1	30.3	62.4	2.6	6.50	
Approximately Optimal Priorities	9.8%	4.0%	1.5%	2.5%	12.1	30.0	78.6	4.9	3.50	
Approx. Opt. Pareto Improving Priorities	4.6%	2.3%	2.2%	0.1%	15.8	36.8	63.5	3.6	2.66	
Pre-2014 Priorities (Baseline)	-	-	-	-	6.7	38.6	66.5	1.1	7.26	
Greedy Priorities	19.9%	17.7%	38.0%	-20.3%	22.5	42.5	51.1	2.5	6.64	
Approximately Optimal Priorities	105.9%	72.7%	89.2%	-16.5%	33.9	75.8	86.3	4.9	2.63	
Approx. Opt. Pareto Improving Priorities	69.1%	49.9%	68.0%	-18.1%	23.7	54.5	75.6	4.3	2.85	
Pre-2014 Priorities (Baseline)	-	-	-	-	18.2	51.4	94.7	1.7	6.84	
Greedy Priorities	-1.9%	-2.6%	-3.7%	1.1%	13.2	45.2	104.7	1.6	6.50	
Approximately Optimal Priorities	-5.9%	-7.1%	-10.4%	3.3%	4.0	33.2	126.0	0.5	6.49	
Approx. Opt. Pareto Improving Priorities	-0.5%	-2.5%	-3.5%	1.0%	14.8	41.6	113.3	1.5	6.63	
Pre-2014 Priorities (Baseline)	-	-	-	-	13.4	84.0	148.9	1.1	6.88	
Greedy Priorities	2.1%	1.5%	2.5%	-1.0%	17.6	82.0	152.2	1.7	6.69	
Approximately Optimal Priorities	14.9%	7.4%	7.2%	0.2%	13.8	78.7	195.0	1.9	3.59	
Approx. Opt. Pareto Improving Priorities	7.3%	3.0%	3.5%	-0.5%	12.0	79.6	177.2	0.9	4.43	

^aTransplants by Donor Type is an annual rate. Prob. (%) is the probability that a patient receives an offer from a young NYRT donor if the donor is offered to any NYRT patients. Years Waited is the mean years a patient waited across offers received from young NYRT donors. An Old NYRT Donor is over age 35, an Expanded Criteria Donor, or had a cardiac death.

discards at the cost of substantially lowering patient welfare. In contrast, using estimated transplant values, we design mechanisms that both reduce discards and increase patient welfare. Moreover, these goals can be achieved while avoiding significant harm to any type of patient. Considering incentives is essential for designing systems that realize these large gains and for accurately predicting outcomes under alternative mechanisms. These findings also point to empirical work's potential to improve dynamic assignment mechanisms in other contexts.

We make several simplifying assumptions that motivate future research. First, we assume that agent beliefs condition on a limited set of variables and we analyze steady state equilibria. Second, we do not allow for patient-level unobserved heterogeneity. Third, we assume that patient and object arrival rates are exogenous, an assumption that may be particularly important to relax in other applications. Finally, while we primarily focus on organ waste and patient welfare as outcomes, clinical outcomes such as life-years gained from transplantation and graft failure rates are likely important to policy makers.

APPENDIX A: ESTIMATION

A.1. Proof of Proposition 1

Substituting $V_i(t) = \tilde{V}_i(t) + O_i(t)$ and $\Gamma_{ij}(t) = \tilde{\Gamma}_{ij}(t) + O_i(t)$ into equation (2):

$$\begin{aligned} & (\rho + \delta_i(t))(\tilde{V}_i(t) + O_i(t)) \\ &= d_i(t) + \delta_i(t)D_i(t) + \lambda \int \pi_{ij}(t) \int \max\{0, \Gamma_{ij}(t) - V_i(t)\} dG dF + (\dot{\tilde{V}}_i(t) + \dot{O}_i(t)) \\ &= (\rho + \delta_i(t))O_i(t) + \lambda \int \pi_{ij}(t) \int \max\{0, \tilde{\Gamma}_{ij}(t) - \tilde{V}_i(t)\} dG dF + \dot{\tilde{V}}_i(t), \end{aligned}$$

where the second equality follows from the differential equation that defines $O_i(t)$ and $\Gamma_{ij}(t) - V_i(t) = \tilde{\Gamma}_{ij}(t) - \tilde{V}_i(t)$. Since the first term on the right-hand side cancels with the terms that include $O_i(t)$ on the left-hand side, $\tilde{V}_i(t)$ satisfies the necessary differential equation. The boundary condition $V_i(T_i) = O_i(T_i)$ is equivalent to the condition that $\tilde{V}_i(T_i) = 0$.

Next, we show that $V_i(t; \pi) - V_i(t; \pi') = \tilde{V}_i(t; \pi) - \tilde{V}_i(t; \pi')$. Equation (2) implies that

$$\begin{aligned} & V_i(t; \pi) - V_i(t; \pi') \\ &= \lambda \int \pi_{ij}(t) \int \max\{0, \Gamma_{ij}(t) - V_i(t; \pi)\} dG dF \\ &\quad - \lambda \int \pi_{ij}(t) \int \max\{0, \Gamma_{ij}(t) - V_i(t; \pi')\} dG dF + (\dot{V}_i(t; \pi) - \dot{V}_i(t; \pi')). \\ &= \tilde{V}_i(t; \pi) - \tilde{V}_i(t; \pi'), \end{aligned}$$

where the second equality follows from the differential equation satisfied by $\tilde{V}_i(\cdot)$ and because $\tilde{V}_i(\cdot) - \tilde{\Gamma}_{ij}(\cdot) = V_i(\cdot) - \Gamma_{ij}(\cdot)$.

A.2. Details on the Estimator

Gibbs' Sampler

Define $y_{ijt} = V(x_i, t) - \Gamma(x_i, z_j, t) - \eta_j - \epsilon_{ijt} = \chi(x_i, z_j, t)\theta - \eta_j - \epsilon_{ijt}$ and $a_{ijt} = 1\{y_{ijt} < 0\}$. The sampler is initialized at any value θ^0 , σ_η^0 and guesses for η_j^0 and y_{ijt}^0 corresponding to observed decisions such that $y_{ijt}^0 \geq 0$ if and only if agent i rejected object j in period t . The sampler then iterates through the sequence

$$\begin{aligned} & y_{ijt}^{s+1} | \theta^s, \eta_j^s, a_{ijt} \\ & \eta_j^{s+1} | y_j^{s+1}, \theta^s, \sigma_\eta^s \\ & \theta^{s+1} | y^{s+1}, \eta^{s+1} \\ & \sigma_\eta^{s+1} | \eta^{s+1}, \end{aligned} \quad (12)$$

where conditioning on the priors and the observables is implicit, y^s and η^s are vectors with components y_{ijt}^s and η_j^s , and y_j^{s+1} is a vector that stacks y_{ijt}^{s+1} across all i, t . The first two steps involve data augmentation to simplify sampling the model parameters. The prior distribution on σ_η is an inverse-Gamma distribution, and the prior for θ is $N(\bar{\theta}, \Sigma_\theta)$. With these priors, the first step samples from a truncated normal, the second and third steps sample from a normal distribution, and the final step samples from an inverse-Gamma.

Computing the Value Function

Given t , for each patient i , the value of continuing is given by equation (7). Using equation (10), the sample analog of the value of continuing is

$$\hat{V}_i(t) = \lambda \int_t^{T_i} \exp(-\rho(\tau - t)) p(\tau | t; x_i) \hat{W}(x_i, \tau; \hat{\theta}) d\tau.$$

We numerically approximate this integral. First, we rewrite $\hat{V}_i(t)$ as follows:

$$\begin{aligned} \hat{V}_i(t) &= \lambda \int_t^{T_i} \exp(-\rho(\tau - t)) p(\tau | t; x_i) \frac{1}{J} \sum_{j=1}^J 1\{c_{ij} = 1\} 1\{s(\tau; x_i, z_j) > s_j^*\} \psi(\hat{P}_{ij\tau}) d\tau \\ &= \lambda \frac{1}{J} \sum_{j=1}^J 1\{c_{ij} = 1\} \int_t^{T_i} \exp(-\rho(\tau - t)) p(\tau | t; x_i) 1\{s(\tau; x_i, z_j) > s_j^*\} \psi(\hat{P}_{ij\tau}) d\tau \\ &= \lambda \frac{1}{J} \sum_{j=1}^J 1\{c_{ij} = 1\} \int_{\underline{\tau}_{ijt}}^{T_i} \exp(-\rho(\tau - t)) p(\tau | t; x_i) \psi(\hat{P}_{ij\tau}) d\tau, \end{aligned}$$

where $\underline{\tau}_{ijt} = \inf\{\tau > t : s(\tau; x_i, z_j) > s_j^*\}$, with $\underline{\tau}_{ijt} = T_i$ if $s(\tau; x_i, z_j) < s_j^*$ for all $\tau \leq T_i$.

For each i and j , we approximate the integral above using $B = 40$ equally spaced points $q^b = \frac{b}{B+1}$ for $b = 1, \dots, B$ on the unit interval. Let $\tau_{ijt}^b = F^{-1}(q_b; \rho, \underline{\tau}_{ijt}, T_i)$, where $F(\cdot; \rho, \underline{\tau}_{ijt}, T_i)$ is the cumulative distribution function of an exponential random variable with parameter ρ that is truncated between $\underline{\tau}_{ijt}$ and T_i . We approximate the value function as

$$\hat{V}_i(t) = \frac{\lambda}{\rho J} \sum_{j=1}^J 1\{c_{ij} = 1\} \frac{1}{B} \sum_{b=1}^B p(\tau_{ijt}^b | t; x_i) \psi(\hat{P}_{ij\tau_{ijt}^b}).$$

This procedure ensures that there are B points of evaluation for each possible donor and patient-time pair. The numerical performance is superior to an alternative that approximates the integral in equation (7) as a sum over a fixed set of draws because some patient, donor, time combinations may have a very small window of availability, $[\tau_{ijt}, T_i]$.

REFERENCES

- ABDULKADIROGLU, A., N. AGARWAL, AND P. A. PATHAK (2017): "The Welfare Effects of Coordinated School Assignment: Evidence From the NYC High School Match," *American Economic Review*, 107 (12), 3635–3689. [41]
- AGARWAL, N., AND P. SOMAINI (2020): "Revealed Preference Analysis of School Choice Models," *Annual Review of Economics*, 12, 471–501. [41]
- AGARWAL, N., I. ASHLAGI, E. AZEVEDO, C. FEATHERSTONE, AND O. KARADUMAN (2019): "Market Failure in Kidney Exchange," *American Economic Review*, 109 (11), 4026–4070. [40]
- AGARWAL, N., I. ASHLAGI, M. REES, P. SOMAINI, AND D. WALDINGER (2021): "Supplement to 'Equilibrium Allocations Under Alternative Waitlist Designs: Evidence From Deceased Donor Kidneys'," *Econometrica Supplemental Material*, 89, <https://doi.org/10.3982/ECTA17017>. [48]
- AGARWAL, N., I. ASHLAGI, P. SOMAINI, AND D. WALDINGER (2018): "Dynamic Incentives in Waitlist Mechanisms," *American Economic Association: Papers and Proceedings*, 108, 341–347. [38,39,48]
- AGUIRREGABIRIA, V., AND J. SUZUKI (2014): "Identification and Counterfactuals in Dynamic Models of Market Entry and Exit," *Quantitative Marketing and Economics*, 12 (3), 267–304. [39,53]
- ARCIDIACONO, P., AND P. B. ELLICKSON (2011): "Practical Methods for Estimation of Dynamic Discrete Choice Models," *Annual Review of Economics*, 3, 363–394. [56]
- ARCIDIACONO, P., AND R. A. MILLER (2011): "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity," *Econometrica*, 79 (6), 1823–1867. [39]
- ARCIDIACONO, P., P. BAYER, J. R. BLEVINS, AND P. B. ELLICKSON (2016): "Estimation of Dynamic Discrete Choice Models in Continuous Time With an Application to Retail Competition," *The Review of Economic Studies*, 83 (3), 889–931. [39]
- BLOCH, F., AND D. CANTALA (2017): "Dynamic Assignment of Objects to Queuing Agents," *American Economic Journal: Microeconomics*, 9 (1), 88–122. [38–40,66,67]
- DANOVITCH, G. M. (2017): *Handbook of Kidney Transplantation* (Sixth Ed.). Philadelphia, PA, USA: Lippincott Williams & Wilkins. [41,42]
- FERSHTMAN, C., AND A. PAKES (2012): "Dynamic Games With Asymmetric Information: A Framework for Empirical Work," *Quarterly Journal of Economics*, 127 (4), 1611–1661. [39,52]
- GELMAN, A., J. B. CARLIN, H. S. STERN, AND D. B. RUBIN (2014): *Bayesian Data Analysis* (Third Ed.). Boca Raton, FL, USA: Chapman & Hall/CRC. [57]
- HART, A., J. M. SMITH, M. A. SKEANS, S. K. GUSTAFSON, D. E. STEWART, W. S. CHERIKH, J. L. WAINRIGHT, A. KUCHERYAVAYA, M. WOODBURY, J. J. SNYDER, B. L. KASISKE, AND A. K. ISRANI (2017): "OPTN/SRTR 2015 Annual Data Report: Kidney," *American Journal of Transplantation*, 17 (S1), 21–116. [51]
- HELD, P. J., F. MCCORMICK, A. OJO, AND J. P. ROBERTS (2016): "A Cost-Benefit Analysis of Government Compensation of Kidney Donors," *American Journal of Transplantation*, 16 (3), 877–885. [38,41]
- HOPENHAYN, H. A. (1992): "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60 (5), 1127–1150. [39,52]
- HOTZ, V. J., AND R. A. MILLER (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *The Review of Economic Studies*, 60 (3), 497–529. [39,50,55]
- HU, Y., AND S. M. SCHENNACH (2008): "Instrumental Variable Treatment of Nonclassical Measurement Error Models," *Econometrica*, 76 (1), 195–216. [57]
- IRWIN, F. D., A. F. BONAGURA, S. W. CRAWFORD, AND M. FOOTE (2012): "Kidney Paired Donation: A Payer Perspective," *American Journal of Transplantation*, 12 (6), 1388–1391. [38,41]
- ISRANI, A. K., N. SALKOWSKI, S. GUSTAFSON, J. J. SNYDER, J. J. FRIEDEWALD, R. N. FORMICA, X. WANG, E. SHTEYN, W. CHERIKH, D. STEWART, C. J. SAMANA, A. CHUNG, A. HART, AND B. L. KASISKE (2014): "New National Allocation Policy for Deceased Donor Kidneys in the United States and Possible Effect on Patient Outcomes," *Journal of the American Society of Nephrology*, 25 (8), 1842–1848. [42,46]
- KALOUPSIDI, M., P. T. SCOTT, AND E. SOUZA-RODRIGUES (2020): "Identification of Counterfactuals in Dynamic Discrete Choice Models," *Quantitative Economics* (forthcoming). [39,53]
- KRUSELL, P., AND A. A. JR. SMITH (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106 (5), 867–896. [39,52]
- LESHNO, J. (2019): "Dynamic Matching in Overloaded Waiting Lists." [38]

- LIU, T. X., C. YANG, AND Z. WAN (2019): "The Efficiency of a Dynamic Decentralized Two-Sided Matching Market." [41]
- MAGNAC, T., AND D. THESMAR (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70 (2), 801–816. [53,56]
- MCCULLOCH, R., AND P. E. ROSSI (1994): "An Exact Likelihood Analysis of the Multinomial Probit Model," *Journal of Econometrics*, 64 (1–2), 207–240. [57]
- MILLER, R. A. (1984): "Job Matching and Occupational Choice," *Journal of Political Economy*, 92 (6), 1086–1120. [59]
- OPTN, ORGAN PROCUREMENT TRANSPLANTATION NETWORK (2011): "Concepts for Kidney Allocation," Technical Report. [39]
- OPTN, ORGAN PROCUREMENT TRANSPLANTATION NETWORK (2014): "Organ Procurement and Transplantation Network Policies," Technical Report. [42]
- PAKES, A. (1986): "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 54 (4), 755–784. [51,59]
- RAO, P. B. L. S. (1992): *Identifiability in Stochastic Models: Characterization of Probability Distributions*. New Delhi, India: Academic Press, Inc. [57]
- REELING, C. J., AND V. VERDIER (2020): "Welfare Effects of Dynamic Matching: An Empirical Analysis." [41]
- ROTH, A. E., T. SONMEZ, AND M. U. ÜNVER (2004): "Kidney Exchange," *The Quarterly Journal of Economics*, 119 (2), 457–488. [40]
- ROTH, A. E., T. SÖNMEZ, AND M. U. ÜNVER (2007): "Efficient Kidney Exchange: Coincidence of Wants in Markets With Compatibility-Based Preferences," *American Economic Review*, 97 (3), 828–851. [40]
- RUST, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55 (5), 999–1033. [50,51,59]
- SCIENTIFIC REGISTRY OF TRANSPLANT RECIPIENTS SRTR (2015): "Kidney-Pancreas Simulated Allocation Model," Technical Report. [40,70]
- SU, X., AND S. A. ZENIOS (2004): "Patient Choice in Kidney Allocation: The Role of the Queueing Discipline," *Manufacturing & Service Operations Management*, 6 (4), 280–301. [38,40,66]
- USRDS, UNITED STATES RENAL DATA SYSTEM (2018): *States Renal Data System, 2018 USRDS Annual Data Report: Epidemiology of Kidney Disease in the United States, National Institutes of Health*. Bethesda, MD: National Institute of Diabetes and Digestive and Kidney Diseases. [41]
- VAN DER VAART, A. W. (2000): *Asymptotic Statistics*. Cambridge, UK: Cambridge University Press. [57]
- WALDINGER, D. (2020): "Targeting In-Kind Transfers Through Market Design: A Revealed Preference Analysis of Public Housing Allocation." [41]
- WATSON, C. J. E., R. J. JOHNSON, AND L. MUMFORD (2020): "Overview of the Evolution of the UK Kidney Allocation Schemes," *Current Transplantation Reports*, 7, 140–144. [70]
- WEINTRAUB, G., C. L. BENKARD, AND B. VAN ROY (2008): "Markov Perfect Industry Dynamics With Many Firms," *Econometrica*, 76 (6), 1375–1411. [39,52]
- WOLFE, R. A., V. B. ASHBY, E. L. MILFORD, A. O. OJO, R. E. ETTENGER, L. Y. AGODOA, P. J. HELD, AND F. K. PORT (1999): "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. [38,41]
- WOLFE, R. A., K. P. MCCULLOUGH, D. E. SCHAUBEL, J. D. KALBFLEISCH, S. MURRAY, M. D. STEGALL, AND A. B. LEICHTMAN (2008): "Calculating Life Years From Transplant (LYFT): Methods for Kidney and Kidney-Pancreas Candidates," *American Journal of Transplantation*, 8 (4p2), 997–1011. [59]
- WOLPIN, K. I. (1984): "An Estimable Dynamic Stochastic Model of Fertility and Child Mortality." *Journal of Political Economy*, 92 (5), 852–874. [59]
- ZENIOS, S. A. (2004): "Models for Kidney Allocation," in *Operations Research and Health Care*. International Series in Operations Research & Management Science, Vol. 70, ed. by M. L. Brandeau, F. Sainfort, and W. P. Pierskalla. Boston, MA: Springer, 537–554. [40]
- ZHANG, J. (2010): "The Sound of Silence: Observational Learning in the US Kidney Market," *Marketing Science*, 29 (2), 315–335. [40]

Co-editor Aviv Nevo handled this manuscript.

Manuscript received 28 January, 2019; final version accepted 4 September, 2020; available online 10 September, 2020.