NBER WORKING PAPER SERIES

AN EMPIRICAL FRAMEWORK FOR SEQUENTIAL ASSIGNMENT: THE ALLOCATION OF DECEASED DONOR KIDNEYS

Nikhil Agarwal
Itai Ashlagi
Michael A. Rees
Paulo J. Somaini
Daniel C. Waldinger

Working Paper 25607
http://www.nber.org/papers/w25607

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2019

We are grateful to Alex Garza, Sarah Taranto, and Jennifer Wainwright of UNOS for facilitating access to the data. We thank Peter Arcidiacono, Lanier Benkard, Liran Einav, Gautam Gowrisankaran, Jacob Leshno, Anna Mikusheva, Michael Ostrovsky, Ariel Pakes, Parag Pathak, Peter Reiss, Andy Skrzypacz, Gabriel Weintraub, Mike Whinston for helpful discussions. Layne Kirshon and Khizar Qureshi provided competent research assistance. The authors acknowledge support from the National Science Foundation (SES-1729090, SES-1254768), the National Institutes of Health (R21-DK113626), 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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

Nikhil Agarwal, Itai Ashlagi, Michael A. Rees, Paulo J. Somaini, and Daniel C. Waldinger

NBER Working Paper No. 25607
February 2019
JEL No. C51,D47,I10

ABSTRACT

An organ transplant can improve a patient's life while also realizing substantial savings in healthcare expenditures. Like many other scarce public resources, organs from deceased donors are rationed to patients on a waitlist via a sequential offer mechanism. The theoretical trade-offs in designing these rationing systems are not well understood and depend on agent preferences. This paper establishes an empirical framework for analyzing waitlist systems and applies it to study the allocation of deceased donor kidneys. We model the decision to accept an organ or wait for a more preferable organ as an optimal stopping problem, and develop techniques to compute equilibria of counterfactual mechanisms. Our estimates show that while some types of kidneys are desirable for all patients, there is substantial match-specific heterogeneity in values. We then evaluate alternative mechanisms by comparing their effect 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 as 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 a 14.2 percent increase in donor supply.

Nikhil Agarwal
Department of Economics, E52-460
MIT
50 Memorial Drive
Cambridge, MA 02142
and NBER
agarwaln@mit.edu

Paulo J. Somaini
Stanford Graduate School of Business
Knight Management Center
655 Knight Way E372D
Stanford, CA 94305
and NBER
soma@stanford.edu

Itai Ashlagi
Management Science and Engineering
Huang Engineering Center, 262
Stanford University
475 Via Ortega
Stanford, CA 94305
iashlagi@stanford.edu

Daniel C. Waldinger
Furman Center, New York University
Wilf Hall, Room 218
139 Macdougal Street
New York, NY 10012
danielwaldinger@nyu.edu

Michael A. Rees
University of Toledo
Department of Urology
3000 Arlington Avenue
Mail Stop 1091
Toledo, OH 43614
michael.rees2@utoledo.edu

A data appendix is available at http://www.nber.org/data-appendix/w25607

1 Introduction

As of January 7, 2019, there were 94,971 patients on the kidney waiting list in the United States, while only 13,483 deceased donor transplants were performed in 2018.[1] Each transplant improves the expected quality and length of a transplanted patient's life while saving hundreds of thousands of dollars in expected dialysis costs (Wolfe et al., 1999; Irwin et al., 2012; Held et al., 2016). Yet, approximately 20 percent of medically suitable organs extracted for transplantation are discarded. The allocation of deceased donor kidneys does not use money because of ethical considerations and legal restrictions,[2] making traditional price-based market-clearing mechanisms infeasible. Instead, available kidneys are allocated through a centralized waitlist. Given these constraints, it is essential to find mechanisms that efficiently allocate resources, minimize waste, and achieve equitable outcomes.[3] Similar considerations motivate the use of waitlist systems to allocate other deceased donor organs, public housing, long-term care, child-care, and child-adoption.

Previous research and guidance on waitlist design either ignores dynamic incentives or assumes specific forms of market primitives. Theoretical approaches to designing dynamic assignment mechanisms have found that even qualitative trade-offs are sensitive to the primitives of the market.[4] 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,[5] do not allow decisions to respond to changes in the system. Moreover, recent empirical advances in analyzing allocation systems are restricted to static choice settings.

This paper develops an empirical framework for analyzing waitlist mechanisms that sequentially assign objects to forward-looking agents, and applies these methods to study the deceased donor kidney allocation system in the U.S. We make several methodological and empirical contributions. First, we develop a method for estimating agent preferences using typical administrative data, and apply it to the kidney waitlist data from New York to es-

[1] Source: https://optn.transplant.hrsa.gov/data/view-data-reports/national-data/

[2] The National Organ Transplantation Act (NOTA) makes it illegal to obtain human organs for transplantation by compensating donors.

[3] These goals are articulated by the Organ Procurement and Transplantation Network (OPTN), a contractor for the Health Resources & Services Administration (HRSA), in their policy document titled "Concepts for Kidney Allocation" (OPTN, 2011). A committee that was charged with reforming the allocation system adopted a new mechanism in 2014. We discuss these reforms in greater detail below.

[4] Agarwal et al. (2018b) compare the results in Su and Zenios (2004), Leshno (2017), Arnosti and Shi (2017), and Bloch and Cantala (2017) and show by example that optimal design depends on the nature of preference.

[5] For example, Kidney Pancreas Simulated Acceptance Module (KPSAM) used by the kidney allocation committee to evaluate various proposed mechanisms prior to the reforms enacted in 2014 assumes that acceptance decisions on the kidney waitlist do not depend on mechanism used, thereby ignoring differences in dynamic incentives generated by various mechanisms. Similar methods are used by the organ allocation agencies in the U.K., Scandinavia and France.

timate payoffs from various types of transplants. This step is based on an optimal stopping problem faced by a patient when she is offered a kidney. Second, we define a notion of steady state equilibrium that is amenable to computation and counterfactual analysis of a broad class of mechanisms. Finally, we use these techniques to compare alternative mechanisms on key outcome measures such as efficiency, equity, and organ waste.

Our empirical results imply that the new kidney allocation system implemented in 2014 leaves significant room for improvement on both average patient welfare and discard rates. This system, which was designed with the assistance of simulations that ignored dynamic incentives, mostly resulted in redistribution relative to the prior mechanism. In fact, our results suggest that these simulations yield biased and attenuated results, highlighting the importance of dynamic incentives and their equilibrium effects. We use our estimates to derive a mechanism that can generate large welfare gains while also reducing waste.

Our empirical application studies deceased donor kidneys allocated in the New York City area (henceforth, NYRT) between 2010 and 2013. The allocation mechanism used to match deceased donor kidneys with patients relies on a coarse point system based on donor and patient characteristics and the patient's waiting time. As soon as an organ becomes available, it is offered to patients on the waitlist in decreasing order of these priority points. The decision of whether or not to accept an offer remains with the patient and the transplant surgeon.[6] Each organ is allocated to the highest-priority biologically compatible patient who accepts it. The patient is removed from the waitlist once she is transplanted. Otherwise, she may remain on the waitlist and may choose to accept the next organ she is offered. The priority system does not depend on whether a patient has refused previous offers. Even though the timing and quality of future offers are uncertain, it can be optimal to turn down an offer to wait for a more suitable one.

We begin by documenting several key descriptive facts using rich administrative data on patient and donor characteristics and acceptance decisions. Patients on the NYRT waitlist face extreme scarcity. While 1,400 patients join the waitlist each year, fewer than 400 deceased donor kidneys are recovered in NYRT. These donors vary widely in quality; some are accepted immediately, while others are rejected by every patient and discarded. The chance a patient is high on the list for a particular type of organ, and therefore the chance of being offered desirable organs, increases with waiting time. As a result, patients have an incentive to reject offers and wait for a better kidney. Indeed, Agarwal et al. (2018b) document that acceptance rates are higher for patients who are less likely to receive offers. Therefore, consistent with dynamic considerations, patients with a higher option value of waiting are more likely to refuse an offer for an organ.

Motivated by these descriptive facts, we model an agent's decision to accept an offer as a

[6] We refer to the decision-maker as the patient because our data does not directly identify cases in which a surgeon makes a decision on behalf of her patient. This approach is reasonable if each surgeon acts in the best interest of each of her patients.

continuous-time optimal stopping problem. She accepts the current offer if the value from the object is higher than the expected value of continuing to wait. The distribution of potential 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. Our technique adapts methods for inverting conditional choice probabilities in this continuous-time problem (Hotz and Miller, 1993; Arcidiacono and Miller, 2011; Arcidiacono et al., 2016) to suit dynamic assignment mechanisms. This approach eases the computation of the continuation values relative to full solution methods (Pakes, 1986; Rust, 1987, for example) because the distribution of future offers depends on all characteristics that influence priorities in the mechanism, and is therefore high-dimensional.

The estimated values from transplants are intuitive. All patients value certain characteristics that are correlated with organ quality – for instance, younger donors are preferred by all patients, as are immunologically similar matches. We also estimate substantial donor-level unobserved heterogeneity; including this feature significantly improves the fit of our estimated choice probabilities. In addition, there is significant match-specific heterogeneity in values. For example, we find that older patients place less value on younger donors as compared to younger patients. This match-specific preference heterogeneity creates scope for the design of the allocation mechanism to improve match quality. Moreover, our descriptive results show that the current mechanism produces significant mismatch on several dimensions. For instance, young patients are often allocated organs from older donors, while many old patients receive organs from young donors. Motivated by these facts, we conclude the paper by analyzing equilibrium allocations of the pre- and post-2014 U.S. kidney allocation systems, benchmark mechanisms from the theoretical literature, and welfare-maximizing mechanisms that exploit the rich heterogeneity across donors and patients.

Predicting assignments in these counterfactual mechanisms requires us to solve two technical issues. First, we need to formulate an equilibrium notion that is both tractable and consistent with our estimation procedure. Computing counterfactuals is challenging because it involves solving a dynamic game with many players. An unrestricted state-space could include the composition of the entire waitlist, resulting in a system with an extremely high-dimensional stationary distribution. To make progress, we develop a notion of a steady-state equilibrium in the spirit of previous approaches to simplifying this task (Hopenhayn, 1992; Weintraub et al., 2008; Fershtman and Pakes, 2012). Our approach computes a steady-state distribution of types and a steady-state queue length. This allows us to find a tractable computational algorithm that iterates between solving the value function using backwards induction and calculating the steady-state composition of the queue by forward simulation.

Second, we need to ensure that assignments under counterfactuals of interest are indeed identified. In dynamic models such as ours, counterfactuals may not be invariant to normalizing the payoff of an action because they arbitrarily restrict payoffs across states (see Aguirre-

gabiria and Suzuki, 2014; Kalouptsidi et al., 2015). Our estimates normalize the payoff of never receiving an assignment to zero. We formally show that our normalization is appropriate for the mechanism design counterfactuals we consider if the value of declining all offers remains fixed. In our empirical context, this assumption is satisfied if the value of remaining on dialysis until death, and the value and opportunity of receiving a living-donor transplant do not change when the deceased donor allocation system is redesigned. We argue that both these assumptions are reasonable.

Our framework allows us to compare, for each patient, the welfare effect of a change in a mechanism to an equivalent change in deceased donor supply (arrival rates) under the existing mechanism. This quantity has the advantage of being interpretable and independent of an arbitrary choice of units during estimation. We then aggregate these equivalent changes in donor supply across patients as a summary of the welfare effects.

We start by comparing assignments generated by mechanisms that are used in practice to two benchmarks recommended for the kidney allocation system in the theoretical literature. Specifically, we compare the pre- and post-2014 assignment systems with the first come first served (FCFS) (Bloch and Cantala, 2017) and last come first served (LCFS) (Su and Zenios, 2004) waitlists. These benchmarks encode notions of procedural fairness and/or perform well under specific forms of agent preferences.

The efficiency properties of various mechanisms turn on whether selective agents generate a net positive or negative externality on others. When an agent declines an object, she allows others to receive an assignment earlier, generating a positive externality. However, by refusing an object and remaining on the list, she generates a negative externality as she takes away future offers. Which force is more important depends on preferences and the mechanism. Bloch and Cantala (2017) show that FCFS induces agents to be selective and produces better assignments when preferences are highly heterogeneous. Selective agents only accept high idiosyncratic match value objects but allow others to receive an offer for a typical object earlier instead of later, generating a positive externality. However, FCFS performs poorly when all agents value each object identically because selective agents pass on only low quality objects onto those lower on the list. This fact creates strong incentives to wait, resulting in organ waste. In fact, Su and Zenios (2004) show that in this environment LCFS improves social surplus by reducing waste.[7] It does so by providing incentives to accept offers quickly as waiting results in reduction in priority and therefore lower quality offers in the future.

Our results show that previously used mechanisms and commonly studied theoretical benchmarks can yield either high average patient welfare or low discards, but not both. The reforms in 2014 resulted primarily in redistribution from older patients to younger patients, with little improvement in the welfare of the average patient. In fact, both the pre- and

[7] Observe that when all agents place an identical value on each object, social surplus is a function only of waste.

post-2014 mechanisms yield patient welfare and organ discard rates within 1.9 percent of the benchmark first come first served mechanism. Last come first served is able to dramatically reduce organ discard rates (27 percent), but at the cost of lowering patient welfare by 34 percent due to poor match quality.

To systematically find reallocation gains beyond those afforded by these previous recommendations, we calculate optimal assignments and optimal offer rates. In the former, the designer has full information about agent preferences and can force agents to accept organs; it is therefore a theoretical bound on the maximum welfare that can be achieved under any mechanism. In the latter, the designer uses only observable characteristics and cannot force agents to accept offers. We use the solution to this problem to design an alternative to the current mechanism. We also try a greedy approach to improving efficiency by offering organs to patients that are predicted to benefit from them the most.

Our results show significant scope for improving assignments. The theoretical bound on the maximum possible gains from improved assignments is equivalent to a 28.1 percent increase in donor supply. While substantial, implementing this assignment requires full information about the value of an organ to all currently registered patients and the ability to dictate allocations. Even with full information, it may be politically infeasible to achieve these gains in practice because such a mechanism would make take-it-or-leave-it organ offers to patients with irreversible organ failure.

We then design a waitlist mechanism aimed at increasing patient welfare. To do this, we solve for offer rates as a function of agent and object observables that maximize patient welfare, but cannot dictate acceptance behavior. Using this solution, we design a scoring mechanism can achieve an increase in the average patient's welfare equivalent to a 14.2 percent increase in donor supply. This increase is approximately half of the potential gains from any possible allocation reform. This mechanism also increases transplant rates by 7.8 percent, and as a result, equilibrium queue lengths and waiting times are much shorter than under the pre-2014 mechanism.

We then consider mechanisms that aim to achieve welfare gains without substantially hurting any patient type. Specifically, we solve of Pareto improving offer rates that maximize patient welfare subject to making no patient type worse off. We then approximate this mechanism using a scoring rule. The importance of the status quo and distributional constraint in this exercise substantially influences the solution. Patient welfare increases by an equivalent of only 9.1 percent.

Our solutions are able to reduce discards and increase patient values by simultaneously considering heterogeneity in patients' transplant values and dynamic incentives. They reallocate offers to patients who not only have high values for specific organs, but also are very likely to accept them. This allows the mechanism to increase the quantity of transplants without lowering match value. Explicitly considering incentives is crucial for these gains. A mechanism which naively prioritizes patients based on predicted transplant value only marginally

improves on the pre-2014 mechanism, increasing patient welfare and transplant rates by 1.2 and 0.3 percent, respectively.

These results point to the significant scope of using our empirical framework for improving dynamic assignment mechanisms. Previously used empirical approaches for reforming organ allocation systems are unable to identify these gains. Moreover, a naive comparison that ignores incentives and holds choice probabilities fixed, as done in KPSAM (SRTR, 2015), predicts much smaller differences between these mechanisms and therefore understates the trade-off between match quality and organ discards. This finding suggests that recommendations need to account for equilibrium responses, be specific to the primitives of the market, and incorporate the desired objectives of the policy maker.

Related Literature

The design of scoring rules for organ allocation has been an active area of research in the Medical and Operations Research communities (Zenios et al., 2000; Su et al., 2004; Zenios, 2004; Kong et al., 2010; Su and Zenios, 2006; Bertsimas et al., 2013). However, this previous research, as well as the KPSAM acceptance module (see SRTR, 2015) used by the Scientific Registry of Transplant Recipients (SRTR) to simulate the effects of various allocation systems, does not empirically model patients' dynamic incentives to accept or reject an organ offer. As a result, the current approach taken by SRTR assumes that acceptance decisions do not depend on the waitlist mechanism. Empirical evidence in Agarwal et al. (2018b) suggests that ignoring dynamic incentives may result in biased predictions.

Within economics, there is a large body of theoretical work (see Roth et al., 2004, 2007, for example) and a more recent empirical literature (see Agarwal et al., 2018a) studying the design of living donor kidney exchange markets. Despite the growth in the living donor markets, deceased donors continue to be the primary source of kidneys for transplantation, enabling approximately 70 percent of all kidney transplants in the United States.[8]

Our paper is related to Zhang (2010), which uses a dynamic model to study how patients learn about the quality of an organ. It shows that patients lower on the list are more likely to refuse an organ if patients that are higher have refused it. The paper argues that this pattern is most consistent with a parametric model of observational learning. Our approach abstracts away from learning,[9] but allows for unobserved donor heterogeneity to capture correlation in acceptance behavior. We do this to focus on allocation issues and equilibrium responses when simulating changes to the offer system.

The methods in this paper contribute to the growing literature on empirical approaches for analyzing centralized assignment systems (see Agarwal, 2015; Fack et al., 2015; Abdulka-

[8] In 2017, organs from deceased donors accounted for 14,038 out of 19,849 kidney transplants. Source: https://optn.transplant.hrsa.gov/data/view-data-reports/national-data/.

[9] Zhang (2010) uses data from 2002. Anecdotal evidence suggests that the information available to patients about donors was dramatically better and standardized during our sample period (2010-2013). This fact significantly reduces the scope for observational learning.

diroglu et al., 2017; Agarwal and Somaini, 2018, for example). These previous approaches have focused on static assignment mechanisms. To our knowledge, the only exceptions are Waldinger (2017) and Reeling and Verdier (2018). Waldinger (2017) estimates a model of public housing choice in which agents face a two-stage decision with a portfolio choice problem in the first stage.[10] Reeling and Verdier (2018) analyze a dynamic allocation mechanism in which agents can enter a lottery for a good each period, and apply it to study the allocation of bear hunting licenses. In contrast, our methods pertain to an optimal stopping rule that differs from these settings.

This distinction between static and dynamic assignment systems is important because the theory of static allocation systems, e.g. mechanism design approaches to school choice (Abdulkadiroglu and Sönmez, 2003), is comparatively well-developed. Abdulkadiroglu et al. (2017) show that, at least in New York City, there is little difference between various well-coordinated school choice systems. In contrast, Leshno (2017), Bloch and Cantala (2017), Arnosti and Shi (2017), and Su and Zenios (2004) arrive at different conclusions about which sequential offer system performs the best. Their results depend on the nature of preference heterogeneity. Therefore, estimating these primitives is essential when designing dynamic allocation mechanisms.

Our work builds on the estimation of dynamic discrete choice models (Wolpin, 1984; Pakes, 1986; Rust, 1987; Hotz and Miller, 1993), particularly recent developments in continuous-time versions of these models (Arcidiacono et al., 2016), and the estimation of dynamic games (Bajari et al., 2007; Pakes et al., 2007; Aguirregabiria and Mira, 2007; Pesendorfer and Schmidt-Dengler, 2008). Additionally, we employ a model of beliefs and an equilibrium notion that resembles concepts aimed at making the analysis of dynamic games tractable (Hopenhayn, 1992; Weintraub et al., 2008; Fershtman and Pakes, 2012). We discuss the relationship to this literature when we develop our approach.

Overview

Section 2 describes the kidney allocation system, the available data and documents key facts from the market. Sections 3 and 4 model the optimal stopping problem from the perspective of each agent and detail our estimation methods. Section 5 describes our parameter estimates. Section 6 defines a steady-state equilibrium and summarizes results on existence; presents our approach to welfare comparisons; and outlines the algorithm for computing an equilibrium. Section 7 describes the mechanisms we compare and presents the results. Section 8 concludes.

[10] This work is also related to a literature that estimates preferences for public housing to answer questions about how to design an allocation mechanism (see Geyer and Sieg, 2013; Sieg and Yoon, 2016; van Ommeren et al., 2016; Thakral, 2016). A key difference is that these approaches are based only on final assignments instead of detailed choice data.

2 Background, Data and Descriptive Evidence

This section starts with the basics of kidney transplation before describing the allocation system. Next, we detail our data sources and the information available for the study. Finally, we present key descriptive facts to motivate our model and empirical exercises.

2.1 Basics of Kidney Transplantation

As of December 31, 2016, there were 726,331 cases of End-Stage Renal Disease (ESRD) in the United States (USRDS, 2018). Medicare provides near universal coverage to ESRD patients; unlike traditional Medicare that is age-specific, the Medicare ESRD program is the only disease-specific entitlement program. This program cost the federal government $35.4 billion in 2016, accounting for 7.2 percent of overall Medicare paid claims (USRDS, 2018). This program therefore accounts for approximately 1 percent of the federal budget. Transplantation is the best treatment for ESRD, as it improves both the quality and length of life while saving an estimated $270,000 over the life of a transplanted patient (Wolfe et al., 1999; Irwin et al., 2012; Held et al., 2016; USRDS, 2018).

A kidney from a deceased donor is considered transplantable to a patient if they are biologically compatible. A donor's organ is considered incompatible if the patient has a pre-existing immune response to proteins on the organ's cells. A biologically incompatible patient's immune system will recognize and attack the transplanted organ, resulting in graft failure.[11] Following transplantation, medications allow transplant physicians to limit new immune responses to foreign protein types, but pre-existing immune responses lead to immediate loss of the transplanted organ if not avoided. The specific form of incompatibility is not important for the purposes of this study.

[11] The immune system tags foreign objects (antigens) with antigen-specific antibodies so that white blood cells (leukocytes) can defend against them. Each donor has blood-type antigens and up to 6 specific types of human leukoctye antigen (HLA) proteins out of a set of hundreds of possible types that are relevant for kidney transplantation. Some patients have pre-existing antibodies to some subset of these antigens. A transplant recipient's immune system will immediately attack the donor's kidney and reject the organ if the recipient has an antibody to any one of the donor antigens. A recipient is tissue-type compatible with a donor's kidney if she has no pre-existing antibodies corresponding to the donor's antigens Danovitch (2009), even if the donor and the recipient do not have the same antigens. A transplant between certain incompatible patient-donor pairs has become possible due to development of desensitization technologies (see Orandi et al., 2014), but compatible transplants are preferred. Following transplantation, the immune system will attack any foreign antigen if such an attack is not attenuated with immunosuppressive medications. Pre-existing donor-specific antibodies cause immediate rejection of a transplanted organ and cannot be prevented using immunosuppression. Thus, transplant physicians measure pre-existing immune responses and avoid them, whereas they prevent future immune responses by treating a transplanted patient with immunosuppressive medications that prevent an immune response to newly-identified foreign antigens. Also for this reason, an organ from a donor that has some antigens in common with the recipient is more desirable because the patient is less likely to develop a new immune response to the organ after transplantation.

Because of biological compatibility, the transplantation possibilities available to patients differ based on their immune systems. In the United States, kidneys from deceased donors are largely allocated to patients of the same blood type to maintain a balance in transplantation opportunities across blood types. However, some patients have immune systems that react to a broader range of tissues, even from donors with the same blood type. These patients have 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 she is tissue-type incompatible. This measure is calculated from blood tests conducted to determine the set of human protein types (antigens) to which a patient has had a pre-existing immune response.

The benefits of a transplantation can substantially differ, even conditional on biological compatibility. Organs from donors that are younger, died of brain death under controlled circumstances, and did not suffer from diabetes, strokes or hypertension usually function longer. Measured kidney function, time from removal of the organ to anticipated transplantation, size and weight match between the patient and the donor is also considered important. A donor with tissue proteins that are similar to the patient's is considered preferable because the likelihood of the patient developing an immune response post-transplantation is reduced. In some cases, the donor may have an infectious disease that the patient is at risk of getting if she proceeds with a transplant. There are a number of other factors that influence the medical benefits from specific organs and to specific patients. We refer the reader to Danovitch (2009) for further details about kidney biology.

2.2 The Allocation of Deceased Donor Kidneys

Assignment of organs from a potential donor begins after death is declared (brain or cardiac) and necessary consent for donation has been obtained. The Organ Procurement Organization (OPO) in the donor's area obtains information about the donor from tests and the donor's medical history. This information is entered into a system, called UNet, that is used to coordinate across transplant centers. The OPO staff use UNet to determine the order in which patients will be offered each of the donor's organs, to transmit information about the donor to the transplant centers, and to record accept/reject decisions. OPO staff usually contact the surgeons for several potential recipients simultaneously to solicit their decisions. This process can take place while the donor is on life-support and before the potential donor's organs have been extracted in order to maintain organ viability. Once a kidney has been recovered from the donor, transplant surgeons or patients that were potentially interested in receiving that kidney may decline based on any new information discovered during medical testing or the physical examination of the kidney. These final decisions need to be made without delay, usually within an hour. A final compatibility blood test is then conducted using samples from the donor and all patients that have accepted an organ. The donor's

kidneys are then allocated to the highest priority patients on the waitlist that were willing to accept the organs.

When offering organs, UNet first excludes patients who are not biologically compatible with the donor. This exclusion is based on a detailed patient immunological profile that is submitted when the patient is registered on the waitlist. Specifically, the allocation system requires patients to list unacceptable donor antigens, i.e. donor protein types with which the patient's immune system is likely to react. The allocation system runs a "virtual crossmatch" with these data, which we mimic in our analysis.

Next, UNet screens out patients that have listed pre-set exclusion criteria within the system. These criteria allow patients to automatically exclude kidneys that are transplantable but undesirable because of donor characteristics such as age, health conditions, and kidney function. UNet then orders the remaining set of patients first by priority type, and then within priority type by priority points. Finally, it breaks ties in order of time waited.

The priority and points system is motivated by both equity and efficiency concerns. The system in place during our sample period (2010 to 2013) first offers kidneys to patients with a perfect tissue-type match, then to patients from the local OPO in which the organs were recovered, then regionally, and finally nationally. Within each priority group, the points system is based on tissue type similarity, whether or not the patient is pediatric, patient sensitization, and waiting time. The detailed priority system is described in policy section 8 in OPTN (2014).

New kidney allocation rules were implemented on December 4, 2014. This new system gives greater priority to the healthiest patients for the most desirable donors, increases priority for extremely hard to match patients, and reduces emphasis on wait time. Israni et al. (2014) discusses this system and the rationale for the changes. We refer the reader to OPTN (2017) for a detailed description of the priorities and points used.

2.3 Data Sources

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 US, submitted by the members of the Organ Procurement and Transplantation Network (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. For tractability, we restrict attention to data on the kidney waitlist and the acceptance decisions of all patients in the New York Organ Donor Network (NYRT) between January 1st, 2010 and December 31st, 2013.[12] NYRT is the largest donor service area (DSA),

[12]We end our sample in 2013 to rule out anticipatory effects and to avoid modeling transition dynamics as agents ancitipate the new system introduced in December 2014. Reports from the United Network for Organ

in terms of number of patients, that used the standard allocation rules in the United States prior to 2014.[13]

The primary dataset on the waitlist, the Potential Transplant Recipient (PTR) dataset, contains the offers made and accept/reject decisions. This dataset is drawn from the records generated by UNet, which is the backbone software system used to coordinate offers and decisions. In addition, we obtained detailed information on patient and donor characteristics from the Standard Transplantation Analysis and Research (STAR) dataset. Fields in the STAR dataset are populated based on information gathered in UNet as well as forms submitted by transplant centers after a transplant is performed.

2.4 Descriptive Analysis

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 a very large number of 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 1 describes our patient sample. A total of 9,917 patients were registered with NYRT at some point between 2010 and 2013. Panel A shows the state of the waitlist on January 1st of each year in our sample and summarizes a subset of important patient characteristics. Our dataset includes rich information on patient health status, including indicators of patient health (e.g. body mass index, age, total serum albumin), and medical history (e.g. diabetes, years on dialysis). The average patient on the list has waited for a little over two years, with the average waiting time increasing over time. Recall that CPRA measures the probability of tissue-type mismatch with a randomly chosen donor. The average CPRA is about 12 percent, which indicates that there is more than a one-in-ten chance that a patient is tissue-type incompatible with a randomly chosen donor. The standard deviation is large because there are many patients with extremely low or extremely high CPRA. The allocation mechanism awards priority and points to high CPRA patients because of equity concerns.

Sharing (UNOS) that track transplantation rates after the adoption of the new system show the existence of short-term transition dynamics (termed "bolus-effects" in these reports) immediately following the reform (Wilk et al., 2017).

[13] See Hart et al. (2017) for a map of DSAs in the United States. As mentioned above, except in cases of perfect tissue-type matching, allocation takes place based on geography, with DSAs constituting the smallest unit. A little less than half of DSAs used rules that were different from the baseline rules. We identified the DSAs that use non-standard rules via a special request for administrative documentation on the various rules in use.

Panel B describes the patients that enter and exit the waitlist during our sample. The list is growing; the number of new patients joining exceeds the number leaving. Panel C shows the reasons for departure from the list and the characteristics of patients by the stated reason for departure. The most common reason for depature is receiving a deceased donor transplant. The average patient has waited for 3.08 years before receiving a deceased donor. The second most common reason is that the patient either dies or becomes "too sick to transplant." These patients are 61.5 years old on average, compared to 54.3 years for patients who received deceased donor transplants. The third most common departure reason is receiving a live 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.

Table 2 summarizes the rich set of donor characteristics used in our study. These include donor age, cause of death, relevant medical history (diabetes, hypertension), and the leading indicator of donor kidney function (donor creatinine). Panel A presents the statistics for donors recovered within NYRT during our sample period. Just under 200 donors are recovered from the NYRT area each year, which is only one-seventh of the number of patients joining the waitlist in NYRT. Therefore, there is extreme scarcity in organ supply within the region. The refusal rate remains high despite this scarcity. Across donors, the mean number of biologically compatible offers that met the pre-set screening criteria is over 430, but the median is much lower, at 27. This skewed distribution arises because undesirable kidneys are rejected by many patients, while desirable kidneys are accepted quickly. Indeed, over 20% of donors have at least one of their viable kidneys discarded. Organs from these donors were refused by an average of almost 1,500 patients. The mean and median amongst the group of donors for which both of their viable kidneys were accepted are 123.1 and 15 respectively.[14] Our observable donor covariates which should predict organ quality – including age, cause of death, and donor creatinine – are correlated with discards in the expected ways. The last two sets of columns show that a number of donors recovered in NYRT were ultimately offered to patients elsewhere. For these donors, only about 0.3 kidneys were transplanted per donor, indicating that most of these kidneys were discarded.

In addition to donors recovered within the local area, NYRT patients are also offered donors from other parts of the country. Indeed, panel B shows that a total of 1,470 donors were offered to patients registered with NYRT in the average year. Because most of these donors were recovered elsewhere in the country but offered to NYRT patients after a large number of refusals, these donors are likely to be undesirable. It is therefore not surprising that they see a very large number of offers and high discard rates. Again, the relatively poor quality of these donors is captured by our observable characteristics. Compared to donors recovered in NYRT, the average non-NYRT donor offered to NYRT patients is older, less likely to

[14]The number of transplanted kidneys amongst donors with no discards is less than two because some donors have only one viable kidney for donation.

have died of head trauma, more likely to be diabetic or hypertensive, more likely to be an undesirable donor (ECD) or be donating after cardiac death (DCD), and more likely to have a high creatinine level. The average donor offered to NYRT patients ultimately donates only 0.75 kidneys.

Waitlist, Offers and Acceptance Rates

We now describe the offer and acceptance rates in the data. Throughout the paper, we only consider blood and tissue type compatible offers. As a reminder, we refer to the decision-maker as the patient for simplicity of exposition, assuming that any decisions made by surgeons are in the best interest of each individual patient. Overall, patients receive many offers and reject most of them. This is because desirable kidneys are accepted quickly, while less desirable kidneys are offered to many patients before being accepted or discarded. A patient's likelihood of receiving a high-quality offer rises as her waiting time increases, creating a strong incentive for patients to wait for attractive offers.

While position on the list increases with waiting time, the kidney allocation system significantly differs from a first come first served waitlist. Figure 1a plots waiting time against position on the list for offers where the donor met the patient's screening criteria. The horizontal axis is the (donor-specific) position of a patient, and the vertical axis is the average waiting time for patients at that position. Mean waiting time falls quickly as we go down the list. However, the system is not well approximated by a first come first served queue. We calculated the fraction of times 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 percent. It would be 100 percent in a first come first served system, and 50 percent in a system that randomly ordered patients for each donor. This summary shows the importance of the points and priorities other than for waiting time in determining position on the list.

Acceptance rates suggest that patients who are higher on the list have access to higher quality donors. Figure 1b shows the fraction of offers accepted by the position of the patient. The acceptance rate in the first few positions is much higher than later positions, but still only between 15 and 20 percent. Only about 1 percent of offers are accepted in lower positions, and this rate falls even farther below the top 20 positions. This sharp decline near the top occurs for two reasons. First, for some donors, the first few offers are made to patients with a perfect tissue-type match. These offers have a high acceptance rate because perfect tissue-type matches are rare and extremely valuable. The second reason, which is the predominant one, is that lower positions see a negatively selected set of kidneys because desirable kidneys are likely to be accepted near the top of the list. Both reasons give patients a strong incentive to reject offers and wait to receive a higher-quality donor in the future.

Because lower-quality kidneys are offered to more patients, patients receive many offers, but the overall acceptance rate, as a fraction of offers, is extremely low. Table 3 describes

these overall patterns in offer and acceptance rates. Panel A considers all feasible offers, including offers that did not meet the patient's pre-set criteria. A typical patient receives about 220 offers per year, but only 0.15 percent of offers are accepted. When interpreting these low acceptance rates, it is important to remember that the majority of offers are from very undesirable donors. Offers from desirable donors are more likely to be accepted: kidneys recovered in NYRT are accepted five times more often, and 10.6 percent of offers of a perfect tissue-type match are accepted. The panel also shows that patients with sensitized immune systems – that is, patients with CPRA above 80 percent – receive many fewer offers from transplantable donors even though the system gives them higher priority, and are more likely to accept. Panels B and C restrict to offers that are likely to be more attractive. Panel B studies offers that met patients' pre-set criteria. The typical patient still gets such an offer every three or four days. However, a typical patient only receives about two offers per month from donors recovered within NYRT. Panel C restricts to offers in the first 10 positions, which are even more likely to be for organs from desirable donors. These offers are rare – the typical patient can expect to receive less than one such offer each year – but they are very likely to be accepted.

Taken together, these statistics suggest that the supply of desirable donors in NYRT is scarce and that patients have to wait several years before receiving a transplant from a desirable donor. Moreover, our dataset contains rich information predictive of the likelihood that a donor is refused.

Evidence on Mismatch

Table 4 provides suggestive evidence of mismatch between donors and transplanted patients. Panel A reports the rates at which patients receive any transplant. Pediatric patients are very likely to be transplanted, either with a deceased donor kidney or through a living donor. The priority given to these patients is likely an important contributing factor. Adult patients are less fortunate, but interestingly, among adults there is no signficant gradient in transplant probability with age. In contrast, the chance of receiving a live donor does fall with age. Panel B describes transplanted donors by age for those patients who receive a kidney through the deceased donor waitlist. Pediatric patients are much more likely to receive a transplant from a young donor as compared to older patients. Although there is some assortative matching by age, signs of age mismatch remain. Many patients above the age of 65 continue to receive kidneys from young adults and middle-aged donors. One concern in interpreting these numbers is that a kidney transplanted to an older patient may be undesirable for other reasons. Panel C focuses on a subset of donors with no clear medically undesirable characteristics such as diabetes, cardiac death, high creatinine levels or hepatitis C. The qualitative patterns of age mismatch persist. These patterns motivated some of the 2014 OPTN reforms, which attempted to match healthier patients (typically young adults) to donors whose organs are predicted to last longest.

Another patient characteristic that will be important in the improvements we ultimately identify is whether the patient had already begun dialysis when they joined the waitlist. Some patients with end-stage renal disease who qualify for the waitlist still have marginal kidney function and can avoid dialysis. These patients are relatively healthy compared to patients already on dialysis at registration, but they are still likely to need dialysis in the future. The last two columns of Table 4 show that outcomes differ substantially for patients on and off dialysis. While patients not yet on dialysis at registration represent 32 percent of our patient sample, they receive only 27 percent of deceased donor transplants. In contrast, they are much more likely to receive a living donor transplants (Panel A): 17 percent of off-dialysis patients received a live donor transplant during our sample, while only 7 percent of on-dialysis patients did. Consistent with being healthier, dialysis patients are much less likely to depart the waitlist because they died or were too sick to transplant. Despite these differences, donors are nearly equally distributed by age among on- and off-dialysis patients (Panel B), though patients not on dialysis at registration are a bit more likely to receive a very young or very old donor.

Evidence on 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. (2018b) present descriptive evidence consistent with dynamic incentives using data from all areas of the United States. They find that highly sensitized patients who are immunologically compatible with fewer donors – and who can therefore expect to receive fewer offers in the future – are more likely than less sensitized patients to accept an offer of a given quality. We replicated their research strategy using data from patients registered in NYRT and found similar patterns. We briefly describe the strategy and results below.

The ideal experiment would compare two identical populations of patients that face different option values for exogeneous reasons. However, we are not aware of such variation in the context of kidney allocations. Instead, Agarwal et al. (2018b) use the likelihood that a patient is biologically compatibile with a randomly chosen donor, as measured by the Calculated Panel Reactive Antibody (CPRA), to study how variation in option values affect acceptance decisions. A patient that is likely to be biologically compatible with a large number of donors should have a high option value of waiting, and therefore be more selective.

We replicate their findings for NYRT and show that, as predicted by the presence of dynamic incentives, CPRA is negatively correlated with offers for compatible organs and positively correlated with acceptance rates (Figures B.1a and B.1b in the Appendix). This pattern is robust to rich controls for patient priority and indicators of the value of an offer, for example, patient and donor characteristics, match characteristics, interactions of CPRA with tissue-

type similarity (Table B.1 in the Appendix).

The main concern is that immune sensitization also influences the value from a transplant. Patients develop sensitive immune systems primarily through blood transfusions and prior transplants. Therefore, these patients are more likely to be frail, making a transplant risky. This risk is less justified unless the donated organ is of high quality. By this argument, sensitized patients should be less likely to accept offers. The fact that we find the opposite correlation strengthens our argument for the importance of dynamic incentives.

3 A Model of Decisions in a Waitlist

This section presents a model of agents' decisions in a waitlist mechanism that will form the basis of our empirical strategy. We begin by defining a class of sequential assignment mechanisms and the primitives governing agents' decisions while on the waitlist. Agents and objects arrive according to exogenous processes, and each object is offered to agents on the waitlist in order of an agent-object-specific priority score until the object is accepted by an agent or rejected by everyone. We then provide assumptions on agents' payoffs and beliefs, as well as the evolution of the state space, which lead to a tractable optimal stopping problem from the agent's perspective. Though the model is motivated by the structure of our application, it may be useful in other settings in which items are offered sequentially to agents, including other organ allocation settings.

3.1 Notation and Preliminaries

Consider a sequential assignment mechanism in which objects indexed by j are offered to agents indexed by i waiting on a list. Let x_i denote observed characteristics of agent i; let z_j and η_j respectively denote observed and unobserved characteristics of object j; and let t_i denote the amount of time the agent has been waiting on the list. We assume that agents observe both η_j and z_j. The model does not include unobserved agent heterogeneity. We discuss this restriction in Section 4.2.

Objects may be incompatible with some agents. Let $c_{ij} = 1$ if object j is compatible with agent i, and 0 otherwise. Incompatibility can arise due to biological reasons in the organ allocation context but they may arise due to other restrictions (e.g. legal) in other contexts.

Time is continuous. Objects and agents arrive at poisson rates λ and γ, respectively. The characteristics of each arriving agent x are independent and identically distributed (iid). Similarly, each object's characteristics (z, η) are drawn iid from the cumulative distribution function (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.

3.2 Mechanisms and Primitives

3.2.1 Mechanisms

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

- 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. For example, ties could be broken either uniformly at random or by waiting time.[15]
- 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. A mechanism does not make offers to agents that are known to be incompatible with the object.
- Step 3 (Assignment): The object(s) are allocated to agents with the highest q_j priorities for whom $a_{ij} = 1$, where q_j is the number of copies of the object available. An object cannot be allocated to an incompatible agent.
- Step 4 (Arrivals and Departures): An agent is removed from the waitlist once an object has been assigned to her. Other agents may exogenously join or leave the list.

Within the set of general offer-based waitlist systems, the primary restriction is on the order in which offers are made. Specifically, we assume that an agent's priority does not depend on the other agents in the market. This restriction is adequate for estimation using the deceased donor kidney allocation system in place during our sample period. Moreover, such mechanisms are a natural class to consider because they are simple and transparent to implement. Indeed, all deceased-donor organ allocation mechanisms as well as systems considered by the kidney allocation committee during their deliberations prior to the 2014 reform were offer mechanisms based on priority scores.[16] In counterfactual analysis, we will compare assignments that result from various mechanisms that obey this structure to benchmark optimal assignments.

Typical administrative datasets from such assignment systems contain information on all characteristics used to determine the priority score because the characteristics are used to make offers. This allows a researcher to calculate the order in which any object would be offered. Our empirical exercises required us to develop computer code for this purpose, and

[15] If ties are broken by t_i, it must be that no two agents have the same value. Since time is continuous and agent arrivals are governed by a poisson process, simultaneous arrivals will be zero-probability events, and t_i strictly orders any two patients with probability one.

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

we were able to verify the output of our code using administrative records of the offers that were made during our sample period.

One complication in our setting is that organs must be allocated within a certain time frame that depends on the condition of the organ and various logistical factors. Limited manpower at the Organ Procurement Organization (OPO) can limit the number of patients that can be contacted and offered the organ. We treat the maximum number of offers that can be made for each object as exogeneous.

3.2.2 Payoffs

There are three types of primitive payoffs in the model. The first is the (expected) net present value of agent i being assigned a compatible object j after waiting t periods, denoted $\Gamma_{ij}(t)$. In our application, this term captures the value placed by patients and surgeons on transplants from various types of organs. The second is the expected net-present value from from departure without an assigment, $D_i(t)$. In our application, departures occur due to living donor transplants, death, or transfers to other listing centers (Table 1). We view $D_i(t)$ as incorporating any of those reasons.[17] Finally, agents incur an expected flow payoff while waiting on the list, $d_i(t)$. In our application, $d_i(t)$ is best interpreted as the payoff from living without a functioning kidney, which includes dialysis for most patients.

Two economic implications of 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 and the flow costs of remaining on dialysis. Second, we assume that agents only value their own outcomes and not those of others. This assumption is commonly made in theoretical and empirical work on assignment mechanisms (e.g. Abdulkadiroglu and Sönmez, 2003; Abdulkadiroglu et al., 2017). This restriction may be violated if surgeons value the outcomes of other patients, especially those that they might be treating. NYRT has a total of 10 transplant hospitals staffed with many more kidney transplant surgeons. This limits common agency problems that surgeons might face. The payoffs can be interpreted as accruing to the patients if surgeons act in the best interest of each of their patients.

Our empirical framework makes the following assumptions on these payoffs:

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

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

[17] We can represent the value from a departure as a weighted average over the value from the various events, i.e. $D_i(t) = \sum_k p_{ik}(t) D_{ik}(t)$ where k denotes the type of depature (e.g. obtaining a live donor, death etc.) and $p_{ik}(t)$ is the probability of each type of depature conditional on a departure occuring. The formulation is agnostic about the sources of these payoffs. For example, the net present value of death can include any bequest motives.

(ii) The random variables ε_{ijt} are iid with a known, non-atomic distribution with cdf G.

(iii) The expected flow payoffs from waiting $d_i(t)$ and the expected payoff from departing without an assignment $D_i(t)$ depend only on (x_i, t).

Restrictions on ε_{ijt} imposed in Assumptions 1(i) and 1(ii) are common in the dynamic discrete choice literature. They will allow us to use an approach based on an inversion technique due to Hotz and Miller (1993). The comparison with other methods and specific functional form assumptions on G, $\Gamma(\cdot)$, and the distribution of η_j are discussed in Section 4 below.

Assumption 1 implies that the model excludes agent-level unobserved heterogeneity. We discuss analytical challenges in relaxing this assumption once we have laid out our estimation approach in Section 4.2.

3.2.3 Arrivals and Departures

Agents arrive stochastically, and may depart the list prior to assigment. We make the following assumption on the arrival and departure processes:

Assumption 2. *(i) Departures prior to assignment and arrivals are governed by Poisson processes that are independent of the waitlist composition and design.*

(ii) The departure rate of agent i 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 100-th birthday. T_i therefore corresponds to the waiting time for a patient on the day she turns 100 years of age.[18]

The primary economic restriction for our purposes is that departures prior to assignment and arrivals do not depend on the design of the kidney waitlist. Table 1 shows that the most common reason for departure without a deceased donor transplant is death or patients becoming "too sick to transplant." It seems safe to assume that these events are not responsive to the design of the kidney waitlist. The second most common reason is receiving a living donor transplant. Departures due to this reason are exogeneous if the design of the kidney 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 are medically superior to deceased donors

[18] It is straightforward to extend the framework to allow for agents that could remain on the list forever, $T_i = \infty$. This generalization will primarily change computational techniques and require that the value function for each patient approaches a constant. We restrict our attention to the finite time-horizon case for simplicity of exposition.

and produce better transplant outcomes in terms of patient and graft survival.[19] Finally, a minority of patients depart for other reasons, in most cases for undisclosed reasons or because they move residences.

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 kidney function is sufficiently poor.[20] Therefore, it is in a patient's interest to join the waitlist as soon as possible. This feature of the priority system motivates our assumption that arrivals do not depend on the state of the waitlist or the allocation mechanism. In counterfactuals, we consider priority systems that do not change how waiting time is calculated.

3.3 Individual Agent's Problem

Agents on the waitlist who receive an offer of an object must decide whether to accept it or wait for a future offer. This results in an optimal stopping problem from the perspective of the agent (Pakes, 1986; Rust, 1987). We follow a common estimation strategy in dynamic games by considering an agent's optimal decision rule taking the distribution of actions of other agents as observed in the data (Pakes et al., 2007; Bajari et al., 2007). Solving for counterfactuals will require a notion of equilibrium, which we discuss in Section 6. This section starts by describing a general formulation of this single-agent problem before making simplifying restrictions.

3.3.1 Beliefs

To make an informed decision about whether to accept an offer, an agent must form beliefs about the organs she may be able to obtain in the future if she declines the current offer. Recall that the kidney waitlist offers organs to patients in sequence of their priority scores

[19] Living donor superiority is partly driven by the higher medical quality of living donor kidneys, and also by the fact that living donation allows for a better planned transplant. For example, patients receiving a living donor transplant can proactively start immunotherapy. OPTN and SRTR (2011) report the rate of late graft failure for transplanted patients by donor type (living or deceased). This measures the time at which half of the transplanted patients are alive with the kidney still functioning. The rate of late graft failure for adult patients transplanted in 1991 was 10.1 years for deceased donor kidneys, compared to 15.8 years for living donor kidneys. Some of this difference may be due to selection into who receives each type of transplant. Hart et al. (2017) compare outcomes following living and deceased donation by patient age and primary diagnosis using the chances of graft failure 10 years after transplantation. This statistic for adults transplanted with a deceased donor kidney in 2005 is 52.8%, whereas it is only 37.3% for those that received a living donor. They report that 5year patient survival differences for living donor and deceased donor recipients are large even when broken down by patient age or primary diagnosis (compare figures KI 79 and KI 80 with figures KI 82 and KI 83 from Hart et al. (2017)).

[20] This feature is shared for all DSAs, including NYRT, that used standard allocation rules during our sample period. A patient was qualified for registration once they had begun dialysis or had a glomerular filtration rate (GFR) below 20mL per minute.

as long as the kidney remains viable for transplantation. The organs are allocated to the highest priority agents that are compatible with and accept the organ. At the end of the allocation process, each organ j effectively has a cutoff priority s_j^*, such that only an agent with priority at least s_j^* would have received the organ if she accepted it. This score depends on the decisions of all agents on the list at the time, their compatibility, and the number of offers that can be made for the object. Agent i can expect to obtain a compatible kidney as long as her score, s_{ijt}, exceeds s_j^*. Therefore, it is sufficient for an agent to form beliefs over the probability distribution of s_j^* in order to decide which organs are likely obtainable in the future.

The beliefs about the distribution of the cutoffs priority s_j^* depend on the quality of the organ and the information an agent may have about the competitive environment. Let

$$H\left(s;\mathcal{F}_{i,t},z_j,\eta_j\right)=\mathbb{P}\left(s_j^*<s\middle|\,\mathcal{F}_{i,t},z_j,\eta_j\right) \tag{2}$$

denote the belief for the CDF of s_j^* given the information set $\mathcal{F}_{i,t}$ and the organ characteristics z_j and η_j. Therefore, if agents know the scoring rule, then agent i believes that the probability of receiving an offer for organ j is given by $H\left(s\left(t_i;x_i,z_j\right);\mathcal{F}_{i,t},z_j,\eta_j\right)$.

In principle, the information set can include the history of offers previously received (and rejected) by the agent as well as identities of the other agents on the waitlist at a given point in time. However, there are several reasons, discussed below, why beliefs are unlikely to be sensitive to such information. We therefore assume that an agent's beliefs des not depend on such fine information:

Assumption 3. *Each agent i's believes that the probability that an object with characteristics (z_j,η_j) is compatible and will be available to her after a waiting time of t is*

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

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

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. The primary threat to this restriction is that some surgeons may be treating multiple patients on the kidney waitlist or may learn about other patients from their colleagues. This concern is limited by the fact that the NYRT area has many transplant hospitals and surgeons. Second, it abstracts away from inference about the likelihood of receiving future offers based on past offers. This restriction is motivated by the institutional features and empirical observations discussed below in Section 3.3.4. Finally, it assumes that the probability that 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.

Our assumption on beliefs is a reasonable approximation if assessments about which organs are likely to be available are based primarily on a surgeon's extensive experience treating patients. It also eases the analysis relative to beliefs of the form in equation (2) because it dramatically reduces the dimension of the information set, and therefore the state space, in the dynamic problem. It is well known that this curse of dimensionality can complicate analysis and estimation in dynamic models (see Pakes and McGuire, 2001, for example).

The simplification of the state space in Assumption 3 is similar in spirit to equilibrium concepts that have been introduced to make the estimation of dynamic games tractable. These concepts simplify the state space by abstracting away from aggregate uncertainty in the composition of competitor types (Hopenhayn, 1992; Weintraub et al., 2008) or by modeling beliefs as being based on past experience (Fershtman and Pakes, 2012). Despite these simplifications, the state space continues to be quite rich because, in addition to aspects that influence payoffs, it contains all characteristics that influence priorities or determine whether or not any given object j is compatible.

3.3.2 Value functions

We assume that agents make optimal accept/reject decisions by comparing the net present value of an object to the value of waiting. Holding the strategies of other agents fixed, she 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:

$$(\rho + \delta_i(t)) V_i(t) = d_i(t) + \delta_i(t) D_i(t) + \lambda \int \pi_{ij}(t) \mathbb{E} \max\{0, \Gamma_{ij}(t) - V_i(t)\} \mathrm{d}F + \dot{V}_i(t), \quad (3)$$

where the operator $\mathbb{E}$ takes expectations 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.

This expression can be derived by considering an agent's value of waiting at time t for an infinitesimal duration Δt. In the event that no object arrives during this period, the agent incurs flow payoffs from dialysis $d_i(t)\Delta t$ and may depart exogenously with probability $\delta_i(t)\Delta t$, incurring a payoff of $D_i(t)$ if such a departure offers. In the event that an object arrives using this period, which occurs with probability $\lambda \Delta t$, the object's characteristics are drawn from the CDF F. The integral calculates the expected increment in the agent's value function for 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 $\mathbb{E} \max\{0, \Gamma_{ij}(t) - V_i(t)\}$ from an offer for object j. In the limit as $\Delta t \to 0$, the probability that both departures and

object arrivals occur within the interval Δt tends to zero, yielding the differential equation above.[21]

Equation (1) shows an intuitive result that values depend on the flow payoffs while waiting on the list, the possibility of and value from exogeneous departures, and the option value of potential offers. The differential equation defining $V_i(t)$ has a unique solution that is determined by the terminal condition $V_i(T_i) = D_i(T_i)$ because the probability of receiving additional offers in the remaining time vanishes as $t \to T_i$.

3.3.3 Normalization and Simplifying the Value Function

A typical dataset from a sequential assigment mechanism such as ours only contains information about accept/reject decisions. As is well understood, data on actions alone may not be sufficient for identifying 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 Kalouptsidi et al. (2015) point out that such normalizations may not be innocuous because they arbitrarily restrict payoffs from specific actions across various states. These restrictions can affect the validity of answers to certain counterfactual predictions. This fact poses a potentially serious problem for empirical analysis if one is interested in answering questions that depend on primitives that are not identified from choice data.

Fortunately, the counterfactuals involving changes in the mechanism are identified in our model. Intuitively, in any waitlist mechanism, the trade-offs between accepting an offer and waiting should only depend on payoffs relative to the value of never receiving an assignment. Assumptions 1(iii) and 2(i) together imply that the value of never receiving an assignment does not depend on the mechanism. This discussion suggests normalizing the value of refusing all offers, irrespective of the state.

[21] Specifically, the discretized version of the Hamilton-Jacobi-Bellman Equation defining the value of waiting at time t is:

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

where $\lambda_i(t) = \lambda \int \pi_{ij}(t)\,\mathrm{d}F$ 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 term is the flow payoff from remaining on dialysis. The second term is the expected probability of departure multiplied by its value. The third term represents the value for a kidney arriving. The fourth term denotes the value of waiting in period $t + \Delta t$ in the case when no offer arrives and an exogenous departure does not occur. 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 $\Delta t \to 0$ under mild continuity conditions yields the differential equation above.

Formally, the value of refusing all offers, and therefore never being assigned, 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 $O(T_i) = D_i(T_i)$. Under Assumptions 1(iii) and 2(i) this value does not depend on the waitlist, and therefore future offer probabilities $\pi_{ij}(t)$. Measuring $V_i(t)$ and $\Gamma_{ij}(t)$ relative to $O_i(t)$ suffices for analyzing decisions and differences in welfare under the current and alternative mechanisms. Appendix B.1 formally shows that this is the case.

With this in mind, we normalize the net present value of refusing all offers in the future, $O_i(t)$, to zero at all t. This normalization implies that $d_i(t) + \delta_i(t) D_i(t) = 0$ for all t and that $D_i(T_i) = 0$. Equation (3) now simplifies to

$$(\rho + \delta_i(t)) V_i(t) = \lambda \int \pi_{ij}(t) \mathbb{E} \max\{0, \Gamma_{ij}(t) - V_i(t)\} \, \mathrm{d}F + \dot{V}_i(t) \tag{4}$$

As can be seen, the advantage of this particular normalization is that the set of primitives that need to be estimated is greatly reduced. We no longer need to estimate the flow payoffs from remaining on the list or the net present value of departing. Going forward, we interpret $\Gamma_{ij}(t)$ and $V_i(t)$ as values relative to never receiving an assignment.

The solution to differential equation above is

$$V(t; x_i) = \int_t^{T_i} \exp(-\rho(\tau - t)) \, p(\tau | t; x_i) \left(\lambda \int \pi(\tau; x_i, z, \eta) \, \tilde{\psi}(\tau, x_i, z, \eta) \, \mathrm{d}F_{z,\eta} \right) \mathrm{d}\tau, \tag{5}$$

where

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

is the probability that agent i does not exogenously depart before τ conditional on being on the list at t and

$$\tilde{\psi}(\tau, x_i, z, \eta) \equiv \mathbb{E} \max\{0, \Gamma(\tau, x_i, z, \eta) + \varepsilon_{ijt} - V(\tau; x_i)\}$$

is the incremental value to agent i of receiving an offer of an object with characteristics (z, η) at time τ, with expectations taken over ε_{ijt}. We have explicitly reintroduced agent and object characteristics into the notation because this equation will form the basis of our empirical strategy. This solution is based on the boundary condition $\lim_{t \to T_i} V(t; x_i) = O(T_i; x_i) = 0$ because the probability of receiving an offer after t vanishes as $t \to T_i$. This equation also clarifies that the model and approach is readily extended to the case when T_i is infinite. To do this, one would replace the condition that $V_i(T_i) = 0$ with the condition $\lim_{t \to \infty} V_i(t) = 0$.

3.3.4 Descriptive Evidence and Institutional Support for Assumption 3

In addition to the tractability provided by Assumption 3, the institutional and empirical features of our setting make it an appealing model of beliefs and acceptance behavior. To begin with, there are several reasons why recent history should have limited predictive power for future offers. First, it is unrealistic that surgeons have detailed information about patients on the waitlist that are not under their care. Privacy concerns preclude surgeons from obtaining such information. Second, the set and order of patients on the waitlist varies significantly across donors, limiting the ability to predict future offers based on recent experience. We calculated the fraction of times any two patients are prioritized in the same order for two randomly chosen donors. We find that the priorities do not preserve the order of patients 18.5% of the time. A random order would place this number at 50%. This calculation focuses only on cases where both patients are compatible with the two donors. Including incompatibility would indicate even less persistence. Last, but not least, we directly test for autocorrelation in priority-score cutoffs s_j^* across organs ordered by the date on which they arrived. One would expect non-zero serial correlation across these cutoffs if the offers a patient observed contained information about the likelihood of receiving future offers. We were not able reject the null hypothesis of zero rank autocorrelation across a range of partitions of donors (Table B.2 in the Appendix).[22] Taken together, the evidence suggests that recent offers are not particularly predictive of future cutoffs.

We also directly test whether recent offer history predicts current acceptance behavior, and fail to find evidence that it does. Suppose that a patient sees a string of unexpectedly frequent or high-quality offers. If the patient updates her beliefs based on recent history, she should infer that she faces relatively little competition from other patients on the waitlist and can expect a better offer set in the near future. As a result, the patient should become more selective – less likely to accept an organ of a particular quality – if her recent offer history is better. In contrast, if the patient does not update her beliefs based on recent history, recent offers should have no predictive power for current acceptance behavior.

We test this hypothesis by constructing an offer-specific variable for the number of years since the patient's most recent offer. We include this variable as an additional predictor in the acceptance regressions presented in Section 2. Under the null hypothesis of no updating, the coefficient on time since last offer should be zero. Under the alternative of updating based on recent history, we expect a positive coefficient because patients who have waited a long time since their last offer expect lower continuation values, and should therefore be more likely to accept an offer of a given quality.

Table 5 presents coefficient estimates from several specifications that include measures of the patient's recent offer history. The first two columns include time since last offer as a predictor in the standard acceptance regressions. Without controlling for current offer

[22] In fact, the p-values of the test statistic across relatively fine partitions are close to uniformly distributed.

characteristics, the estimated coefficient on time since last offer is positive and statistically significant; however, with additional controls the coefficient becomes much smaller and statistically insignificant. This finding suggests that recent offer rates do not predict current acceptance behavior.

Of course, time since last offer may not fully capture the value of a patient's recent offers, leading to a lack of power or omitted variables bias.[23] For example, beliefs may be formed not only based on the most recent offer, but the last several, and a test based on the time since the previous offer may be underpowered. Columns (3) and (4) show that when time since last offer is averaged across a patient's two and five most recent offers, the coefficient estimate on average time since last offer remains positive but statistically insignificant. There may also be other reasons why time since last offer is not the most relevant proxy for the value of a patient's recent offers. Column (5) includes controls for donor characteristics of the previous offer. This has almost no impact, and the coefficient estimates on donor characteristics are statistically insignificant. Finally, it is possible that our test does not consider the relevant set of offers. Column (6) restricts the analysis to offers from ideal donors, and column (7) restricts to donors recovered in NYRT. In these two specifications, the time since last offer coefficient is statistically insignificant and in fact becomes negative. Taken together, there is no evidence that patients adjust their acceptance behavior based on the recent offers they have experienced. We are therefore comfortable with the restrictions embedded in Assumption 3.

4 Estimation

The key primitive 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. The two leading techniques for estimating dynamic choice models of this type are the conditional choice probabilities (CCP) approach (Hotz and Miller, 1993; Aguirregabiria and Mira, 2007; Arcidiacono and Miller, 2011) and the full solution or nested fixed point approach (Miller, 1984; Wolpin, 1984; Pakes, 1986; Rust, 1987). We employ the CCP approach because it affords a computationally tractable estimator that allows us to use detailed knowledge of the mechanism. This section begins by laying out our preferred approach and the empirical specification before discussing alternatives in Section 4.2.

[23] Note that many sources of bias, such as patient unobserved heterogeneity and measurement error, would bias our estimates toward finding a positive coefficient on time since last offer. For example, suppose our controls for patient priority were imperfect. In this case, some patients would be unobservably higher-priority and, as a result, more selective. For these patients, we would observe lower times since last offer and lower acceptance rates for a given kidney, generating a positive omitted variables bias.

4.1 A CCP Approach for Sequential Assignments

Although equation (5) is an integral equation, it expresses the value function in terms of three sets of parameters. The first is the distribution of difference between the value of accepting an offer and declining:

$$\Gamma(t, x, z, \eta) + \varepsilon - V(t; x).$$

The second is the distribution of offers an agent receives in the future, which is a function of the object arrival rate γ, the distribution of object characteristics F, and the offer probabilities π. The final set of parameters are the departure rates $\delta(t; x)$ and the parameter governing time-preferences ρ. The value function can be recovered if these parameters can be estimated. The differences in the first set of parameters can then be used to obtain $\Gamma(\cdot)$.

We estimate the model in four steps. First, we estimate departure rates using observed patient departures. Second, we estimate conditional choice probabilities from patient accept/reject decisions. We use these estimates and the Hotz-Miller inversion to solve for the differences $\Gamma(t, x, z, \eta) + \varepsilon - V(t; x)$ and $\tilde{\psi}(\tau, x, z, \eta) = \mathbb{E} \max\{0, \Gamma(t, x, z, \eta) + \varepsilon - V(t; x)\}$ in equation (5). Third, we compute the integral in equation (5) using the empirical distribution of donor types and offer probabilities 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 by evaluating equation (5).

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 percent per year. Our results are robust to using an annual discount rate of 10 percent. For modest discount rates, most of the discounting of future offers is due to the term $\delta(t; x)$, which is estimated at approximately 16% per year for the average patient.

Step 1: Estimating Departure Rates

A patient's continuation value on the waiting list depends on how long she can expect to continue waiting before an exogenous departure. Our dataset contains information on how long each patient is observed on the list without a transplant, and their reason for departure. We can therefore construct a censored measure of the length of time a patient would remain on the list without a transplant. Censoring occurs if the patient is transplanted, or if she is still on the list at the end of the sample period. These censored measures can be used to estimate departure rates independently of payoffs because Assumption 2 implies that, conditional on patient characteristics, departure from the list prior to assignment is exogenous.

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

$$\delta(t; x_i) = \delta_1 \exp(\delta_2 t) \exp(x_i \beta), \tag{6}$$

where $\delta_1 \exp(\delta_2 t)$ is the baseline hazard function for the Gompertz model and x_i are observed patient covariates. The parametric form for the baseline hazard function has the advantage of allowing for a simple expression for the survival function $p(\tau|t; x_i) \equiv \exp\left(-\int_t^\tau \delta(\tau'; x_i)\, d\tau'\right)$. It turns out that the estimated model yields a survival curve similar to the semi-parametric Cox proportional hazards model.

Step 2: CCP Representation and Gibbs Sampling

Consider the probability that agent i refuses an offer of kidney j at time t. Assumption 1(i) implies that this probability is given by

$$P_{ijt} = G\left(V(x_i, t) - \Gamma(x_i, z_j, \eta_j, t)\right),$$

where G is the CDF of ε_{ijt}. This quantity is referred to as the conditional choice probability (CCP) of refusing an offer, given x_i, z_j, η_j, and t. For now, assume that P_{ijt} is known. Proposition 1 of Hotz and Miller (1993) shows that for any known distribution G that satisfies Assumption 1(ii), there is a known function ψ such that

$$\psi(P_{ijt}) = \mathbb{E}\max\left\{0, \Gamma(x_i, z_j, \eta_j, t) + \varepsilon_{ijt} - V(x_i, t)\right\},$$

where the dependence of ψ on G has been suppressed for simplicity of notation. Substituting $\psi(P_{ijt})$ for $\tilde{\psi}(\tau, x_i, z_j, \eta_j)$ in equation (5), the value function can be re-written in terms of the CCPs as

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

Therefore, if P_{ijt} can be estimated, the only remaining unknowns in this equation are $\pi_{ij}(\tau)$ and F. The value of a transplant $\Gamma(x_i, z_j, \eta_j, t) = V(x_i, t) - G^{-1}(P_{ijt})$ can therefore be written in terms of the remaining unknowns, P_{ijt}, $\pi_i(\cdot)$ and F, without directly solving the integral equation (5).

In our application, we assume that $\varepsilon_{ijt} \sim N(0, 1)$.[24] The variance of this term serves as our scale normalization. Further, we assume that $\Gamma(\cdot)$ is additively separable in η_j and approximate

$$V(x_i, t) - \Gamma(x_i, z_j, \eta_j, t) = \chi(x_i, z_j, t)\, \theta + \eta_j, \tag{8}$$

where $\chi(\cdot)$ is a flexible set of functions with interactions between its arguments. We include dummies in x_i and z_j for categorical variables and piecewise linear splines for their continuous elements, as well as piecewise linear splines in t. The bases in these categorical variables and

[24] Therefore, $G = \Phi$ is the CDF of the standard normal. This give us a simple expression for evaluating $\psi(P_{ijt})$ because in this case $\psi(P) = \phi\left(\Phi^{-1}(P)\right) - (1 - P)\,\Phi^{-1}(P)$.

splines are interacted with each other. Donor unobserved heterogeneity is parameterized as $\eta_j \sim N(0, \sigma_\eta)$, with a variance to be estimated.

Because agents make an accept/reject decisions, they solve a binary choice problem. We estimate the parameters (θ, σ_η) using a Gibbs' sampler (McCulloch and Rossi, 1994; Gelman et al., 2014), which yields a posterior distribution with a mean that is asymptotically equivalent to the maximum likelihood estimator (see van der Vaart, 2000, Theorem 10.1 (Bernstein-von-Mises)).[25]

Identification of these parameters is intuitive. 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 kidneys offered to patients. The correlation between the number of offers before the acceptance of the first and second kidneys reveals the importance of unobserved donor quality. If σ_η^2 is large, then conditional on the observables x_i, z_j, and t, an early acceptance of the first kidney from a donor indicates that the second acceptance should soon follow. In constrast, if σ_η^2 is small, then the position of the first acceptance should have little information about the second. The intuition is similar to those for results on the identification of measurement error models (see Kotlarski's theorem in Rao, 1992; Hu and Schennach, 2008).

Step 3: Simulating the Mechanism

With estimates of the CCP parameters (θ, σ_η^2) and departure rates $\delta_i(t)$, our next objective is to use equation (7) to calculate each patient's value function and recover transplant values. To do this, we only need to estimate the inner integral in equation (7),

$$\begin{aligned} W(x_i, t; \theta_0) &= \int \pi_j(t; x_i)\, \psi(P_{ijt})\, \mathrm{d}F, \\ &= \mathbb{E}\left[\mathbb{P}(c_{ij} = 1 | z_j, x_i)\, 1\left\{s_{ijt} > s_j^*\right\} \psi(P_{ijt}) \middle| x_i, t\right] \end{aligned}$$

because ρ is fixed and we have consistent estimates of $p(\tau; t, x_i)$, θ_0 and λ_0. Expectations in this expression are taken over donor characteristics (z_j, η_j) drawn from F; and the priority-score cutoff s_j^* drawn from the conditional distribution of cutoffs $H(\cdot; z, \eta_j)$ given (z_j, η_j). The second equality is implied by the definition of $\pi_{ij}(t)$ given in Assumption 3. As a notational reminder, $c_{ij} = 1$ if agent i is compatible with object j and s_{ijt} is the priority score of agent i for object j at time t.

We estimate this quantity by first determining the set of donors that patient i would have been offered had the donor arrived when the patient had waited for t periods. Recall from

[25] The Gibbs' sampler obtains draws of θ and σ_η from a sequence of conditional posterior distributions using a Markov chain given dispersed priors and an initial set of parameters $\left(\theta^0, \sigma_\eta^0\right)$. 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 B.2.

our discussion in Section 3.3.1 that an agent receives an offer for object j if she is compatible and her priority score exceeds s_j^*. Therefore, we constuct the sample analog

$$\hat{W}\left(x_i, t; \hat{\theta}\right) = \frac{1}{J} \sum_{j=1}^{J} \hat{P}\left(c_{ij} = 1\right) 1\left\{s_{ijt} > s_j^*\right\} \psi\left(\hat{P}_{ijt}\right) \tag{9}$$

where the P_{ijt} is replaced with $\hat{P}_{ijt} = G\left(\chi\left(x_i, z_j, t\right)\hat{\theta} + \eta_j\right)$, and j indexes a donor in our sample and the observed threshold priority for that donor, s_j^*. Because the summand in this expression conditions on s_j^*, we draw η_j from the posterior distribution given the observed accept/reject decisions of all patients offered donor j. Knowledge of the mechanism allows us to directly compare the patient's priority score s_{ijt} with s_j^*. Further, our dataset contains rich information on donor proteins and patient immune system characteristics that allow us to accurately estimate the probability with which $c_{ij} = 1$.[26]

It is worth emphasizing the importance of Assumption 3 in substituting the expectation with the sample average. Under a richer information set $\mathcal{F}_{i,t}$ (as defined in equation 2) that conditioned on the history of offers received by a patient or the configuration of the list, we would only be able to use the subset of donors that were offered to a patient with exactly the same information set when calculating the second approximation above. It is easy to see why this would restrict the sample size and limit our ability to accurately estimate W.

Theorem 1 in Appendix B.4 shows that, for each x_i and t, $\hat{W}\left(x_i, t; \hat{\theta}\right)$ is a $\sqrt{J}$-consistent estimator of $W\left(x_i, t; \theta_0\right)$ under conditions formalized in Assumption 4, also in Appendix B.4. The main requirement is on the serial dependence of the value of offers to a patient. Specifically, we require that the dependence of a potential future offer on the organ that has arrived today diminishes with the time-horizon for the future offer. The other conditions for the result are technical regularity conditions on the primitives, and the use of a well-behaved estimator for θ_0.

Step 4: Estimating Γ

The final step recovers $V\left(t; x_i\right)$ and $\Gamma\left(t, x_i, z_j, \eta_j\right)$. We estimate $V\left(t; x_i\right)$ for time t and agent i by numerically integrating the expression in equation (7). To do this, we evaluate the integrand at a large number of points. We substitute the sample analog for $W\left(x_i, \tau; \theta_0\right)$ in equation (9), the estimated departures model for $p\left(\tau | t; x_i\right)$, and the observed donor arrival rate for λ. Details of this procedure for computing $\hat{V}\left(t; x_i\right)$ are provided in Appendix B.2.

[26] As mentioned in Section 2.2 a crossmatch is conducted using blood from the donor and patient in case the virtual crossmatch yielded a false negative. We observe the rate of positive crossmatches in the data using instances where a kidney was accepted because of a negative crossmatch, but the transplant did not occur because the final crossmatch was positive. We use this conditional probability of a positive crossmatch for $\hat{P}\left(c_{ij} = 1\right)$ in the expression above. Further details are provided in Appendix B.3.

Once $\hat{V}(t; x_i)$ is calculated, we recover $\Gamma(\cdot)$ by inverting G :

$$\hat{\Gamma}(t, x_i, z_j, \eta_j) = \hat{V}(t; x_i) - G^{-1}\left(\hat{P}_{ijt}\right).$$

This quantity can be calculated for any value of (t, x_i, z_j) .

4.2 Discussion

Unobserved Heterogeneity

The model omits patient unobserved heterogeneity. Arcidiacono and Miller (2011) develop CCP approaches that allow for time-invariant unobserved heterogeneity with discrete types. There are three main complications in applying our methods. First, both the departure rates $\delta_i(t)$ and choice probabilities would ideally depend on unobserved heterogeneity and would need to be simultaneously estimated. Second, a solution to the well-known initial conditions problems needs to be developed for the agents already on the waiting list at the begining of the sample. Third, one may further argue that unobserved heterogeneity may be time varying if unobserved patient health is stochastic. This last extension may require significant revision to the model.

We believe that abstracting away from unobserved heterogeneity still yields useful results because our dataset contains a very rich set of patient characteristics. Nonetheless, we explored specifications in which unobserved heterogeneity was included in the CCP specification only. These specifications yielded qualitatively similar results for the counterfactual analysis (see Appendix E for details).

Our model does include donor-level unobserved heterogeneity through η. One motivation for doing this is the pattern of sharply declining acceptance rates by position documented in Figure 1b. The observable characteristics included in the model do not explain all of this sharp decline or the composition of offers, especially at lower positions on the waitlist. We will compare model fit with and without such unobserved heterogeneity in the next section.

Comparison with Full Solution Approaches

An alternative to the CCP approach is to use a full solution or nested fixed point approach. The full solution approach would parametrize $\Gamma(\cdot)$ directly, say, as $\chi(x, z, t, \eta)\theta_\Gamma$, recover the value function by solving the fixed point in equation (5), and optimizing a statistical loss function (e.g. maximimum likelihood) to obtain an estimate $\hat{\theta}_\Gamma$. Compared to the CCP approach outlined above, the primary advantage of the full solution approach is to avoid directly parametrizing the (endogenous) difference in equation (8), $\Gamma(\cdot) - V(\cdot)$, in terms of

$\chi(\cdot)$. Aside from the appeal of using an implied functional form for $V(\cdot)$ that is consistent with the primitives, the approach also uses the data more efficiently.

However, this full solution approach is well known to be computationally burdensome when the state space is large (see Hotz and Miller, 1993; Arcidiacono and Miller, 2011). The high dimensionality of the state space remains a problem in the deceased donor allocation context despite the simplified model of beliefs. Specifically, the simulation in equation (9) computes the compatibility and priority score for each patient and donor using all the variables that enter the assignment mechanism and detailed information about the immune response of a patient to each potential donor. A full solution or nested fixed point method would require us to separately solve for the value function using equation (5) for each patient at every single offer. Moreover, it would require such a solution to be computed for each guess of θ_Γ used within an optimization routine.[27]

The CCP approach avoids these complications and makes it computationally feasible to respect the details of the mechanism. The main cost is loss in statistical efficiency because we directly estimate the difference in equation (8), and parametrize this difference in terms of the functions $\chi(\cdot)$. The latter cost is mitigated by using flexible functional forms. Nonetheless, our implementation does impose some continuity across continuous states and limits interactions because a fully non-parametric approach is prohibitive given the dimensionality of the state space.

5 Parameter Estimates

This section describes our estimates of patient departure rates, conditional choice probabilities, the value function, and the value of transplantation.

5.1 Departure Rates

Table 6 presents estimates from hazard models of departures from the kidney waitlist prior to transplantation. We estimated models under different parametric assumptions about the baseline hazard functions. We used all of the patient-specific variables included in the CCP model.[28] Across specifications, we estimate an increasing baseline hazard of departure,

[27] Essentially no two patients are identical because the mechanism awards points whenever a donor and a patient have overlapping antigens, and because immune responses to donors are idiosyncratic. There are about 2.85 million offers in our dataset, making this problem extremely computationally expensive. One approach to simplifying the problem, as done in our counterfactual analysis, is to evaluate the value function for each patient on a discrete grid of times. Because we have just under 10,000 patients, even a grid with 100 points for each patient would require solving for approximately 1 million instances of the value function at each guess of θ_Γ.

[28] The hazard model specifications include the CPRA variable from the CCP model, but omits a dummy for $CPRA > 0.8$. Because priority is discontinuous in CPRA at 0.8, it is important to allow acceptance behavior

consistent with patients becoming less healthy over time. Column (1) presents a Gompertz model, in which the estimated value of $\log(\delta_1)$ is -5.151 per day, indicating an increasing baseline hazard. Other coefficient estimates reveal significant heterogeneity across patients in their departure rates. For example, diabetic patients depart at higher rates. Patients with blood type A are more likely to depart than patients with blood type O, perhaps because of better living donor transplant opportunities. Among pediatric patients, older patients are less likely to depart, perhaps because they are better able to tolerate dialysis. However, as patients get older, the departure rates begin to increase with age. Columns (2) and (3) estimate corresponding Weibull and Cox proportional hazards models with the same set of patient covariates. Point estimates for these coefficients remain stable across assumptions on the baseline hazards, including the semi-parametric Cox proportional hazards model. Figure B.2 in the appendix compares the estimated survival curves from columns (1) and (3) and shows that the Gompertz hazards model yields a survival curve that is very similar to the non-parametric Cox proportional hazards model. We use the Gompertz hazard model with patient covariates reported in column (1) as our preferred specification.

5.2 Estimated CCPs

We estimated three specifications for the conditional choice probability of accepting an offer. All specifications include the rich set of patient and donor observed characteristics summarized in Tables 1 and 2. The first specification includes all of these baseline variables, but does not include 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. We explain the choice of baseline characteristics, and then describe the estimates.

The baseline characteristics common across specifications, as well as linear splines and interactions among these variables, were chosen by surveying the medical literature. Specifically, we use covariate and spline specifications from the KPSAM model, which was used by the kidney allocation committee to predict the outcomes of various allocation systems.[29] We also include any covariates that were part of the survival models for kidney transplant patients used in Wolfe et al. (2008). Following KPSAM and our earlier observation that donors from other DSAs are less desirable, we include interactions of donor and patient characteristics

to be discontinuous at this point. Departures without a transplant, however, should not be discontinuous in CPRA. Moreover, specifications of the departures model that included this variable estimated a statistically insignificant coefficient.

[29] We obtained the KPSAM module from the Scientific Registry of Transplant Recipients (SRTR), which contains the specification of the KPSAM acceptance model. Our dataset contained all but one of the variables used in that model. Visit https://www.srtr.org/requesting-srtr-data/simulated-allocation-models/ for a description of the various simulated allocation models and the procedure to request these modules.

with whether the donor was recovered in NYRT. In addition, we include patient-donor specific variables that capture match-specific heterogeneity, for example whether there are two DR antigen mismatches. We specify piecewise linear splines for continuous covariates and interact them with a variety of indicator variables.

Table 7 presents select parameter estimates from the three specifications (see Table B.3 for the full specification). The estimated coefficients on observed donor characteristics are intuitive and fairly robust across specifications. For example, offers from donors older than 50 years of age are less likely to be accepted than offers from 35 to 50 year old donors, and kidneys from younger donors are even more likely to be accepted. These differences are larger for donors recovered in NYRT. A perfect tissue type match is very desirable, much more so than a young donor. Also intuitive are our estimates that offers of kidneys with more antigen mismatches (A, B or DR) are less desirable and that regional and national offers are less likely to be accepted.

We also estimate significant patient-level and match-specific heterogeneity in acceptance rates. Consistent with our discussion in Section 2, patients with more sensitive immune systems (higher CPRA) are more likely to accept an offer. Acceptance rates also depend on patient age. Among pediatric patients, older patients are more likely to accept an offer, which is intuitive because size compatibility is important for younger pediatric patients. 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, and suggests potential welfare gains from improved matching on patient and donor age.

In the second specification, the estimated standard deviation of donor unobserved heterogeneity is 1.03. The implied change in acceptance rate from a one standard deviation increase in η is therefore half the difference between a kidney recovered inside NYRT and one recovered outside NYRT. The third specification shows that acceptance rates fall rapidly with waiting time in the first few years before starting to increase after year three. Most other parameter estimates, including the standard deviation of unobserved donor-level heterogeneity, are similar in the second and third specifications.

Figure 2 describes the fit of these models. The first panel mimics the fit of acceptance rates by position described in Figure 1b, but includes offers that did not meet pre-set screening criteria. This panel shows that including donor unobserved heterogeneity much better captures the sharp decline in acceptance rates. The specification that excluded donor unobserved heterogeneity does not even accurately capture the average acceptance rate across the first 20 positions. Instead, it implies a more gradual decline as the composition of donor observables changes moving down the list. The second panel of Figure 2 describes the fit of acceptance rates by time waited. Not surprisingly, the third specification does the best job of capturing changes in acceptance rates over time.

Taken together, we feel comfortable with the fit of the CCPs in the third model and use that as our preferred specification. All results that follow use estimates from this specification.

5.3 Estimated Value of Organ Offers

Below, we develop an interpretable measure for comparing values across patients in terms of an equivalent change in donor arrival rates (supply). We then use these units to describe our estimates.

Sequential assignment mechanisms can re-assign offers from some patients to others. Motivated by this fact, consider the proportional increase in a patient's value from a one-time offer for kidney j at the time of registration:

$$EV_{ij} \equiv \frac{\mathbb{E}\max\left\{0, \Gamma\left(0, x_i, z_j, \eta_j\right) + \varepsilon_{ij} - V\left(0; x_i, \lambda\right)\right\}}{V\left(0; x_i, \lambda\right)},$$

where the dependence of V on the object arrival rate λ is re-introduced for clarity. The numerator is the difference between the value with and without an additional one-time offer for object j at time 0, and the denominator is the baseline value.

Instead of a one-time offer, the same change in value can be generated by an increase in the donor arrival rate. Specifically, let λ_{ij} be defined such that

$$V\left(0; x_i, \lambda_{ij}\right) = V\left(0; x_i, \lambda\right) + \mathbb{E}\max\left\{0, \Gamma\left(0, x_i, z_j, \eta_j\right) + \varepsilon_{ij} - V\left(0; x_i, \lambda\right)\right\}.$$

Using this expression, we can rewriteEV_{ij} as

$$\begin{aligned} EV_{ij} &= \frac{V\left(0; x_i, \lambda_{ij}\right) - V\left(0; x_i, \lambda\right)}{V\left(0; x_i, \lambda\right)} \\ &\approx \frac{\lambda_{ij} - \lambda}{\lambda}, \end{aligned}$$

where the approximation follows because, holding behavior and offer probabilities fixed, $V\left(t; x_i, \lambda\right)$ is approximately linear in λ (see equation 5). This approximation is appropriate for small changes in λ, so that offer probabilities do not change substantially.

The equivalent change in donor arrivals is, by definition, invariant to the scale of utility units across agents and will be used to report welfare effects going forward. It equates an additional one-time offer to the value of an alternative policy that is able to marginally increase the organ supply. This feature makes this quantity similar in spirit to Equivalent Variation at the time of registration, yielding a measure of payoffs that is interpretable across agents for small changes in the environment.

Figures 3 – 5 describe our estimates in these units. The results are based on our preferred specification, which includes donor unobserved heterogeneity and time waited. The plots

show how the value of an organ offer varies across specific patient and donor characteristics, holding all remaining characteristics fixed.

Healthier patients, whether measured by patient age or by dialysis status at registration, prefer younger donors. Figure 3 shows that patients across all age groups prefer younger adult donors. However, the relative value of a young donor decreases with patient age. For a 70 year old patient, a single offer from an 18 to 35 year old donor is as valuable as a 4 to 5 percent increase in overall donor supply, while a donor over 50 years old is half as valuable. In contrast, for a 20 year old patient, a donor over 50 years old is only one twentieth as valuable as an 18 to 35 year old donor.[30] This pattern is consistent with the CCP estimates and with the intuition that older donors should place less value on an organ that is likely to function for a very long period of time. Similarly, Figure 4 shows that patients off dialysis at registration place a relatively higher value for younger donors

In addition to donor age and patient health, a perfect tissue type match is especially important for patients (Figure 5). Such an offer from a young donor is worth a 32 percent increase in overall donor supply for a representative patient; from a donor over age 50, it is still worth a 17 percent increase in supply. This result is also intuitive because an organ with a perfect tissue type match is less likely to cause an adverse immune response, thereby increasing the life-years afforded by the transplant. Offers which are not perfect tissue type matches are an order of magnitude less valuable.

Detailed estimates for V and Γ (in ε_{ijt} units) are presented in Tables B.3 and B.3 in the Appendix. The reported co efficients are obtained from a projection of these quantities on $\chi(\cdot)$. The projection coefficients for both Γ and V are intuitive. For example, younger donors are more valuable, while donors with tissue type mismatches are less valuable. Similarly, donors recovered outside NYRT are less desirable.

6 Steady State Equilibria and Welfare Comparisons

6.1 Equilibrium Concept

We now define an equilibrium concept for counterfactual analysis. The concept is intended to capture a large pool of agents waiting for offers and making optimal decisions. Agents have type $x \in \mathcal{X}$, and objects have type $z \in \zeta$, where we henceforth include the unobserved donor characteristic η in z for notational simplicity. For computational reasons, we will treat $\mathcal{X}$ and ζ as finite sets.

[30] Older patients place higher values on all types of offers, in terms of equivalent supply increases. This is mainly because older patients exogenously depart more quickly, and therefore place a higher value on a single current offer relative to a proportional increase in the stream of future offers.

To compactify notation, albeit with a slight abuse, the rest of the paper replaces subscripts that index individuals i and objects j with 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, $\sigma_x : \mathbb{R} \times \mathbb{R}_+ \to \{0, 1\}$, indexed by $x \in \mathcal{X}$. The first element of the domain 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 composition of the queue using a single steady state. Specifically, the queue composition will be governed by a probability density function, m, defined on the set $\mathcal{X} \times [0, T]$, where T is the maximum wait time.[31] This density governs the distribution of agents of each type and how long they have waited. We write $m_x(t)$ to denote the density evaluated at (x, t). The length of the queue 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 for type x of declining the object and following the optimal strategy given π^* after t.

2. Consistent beliefs: For each (t, x, z), the beliefs $\pi^*(t; x, z)$ is consistent with equilibrium offer probabilities. In particular, for mechanisms that uses 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^*.[32]

3. Steady State: m^* and N^* satisfy the balance conditions

   (a) For each $x \in \mathcal{X}$, $m_x^*(t)$ and N^* satisfy

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

[31] We assume that the density m is defined with respect to the Lebesgue measure on the Borel sets formed from $\mathcal{X} \times [0, T]$, where $\mathcal{X}$ is a finite set. Our counterfactuals continue to assume that no agent lives past one hundred years of age. Consequently, we set T to one hundred years and restrict $m_x(t) = 0$ for all t such that a agent of type x would be more than one hundred years old.

[32] We do not restrict the queue length $N^*$to be an integer and therefore round it to the nearest integer when calculating π.

where γ_x is arrival rate of an agent of type x, and $\kappa_x(t)$ is the equilibrium departure rate of an agent of type x at waiting time t.

(b) 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 her beliefs, assuming that she continues to make optimal decisions in the future. The value from declining an offer is given by the Hamilton-Jacobi-Bellman equation defined in Section 3.3. The second condition imposes that agents have correct beliefs. In the specific case of a mechanism based on the scoring rule s, it writes agent beliefs about future offers as the product of the steady state distribution of cutoffs and exogenous compatibility realizations. The distribution H_z^* governs the cutoffs that arise when agents use strategies σ^* and N^* agents are drawn from a distribution governed by m^*.[33] The final condition determines the composition of agent characteristics on the list. The left-hand side in part (a) is the change in the density of agents of type x who have waiting time t. The right-hand side term is the rate of departure for those agents. Departures occur for both exogenous reasons and because agents are removed from the waiting list once they are assigned; that is, $\kappa_x(t)$ is the sum of $\delta_x(t)$ and the equilibrium rate at which agents of type x are assigned at time t. The strategy σ^* and the offer rate of objects, given by π^*, determine the endogenous departures. The agent arrival rate γ_x is exogenous, and in the context of our application, it will only be positive for agents with zero waiting time since patients begin to accumulate waiting time once they enter the queue. Part (b) ensures that m^* integrates to one.

This equilibrium concept abstracts away from transitional dynamics in the size and composition of the queue. For example, in our empirical context, it is possible that several patients join the queue in close succession before any organs arrive. One approach may be to model these dynamics by assuming that the queue length and composition follow a Markov process. In principle, one could solve for the stationary distribution over the queue length and queue compositions. However, this distribution is extremely high dimensional, and it would make counterfactual exercises computationally intractable.

As in Assumption 3, we view our equilibrium concept as an approximation to the behavior of the system. For example, although the mass of agents of type x is discrete, part 3(a) implies that the ratio $m_x(t)/m_x(0)$ is given by the survival curve implied by the equilibrium departure rates $\kappa_x(t)$. Similarly, N^* is assumed to be the expected queue length. In fact, Theorem 2 in Appendix C uses a concentration inequality to show that the stationary distribution of the queue length concentrates mass on N^*. Our result implies that if N^* is 5000, the weight placed by the stationary distribution on queue lengths that deviate by more than 5 percent is at most 0.5 percent. Hence, we expect that our equilibrium notion will be a good approximation for the behavior of the waitlist.

[33] In contrast with a direct continuum approximation with no aggregate uncertainty, this specification allows for H_z^* to be non-degenerate.

Finally, we prove the existence of a steady state equilibrium for sequential assignment mechanisms that use a priority score in Theorem 3. The challenge in showing existence arises because the strategies, beliefs and composition are a function of time, which is a continuous variable. We use the Brower-Schauder-Tychonoff fixed point theorem (Corollary 17.56, Aliprantis and Border, 2006) for general Banach spaces to prove existence. The primary assumptions are technical regularity conditions imposing bounds and Lipschitz continuity of 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 set of counterfactuals and specifications that are considered below.

6.2 Welfare Comparisons

Given a mechanism $\mathcal{M}$ and a donor arrival rate λ, let the **steady-state value** of an agent of type x be the integrated value function over current and future generations:

$$\bar{V}_x^{\mathcal{M}}(\lambda) = \frac{\gamma_x}{\rho}\int_0^{\infty} \exp(-\rho\tau)\, V_x^{\mathcal{M}}(0;\lambda)\,\mathrm{d}\tau + \int_0^{T} N^* m_x^*(\tau)\, V_x^{\mathcal{M}}(\tau;\lambda)\,\mathrm{d}\tau. \tag{10}$$

The first term represents the discounted value of agents of type x that are expected to arrive in the future, and the second term represents the value of agents presently on the waiting list. The first term is weighted by the net present value of the total mass of future arrivals $\frac{\gamma_x}{\rho}$, and the second term is weighted by the equilibrium measure $N^* m_x(\tau)$ that governs the mass of agents of type x that have waited for τ periods. By considering the welfare of both current and future generations, we take the view of a social planner who is interested in the net present value of payoffs generated by all future assignments. Our tables will also report an alternative measure that only considers the value at the time of entry, $V_x^{\mathcal{M}}(0;\lambda)$. The theoretical discussion in this subsection applies to this alternative as well.

Let $\mathcal{M}_0$ be the baseline mechanism used during our sample period and let λ_0 be the observed donor arrival rate. For any mechanism $\mathcal{M}$, define $\lambda_x(\mathcal{M})$, the equivalent donor arrival rate for agents of type x, as 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 5.3, 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 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.

A common challenge in comparing the aggregate welfare effects of various mechanisms in

environments without transfers is the lack of a clear numeraire good that can be used to justify a comparison based on a Kaldor-Hicks criterion. The measure $EV_x(\mathcal{M})$ is type-specific, and therefore does not make comparisons across patients of different types. We will summarize welfare effects by reporting equally weighted averages of $EV_x(\mathcal{M})$. This aggregation treats equivalent changes in donor arrivals for different patient types as equal. Of course, alternative welfare weights are easily entertained within this framework. Our results will also report distributional effects on key subgroups in order to allow the reader to independently assess the welfare effects for alternative weights on these subgroups.

6.3 Computing Equilibria

We compute steady-state equilibria using an algorithm that iterates between computing the value function and optimal decisions, and the steady-state composition of the waitlist. A detailed description with expressions for each step of the procedure and pseudocode is provided in Appendix D. The following discussion provides a simplified description of the key steps.

In addition to the primitives, the algorithm uses a discrete time grid $t = t_0, \ldots, t_l, t_{l+1}, \ldots, T$, abitrary initial beliefs π^0, and a sample of patients and donors as inputs. An equilibrium is computed by iterating through the following steps for $k \geq 1$:

1. Compute the value function $V_x^k(t_l)$, given beliefs π^{k-1}, via backwards induction from the value of waiting in the next period $V_x^k(t_{l+1})$:

$$V_x^k(t_l) = \int_{t_l}^{t_{l+1}} \exp\left(-\rho(\tau - t_l)\right) p_i(\tau|t_l)\, \lambda \int \pi^{k-1}(t;x,z)\, \mathbb{E} \max\left\{V_x^k(t_{l+1}), \Gamma(\tau;x,z) + \varepsilon\right\} \mathrm{d}F \mathrm{d}\tau.$$

   The inner integral in the above expression is approximated using sampled subset of donor types. This procedure starts with $V_x^k(T) = 0$ for all x and k, and working backwards to compute $V_x^k(t)$ for all t. This calculation also yields patient strategies $\sigma_x^k(\Gamma, t) = 1\left\{\Gamma \geq V_x^k(t)\right\}$ and departure rates $\kappa_x^k(t)$.

2. Compute the queue composition m^k via forward simulation:

$$m_x^k(t_l) \propto \gamma_x \exp\left(-\int_0^{t_l} \kappa_x^k(\tau)\, \mathrm{d}\tau\right).$$

3. Compute $\pi^k(t;x,z)$, which is the probability that an agent of type x is offered an object of type z using the queue composition and the accept/reject strategies $\sigma_x^k(\Gamma, t)$.

4. For step $k > 1$: Terminate if the change in value functions and queue length/compositions between iterations – $\sup_{x,l}\left|V_x^k(t_l) - V_x^{k-1}(t_l)\right|$, $\sup_{x,l}\left|m_x^k(t_l;x) - m_x^{k-1}(t_l)\right|$, $N^k - N^{k-1}$

– are uniformly below a chosen tolerance level. If these conditions are not satisfied, repeat steps 1-4.

If this algorithm terminates, the resulting accept/reject rules yield an equilibrium (up to the threshold tolerance). Because the equilibrium we compute may not be unique, we tried different starting values for π^0. Our experiments at the estimated parameters do not indicate that multiplicity is a concern in our setting.

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. Further, we discretize time into quarters for the first 15 years after registration, then every 2 years until year 25, and every 25 years thereafter. As discussed in Appendix E, the results reported below are not sensitive to a larger set of types used to calculate equilibria of scoring mechanisms and to finer partitions after the first few years since the probability that a patient survives without a transplant falls dramatically.

7 Evaluating Design Trade-Offs

7.1 Alternative Mechanisms

A new deceased donor organ allocation system was adopted in December 2014. The mandate of the kidney committee, as laid out by the U.S. Department of Health and Human Services, was to find mechanisms that balanced the goals of providing equitable outcomes for patients, efficiently using available organs, and minimizing organ waste. This section compares mechanisms aimed at achieving these goals to the two mechanisms used prior to and following the 2014 re-design.[34]

We start by computing four benchmarks against which we compare simpler mechanisms motivated by theory and practice.

- **Optimal Assignments:** This problem bounds the welfare gain achievable by any mechanism. To do this, we maximize the steady-state average welfare of all agents subject only to feasibility constraints. The program solves for an assignment policy $a(\varepsilon; x, z, t) \in \{0, 1\}$ under full information about preferences and rational expectations (but no foresight) about object arrival and agent departure processes. The policy

[34] Our calculations only change the allocation mechanism in NYRT. Evaluating a nationwide change would require us to use data on decisions made by all patients in the US, which is burdensome due to the patchwork of variants on points used in approximately half 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 the policy function of patients from the rest of the US, which governs offers for non-local donors to patients in NYRT, remains fixed.

can dictate assignments to agents as a function of the agent type, the object type, time waited, as well as the preference shocks ε for all agents.

We choose a to maximize the average equivalent donor supply increase, $\sum_x \frac{\bar{V}_x^a(\lambda_0)}{\bar{V}_x^{\mathcal{M}_0}(\lambda_0)}$, where $\bar{V}_x^a(\lambda_0)$ is the steady-state value for type x under assignment policy a. The feasibility constraints ensure that the total rate of assignment does not exceed the total arrival rate of objects. Specifically, for each z, we impose (a time-discretized version of) the constraint

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

The term $\mathbb{P}(c=1, a=1|x,z,t)$ is the probability that an object of type z is compatible ($c=1$) and assigned ($a=1$) to a randomly chosen agent of type x that has waited for time t. In addition, we require that the composition of the waitlist be in steady state (Definition 1, part 3). These constraints are described in detail in Appendix D.2.1.

- **Optimal Offer Rates:** We solve for welfare maximizing offer rates in which the designer can choose offer rates, but cannot dictate assignments or condition on offer rates refused in the past. The offer rates $\pi_{xz}(t)$ depend on observable characteristics of patients, but not unobserved preference shocks ε. We solve for steady-state equilibria (Definition 1), implying that agents make optimal decisions given these offer rates. Justifications for only not allowing the designer to dictate assignments include respect for patient and doctor discretion and the concern that forcing assignments or penalizing terminally ill patients may be politically infeasible. In addition, a designer may not have full information about idiosyncratic preferences, and may therefore be unable to implement the optimal assignment solution even if she can dictate assignments.

  Formally, we choose π to 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, so that the steady-state rate of assignment implied by offer and acceptance rates does not exceed the arrival rate of objects. In this problem, the feasibility constraint from the optimal assignment problem is modified by replacing the term $\mathbb{P}(c=1, a=1|x,z,t)$ with the probability π_{xzt} of making an offer to an agent of type x that has waited for time t multiplied by the probability that this offer is the last offer than can be made. Ignoring, for the moment, the limit on the maximum number of offers that can be made, for each donor type z, we impose the constraint that

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

The left-hand side is the expected number of objects assigned under offer rates π and the right-hand side is the total number of objects available. The mathematical problem we solve is formally described in Appendix D.2.2.

Because this constraint is placed only on expected quantities, the offer rates may not be implementable for any particular instance of objects of type z. The solution provides an upper bound on welfare under any offer mechanism that does not condition on past behavior or have information about idiosyncratic preferences.

- **Approximately Optimal Priorities:** Although the offer rates calculated above may not be implementable, we can use the solution to implement very similar offer rates in equilibrium using a scoring rule. Specifically, we set $s_{xz}(t)$ to the values of $\pi_{xz}(t)$ that solve the optimal offer rates. Therefore, types with higher values of π are accorded higher priority. This mechanism can be implemented and described using the existing organ allocation infrastructure. We compute an equilibrium for these priorities using the algorithm outlined in Section 6.3.

- **Approximately Optimal Pareto Improving Priorites:** In practice, policy makers may be constrained to reforms that do not systematically disadvantage specific groups of patients. We design a priority system that aims to improve the welfare of the average patient without significantly hurting many patient types. To do this, we use a procedure similar to the one used to derive approximately optimal priorities. Specifically, we first solve for offer rates that are defined identically to the optimal offer rates except that they are constrained to make no agent type worse off than under pre-2014 priorities at the time they join the waitlist.[35] We then solve for the equilibrium allocation under a scoring function $s_{xz}(t)$ that is equal to the values of $\pi_{xz}(t)$ that are the solution to this problem.

  It is worth noting that the resulting mechanism may not result in a strict Pareto improvement in equilibrium. However, we expect that fewer agents will be significantly worse off under these priorities relative to the approximately optimal priorities. Our solutions will allow us to quantify these effects.

We compare these benchmarks to a set of waitlist priority rules motivated by theory and practice. These mechanisms use scoring rules $s_{xz}(t)$ to order patients waiting for a transplant and break ties uniformly at random among patients with the same score. The focus on priority rules is motivated by the design parameters considered by the kidney allocation committee.[36]

[35] During the 2014 reforms to the deceased donor kidney allocation mechanism, the kidney committee conducted simulations to verify that there would be no adverse impacts by race, age group, geography, or CPRA.

[36] Our examination of the meetings of the kidney allocation committee indicate that they did not consider

Priority rules are simpler to describe and implement than mechanisms in which assignment probabilities are determined. Equilibria for this class of mechanisms can be computed using the algorithm described in Section 6.3.

- **Post-2014:** In December 2014, the kidney allocation mechanism switched to a system that awards greater priority to patients who are extremely difficult to match (high CPRA), and also prioritizes healthier patients for higher quality donors. The rationale for the first change was that high CPRA patients have few opportunities for transplantation and are likely to accept most organs. Therefore, giving them additional priority could reduce organ waste and achieve more equitable outcomes for sensitized patients. The second change intended to offer high-quality kidneys to patients likely to benefit from them most, and in particular to reduce age mismatch.

- **First Come First Served (FCFS):** One concern of the kidney committee has been to maintain a transparent and procedurally fair offer system. FCFS is a procedurally fair and commonly used mechanism which offers objects to agents in the order they joined the waiting list. FCFS also has attractive efficiency properties: Bloch and Cantala (2017) show that it maximizes agent welfare within a broad class of primitives when values for an object are drawn i.i.d. across agents. This result is driven by the fact that FCFS encourages agents to be selective and choose only objects with high match-specific values. We approximate this system by finely discretizing time on a grid $t_0, t_1, \ldots, t_L$ and set $s_{xz}(t) = l$ if $t \in [t_{l-1}, t_l)$.

- **Last Come First Served (LCFS):** Another goal of the kidney committee has been to minimize organ waste. Last come first served provides strong incentives to accept organs because agents that refuse an offer are demoted if other agents arrive in the future. Su and Zenios (2004) show that a last come first served system both maximizes welfare and minimizes organ waste when there is agreement across agents on the values of various objects, that is, if preferences are vertical. This is because social welfare depends not on who is assigned the object, but only on the fraction of objects allocated. We approximate this system by finely discretizing time on a grid $t_0, t_1, \ldots, t_L$ and set $s_{xz}(t) = L - l$ if $t \in [t_{l-1}, t_l)$.

- **Greedy Priorities:** This mechanism attempts to maximize welfare by using a greedy procedure which offers organs to patients in order of predicted match value. For each donor, it divides patients into 10 equally sized bins based on our predicted welfare gain from assignment, measured in donor supply units (EV_{xz}). $s_{xz}(t)$ is set to 10 for those in the highest bin and to 1 for those in the lowest bin.

alternatives in which doctors were mandated to accept particular organs, or in which past rejections were used to update priorities.

Some of the mechanisms we consider, particularly the approximately optimal and greedy priorities, condition finely on patient and donor characteristics. One concern with such fine conditioning is that agents may manipulate the characteristics they report to the mechanism. We therefore caution against interpreting our results as describing readily implementable mechanisms. However, it is worth noting that such manipulation would likely involve forging medical records. In addition, most patient characteristics in our model are already used by the pre- or post-2014 mechanism.[37] Another concern is that a finely tuned priority system is complicated to understand for market participants. Addressing this issue is beyond the scope of this paper as it requires a well-articulated and practically relevant notion of mechanism complexity.

7.2 Comparing Mechanisms

We now describe the equilibrium outcomes under each mechanism described in the previous section. Our analysis shows that previous theoretical and practical recommendations can either increase patient welfare or organ discard rates, but not both. Fortunately, the optimal mechanisms described above can improve on both these goals while also ensuring that almost no patient type is significantly worse off.

The priority systems that have been in place in the U.S. since 2010 are very similar to the benchmark first come first served mechanism in terms of average patient welfare and transplant rates. In particular, the 2014 reforms did little to improve average patient welfare, and primary resulted in redistribution towards younger and more highly sensitized patients. Indeed, we find a small decline for patients above age 50 after the 2014 reforms and gains for younger patients (Table 8). These welfare changes are small on average, and only a minority (16 percent) of patients, most of whom are young, benefit from the change. In addition, the pre- and post-2014 mechanisms yield patient welfare and organ discard rates within 1.9 percent of the benchmark first come first served mechanism (Figure 6). Consistent with these results, waiting times, queue lengths, and the quality of the average donor transplanted are similar across the three mechanisms.

These results suggest that the priorities implemented in the pre- and post-2014 mechanisms primarily result in redistribution relative to the FCFS benchmark. Most patients marginally prefer FCFS to the two mechanisms that have been used in practice (Table 8). As mentioned earlier, Bloch and Cantala (2017) show that FCFS results in desirable equilibrium allocations

[37] In fact, all but two of the patient characteristics in our model were used to determine priority in either the pre- or the post-2014 kidney allocation system. The patient characteristics in our model include those used to calculate the EPTS score in the post-2014 mechanism (see https://optn.transplant.hrsa.gov/resources/allocation-calculators/epts-calculator/), CPRA, total serum albumin, and body mass index. Only the last two characteristics are not used in the post-2014 mechanism. These characteristics were implicitly used in design proposals based on a model of estimated life-years from transplantation because this model includes these variables.

when preferences are heterogeneous. Patients that have a strong preference for a particular type of kidney should be willing to wait for those organs. For example, even though the pre-2014 system did not explicitly prioritize young patients for kidneys from young donors, a significant degree of age matching was discernible (Table 4). For this reason, a relatively coarse priority system may not be able to sufficiently improve welfare relative to FCFS.

In contrast to the first three mechanisms, last come first served (LCFS) substantially reduces organ discard rates at the cost of lower patient welfare. Discard rates fall by 26.6 percent under LCFS relative to the pre-2014 mechanism, and the steady state queue length falls to 2,694 from 4,992. This increase in transplant rates comes at a significant welfare cost because patients accept organs that are poorly matched to them: the welfare of the average patient falls by an equivalent of a 33.6 percent reduction in donor supply (Figure 6). This result is also reflected in donor characteristics, with patients accepting organs from donors that are older, more likely to be hypertensive, and less likely to have died from head trauma (Table 8).

The stark difference between LCFS and the other systems occurs because LCFS dramatically changes agents' incentives, penalizing them for rejecting poorly matched offers instead of rewarding them with higher priority in the future. A model that ignores these incentives finds that the allocations are largely insensitive to the mechanism. Panel B in Table 8 shows that changes in predicted discard rates, queue lengths, and donor characteristics are similar across the benchmark mechanisms when acceptance probabilities do not adjust to the new equilibrium. Thus, it is essential to consider incentives when predicting the effects of alternative allocation mechanisms.

These results also highlight an important trade-off between organ discard rates and match quality identified in the theoretical literature. Su and Zenios (2004) show, using a model in which agents have identical preferences over objects, that LCFS reduces selectivity because agents expect to receive lower quality offers in the future. This force encourages agents to accept mismatched offers, which reduces organ discards but also lowers match quality. In contrast, FCFS increases selectivity as agents retain their priority when they decline an offer, but also increases discard rates when there is a limit to the number of offers that can be made. Bloch and Cantala (2017) use a model with highly heterogeneous agent preferences to show that the benefits of improved match quality due to increased selectivity can make FCFS a desirable mechanism. Our empirical findings weigh in favor of models that emphasize heterogeneity in match value: in terms of patient welfare, the loss from lower match quality under LCFS substantially outweighs the gain from fewer discards.

However, these theoretical benchmarks are far from optimal; a mechanism designed using estimated preferences as inputs can substantially increase both patient welfare and transplant rates, and many of these gains are possible without significantly disadvantaging specific patient groups.

The upper bound on welfare implied by an optimal assignment is large: a designer who has

full information about preferences and can dictate assignments can improve patient welfare by an equivalent of a 28.1 percent increase in donor supply. At the same time, this designer can reduce discard rates by 13.6 percent. Table 8 shows that the vast majority (83 percent) of patient types are better off at registration under these assignments than under the pre-2014 mechanism. However, not all patients are better off even in this ideal case; there are trade-offs between efficiency and distributional objectives.

The approximately optimal priorities achieve half of these potential gains. Patient welfare increases by 14.2 percent, and discard rates fall by 7.8 percent (Table 8). Since patients are transplanted at higher rates, they spend less time on the waiting list, and queue lengths fall from almost 5,000 to below 4,400. These results suggest that large improvements may be possible by re-designing priority rules. In fact, the approximately optimal priorities mechanism achieves most of the gains possible under any offer mechanism. The optimal offer rates, which place an upper bound on welfare from any offer mechanism, performs only marginally better: it increases patient welfare by 17.4 percent and reduces discards by 7.5 percent. While these gains are large, one drawback of approximately optimal priorities is that gains are concentrated amongst a minority of agents: only 42 percent of patient types are better off at registration than they were under the pre-2014 mechanism. The next section discusses how the approximately optimal priorities mechanism is able to find these improvements, and in particular how gains are distributed across patients.

A significant fraction of these gains can be achieved while respecting the types of distributional constraints faced by policy makers. The approximately optimal Pareto improving priorities mechanism increases patient welfare by 9.1 percent and reduces discards by 4.5 percent. As under approximately optimal priorities, this mechanism increases welfare through a combination of higher transplant rates and improved match quality. Recall that because this mechanism approximates the optimal Pareto improving offer rates with a scoring rule, constraint that no type should be worse off may not be exactly satisfied. In fact, most patient types (85 percent) are slightly worse off at registration under our approximation than they were under pre-2014 priorities. However,these patients experience small welfare losses: 99.7 percent of patients are no more than 5 percent worse off at registration. Thus, the very strong requirement that no patient type be significantly worse off ex-ante can be approximately satisfied while achieving a substantial improvement in average patient welfare and transplant rates.

The mechanism designer is only able to achieve these gains by considering patient incentives. The greedy priorities mechanism, which offers organs to patients based on predicted match value alone, only marginally improves on the pre-2014 mechanism. It increases the average patient's welfare by 1.2 percent and reduces organ discards by 0.3 percent (Figure 6). This small improvement is surprising given that greedy priorities incorporates the rich estimated heterogeneity in preference estimates, emphasizing the value of explicitly incorporating incentive constraints into the mechanism design problem.

It is worth noting that our computed steady state queue length for the pre-2014 mechanism is 4,942 (Table 8, panel A), which is only slightly larger than the queue length of 4,632 on January 1, 2013 (Table 1, panel A). This similarity is remarkable for two reasons. First, this moment of the data is not targeted in the CCP approach but matches our model.[38] A slightly longer steady state should be expected as the kidney waiting list was growing during our sample period. Second, it suggests that our sample is likely close to the steady state. Indeed, panel A in Table 1 shows that the rate of growth in the length of the waiting list in NYRT declined during our sample period.

Finally, Appendix E assesses the sensitivity of our results to three variations. First, we estimate a model with a limited form of patient unobserved heterogeneity. Second, we evaluate a model with a discount factor of 10 percent per year. Third, we re-compute the counterfactuals using a sample with 1000 and 1500 types of patients and donors respectively instead of the 300 and 500 types used in the baseline results. Our main qualitative conclusions are robust to these variations. For example, we robustly find that the pre- and post-2014 mechanisms are both similar in overall welfare and outcomes to the first come first served mechanism.

7.3 Optimal Offer Mechanism: Sources of Welfare Gains

Why is the approximately optimal priorities able to substantially outperform benchmark mechanisms? The answer is closely tied to dynamic incentives. Compared to the pre-2014 benchmark, approximately optimal priorities offers high quality organs to patients who not only have high transplant values (in terms of equivalent increases in donor supply), but are also are likely to accept the organs they are offered in equilibrium. This allows the mechanism to reduce discards while dramatically increasing the welfare of some patients.

In particular, the approximately optimal priorities reallocate desirable donors from patients on dialysis to patients off dialysis at registration. Table 10 compares transplant rates, offer probabilities, and waiting times for an offer under the pre-2014, greedy, approximately optimal, and approximately optimal Pareto improving priorities, separately by patient age and dialysis status at registration.[39] The first set of columns compares the number of transplants per year from each donor type received by each patient group. Under baseline priorities, older off-dialysis patients (Panel B) receive about 11 transplants per year from young NYRT donors, compared to more than 37 under approximately optimal priorities. Meanwhile, older on-dialysis patients (Panel D) receive 30 kidneys each year from young healthy donors under

[38] The steady-state length of the waiting list N is approximately the ratio of the arrival rate of patients γ to the average rate of departure (due to transplantation or otherwise) of each patient $\bar{\kappa}$. This is because, for large N, the queue must satisfy the detail balance condition $\alpha = N\bar{\kappa}$. The choice probabilities, our simulation of the mechanism and the estimated departure process without a transplant directly influence $\bar{\kappa}$. That our results on queue lengths are close is reassuring for the methodology laid out in this paper.

[39] While we do not report reallocation patterns under the (infeasible) optimal offer rates mechanism, they are very similar to the approximately optimal priorities.

baseline priorities, but only 25 under approximately optimal priorities. A similar reallocation occurs from young on-dialysis to young off-dialysis patients. In contrast, transplant rates from non-NYRT donors and low quality NYRT donors increase across all patient types, driving an overall increase in transplants. Approximately optimal Pareto improving priorities yield similar qualitative patterns to those in the unconstrained approximately optimal priorities, but the differences across patient groups are less stark. This is because the mechanism approximates offer rates that are constrained to make each patient type better off ex-ante; therefore, it cannot reallocate as many desirable organs across groups.

This reallocation is achieved by offering high-quality donors to off-dialysis patients. The second group of columns in Table 10 reports the probability that a patient of each type receives an offer from a donor of each type, conditional on the donor being offered to some NYRT patients. These results capture the incentives faced by each patient type by quantifying the stream of offers they receive. The probability that a young off-dialysis patient receives an offer from a young NYRT donor rises from 8.2 percent at baseline to 13.5 percent under approximately optimal priorities; the same probability for a young on-dialysis patient falls from 6.0 percent to 4.6 percent. On-dialysis patients see a uniform decrease in offer probability across donor types under the approximately optimal priorities, and consistent with this, they are worse off. In contrast, among off-dialysis patients, older patients see an increase in offer probabilities for all donor types, while younger patients see a slight decrease for less desirable NYRT donors. Thus the approximately optimal priorities perform age/quality matching in addition to reallocating organs to off-dialysis patients.

A natural question is why reallocation from on-dialysis to off-dialysis patients maximizes patient welfare. Table 11 shows that it is beneficial to reallocate organ offers to off-dialysis patients because they not only value each offer more, but are also more likely to accept. Panel A reports the average value of an offer from a randomly chosen donor to different groups of patients. To illustrate potential reallocation gains, patient types are weighted according to their steady-state proportions under pre-2014 priorities. Off-dialysis patients benefit more than on-dialysis patients from an offer from each donor group. An offer from a young, healthy NYRT donor is equivalent to a donor supply increase of 32.6 percent for an older off-dialysis patient, but only 18.1 percent for an older on-dialysis patient. The difference is even starker for non-NYRT donors, who are less desirable (equivalent to 0.9 and 0.4 percent increases in donor supply, respectively). This is in part because off-dialysis patients have low offer rates in the baseline mechanism, so each additional offer is quite valuable. In addition, healthy donors are valued relatively more highly by young off-dialysis patients, which explains the age/quality matching under approximately optimal priorities. Panel B shows that off-dialysis patients are also more likely to accept an offer from a randomly chosen donor in the baseline mechanism. For example, off-dialysis patients would accept 1.2 percent of non-NYRT donors, while on-dialysis patients would accept 0.8 to 0.9 percent of them. Differences are similar for other donor types. Because of the limit on the number of offers than can be made, offering

the organs to patients more likely to accept increases transplant rates, and therefore expected welfare conditional on match quality. Approximately optimal priorities performs much better than greedy priorities in part because greedy priorities only considers transplant values and not dynamic incentives, and therefore misses many of these potential gains.

Finally, one concern is that using dialysis status to determine priorities may result in agents trying to game the system. While addressing this issue is beyond the scope of this paper, we make two observations. First, the post-2014 mechanism implicitly uses dialysis status to determine priority through the Expected Post Transplant Survival model. Second, during our sample period, patients were allowed to register on the list as soon as kidney function was sufficiently low (below a glomerular filteration rate of 20), but dialysis may not be necessary for all patients (for example, if the glomerular filteration rate is above 15). Therefore, if gaming is of concern, then priority could be alternatively awarded based on measured kidney function.

8 Conclusion

Previous reforms of organ allocation systems have been assisted by a simple simulation model that does not account for the dependence of accept/reject rules on agents' incentives in various mechanisms. This paper develops an empirically tractable method for estimating the value of various assignments in dynamic assignment systems, shows how to compute equilibria of counterfactual mechanisms, and computes optimal solutions. Our results show that accounting for changes in agents' incentives is important in predicting outcomes.

Moreover, we find that there is significant scope for further improving the deceased donor kidney allocation mechanism. There exist mechanisms that improve the average patient's welfare by as much as the equivalent of a 14.2 percent increase in donor arrival rates while also reducing organ waste. An allocation system that also ensures that no patient type is worse off can realize the majority of these gains. In contrast, the reforms implemented in 2014 were mostly redistributive, and both the pre- and the post-2014 systems are similar to a first come first served system. Alternatively, if our measure of patient welfare is of no concern and the only goal is to reduce organ discards, it is possible to do so by 26.6 percent using a last come first served system.

These findings suggest significant scope for using empirical approaches to improve dynamic assignment mechanisms. While the specific objective function for a designer may differ from those considered considered here, our methods can inform design once a well-defined objective and a set of acceptable mechanisms has been specified. Our approach can be used to both evaluate outcomes from specific proposed mechanisms and to find optimal solutions.

Our approach makes several simplifying assumptions. First, and foremost, we assume that agent beliefs are passive and we analyze steady state equilibria. Relaxing these assumptions is

an important direction for future work. Second, we do not allow for patient-level unobserved heterogeneity. Including this feature, potentially with time-varying components, may be important when applying these methods to other empirical contexts or data environments. Third, we assume that arrival rates of patients and objects are exogenous, an assumption reasonable in our setting, but one that may be particularly important to relax in other applications. Finally, the outcomes we primarily focus on are organ waste and overall patient welfare. Nonetheless, our framework can be used to evaluate the effects of various mechanisms on other well-defined criteria that may be important for policy. These include equity or effects on predicted medical outcomes. Enriching this framework to incorporate these features is left for future work.

Empirical approaches for evaluating dynamic assignment systems are particularly important because previous theoretical approaches have not provided sharp guidance on these mechanisms. Models that emphasize heterogeneity in match value find that systems similar to first come first served are most efficient (Bloch and Cantala, 2017), while those that emphasize heterogeneity in object quality suggest that last come first served be adopted (Su and Zenios, 2004). Our results in this context find significant heterogeneity in match value and therefore favor mechanisms of the former kind because they induce agent selectivity. Nonetheless, we find that targeting offers using observable characteristics and equilibrium considerations promises additional improvements. These findings demonstrate the large scope for empirical work on understanding these designs more broadly than in the context of deceased donor kidney allocation.

References

Abdulkadiroglu, Atila and Tayfun Sönmez, "School Choice: A Mechanism Design Approach," *American Economic Review*, may 2003, *93* (3), 729–747.

_ **, Nikhil Agarwal, and Parag A. Pathak**, "The Welfare Effects of Coordinated School Assignment: Evidence from the NYC High School Match," *American Economic Review, forthcoming*, 2017.

Agarwal, Nikhil, "An Empirical Model of the Medical Match," *American Economic Review*, 2015, *105* (7), 1939–78.

_ **and Paulo Somaini**, "Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism," *Econometrica*, 2018, *86* (2), 391–444.

_ **, Itai Ashlagi, Eduardo Azevedo, Clayton Featherstone, and Omer Karaduman**, "Market Failure in Kidney Exchange," 2018.

_ , _ , **Paulo Somaini, and Daniel Waldinger**, "Dynamic Incentives in Waitlist Mechanisms," *American Economic Association: Papers and Proceedings*, 2018.

Aguirregabiria, Victor and Junichi Suzuki, "Identification and counterfactuals in dynamic models of market entry and exit," *Quantitative Marketing and Economics*, sep 2014, *12* (3), 267–304.

_ **and Pedro Mira**, "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, jan 2007, *75* (1), 1–53.

Aliprantis, Charalambos D. and Kim C. Border, *Infinite dimensional analysis : a hitchhiker's guide*, 3rd ed., Springer, 2006.

Arcidiacono, Peter and Robert A. Miller, "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity," *Econometrica*, 2011, *79* (6), 1823–1867.

_ **, Patrick Bayer, Jason R. Blevins, and Paul B. Ellickson**, "Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition," *The Review of Economic Studies*, jul 2016, *83* (3), 889–931.

Arnosti, Nick and Peng Shi, "How (Not) to Allocate Affordable Housing," feb 2017.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin, "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 2007, *75* (5), 1331–1370.

Bertsimas, Dimitris, Vivek F Farias, and Nikolaos Trichakis, "Fairness, efficiency, and flexibility in organ allocation for kidney transplantation," *Operations Research*, 2013, *61* (1), 73–87.

Bloch, Francis and David Cantala, "Dynamic Assignment of Objects to Queuing Agents," *American Economic Journal: Microeconomics*, feb 2017, *9* (1), 88–122.

Canone, Clement, "A short note on Poisson tail bounds," 2017.

Connault, Benjamin, "Hidden Rust Models," 2016.

Danovitch, Gabriel M, *Handbook of kidney transplantation*, Lippincott Williams & Wilkins, 2009.

Fack, Gabrielle, Julien Grenet, and Yinghua He, "Beyond Truthtelling," 2015.

Fershtman, Chaim and Ariel Pakes, "Dynamic Games With Asymmetric Information: A Framework for Empirical Work," *Quarterly Journal of Economics*, 2012, pp. 1611–1661.

Gelman, Andrew, John B Carlin, Hal S Stern, and Donald B Rubin, *Bayesian data analysis*, third ed., Chapman & Hall/CRC Boca Raton, FL, USA, 2014.

Geyer, Judy and Holger Sieg, "Estimating a model of excess demand for public housing," *Quantitative Economics*, nov 2013, *4* (3), 483–513.

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, "OPTN/SRTR 2015 Annual Data Report: Kidney," *American Journal of Transplantation*, jan 2017, *17* (S1), 21–116.

Hayashi, Fumio., *Econometrics.*, Princeton University Press, 2001.

Held, Philip J, F McCormick, A Ojo, and John P Roberts, "A cost-benefit analysis of government compensation of kidney donors," *American Journal of Transplantation*, 2016, *16* (3), 877–885.

Hopenhayn, Hugo A., "Entry, Exit, and firm Dynamics in Long Run Equilibrium," *Econometrica*, sep 1992, *60* (5), 1127.

Hotz, V J and R A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

Hu, Yingyao and Susanne M. Schennach, "Instrumental Variable Treatment of Nonclassical Measurement Error Models," *Econometrica*, 2008, *76* (1), 195–216.

Irwin, F D, A F Bonagura, S W Crawford, and M Foote, "Kidney paired donation: a payer perspective," *American Journal of Transplantation*, 2012, *12* (6), 1388–1391.

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, "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*, aug 2014, *25* (8), 1842–1848.

Kalouptsidi, Myrto, Paul T Scott, and Eduardo Souza-Rodrigues, "Identification of Counterfactuals in Dynamic Discrete Choice Models," *NBER Working Paper Series*, 2015, (21527).

Kong, N, A J Schaefer, B Hunsaker, and M S Roberts, "Maximizing the efficiency of the US liver allocation system through region design," *Management Science*, 2010, *56* (12), 2111–2122.

Leshno, Jacob, "Dynamic Matching in Overloaded Waiting Lists," 2017.

Magnac, Thierry and David Thesmar, "Identifying Dynamic Discrete Decision Processes," *Econometrica*, mar 2002, *70* (2), 801–816.

McCulloch, Robert and Peter E Rossi, "An exact likelihood analysis of the multinomial probit model," *Journal of Econometrics*, sep 1994, *64* (1-2), 207–240.

Miller, Robert A., "Job Matching and Occupational Choice," 1984.

of Transplant Recipients SRTR, Scientific Registry, "Kidney-Pancreas Simulated Allocation Model," Technical Report 2015.

OPTN, "Organ Procurement and Transplantation Network Policies," Technical Report 2014.

_ , "Organ Procurement and Transplantation Network Policies," Technical Report 2017.

OPTN, Organ Procurement Transplantation Network, "Concepts for Kidney Allocation," Technical Report 2011.

_ **and Scientific Registry of Transplant Recipients SRTR**, "OPTN/SRTR 2010 Annual Data Report," Technical Report, Department of Health and Human Services, Health Resources and Services Administration, Healthcare Systems Bureau, Division of Tranplantation, Rockville, MD 2011.

Orandi, B J, J M Garonzik-Wang, Allan B Massie, Andrea A Zachary, J R Montgomery, K J Van Arendonk, Mark D Stegall, S C Jordan, J Oberholzer, T B Dunn, and Others, "Quantifying the risk of incompatible kidney transplantation: a multicenter study," *American Journal of Transplantation*, 2014, *14* (7), 1573–1580.

Pakes, Ariel, "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, jul 1986, *54* (4), 755.

_ **and Paul McGuire**, "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'curse' of Dimensionality," *Econometrica*, sep 2001, *69* (5), 1261–1281.

_ **, Michael Ostrovsky, and Steven Berry**, "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)," *The RAND Journal of Economics*, jun 2007, *38* (2), 373–399.

Pesendorfer, Martin and Philipp Schmidt-Dengler, "Asymptotic Least Squares Estimators for Dynamic Games," *Review of Economic Studies*, jul 2008, *75* (3), 901–928.

Rao, Prakasa B. L. S., *Identifiability in Stochastic Models: Characterization of Probability Distributions*, New Delhi, India: Academic Press, Inc., 1992.

Reeling, Carlson J and Valentin Verdier, "Welfare Effects of Dynamic Assignment: Evidence from the Michigan Bear Lottery," 2018.

Roth, Alvin E., Tayfun Sonmez, and M. Utku Ünver, "Kidney Exchange," *The Quarterly Journal of Economics*, 2004, *119* (2), 457 – 488.

Roth, Alvin E, Tayfun Sönmez, and M. Utku Ünver, "Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences," *American Economic Review*, jun 2007, *97* (3), 828–851.

Rust, J, *Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher*, Vol. 55, Econometrica, 55:999â1033, 1987.

Sieg, Holger and Chamna Yoon, "Waiting for Affordable Housing," 2016.

Su, X and S A Zenios, "Recipient choice can address the efficiency-equity trade-off in kidney transplantation: A mechanism design model," *Management Science*, 2006, *52* (11), 1647–1660.

_ , _ , and G M Chertow, *Incorporating recipient choice in kidney transplantation*, Vol. 15 2004.

Su, Xuanming and Stefanos Zenios, "Patient Choice in Kidney Allocation: The Role of the Queueing Discipline," *Manufacturing & Service Operations Management*, oct 2004, *6* (4), 280–301.

Thakral, Niel, "The Public-Housing Allocation Problem: Theory and Evidence from Pittsburgh," 2016.

USRDS, United States Renal Data System, *2018 USRDS annual data report: Epidemiology of kidney disease in the United States*, National Institutes of Health, National Institute of Diabetes and Digestive and Kidney Diseases, Bethesda, MD, 2018.

van der Vaart, A. W., *Asymptotic Statistics*, Cambridge University Press, 2000.

van Ommeren, Jos, Arno van der Vlist, Jos van Ommeren, and Arno van der Vlist, *Journal of urban economics.*, Vol. 92, Academic Press, 2016.

Waldinger, Daniel, "Targeting In-Kind Transfers Through Market Design: A Revealed Preference Analysis of Public Housing Allocation," 2017.

Weintraub, Gabriel, C. Lanier Benkard, and Ben van Roy, "Markov Perfect Industry Dynamics With Many Firms," *Econometrica*, 2008, *76* (6), 1375–1411.

Wilk, Amber R, John Beck, and Anna Y. Kucheryavaya, "The Kidney Allocation System," Technical Report, United Network for Organ Sharing 2017.

Wolfe, R. A., K. P. McCullough, D. E. Schaubel, J. D. Kalbfleisch, S. Murray, M. D. Stegall, and A. B. Leichtman, "Calculating Life Years from Transplant (LYFT): Methods for Kidney and Kidney-Pancreas Candidates," *American Journal of Transplantation*, apr 2008, *8* (4p2), 997–1011.

Wolfe, R.A., V.B. Ashby, E.L. Milford, A.O. Ojo, R.E. Ettenger, and Others, "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*, 1999, *341* (23), 1725–1730.

Wolpin, Kenneth I., "An Estimable Dynamic Stochastic Model of Fertility and Child Mortality," 1984.

Zenios, S A, "Models for Kidney Allocation," 2004, *70*, 537–554.

_ , G M Chertow, and L M Wein, *Dynamic allocation of kidneys to candidates on the transplant waiting list*, Vol. 48 2000.

Zhang, J, *The sound of silence: Observational learning in the US kidney market*, Vol. 29 2010.

Table 1: Patient Characteristics

	Patient Stocks, Arrivals, and Departures							
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Panel A: Patient Stocks								
	January 1, 2010		January 1, 2011		January 1, 2012		January 1, 2013	
Number of Patients	4018		4252		4481		4632	
Years on List	2.00	1.81	2.20	1.84	2.27	1.88	2.37	1.94
Years on Dialysis	3.24	3.64	3.32	3.76	3.23	3.73	3.24	3.62
Prior Transplant	14.7%	35.4%	14.3%	35.1%	14.1%	34.9%	13.9%	34.6%
Current Age	53.6	13.4	54.0	13.3	54.0	13.4	54.1	13.4
Calculated Panel Reactive Antibodies (CPRA)	9.6%	25.8%	11.0%	26.8%	13.2%	29.1%	14.2%	30.1%
Body Mass Index (BMI) at Arrival	27.8	5.9	27.5	5.7	27.7	6.0	27.9	5.8
Total Serum Albumin	4.0	0.6	4.0	0.6	4.0	0.6	4.1	0.6
Diabetic Patient	39.7%	48.9%	39.9%	49.0%	40.0%	49.0%	40.4%	49.1%
Body Mass Index (BMI) at Arrival	27.8	6.0	27.9	5.9	27.9	6.0	28.0	5.9
On Dialysis at Arrival	76.1%	42.6%	73.5%	44.1%	71.3%	45.3%	69.6%	46.0%
Panel B: Patient Arrivals and Departures								
	Year 2010		Year 2011		Year 2012		Year 2013	
Number of Patients Arriving	1434		1563		1549		1353	
Number of Patients Departing	1145		1274		1325		1278	
Age at Departure	54.3	15.4	54.6	15.0	54.7	15.2	54.3	15.0
CPRA at Departure	9.7%	25.6%	11.4%	26.8%	12.8%	28.4%	13.5%	29.6%
Years on Dialysis at Departure	3.49	3.76	3.61	3.74	3.50	4.00	3.33	3.76
Panel C: Departures by Reason								
	Received Deceased Donor Transplant		Received Live Donor Transplant		Died or Too Sick to Transplant		Departed for Other Reason	
Number of Patients	2141		1009		1091		781	
Years on Dialysis	4.04	3.90	1.13	2.19	4.60	4.01	3.43	3.72
Years on Waitlist	3.08	2.21	0.93	1.11	3.05	2.07	2.85	1.93
Age at Departure	54.3	15.2	47.7	15.4	61.5	12.2	53.7	14.1

Notes: 9,917 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.

Table 2: Donor Characteristics

	Donors									
	All		Any Kidney(s) Discarded				Last Offer Category			
			Yes		No		Local or Perfect Tissue Type Match		Non-Local, Some Tissue Type Mismatch	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Panel A: Donors Recovered in NYRT, By Number of Organs Allocated or Category of Last Offer										
Number of Donors Per Year	196		44.75		151.25		161		35	
Median Number of Offers per Donor	27.0		496.0		15.0		14.0		1217.5	
Number of Offers per Donor	431.3	1422.0	1472.7	2662.0	123.1	341.6	72.5	136.6	2081.4	2822.5
Number of Kidneys Transplanted per Donor	1.55	0.78	0.35	0.48	1.90	0.42	1.81	0.51	0.35	0.73
Donor Age	43.4	17.8	56.0	14.2	39.7	17.1	41.0	17.0	54.6	17.4
Cause of Death -- Head Trauma	26.1%	44.0%	11.7%	32.3%	30.4%	46.0%	29.3%	45.6%	11.4%	31.9%
Cause of Death -- Stroke	43.1%	49.6%	60.9%	48.9%	37.9%	48.5%	39.6%	48.9%	59.3%	49.3%
Diabetic Donor	13.8%	34.5%	25.1%	43.5%	10.4%	30.6%	10.6%	30.8%	28.6%	45.3%
Hypertensive Donor	37.2%	48.4%	60.9%	48.9%	30.2%	46.0%	31.8%	46.6%	62.1%	48.7%
Expanded Criteria Donor (ECD)	29.7%	45.7%	58.1%	49.5%	21.3%	41.0%	23.3%	42.3%	59.3%	49.3%
Donation after Cardiac Death (DCD)	9.1%	28.7%	12.3%	32.9%	8.1%	27.3%	8.7%	28.2%	10.7%	31.0%
Donor Creatinine	1.3	1.5	1.5	1.2	1.3	1.6	1.2	1.4	1.8	1.9
Panel B: All Donors, By Number of Organs Allocated or Category of Last Offer										
Number of Donors Per Year	1465.75		906.75		559		201.25		1264.5	
Median Number of Offers per Donor	725.0		1200.0		239.0		15.0		987.5	
Number of Offers per Donor	1617.3	2560.9	2196.6	2990.7	677.7	1123.7	87.6	160.6	1860.8	2677.0
Number of Kidneys Transplanted per Donor	0.75	0.90	0.21	0.41	1.64	0.76	1.76	0.54	0.59	0.84
Donor Age	48.0	18.5	52.6	15.7	40.6	20.2	41.1	17.2	49.1	18.4
Cause of Death -- Head Trauma	19.2%	39.4%	15.4%	36.1%	25.4%	43.6%	28.3%	45.1%	17.8%	38.2%
Cause of Death -- Stroke	48.1%	50.0%	55.1%	49.7%	36.8%	48.2%	40.0%	49.0%	49.4%	50.0%
Diabetic Donor	20.4%	40.3%	25.0%	43.3%	12.8%	33.5%	10.3%	30.4%	22.0%	41.4%
Hypertensive Donor	53.6%	49.9%	63.1%	48.3%	38.3%	48.6%	33.2%	47.1%	56.9%	49.5%
Expanded Criteria Donor (ECD)	44.7%	49.7%	54.5%	49.8%	28.8%	45.3%	23.5%	42.4%	48.0%	50.0%
Donation after Cardiac Death (DCD)	12.6%	33.2%	13.9%	34.6%	10.4%	30.6%	9.3%	29.1%	13.1%	33.8%
Donor Creatinine	1.5	1.3	1.6	1.2	1.5	1.5	1.2	1.3	1.6	1.3

Notes: Panel A consists of all deceased kidney donors (784) recovered in NYRT and offered to NYRT patients between January 1st, 2010 and December 31st, 2013. Panel B includes all donors (5,683) 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 pre-determined 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.

Table 3: Rates of Receiving and Accepting Offers

	Number of Patients	Offer & Acceptance Rates					
		All Donors		NYRT Donors		Perfect Tissue Type Match	
		Annual Rate	% Accepted	Annual Rate	% Accepted	Annual Rate	% Accepted
				Panel A: All Offers			
All	9917	218.4	0.15%	40.1	0.73%	0.094	10.6%
On Dialysis at Registration	6703	217.2	0.16%	42.1	0.76%	0.090	11.3%
Not on Dialysis at Registration	3214	220.9	0.12%	36.1	0.67%	0.104	9.4%
Peak CPRA < 0.80	8864	234.5	0.10%	42.8	0.49%	0.093	9.3%
Peak CPRA >= 0.80	1053	83.1	1.03%	17.3	4.39%	0.103	15.8%
				Panel B: Offers that Met Screening Criteria			
All	9917	104.1	0.34%	23.5	1.35%	0.051	21.8%
On Dialysis at Registration	6703	102.7	0.37%	24.8	1.39%	0.042	24.6%
Not on Dialysis at Registration	3214	107.1	0.27%	20.6	1.25%	0.069	17.5%
Peak CPRA < 0.80	8864	112.3	0.22%	25.1	0.89%	0.049	19.7%
Peak CPRA >= 0.80	1053	35.8	2.51%	9.8	8.11%	0.068	29.0%
				Panel C: Offers Within the First 10 Positions that Met Screening Criteria			
All	9917	0.8	25.38%	0.8	26.68%	0.021	48.3%
On Dialysis at Registration	6703	0.8	25.42%	0.8	26.68%	0.017	48.5%
Not on Dialysis at Registration	3214	0.8	25.27%	0.8	26.71%	0.029	47.9%
Peak CPRA < 0.80	8864	0.8	15.79%	0.8	16.61%	0.019	51.1%
Peak CPRA >= 0.80	1053	0.9	44.10%	0.9	46.24%	0.038	43.6%

Notes: there were 2,850,572 offers made to NYRT patients between January 1st, 2010 and December 31st, 2013. Panel C 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 pre-determined criteria for acceptable donors. “Annual Rate” columns report annual offer rates computed by patient and then averaged across patients. Peak CPRA is the highest level of Calculated Panel Reactive Antibodies (CPRA) recorded for each patient during the period of study.

Table 4: Evidence on Mismatch

	Patient Age					On Dialysis at Registration	
	0-17	18-34	35-49	50-64	65+	Yes	No
Share of Patient Population	2.2%	11.1%	27.2%	41.5%	18.0%	67.7%	32.3%
Share of Deceased Donor Transplants	4.7%	9.6%	26.1%	41.9%	17.7%	72.6%	27.4%
Share of Ideal Donor Transplants	6.7%	9.8%	26.6%	39.8%	17.1%	73.1%	26.9%
Panel A: Outcomes of all Patients							
Received Deceased Donor Transplant	47.9%	19.2%	21.3%	22.4%	21.9%	23.8%	18.9%
Received Live Donor Transplant	19.4%	18.6%	11.5%	8.9%	6.5%	7.2%	17.2%
Still Waiting	28.9%	48.6%	50.4%	48.8%	44.0%	47.7%	48.3%
Died or Too Sick to Transplant	1.9%	4.1%	6.8%	12.9%	20.0%	13.4%	6.9%
Departed for Another Reason	1.9%	9.5%	10.1%	7.0%	7.6%	7.8%	8.7%
Panel B: Donor Age and Quality							
Donor is Ideal	97.0%	70.2%	69.7%	64.9%	66.2%	68.9%	67.1%
Donor Age 0-17	23.8%	14.6%	7.2%	4.3%	4.5%	6.6%	8.0%
Donor Age 18-35	73.3%	37.6%	26.7%	17.3%	12.9%	22.8%	25.4%
Donor Age 35-49	3.0%	30.2%	34.9%	28.7%	17.7%	28.1%	25.2%
Donor Age 50-64	0.0%	16.6%	29.2%	42.1%	52.0%	36.7%	34.4%
Donor Age >= 65	0.0%	1.0%	2.0%	7.6%	12.9%	5.7%	7.0%
Panel C: Donor Age, Ideal Donors							
Donor Age 0-17	23.5%	17.4%	9.5%	5.8%	5.6%	8.5%	10.7%
Donor Age 18-35	73.5%	38.9%	23.7%	17.2%	11.2%	23.3%	24.9%
Donor Age 35-49	3.1%	26.4%	33.9%	22.5%	14.7%	24.1%	21.1%
Donor Age 50-64	0.0%	16.0%	30.6%	45.1%	51.8%	37.3%	34.5%
Donor Age >= 65	0.0%	1.4%	2.3%	9.4%	16.7%	6.8%	8.9%

Panel A sample is all NYRT patients on the waiting list between January 1st, 2010 and December 31st, 2013. For other panels, the sample is NYRT patients who received deceased donor transplants between 2010 and 2013. An ideal donor has no history of diabetes; is non-DCD; has creatinine below 3; and is Hepatitis C negative.

Table 5: Acceptance Rate by Past Offer Rates

	Dependent Variable: Current Offer Accepted						
	All Offers					Ideal Donors	NYRT Donors
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Time Since Last Offer (Years)	0.0372	0.00227			0.00217	-0.000324	-0.00948
	(0.00875)	(0.00787)			(0.00786)	(0.00505)	(0.00663)
Time Since Last Two Offers (Years)			0.00646				
			(0.00880)				
Time Since Last Five Offers (Years)				0.0151			
				(0.0118)			
Last Offer Donor Age					-0.00000369		
					(0.00000304)		
Last Offer from Diabetic Donor					-0.0000854		
					(0.000116)		
Last Offer from Expanded Criteria Donor					-0.000174		
					(0.000116)		
Last Offer from Donation after Cardiac Death					0.00000289		
					(0.000119)		
Variables Affecting Priority		X	X	X	X	X	X
Patient Characteristics		X	X	X	X	X	X
Donor and Match Characteristics		X	X	X	X	X	X
Listing Center Fixed Effects		X	X	X	X	X	X
Observations	2793098	2831262	2831262	2821641	2831262	1083686	526233
R-squared	0.098	0.097	0.097	0.099	0.099	0.133	0.116
Mean Acceptance Rate	0.15%	0.15%	0.15%	0.15%	0.15%	0.25%	0.73%
Mean Time Since Last N Offers	0.005	0.005	0.005	0.005	0.005	0.012	0.023
S.D. Time Since Last N Offers	0.016	0.016	0.011	0.008	0.016	0.028	0.048

Notes: estimates from a linear probability model of offer acceptance as a function of the patient's recent offer history. Time Since Last N Offers measures the average number of years since the patient's previous offers, averaged over their last N offers. Time Since Last Offer, Including Inactive Periods counts inactive days as well as active days on the waitlist. Column (1) considers all offers and includes no controls for current offer characteristics. Columns (2) - (7) control for current patient, donor, and match characteristics. Column (5) includes controls for donor characteristics of the patient's previous offer. Column (6) restricts to offers from ideal donors, and Column (7) restricts to NYRT donors. Controls are as described in the notes for Appendix Table B.1. An ideal donor has no history of diabetes; is non-DCD; has creatinine below 3; and is Hepatitis C negative.

Table 6: Survival Model Estimates

	Gompertz (1)	Weibull (2)	Cox (3)
Diabetic Patient	0.112	0.104	0.115
	(0.0334)	(0.0334)	(0.0334)
Bloodtype A Patient	0.147	0.115	0.154
	(0.0436)	(0.0435)	(0.0436)
Bloodtype O Patient	0.0146	0.0145	0.0132
	(0.0390)	(0.0390)	(0.0391)
Calculated Panel Reactive Antibodies (CPRA)	-0.000601	-0.000692	-0.000775
	(0.00150)	(0.00150)	(0.00150)
CPRA = 0	0.182	0.171	0.173
	(0.0737)	(0.0736)	(0.0738)
CPRA - 80 if CPRA>=80	-0.0182	-0.0154	-0.0172
	(0.00652)	(0.00652)	(0.00652)
Age (at Registration)	-0.0261	-0.0213	-0.0210
	(0.0152)	(0.0153)	(0.0153)
Age - 18 if Age>=18	0.0230	0.0194	0.0186
	(0.0186)	(0.0187)	(0.0187)
Age - 35 if Age>=35	-0.00964	-0.0120	-0.0104
	(0.00960)	(0.00960)	(0.00960)
Age - 50 if Age>=50	0.0232	0.0230	0.0239
	(0.00723)	(0.00722)	(0.00724)
Age - 65 if Age>=65	0.0243	0.0234	0.0241
	(0.00923)	(0.00922)	(0.00925)
Prior Transplant	0.0205	0.0329	0.0289
	(0.0547)	(0.0545)	(0.0547)
Body Mass Index (BMI)	-0.0131	-0.0123	-0.0131
	(0.00643)	(0.00642)	(0.00643)
Missing BMI	0.000107	0.132	-0.0477
	(0.199)	(0.199)	(0.200)
BMI >= 18.5	-0.0653	-0.0743	-0.0675
	(0.105)	(0.106)	(0.106)
BMI >= 25	-0.0140	-0.0177	-0.0133
	(0.0492)	(0.0492)	(0.0492)
BMI >= 30	0.0285	0.0223	0.0269
	(0.0595)	(0.0594)	(0.0595)
Total Serum Albumin	-0.184	-0.180	-0.175
	(0.0543)	(0.0543)	(0.0543)
Missing Total Serum Albumin	-0.675	-0.598	-0.621
	(0.187)	(0.187)	(0.187)
Total Serum Albumin >= 3.7	-0.0836	-0.0817	-0.0871
	(0.0589)	(0.0590)	(0.0589)
Total Serum Albumin >= 4.4	0.0422	0.0308	0.0415
	(0.0507)	(0.0506)	(0.0507)
On Dialysis at Registration	-1.087	-1.067	-1.086
	(0.0659)	(0.0660)	(0.0659)
Log Years on Dialysis at Registration	0.108	0.109	0.107
	(0.0103)	(0.0103)	(0.0103)
Log Years on Dialysis at Registration * Over 5 Years	-0.129	-0.123	-0.137
	(0.105)	(0.105)	(0.105)
Constant	-5.151	-4.702	
	(0.328)	(0.337)	
Constant (delta(2))	0.000106		
	(0.0000207)		
Constant (delta)		-0.0635	
		(0.0142)	
Observations	9917	9917	9917

Table 7: Conditional Choice Probability of Acceptance (select co-efficients)

	Base Specification		Unobserved Heterog.		Waiting Time + UH	
	(1)		(2)		(3)	
Calculated Panel Reactive Antibody (CPRA)	0.60	(0.05)	0.67	(0.06)	0.54	(0.08)
Donor Age < 18	0.30	(0.10)	-0.10	(0.18)	-0.07	(0.19)
Donor Age 18-35	0.60	(0.12)	0.09	(0.18)	0.14	(0.18)
Donor Age 50+	-0.83	(0.15)	-0.81	(0.20)	-0.85	(0.21)
Expanded Criteria Donor (ECD)	-0.15	(0.02)	-0.51	(0.08)	-0.54	(0.08)
Donation from Cardiac Death (DCD)	-0.11	(0.02)	-0.49	(0.08)	-0.49	(0.07)
Perfect Tissue Type Match	2.28	(0.31)	2.79	(0.42)	2.75	(0.43)
2 A Mismatches	-0.08	(0.01)	-0.01	(0.02)	-0.01	(0.02)
2 B Mismatches	0.06	(0.02)	0.03	(0.02)	0.02	(0.02)
2 DR Mismatches	-0.07	(0.02)	-0.06	(0.02)	-0.06	(0.02)
Regional Offer	-1.34	(0.05)	-2.69	(0.17)	-2.75	(0.17)
National Offer	-1.50	(0.04)	-2.92	(0.11)	-2.95	(0.12)
Non-NYRT Donor, NYRT Match Run	1.21	(0.02)	1.97	(0.05)	2.01	(0.06)
Patient on Dialysis at Registration	-0.04	(0.01)	-0.12	(0.02)	-0.11	(0.02)
Log Waiting Time (years)					-0.01	(0.05)
Log Waiting Time * Over 1 Year					-0.10	(0.06)
Log Waiting Time * Over 2 Years					-0.17	(0.12)
Log Waiting Time * Over 3 Years					0.42	(0.11)
Log Dialysis Time at Registration (Years)	0.05	(0.00)	0.05	(0.00)	0.05	(0.01)
Log Dialysis Time at Registration * Over 5 years	0.48	(0.03)	0.42	(0.03)	0.41	(0.04)
NYRT Donor * Donor Age < 18	-0.05	(0.06)	0.22	(0.20)	0.23	(0.22)
NYRT Donor * Donor Age 18-35	0.10	(0.04)	0.11	(0.13)	0.14	(0.14)
NYRT Donor * Donor Age 50+	-0.42	(0.03)	-0.62	(0.11)	-0.62	(0.13)
Patient Age * Donor Age < 18	-0.01	(0.00)	0.00	(0.00)	0.00	(0.00)
Patient Age * Donor Age 18-35	-0.02	(0.00)	0.00	(0.01)	0.00	(0.01)
Patient Age * Donor Age 50+	0.02	(0.00)	0.02	(0.01)	0.02	(0.01)
Patient Age - 35 if Age >= 35 * Donor Age 18-35					0.00	(0.01)
Patient Age - 35 if Age >= 35 * Donor Age 50+					0.00	(0.01)
Donor Unobservable Std. Dev.			0.98	(0.21)	1.00	(0.23)
Idiosyncratic Shock Std. Dev.	1.00		1.00		1.00	
Acceptance Rate	0.150%		0.150%		0.150%	
Number of Offers	2850572		2850572		2850572	

Table 8: Outcomes in Various Mechanisms

	Waitlist			Characteristics of Transplanted Donors				Change in Equivalent Donor Supply (Mean)		Fraction with $\Delta V_x(0)$	
	Queue Length	Reduction in Discard Rate	Years on Waitlist	Age	Died of Head Trauma	Cardiac Death (DCD)	Hyper-tensive	At Listing $V_x(0)$	Steady State $\bar{V}_x$	> 0%	> -5%
Panel A: Steady State Equilibrium, All Patients											
Pre-2014 Priorities	4992.2	--	2.52	43.8	25.6%	11.6%	38.1%	--	--	--	--
Post-2014 Priorities	4913.8	0.5%	2.49	43.8	25.5%	11.6%	38.2%	-0.3%	-0.9%	35.0%	98.7%
First Come First Served	5095.6	-1.4%	2.55	43.7	25.5%	11.6%	38.0%	0.6%	1.2%	86.3%	95.3%
Last Come First Served	2694.3	26.6%	2.45	45.9	24.3%	11.7%	44.5%	-21.8%	-33.6%	2.0%	9.7%
Greedy Priorities	4917.4	0.3%	2.53	43.6	25.7%	11.4%	37.9%	1.7%	1.2%	62.0%	79.3%
Optimal Assignment	3815.6	13.6%	2.17	44.1	25.4%	11.6%	40.6%	44.4%	28.1%	83.3%	96.3%
Optimal Offer Rates	4438.3	7.5%	2.40	44.1	25.5%	11.6%	40.4%	27.4%	17.4%	44.3%	62.0%
Approximately Optimal Priorities	4358.3	7.8%	2.38	44.2	25.3%	11.5%	40.5%	24.2%	14.2%	41.7%	61.0%
Approx. Opt. Pareto Improving Priorities	4560.9	4.5%	2.41	44.0	25.5%	11.5%	39.5%	15.0%	9.1%	14.7%	99.7%
Panel B: Estimated Choice Probabilities from Pre-2014 Mechanism, All Patients											
Pre-2014 Priorities	5387.6	--	2.86	43.3	25.8%	11.5%	36.5%	--	--	--	--
Post-2014 Priorities	5337.8	0.0%	2.86	43.3	25.8%	11.5%	36.5%	--	--	--	--
First Come First Served	5460.6	-0.1%	2.84	43.3	25.8%	11.5%	36.5%	--	--	--	--
Last Come First Served	5175.0	3.6%	5.39	43.6	25.5%	11.3%	37.7%	--	--	--	--
Greedy Priorities	5264.2	0.9%	2.89	43.4	25.7%	11.4%	36.8%	--	--	--	--

Table 9: Outcomes in Various Mechanisms, by Patient Group

	Waitlist	Characteristics of Transplanted Donors				Change in Equivalent Donor Supply (Mean)		Fraction with $\Delta V_x(0)$	
	Years on Waitlist	Age	Died of Head Trauma	Cardiac Death (DCD)	Hyper-tensive	At Listing $V_x(0)$	Steady State $\bar{V}_x$	> 0%	> -5%
					Panel A: Not on Dialysis at Registration, Age 0-49				
Pre-2014 Priorities	2.40	38.3	32.2%	13.4%	27.5%	--	--	--	--
Post-2014 Priorities	2.37	39.3	33.8%	13.4%	28.5%	0.9%	0.3%	55.3%	97.4%
First Come First Served	2.43	41.0	29.2%	12.7%	35.2%	-5.5%	-4.6%	68.4%	84.2%
Last Come First Served	2.33	42.6	26.6%	13.0%	40.7%	-32.0%	-39.9%	0.0%	7.9%
Greedy Priorities	2.43	38.0	33.4%	11.6%	29.0%	1.2%	1.5%	71.1%	78.9%
Optimal Assignment	2.07	39.5	30.7%	11.9%	33.3%	43.9%	30.7%	81.6%	92.1%
Approximately Optimal Priorities	2.20	38.6	31.4%	11.3%	33.1%	20.4%	13.1%	68.4%	84.2%
Approx. Opt. Pareto Improving Priorities	2.34	37.7	30.9%	12.3%	33.7%	6.1%	4.6%	18.4%	100.0%
					Panel B: Not on Dialysis at Registration, Age 50+				
Pre-2014 Priorities	2.51	47.1	21.6%	12.0%	44.6%	--	--	--	--
Post-2014 Priorities	2.48	46.6	21.8%	11.8%	44.1%	-0.7%	-1.3%	25.4%	100.0%
First Come First Served	2.52	45.6	23.4%	12.0%	40.9%	3.7%	4.1%	95.8%	100.0%
Last Come First Served	2.52	48.4	22.1%	12.3%	48.7%	-30.3%	-41.5%	0.0%	4.2%
Greedy Priorities	2.58	44.7	23.9%	11.8%	38.0%	5.2%	5.8%	59.2%	71.8%
Optimal Assignment	1.87	44.5	25.3%	12.4%	38.7%	96.9%	65.7%	95.8%	100.0%
Approximately Optimal Priorities	2.10	44.7	25.7%	12.7%	37.5%	70.0%	45.7%	74.6%	88.7%
Approx. Opt. Pareto Improving Priorities	2.33	46.1	23.9%	13.5%	39.3%	37.7%	25.9%	38.0%	100.0%
					Panel C: On Dialysis at Registration, Age 0-49				
Pre-2014 Priorities	2.53	40.2	27.5%	11.5%	34.1%	--	--	--	--
Post-2014 Priorities	2.50	40.3	27.2%	11.5%	33.7%	0.0%	-0.6%	43.7%	98.6%
First Come First Served	2.57	40.9	27.0%	11.0%	35.8%	-1.0%	-0.2%	84.5%	93.0%
Last Come First Served	2.49	42.7	26.9%	11.3%	41.1%	-20.3%	-31.3%	1.4%	9.9%
Greedy Priorities	2.49	40.5	27.3%	10.7%	35.4%	-0.7%	-1.8%	70.4%	90.1%
Optimal Assignment	2.31	41.8	26.8%	11.0%	41.1%	6.5%	0.9%	71.8%	93.0%
Approximately Optimal Priorities	2.51	41.5	25.8%	10.9%	41.4%	-7.4%	-7.9%	14.1%	36.6%
Approx. Opt. Pareto Improving Priorities	2.46	39.1	28.1%	10.3%	36.2%	-1.0%	-2.4%	0.0%	98.6%
					Panel D: On Dialysis at Registration, Age 50+				
Pre-2014 Priorities	2.56	46.1	24.3%	10.9%	40.8%	--	--	--	--
Post-2014 Priorities	2.53	46.1	23.6%	10.9%	41.1%	-0.7%	-1.3%	29.2%	98.3%
First Come First Served	2.59	45.3	24.5%	11.4%	38.7%	1.7%	2.3%	87.5%	97.5%
Last Come First Served	2.42	47.2	23.4%	11.3%	45.3%	-14.4%	-28.4%	4.2%	13.3%
Greedy Priorities	2.54	46.6	23.3%	11.5%	42.1%	1.1%	0.2%	55.8%	77.5%
Optimal Assignment	2.25	46.4	23.1%	11.3%	43.7%	35.8%	21.2%	83.3%	97.5%
Approximately Optimal Priorities	2.48	47.2	22.9%	11.2%	44.4%	17.1%	9.1%	30.0%	51.7%
Approx. Opt. Pareto Improving Priorities	2.45	47.6	23.2%	10.8%	43.2%	13.9%	7.4%	8.3%	100.0%

Table 10: Features of Various Mechanisms

	Number of Transplants per Year				Offer Probability (%)				Mean Wait Time at Offer (Years)			
	NYRT Donor			Non NYRT Donor	NYRT Donor			Non NYRT Donor	NYRT Donor			Non NYRT Donor
	Young	Old	Low Quality Donor		Young	Old	Low Quality Donor		Young	Old	Low Quality Donor	
Panel A: Patients not on Dialysis at Registration, Age 0-49												
Pre-2014 Priorities (Baseline)	21.8	9.7	7.5	50.2	8.2	9.8	31.5	48.0	2.81	4.54	3.63	3.52
Greedy Priorities	19.2	11.5	7.9	48.5	6.8	9.6	28.5	41.8	3.63	4.04	3.21	2.96
Approximately Optimal Priorities	20.3	12.2	8.5	67.1	13.5	14.3	29.5	48.5	3.04	3.39	2.78	2.86
Approx. Opt. Pareto Improving Priorities	18.1	12.3	6.4	56.6	9.6	9.7	23.8	46.6	2.80	3.43	2.93	3.01
Panel B: Patients not on Dialysis at Registration, Age 50+												
Pre-2014 Priorities (Baseline)	10.9	20.3	18.7	87.6	5.8	9.9	30.9	43.8	4.43	4.93	3.86	3.82
Greedy Priorities	22.3	24.0	16.3	72.6	4.2	6.7	22.8	31.7	4.64	4.80	3.62	3.31
Approximately Optimal Priorities	37.6	38.3	28.9	119.0	17.1	18.3	38.0	49.3	2.51	2.65	2.42	2.74
Approx. Opt. Pareto Improving Priorities	23.4	29.0	25.4	98.5	8.2	9.7	27.9	43.4	3.50	3.55	2.84	3.01
Panel C: Patients on Dialysis at Registration, Age 0-49												
Pre-2014 Priorities (Baseline)	30.5	24.4	17.8	107.2	6.0	8.9	31.6	44.2	4.09	4.72	3.80	3.75
Greedy Priorities	23.7	25.6	17.5	113.5	7.3	11.7	35.9	53.0	3.90	4.15	3.33	3.17
Approximately Optimal Priorities	7.8	14.5	14.1	141.0	4.6	6.5	20.9	41.4	3.93	4.32	3.55	3.58
Approx. Opt. Pareto Improving Priorities	23.7	22.2	14.2	125.0	6.4	7.2	22.5	42.1	3.69	4.25	2.98	3.19
Panel D: Patients on Dialysis at Registration, Age 50+												
Pre-2014 Priorities (Baseline)	30.0	41.2	42.1	172.1	5.1	9.8	31.4	46.4	4.40	4.86	3.83	3.80
Greedy Priorities	28.0	34.3	44.5	188.2	5.7	10.7	34.6	51.3	3.82	4.15	3.34	3.17
Approximately Optimal Priorities	25.1	32.0	47.2	230.7	5.6	8.8	26.5	46.5	3.77	3.97	3.40	3.46
Approx. Opt. Pareto Improving Priorities	27.3	35.6	47.3	214.7	5.4	9.8	33.7	50.0	3.27	3.55	2.72	2.90

Notes: Number of Transplants per Year measures the number of kidneys from each donor group transplanted into each patient group. Offer Probability is the probability that a patient receives an offer from a donor in each group conditional on the donor being offered to some patients. Mean Wait Time at Offer reports the mean wait time of a patient receiving a randomly selected offer.

Table 11: Sources of Welfare Gains

	NYRT Donor			Non NYRT Donor
	Young	Old	Low Quality Donor	
	Panel A: Value of an Offer			
Patients Not on Dialysis at Registration, Age 0-49	36.7	15.1	2.7	0.7
Patients Not on Dialysis at Registration, Age 50+	32.6	18.6	5.1	0.9
Patients on Dialysis at Registration, Age 0-49	13.6	5.2	0.8	0.2
Patients on Dialysis at Registration, Age 50+	18.1	7.9	2.2	0.4
	Panel B: Percentage Willing to Accept			
Patients Not on Dialysis at Registration, Age 0-49	34.7	16.4	5.3	1.2
Patients Not on Dialysis at Registration, Age 50+	26.7	18.5	6.2	1.2
Patients on Dialysis at Registration, Age 0-49	26.4	14.2	3.2	0.8
Patients on Dialysis at Registration, Age 50+	25.6	15.7	5.3	0.9

Notes: in each cell, quantities are averaged over randomly drawn donors and patients according to their proportions in the Pre-2014 Priorities (Baseline) mechanism. Panel A reports the average equivalent change in donor supply from an offer from each donor type to each patient type. All numbers are reported in percentage (%) units.

Figure 1: Offer Statistics by Position on Waiting List

(a) Waiting Time by Position

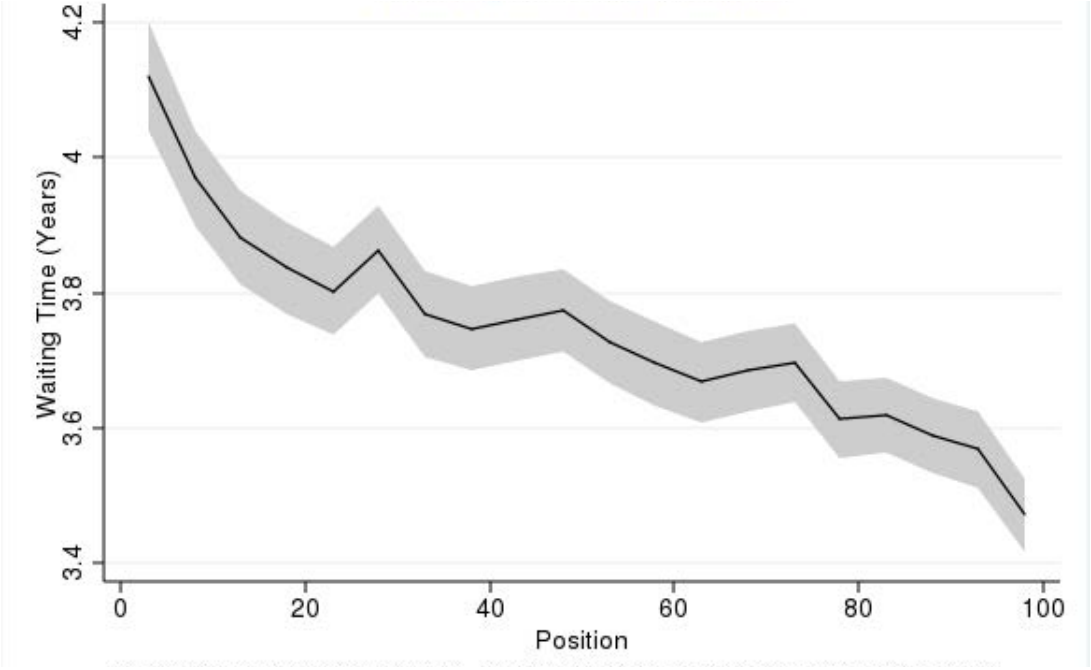

(b) Acceptance Rate by Position

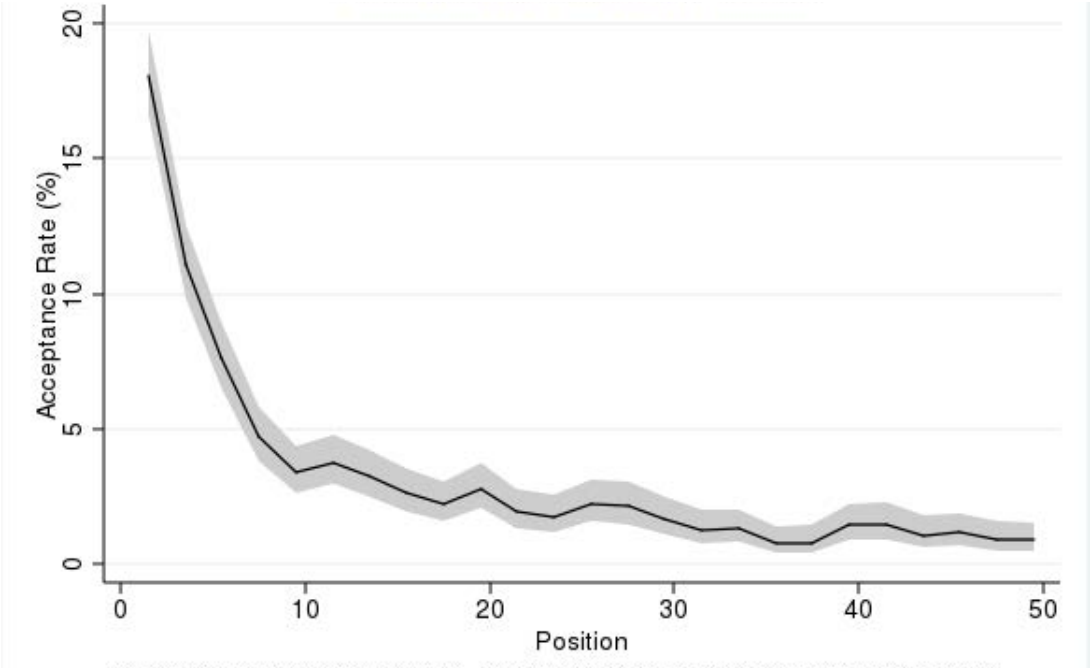

Note: Sample is offers made to NYRT patients between 2010 and 2013, excluding offers that did not meet a patient's pre-set donor screening criteria. Positive crossmatches are counted as acceptances. In each figure, the black line plots the mean among offered patients in each position group, and the shaded region represents pointwise 95 percent confidence intervals.

Figure 2: Model Fit

(a) Fit by position

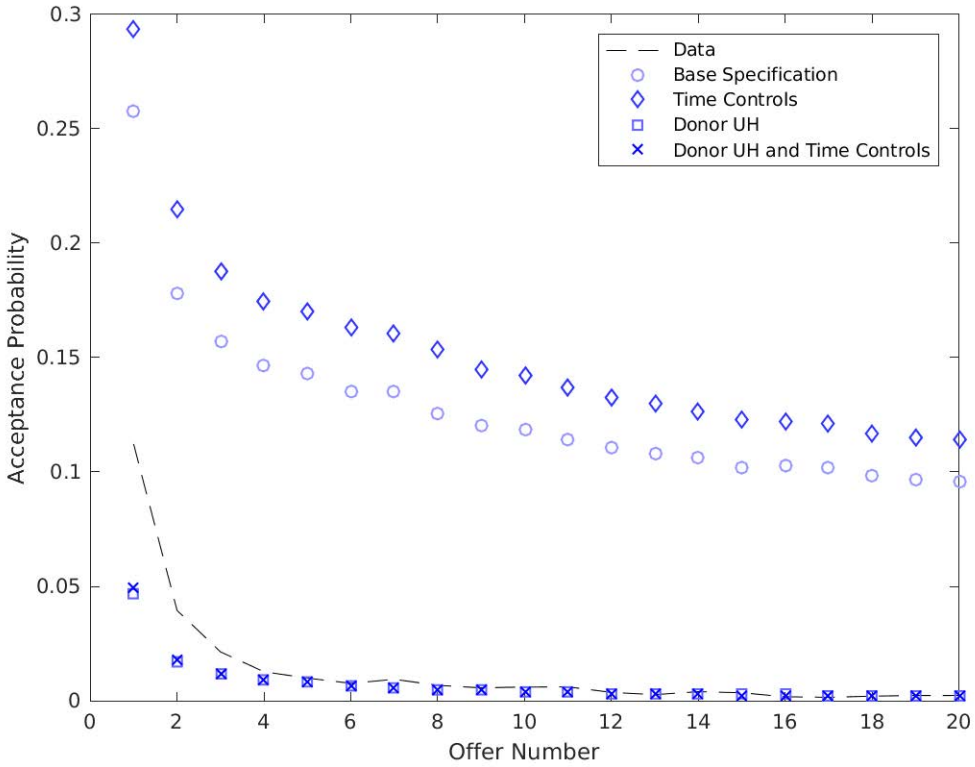

(b) Fit by wait time

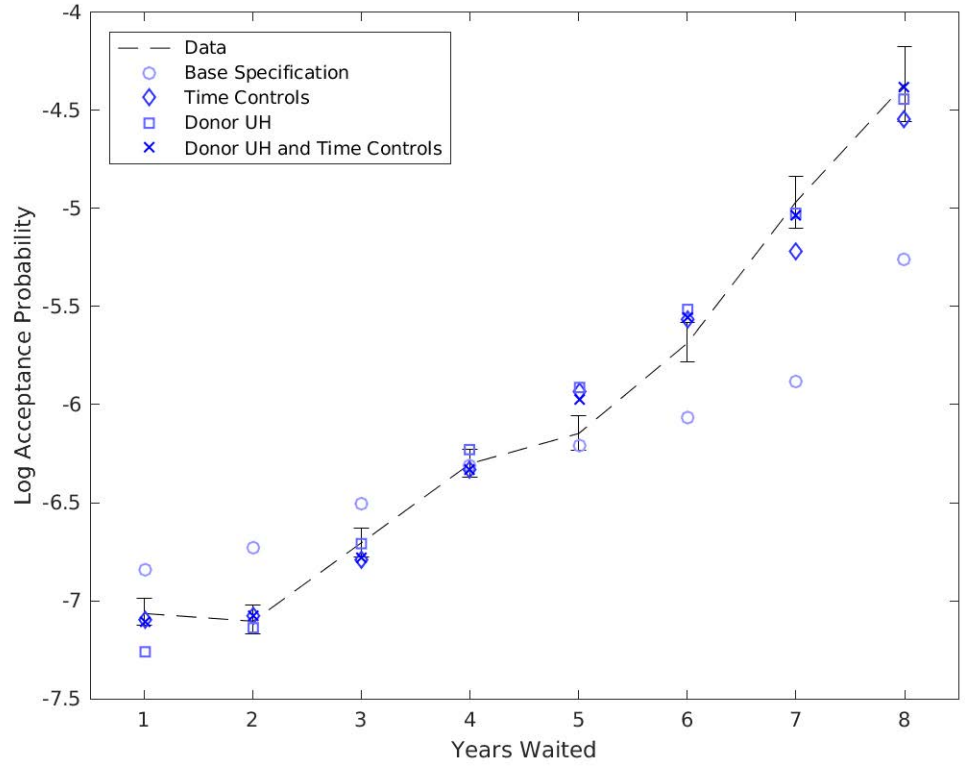

Figure 3: Value of an Organ Offer by Patient and Donor Age

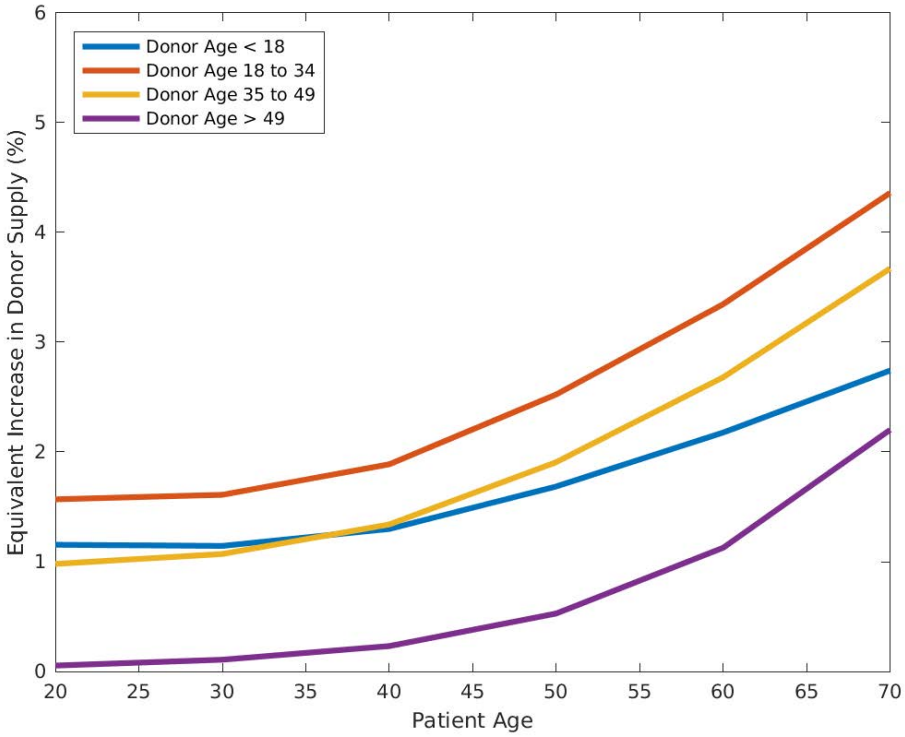

Notes: Patient and donor characteristics other than those explicitly varied are set to median values from the NYRT sample. Value functions are computed at registration.

Figure 4: Patient Age and Dialysis

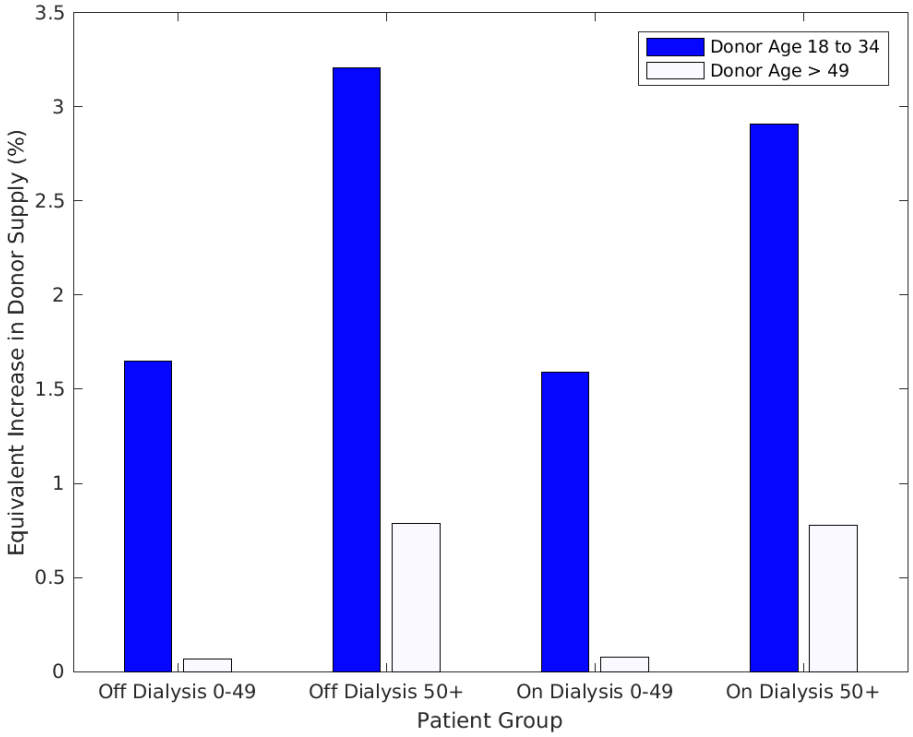

Notes: Patient and donor characteristics other than those explicitly varied are set to median values from the NYRT sample. Value functions are computed at registration.

Figure 5: Value of an Organ Offer by Match Quality

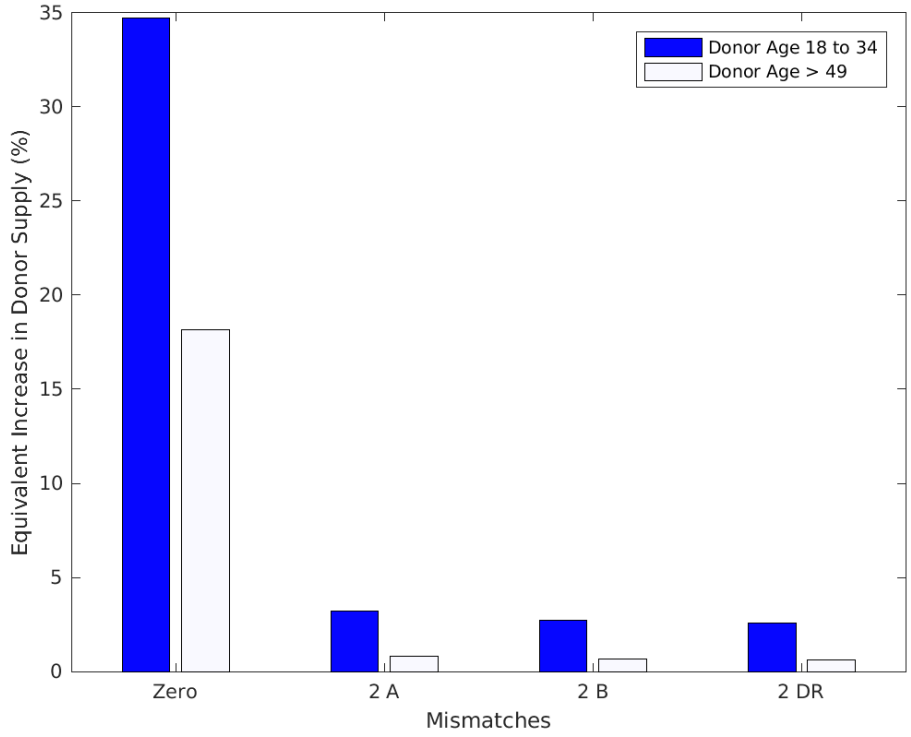

Notes: Patient and donor characteristics other than those explicitly varied are set to median values from the NYRT sample. Value functions are computed at registration.

Figure 6: Welfare and Organ Utilization under Alternative Mechanisms

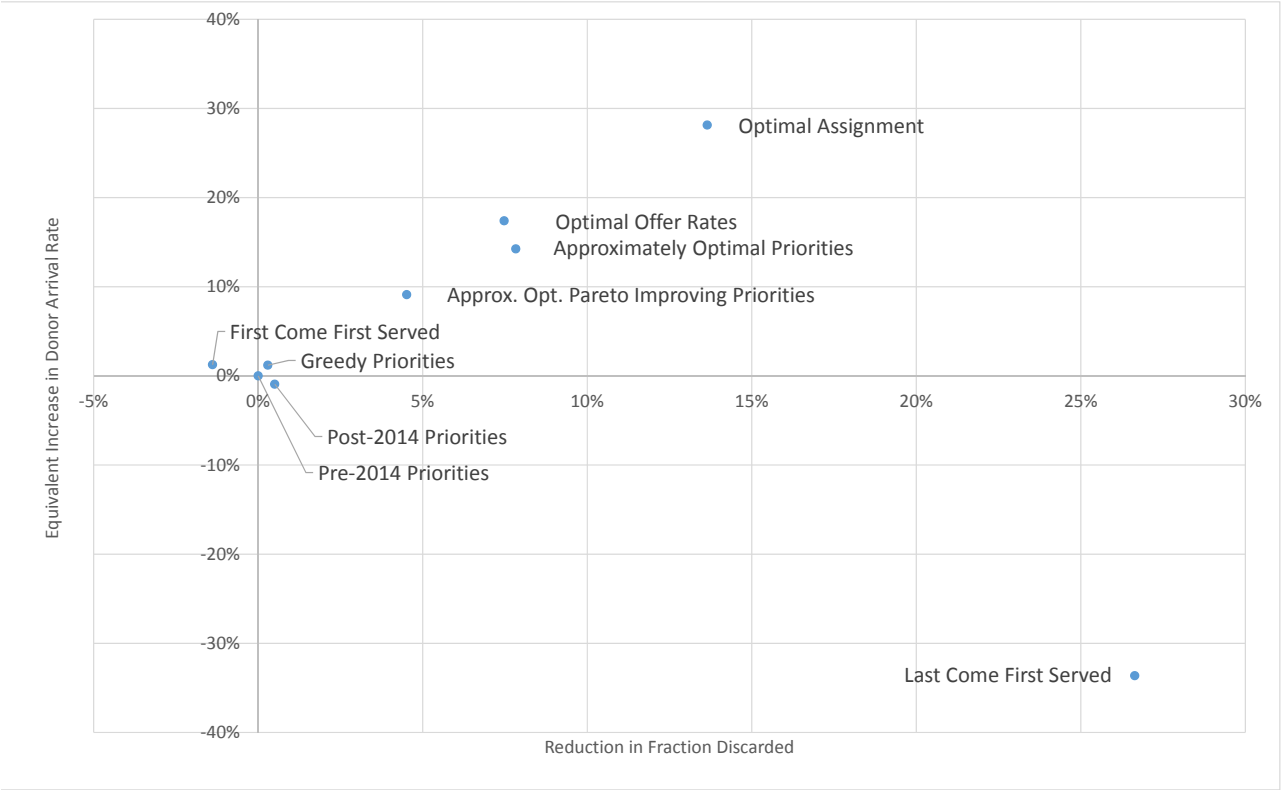

Notes: Results are reported relative to pre-2014 priorities and based on each mechanism's steady state equilibrium. The reduction in fraction discarded is defined as the change in the number of kidneys rejected by all NYRT patients to whom the organ was offered.