Empirical Essays on Dynamic Allocation Mechanisms

by

Daniel Cane Waldinger

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Economics
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2018

Author . . . Signature redacted . . . .
Department of Economics
May 15, 2018

Certified by . . . Signature redacted . . . .
Nikhil Agarwal
Castle Krob Career Development Assistant Professor of Economics
Thesis Supervisor

Certified by . . . . Signature redacted . . . .
Parag Pathak
Jane Berkowitz Carlton and Dennis William Carlton Professor of Economics
Thesis Supervisor

Accepted by . . Signature redacted . . . .
Ricardo Caballero
Ford International Professor of Economics
Chairman, Departmental Committee on Graduate Studies

Empirical Essays on Dynamic Allocation Mechanisms

by

Daniel Cane Waldinger

Submitted to the Department of Economics
on May 15, 2018, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This thesis contains three chapters which empirically study how dynamic decision making affects the allocation of public resources.

In the first chapter, I study the problem of allocating public housing. In the U.S., public housing authorities (PHAs) allocate apartments using a wide range of choice and priority rules. I evaluate how these allocation mechanisms affect the efficiency and redistribution achieved through assignments. Using waiting list data from Cambridge, MA, I estimate a structural model of public housing preferences, finding substantial heterogeneity in applicant outside options and preferred apartment types. Counterfactual simulations suggest that the range of mechanisms used by PHAs involves a significant trade-off between efficiency and redistribution. However, some commonly used mechanisms are never optimal.

In the second chapter, joint with Nikhil Agarwal, Itai Ashlagi, Michael Rees, and Paulo Somaini, I study the allocation of deceased donor kidneys. In the U.S., patients on the kidney waiting list are offered organs in order of priority, and may decline an offer without penalty. This paper establishes an empirical framework for analyzing the design of these waiting lists. We model the decision to accept an organ as an optimal stopping problem and use waiting list data to estimate the value of accepting various kidneys. We then show how to solve for counterfactual equilibria under different priority rules, and search for mechanisms that improve the match quality of transplants and reduce organ waste.

In the third paper, joint with Sydnee Caldwell and Scott Nelson, I investigate how beliefs about risky future income influence households' financial decisions. We quantify one contributor to income uncertainty by surveying low-income tax filers' expectations of and uncertainty about their tax refunds, and link the survey with administrative tax and credit report data. Households face substantial refund uncertainty, and both refund expectations and surprises influence financial behavior. Households borrow in anticipation of their tax refunds, and this pattern is less pronounced for more uncertain households, consistent with precautionary behavior. Surprisingly, positive refund surprises induce higher debt levels by relaxing down-payment collateral constraints.

Thesis Supervisor: Nikhil Agarwal
Title: Castle Krob Career Development Assistant Professor of Economics

Thesis Supervisor: Parag Pathak
Title: Jane Berkowitz Carlton and Dennis William Carlton Professor of Economics

Acknowledgments

I am incredibly fortunate to have had wonderful family, friends, classmates, and advisors/mentors throughout my life and academic career.

I have been most fortunate in what I least control – my amazing parents, Jennifer and Robert, and my most excellent brother David. Their warmth, love, and support made me into a functional human and sustained me through stressful periods. My parents' love of learning and commitment to social good fueled my own intellectual curiosity and interest in public policy. David played the fraternal role of chief interlocutor, and we honed our rhetorical and logical skills through endless arguments. Some of my clearest early academic memories are when my parents coached me on class presentations in elementary and middle school. These skills paid serious dividends twenty years later during my job talks.

My closest friends do not study economics, but they have had an enormous influence on my personal and intellectual development: Seth Green, Aviva Rosman, Colin Gaw, Eve Addai, Madeleine Abromowitz, Matt Woolf, Alex Dietz, and many others.

Several mentors facilitated my path to economics gradate school. Rick Victor gave me my first economics research job and exemplified how economic analysis can directly inform public policy. Undergraduate mentors Grace Tsiang, Victor Lima, Susanne Schennach, and Hugo Sonnenschein encouraged me to pursue economic research and gave me confidence in my own abilities. Working for Susan Athey and Denis Nekipelov at Microsoft Research, I developed empirical skills as well as a taste for structural modeling and the market design approach to economic analysis.

My (mostly) wonderful experience in graduate school has had a lot to do with terrific classmates and colleagues. I have shared MIT's intellectual and emotional vicissitudes with office mates Otis Reid, Christina Patterson, Benjamin Marx, Jon Petkun, Ray Kluender, and Michael Stepner. I am especially grateful to Alonso Bucarey for being an awesome friend and colleague as we wrote our job market papers. Some of my colleagues were also roommates – Samuel Grondahl, Ariel Zucker, Scott Nelson, and Andrey Fradkin. Others were cohabitants of the computer lab second year (Chishio Furukawa, Alan Olivi, Roman Zarate). The physical writing of my thesis is in large part thanks to coffee shop paper-writing comrades Robert Manduca and Roseanna Sommers.

For the thesis itself, and for what will come after, I owe more to Nikhil Agarwal than any other human being. As my advisor and coauthor for nearly four years, he has had a tremendous impact on my research and on my general approach to economics, both through his incredibly committed advising and by his own example. I strive to match his focus and thoughtfulness in research, his passion for learning, his commitment to his colleagues and students, and his conviction in the importance of our work. I also owe an incredible debt to my advisors Parag Pathak, Michael Whinston, and Amy Finkelstein for their feedback, support, and encouragement. They have not only pushed my research far beyond what I could have imagined, but also given me invaluable mentorship. Their inspiring examples will shape me as a researcher, thinker, and teacher for the rest of my career.

Many others have provided critical feedback and mentorship. Jacob Leshno not only influenced me through his work on dynamic mechanism design, but through his generously shared time and advice at MSR and afterward. Ingrid Gould Ellen helped me connect my thesis to practical policy issues.

Countless others deserve mention yet are not named. To all of them, I am grateful.

Contents

List of Figures

List of Tables

Chapter 1

Targeting In-Kind Transfers Through Market Design: A Revealed Preference Analysis of Public Housing Allocation

1.1 Introduction

In the United States, 1.2 million low-income households live in public housing. Tenants receive a permanent, place-based entitlement to a rent subsidy that can exceed $10,000 per year. However, this assistance is rationed – in 2012, there were at least 1.6 million additional households on public housing waiting lists nationwide (Collinson et al., 2015). Public Housing Authorities (PHAs) in each city have wide discretion over how to allocate available apartments and differ in the choice afforded to applicants and the priority systems used. Despite the range of policies, there is little empirical or theoretical work on how to design efficient dynamic allocation mechanisms when redistribution is also an important goal.

The range of choice and priority systems used in public housing allocation may have large welfare and distributional impacts. In cities such as New York City and Philadelphia, applicants may choose their preferred housing development; in other cities such as Los Angeles and Miami, applicants do not have any choice over where they are assigned. Theoretical work has shown that allowing choice can provide good match quality for those who receive apartments (Bloch and Cantala, 2017a; Leshno, 2017; Thakral, 2016). However, removing choice may induce applicants with good outside options to reject mismatched offers and self-select out of the public housing program, improving targeting (Arnosti and Shi, 2017; Nichols and Zeckhauser, 1982). PHAs also differ in whether priority is given to more or less economically disadvantaged households. These priorities directly affect targeting through observed characteristics that predict disadvantage, but may also limit the ability of applicants to self-select based on unobserved differences. Ultimately, the effects of these policies on efficiency and redistribution are an empirical question; they depend on the characteristics of public housing applicants, and the degree of heterogeneity in outside options and preferred apartment types.

This paper provides empirical evidence on the roles of choice and priority in public housing allocation using application data from the Cambridge Housing Authority (CHA), which administers public housing in Cambridge, MA. Based on applicants' submitted development choices, I estimate a structural model of public housing demand that quantifies heterogeneity in applicants' preferred developments and in their overall values of living in Cambridge public housing. In counterfactual simulations, I use the structural model to evaluate the welfare and distributional impacts of mechanisms used by PHAs in other U.S. cities. I find that when applicants have choice over where they are assigned, tenants value their assignments (relative to their outside options) as much as they would value cash transfers of $6,500 per year. The CHA could house more disadvantaged applicants by either removing choice or simply prioritizing lower-income applicants, as is done in other cities. Both policies result in lower tenant welfare per dollar spent on the public housing program, but prioritizing lower-income applicants improves targeting without lowering tenant welfare. As a result, some combinations of choice and priority are strictly dominated in Cambridge in a broad class of social welfare functions.

While studies of other centralized matching markets have used choice data to analyze the behavior and preferences of agents, this type of data is novel in the public housing context. The application data from Cambridge provide a direct measure of which households applied for Cambridge public housing and contain rich development choice information. During the period of study, the CHA allowed applicants to choose their preferred development in a two-stage process, which I refer to as the *Cambridge Mechanism.* In the first stage, an applicant made an initial choice of up to three developments. The initial choice formed the applicant's choice set in the second stage, when the applicant made a final choice after learning their position on the waiting list for each development in their choice set. This position information allowed applicants to update their beliefs about waiting time before making their final choices, and I provide descriptive evidence that applicants' final choices were responsive to this information.

The Cambridge Mechanism does not induce applicants to directly reveal their preferred housing developments. Instead, applicants face a trade-off between being housed in their preferred development and being housed more quickly. I propose a model of development choice that captures this trade-off. Each applicant compares the indirect flow utility from living in each public housing development to their outside option and chooses their preferred distribution of assignments and waiting times at each stage of the application process, understanding that their initial choice may affect the conditions under which the final choice is made. The resulting two-stage decision problem is a generalized version of the portfolio choice problem considered in Chade and Smith (2006). An eligible household applies if some public housing development is preferred to its outside option.

I interpret the distribution of indirect flow utilities using a model that allows applicants to have heterogeneous tastes for public housing developments and unobservably different outside options. Households receive utility from consuming housing and a numeraire, and maximize utility subject to a budget constraint. If utility is additively separable in housing and the numeraire, the difference in flow payoffs between living in each public housing development and the outside option is naturally decomposed into two parts. The first is the household's *value of assistance*, a common component across developments which captures the household's value of the homogeneous aspects of public housing. The second is the household's *match value* for the specific development, which captures the heterogeneous aspects of public housing and determines an applicant's preferred developments. In estimation, I make an assumption on the functional form of utility and impose a restriction on differences in the value of assistance. Specifically, I assume that unobserved differences in the value of assistance are driven by heterogeneous outside options rather than heterogeneous tastes for public housing itself. These assumptions lead to a natural parameterization of the value distribution and allow welfare gains from assignments to be compared to the value of cash transfers.

The two types of preference heterogeneity – values of assistance and match values – are closely related to the market design trade-off between providing good match quality for tenants and targeting the most disadvantaged applicants. Values of assistance determine which applicants a PHA would like to house, while match values determine how a PHA should match a fixed set of applicants to available apartments. They also determine how applicants will behave under different allocation mechanisms. Holding match values fixed, applicants with higher values of assistance will accept apartment offers from more developments and select developments with shorter waiting times. Holding the value of assistance fixed, applicants with high match values for specific developments will be willing to wait longer for those developments. A mechanism which induces applicants to reject mismatched offers may house more applicants with high values of assistance, with the potential cost that tenants enjoy lower match values from their assignments. The effect of allocation policy on targeting, match quality, and total welfare depends on the distribution of heterogeneity in each dimension.

The application data and structure of the Cambridge Mechanism provide crucial information about both types of preference heterogeneity. Application rates by income and demographic groups are particularly informative about observable differences in values of assistance. Combining American Community Survey (ACS) data with the Cambridge waiting list data, I estimate that lower-income and non-white households are much more likely to apply for public housing than other eligible households. However, some very low-income households did not apply, while some of the highest-income eligible households did, suggesting that there are also unobserved differences in values of assistance or match values. The initial development choices of applicants are informative about these unobserved differences. Since applicants choose up to three lists, initial choices reveal not only which developments are more likely to be chosen overall, but also which developments tend to be chosen together. These patterns reveal match value heterogeneity that can be predicted by observed applicant and development characteristics, as well as unobserved heterogeneity in tastes. With parametric restrictions, initial choices also separate match value heterogeneity from unobserved heterogeneity in values of assistance. The final choice stage informs sensitivity of development choices to waiting times since applicants receive new information before making their choices. This allows me to estimate a discount factor in addition to the parameters governing applicants' flow payoffs.

I estimate the development choice model by matching observed choice patterns to those predicted by the model using the method of simulated moments (McFadden, 1989; Pakes and Pollard, 1989). Implementing the procedure requires two preliminary steps. First, to measure application rates by income and demographic groups, I estimate the distribution of potential applicants – including eligible households who did not apply – by combining ACS data with administrative data on current public housing tenants in Cambridge. Second, I estimate applicants' beliefs about how each sequence of development choices affects the distribution of assignments and waiting times in the Cambridge Mechanism. Estimating beliefs presents a challenge because the Cambridge Mechanism created interdependence in the waiting time distributions across lists. As a result, the beliefs of sophisticated applicants are high-dimensional while data on realized waiting times are sparse. I overcome this problem by assuming that applicants' beliefs match the long-run steady state distributions that the Cambridge Mechanism would generate given observed frequencies of applicant arrivals and departures, apartment vacancies, and initial and final choices of applicants. This assumption allows me to exploit knowledge of the Cambridge Mechanism and construct the high-dimensional belief objects by simulation, using the data to estimate a lower-dimensional set of parameters governing simulation inputs.

Given these inputs, simulating the development choice model presents a computational challenge because the two-stage development choice problem is computationally burdensome to solve and does not yield closed-form choice probabilities.

Standard simulation techniques would re-solve the model at each proposed value of the parameter vector. This is computationally prohibitive in my application. I use a technique proposed by Ackerberg (2009) that combines a change of variables with importance sampling and allows me to solve the development choice model once. The optimization procedure re-weights simulation draws at new parameter values and minimizes the objective function over a grid of discount factors.

Estimates imply that applicants are moderately patient and exhibit substantial heterogeneity in values of assistance and match values. The point estimate of the annual discount factor is between 0.90 and 0.92, suggesting that development choices will be sensitive to equilibrium waiting times in mechanisms that allow choice. While observed characteristics strongly predict the value of assistance – particularly income and race – households also have unobserved differences in their outside options. Conditional on observed characteristics, the standard deviation of a household's outside option amounts to several thousand dollars of annual unobserved income. Applicants have strong preferences for specific developments, and would require a median cash transfer of $1,435 per year to provide the same welfare increase as moving from their second choice development to their first choice. Given such large heterogeneity in match values and values of assistance, 31 percent of applicants would accept any development, while an equal share would only be willing to live in three or fewer developments. Applicants that would accept any development have much lower observed incomes than other applicants as well as unobservably worse outside options. As a result, a development choice system that induces offer rejections will filter out applicants with better outside options but have large welfare costs in terms of match quality.

Given these estimates, I consider how the development choice and priority systems used by other PHAs would perform in Cambridge. Because computing the equilibrium of the two-stage Cambridge Mechanism is challenging, the counterfactuals focus on a simpler class mechanisms in which applicants make choices in one stage. The Cambridge Mechanism is closest to a one-stage mechanism in which applicants apply for one development and all eligible households living or working in Cambridge have equal priority. I consider what would happen if the CHA moved to other development choice systems, including ones that induce offer rejections. I also consider priority systems that offer apartments to either lower- or higher-income applicants before others. To show what could be achieved if incentive compatibility constraints were relaxed, I also analyze a full-information benchmark in which the social planner knows applicants' preferences but has limited foresight about future apartment vacancies and applicant arrivals and departures.

Under the current priority system in Cambridge, the range of development choice systems used in practice would have large effects on match quality, targeting, and total welfare. Removing choice would reduce the average value of an assigned unit, measured in equivalent cash transfers, from \$6,956 to \$5,399 per year. Match quality would fall dramatically; the fraction of tenants living in their first choice developments would fall from 43 percent to 11 percent. Since lower-income applicants are more likely to accept a mismatched apartment offer, the incomes of tenants would fall from \$20,509 to \$17,535, and tenants would have worse outside options conditional on their observed characteristics. Since lower-income tenants pay lower rents in public housing, cost-adjusted welfare gains fall even more than welfare per assigned unit. Based on a conservative estimate of the cost of maintaining each Family Public Housing apartment, cost-adjusted welfare gains would by fall 30 percent if the CHA gave applicants no choice over their assignment instead of allowing them to choose their preferred development. In contrast, the effects of prioritizing higher- or lower-income applicants are mainly distributional: equivalent variation per apartment allocated and match quality are similar across priority systems, but income-based priorities would dramatically change tenant incomes. As a result, cost-adjusted welfare gains are larger when higher-income applicants are prioritized.

The measure used to summarize welfare gains from assignments – equivalent cash transfers – implicitly places equal value on cash transfers to households of different incomes. To conclude the paper, I ask which allocation mechanism should be used depending on one's taste for income redistribution. I argue that social welfare weights should be monotone in the value of a household's outside option. Following Atkinson (1970), I consider a class of social welfare functions with "constant relative inequality-aversion" in which the strength of one's taste for redistribution is summarized by a single parameter. The value to each tenant of their assignment, measured in equivalent cash transfers, is transformed by a function that depends on the value of the household's outside option and the planner's degree of inequality aversion. This class of functions captures a wide range of distributional preferences and has attractive properties for making interpersonal welfare comparisons. In addition, welfare gains from each counterfactual allocation can be adjusted for changes in total rent payments, allowing mechanisms to be compared in terms of welfare gains per dollar of public expenditure.

Within this class of social welfare functions, several combinations of choice and priority are on the frontier of efficiency and redistribution among the mechanisms considered. With a low taste for redistribution, it is best to prioritize high-income applicants, since they are cheapest to house, and ask applicants to choose their preferred development. With a moderate taste for redistribution, one should prioritize applicants equally and still allow choice. With a high taste, one should prioritize low-income applicants, and eventually also limit applicants' ability to choose their preferred development. In the latter case, some applicants will self-select out by rejecting mismatched offers, improving targeting on unobserved as well as observed characteristics. The one-stage mechanism closest to the Cambridge Mechanism, choosing one development with equal priority, performs well under a moderate taste for income redistribution. A social planner would choose this mechanism if it equally valued transferring two dollars to a household earning \$20,000, and transferring one dollar to a household earning \$10,000.

Although the preferred mechanism depends on distributional preferences, certain combinations of choice and priority used in other cities are strictly dominated for the Cambridge population. In particular, it is never optimal to prioritize higher-income applicants while not allowing choice. Intuitively, prioritizing lower-income applicants yields a targeting improvement comparable to removing choice, but does so without lowering match quality. Inducing offer rejections is a policy of last resort to improve targeting once observed characteristics have been used. This implies that mechanisms used in other cities would not perform well in Cambridge. For example, Los Angeles prioritizes higher-income applicants but does not give applicants choice. In Cambridge, there would be a better policy whether one has a high or a low value of redistribution.

The paper proceeds as follows. Section 1.1 discusses related literature. Section 2 provides institutional background on the public housing program, discusses allocation policies used in practice, and describes the CHA dataset. Section 3 presents descriptive facts about Cambridge public housing developments, applicants, and their choices. Section 4 proposes a model of household preferences and development choice. Section 5 describes the estimation procedure used to recover the distribution of preferences for public housing developments. Section 6 presents the estimation results, and Section 7 presents results from counterfactual simulations. Section 8 concludes.

1.1.1 Related Literature

This paper is related to several literatures on means-tested housing assistance, dynamic market design, and the economics of in-kind transfers.

The empirical papers most closely related to this work estimate demand for public housing using data on assignments (Geyer and Sieg, 2013; Sieg and Yoon, 2016b; Van Ommeren and Van der Vlist, 2016). To my knowledge, this paper is the first to use individual-level waiting list data to estimate demand for public housing. Other empirical work has argued that there is substantial misallocation in the public and rent-controlled housing sectors (Glaeser and Luttmer, 2003; Thakral, 2016). Consistent with this work, I find that public housing allocation policy can dramatically affect how tenants are matched to apartments. A complementary literature evaluates the causal effects of receiving housing assistance, and has found that receiving housing assistance and living in higher socioeconomic status neighborhoods as a child leads to improved economic outcomes as adults (Andersson et al., 2016; Chetty et al., 2015; Kling et al., 2007; Ludwig et al., 2013). The subjective values for public housing estimated in this paper may include households' beliefs about the program's long-term benefits in addition to immediate changes in disposable income and housing and neighborhood quality.

The market design trade-off between match quality and targeting has been studied in the theoretical literature on one-sided dynamic assignment (Arnosti and Shi, 2017; Bloch and Cantala, 2017a; Leshno, 2017; Thakral, 2016). Arnosti and Shi (2017) show that the relationship between match quality and total welfare is theoretically ambiguous and depends on the distribution of agent preferences. This paper provides empirical evidence on these primitives and their implications for allocation policy. The trade-off between match quality and targeting is also connected to a literature on targeting and ordeals in public assistance programs (Akerlof, 1978; Nichols and Zeckhauser, 1982). This literature has highlighted the tension between providing valuable assistance to those who receive it ("productive efficiency") and restricting assistance to the households which need it most ("targeting efficiency"). Several recent papers have studied this idea empirically in the context of means-tested transfer programs of homogeneous items (Alatas et al., 2016; Deshpande and Li, 2017; Lieber and Lockwood, 2017). This paper explores a related trade-off created by the heterogeneous nature of public housing and its limited supply.[1] I also analyze how applicant priorities, a version of the tags considered in Akerlof (1978), interact with the screening properties of development choice in public housing allocation.

The structural model and estimation procedure used in this paper draw on techniques in discrete choice demand estimation (Berry et al., 2004; McFadden, 1973, 1989; Pakes and Pollard, 1989). My implementation of the method of simulated moments uses a change of variables and importance sampling technique proposed by Ackerberg (2009) to reduce the computational burden in estimation. This paper also joins a growing literature on revealed preference analysis in centralized matching markets (Abdulkadiroglu et al., 2017a; Agarwal, 2015; Fack et al., 2015a; Hastings et al., 2009; He, 2017; Narita, 2016). Along with Agarwal et al. (2017), this paper is among the first to conduct revealed preference analysis using the choices of agents in a dynamic assignment mechanism.

[1]The fact that public housing involves an in-kind transfer of housing rather than cash may also sacrifice productive efficiency by distorting the housing consumption of those who receive assistance. Given that only one quarter of eligible households applied for Cambridge public housing during the period of study, the targeting gains from public housing may be large compared to a cash transfer of equal value.

1.2 Institutional Background and Data

Section 1.2.1 provides an overview of the U.S. public housing program, surveys allocation policies used in practice, and discusses the design trade-offs these policies entail. Section 1.2.2 describes the Cambridge Housing Authority and the mechanism it used to allocate public housing during the period of study. Section 1.2.3 describes the applicant dataset and sample criteria.

1.2.1 Public Housing in the U.S.

The U.S. public housing program subsidizes the rents of 1.2 million low-income households at an annual cost of $8-10 billion. A Public Housing Authority (PHA) in each city maintains the stock of public housing developments located in its jurisdiction using funds allocated by Congress and distributed by the U.S. Department of Housing and Urban Development (HUD). A public housing tenant pays 30 percent of pre-tax income toward rent, and is permanently entitled to assistance as long as it complies with the terms of its lease and remains in its assigned apartment. Public housing and its private market counterpart, the Housing Choice Voucher program, are unusual in their benefit generosity: in 2013, participants received an average annual subsidy of $8,000.[2]

Due to the combination of limited federal funding, generous per-household benefits, and broad eligibility criteria, demand for public housing greatly exceeds supply. Congress does not set funding levels to assist all eligible households, but rather to maintain existing services. New public housing is not being built.[3] The income limit

[2]Based on per-household subsidy from tenant-based vouchers reported in HUD Congressional Justification for FY2015, available at https://www.hud.gov/sites/documents/FY15CJ_PUB_HSNG_CAPTL_FND.PDF. In 2013, the public housing program served a population with similar incomes.

[3]The majority of new affordable housing is built through the Low-Income Housing Tax Credit

for eligibility is 80 percent of Area Median Income (AMI), which includes lower-middle income households as well as the very poorest. As a result, in 2012 there were approximately 1.6 million households on public housing waiting lists nationwide, and nearly 3 million applicants on voucher waiting lists.[4]

Public Housing Allocation Mechanisms and Design Trade-Offs

The limited supply of public housing creates a dynamic assignment problem for each PHA. When tenants move out, the PHA must assign vacant apartments to applicants on a waiting list. PHAs have substantial autonomy over allocation policy. In particular, they control how applicants are ordered on the waiting list and whether applicants can choose the developments to which they are assigned. These policy levers – the *priority system* and *development choice system* – can affect which types of applicants receive assistance and whether they are matched to their preferred developments. To my knowledge, there is no resource that systematically documents the current waiting list policies of each of the 3,300 U.S. PHAs. To summarize allocation policies used in practice, I examined most recent available administrative plans of 24 PHAs falling into two categories: (1) those with the largest public housing stocks, and (2) those with public housing stocks and city populations similar to Cambridge, MA. The priority and development choice systems used by these PHAs are summarized in Table 1.1.

The allocation policies of surveyed PHAs share several common features. Ap-

(LIHTC), a federal tax expenditure that subsidizes private sector construction of new affordable housing. This program is administratively separate from the public housing and voucher programs and has a different rent payment structure, so that tenants with very low incomes receive a smaller effective rent subsidy than in public housing.

[4]Public and Affordable Housing Research Corporation (PAHRC), 2015. "Value of Home: 2015 PHARC Report." Based on PAHRC tabulation of the Public Housing Agency Homelessness Preferences Survey, 2012. `https://www.housingcenter.com/sites/default/files/waiting-list-spotlight.pdf`

plicants are ordered on a waiting list by priority and then by date of application. If applicants are allowed to choose a subset of developments to which they can be assigned, they are placed on waiting lists for their chosen developments. PHAs offer apartments to applicants living or working in the jurisdiction before other applicants. There are also federally mandated need-based priorities for certain groups, including households displaced by natural disasters, victims of domestic violence, and veterans. Apartments are offered to applicants at the top of the waiting list first; if an applicant rejects without good cause, they are removed from the waiting list and the next applicant is offered the apartment. A few PHAs allow one or two rejections before the applicant is removed from the waiting list, but most do not.

Despite these similarities, the development choice and priority systems used by PHAs exhibit important differences. The key difference across priority systems is whether households with higher or lower socioeconomic status are given priority. Some PHAs, including New York City and Los Angeles, give priority to households with a working member, that are economically self-sufficient, or that have incomes above 30 percent of the Area Median Income (AMI), a regional income benchmark that adjusts for household size. Others do just the opposite – the Seattle Housing Authority prioritizes households below 30 percent AMI, and several other PHAs prioritize households that are severely rent burdened or at risk of being displaced. Still other PHAs, including the Cambridge Housing Authority, treat all applicants living or working in the jurisdiction equally. Income-based priorities can have a large impact on the income distribution among public housing tenants. This will determine whether housed applicants have the highest values of living in public housing and, since lower-income households pay less rent, the fiscal cost of the public housing program. They also make it harder for applicants to obtain assistance who are not prioritized but have unusually high values of living in public housing.

The range of development choice systems across PHAs is equally wide. A development choice system gives each applicant a choice set consisting of certain *subsets* of developments from which the applicant can receive offers. Several PHAs, including those in New York City, Seattle, and New Haven as well as Cambridge, require applicants to choose a limited number of developments ("Limited Choice"). As noted in the dynamic market design literature, asking applicants to commit to their preferred options tends to achieve good match quality. Applicants will choose their preferred combinations of assignments and waiting times, and applicants with the highest values of over-subscribed developments will be more likely to apply for and occupy them. Other PHAs do not allow applicants to choose developments ("No Choice"); in Miami, Los Angeles, and Minneapolis, applicants must accept the first offer from any development. Such a mechanism will generate mismatch between tenants and their assigned apartments, but mismatched offers may filter out applicants with good outside options, allowing applicants to self-select into public housing based on both observed and unobserved characteristics. Other PHAs use intermediate development choice systems. Chicago allows applicants to select a neighborhood but not a specific development, which reduces spatial mismatch but may still induce offer rejections. In Boston, applicants may choose any subset of developments ("Any Subset"), allowing them to hedge against waiting time uncertainty. Philadelphia and Baltimore present applicants with a hybrid option ("Limited or All"): either commit to a few developments, or accept the first available apartment offer.

PHAs combine development choice and priority systems in different ways. Los Angeles uses No Choice, but prioritizes applicants that are economically self-sufficient (High SES). Seattle does the reverse, allowing Limited Choice while prioritizing Low SES applicants. Minneapolis uses both development choice (No Choice) and priorities (Low SES) to maximize targeting, while New Haven prioritizes higher-income applicants and provides choice. In counterfactuals, I ask what would happen if the Cambridge Housing Authority adopted different combinations of development choice and priority systems used in practice.

1.2.2 The Cambridge Housing Authority

The Cambridge Housing Authority (henceforth, CHA) administers the Public Housing and Housing Choice Voucher programs in Cambridge, MA. Its public housing stock consists of about 2,450 apartments, evenly split between the Elderly/Disabled and Family Public Housing programs. Although Cambridge is a low-poverty area compared to a nationally representative sample of public housing sites, Cambridge public housing tenants are comparable to those nationwide in terms of socioeconomic status and demographics. In 2014, 74 percent of Cambridge public housing tenants earned less than 30 percent AMI and 48 percent were headed by an African American, compared to 72 percent and 48 percent nationwide.

During the period of study – January 1st, 2010 to December 31st, 2014 – the CHA employed a site-based waiting list system to allocate public housing. The waiting list for vouchers was closed from 2008 until 2016, while public housing waiting lists were open from 2008 until 2015. For this reason, I study the public housing program in isolation. The CHA used a two-stage development choice system for public housing, which I will refer to as the *Cambridge Mechanism*.[5]

[5] Every year, each housing authority is required to publish an Admissions and Continued Occupancy Policy (ACOP). The CHA's most recent ACOP for federal public housing can be found here: http://cambridge-housing.org/civicax/filebank/blobdload.aspx?BlobID=23535

The Cambridge Mechanism

In the Cambridge Mechanism, applicants select their preferred development – they have Limited Choice – and all applicants with a household member living or working in Cambridge receive Equal Priority. The development choice system shares features with those used in New York City, Seattle, and New Haven; the priority system is similar to those used in Chicago, Philadelphia, and Boston.

A key difference between the Cambridge Mechanism and many other development choice systems is that applicants choose their preferred development in two stages.[6] At initial application, a household is assigned a program (Elderly/Disabled or Family) and bedroom size and makes an initial choice of up to three developments from 9 to 13 alternatives. Each development is a building or complex in a distinct geographic location, and apartments with the same number of bedrooms are mostly homogeneous within a development. The initial choice forms the applicant's choice set later on, and the applicant is placed on a waiting list for each chosen development. At a later date, the CHA sends the applicant a letter asking them to make a final development choice. The letter informs the applicant of its current position on each list in its choice set, allowing the applicant to make its final choice based on new information. Appendix 1.10.2 provides a formal description of the Cambridge Mechanism, including when the CHA sends these letters and how it calculates list position. After making its final choice, the applicant remains on the waiting list for that development until the CHA makes a single, take-it-or-leave-it offer of an apartment. If the applicant rejects, it is removed from the waiting list and cannot reapply for one year. The applicant may also be removed if it fails to attend its screening

[6]The New York City Housing Authority uses a similar two-stage development choice system. Applicants first choose a preferred borough, and later choose their preferred development from a subset of the developments in that borough.

appointment, produce required documentation, or respond to mail from the CHA.

1.2.3 Dataset and Sample Selection

The main dataset used in this paper, provided by the CHA, contains anonymized records of all applicants for Cambridge public housing who were active on a waiting list between October 1st, 2009 and February 26th, 2016. The CHA maintains a database of applicants to manage its waiting lists and comply with HUD regulations. For each applicant, the dataset records household characteristics, development choices, and the timing and outcome of all events during the application process.

For analysis, I restrict my sample to applicants who had priority for Cambridge public housing; who applied for 2 and 3 bedroom apartments in the Family Public Housing program; and who submitted an application between 2010 and 2014. Non-priority applicants had virtually no chance of being housed, and are therefore excluded. Family Public Housing applicants are a more homogeneous group than Elderly/Disabled applicants. I restrict to 2 and 3 bedroom apartments for sample size; there are few apartments and applicants for other bedroom sizes. Analyzing new applications between 2010 and 2014 avoids selection issues because not all pre-2010 applicants were still on the waiting list in 2010. These restrictions produce a sample of 1,752 applicants. After omitting 26 irregular applications, 1,726 applicants remain.

To estimate the distribution of potential applicants during the sample period, I augment the CHA applicant dataset with a sample of eligible households from the American Community Survey (ACS). I also use data provided by the CHA on Cambridge public housing tenants between 2012 and 2014. Appendix 1.10.1 provides details of the CHA and ACS datasets, and Section 1.5.1 explains how they are used to estimate the distribution of potential applicants.

1.3 Descriptive Evidence

This section presents descriptive statistics of Cambridge public housing applicants and their development choices. These facts illustrate the key economic forces that will be quantified in the structural model. Cambridge public housing developments differ in size, location, and expected waiting time. The decision to apply and applicants' initial development choices reveal heterogeneity in values of assistance and match values. While observed characteristics strongly predict who applies and which developments they prefer, much choice behavior is left unexplained. Final choices reveal that applicants are sensitive to waiting time information, and will choose a less preferred development in exchange for a shorter expected waiting time.

1.3.1 Cambridge Public Housing Developments

During the period of study, applicants for Family Public Housing in Cambridge chose among thirteen developments located throughout the city. The location of each development is shown in Figure 1-1. There are 3 developments in East Cambridge, 3 in North Cambridge, and 7 near Central Square. Table 1.2 displays characteristics of these developments. The smallest developments contain just a few apartments that blend in with the surrounding housing stock,[7] while the largest developments are complexes of several buildings containing hundreds of apartments. Developments also have different expected waiting times. Average waiting times for housed applicants range from 1.58 to 3.75 years across developments, with smaller developments

[7]The "Scattered" waiting list represents three lists: one for scattered sites in Mid-Cambridge (Central), one for East-Cambridge, and one for River Howard Homes (Central).

tending to have longer waits. As a result, some applicants faced a trade-off between their preferred assignment and a shorter expected wait. Developments are less heterogeneous in terms the characteristics of their tenants, with similar average incomes and proportions of African American tenants.[8]

1.3.2 Application Decisions and Initial Development Choices

Application rates by income and demographic groups reveal which types of households value public housing the most relative to their outside options. The first two columns of Table 1.3 show that only one in four eligible households actually applied for Cambridge public housing during the sample period. Those who did apply had much lower incomes and were more likely to be non-white and to already live in Cambridge. The average income of eligible households is \$42,219, while that of applicants is \$18,477. This is to be expected; since rent is 30 percent of pre-tax income, a lower-income household sees larger increases in housing quality and disposable income in public housing compared to its outside option. Differences by race are also striking: half of applicant households are headed by an African American, while only one in five eligible households are. Although income and race strongly predict who applies, they are not perfectly predictive. Figure 1-2 shows that while application rates fall steadily as income rises, some of the lowest-income households did not apply and some high-income households did. Similarly, 25 percent of African American headed households did not apply.

The remaining columns of Table 1.3 show that most applicant characteristics are stable over time, but there are a couple of moderate trends. The rate of new

[8]There are outliers. For example, Roosevelt Mid-Rise has an unusually low average tenant income and a small fraction of African American tenants. This is because it is a mixed development, with some apartments for Elderly and Disabled households. Its tenants are older, and as a result have lower incomes and are more likely to be white.

applications fell from 415 per year in 2011 to 347 in 2014.[9] Over time, new applicants had higher incomes and were more likely to work in Cambridge and have a white head of household. Applicant income growth is consistent with median income growth in the Boston area following the Great Recession. Despite the fact that only one in four eligible households applied for public housing during the sample period, there were five applicants for each of the 327 apartment vacancies. Demand greatly exceeded supply in this market.[10]

Initial development choices suggest that applicants have strong tastes for specific developments and that their preferences are correlated with observed characteristics. Table 1.4 presents statistics from initial development choices for all applicants and broken out by household income and neighborhood of current residence. Applicants that already live in Cambridge are much more likely to select developments in their own neighborhoods. The majority of applicants (84 percent) exhaust their initial choice set and select three housing developments. This rate is lower for applicants with incomes over \$32,000: only 78 percent select three lists, compared to 85 percent for lower-income applicants. Higher-income applicants also select developments with slightly longer average waiting times. These patterns are consistent with a model in which applicants with better outside options are more selective in their development choices. However, the fact that these differences are not larger suggests the presence of unobserved heterogeneity in values of assistance.[11] Similarly, specific chosen

[9]The CHA closed its Family Public Housing waiting lists during the second and third quarters of 2010. As a result, 2010 saw fewer new applications than subsequent years.

[10]The number of vacancies is below the long-run average because the CHA began renovating its public housing stock during the sample period. For a plausible upper bound on the long-run average, an annual turnover rate of 10 percent per unit would raise the expected number of vacancies to 540 over a five year period.

[11]Note that higher-income households who applied for Cambridge public housing are already a selected sample. This should mute any correlation between applicant characteristics and the selectivity of their development choices.

developments are not fully predicted by observed characteristics. The structural model will quantify heterogeneity in both values of assistance and match values, as a function of both observed and unobserved characteristics.

1.3.3 Response to Waiting Time Information

This section presents quasi-experimental evidence that applicant choices are sensitive to information about waiting time. Between 2010 and 2014, Cambridge sent final choice letters to applicants who were near the top of the list for one of their initial choice developments. The letter informed applicants of their position on each list and asked them to make a final development choice. Because of fluctuations in relative list lengths over time, and also due to Cambridge's algorithm for calculating list position and sending final choice letters, applicants who made the same initial development choices but applied on different dates were given different position information when they made their final choices. Final choices are sensitive to this information: when an applicant is told a lower list position for one development relative to the others in their choice set, they are more likely to pick that development.

To test the null hypothesis of no response to waiting time information, I run a conditional logistic regression that predicts an applicant's final choice as a function of list position or expected continued waiting time. The sample is applicants who made a final choice during the period of study, and the outcome is which development they chose. Since each applicant chose their choice set at initial application, I include as controls fixed effects for the interaction between each development and choice set. This isolates the natural experiment in which applicants who made the same initial choices – and whose development preferences are therefore drawn from the same distribution – are told different waiting times for the same alternatives.

Table 1.5 displays coefficient estimates and implied marginal effects from the conditional logistic regressions of final choice on waiting time information with no controls; with development fixed effects; and with the full set of development and choice set interactions. For each set of controls, the specification is run for both list position and expected continued waiting time. Except for Column (2), coefficient estimates are precise and show a negative response to list position and continued waiting time. The response grows stronger with additional controls. The implied elasticities are large: with full controls, the elasticity of final choice is -1.1 with respect to list position and -4.1 with respect to continued waiting time.

For a test of the null hypothesis of no response to be valid, position information must be uncorrelated with development preferences among applicants with the same choice set who made a final choice. Two conditions are sufficient for this assumption to hold. The first is that the development preferences of applicants who applied on different dates but made the same initial choice are drawn from the same distribution. This would not be true if applicants anticipated fluctuations in waiting times, since this would influence initial choices. However, given that waiting time fluctuations are determined by randomness in when apartments become vacant and the decisions of other applicants, these fluctuations would have been difficult to predict or influence. The second condition is that response to the final choice letter is uncorrelated with the specific information in the letter, conditional on the elapsed time since application. This will be true if applicants become unresponsive for exogenous reasons.

These results simply establish the existence of a response. In structural estimation, moments based on responsiveness to waiting time information will separate applicants' rate of time preference from heterogeneity in their values of specific developments.

1.4 Model of Preferences and Development Choice

Section 1.4.1 presents a development choice model which predicts how eligible households behave at each stage of the application process given the structure of the Cambridge Mechanism. This model allows me to recover the distribution of preferences for Cambridge public housing developments based on the application decisions and development choices of eligible households. Section 1.4.2 provides a micro-foundation of preferences that links development preferences to households' outside options. In counterfactuals, I use this model to quantify the welfare and distributional impacts of alternative waiting list policies.

1.4.1 Development Choice Model

The development choice model provides a rational benchmark through which to interpret the application decisions of eligible households and development choices of applicants. In particular, it captures the trade-off applicants may face between spending less time on the waiting list and being assigned to their preferred housing development.

Knowing the structure of the Cambridge Mechanism, applicants solve a single-agent problem and choose their preferred distribution of assignments and waiting times given their information at each stage of the application process. They have limited information about the state of the waiting list when making their initial choices, but update their beliefs based on the position information in their final choice letters. Because applicants make development choices in two stages and receive new information in the second stage, the Cambridge Mechanism generates a portfolio choice problem. I assume that applicants are sophisticated and solve this choice problem backwards, anticipating that the full set of developments in their initial choice may jointly affect the timing of and position information received in the final choice stage.

The following sections specify the sequence of decisions; information and beliefs about how choices affect future states; payoffs; and the resulting portfolio choice problem.

Sequence and Timing of Decisions

An eligible household, indexed by i, makes decisions in the following sequence:

1. *Application Decision*: Household i receives the opportunity to apply on a random date.

2. *Initial Choice*: If i applies, it immediately chooses up to three developments, denoted $C \subset \{1, ..., J\}$ with $|C| \leq 3$. These developments form i's choice set in the final choice stage, and i is placed on a waiting list for each development in its initial choice.

3. *Final Choice*: At a later date, i receives a letter containing i's position on the waiting list for each development in its choice set. The letter asks i to make a final choice $f \in C$. Let s denote the number of years between initial application and the final choice letter, and let $p \equiv \{p_j\}_{j \in C}$ denote the vector of list positions. If i responds to the letter and chooses development f, it remains on the waiting list until it receives a take-it-or-leave-it apartment offer in f.

Household i may become unresponsive at any point during the application process and is removed from the waiting list if this occurs. I will assume that attrition is exogenous to the model; that an applicant cannot anticipate the date it will be removed; and that removal occurs at a poisson rate α that is equal across applicants.

Applicants may not fully anticipate the possibility of attrition, and have a subjective attrition probability $\tilde{\alpha} \leq \alpha$.

Applicant Information and Beliefs

An applicant's optimal initial and final choices will depend on its beliefs about how each possible choice affects the joint distribution of assignments and continued waiting times. Based on institutional features of the Cambridge Mechanism as well as descriptive evidence, I assume that applicants do not know the state of the queue when they first apply, but update their beliefs about continued waiting times based on the position information in their final choice letters.[12] When applicant i makes its initial choice, it does so with beliefs about the likely date s and position information p at the final choice stage, which are unknown and whose joint distribution depends on i's initial choice. Let $G_C(s, p)$ denote the probability that the final choice letter is sent less than s years after initial application and that the applicant's list position is no greater than p_j for each development $j \in C$. At the final choice stage, s and p are realized, and i updates its beliefs about the continued waiting time for each development $j \in C$. Let $F_{j,C}(t \mid p)$ denote the probability that continued waiting time for list $j \in C$ is less than t years given position vector p. Importantly, these distributions depend on the full set of lists C in an applicant's initial choice. Due to the algorithm by which the CHA sent out final choice letters, described in Appendix 1.10.2, the full set of lists in C could affect the date and information at the final choice stage. In addition, because applicants make their final choices based on new position information, the full set of list positions p may be informative about the

[12]Descriptive evidence from the CHA dataset suggests that applicants are unaware of short and medium-term fluctuations in list lengths. It is also consistent with the information they are given at initial application, and with conversations with the CHA. The CHA generally knew which developments had longer waiting times than others but was unaware of fluctuations.

expected continued waiting time for each development $j \in C$.

Preferences over Assignments and Waiting Times

Household i receives a payoff that is realized continuously over time and depends on where it lives. In particular, i's per-period indirect flow utility from living in development j is v_{ij}, and its indirect flow utility from not living in Cambridge public housing is v_{i0}. I will refer to these indirect flow utilities as flow payoffs. Assignments are believed to be permanent, and anticipated flow payoffs are not time-dependent. This rules out learning about characteristics of the developments over time or changing household circumstances. When making development choices, the household discounts future payoffs at exponential rate $\rho = r + \tilde{\alpha}$. This includes both the household's rate of time preference r, and their beliefs about the rate at which they will exogenously depart the waiting list, $\tilde{\alpha}$. There is no direct cost of remaining on the waiting list, and no fixed cost of beginning or continuing the application process. The present discounted value to i of being assigned to development j in t years is

$$e^{-\rho t} \frac{1}{\rho}(v_{ij} - v_{i0}) \ .$$

Choice Problem

Given beliefs and payoffs, an applicant solves the two-stage development choice problem backwards. In the final choice stage, applicant i with initial choice C learns its list positions p and solves

$$\max_{j \in C} \frac{1}{\rho} E\left[e^{-\rho T_j} \mid p\right](v_{ij} - v_{i0})$$

$$= \max_{j \in C} \int \frac{1}{\rho} e^{-\rho T_j} (v_{ij} - v_{i0}) dF_{j,C}(T_j \mid p) \ .$$

Anticipating the final choice stage, applicants make their initial choices to maximize the expected discounted value of the final choice:

$$\max_{C \in \{0,1,\ldots,J\}^3} E \left[e^{-\rho S} \max_{j \in C} \frac{1}{\rho} E \left[e^{-\rho T_j} \mid P \right] (v_{ij} - v_{i0}) \right]$$

$$= \max_{C \in \{0,1,\ldots,J\}^3} \int e^{-\rho S} \max_{j \in C} \left[\int \frac{1}{\rho} e^{-\rho T_j} (v_{ij} - v_{i0}) dF_{j,C}(T_j \mid P) \right] dG_C(S, P) \ .$$

Finally, since there is no direct cost of applying or remaining on the waiting list, an eligible household applies for public housing if and only if some development is preferred to their outside option: $\max_j v_{ij} > v_{i0}$. Applicants will also continue the application process if they have not already been removed for exogenous reasons. As a result, counterfactual mechanisms will affect development choices and waiting times, but not which households apply or when they would depart before being offered an apartment.

1.4.2 Utility Model

Because development choices depend on a household's value of living in each development relative to their outside option, my empirical strategy will estimate the distribution of $v_i = (v_{i1} - v_{i0}, ..., v_{iJ} - v_{i0})$. This section provides a micro-foundation of payoffs that explicitly links these payoff differences to the value of a household's outside option. The key assumptions are that utility is additively separable in housing and non-housing consumption, and that unobserved differences in the value of living in public housing are driven by outside options. In estimation, I add a restriction on the functional form of utility to parameterize the distribution of $v_{ij} - v_{i0}$ and to compare changes in utility to equivalent cash transfers.

Micro-Foundation of Flow Payoffs

Household i receives utility from consumption of housing h and a numeraire c. The utility function is additively separable in the two goods:

$$u(c,h) = u_1(c) + u_2(h)\,.$$

Both u_1 and u_2 are strictly increasing, concave functions. The household has three characteristics: observed income y_i; unobserved income η_i; and development-specific preferences summarized in hedonic indices $h_i = (h_{i1}, ..., h_{iJ})$. Outside of public housing, a household chooses how much to spend on each good given its budget $y_i + \eta_i$. The prices of both goods are normalized to one. The household's indirect flow utility from its outside option is

$$v_{i0} \equiv \max_{c,h} \quad u_1(c) + u_2(h) \qquad s.t. \qquad c + h \leq y_i + \eta_i \tag{1.1}$$

$$= v_0(y_i + \eta_i)\,. \tag{1.2}$$

One can think of unobserved income as capturing resources that relax or tighten the household's budget constraint, shifting the value of its outside option. An extensive literature has shown that social ties and alternative living arrangements are an important economic resource for many low-income households (Desmond and An, 2015; Stack, 1974). By modeling these resources as part of the budget constraint, I assume that they are substitutable between housing and the numeraire.

In public housing, household i only has access to observed income y_i. Because it is assigned to a particular apartment, it does not choose how much to spend on housing and the numeraire. Instead, pays a fixed fraction τ (30%) of income in rent, spends the remainder on the numeraire, and enjoys housing consumption h_{ij} in development j. The indirect flow utility from living in development j is

$$v_{ij} \equiv u_1((1-\tau)y_i) + u_2(h_{ij}) \,. \tag{1.3}$$

The difference in flow payoffs is given by

$$v_{ij} - v_{i0} = \underbrace{u_1((1-\tau)y_i) - \overbrace{v_0(y_i+\eta_i)}^{\text{outside option}}}_{\text{value of assistance}} + \underbrace{u_2(h_{ij})}_{\text{match value}} \,. \tag{1.4}$$

This expression decomposes the difference in flow payoffs into two components: the household's value of assistance and its match value. The value of assistance is common across developments and depends only on household i's observed and unobserved income. It can be thought of as the household's value of the homogeneous aspects of Cambridge public housing. The match value depends on i's taste for the characteristics of development j; it comes from the heterogeneous nature of public housing. These two terms capture the mechanism design trade-off between providing better match quality for housed applicants and housing applicants who want public housing the most. A mechanism that does not give applicants choice over their assignment may induce low-value applicants to reject mismatched offers. If this occurs, more high-value applicants will be housed, with the potential cost that tenants enjoy lower match values.

This utility model embeds two key assumptions. The first is that utility is additively separable in housing and the numeraire. This rules out complementarity between housing and non-housing consumption, and assumes that the match quality a tenant enjoys from their apartment does not affect the value of consuming other goods. The second assumption is that unobserved income is only available outside of public housing, and that it is substitutable between housing and the numeraire. This implies that differences in the value of assistance are driven by households' outside options rather than the value of public housing itself, and that the value of the outside option determines the value of cash transfers. These two assumptions will make it possible to separately identify the value of the outside option from the financial and other benefits of living in public housing.[13]

1.5 Empirical Strategy

This section describes the three steps in my estimation procedure. First, I estimate the distribution of potential applicants for Cambridge public housing, including eligible households who did not apply. Second, I estimate applicants' beliefs about how their choices affect payoffs through the distribution of assignments and waiting times. Third, given beliefs and the distribution of potential applicants, I estimate preferences over assignments and waiting times by matching application decisions and development choices using the method of simulated moments (McFadden, 1989; Pakes and Pollard, 1989). Solving the two-stage development choice problem is computationally expensive, and a change of variables and importance sampling technique proposed by Ackerberg (2009) reduces the computational burden. The final subsection shows how estimates from the utility model can be interpreted in terms of equivalent cash transfers.

[13] One would ideally obtain additional data on households' outside options to separate unobserved differences in outside options and taste for public housing, but such data were not available for this study.

1.5.1 Distribution of Potential Applicants

The first decision an eligible household makes is whether to apply for public housing at all. Application rates by income and demographic groups will be informative about heterogeneity in the value of assistance. To measure application rates, I need to estimate the distribution of characteristics of all households that could have applied for Cambridge public housing during the sample period. This includes households that did apply and also *eligible non-applicants* – eligible households that did not apply and were not already Cambridge public housing applicants or tenants at the beginning of 2010. This section outlines the statistical procedure used to estimate the distribution of potential applicants.

Estimating the distribution of potential applicants is not straightforward. The CHA dataset contains information on households who applied during the sample period, but it does not contain households that could have applied but did not. Survey data can identify households whose characteristics made them eligible for Cambridge public housing. However, some eligible households were already Cambridge public housing tenants, and others were on the waiting list but applied before 2010. These households were not potential applicants during the sample period, and survey data do not distinguish them from households that could have applied.[14]

My approach is to combine a sample of eligible households from the American Community Survey (ACS) with the CHA dataset to determine the distribution of characteristics among eligible non-applicants. I do this by assigning a probability to each household in the ACS for whether it appears in the CHA dataset, either as a tenant or as a past or current applicant. I specify these probabilities as a parametric

[14]The American Community Survey (used here) does ask whether a household receives housing assistance. However, a number of studies including Meyer and Mittag (2015) have shown that these questions tend to understate program participation. To my knowledge, no large survey asks households whether they are on a *waiting list* for public housing.

function of household characteristics, and estimate the parameters by matching the characteristics of households in the CHA dataset using minimum distance. One minus each probability is an estimate of the probability that the corresponding ACS household could have applied for Cambridge public housing during the sample period, but did not. Using these probabilities, I draw a sample of eligible non-applicants and combine it with the applicant sample. This procedure is consistent with a model in which households become eligible for public housing once, choose whether to apply, and exit the waiting list or tenancy when they are no longer eligible. Though this model abstracts from the possibility that households might re-apply for public housing, it captures the key idea that households with higher values of living in public housing should be more likely to apply.

The ACS publishes a 5 percent sample of U.S. households covering 2010 through 2014, the same period covered by the CHA applicant dataset.[15] It contains information on household structure and economic and demographic characteristics that determine eligibility and priority for Cambridge public housing. In particular, I observe whether each ACS household lives or has a member working in Cambridge; whether it meets the income and asset tests; and whether its household structure qualifies it for a two or three bedroom apartment in Family Public Housing.

I estimate parameters of the probability model by minimum distance. Households are indexed by $b = 1, ..., B$ and have observed household characteristics Z_b. The ACS assigns each surveyed household a weight w_b based on household b's inverse probability of being sampled – in other words, w_b is the expected number of households that b represents. The estimator chooses a parameter vector θ_{ACS}, which determines the probability that each household appears in the CHA dataset given their charac-

[15]Samples from the ACS can be downloaded here: https://usa.ipums.org/usa-action/variables/group

teristics through a probit link function. θ_{ACS} is chosen to match the total number of households in the CHA dataset; the number of households in six income groups; and the numbers of households from Cambridge and with African American or Hispanic household heads. Denote statistics from the Cambridge dataset by m_{data}, and denote the contribution of each ACS household to the same statistics by m_b. The minimum distance estimator solves

$$\min_{\theta_{ACS}} \ (m_{ACS}(\theta_{ACS}) - m_{data})'(m_{ACS}(\theta_{ACS}) - m_{data})$$

where

$$m_{ACS}(\theta_{ACS}) \equiv \sum_{b=1}^{B} p(Z_b, \theta_{ACS}) w_b m_b$$

$$p(Z, \theta) = \Phi(Z'\theta)$$

1.5.2 Belief Distributions over Assignments and Waiting Times

The information about preference heterogeneity contained in applicants' development choices depends on their beliefs about how choices affect payoffs. An applicant solving the two-stage development choice problem of Section 1.4.1 has beliefs about how each initial choice affects the date and position information at the final choice stage, and about continued waiting times for each development given list positions:

$$\{G_C(S,P) \ \ , \ \ \{F_{j,C}(T_j \mid p)\}_{j,p}\}_{C \in \mathcal{C}}$$

Because the final choice stage of the Cambridge Mechanism generates interdependence in waiting times across developments, each possible initial choice may induce a different set of distributions over final choice states and continued waiting times.

A major challenge is that data on realized waiting times are sparse, while the beliefs of sophisticated applicants are high-dimensional. To address this issue, I assume that applicants have beliefs of a particular form: their beliefs are consistent with the long run steady state distributions that the Cambridge Mechanism would generate given empirical vacancy rates, applicant arrival and departure rates, and initial and final choice frequencies. These empirical quantities can be estimated directly from application data. Combining these estimates with knowledge of the Cambridge Mechanism, I simulate steady state outcomes which quantify interdependence across lists and the option value of the timing and information of the final choice stage. I assume that applicants have these beliefs when simulating the model in the final step of estimation.

The rest of this section describes the model of the Cambridge Mechanism, the construction of simulation inputs, and the construction of belief distributions from simulation outputs.

Structure of Simulation Inputs

Appendix 1.10.2 provides a formal model of the Cambridge Mechanism. This section explains the structure placed on inputs that determine assignments. Each day, the following steps occur:

1. New applicants enter the queue and make their initial development choices.

2. Vacant apartments are offered to applicants who have already made their final choices.

3. If the number of applicants on a list who have made their final choices falls below a threshold, the CHA sends final choice letters to a group of applicants

on that list. Each letter tells the applicant their current list positions and asks them to make a final choice.

4. Applicants that do not respond to a final choice letter or to an apartment offer are removed from all waiting lists.

Given this structure, outcomes in the Cambridge Mechanism are determined by apartment vacancies, arrival and departure dates of applicants, initial and final choices of applicants, and the CHA's policy for sending final choice letters. Vacancies, applicant arrivals and departures, and initial choices do not depend on the state of the waiting list and are modeled as independent exogenous processes; however, the CHA's policy for sending final choice letters and the final choices of applicants do depend on the current state of the waiting list. I therefore place the following structure on inputs:

- Calendar time is indexed in days by $t \in \{1, ..., T\}$. Each list $j \in \{1, ..., J\}$ represents a development and bedroom size. There are S_j apartments represented by list j.

- **Apartment Vacancies:** each vacancy $\nu \in \{1, ..., V\}$ is associated with a calendar date t_ν and a waiting list j_ν. Vacancies occur independently on each list at poisson rates. Vacancy rates were unusually low during the period of study; according to the CHA, the long-run vacancy rate per apartment is once every 10 years, so the vacancy rate of list j is set to $0.1 * S_j$.

- **Applicant Arrivals and Exogenous Departures:** each applicant $i \in \{1, ..., N\}$ arrives on date t_i and becomes unresponsive after date r_i if it has not been housed. Applicants arrive according to a poisson process with arrival rate α.

Each applicant becomes unresponsive immediately with probability a_0, and departs at an exponential rate a_1 thereafter.

- **Initial Choices:** applicant i makes an initial choice $C_i \subset \{1, ..., J\}, |C_i| \leq 3$ upon arrival. Since applicants do not know the state of the waiting list when they apply, their initial choices are independent of the current state.

- **Final Choice Letters:** the CHA sends final choice letters according to a rule that depends on the state of each waiting list. For each list j, there is a sequence of trigger and batch size policies $\{(L_{j,l}, K_{j,l})\}_{l=1}^{L}$ for sending letters. Each day, if fewer than $L_{j,l}$ applicants on list j have made a final choice, this triggers a batch of final choice letters to the next $K_{j,l}$ applicants on list j who have not yet made a final choice. After batch l of final choice letters is sent on list j, pair $(L_{j,l+1}, K_{j,l+1})$ becomes the next trigger and batch policy.

- **Final Choices:** applicants who respond to the final choice letter make their final choice based on their list positions. I use a reduced form model to capture the sensitivity of the final choice to this information. Applicant i selects list $j \in C_i$ with probability

$$\frac{\exp(\beta p_{ij} + \xi_j)}{\sum_{m \in C_i} \exp(\beta p_{im} + \xi_m)}$$

where p_{im} is applicant i's position on list m and ξ_m is a fixed effect for list m.

Construction of Simulation Inputs

The parameters governing inputs are estimated as follows. The annual probability each apartment becomes vacant is calibrated to 10 percent per year.[16] The applicant arrival rate is simply the mean number of applicants per year during the period of study. Initial choice probabilities are also taken directly from the data. Departure parameters were estimated by non-linear least squares using response to the final choice letter as a function of time since application. The coefficients of the final choice model were estimated using the specification in Column (3) of Table 1.5, replacing continued waiting time with the list position number. Each list has its own distribution of trigger and batch policies, the empirical distribution for the list during the sample period. Sequences of trigger and batch policies are drawn with replacement from their empirical distributions on each list during the period of study.

Given these parameters, I draw sequences of inputs and run the Cambridge Mechanism until it reaches a steady state. Sequences of apartment vacancies and applicant arrival and departure dates are drawn independently. Each applicant's departure date equals its arrival date with probability a_0 and follows an exponential distribution with mean $\frac{1}{a_1}$ years otherwise. The applicant's initial choice is drawn with replacement from the empirical distribution. Finally, I draw a random number for each applicant that determines which final choice it will make given the choice probabilities implied by its list positions.

[16]Due to renovations, the empirical vacancy rate during the sample period was below the long-run average. This approach also assumes an equal vacancy rate per apartment across developments. In principle one could estimate a development-specific vacancy rate based on observed tenant move-outs or the composition of tenants; however, the CHA tenant data do not cover a long enough period for this approach to be effective.

Construction of Belief Distributions from Simulation Outputs

To construct the relevant distributions from simulation results, I consider what would have happened in the simulation to an additional applicant given each choice the applicant could have made at each stage in the development choice process. For each initial choice, I take the final choice states that would have resulted from that initial choice on a random sample of application dates in the simulation as the distribution $\hat{G}_C(s, p)$. To model the continued waiting time distributions given position information in the final choice stage, $F_{j,C}(T_j \mid p)$, I use a model of continued waiting time that is flexible across initial choices and parametric in list position. For each list j and initial choice C, continued waiting time follows a beta distribution whose parameters depend on current list positions. These distributions are estimated separately for each (j, C) pair using a sample of continued waiting times in the simulation. Appendix 1.10.2 provides details of how these distributions were constructed.

1.5.3 Preferences over Assignments and Waiting Times

Given the distribution of potential applicants and their beliefs, I estimate the discount factor and parameters governing the distribution of flow payoffs using the method of simulated moments. This section describes the parameterization of flow payoffs, the moments used in estimation, and the construction and minimization of the objective function.

Parameterization of Flow Payoffs

For estimation, I choose a homothetic utility function:

$$u(c, h) = \gamma \log c + (1 - \gamma) \log h \, .$$

Here γ is the fraction of a household's disposable income that it would spend on the numeraire if unconstrained. I also parameterize the distribution of unobserved income η_i and tastes for specific development characteristics h_i. Let Z_i represent observed household characteristics other than income; let X_j represent observed development characteristics; and let X_{ij} represent interactions between applicant and development characteristics. Flow payoffs take the form

$$v_{ij} - v_{i0} = \delta_j + \underbrace{\phi_1 \log y_i - \overbrace{\phi_2 \log(y_i + \eta_i)}^{\text{outside option}} + g(Z_i)}_{\text{value of assistance}} + \underbrace{\sum_k X_{ijk}\beta_k^o + \sum_m X_{jm}\nu_{im}\beta_m^u}_{\text{match value}} + \epsilon_{ij}, \tag{1.5}$$

where δ_j is a development fixed effect that is common across applicants and (ν_i, ϵ_i) are individual-specific taste parameters not observed by the econometrician. Note that $\phi_1/\phi_2 = \gamma$. The unobserved characteristics are parameterized as

$$\eta_i \overset{iid}{\sim} TN(0, \sigma_\eta^2, -y_i, \infty) \qquad \nu_{im} \overset{iid}{\sim} N(0,1) \qquad \epsilon_{ij} \overset{iid}{\sim} N(0,1) \tag{1.6}$$

In addition to placing parametric structure on the unobservables, this parameterization adds development fixed effects and demographic shifters to equation 1.4. The development fixed effect δ_j captures the component of development quality that is common across households, and can include both observed and unobserved characteristics of the development. The value of assistance may depend on other household characteristics Z_i in addition to income. Unobserved income is parameterized so that at each observed income y_i, total income $y_i + \eta_i$ has full support on the positive real line and has a conditional expectation that increases in y_i. The matching type contains standard terms in discrete choice demand estimation: tastes for observed development characteristics that depend on observed and unobserved household characteristics (v_{im}), and idiosyncratic tastes for each development (ϵ_{ij}).

The parametric restrictions in equation 1.6 assume independence between values of assistance and match values conditional on observed characteristics, and also place restrictions on the correlation structure of match values across developments. These assumptions are not innocuous for separating unobserved heterogeneity in values of assistance and match values. As a check for sensitivity to restrictions on match value heterogeneity, in Section 1.6.3 I examine robustness of parameters governing the value of assistance to adding random coefficients for development size and location.

Moments and Objective Function

The parameters to be estimated are the discount factor and the parameters governing flow payoffs:

$$\theta \equiv \{\rho, \delta, g(.), \phi, \beta, \sigma_\eta\} .$$

I estimate θ based on moment conditions

$$E[(m_i - E(m_i \mid Z_i, \theta_0)) \mid Z_i] = 0 ,$$

where θ_0 is the true parameter vector, m_i contains features of household decisions, and Z_i contains household characteristics and choice conditions that are determined outside the model. The method of simulated moments captures these conditions in a set of moments, indexed by $q \in \{1, ..., Q\}$, for specific choice features $m_i^{(q)}$ and household characteristics $Z_i^{(q)}$:

$$\hat{g}^{(q)}(\theta) = \frac{1}{N} \sum_{i=1}^{N} \left(m_i^{(q)} - \hat{E}[m_i^{(q)} \mid Z_i, \theta]\right) Z_i^{(q)} .$$

In estimation, the conditional expectation $\hat{E}(m_i \mid Z_i, \theta)$ is estimated by simulation, and the parameter estimate $\hat{\theta}_{MSM}$ is chosen to solve

$$\min_{\theta}\ \hat{\mathbf{g}}(\theta)'\, A\ \hat{\mathbf{g}}(\theta)$$

where $\hat{\mathbf{g}}(\theta) \equiv (\hat{g}^{(1)}(\theta), ..., \hat{g}^{(Q)}(\theta))'$ and A is a symmetric, positive-definite weight matrix. I match the following choice features $(m_i^{(q)})$ and applicant characteristics $(Z_i^{(q)})$ in the data to those predicted by the simulated model:

1. **Application Rates** by income and demographic groups:

$$m_i^{(q)} = 1\{C_i \neq \emptyset\}; \qquad Z_i^{(q)} = 1\{(y_i, Z_i) \in \mathcal{Y}^{(q)} \times \mathcal{Z}^{(q)}\}$$

2. **Development Shares** among applicants' initial and final choices: for each list j,

$$m_i^{(q)} = 1\{j \in C_i\}, \qquad 1\{j = f_i\}; \qquad Z_i^{(q)} = 1$$

3. **Covariances** between applicant characteristics and characteristics of their initial development choices:

$$m_i^{(q)} = 1\{C_i \neq \emptyset\}\frac{1}{|C_i|}\sum_{j \in C_i} X_j^{(q)}; \qquad Z_i^{(q)} = 1\{(y_i, Z_i) \in \mathcal{Y}^{(q)} \times \mathcal{Z}^{(q)}\}$$

4. **Means and Variances** of chosen development size and location within and between applicants:

$$m_i^{(q)} = \frac{1}{|C_i|}\sum_{j \in C_i} X_j^{(q)}, \qquad \left(\frac{1}{|C_i|}\sum_{j \in C_i} X_j^{(q)}\right)^2, \qquad \frac{1}{|C_i|}\sum_{j \in C_i}\left(X_j^{(q)}\right)^2; \qquad Z_i^{(q)} = 1$$

5. **Means and Variances of Chosen Waiting Times** within and between applicants, by income and demographics. Let $\bar{T}_j$ be the expected waiting time for development j from initial application if an applicant's initial choice was only j. I treat this as another development characteristic and construct moments analogous to those for other development characteristics:

$$m_i^{(q)} = \frac{1}{|C_i|}\sum_{j\in C_i}\bar{T}_j, \qquad \left(\frac{1}{|C_i|}\sum_{j\in C_i}\bar{T}_j\right)^2, \qquad \frac{1}{|C_i|}\sum_{j\in C_i}\left(\bar{T}_j\right)^2;$$

$$Z_i^{(q)} = 1\{(y_i, Z_i) \in \mathcal{Y}^{(q)} \times \mathcal{Z}^{(q)}\}$$

6. **Final Choice Moments**: for all of these, $Z_i^{(q)} = 1$.

   - The fraction of eligible households who made a final choice:

   $$m_i^{(q)} = 1\{f_i \neq \emptyset\}$$

   - The mean expected *continued* waiting time of final choices, given an applicant's position information:

   $$m_i^{(q)} = 1\{f_i \neq \emptyset\}t_{f_i}$$

   - The *relative price index*, as an expected continued waiting time ratio, of the final choice compared to other developments in each applicant's choice set. If $C = \{j, k, m\}$, and the expected continued waiting times for the developments are $\{t_j, t_k, t_m\}$, then the relative price index for development j is defined

   $$R_{j,C} = \frac{1}{2}\left[\frac{t_j}{t_k}/\bar{r}_{jk,C} + \frac{t_j}{t_m}/\bar{r}_{jm,C}\right]$$

where $\bar{r}_{jk,C}$ is the mean continued waiting time ratio between developments j and k for applicants who made a final choice from choice set C. The resulting moments are

$$m_i^{(q)} = 1\{f_i \neq \emptyset\}R_{f_i,C_i}, \qquad 1\{f_i \neq \emptyset\}1\{R_{f_i,C_i} > 1\};$$

The relative price index captures whether an applicant faced a high or a low "price" for its final choice f_i, compared to other applicants who made their final choice *from the same choice set* C_i. This isolates the natural experiment created by the Cambridge Mechanism, where applicants who made the same initial choices are given different waiting time information when they make their final choices.

- The average and maximum difference in expected continued waiting time between the chosen and alternative developments:

$$m_i^{(q)} = 1\{f_i \neq \emptyset\}\left(t_{f_i} - \frac{1}{2}\left[t_k + t_m\right]\right), \qquad 1\{f_i \neq \emptyset\}\left(t_{f_i} - \min\{t_k, t_m\}\right).$$

It is useful to consider which moments are most informative about which parameters. Application rates by income and demographic groups reveal heterogeneity in the value of assistance. Since low-income and non-white households are more likely to apply for public housing, these groups value living in public housing more on average. However, not all of these households apply, and the rate at which application rates fall with income reveals unobserved differences in values of assistance and/or match values. Initial choices reveal heterogeneity in match values by arguments similar to those in Berry et al. (2004). Covariances between applicant and chosen development characteristics – for example, between an applicant's neighborhood of current residence and the neighborhoods of its chosen developments – reveal which applicants systematically prefer which types of developments. The second moments of chosen development characteristics capture unobserved differences in match values that depend on development characteristics. In addition, the number of and expected waiting times for initially chosen developments reveal unobserved heterogeneity in the value of assistance. Some high income applicants initially choose developments with short waiting times, while others choose long waiting times or select just one or two developments. To the extent that this cannot be explained by observed applicant or development characteristics, or idiosyncratic taste shocks, these differences in behavior suggest that households differ in how much they want public housing overall. Development shares reveal which developments are more desirable conditional on observed characteristics. Finally, combined with the other moments, moments capturing the sensitivity of the final choice to waiting time information separate the discount factor from heterogeneity in flow payoffs.

Change of Variables and Importance Sampling

Estimating the conditional expectation $E[m_i \mid Z_i, \theta]$ presents a computational challenge because the two-stage development choice problem is computationally burdensome to solve. A standard simulation procedure would draw unobserved characteristics $\{(\eta_{is}, \nu_{is}, \epsilon_{is})\}_{s=1,\ldots,S}^{i=1,\ldots,N}$ once, re-solve the development choice problem at each proposed value of θ given the implied flow payoffs for each simulation draw, and construct the conditional expectations

$$\hat{E}[m_i \mid Z_i, \theta] = \frac{1}{S}\sum_{s=1}^{S} m_{is}(\theta)\,.$$

This approach was computationally prohibitive in my setting because the development choice problem would have to be re-solved thousands of times for each simulation draw. To alleviate this problem, I use a technique proposed by Ackerberg (2009) that combines a change of variables with importance sampling. The key insight is that the optimal sequence of choices for an applicant depends only on their flow payoffs $v_i = \{v_{i0}, v_{i1}, ..., v_{iJ}\}$ and discount factor ρ. The technique draws flow payoffs $\{v_i^s\}_{s=1,...,S}^{i=1,...,N}$ from an initial (proposal) distribution $g(. \mid Z_i)$; computes the optimal sequence of choices, yielding features $m(v_i^s, \rho)$; and re-weights the simulation draws according to the density implied by proposed values of θ:

$$\hat{E}[m_i \mid Z_i, \theta] = \frac{1}{S}\sum_{s=1}^{S} m(v_i^s, \rho)\frac{p(v_i^s \mid Z_i, \theta)}{g(v_i^s \mid Z_i)}.$$

Because flow payoffs were drawn from $g(. \mid Z_i)$, each term in the sum is an unbiased estimate of the true conditional expectation at θ. Evaluating the objective function at proposed values of θ amounts to re-weighting the simulation draws. An additional computational benefit is that the objective function has an analytical gradient in $\theta \setminus \{\rho\}$ when $p(. \mid Z_i, \theta)$ is differentiable in θ. An outer grid search over the discount factor minimizes the objective function in θ.

Details of the simulation, optimization procedure, weight matrix, and standard errors are provided in Appendix 1.10.2. The optimal weight matrix performed poorly in my application because the moment functions are highly collinear; I used a diagonal weight matrix instead. Standard errors account for sampling error in applicant decisions and simulation error from estimating the conditional expectation $\hat{E}[m_i \mid Z_i, \theta]$. They do not yet account for estimation error in the distribution of potential applicants or their beliefs.

1.5.4 Equivalent Cash Transfers

The micro-foundation of preferences provides a way to interpret estimates from the utility model in terms of equivalent cash transfers. I use the concept of equivalent variation (EV), the cash transfer that would produce a welfare change equal to that of a public housing assignment or re-assignment. In counterfactuals, I use this concept to quantify welfare gains from alternative policies and to make interpersonal comparisons based on the social value of cash transfers to different types of households.

If household i is assigned to development j, one can measure the welfare gain from their assignment as the cash transfer EV_{ij} that would make i equally well-off outside of public housing. This value is defined implicitly by

$$v_{ij} - v_{i0} \quad = \quad v_0(y_i + \eta_i + EV_{ij}) - v_0(y_i + \eta_i)\,, \tag{1.7}$$

where $v_0(.)$ is the indirect utility function defined in equation 1.1. Concavity of v_0 implies that a household's equivalent cash transfer is increasing in their total income $y_i + \eta_i$, holding the change in flow payoffs $v_{ij} - v_{i0}$ fixed. This is intuitive – higher-income households should have greater willingness to pay for the same change in housing quality, for example. Conversely, holding $y_i + \eta_i$ fixed, EV is convex in the change in flow payoffs $v_{ij} - v_{i0}$. As a result, households with high indirect flow utility from their assignments require large equivalent transfers.

Under homotheticity, EV has the following closed form expression:

$$EV_{ij} \quad = \quad (y_i + \eta_i)\left(\exp^{v_{ij}-v_{i0}} - 1\right)\,. \tag{1.8}$$

One can use similar logic to quantify the value of living in one public housing development instead of another. Imagine giving an applicant a choice between living in two developments, A and B. The applicant can either live in development A at their current income, or live in development B and receive a (possibly negative) transfer each year. The transfer $EV_{i,AB}$ that would make household i indifferent between the two options is defined by

$$v_{iA} - v_{iB} = u_1((1-\tau)y_i + EV_{i,AB}) - u_1((1-\tau)y_i), \tag{1.9}$$

where u_1 is utility from the numeraire as defined in equation 1.3. Equation 1.9 differs from equation 1.7 because in public housing, disposable income can only be spent on the numeraire. The EV measure still depends on the household's disposable income, which is $(1-\tau)y_i$ instead of $y_i + \eta_i$. The transformation depends on its sub-utility function over the numeraire $u_1(.)$ rather than the indirect utility function $v_0(.)$. With homothetic preferences, the closed form expression is

$$EV_{i,AB} = (1-\tau)\, y_i \left(\exp^{\frac{v_{iA}-v_{iB}}{\gamma}} - 1\right). \tag{1.10}$$

1.6 Estimation Results

This section presents estimates of the distribution of potential applicants, applicants' waiting time beliefs, and preferences over assignments and waiting times.

1.6.1 Eligible Population

Appendix Table 1.12 presents coefficient estimates from the probit model predicting the probability that an ACS household was in the CHA dataset as an applicant or tenant. The point estimates reinforce the discussion of application rates in Section

1.3.2, with lower-income and non-white households much more likely to appear in the CHA dataset as either applicants or tenants. Though the individual coefficient estimates have a great deal of sampling error from the ACS, Figure 1-2 shows that the pattern of steadily falling application rates by income is consistent across estimates from bootstrapped ACS samples. In addition, the 90 percent confidence interval for the coefficient estimate on an African American household head does not contain zero.

1.6.2 Applicant Beliefs

Selected parameters governing inputs to the Cambridge Mechanism simulation are shown in Table 1.6. The annual vacancy rate per unit is calibrated to 10 percent, implying an average of 108 apartment vacancies per year. The applicant arrival rate was 345 per year during the sample period. Based on response to final choice letters, 24.3 percent of applicants become unresponsive immediately, and attrition occurs at an annual rate of 24.5 percent thereafter. Coefficients from the final choice model are also shown. Consistent with the analysis in Section 1.3.3, applicants are less likely to choose a development with a higher list position.

Table 1.13 shows the mean and standard deviation of average waiting times for each development in the simulation, and compares them to means in the data. Simulated waiting times are constructed by averaging realized waiting times across applicants housed during the simulation. Simulated waiting times match observed waiting times qualitatively. The largest developments – Jefferson Park, Newtowne Court, Putnam Gardens, and Washington Elms – have simulated average waiting times between 1.0 and 3.2 years. The smaller developments, including Mid and East Cambridge, Lincoln Way, and Jackson Gardens, have longer simulated waiting times of 3.9 to 6.2 years. Although the simulation captures which developments have longer waiting times, the simulated average waiting times are more dispersed than those observed in the data. The main reason for this is that the Cambridge Mechanism was not in steady state during the sample period. List closures before and during the sample period allowed some applicants to be housed quickly. In addition, since some developments housed only a few applicants, observed average waiting times have considerable sampling noise. Since applicants had limited information about list closures and current and future fluctuations in list lengths, a reasonable policy would have been to form beliefs based on the long-run distribution of outcomes generated by the Cambridge Mechanism in steady state.

1.6.3 Preferences over Assignments and Waiting Times

I estimated two specifications of the development choice model. Both specifications estimate fixed effects for each public housing development, for the race/ethnicity of the household head, and for whether the household currently lives in Cambridge. They include the two terms that depend on income: the value of non-housing consumption while in public housing, and the value of the household's outside option. They also include an indicator for whether an applicant lives in the same neighborhood as each development. Finally, both specifications include the random effect corresponding to unobserved income available outside public housing. Specification (2) adds random coefficients for development size and location. Specifications with random coefficients were less robust but provide a check for sensitivity to restrictions on match value heterogeneity. For counterfactuals, I use the more stable estimates from Specification (1). I first summarize the parameter estimates, and then describe features of the preference distribution that will be relevant for counterfactuals.

Parameter Estimates

Applicants discount the future at a moderate rate, and are therefore willing to trade a shorter waiting time for a preferred assignment. The first row of Table 1.7 shows the estimated annual discount factor, which is 0.90 in Specification (1) and 0.92 in Specification (2). The estimates suggest that applicants do not fully anticipate that they might exit the queue before being housed. In Specification (1), standard errors reject discount rates close to one at a reasonable confidence level.

The parameter estimates governing the value of assistance (Panel A of Table 1.7) show that while income and demographic variables strongly predict the value of public housing, there are also large unobserved differences. Households would like to spend just over 60 percent of income on non-housing consumption; the point estimate on observed income is 0.60 in Specification (1) and 0.63 in Specification (2). These estimates are consistent with high rent burdens among very low-income households and imply that the value of assistance falls rapidly with observed income. Consistently across the three specifications, households with a non-white head have higher values of living in public housing, especially African American headed households. Finally, unobserved income makes a substantial contribution to welfare. The point estimates of the scale parameter of the truncated normal distribution are between \$7,280 and \$7,540. For households with high observed incomes, the scale parameter is close to the standard deviation of the distribution of unobserved incomes; for households with low observed incomes, the standard deviation is still several thousand dollars.

The parameters governing match values (Panel B) show substantial heterogeneity in which developments are preferred. Location is an important source of predictable heterogeneity, and applicants who already live in Cambridge prefer to remain in their neighborhoods. However, a substantial component of match values are explained by idiosyncratic tastes, with estimated standard deviations of 0.115. Adding random coefficients for development size and location in Specification (2) has some effect on the coefficients governing the value of assistance, but the qualitative patterns are similar.

Features of the Preference Distribution

In counterfactuals, this paper considers the welfare and distributional consequences of allocation policy, focusing on the trade-off between matching applicants to their preferred apartments an identifying the most disadvantaged households. This section summarizes two features of the preference distribution that will drive these counterfactuals: the value of assigning each applicant to their preferred development, and the number of developments for which applicants would accept a take-it-or-leave-it offer. I report statistics based on a sample of applicants drawn from the preference distribution estimated in Specification (1). The features are summarized for all eligible households, and for two sub-groups with high values of assistance: African American households, and households with less than $15,000 of observed annual income.

There are large welfare gains from matching applicants to their most preferred developments. Table 1.8 displays medians and means of the equivalent variation (EV) from moving an applicant from a lower-ranked choice to their first choice. Since this exercise involves a comparison between two public housing developments, EV is calculated using equation 1.10. Across all applicants, the median EV between an applicant's second and first choice is 8.6 percent of observed income, or $1,425 per year, which equates to a monthly payment of $119. The mean is even larger, driven by a long right tail in the distribution. The proportional values are similar among African American and low-income households, but the dollar values are much lower for low-income households. EV from moving an applicant from their last choice to their first choice development is very large, with a median of $1,679 per month across all applicants and $578 per month among low-income applicants. A mechanism that provides lower match quality will have a substantial welfare cost.

Most applicants are only willing to live in some developments, and applicants with worse outside options are more willing to accept mismatched offers. As a result, removing choice would induce many applicants to reject mismatched offers, improving targeting on both observed and unobserved characteristics. Table 1.9 tabulates applicants by the number of developments they find acceptable, showing the total and observed incomes of each group. Some applicants are quite selective – 33 percent would only be willing to live in three or fewer developments – while 31 percent of applicants would be willing to live in any development. The latter group has much lower observed and unobserved incomes than other applicants. The patterns are qualitatively similar for African American and very low-income households, but applicants are less selective in these groups. 56 percent of very low-income applicants and 35 percent of African American applicants would accept any development.

Because the model fits substantial preference heterogeneity in both match values and values of assistance, mechanisms that affect match quality and targeting may have large welfare and distributional consequences. A development choice system that does not allow applicants to choose their preferred development will induce many applicants to reject offers, but the welfare loss from lower match quality for those who are housed will be substantial.

1.7 Evaluating Design Trade-Offs

Using the estimates from Section 1.6, I consider how the development choice and priority systems commonly used to allocate public housing would perform in Cambridge. I begin by analyzing the effects of these mechanisms on total welfare and the distribution of housed applicants, and then show how one can apply social welfare weights to decide which mechanism to use depending on one's taste for income redistribution. This exercise has non-trivial implications for which mechanisms the CHA should use, ruling out some combinations of choice and priority within a broad class of social welfare functions.

Section 1.7.1 defines a class of one-stage choice mechanisms that incorporates the range of development choice and priority systems used in practice, and describes the specific mechanisms considered. Section 1.7.2 presents results from counterfactual simulations of these mechanisms and compares them to the Cambridge Mechanism and to a full information benchmark in which the housing authority knows applicants' preferences.

1.7.1 Space of Mechanisms

This section formalizes a simple class of mechanisms that captures the key features of public housing choice and priority systems used in practice. In this class, applicants make development choices in one stage at initial application, and are ordered on the waiting list by priority group and then application date. Compared to the two-stage development choice mechanism used by the CHA, one-stage choice greatly simplifies equilibrium computation. It is also more common in practice. To isolate the long-run impacts of policy changes, I analyze counterfactual equilibria in long-run steady state.

This rest of this section formalizes one stage choice mechanisms, defines equilibrium, explains how allocations are evaluated, and describes the mechanisms explored in counterfactual simulations.

One-Stage Choice Mechanisms

A one-stage choice mechanism φ is defined by two objects:

1. A **development choice system** $\mathcal{C}_\varphi \subseteq 2^{\{1,\dots,J\}}$. Each element of $\mathcal{C}_\varphi$ is a subset of developments from which the applicant may receive apartment offers.

2. A **priority system** $\psi_\varphi : \mathcal{Z} \longrightarrow \{1, ..., B\}$ maps applicant characteristics to a priority group. Applicant i has higher priority than applicant i' in φ if $\psi_\varphi(Z_i) < \psi_\varphi(Z_{i'})$.

The mechanism operates on sequences of apartment vacancies, applicant arrivals, and exogenous applicant departures. Each vacancy $\nu \in \{1, ..., V\}$ has a date t_ν and development j_ν. Each applicant $i \in \{1, ..., N\}$ has arrival date t_i, departure date r_i, observed characteristics Z_i, and payoff vector $v_i = (v_{i0}, v_{i1}, ..., v_{ij})$. The mechanism φ runs according to the following algorithm. On each date t,

(i) Each arriving applicant ($t_i = t$) chooses a set of developments $C_i \in \mathcal{C}_\varphi$ and is placed on the waiting list for each development $j \in C_i$. On each list, applicants are ordered lexicographically by $(\psi_\varphi(Z_i), t_i)$.

(ii) Each vacancy ν with $t_\nu = t$ is offered to the first applicant on list j_ν. If the applicant accepts, it is housed and removed from all lists $j \in C_i$. If the applicant rejects, it is removed from all waiting lists and cannot reapply. This step is repeated until an applicant accepts or the waiting list is empty. If the latter occurs, the vacancy is held until the next day.

(iii) Departing applicants ($r_i = t$) are removed from all lists $j \in C_i$.

Development Choice Problem, Information, and Equilibrium

In one stage choice mechanisms, an applicant's choice problem is simpler than in a two-stage mechanism. The applicant simply considers, for each possible subset of developments it can choose, which development is likely to arrive first, and the distribution of waiting times for the first arrival. Let T_j be the random variable for the waiting time for development j if an applicant were only on the waiting list for j. The realization of T_j will depend on applicant i's date of application. The joint distribution $F_{T_1,\ldots,T_J}$ may depend on the applicant's priority $\psi_\varphi(Z_i)$. The applicant solves the following choice problem:

$$\max_{C \in \mathcal{C}_\varphi} \sum_{j \in C} w_j^C(\psi_\varphi(Z_i))(v_{ij} - v_{i0}) \tag{1.11}$$

$$w_j^C(\psi_\varphi(Z_i)) \equiv \frac{1}{\rho} E_{\psi_\varphi(Z_i)} \left[e^{-\rho T_j} \mid T_j = \min_{k \in C_i} T_k \right] \times P_{\psi_\varphi(Z_i)} \left[T_j = \min_{k \in C_i} T_k \right]$$

As in the Cambridge Mechanism, applicants do not know the state of the queue when they apply, but they do know the distribution of outcomes that they face for each possible choice $C \in \mathcal{C}_\varphi$ given their priority group $\psi_\varphi(Z_i)$. As a result, an applicant's beliefs do not depend on its application date. In equilibrium, beliefs are consistent with the distributions generated by the mechanism in long-run steady state given the distribution of potential applicants, the preference distribution $p(v_i \mid Z_i, \hat{\theta}_{MSM})$, and given that applicants choose developments according to equation 1.11.

In the counterfactual simulations, the exogenous departure model is the same as in the Cambridge Mechanism simulation, as are vacancy rates. Applicant arrivals are generated using the distribution of potential applicants and preferences estimated in Section 1.6, and choices are computed given applicants' preferences and beliefs. As before, potential applicants choose to apply if any development is preferable to their outside option. Appendix 1.10.3 provides details of how the equilibrium is computed. The algorithm iteratively updates applicant choices and their implied steady state waiting time distributions until a fixed point is reached between choices and beliefs.

Evaluating Allocations

Given sequences of inputs, a mechanism φ produces an eventual assignment $j_\varphi(i) \in \{0, 1, ..., J\}$ for each applicant, with $j_\varphi(i) = 0$ if applicant i is not assigned an apartment. A natural way to summarize the welfare and distributional impacts of a mechanism is to average characteristics of assigned applicants and their values over assigned apartments. In long-run steady state, if applicants vacate apartments at an exogenous, poisson rate, then this provides an estimate of the mean characteristics of public housing tenants at any given time. A social planner interested in maximizing the expected discounted sum of future payoffs would be interested in these statistics. To summarize welfare, I use equivalent cash transfers as a baseline measure:

$$W(\varphi) = \frac{1}{\sum_{i=1}^{N} 1\{j_\varphi(i) \neq 0\}} \sum_{i=1}^{N} EV_{i,j_\varphi(i)} \tag{1.12}$$

where $EV_{i,j_\varphi(i)}$ is as defined in equation 1.8. To summarize characteristics of housed applicants, one can do the same for transformations of applicant characteristics:

$$\frac{1}{\sum_{i=1}^{N} 1\{j_\varphi(i) \neq 0\}} \sum_{i=1}^{N} h(Z_i, v_i, j_\varphi(i)) \tag{1.13}$$

To incorporate social welfare weights into welfare calculations, one can transform equivalent variation from assignments by a function $f(Z_i, v_i, EV)$ that depends on applicant characteristics:

$$W(\varphi; f) = \frac{1}{\sum_{i=1}^{N} 1\{j_\varphi(i) \neq 0\}} \sum_{i=1}^{N} f(Z_i, v_i, EV_{i,j_\varphi(i)}) \tag{1.14}$$

In particular, this formulation allows a social planner to have different marginal values of transferring one dollar to different households.

Finally, one can compare welfare gains from different mechanisms adjusting for the total cost of the public housing program under each. This is important when mechanisms affect the income distribution of housed applicants; since rent in public housing is proportional to a tenant's income, the CHA will receive lower rent payments if it houses lower-income applicants. Administrative documents from the CHA suggest that the cost of maintaining each Family Public Housing apartment was close to $c \equiv \$14,300$ per year.[17] Subtracting tenant rent payments from this cost measure provides a reasonable lower-bound on the full economic cost of the public housing program in Cambridge. Adjusted for cost, welfare gains are

$$\tilde{W}(\varphi; f) = \frac{\sum_{i=1}^{N} f(Z_i, v_i, EV_{i,j_\varphi(i)})}{\sum_{i=1}^{N} 1\{j_\varphi(i) \neq 0\}(c - 0.3y_i)} \tag{1.15}$$

Simulated Mechanisms

The mechanisms used by the 24 surveyed PHAs in Section 1.2 can be modeled using six development choice systems and three priority systems. I computed the counterfactual equilibrium that would arise in Cambridge under each combination. The development choice systems are

[17] http://www.cambridge-housing.org/civicax/filebank/blobdload.aspx?BlobID=22801

1. **Choose One**: $\mathcal{C} = \{\{1\}, ..., \{J\}\}$. Applicants must select one development. This choice system is closest to those used in Cambridge, New York City, New Haven, and Seattle, which allow applicants to select a limited number of developments.

2. **Choose Any Subset**: $\mathcal{C} = 2^{\{1,...,J\}}$. Applicants may choose any subset of developments, as in Boston and San Antonio.

3. **Choose All or One**: $\mathcal{C} = \{\{1\}, ..., \{J\}, \{1, ..., J\}\}$. Applicants may either wait for their preferred development or take the first available offer from any development. This choice system approximates the policies used in Philadelphia, Baltimore, and Newark.

4. **Choose Neighborhood**: $\mathcal{C} = \{C_{\text{north}}, C_{\text{east}}, C_{\text{central}}\}$. Applicants choose a neighborhood from which to receive an apartment offer. Importantly, an applicant cannot choose to wait for their most preferred development.

5. **Choose All or Neighborhood**: $\mathcal{C} = \{C_{\text{north}}, C_{\text{east}}, C_{\text{central}}, \{1, ..., J\}\}$. Applicants may either choose a neighborhood or receive the first offer city-wide. Chicago uses this development choice system for Family Public Housing.

6. **No Choice**: $\mathcal{C} = \{\{1, ..., J\}\}$. Applicants must accept the first available apartment in any development; they have no choice over their assignment.

For priority systems, I model priority for higher socioeconomic status households as a priority for higher-income applicants, and lower socioeconomic status or need-based priorities as a priority for low-income applicants:

1. **Equal Priority**: Applicants are treated equally and ordered only by application date. Apart from emergency priorities that affect few applicants, several

PHAs, including the CHA, use equal priority.

2. **Low-Income Priority**: Applicants below 30% AMI are offered apartments first. Among the 24 sampled PHAs, only Seattle uses this exact policy. However, several PHAs used "need-based" priorities for households that were severely rent burdened, faced involuntary displacement, or were referred by other agencies that provide public assistance.

3. **High-Income Priority**: Applicants above 30% AMI are offered apartments first. This is the explicit policy in New York City and New Haven, and also captures priorities for working or economically self-sufficient households used by several other PHAs.

1.7.2 Welfare and Distributional Impacts of Allocation Policy

I begin by analyzing the effect of development choice systems under equal priorities and then consider the effects of prioritizing higher- or lower-income applicants. Finally, I show how distributional preferences determine which mechanism should be adopted in Cambridge. In all cases, results are reported by averaging payoffs and characteristics of housed applicants over apartments allocated in the simulated equilibrium of each mechanism, as in equations 1.12 - 1.15.

Effect of Development Choice under Equal Priority

The range of development choice systems used in practice would have large welfare and distributional impacts in Cambridge. To begin, compare Columns (1) and (6) of Table 1.10, which show the allocations from "Choose One," which forces applicants to choose their preferred development, and "No Choice," which does given applicants any choice over their assignment (other than the option to reject an apartment offer and leave the waiting list). Under "Choose One," the average housed applicant values their assignment as much as a cash transfer of $6,956; under "No Choice," the value falls to $5,399. Part of this welfare loss is driven by a reduction in match quality. While 43 percent of housed applicants are assigned to their first choice development under "Choose One," only 11 percent are under "No Choice." By inducing applicants with higher incomes and better outside options to reject mismatched offers, "No Choice" substantially improves targeting. The mean observed income of housed applicants falls from $20,509 to $17,535, and housed applicants also have worse outside options conditional on their observed characteristics. Due to lower tenant incomes, the CHA would receive lower rent payments and therefore incur a higher cost per unit under "No Choice." Adjusted for cost, "Choose One" produces 84 cents of welfare gains per dollar spent, while "No Choice" produces only 59 cents, a 30 percent decrease.

The other development choice systems produce allocations in between "Choose One" and "No Choice" in terms of match quality, targeting, and total welfare. "Choose Any Subset" and "Choose All or One," which allow applicants to select several developments as a hedge against waiting time uncertainty, have small effects on the allocation. This is because in equilibrium, waiting time uncertainty is small relative to differences in average waiting times across developments. Applicants that choose several developments are very likely to be housed in the development with the shortest expected waiting time, and would have picked that development under "Choose One." In contrast, "Choose Neighborhood" and "Choose All or Neighborhood," which allow applicants to choose their neighborhood but not a specific development, do impact assignments. Section 1.6.3 documented that many applicants would only accept one or a few developments; in Cambridge, each neighborhood contains at least three developments. As a result, neighborhood choice still induces many applicants to reject offers, lowering match quality while improving targeting.

Effect of Income-Based Priorities

Prioritizing higher- or lower-income applicants can dramatically affect targeting with little change in match quality or in applicants' values of their assigned apartments. Columns (1) - (6) of Table 1.11 summarize allocations under the three priority systems – "Low-Income Priority," "High-Income Priority," and "Equal Priority" – each under "Choose One" and "No Choice." Each choice system produces similar values of assigned apartments under the three priority systems, measured in equivalent cash transfers as defined in equation 1.8. Priority also has little effect on match quality. Under "Choose One," applicants are equally willing to wait for their preferred developments under each priority system. With "No Choice," applicants are equally likely to be offered a mismatched apartment, and although low-income applicants are more willing to accept mismatched offers, the overall effect on match quality is small.

As one would expect, income-based priorities most impact the incomes and outside options of housed applicants. Under "Choose One," average incomes are $27,950 under "High-Income Priority" and $13,581 under "Low-Income Priority." Due to the change in rents paid by tenants, priorities dramatically affect welfare gains per dollar spent. Under "High-Income Priority, Choose One," applicants value their assignments more than the cost of housing them; in contrast, they value it only two-thirds as much under "Low-Income Priority, Choose One."

Table 1.11 also illustrates how the priority and development choice systems interact. When higher-income applicants receive priority, development choice has a large effect on targeting – applicants' observed incomes fall by more than 25 percent moving from "Choose One" to "No Choice," driven by the fact that higher-income applicants are more likely to reject mismatched offers. When lower-income applicants are prioritized, moving from "Choose One" to "No Choice" provides much smaller targeting gains. Using observed characteristics in allocation policy affects the ability of choice design to screen on unobserved (or unused) characteristics.

Incorporating a Preference for Redistribution

Measuring welfare gains in terms of equivalent cash transfers implicitly places equal value on transferring resources to households at different points in the income distribution. A housing authority or social planner with a taste for redistribution would prefer to transfer dollars to a lower-income household. This section incorporates social welfare weights into comparisons among allocation mechanisms and discusses implications for the policies of the CHA and other PHAs.

In the preference model presented in Section 1.4.2, a social planner with a distaste for inequality or a preference for transferring resources to households with higher marginal utilities of income should apply higher social welfare weights to households with worse outside options. A household's utility from its outside option is determined by its total income outside of public housing, $\tilde{y}_i \equiv y_i + \eta_i$. Any monotonically increasing function $f(\tilde{y}_i)$ corresponds to a social welfare function that dislikes income inequality. To capture these social preferences in one dimension, I consider a class of social welfare functions proposed by Atkinson (1970):

$$f(\tilde{y}_i, EV; \lambda) = \begin{cases} \frac{1}{1-\lambda}\left[(\tilde{y}_i + EV)^{1-\lambda} - \tilde{y}_i^{1-\lambda}\right] & \lambda \neq 1 \\ \log(\tilde{y}_i + EV) - \log(\tilde{y}_i) & \lambda = 1 \end{cases}$$

This class of functions captures "constant relative inequality-aversion." It implies that the social value of transferring one dollar to a household with 1 percent lower income is approximately λ percent greater. An inequality-aversion parameter of $\lambda = 0$ implies no taste for redistribution; $\lambda = \infty$ corresponds to a social welfare function that only cares about welfare changes for the agent who is worst off. In addition to capturing a wide range of social preferences, this class has desirable properties. For $\lambda > 0$, social welfare increases whenever resources are transferred from higher- to lower-income households, and for any $\lambda \in \mathbb{R}$ income distributions are ranked identically if all incomes are multiplied by a constant. Within this class of social welfare functions, one can use equation 1.15 to determine which mechanism should be used given a PHA's degree of inequality aversion.

Figure 1-3 shows that under the current CHA priority system ("Equal Priority"), applicants should always have some choice over where they live for any $\lambda > 0$. The figure plots the cost-adjusted welfare measures from equation 1.15 for each mechanism, normalized by welfare under "Equal Priority, Choose One" at a range of inequality aversion parameters. Consistent with Table 1.10, "Choose One" is preferred with low inequality aversion. With high inequality aversion, "Choose Neighborhood" is preferred and actually performs better than "No Choice". In addition, "Choose One or All" and "Choose Any Subset" perform slightly better than "Choose One" with moderate inequality aversion. Appendix Figure 1-6 shows that under "Low-Income Priority," "No Choice" is preferred to "Choose Neighborhood" with high inequality aversion. When low-income applicants are given priority, completely removing choice maximizes the rates at which the most disadvantaged applicants are housed. Figure 5 repeats this exercise for priority systems under the "Choose One" development choice system. When allowing choice, to maximize cost-adjusted welfare the CHA should prioritize high-income applicants since they can be housed at a low cost. As inequality aversion increases, the CHA should begin prioritizing low-income applicants; the social value of welfare gains to lower-income households outweighs the additional cost of housing them. "Equal Priority" is preferred under moderate levels of inequality aversion.

Many of the mechanisms used by PHAs are strictly dominated in the Cambridge setting; there is a better policy for any social welfare function in the class considered. Figure 6 plots the mechanisms which form the upper envelope of the 18 mechanisms considered so far. Under low inequality aversion, the CHA should prioritize higher-income households and allow choice. If the CHA wishes to improve targeting, it should first prioritize low-income applicants but allow choice, and then, if its taste for redistribution is sufficiently high, remove choice as well. Prioritizing low-income applicants targets disadvantaged households without distorting match quality, and as a result, removing choice is a policy of last resort. A mechanism such as the one used in Los Angeles, which combines "No Choice" with priority for economically self-sufficient households, is strictly sub-optimal in Cambridge within this class of social welfare functions.

Finally, the Cambridge Mechanism is likely to perform well under moderate inequality aversion. As discussed in the next section, the mechanism "Equal Priority, Choose One" is similar to the Cambridge Mechanism, and is nearly optimal among the mechanisms considered at inequality aversion parameters between 0.9 and 1.4. If the CHA chose a welfare maximizing mechanism using this class of social welfare functions, they placed equal social value on transferring between \$1.90 and \$2.60 to a household earning \$20,000 per year, and transferring one dollar to a household earning \$10,000 per year.

The Cambridge Mechanism and a Full-Information Benchmark

The development choice systems analyzed in the previous sections abstracted from the two-stage decision problem in the Cambridge Mechanism. The effect of providing new waiting time information in the second stage may impact total welfare and the distribution of housed applicants. Column (7) of Table 1.10 summarizes the allocation that the Cambridge Mechanism would produce if applicants had the waiting time beliefs estimated in Section 1.6.2 and the same preference distribution as in the other counterfactuals. Since this computation does not enforce consistency between choices and implied waiting times, the allocation should be viewed as an approximation to the actual equilibrium that the Cambridge Mechanism would generate in steady state. Qualitatively, the Cambridge Mechanism is close to "Equal Priority, Choose One," providing good match quality for tenants and targeting applicants with slightly worse outside options than the general applicant pool. Due to some inconsistencies between the estimated preference distribution and the belief model, the Cambridge Mechanism performs even better than one-stage choice mechanisms.[18] The average value of assignments is \$8,238, or 93 percent of program cost, and 51 percent of housed applicants are assigned to their first choice development.

Another important question is how well the CHA could do if it obtained more information about applicants. Columns (8) and (9) of Table 1.10 provide a lower bound on the welfare and targeting gains that would be possible if the social planner knew current applicants' preferences and outside options, but did not know when

[18]The initial choice shares of a couple of developments were not matched perfectly in structural estimation. These developments are under-subscribed in the counterfactual simulation of the Cambridge Mechanism, but applicants believe at the initial choice stage that those developments have long waiting times. In equilibrium, applicants would substitute toward the under-subscribed developments in the initial choice stage, leading to lower match quality. This does not occur in the simulation because the equilibrium is not recomputed.

applicants would arrive and depart in the future. The results show that private information sharply limits what can be achieved. The social planner maximizes the equivalent variation from assignments in Column (8) and minimizes the outside options of housed applicants in Column (9). In both cases, the planner uses a greedy algorithm, housing the applicant with the highest social value when an apartment becomes available without taking dynamic considerations into account. In the welfare-maximizing allocation, assignments are valued 61 percent more highly than under Choose One. The social planner achieves this by selecting non-white households, which have high values of assistance, with moderately high incomes that make them require large equivalent cash transfers. The targeting-maximizing allocation sacrifices match quality and the value of assistance in order to house applicants with the worst outside options. Many PHAs already use need-based priorities that affect a small set of applicants. For example, some PHAs prioritize victims of domestic violence, the homeless, or households that are severely rent burdened or have been involuntarily displaced. An important question for future research is whether PHAs could obtain additional information about applicants that strongly predicts their outside options or preferred developments.

1.8 Conclusion

The allocation of scarce public resources often involves trading off efficiency and other policy goals, such as fairness or redistribution. This paper empirically studies such a trade-off in the allocation of public housing. Using data on the choices of public housing applicants in Cambridge, MA, I estimate a structural model of preferences for public housing that quantifies heterogeneity in applicants' preferred developments and in their overall values of obtaining assistance. The empirical strategy exploits a trade-off faced by applicants between shorter waiting times and preferred assignments as well as the structure of the allocation mechanism used in Cambridge. I use the estimated model to simulate counterfactual equilibria under allocation mechanisms that housing authorities use in different U.S. cities, focusing on welfare gains to tenants and whether the most economically disadvantaged applicants receive assistance.

In Cambridge, applicants exhibit substantial heterogeneity in their preferred developments and outside options. As a result, the range of choice and priority systems used in practice would dramatically affect efficiency and redistribution. Mechanisms allowing applicants to choose their preferred development provide large welfare gains to tenants, comparable to cash transfers of more than $6,500 per year. Mechanisms that do not allow choice would induce many applicants to reject mismatched apartment offers, allowing more disadvantaged applicants to be housed. This would lower match quality for tenants, and cost-adjusted welfare gains would fall by 30 percent. The CHA could achieve the same increase in targeting without lowering tenant welfare by prioritizing lower-income applicants and allowing choice. As a result, some of the mechanisms used in other cities are strictly dominated in Cambridge within a broad class of social welfare functions. Prioritizing higher-income applicants without allowing choice, as is done in some cities, is never optimal.

These findings yield concrete policy takeaways for housing authorities. A number of papers beginning with Nichols and Zeckhauser (1982) have argued that ordeals can increase the efficiency of public programs by more effectively targeting intended beneficiaries. My results suggest that PHAs should be hesitant to use choice restrictions as an ordeal. Because choice restrictions impose a large cost on tenants, a policy maker should only use them with very strong preferences for redistribution. In addition, PHAs already collect household information that is highly predictive of need, and they can use this information to improve targeting without creating inefficient matches for tenants. PHAs should only use choice restrictions as a last resort after establishing priorities based on these observed characteristics.

This study also raises a number of questions for future work on the design of dynamic allocation mechanisms and government-provided housing benefits. Optimal dynamic mechanisms in settings like public housing allocation are an open theoretical question. Combined with the revealed preference methods developed here, theoretical insights into optimal mechanisms could provide policy guidance for PHAs and other organizations which allocate scarce resources over time. Another direction for future work is to study how housing assistance benefits should themselves be designed. Would it be better to provide less generous public housing benefits but cover more eligible households? Should housing assistance be provided in-kind through public housing, or through private market subsidies as in the Housing Choice Voucher program? Should the government provide housing-specific subsidies at all? The revealed preference methods developed here, ideally in combination with evidence on the causal effects of different program designs, may prove useful for answering such questions.

1.9 Tables and Figures

Figure 1-1: Locations of Cambridge Family Public Housing Developments

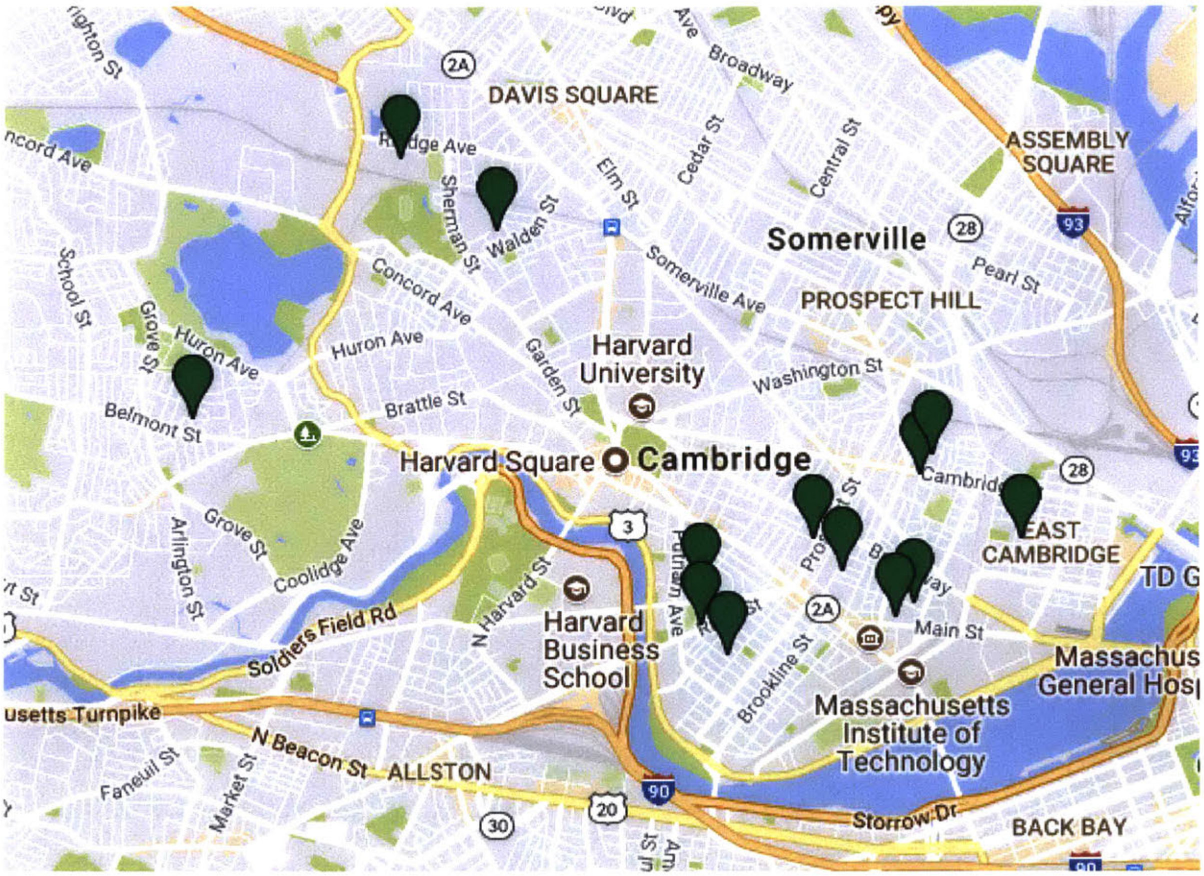

Figure 1-2: Application Rates by Income

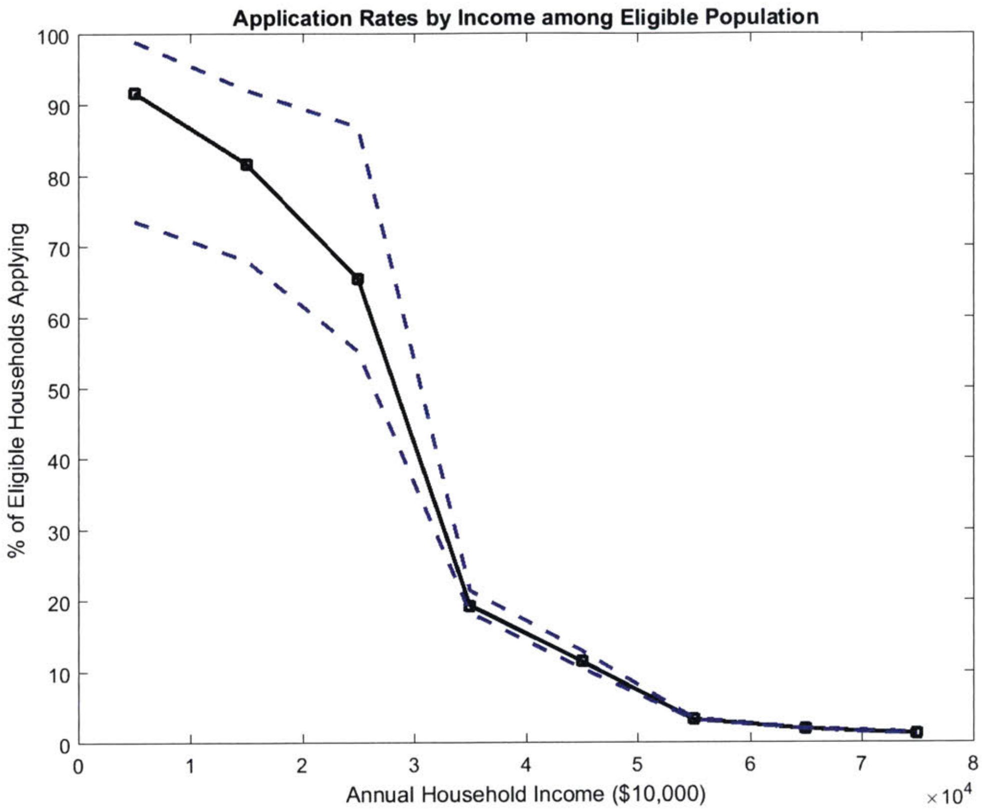

The estimated fraction of eligible households that applied for Family Public Housing in Cambridge between 2010 and 2014, by $10,000 income groups. For each group, the number of applicants is divided by the number of eligible households as estimated in Section 1.6.1. The dotted lines give point-wise 90 percent confidence bands obtained from a bootstrap that re-samples the set of eligible ACS households with replacement.

Figure 1-3: Welfare Effects of Development Choice

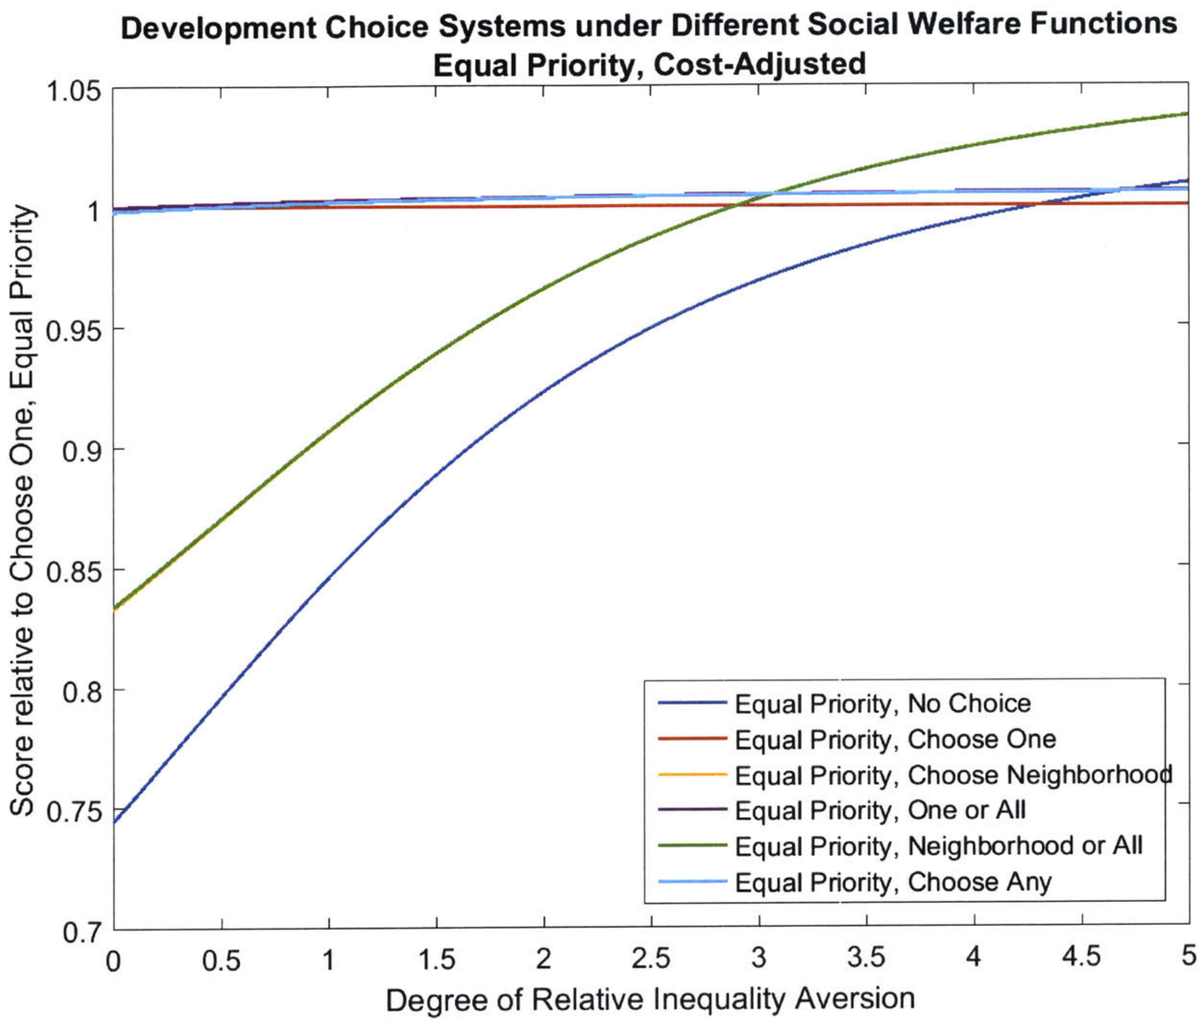

Comparison of cost-adjusted welfare gains produced by development choice systems used in practice, defined in Section 1.7.1. Applicants have Equal Priority in all mechanisms. Each point on the x-axis corresponds to a degree of relative inequality aversion. At each point, cost-adjusted welfare gains from each mechanism are normalized by the value for Equal Priority, Choose One.

Figure 1-4: Welfare Effects of Priority

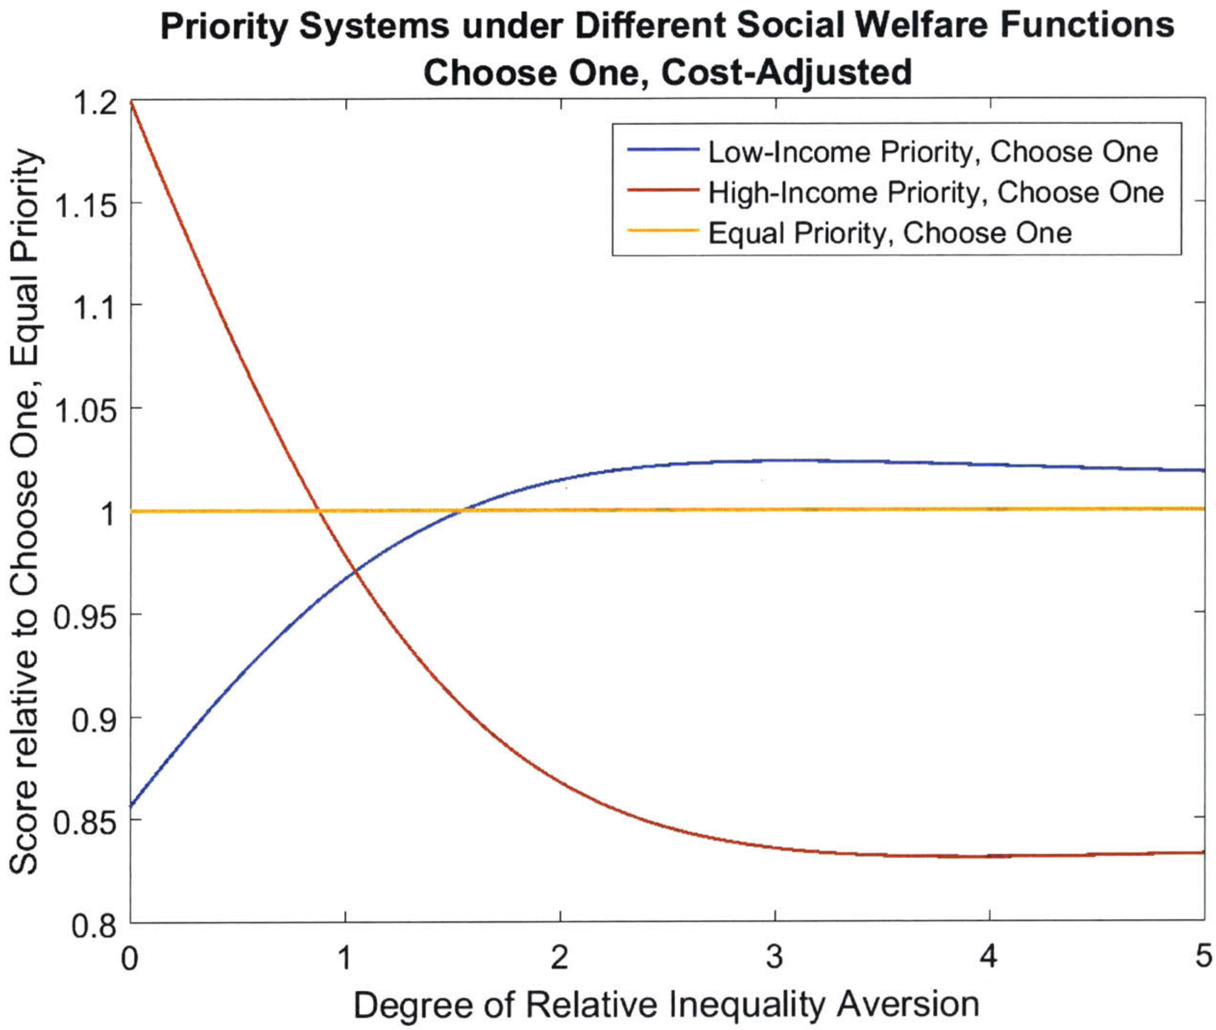

Comparison of cost-adjusted welfare gains produced by different priority systems used in practice. Low-Income Priority offers apartments to applicants below 30% AMI before other applicants, while High-Income Priority first offers apartments to applicants above 30% AMI. Applicants choose one development in all mechanisms. Each point on the x-axis corresponds to a degree of relative inequality aversion. At each point, cost-adjusted welfare gains from each mechanism are normalized by the value for Equal Priority, Choose One.

Figure 1-5: Preferred Choice and Priority Systems

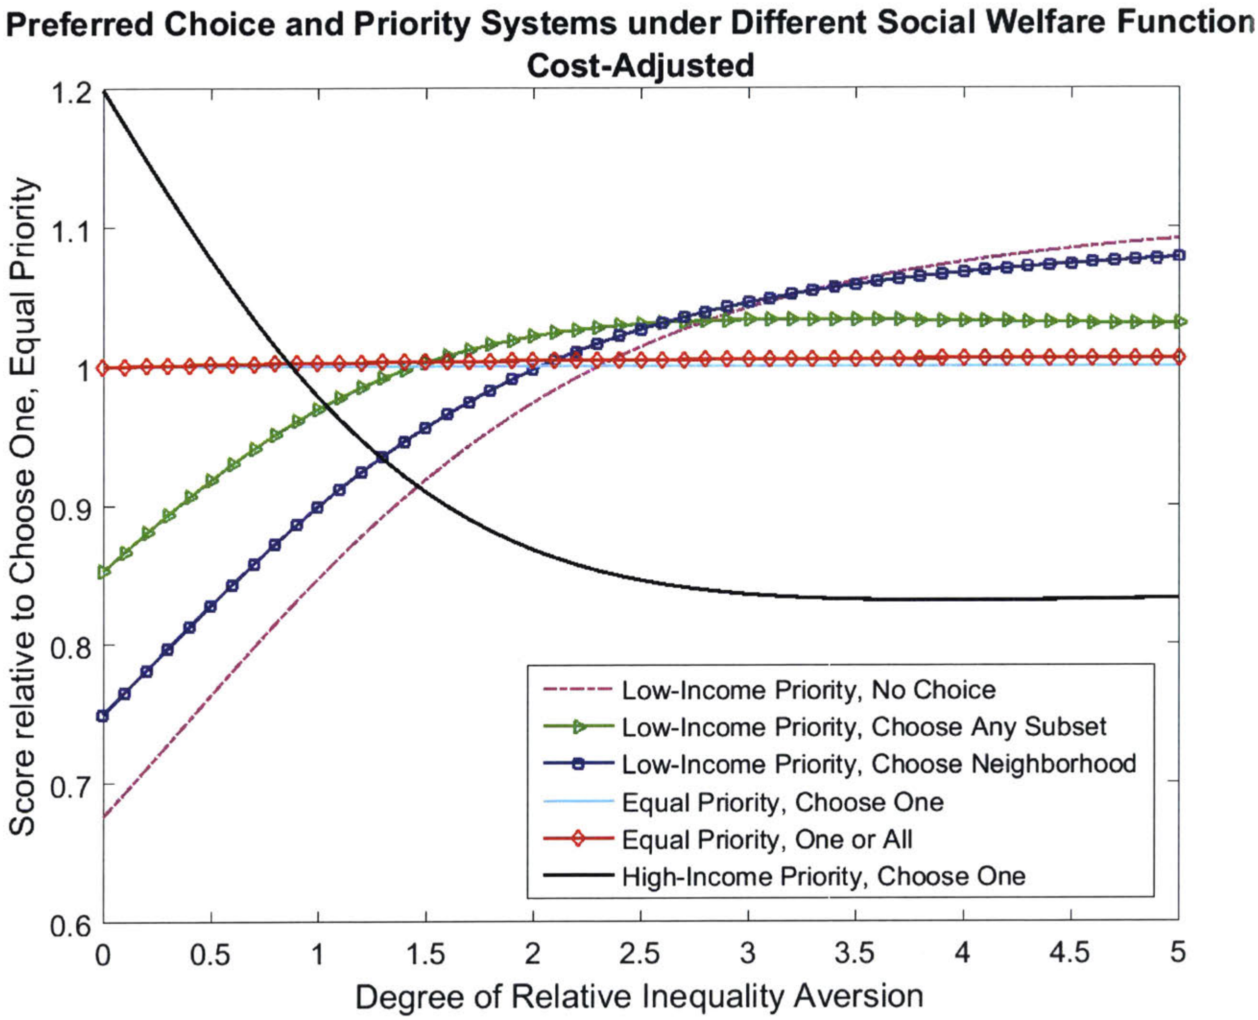

Cost-adjusted welfare gains from choice and priority systems that perform well for different degrees of relative inequality aversion. Each point on the x-axis corresponds to a degree of relative inequality aversion. At each point, cost-adjusted welfare gains from each mechanism are normalized by the value for Equal Priority, Choose One.

Table 1.1: Allocation Policies Used in Practice

Public Housing Authority (PHA) Jurisdiction	City Population in 2016	Number of Public Housing Units in 2013	Priority System	Development Choice System
Panel A: PHA's with Largest Public Housing Stock				
New York City, NY	8,537,673	175,000	Mixed	Limited Choice
Chicago, IL	2,704,958	21,150	Equal	Limited or All
Philadelphia, PA	1,567,872	15,000	Equal	Limited or All
Baltimore, MD	614,664	11,250	High SES	Limited or All
Boston, MA	673,184	10,250	Equal	Any Subset
Cleveland, OH (Cuyahoga Metro Area)	385,809	10,000	High SES	Limited Choice
Miami, FL	453,579	9,400	Equal	No Choice
Washington, D.C. *	681,170	8,350	--	--
Newark, NJ	281,764	7,750	High SES	Limited or All
Los Angeles, CA	3,976,322	6,900	High SES	No Choice
Seattle, WA	704,352	6,300	Low SES	Limited Choice
Minneapolis, MN	413,651	6,250	Low SES	No Choice
San Antonio, TX	1,492,510	6,200	Low SES	Any Subset
Panel B: PHA's comparable to Cambridge, MA (2000-3000 public housing units, 100-200K population)				
Cambridge, MA	110,650	2,450	Equal	Limited Choice
Rochester, NY *	114,011	2,500	Equal	No Choice
New Haven, CT	129,934	2,600	High SES	Limited Choice
Columbia, SC	134,209	2,140	Equal	No Choice
Dayton, OH	140,489	2,750	High SES	Any Subset
Syracuse, NY *	143,378	2,340	High SES	No Choice
Bridgeport, CT *	145,936	2,600	Equal	--
Kansas City, KS	151,709	2,050	Mixed	No Choice
Macon, GA *	152,555	2,250	High SES	No Choice
Providence, RI	179,219	2,600	Equal	No Choice
Worcester, MA *	184,508	2,470	Low SES	No Choice
Augusta, GA *	197,081	2,250	Equal	No Choice
Yonkers, NY	200,807	2,080	Equal	Any Subset

Notes: features of allocation mechanisms used by PHAs in 25 cities. PHAs were chosen based on city population and/or the size of their public housing stocks. * indicates that the PHA's administrative plan was not available online. In these cases, information was gleaned from the PHA website and application forms. A High SES priority system favors households above 30% of Area Median Income (AMI), or which are economically self-sufficient or have a working member. A Low SES priority system prioritizes households below 30% AMI, or which are severely rent burdened or have been involuntarily displaced. A Mixed priority system prioritizes some (but not all) households of both types, and an Equal priority system does not prioritize households based on socioeconomic status. Under Limited Choice, applicants must choose a small number of developments from which to receive offers. Under Any Subset, applicants may choose any subset of the developments. Under No Choice, applicants must accept the first available apartment in any development. Under Limited or All, applicants may either commit to taking the first available apartment or select a limited number of developments. In Chicago, applicants for Family Public Housing may select a specific neighborhood, but not developments within a neighborhood.

Table 1.2: Developments

	Family Public Housing Developments						
List Name	Mean Waiting Time	# Housed Applicants	# Units	Neighborhood	Tenant Income	% Black Tenants	Applicant Income
Roosevelt Mid-Rise	1.58	18	77	East	$ 18,370	41%	$ 13,930
Woodrow Wilson	1.98	2	68	Central	$ 21,181	75%	$ 15,662
Jefferson Park	2.16	62	284	North	$ 27,982	62%	$ 16,025
Newtowne Court	2.33	95	268	Central	$ 23,368	62%	$ 16,619
Washington Elms	2.92	26	175	Central	$ 31,795	61%	$ 16,237
Putnam Gardens	2.98	36	122	Central	$ 22,460	60%	$ 16,896
Corcoran Park	3.05	45	153	North	$ 26,968	65%	$ 17,923
Scattered	3.52	11	88	N/A	$ 25,480	63%	$ 17,064
Roosevelt Low-Rise	3.55	21	124	East	$ 28,929	63%	$ 18,040
Lincoln Way	3.72	2	70	North	$ 32,528	62%	$ 17,960
Jackson Gardens	3.75	9	45	Central	$ 22,352	47%	$ 17,322

Notes: characteristics of CHA Family Public Housing developments available between 2010 and 2014. Each list reflects a development or collection of units in Cambridge Family Public Housing. Roosevelt Mid-Rise contains both Family and Elderly/Disabled units. Mean Waiting Time is the mean waiting time for applicants who were housed during the sample period. Tenant characteristics reflect active tenant certifications on January 1st, 2014. Applicant characteristics reflect all applicants who selected the list as an initial choice. The "Scattered" list aggregates three lists that were available until July 2013: Mid Cambridge, East Cambridge, and River Howard Homes. In July 2013, the CHA combined Mid Cambridge, River Howard Homes, and Woodrow Wilson with Putnam Gardens, and also combined East Cambridge with Roosevelt Low-Rise. Only Putnam Gardens and Roosevelt Low-Rise were options thereafter, reflecting units from the combined lists.

Table 1.3: Applicants

	Characteristics of Eligible Population and Applicants						
	All		by Year of Initial Application				
	Eligible	Applied	2010	2011	2012	2013	2014
# Applicants	6828	1726	183	415	407	371	347
Income ($)	42,219	18,477	17,138	17,971	18,718	18,191	19,835
2 Bedrooms	76.5%	69.8%	69.9%	68.9%	69.8%	68.2%	72.6%
3 Bedrooms	23.4%	29.8%	28.4%	30.8%	30.0%	31.8%	27.1%
Lives in Cambridge	49.3%	57.4%	61.7%	55.2%	62.4%	52.6%	57.1%
Works in Cambridge	55.2%	39.7%	28.4%	36.6%	39.8%	44.7%	44.1%
Age Youngest Member	10.5	8.5	8.2	8.2	8.2	8.7	9.0
Age Oldest Member	40.0	36.7	34.7	35.7	36.6	37.7	37.7
# Children	1.25	1.27	1.25	1.39	1.27	1.24	1.16
Child Under 10	60.8%	60.8%	56.8%	56.6%	62.9%	62.0%	64.8%
Household Head Head White	55.2%	36.2%	37.2%	32.3%	38.8%	38.8%	34.3%
Household Head Head Black	19.6%	50.3%	55.7%	54.7%	47.7%	46.6%	49.3%
Household Head Head Hispanic	17.9%	19.2%	17.5%	20.2%	17.2%	20.8%	19.9%

Notes: the applicant sample consists of Family 2-3 bedroom priority applicants who made their initial development choices between 2010 and 2014. Application date is defined as the first date an applicant appears on a waiting list in the status log. Family Public Housing waiting lists were closed during the second and third quarters of 2010. The eligible population is estimated using the 2010-2014 American Community Survey (ACS). Households already living in Cambridge public housing, as well as households that applied before 2010 and were still on the waiting list during the sample period, are not counted as eligible.

Table 1.4: Initial Development Choices

		Initial Choices of Applicants						
		Selectivity				Location		
Sub-Group	Number of Applicants	2 Initial Choices	3 Initial Choices	Mean Waiting Time (Years)	Number of Units	# Central Cambridge	# East Cambridge	# North Cambridge
All	1726	12.1%	84.1%	2.89	145	1.50	0.51	0.79
$0 - 8,000	466	11.2%	85.0%	2.86	148	1.50	0.52	0.79
$8,000 - 16,000	411	10.7%	85.6%	2.87	145	1.51	0.54	0.77
$16,000 - 32,000	555	10.8%	85.2%	2.89	145	1.50	0.50	0.82
Over $32,000	294	17.7%	78.2%	2.98	142	1.48	0.49	0.77
Central Cambridge	521	9.8%	85.8%	2.89	141	1.68	0.50	0.63
East Cambridge	131	12.2%	84.0%	2.94	136	1.46	0.87	0.47
North Cambridge	338	19.2%	76.9%	2.93	147	1.26	0.37	1.11
Outside Cambridge	736	10.3%	86.1%	2.87	150	1.49	0.52	0.82

Notes: characteristics of initial choices, by applicant characteristics. Initial choice characteristics are first averaged across each applicant's chosen developments, and then averaged across applicants. Sample is Family 2-3 bedroom priority applicants who made their initial choices between 2010 and 2014. Neighborhood is based on the zip code of the applicant's contact address. East contains zip codes 02141 and 02142; Central contains 02139; North contains 02138 and 02140; and Outside Cambridge contains all other zip codes.

Table 1.5: Final Development Choice

Sensitivity of Final Development Choice to Waiting Time Information						
	No Controls		Development Controls		Choice Set Controls	
	(1)	(2)	(3)	(4)	(5)	(6)
Position on Waiting List	-0.0175		-0.0191		-0.0259	
	(0.0031)		(0.0036)		(0.0063)	
Expected Waiting Time (Years)		-0.0639		-4.051		-4.992
		(0.279)		(0.755)		(1.319)
Development FE's			X	X		
Development - Choice Set FE's					X	X
Implied Own-Price Elasticity	-0.657	-0.029	-0.747	-3.511	-1.125	-4.087
	(0.145)	(0.128)	(0.175)	(0.669)	(7.677)	(2.121)
Observations	573	573	573	573	343	343

Notes: estimates from a conditional logistic regression of final development choice on waiting time information from the applicant's final choice letter. Sample is applicants who made a final development choice between 2010 and 2014. List position is calculated for each applicant/list on the date the Cambridge Housing Authority sent the final choice letter. Continued waiting time is estimated from realized waiting times after applicants made their final choices. Columns (1) and (2) have no controls. Columns (3) and (4) include fixed effects for each development. Columns (5) and (6) include as fixed effects a full set of interactions between the development and the applicant's choice set.

Table 1.6: Inputs to Waiting Time Simulation

Parameter	Value
Apartment Vacancies	
Annual Vacancy Rate per Unit	0.10
Annual Vacancy Rate Total	108
Applicant Arrivals and Departures	
Daily Applicant Arrival Rate	0.945
Annual Applicant Arrival Rate	345
Instant Departure Probability	0.243
Annual Departure Rate	0.245
Final Choice Model	
List Position Coefficient	-0.019
Fixed Effects	
Corcoran Park	0.347
East Cambridge	-0.130
Jackson Gardens	0.292
Jefferson Park	-0.434
Lincoln Way	0.690
Mid Cambridge	0.265
Newtowne Court	0.073
Putnam Gardens	-0.299
River Howard Homes	0.000
Roosevelt Low-Rise	-0.604
Washington Elms	-0.321
Woodrow Wilson	-0.260
Roosevelt Mid-Rise	-0.876

Table 1.7: Parameter Estimates

	No Unobserved Heterogeneity (1)		Unobserved Taste for Size and Location (2)	
Annual Discount Rate	0.900	(0.011)	0.920	(0.079)
Panel A: Value of Assistance				
Head Is Black	0.465	(0.059)	0.598	(0.052)
Head Is Hispanic	0.186	(0.066)	0.141	(0.053)
Lives In Cambridge	0.042	(0.042)	-0.117	(0.061)
Log Of Observed Income	0.605	(0.032)	0.628	(0.028)
Log Of Observed And Unobserved Income	-1.000		-1.000	
Scale of R.E. Unknown Income ($10,000)	0.728	(0.216)	0.754	(0.149)
Panel B: Match Values				
From Same Neighborhood as Development	0.132	(0.045)	0.135	(0.069)
S.D. Unobserved Taste For Development Size			0.031	(0.033)
S.D. Unobserved Taste For Development Central			0.055	(0.046)
S.D. Idiosyncratic Shock	0.116	(0.005)	0.115	(0.009)
S.D. Development Fixed Effects	0.102		0.103	

Notes: estimates from the development choice model specified in Section 5.3. Both specifications include fixed effects for each development. Flow payoff coefficients are normalized by the coefficient on Log of Observed and Unobserved Income, the value of the household's outside option. The Scale of R.E. Unknown Income is the scale parameter of the truncated normal distribution governing the distribution of unknown income.

Table 1.8: Equivalent Variation to Moving from Lower-Ranked to 1st Choice Development

		All Applicants		African American Household Head		Observed Annual Income below $15,000	
		Median	Mean	Median	Mean	Median	Mean
1st Choice instead of 2nd	% Income	8.6%	11.8%	8.2%	11.2%	8.0%	10.8%
	($/year)	1,435	3,026	1,452	3,147	519	854
1st Choice instead of 3rd	% Income	16.5%	19.5%	16.1%	18.7%	15.4%	18.0%
	($/year)	2,957	4,941	3,062	5,195	1,023	1,435
1st Choice instead of Last	% Income	100.6%	106.1%	99.1%	104.2%	97.3%	102.0%
	($/year)	20,157	26,231	21,238	28,279	6,935	8,144

Notes: equivalent variation of re-assigning applicants from a less preferred development to their first choice, averaged across a simulated sample of eligible households that would apply for Cambridge public housing. The simulation uses estimates from Specification (1). "% Income" is the percentage of observed income a household would require to generate the same welfare increase while remaining in the less preferred development.

Table 1.9: Willingness to Accept Mismatched Offers

# Acceptable Developments	All Applicants			African American Household Head			Household Income below $15,000		
	%	Outside Option Value ($)	Observed Income ($)	%	Outside Option Value ($)	Observed Income ($)	%	Outside Option Value ($)	Observed Income ($)
1	17.6	34,083	35,593	13.7	46,292	44,189	6.7	12,784	8,159
2	9.1	31,483	33,652	8.2	40,980	39,951	3.8	12,950	8,714
3	6.7	28,843	31,729	6.2	37,104	37,294	3.4	11,890	7,949
4	4.7	26,722	29,629	4.5	34,269	34,570	2.6	11,523	7,948
5	4.5	25,188	28,244	4.6	30,959	32,086	3.0	11,062	7,820
6	3.8	22,909	25,602	3.7	28,062	27,952	2.7	10,432	7,442
7	3.2	21,724	24,736	3.2	27,737	28,931	2.8	10,979	8,794
8	3.4	19,859	22,927	3.3	25,292	26,696	3.3	9,624	7,741
9	3.2	20,254	23,782	3.4	25,630	27,833	2.6	9,143	7,650
10	3.4	19,268	22,784	3.9	23,337	25,631	3.2	9,911	7,889
11	3.8	17,210	20,527	4.3	20,480	22,164	3.9	9,436	8,164
12	5.2	15,310	18,901	5.7	18,480	20,580	5.8	9,101	8,274
13	31.4	7,629	12,897	35.4	9,022	13,445	56.3	4,859	7,826

Notes: distribution of number of acceptable developments, averaged across a simulated sample of eligible households that would apply for Cambridge public housing. The simulation uses estimates from Specification (1). Outside Option Value includes a household's observed income and their unobserved income outside of public housing.

Table 1.10: Effect of Development Choice

	Common Development Choice Systems						Cambridge	Full Information	
	Equal Priority							Equivalent Variation Maximizing	Targeting Maximizing
	Choose One	Choose Any Subset	Choose All or One	Choose Neighborhood	Choose All or Neighborhood	No Choice			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Welfare Gain and Cost of Allocation									
Equivalent Variation ($)	6,956	6,950	6,950	5,904	5,908	5,399	8,238	9,889	5,189
Cost per Unit ($)	8,238	8,242	8,242	8,768	8,764	9,131	8,871	7,878	11,490
Equivalent Variation per $ Cost to Gvt.	0.84	0.84	0.84	0.67	0.67	0.59	0.93	1.26	0.45
Panel B: Targeting									
Observed Income ($)	20,509	20,497	20,498	18,745	18,756	17,535	18,400	21,711	9,669
Observed and Unobserved Income ($)	16,732	16,701	16,709	14,735	14,758	13,316	14,337	15,995	6,717
% Highest Need Quartile	35.9%	36.0%	36.0%	39.8%	39.8%	42.2%	41.9%	37.3%	65.5%
Panel C: Match Quality									
% Assigned Top Choice	43.3%	42.4%	43.3%	20.0%	20.1%	10.7%	50.6%	30.5%	10.2%
% Assigned Top 3	74.5%	73.5%	74.3%	44.5%	44.6%	29.7%	84.7%	56.7%	27.7%
Panel D: Characteristics of Housed Applicants									
Waiting Time (days)	1134	1140	1135	744	744	456	759	112	77
% Black	54.4%	54.5%	54.4%	55.8%	55.8%	56.4%	56.2%	71.4%	46.3%
% Hispanic	21.1%	21.1%	21.0%	20.4%	20.5%	20.6%	21.1%	18.9%	20.3%
From Cambridge	67.7%	67.5%	67.6%	68.4%	68.4%	67.4%	69.4%	70.9%	66.7%

Notes: statistics averaged across assigned apartments in each counterfactual simulation. Cost per unit is calculated based on the CHA's maintenance and operations costs for Family Public Housing in 2014. Equivalent Variation is calculated as the equivalent cash transfer outside of public housing that would generate the same welfare change for a housed applicant as their assignment. Low-Income Priority first offers vacant apartments to applicants with incomes below 30% AMI; High-Income Priority does the same for applicants above 30% AMI.

Table 1.11: Effect of Priority System

	Common Choice and Priority Systems						Cambridge	Full Information	
	Low-Income Priority		High-Income Priority		Equal Priority			Equivalent Variation Maximizing	Targeting Maximizing
	Choose One	No Choice	Choose One	No Choice	Choose One	No Choice			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Welfare Gain and Cost of Allocation									
Equivalent Variation ($)	6,993	5,407	6,652	5,336	6,956	5,399	8,238	9,889	5,189
Cost per Unit ($)	10,317	10,835	6,006	8,217	8,238	9,131	8,871	7,878	11,490
Equivalent Variation per $ Cost to Gvt.	0.68	0.50	1.11	0.65	0.84	0.59	0.93	1.26	0.45
Panel B: Targeting									
Observed Income ($)	13,581	11,855	27,950	20,580	20,509	17,535	18,400	21,711	9,669
Observed and Unobserved Income ($)	10,832	8,939	23,049	15,712	16,732	13,316	14,337	15,995	6,717
% Highest Need Quartile	46.3%	54.3%	22.4%	36.1%	35.9%	42.2%	41.9%	37.3%	65.5%
Panel C: Match Quality									
% Assigned Top Choice	41.2%	9.1%	42.8%	11.6%	43.3%	10.7%	50.6%	30.5%	10.2%
% Assigned Top 3	73.6%	26.2%	70.5%	30.5%	74.5%	29.7%	84.7%	56.7%	27.7%
Panel D: Characteristics of Housed Applicants									
Waiting Time (days)	670	136	970	447	1134	456	759	112	77
% Black	53.1%	54.9%	57.2%	57.6%	54.4%	56.4%	56.2%	71.4%	46.3%
% Hispanic	18.8%	18.9%	23.2%	21.2%	21.1%	20.6%	21.1%	18.9%	20.3%
From Cambridge	67.6%	66.9%	67.9%	67.7%	67.7%	67.4%	69.4%	70.9%	66.7%

Notes: statistics averaged across assigned apartments in each counterfactual simulation. Cost per unit is calculated based on the CHA's maintenance and operations costs for Family Public Housing in 2014. Equivalent Variation is calculated as the equivalent cash transfer outside of public housing that would generate the same welfare change for a housed applicant as their assignment. Low-Income Priority first offers vacant apartments to applicants with incomes below 30% AMI; High-Income Priority does the same for applicants above 30% AMI.

1.10 Appendix

1.10.1 Datasets

CHA Dataset and Sample Selection

The Cambridge Housing Authority maintains a database of applicants and tenants to manage its programs and comply with HUD regulations. The dataset used in this paper is based on an extract made on February 26th, 2016. It contains anonymized records of all applicants for Cambridge public housing who were active on a waiting list between October 1st, 2009 and February 26th, 2016. This includes all households who submitted an application after October 2009, and a selected sample of households who applied before late 2009 and were still on the waiting list.

For each applicant, I observe household characteristics, development choices, and the timing and outcome of all events during the application process. Household characteristics include family size; the age, gender, and race/ethnicity of each household member; zip code of current residence; and self-reported household income. The data also record whether an applicant had priority. Development choices and waiting list events come from a time-stamped status log that records the status of each application over time. This includes the applicant's initial application date; the date it joined each waiting list; the date it was sent a final choice letter, and if it responded, its final choice; and the date the applicant was offered an apartment. I also observe the date and reason if a household was removed from the waiting list.

From the application data, I construct several objects that allow me to interpret development choices. I infer the set of developments for which each applicant was eligible based on household structure and application date.[19] I observe waiting times

[19]To reduce waiting time uncertainty, CHA merged four small waiting lists with larger lists in

for applicants who were offered apartments, both from initial application and from the date the applicant made its final choice. I also infer the information each applicant received in their final choice letter by computing the applicant's list position on the date CHA sent the letter.

For analysis, I restrict my sample to priority applicants for 2 and 3 bedroom apartments in the Family Public Housing program who submitted an application between January 1st, 2010 and December 31st, 2014. Non-priority applicants had virtually no chance of being housed, so it is unclear how to interpret their development choices. Family Public Housing applicants are a more homogeneous group than Elderly/Disabled households, and families with children are of substantial policy interest. I restrict to 2 and 3 bedroom apartments for sample size; the vast majority of Family Public Housing applicants apply for these units, and data on choices, waiting times, and list positions from each development are sparse for other bedroom sizes. Analyzing new applications between 2010 and 2014 avoids selection issues with pre-2010 applicants since some pre-2010 applicants were no longer on the waiting list at the beginning of the sample period. These restrictions produce a sample of 1,752 applicants. 26 of these applicants selected more than three developments; omitting them leaves 1,726 applicants for structural estimation.

American Community Survey

The American Community Survey (ACS) publishes anonymized, household-level micro-data covering 1 percent of the U.S. population each year. The years 2010-2014 form a 5 percent sample of U.S. households. The survey collects detailed information on each household's structure, geography, and economic and demographic

2013. As a result, an applicant's initial choice set depended on its application date.

characteristics. Data can be downloaded at `https://usa.ipums.org/usa-action/variables/group`.

The ACS contains key household-level information that determines whether a household could have appeared in my applicant sample, which contains applicants with priority for 2 and 3 bedroom apartments in Cambridge Family Public Housing. I begin with the universe of ACS households living in the state of Massachusetts. I then determine whether each household lived or worked in Cambridge.[20] Cambridge has its own city code since its population is greater than 100,000. The *CITY* field identifies whether each household lives in Cambridge, and place of work for each working household member comes from the *PWPUMA00* field. To determine a household's bedroom size, I apply the rule used by the CHA based on the age and gender of each household member and their relation to the household head. I also identify whether households would have been eligible for the Elderly/Disabled or the Family Public Housing program based on the age of the oldest household member. For households composed of three or more generations, I created separate households for the elderly members and the younger members.[21] For income eligibility, I divide the household's total income by the Area Median Income for their household size and survey year. Other characteristics of eligible ACS households, such as the race, ethnicity, and gender of the household head, are determined using ACS demographic variables.

[20]There are tens of thousands of households with veteran status in Massachusetts, so veteran status is not counted to determine which households would have had priority for Family Public Housing in Cambridge. Only a small number of applicants have veteran status, and most already live in Cambridge.

[21]According to the CHA, it is common for Family Public Housing applicants to apply with a two-generation subset of their current multi-generational household.

1.10.2 Estimation Details

Waiting Time Beliefs

This section provides details of the simulation-based procedure to estimate applicant beliefs using knowledge of the Cambridge Mechanism and waiting list data. Since applicants choose developments in two stages, select multiple developments in the first stage, and make choices based on new information in the second stage, the waiting lists for different developments move interdependently. A sophisticated applicant will account for the fact that the combination of developments selected in the first stage will jointly affect the conditions under which they make their final development choice in the second stage. They will also update their beliefs about continued waiting times given their positions on all three lists at the final choice stage. This poses a challenge for estimation since data on realized waiting times given initial choices and final choice states are sparse. A parsimonious model of dependence across lists may not be realistic or feasible.

I assume that beliefs are consistent with the steady-state distributions that the Cambridge Mechanism would generate given applicant arrival and departure rates, initial and final choice frequencies, and empirical vacancy rates. These empirical quantities can be estimated directly from application data. Combining these estimates with knowledge of the Cambridge Mechanism, I simulate steady state outcomes which quantify interdependence across lists and the option value of the timing and information of the final choice stage.

Cambridge Mechanism

Between 2010 and 2014, Cambridge ran its public housing waiting lists according to the following algorithm. Calendar time is indexed $t = 1, ..., T$. Waiting lists are indexed by $j = 1, ..., J$, where a list corresponds to a specific bedroom size apartment (2 or 3 bedrooms) in a specific development. Applicants are indexed $i = 1, ..., N$, vacancies by $\nu = 1, ..., V$. Applicant i has an arrival date t_i and a latent departure date r_i, and makes initial choice C_i. Vacancy ν occurs on date t_ν on list j_ν. For each list j, there is a sequence of trigger and batch size policies $\{(L_{j,l}, K_{j,l})\}_{l=1}^{L}$ for sending final choice letters. If fewer than $L_{j,l}$ applicants on list j have made a final choice, Cambridge sends final choice letters to the next $K_{j,l}$ applicants on list j who have not yet made a final choice. The pair $(L_{j,l+1}, K_{j,l+1})$ become the next trigger and batch policy for list j. x_{ij} is applicant i's list j position in its final choice letter, computed as the total number of applicants on list j with an earlier application date on the date the letter is sent. Finally, the coefficients for the final choice model are $(\beta, \{\xi_j\}_{j=1}^{J})$.

The Cambridge mechanism proceeds as follows. The simulation begins at $t = 0$ with empty lists, no vacant units, and an initial trigger and batch policy $(L_{j,1}, K_{j,1})$ for each list. The following occurs in each period t:

(i) Each applicant i with arrival date $t_i = t$ is added to the lists in its initial choice set ($j \in C_i$).

(ii) Each vacancy ν with $t_\nu = t$ is offered to the first applicant on list j_ν who has made a final choice. Applicant i is housed in j_ν and removed from the waiting list. If no applicants are available, the vacancy is pushed to next period (t_ν is moved to $t_\nu + 1$).

(iii) For each list j, if the number of applicants who are on list j and have made their final choice is less than the current trigger $L_{j,k}$, the following steps occur:

(a) Cambridge sends final choice letters to the first $K_{j,k}$ applicants on list j who have not made their final choice.

(b) Applicant i responds to the final choice letter if $r_i \geq t$

(c) If i responds, it chooses list j with probability

$$\frac{\exp(\beta x_{ij} + \xi_j)}{\sum_{m \in C_i} \exp(\beta x_{im} + \xi_m)}$$

(d) If i does not respond, it is removed from all lists $m \in C_i$

(e) The next trigger and batch policy, $(L_{j,k+1}, K_{j,k+1})$, is drawn for next period

Otherwise, $(L_{j,l}, K_{j,l})$ is held for the next period.

(iv) Each applicant with $t_i = t$ who has already made its final choice is removed from the list.

Inputs to Simulation

Simulation of the Cambridge Mechanism requires a sequence of applicant arrival dates t_i and the initial choice C_i and departure date r_i of each arrival; a sequence of apartment vacancies with dates t_ν on list j_ν; and a sequence of batch and trigger policies $\{L_{j,k}, K_{j,k}\}_{k=1}^{K}$ for each list j. I assume that all sequences are drawn independently and make the following parametric assumptions:

- Applicants arrive at a poisson rate α

- Each applicant departs immediately with a non-zero probability a_1 and at exponential rate a_2 after.
- Applicant choices are drawn uniformly from the empirical distribution in the Cambridge dataset
- Vacancies on each list occur at poisson rate $v_j = 0.1 * S_j$, where S_j is the number of units corresponding to list j. The sequences occur independently across developments and bedroom sizes.
- The sequence of trigger and batch policies is drawn with uniform probability from its empirical distribution in the Cambridge dataset.
- Final choice probabilities are determined by Specification (3) in Table 4, in which the latent utility of each option depends on list position and a development fixed effect.

Given these primitives, I draw inputs for a 500 year simulation and run the Cambridge mechanism. Waiting times converged after about 10 years. I used the last 490 years of the simulation to construct beliefs.

Constructing Belief Objects

The simulation produces the state of all Cambridge waiting lists every day for 490 years. To estimate the relevant distributions governing beliefs, I consider what would have happened to an additional applicant arriving on each simulation date, for each sequence of choice the applicant could have made.

To estimate $\{G_C(S_C, P_C)\}_{C \in \mathcal{C}}$, the distribution of final choice states for each initial choice C, I sample 1000 dates $t_1, ..., t_{1000}$ from the simulation. For every C, I compute the date s_C and position vector p_C that an applicant who applied on date t_s would have received, for $s = 1, ..., 1000$. These states – $\{(s_C^s, p_C^s)\}_{s=1,...,1000}$ – form an empirical measure $\hat{G}_C$.

Constructing beliefs $\{F_{j,C}(. \mid p_C)\}_{j,C,p_C}$ for continued waiting time at final choice is more complicated. There are over 1800 possible (j, C) initial and final choice combinations, and for each combination, each position vector p_C induces a different continued waiting time distribution. Even using the simulation results, there is a limit to how flexibly these distributions can (and should) be estimated. My approach is to specify a hierarchical parametric model for the continued waiting time distribution. I assume that continued waiting time follows a beta distribution

$$T_j \mid j, C, p_C \sim Beta(\alpha_{j,C}(p_C), \beta_{j,C}(p_C))$$

whose parameters depend flexibly on choices j and C and parametrically on positions p_C. For a (j, C) pair with $|C| = 3$, the position vector p_C enters the beta distribution parameters as

$$\alpha_{j,C}(p_C) = \exp\{\pi_1 p_1 + \pi_2 \log(p_1) + \pi_3 \log(p_2) + \pi_4 \log(p_3)\}$$

$$\beta_{j,C}(p_C) = \exp\{\pi_5 p_1 + \pi_6 \log(p_1) + \pi_7 \log(p_2) + \pi_8 \log(p_3)\}$$

where the π parameters are (j, C)-specific. p_1 is the position on list j, and p_2 and p_3 are the other positions. I found that this parametric specification did a good job fitting the distribution of realized waiting times from the simulation. The range of each beta distribution is $[0, \lceil \max T_{j,C} \rceil]$.

The hierarchical parameters of each beta distribution are estimated as follows: for computational speed, I take a 5% sample of application dates from the simulation denoted $\{t_d\}_{d=1,...,D}$. For each initial choice C, I calculate the position vector an applicant would have received in their final choice letter, as well as the continued waiting time for each list. From this dataset of position vectors and continued waiting times $\{p_{C,d}, t_{C,d}\}_{d=1,...,D}$, π and the upper bound of the support of the beta distribution for each $j \in C$ are estimated by maximum likelihood.

Development Preferences

Distribution of Flow Payoffs

For household i, the difference in flow payoffs between living in public housing development j and the outside option is given by

$$v_{ij} - v_{i0} = \delta_j + \phi_1 \log y_i - \phi_2 \log(y_i + \eta_i) + g(Z_i) + \sum_k X_{ijk}\beta_k^o + \sum_m X_{jm}\nu_{im}\beta_m^u + \epsilon_{ij}.$$

where

$$\eta_i \overset{iid}{\sim} TN(0, \sigma_\eta^2, -y_i, \infty) \qquad \nu_{im} \overset{iid}{\sim} N(0,1) \qquad \epsilon_{ij} \overset{iid}{\sim} N(0,1)$$

The parameters governing flow payoffs, along with the discount factor, are

$$\theta \equiv \{\rho, \delta, \beta, g(.), \sigma_\eta, \phi\}$$

Moments

To estimate the parameter vector $\theta = \{\rho, \delta, \beta, g(.), \sigma_\eta\}$, I match the following sets of moments:

- **Application Rates** by income and demographics: I currently use the following characteristics Z_i: an indicator equal to 1 for all households; indicators for annual household income in the ranges of $[X, X + 20,000]$ for X in \$5,000 intervals from \$0 to \$40,000; indicators for whether the household head is black and hispanic; and an indicator for whether the household currently lives in Cambridge. I also match the rate at which all households and households earning \$0-\$20,000 and \$20,000-\$40,000 choose three developments in their initial choice.

- **Development Shares**: There is one moment for the initial and final choice shares of each of the thirteen developments.

- **Covariances** between applicant characteristics and characteristics of their initial development choices. I match the rates at which Cambridge residents select developments in their current neighborhood of residence. There are separate moments for Central, North, and East Cambridge.

- **Means and Variances** of chosen development characteristics within and between applicants. Each of these moments is constructed for development size (# units) and whether the development is in North, East, or Central Cambridge. For households that do not apply, all moments are zero.

- **Means Variances of Chosen Waiting Times** within and between applicants, by income and demographics. The first and second time moments are interacted with income bins for \$0-\$20,000, \$20,000-40,000, and \$40,000+.

- **Final Choice Moments** are as described in the main text.

Importance Sampling and Change of Variables

I estimate the parameter vector θ based on moment conditions

$$E[(m_i - E(m_i \mid Z_i, \theta_0)) \mid Z_i] = 0\,,$$

where θ_0 is the true parameter vector, m_i contains features of household decisions, and z_i are household characteristics. A standard way to simulate $\hat{E}(m_i \mid z_i, \theta)$ in my setting would be the following:

(i) For each sampled household i, draw preference shocks $\{\eta_{is}, \nu_{ims}, \epsilon_{is}\}_{s=1}^{S}$ and realized final choice states given each possible initial choice.

(ii) At each proposed value of θ, compute v_{is} given z_i and the simulation draws . Then calculate the optimal choice at each stage given preferences (ρ, v_{is}) and beliefs. This requires solving the two-stage choice problem for each simulation draw at each proposed value of θ.

(iii) Use choices to construct the conditional expectations

$$\hat{E}(m_i \mid z_i, \theta) = \frac{1}{S}\sum_{s=1}^{S} m_{is}$$

and form moment conditions.

The problem with this procedure is that Step (ii) is computationally expensive. The optimal choice must be calculated for every simulation draw at each value of the parameter vector θ. In my application, Step (ii) takes several minutes for a reasonable number of simulation draws. Furthermore, since the objective function has no analytical gradient, an effective optimization procedure would need to evaluate the objective function thousands of times.

I use importance sampling and a change of variables proposed by Ackerberg (2009) to avoid repeating Step (ii) for each value of θ. The key insight is that an applicant's optimal decision sequence only depends on (ρ, v_i) given a choice environment. This permits a change of variables where instead of drawing $\{\eta_{is}, \nu_{ims}, \epsilon_{ijs}\}_{s=1}^{S}$, I draw (v_{is}, η_{is}) from a proposal distribution $g(v, \eta \mid z_i)$ and compute the optimal choice for each v_{is} once for each value of ρ. Then, to estimate $E(m_i \mid z_i, \theta)$, I re-weight the simulation draws at new parameter vectors $\theta_{-\rho}$:

$$\hat{E}(m_i \mid z_i, \theta) = \frac{1}{S}\sum_{s=1}^{S} m_{is}(\rho, v_{is})\frac{p(v_{is}, \eta_{is} \mid z_i, \theta_{-\rho})}{g(v_{is}, \eta_{is} \mid z_i)}$$

Since the flow payoffs and unknown income are drawn according to $g(. \mid z_i)$, the above formula provides an unbiased estimate of $E(m_i \mid z_i, \theta)$. This formulation has two desirable properties. First and most importantly, once choices $m_{is}(\rho, v_{is})$ are computed, the objective function can be evaluated quickly at each parameter vector θ. Second, the objective function is now differentiable in $\theta_{-\rho}$, which improves the speed and accuracy of optimization. A grid search over ρ minimizes the objective function in a few hours.

My application satisfies the Constant Support assumption required for this simulation procedure to yield valid conditional expectation estimates. Each payoff vector has full support on $\mathbb{R}^J$, and unknown income has full support on $[0, \infty)$ for all household characteristics Z and parameter vectors θ.

Simulation Procedure

Constructing the simulated moments involves the following steps:

1. For each eligible household i, draw S flow payoffs $\{v_{is}, \eta_{is}\}_{s=1}^{S}$ from proposal

distribution $g(. \mid z_i)$

2. Compute the optimal initial choice C_{is} for each simulation draw given v_{is}, waiting time beliefs, and discount factor ρ.

3. Draw the following objects pertaining to the final choice stage:

   - The date and position information of final selection (s_{is}, p_{is}), drawn from the distribution $G_{C_{is}}(S_{C_{is}}, P_{C_{is}})$
   - Whether the simulated applicant makes a final choice. To determine this, I compute the probability that a household would survive until date s_{is}. Each simulation draw makes a final choice with this probability.

4. If the simulation draw makes a final choice, the choice is computed given (ρ, v_{is}) and the continued waiting time distributions $F_{j,C_{is}}(T_j \mid p_{is})$ for $j \in C_{is}$.

This procedure is repeated for each candidate value of ρ. Since initial choices may change as ρ changes, I must draw final choice states and response indicators for each value of ρ, which will determine whether each simulation draw makes a final choice and, if it does, which development is chosen. To minimize simulation error, for each simulation draw I draw one final choice state for each possible initial choice and hold those draws fixed across values of ρ. This way, if a simulation draw v_{is} makes the same initial choice for two different discount factors, it will make its final choice under the same conditions (and will have the same response indicator).

It is worth emphasizing that the flow payoffs $\{v_{is}\}$ are only drawn once. Then, initial and final choices are computed once for each value of the discount factor. These choices yield choice features $m(\rho, v_{is}, x_{is})$ which do not need to be re-calculated. I will often use m_{is} for convenience, keeping in mind that choice features may depend not only on preferences but also on the conditions under which the final choice is made.

Objective Function and Optimization

Because the moments used in estimation are highly correlated, the optimal weight matrix performed poorly. The model failed to match moments key for identifying value of assistance parameters and the discount factor such as overall application rates and the mean waiting times of initial development choices. Instead, I used a diagonal weight matrix with elements inversely proportional to the sampling variance of the corresponding moment functions. I also placed more weight on moments that are important to match precisely such as application rates, variances of chosen development characteristics within and between applicants, and the final choice moments.

The proposal distribution was chosen to broadly fit choice patterns in the data, such as application rates by group. A large value was chosen for σ_ϵ ($\sqrt{2}$). Using a proposal distribution that is moderately dispersed and centered near the estimated distribution limits the variance of the importance sampling weights, and hence simulation error.

The objective function was minimized using the *Knitro* optimization package in Matlab. A gradient-based search over the parameters governing flow payoffs was conducted for a grid of annual discount factors $\beta \in \{1, 0.98, 0.96, ..., 0.5\}$. To limit numerical instability in specifications with several random coefficients, the variance of each random coefficient was constrained to be less than one million.

Inference

The standard errors in Table 6 account for sampling error in the choices of eligible households and simulation error in constructing the simulated moments. They do not correct correct for statistical error in the minimum distance procedure used to estimate the distribution of eligible households, or for statistical error in the estimated distributions governing applicant beliefs.

The asymptotic variance of the method of simulated moments estimator is

$$(G'AG)^{-1}G'A\Omega AG(G'AG)^{-1}$$

where $G = E[\nabla_\theta g_i(\theta_0)]$, $\Omega = E[g_i(\theta_0)g_i(\theta_0)']$, and A is the symmetric positive-definite weight matrix used in estimation. For a consistent estimate of G, I evaluate the gradient of the moment functions at $\hat{\theta}$:

$$\hat{G} = \frac{1}{N}\sum_{i=1}^{N} \nabla_\theta \hat{g}_i(\hat{\theta})$$

Variance in the moment functions comes from two components: sampling error in applicant choice features m_i, and simulation error in $\hat{E}[m_i \mid z_i, \theta]$:

$$\Omega = \Omega_m + \frac{1}{S}\Omega_s$$

The empirical variance of the moment functions evaluated at $\hat{\theta}$ provides a consistent estimate of Ω_m:

$$\hat{\Omega}_m = \frac{1}{N}\sum_{i=1}^{N}\hat{g}_i(\hat{\theta})\hat{g}_i(\hat{\theta})'$$

Ω_s can be estimated consistently by

$$\hat{\Omega}_s = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{S-1}\sum_{s=1}^{S}(m_{is}(\hat{\theta}) - \hat{m}_i(\hat{\theta}))(m_{is}(\hat{\theta}) - \hat{m}_i(\hat{\theta}))'$$

where

$$m_{is}(\hat{\theta}) = m(v_{is}, \hat{\rho})\frac{p(v_{is} \mid z_i, \hat{\theta})}{g(v_{is} \mid z_i)} \otimes h(z_i) \qquad \hat{m}_i(\hat{\theta}) = \frac{1}{S}\sum_{s=1}^{S} m_{is}(\hat{\theta})$$

The variance estimate is

$$(\hat{G}'A\hat{G})^{-1}\hat{G}'A\left(\hat{\Omega}_m + \frac{1}{S}\hat{\Omega}_s\right)A\hat{G}(\hat{G}'A\hat{G})^{-1}$$

1.10.3 Counterfactuals: Computational Details

To compute counterfactual equilibria, I drew one sequence of applicant arrivals along with their departure dates, characteristics, and payoffs, and one sequence of apartment vacancies. For the arrival sequence, I first draw a sequence of characteristics of potential applicants from the distribution estimated in Section 1.5.1, and then drew flow payoffs given those characteristics using the estimates from Specification (1) of the structural model. Apartment vacancies and exogenous departure dates are drawn from the distributions estimated in Section 1.5.2.

These sequences are used to compute counterfactual allocations under all mechanisms. In computing features of the equilibrium and allocation, the first 10 years were discarded to allow the waiting list to approach steady state. All applicants were eligible for all 13 public housing developments, and all waiting lists remained open during the entire simulation. This abstracts from temporary list closures (which are common in practice) in order to focus on the long-run effects of choice and priority in steady state.

To compute equilibria of lottery mechanisms allowing choice, I searched for a fixed point between applicants' choices and the implied weights $\{w_j^C(\psi_\varphi(y_i))\}_{C\in\mathcal{C}_\varphi}^{j=1,\ldots,J}$. The algorithm worked as follows. Iteration q begins with a vector of proposed weights $w^{(q)}$. The following steps then occur:

1. Each applicant's optimal choice is calculated when the applicant believes weights are given by $w^{(q)}$.

2. The waiting list is run, yielding predicted weights $w^{(q)\prime}$ with distance $D^{(q)} = ||w^{(q)\prime} - w^{(q)}||$

3. Weights are updated as a convex combination of the proposed and implied weights:
$$w^{(q+1)} = \lambda^{(q)} w^{(q)\prime} + (1 - \lambda^{(q)}) w^{(q)} .$$

The factor λ determines how aggressively the weights are updated. If $\lambda = 1$, then the weights implied by applicant choices ($r^{(q)\prime}$) are taken as the new proposal. If $\lambda = 0$, the weights are not updated at all. I began with $\lambda^{(0)} = 1$ and lowered it by 50% each time the Euclidean distance between the proposed and implied offer rates was higher than in the previous iteration ($D^{(q+1)} > D^{(q)}$). This algorithm converged quickly, requiring no more than 50 iterations before implied offer rates were less than 0.1% different than proposed rates in every mechanism.

For the Cambridge Mechanism, I did not recompute the equilibrium. Finding a fixed point of choices and implied waiting time distributions in the two-stage development choice problem would have required re-estimating the full waiting time model every iteration, which was computationally prohibitive. Instead, I use the fact that the waiting time model used in estimation was generated by the Cambridge Mechanism to justify simulating outcomes in the Cambridge Mechanism when applicants have the beliefs used in estimation. This can be viewed as an approximation to the long-run equilibrium; given preference estimates, the actual equilibrium may differ if there was misspecification or estimation error in either the waiting time or development choice models.

In the full-information allocations, the social planner uses a greedy algorithm to house applicants from the waiting list. When maximizing equivalent variation from assignments, the planner assigns each vacancy to the applicant with the highest value currently on the waiting list. This is not the strictly optimal policy because each applicant has different values for each development; it may be better to save the highest-value applicant for later and house a lower-value one. Nevertheless, it is still a useful benchmark. The targeting-maximizing allocation also uses a greedy algorithm, assigning each vacancy to the applicant with the worst outside option who is willing to accept the unit.

Table 1.12: Coefficient Estimates Predicting Probability in CHA Dataset

	Point Estimate	90% Confidence Interval
Income $0-$8,000	2.13	[0.71 , 16.83]
Income $8,000-$16,000	1.64	[0.45 , 13.05]
Income $16,000-$32,000	0.64	[-0.14 , 6.67]
Income $32,000-$48,000	-4.98	[-8.6 , -1.29]
Income Above $48,000	-6.15	[-14.69 , -2.19]
African American Household Head	4.98	[2.18 , 15.24]
Hispanic Household Head	-0.38	[-1.49 , 6.1]
Household lives in Cambridge	-2.19	[-7.02 , 0.95]

Notes: coefficient estimates predicting the probability that an eligible household from the American Community Survey was in the CHA dataset. The model uses a probit link function and is estimated by minimum distance. The point estimates use the actual ACS sample. The 90 percent confidence intervals are bootstrapped by re-sampling the ACS with replacement and re-running the estimation procedure.

Table 1.13: Simulated Waiting Times from Initial Application

Simulated Waiting Time Realizations				
	Simulation		Data	
Development	Mean	S.D.	Mean	# Obs.
Corcoran Park	2.74	1.20	3.05	45
East Cambridge	5.11	1.98	3.52	11
Jackson Gardens	6.14	1.84	3.75	9
Jefferson Park	0.98	1.11	2.16	62
Lincoln Way	3.90	2.19	3.72	2
Mid Cambridge	5.35	2.08	3.52	11
Newtowne Court	2.07	0.95	2.33	95
Putnam Gardens	3.25	1.02	2.98	36
River Howard Homes	6.18	2.17	3.52	11
Roosevelt Low-Rise	2.22	0.87	3.55	21
Washington Elms	2.30	1.39	2.92	26
Woodrow Wilson	4.13	1.69	1.98	2
Roosevelt Mid-Rise	5.03	1.85	1.58	18

Notes: realized waiting times are averaged across all housed applicants in each development.

Figure 1-6: Welfare Effects of Development Choice with Low-Income Priority

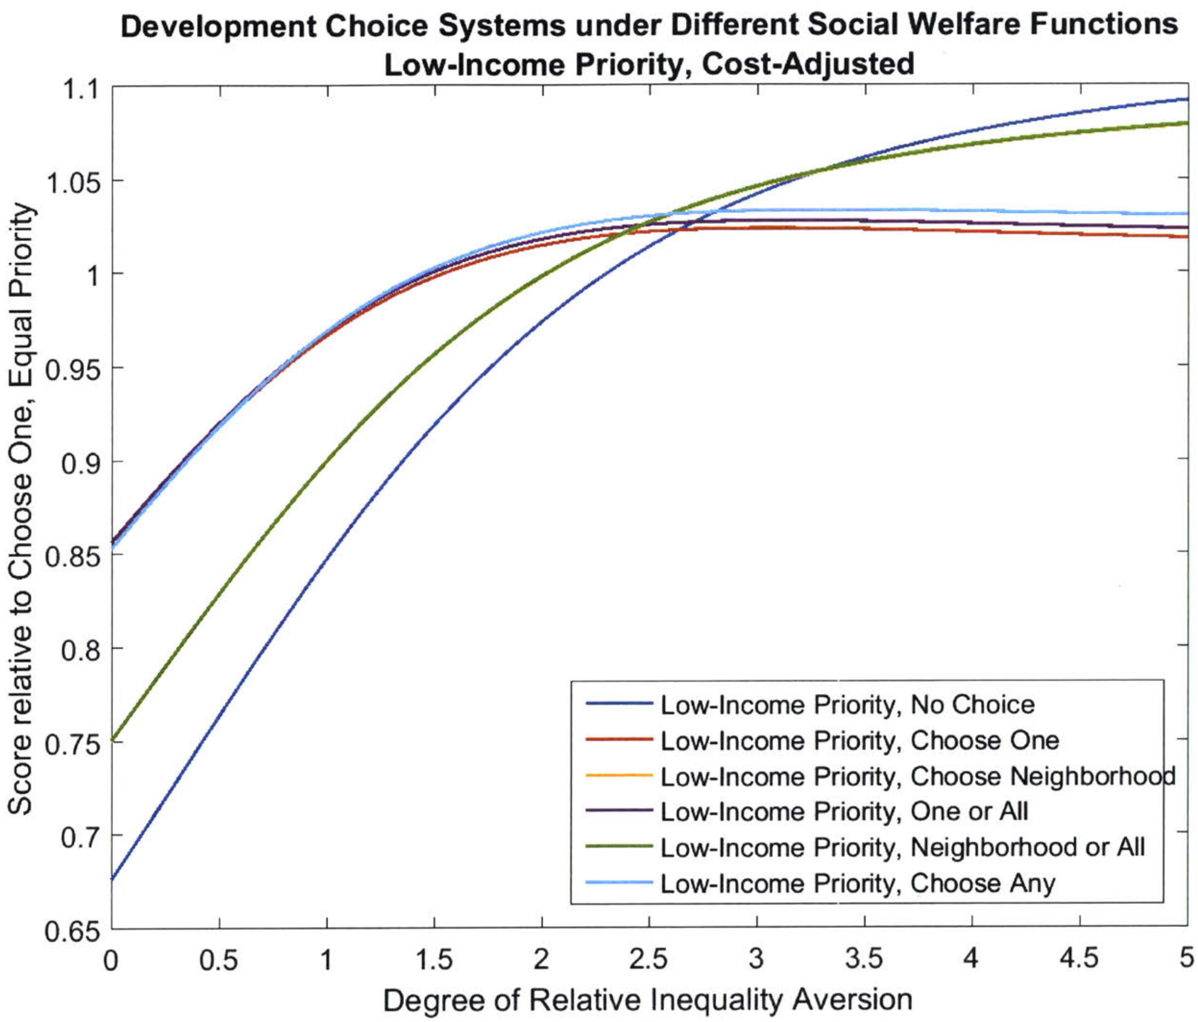

Comparison of cost-adjusted welfare gains produced by development choice systems used in practice, with priority for households with income below 30% AMI. Welfare gains are normalized by the value for Equal Priority, Choose One.

Figure 1-7: Welfare Effects of Choice and Priority, without Cost Adjustment

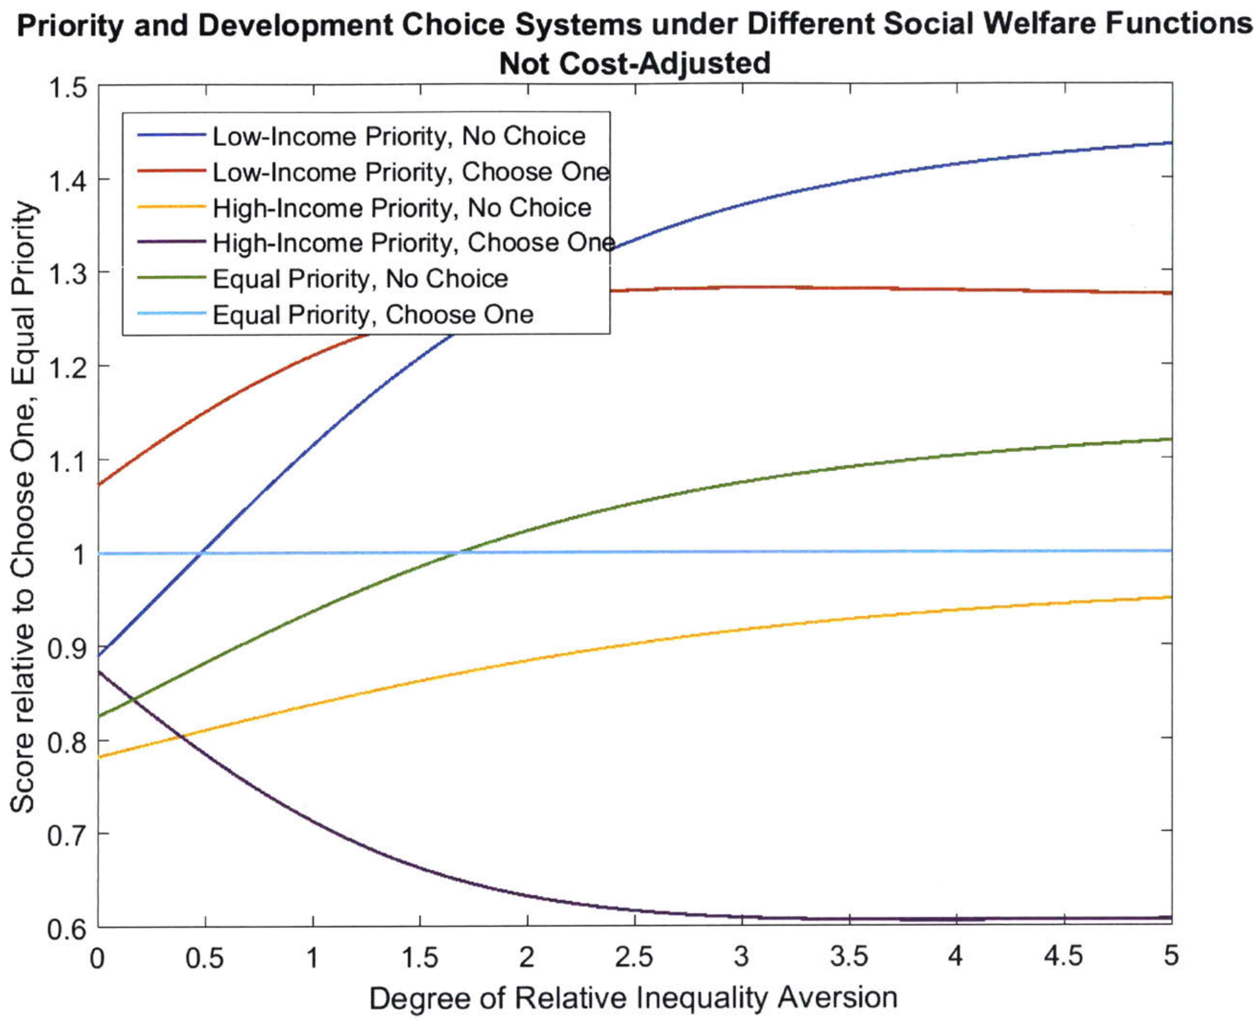

Comparison of welfare gains produced by development choice and priority systems used in practice. Welfare gains are not adjusted for cost, and are normalized by the gains from Equal Priority, Choose One.

Chapter 2

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

2.1 Introduction

As of November 1, 2017, there were 96,464 patients on the national kidney waiting list, but only 13,431 deceased donor transplants were performed in 2016.[1] Each transplant improves the expected quality and length of a transplanted patient's life while saving hundreds of thousands of dollars on dialysis (Held et al., 2016; Wolfe et al., 1999). Yet, approximately 20% of medically suitable organs extracted for transplantation are wasted in a typical year. Efficiently allocating these scarce resources,

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

reducing waste, and achieving equitable outcomes are important design objectives in this context.[2]

The allocation of deceased donor kidneys does not use money because of ethical considerations and legal restrictions,[3] making traditional price-based market-clearing mechanisms infeasible. Similar considerations motivate the use of waitlist systems to ration other deceased donor organs, public housing, nursing homes, child-care, and child-adoption. Theoretical approaches to designing dynamic assignment mechanisms have found that even qualitative trade-offs depend on the distribution of preferences.[4] This paper develops an empirical framework for analyzing waitlist mechanisms that sequentially assign objects and applies it to study the design of the deceased donor kidney allocation system. Previous empirical methods for analyzing assignment systems have been restricted to static choice settings, or have ignored dynamic considerations that agents face.[5]

We make several methodological and empirical contributions. First, we develop a procedure to estimate agent preferences using data from a dynamic assignment system, and apply it to waitlist data from New York to estimate payoffs from various types of transplants. This step is based on an optimal stopping problem faced by

[2]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.

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

[4]Agarwal et al. (2018) compare the results in Su and Zenios (2004), Leshno (2017), Arnosti and Shi (2017), and Bloch and Cantala (2017b) 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.

a patient when she is offered a kidney. Second, we define a notion of steady-state equilibria for a broad class of assignment systems that is ammenable to computation and counterfactual analysis. Finally, we use these techniques to describe how various mechanisms perform on key measures: efficiency, equity, and organ waste.

The existing allocation system 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 waiting list in decreasing order of organ-specific points. The decision of whether or not to accept an offer remains with the patient and the transplant surgeon. The organ is allocated to thc 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. Indeed, Agarwal et al. (2018) provide empirical evidence suggesting that acceptance rates are lower for patients that are less likely to receive offers in the future. Therefore, consistent with dynamic considerations, patients with a higher option value of waiting are more likely to refuse an offer for an organ.

We therefore model an agent's decision to accept an offer as a 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 uses the probability that an organ of a given type is accepted by a patient and detailed knowledge of the mechanism to recover the value of a transplant. This value depends on observed and unobserved patient and donor characteristics. Our technique adapts methods based on inverting conditional choice probabilities in this continuous-time problem (Arcidiacono and Miller, 2011; Arcidiacono et al., 2016; Hotz and Miller, 1993) to suit dynamic assignment mechanisms. The technique eases computation of the continuation value relative to more efficient 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 transplanting various donors to various patients are intuitive. All patients value certain characteristics that are correlated with organ quality – for instance, we estimate that younger donors are preferred by all patients, as are immunologically similar matches. However, there is also significant match-specific heterogeneity in values. For example, we find that older patients place less value on younger donors as compared to younger patients. Our estimates correlate well with predicted life-years gained from each transplant, and the fit of our choice probabilities is significantly better than previous approaches that abstract away from donor unobserved heterogeneity.

The allocation mechanism can affect match quality because there is substantial match-specific preference heterogeneity. Moreover, our descriptive results show that the current mechanism produces significant mismatch on several dimensions. For instance, young patients are often allocated older donors, while many old patients receive young donors. Reallocating young donors to young patients could improve efficiency. There are many other dimensions for reallocation gains. To take a systematic approach to finding these gains, we study a few alternative mechanisms. The first takes a greedy approach to improving efficiency by offering organs to patients that are predicted to benefit from them most. The second attempts to maximize equity by offering organs to patients with the most sensitive immune systems, who have the most limited transplant opportunities. The 2014 reforms incorporated aspects of both these designs. Finally, we consider a benchmark random order waitlist in which patients do not receive priority based on waiting time or other characteristics. This mechanism encodes a form of procedural fairness and has the additional benefit of reducing organ waste. We compare the outcomes from these mechanisms to the pre- and post-2014 assignment systems.

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

Second, we need to ensure that these counterfactuals are indeed identified. In dynamic models such as ours, counterfactuals may not be invariant to normalizations imposed during estimation (see Aguirregabiria 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 possibilities 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.

Related Literature

The design of scoring rules for organ allocation has been an active area of research in the Medical and Operations Research communities (Bertsimas et al., 2013; Kong et al., 2010; Su and Zenios, 2006; Su et al., 2004; Zenios, 2004; Zenios et al., 2000). 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. Our empirical evidence suggests that ignoring dynamic incentives may result in biased predictions.

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,[6] 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

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

approaches for analyzing centralized assignment systems (see Abdulkadiroglu et al., 2017b; Agarwal, 2015; Agarwal and Somaini, 2015; Fack et al., 2015b, for example). These previous approaches have focused on static assignment mechanisms. The only exception to our knowledge is Waldinger (2017). It presents a model of public housing choice in which agents face a two-stage decision with a portfolio choice problem in the first stage.[7] Our methods, in contrast, pertain to an optimal stopping rule that differs from all 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 well-developed. Abdulkadiroglu et al. (2017b) 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 (2017b), 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 (Hotz and Miller, 1993; Pakes, 1986; Rust, 1987; Wolpin, 1984), particularly recent developments in continuous-time versions of these models (Arcidiacono et al., 2016), and the estimation of dynamic games (Aguirregabiria and Mira, 2007; Bajari et al., 2007; Pakes et al., 2007; Pesendorfer and Schmidt-Dengler, 2008). Additionally, we employ a model of beliefs and an equilibrium notion that bears resemblance to concepts aimed at making the analysis of dynamic games tractable (Fershtman and Pakes,

[7]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, 2016a; Thakral, 2016; van Ommeren et al., 2016). A key difference is that these approaches are based only on final assignments instead of detailed choice data.

2012; Hopenhayn, 1992; Weintraub et al., 2008). We discuss the relationship to this literature when we develop our approach.

2.2 Background, Data and Descriptive Evidence

2.2.1 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 biological testing or examinations of the kidney. These final decisions need to be made without much delay, usually within an hour. A donor's kidneys are then allocated to the highest priority patients on the waitlist that were willing to accept the organs.

To minimize the number of patients that are contacted, UNet first uses the available donor characteristics to exclude the set of patients that are not biologically compatible with the donor. This may occur either because of blood-type incompatibility or because the patient's immune system would react negatively to the donor's tissue-type.[8] Next, UNet screens out patients that have listed pre-set exclusion criteria within the system. These are kidneys that are transplantable, but which a patient has determined to be undesirable because of donor characteristics such as age, health conditions, and kidney function measures. UNet then orders the remaining set of patients based first by priority type, and then by the number of points. Finally, it breaks ties in order of time waited.

A patient's priority type and points for each kidney depend on both donor and patient characteristics. Broadly speaking, priorities are based on geography and a few patient and donor characteristics, while the points system takes into account tissue-type similarity, pediatric patients, patient sensitization, and waiting time. The detailed priority system during our sample period (2010 - 2013) is described in policy section 8 in OPTN (2014).[9] There are many different priority types motivated by both equity and efficiency concerns. For example, the highest priority for young and healthy donors (referred to as standard criteria donors in the mechanism) is given

[8]The immune system tags foreign objects (antigens) with antigen-specific antibodies so that white blood cells (leukocytes) can defend against them. A patient's immune system will attack any antigen for which she has an antibody. Each donor has up to 6 specific types of human leukoctye antigen (HLA) proteins out of a set of hundreds of possible types. Each patient has antibodies to some subset of these HLA antigens. A transplant recipient's immune system will attack the donor's kidney and reject the organ if the recipient has an antibody to one of the donor antigens. A recipient is tissue-type compatible with a donor's kidney if she has no antibodies corresponding to the donor's antigens *(Danovitch, 2009).* A transplant between certain incompatible patient-donor pairs has become possible due to development of desensitization technologies (see, e.g., Orandi et al., 2014), but compatible transplants are preferred.

[9]On December 4, 2014, the kidney allocation rules were modified (see Israni et al., 2014; OPTN, 2017, for details). We chose not to study data for a period prior to the reform to rule out anticipatory effects of the change and to avoid using post-reform data because agents may be adapting to the new system. Reports from the United Network for Organ 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).

to patients with perfect tissue-type matches. Such matches are rare but particularly valuable because the organ is likely to function for longer. Priority is also given to patients in the same local area as the donor, to patients with sensitive immune systems,[10] and to pediatric patients. Donors that are not accepted by candidates in the local area are offered to patients in the broader geographic region and then to all patients, nationwide. Within each priority group, patients are awarded points based on the number of years they have waited for a kidney, the extent to which a patient's tissue type matches the donor's, and whether they are pediatric patients. The priority system is simpler for less healthy donors, also referred to as expanded criteria donors (ECD). The system offers these donors according to geography with points for waiting time and for tissue-type similarity. These donors are only offered to patients that have actively decided to consider such offers.

2.2.2 Data and Descriptive Analysis

This study uses data from the Organ Procurement and Transplantation Network (OPTN). The OPTN data system includes data on all donor, 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. 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. NYRT is the largest donor service area (DSA), in terms of number of patients, that

[10]Specifically, patients with a Calculated Panel Reactive Antibodies (CPRA) greater than 80% get higher priority than patients with a CPRA between 21% and 79%, followed by the remaining patients.

used the standard allocation rules in the United States prior to 2014.[11]

The primary dataset on the waitlist, the Potential Transplant Recipient (PTR) dataset, contains the offers made and patient 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 that are collected in 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.

Patients and Donors

Table 2.1 shows the 9,917 patients that were registered with NYRT at some point during our sample period. 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 important indicators of kidney health (e.g. total serum albumin), patient health (e.g. body mass index, age), 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 overall waiting time increasing over time. A variable that will be important in our descriptive analysis is a patient's immune sensitization as measured by the patient's Calculated Panel Reactive Antibodies (CPRA). A patient's CPRA is the percentage of donors in a representative sample with whom she is tissue-type incompatible. This measure is

[11]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.

calculated from blood tests conducted to determine the set of antigens to which a patient has an immune response. The average CPRA is about 12%, which indicates that there is a bit 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.

Panels B describes the patients that join and leave 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 2.2 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 all 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 significant 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. The mean is an artifact of a skewed distribution whereby undesirable kidneys are rejected by a large number of patients. 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 each of their viable kidneys was accepted are 151.25 and 15 respectively.[12] These columns also show that our observable donor covariates (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 were ultimately offered outside NYRT. For these donors, only about 0.3 kidneys were transplanted per donor, indicating that many kidneys were discarded. This is consistent with the hypothesis that donors that are not accepted within the local area are likely to be undesirable.

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. For example, compared to donors recovered from NYRT, the average donor offered is older, less likely to 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

[12]In some cases, donors have only one viable kidney for donation. For this reason, the number of transplanted kidneys amongst donors with no discards is less than 2.

only 0.75 kidneys.

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.

Waitlist, Offers and Acceptance Rates

We now describe the waitlist system, the offers made and the acceptance rates. 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. However, the priority system is not well approximated by a first-come first-served queue. Throughout the paper, we only consider potentially transplantable offers by excluding any offers where the donor had antigens that were unacceptable to the patient, or was blood type incompatible.

Figure 2-2 plots waiting time against position on the list. We only consider patients for whom the donor met the 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. The shaded region depicts the 95% confidence intervals for the means. A decrease in mean waiting time is apparent as we go down the list. This should be expected given that the system awards priority to patients that have waited longer. Nonetheless, these statistics do not directly speak to whether the system is well approximated by a first-come first-served queue. We calculated the fraction of times a pair of patients that were offered the same donor is ordered identically on the list for the next donor they are both offered. We found that this fraction is 81.5%. This fraction would be 100% in a first-come first-served system, while in a random priority system it would be about 50%. Therefore, while waiting time is an important determinant of priority, the deceased donor kidney allocation system is not well approximated by a first-come first-served system.

Figure 2-1 shows the fraction of offers accepted by the (donor-specific) position of the patient. The acceptance rate in the first few positions is much higher than later positions, but is still only between 15 and 20 percent. As we go down the list, only about 1% of offers are accepted. This sharp decline near the top occurs for two reasons. The first reason is that for some donors, the first few offers are made to patients with a perfect tissue-type match. This feature of the priority system generates a high acceptance rate near the top because offers of an organ with a perfect tissue-type match are rare and extremely valuable. The second reason, which is the predominant one, is that desirable kidneys are likely to be accepted near the top of the list and not offered to many patients. In contrast, undesirable kidneys are offered to many patients in an attempt to place them. This causes the overall acceptance rate (as a fraction of offers) to be extremely low as we go down the list because of the changing composition of donor quality. It also results in an extremely large number of offers for each patient, although most offers are of undesirable organs. It is important to remember that the majority of offers are from very undesirable donors when interpreting these low acceptance rates.

Table 2.3 shows the rate at which patients receive offers and the overall acceptance rates. Panel A shows all feasible offers, including offers that did not meet the patient's pre-set criteria. It shows that a typical patient receives about 220 offers per year. On average, the offer rate from donors recovered in NYRT is much lower, about 40 offers per year. Recall that donors that were recovered elsewhere are typically undesirable. The panel also shows that patients with sensitized immune systems – that is, patients with CPRA above 80% – receive many fewer offers from transplantable donors even though the system gives them higher priority. In panel B, we study the offers thatmet the pre-set criteria. The typical patient still gets such an offer every 3 to 4 days. However, a typical patient only receives approximately two offers per month from donors recovered within NYRT. Panel C restricts to offers within the first 10 positions, that is, offers from donors that are likely to be desirable. These offers are rare, and the typical patient can expect to receive less than one such offer each year.

The table also shows interesting patterns of acceptance behavior. The acceptance rates in panel C are much higher than those in panels A and B. Within each panel, we can see that offers from desirable donors are more likely to be accepted. For example, the acceptance rate is higher for kidneys recovered in NYRT, and much higher for kidneys with a perfect tissue-type match. Finally, the last two columns show that acceptance decisions are correlated with kidney quality as predicted by a measure of life-years from transplantation (LYFT) proposed by Wolfe et al. (2008). These LYFT estimates are the predicted median quality-adjusted life-year gains for a patient using survival models of patient life with and without a functioning kidney transplant. The columns show that the average offer has an LYFT gain of about 5 years, but the average accepted offer has a higher LYFT. Additionally, LYFT is higher for the top 10 offers that met the screening criteria. However, these differences are perhaps smaller than expected, and not all differences are consistent. For example, offers in panel B have a lower LYFT than offers in panel A. This suggests that the LYFT model may be missing some characteristics that should be included in the model.

Evidence on Mismatch

Table 2.4 looks at the outcomes from the mechanism and provides suggestive evidence of mismatch between donors and transplanted patients. Panel A shows the outcome in terms of whether or not a patient receives a 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. The chances of receiving a live donor, however, does fall with age. Panel B describes the age of the transplanted donor for those that receive a kidney through the deceased donor waitlist. We see that pediatric patients are very likely to receive a transplant from a young donor. Again, older patients are less fortunate. 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.

Evidence on Response to Dynamic Incentives

An important assumption in our framework is that agents respond to dynamic incentives. One implication of this assumption is that patients for whom the option value of waiting is lower should be less selective. Agarwal et al. (2018) 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 compares two identical population 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. (2018) use the likelihood that a patient is biological compatibility 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 2-6a and 2-6b 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 2.10 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. These risks are less likely to be justified unless the donated organ is of high quality. Taken together, these results are consistent with dynamic incentives being an important driver of acceptance decisions.

2.3 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.

2.3.1 Notation and Preliminaries

Consider a sequential assignment mechanism in which objects (indexed by $j \in \mathcal{J} \subseteq \mathbb{N}$) are offered to agents (indexed by $i \in \mathcal{I} \subseteq \mathbb{N}$) waiting on a list. Let x_i and α_i respectively denote observed and unobserved characteristics of agent i; likewise, 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 all characteristics.

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, α) are independent and identically distributed (iid) according to the cumulative distribution function (CDF) $F_{X,\alpha}$. Similarly, each object's characteristics (z, η) are drawn iid from CDF $F_{Z,\eta}$ 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.

2.3.2 Mechanisms and Primitives

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 $s_{ij} = s(t_i; x_i, z_j)$ is calculated for all agents on the waitlist. Ties in $s(t_i; x_i, z_j)$, if any, are broken using a known tie-breaking rule. For example, ties could be broken either uniformly at random or by t_i.[13]

- 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. An object cannot be allocated to an incompatible agent.

[13]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.

- Step 4 (Arrivals and Departures): An agent is removed from the waitlist once an object has been assigned to her. Other agents may join or leave the list.

Within the set of general offer-based waitlist systems, the primary restriction in the class considered here 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 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, and minimize the complexity of communication between agents and the mechanism designer. 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 priority-based offer mechanisms.[14] In counterfactual analysis, we will compare assignments that result from various mechanisms that obey this structure to benchmark optimal assignments. An analysis of pseudo-market market mechanisms (Hylland and Zeckhauser, 1979), which is not trivial even in static assignment settings with priorities (He et al., forthcoming), is left for future work.

Typical administrative datasets from such assignment systems containinformation on all characteristics used to determine the priority score $s\left(t_i; x_i, z_j\right)$ 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 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 the context of the allocation of deceased donor kidneys is

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

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.

Payoffs

There are three types of primitive payoffs in the model. The first is the (expected) flow payoff from remaining 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. The second is the (expected) net-present value from from departure without an assigment, $D_i(t)$. In our application, departures occur due to live donor transplants, death, or transfers to other listing centers (Table 2.1).[15] We view $D_i(t)$ as incorporating any of those reasons. Finally, we have the (expected) net present value of agent i being assigned a compatible object j after waiting t periods, denoted $\Gamma_{ij}(t)$.

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 et al., 2017b; Agarwal and Somaini, 2015). This re-

[15]In this case, 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.

striction 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 patient.

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_{ij}*:*

$$\Gamma_{ij}(t) = \Gamma(t, x_i, \alpha_i, z_j, \eta_j) + \varepsilon_{ijt}. \tag{2.1}$$

(ii) The random variable ε_{ij} *is independent of* $(t,\ x_i,\ \alpha_i,\ z_j,\ \eta_j)$ *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, α_i, t)*.*

Restrictions on ε_{ij} 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 2.4 below. Our framework can also be applied to other parametric forms for G.

Assumption 1(iii) and the restriction on $\Gamma(\cdot)$ exclude unanticipated time-varying persistent agent-level unobserved heterogeneity. The assumption simplifies the evolution beliefs about the expected payoffs of waiting. We discuss analytical challenges in relaxing this assumption once we have laid out our estimation approach in section 2.4.

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.*

(ii) The departure rate is non-negative, and given by $\delta_i(t) = \delta(t; x_i, \alpha_i)$. Further, each agent has a terminal date $T_i < \infty$ at which departure occurs for sure.

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.[16]

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 2.1 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 patients leave the deceased donor list once they have found a compatible living donor and do not respond to the design of the kidney waitlist. This assumption is motivated by the fact that living donors are medically superior to deceased donors and produce better transplant outcomes in terms of patient and graft survival. Living donor superiority is partly driven by the higher medical quality of living donor kidneys, and also by the fact that living donation

[16]It is fairly 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.

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 with the kidney still functioning. The rate of late graft failure for adult patients transplanted in 1991 was 14.7 years for deceased donor kidneys, compared to 26.6 years for living donor kidneys.[17] 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 they were qualified based on medical criteria.[18] 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.

2.3.3 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

[17]Some of this difference may be due to selection into who receives each type of transplant. Hart et al. (2017) report the unconditional 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 5-year 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)).

[18]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.

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 given the strategies played by agents that generated the data (Bajari et al., 2007; Pakes et al., 2007). Solving for counterfactuals will require a notion of equilibrium, which we discuss in section 2.6. This section starts by describing a general formulation of this single-agent problem before making simplifying restrictions.

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 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 that arrives at time t_i as long as her score, $s\left(t_i; x_i, z_j\right)$, 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.

Let

$$H_j\left(s; \mathcal{F}_{i,t}\right) = \mathbb{P}\left(s_j^* < s \middle| \mathcal{F}_{i,t}, z_j, \eta_j\right) \tag{2.2}$$

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. In this notation, the probability that agent i can obtain a compatible organ j (i.e. $c_{ij} = 1$) is given by $H_j\left(s\left(t_i; x_i, z_j\right); \mathcal{F}_{i,t}\right)$. This representation relies on knowledge of the mechanism as a waitlist that uses scoring rules to order agents in the market.

In principle, these beliefs could depend on the history of offers previously received (and rejected) by the agent as well as information about the other agents on the waitlist at a given 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 belief about the distribution $H_j\left(s; \mathcal{F}_{i,t}\right)$ does not depend on $\mathcal{F}_{i,t}$:

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

$$\pi\left(t_i; z_j, \eta_j, x_i\right) = H\left(s_{ij}; z_j, \eta_j\right) \times \mathbb{P}\left(c_{ij} = 1 \mid z_j, x_i\right),$$

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

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 in mitigated by the large number of transplant hospitals and surgeons in the NYRT area. Second, it abstracts away from inference about the likelihood of receiving future offers based on the agent's past offers. This restriction is motivated by the institutional features and empirical observations discussed below. 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 proteins and blood-types that are known to be incompatible with the patient.

Our assumption on beliefs about s_j^* 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 (2.2) 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).

Descriptive Evidence

In addition to the tractability provided by Assumption 3, the institutional and empirical features of our setting suggest it provides a reasonable 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 of patients 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 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 autocorrelation across a range of partitions of donors (Table 2.11 in the Appendix).[19] Taken together, the evidence suggests that the set of patients competing for an organ is far from perfectly predictive of future cutoffs.

We also test whether recent offer history predicts current acceptance behavior, and fail to find evidence that it does. The logic of our test is as follows. 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 2.2. 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 2.5 presents coefficient estimates from several specifications that include

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

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 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 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.[20] 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 inactive time in the time since last offer variable, with little change in the estimated coefficient. Column (6) 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 (7) restricts the analysis to offers from ideal donors, and column (8) restricts to donors recovered in NYRT. In these two specifications, the time since last offer coefficient is statistically insignificant and actually becomes negative. Taken together, there is no evidence that patients adjust

[20]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 ommitted variables bias.

their acceptance behavior based on the recent offers they have experienced. We are therefore comfortable with the restrictions embedded in Assumption 3.

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 more 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) and by modeling beliefs as being based on past experience (Fershtman and Pakes, 2012). Despite these simplifications, x_i continues to be a high-dimensional object 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. Fortunately, these sources of dependence can be determined using the data. This will allow us to estimate our model tractably even though the state space is high-dimensional.

Value functions

We assume that agents make optimal accept/reject decisions by comparing the net present value of an object with the value of waiting. Holding the strategies of other agents fixed, she decides to remain on the list instead of accepting object j if $\Gamma_{ij}(t) = \Gamma(t; x_i, \alpha_i, z_j, \eta_j) + \varepsilon_{ijt}$ is less than $V_i(t) = V(t; x_i, \alpha_i)$. $V_i(t)$ is the value of continuing to wait conditional on the agent's current waiting time t and her observed and unobserved characteristics (x_i and α_i) that affect payoffs and information. To derive an expression for the value of waiting, consider an agent's value of waiting for another infinitesimal duration Δt given her current waiting time t. With a discount rate of ρ, the discrete approximation to the Hamilton-Jacobi-Bellman Equation defining the value of waiting at time t can be written as:

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

where $\lambda_i(t) = \lambda \int \pi_{ij}(t)\,\mathrm{d}F$ is the rate at which agent t expects to receive an offer, the operator $\mathbb{E}$ takes expectations over the idiosyncratic payoff shocks ε_{ijt} in equation (2.1), and, with a slight abuse of notation, $\pi_{ij}(t) = \pi(t; x_i, z_j, \eta_j)$. The leading fraction represents discounting due to time preferences. The first term is the flow payoff from remaining on dialysis during the Δt periods following t. The second term is the expected probability of departure during this period 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.[21] Taking the limit as $\Delta t \to 0$ under mild continuity conditions yields the differential equation

$$
(\rho + \delta_i(t))V_i(t) = d_i(t) + \delta_i(t)D_i(t) + \lambda \int \pi_{ij}(t)\,\mathbb{E}\max\{0, \Gamma_{ij}(t) - V_i(t)\}\,\mathrm{d}F + \dot{V}_i(t). \tag{2.3}
$$

This 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 as $t \to T_i$, the probability of receiving additional offers in the remaining time vanishes. This equation shows an intuitive result that values depend on the flow payoffs while waiting on the list, the

[21]The last term, which has not been fleshed out, 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.

possibility of and value from exogeneous departures, and the option value of potential offers.

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 a particular action in each state must be normalized (Magnac and Thesmar, 2002). However, Aguirregabiria and Suzuki (2014) and Kalouptsidi et al. (2015) point out that such normalizations may not be innocuous because answers to some counterfactuals can be sensitive to the normalization used. 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 under Assumptions 1(iii) and 2(i). Intuitively, the trade-offs between accepting an offer and waiting should only depend on payoffs relative to the value of not being assigned. This quantity is defined by the differential equation

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

and the terminal condition $O(T_i) = D_i(T_i)$. Under Assumptions 1(iii) and 2(i), this value of not being assigned does not depend on the waitlist. Therefore, measuring $V_i(t)$ and $\Gamma_{ij}(t)$ relative to $O_i(t)$ should sufficient to analyze decisions and welfare under the current and alternative mechanisms. Appendix 2.8.2 formally shows that this is indeed 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. It is straightforward to show that 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 (2.3) now simplifies to

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

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

$$V(t; x_i, \alpha_i) = \int_t^{T_i} \exp(-\rho(\tau - t)) \, p(\tau | t; x_i, \alpha_i) \left(\lambda \int \pi(\tau; x_i, z, \eta) \, \tilde{\psi}(\tau, x_i, \alpha_i, z, \eta) \, \mathrm{d}F_{z,\eta} \right) \mathrm{d}\tau, \tag{2.4}$$

where

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

is the probability that agent i departs the list prior to τ conditional on being on the list at t and

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

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 $\lim_{t \to T_i} V(t; x_i, \alpha_i) = O(T_i; x_i, \alpha_i) = 0$ because the probability of receiving an additional offer vanishes as $t \to T_i$.

Discussion

In addition to the normalization, the simple form of the value function and the agent's problem results from two key features of the model. First, the agent's decisions can be seen as an optimal stopping problem because agents are removed from the list once they obtain an assignment. That T_i is finite is inessential because an identical expression holds for a boundary condition of the form $\lim_{t \to \infty} V_i(t) = 0$ by setting T_i in equation (2.4) to infinity. Learning about the likelihood of receiving future offers or costs of considering an offer could also be incorporated within the framework of an optimal stopping problem, but would significantly complicate the model and subsequent analysis. We leave these extensions for future work.

Second, Assumptions 1-3 and the form of the mechanism together imply that there are no unforeseen state transitions. Beliefs and payoffs only depend on time-invariant characteristics (x_i, α_i) and a deterministically evolving state t. This may not be appropriate in some sequential assignment settings, including the allocation of deceased donor livers in which a patient's current health status also determines priority. It may possible to use techniques in Arcidiacono et al. (2016) to extend the model to incorporate stochastic changes in patient health status in a continuous time model as long as these transitions are observed discrete state jump processes. This modification would introduce incentives for waiting based on potential changes in unobserved health-status, further complicating the analysis.

2.4 Estimation

Equation (2.4) expresses the endogeneous value function $V(t; x_i, \alpha_i)$ as a solution to a fixed point problem. The two leading techniques for estimating dynamic choice models of this type are the conditional choice probabilities (CCP) approach (Aguirregabiria and Mira, 2007; Arcidiacono and Miller, 2011; Hotz and Miller, 1993) and the full-solution or nested fixed point approach (Miller, 1984; Pakes, 1986; Rust, 1987; Wolpin, 1984). 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 2.4.2.

2.4.1 A CCP Approach for Sequential Assignments

The primitives that we estimate are patient departure rates, $\delta(t; x_i)$; the distribution of donor types, $F_{z,\eta}$; and transplant values $\Gamma(t, x_i, z, \eta)$. These primitives are estimated in four steps, detailed below. First, we estimate departure rates offline 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 $\tilde{\psi}(\tau, x_i, \alpha_i, z, \eta)$ in equation (2.4). Third, we estimate the distribution of donor types, F, and offer probabilities, π, based on the set of donors and priority score cutoffs observed in our data. In the final step, we recover transplant values $\Gamma(t, x_i, z, \eta)$ by solving for each patient's value function at each date by evaluating equation (2.4).

There are two types of primitives that we do not estimate. First, our empirical specifications abstract away from patient-level unobserved heterogeneity α_i, and we omit this term from the notation going forward. We are currently working on including this source of heterogeneity. However, since our dataset contains rich patient-level information, including all covariates that influence priority on the waiting list and several other payoff relevant characteristics, estimates without patient-level unobserved heterogeneity are likely to be good approximations. Second, we set the discount rate ρ to a fixed value of 5% per year. As is well known, time preferences are not identified from observed choices alone in dynamic discrete choice models (Magnac and Thesmar, 2002). We are currently assessing robustness of our results to larger values of the discount rate.

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

$$\delta_i(t) = \delta_0(t) \exp(x_i \beta), \tag{2.5}$$

where $\delta_0(t)$ is a baseline hazard function and x_i are observed patient covariates. We include the same set of patient covariates x_i that we use in $\Gamma(\cdot)$. The Gompertz hazards model assumes that

$$\delta_0(t) = \delta_1 \exp(\delta_2 t)$$

which allows for the hazard rate to change with t. This parametric approach has the advantage of allowing for a simple expression for the survival function $p(\tau|t;x_i)$.

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

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

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, β_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

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

where the dependence of ψ on G has been suppressed for simplicity. Substituting $\psi(P_{ijt})$ for $\tilde{\psi}(\tau, x_i, \alpha_i, z, \eta)$ in equation (2.4), the value function can be re-written in terms of the CCPs as

$$V(t;x_i) = \int_t^{T_i} \exp(-\rho(\tau - t))\, p(\tau|t;x_i) \left(\lambda \int \pi_{ij}(\tau)\, \psi(P_{ijt})\, \mathrm{d}F\right) \mathrm{d}\tau. \tag{2.6}$$

Therefore, if P_{ijt} can be estimated, the only remaining unknowns in this equation are $\pi_{ij}(\tau)$ and F. We can then recover $\Gamma(x_i, z_j, \eta_j, t)$, the value of a transplant, without directly solving the integral equation (2.4).

In our application, we assume that $\varepsilon_{ijt} \sim N(0,1)$.[22] The variance of this term serves as our scale normalization. Further, we assume that $\Gamma(\cdot)$ is additively separable in η_j and approximate

$$V(x_i, t) - \Gamma(x_i, z_j, \eta_j, t) = \chi(x_i, z_j, t)\,\theta + \eta_j, \tag{2.7}$$

where $\chi(\cdot)$ is a flexible set of functions with interactions. We include dummies in x_i and z_i for categorical variables and piecewise linear splines for their continuous elements, as well as linear splines in t. The bases in these categorical variables and 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 decision, they solve a binary choice problem. We estimate the parameters (θ, σ_η) using a Gibbs' sampler (McCulloch and Rossi, 1994), 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)).[23]

Identification of these parameters is intuitive. The parameter θ is identified by based on 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

[22]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 $\psi(P) = \phi\left(\Phi^{-1}(P)\right) - (1-P)\,\Phi^{-1}(P)$.

[23]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 2.8.2.

have two kidneys offered to patients. The correlation between the number of offers made for 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 Hu and Schennach, 2008; Rao, 1992).

Step 3: Simulating the Mechanism

With an estimate of (θ, σ_η^2), our next objective is to use equation (2.6) to recover $\Gamma(x_i, z_j, \eta_j, t; \theta_0) = V(x_i, t; \theta_0) - \chi(x_i, z_j, t)\,\theta_0$. To do this, we only need to estimate the inner integral in equation (2.6),

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

because ρ is fixed and we have consistent estimates of $p(\tau; t, x_i)$, θ and λ. Expectations, in this expression, are taken over donor characteristics (z, η) drawn from F and the cutoffs s^*. The second equality is implied by the definition of $\pi_{ij}(t)$ given in Assumption 3.

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 our discussion in Section 2.3.3 that an agent receives an offer for object j if she is compatible $c_{ij} = 1$ and her priority score exceeds s_j^*. Therefore, we constuct the sample analog

$$\hat{W}\left(x_i, t; \hat{\theta}\right) = \frac{1}{J}\sum_{j=1}^{J} 1\left\{c_{ij} = 1\right\} 1\left\{s\left(t; x_i, z_j\right) > s_j^*\right\} \psi\left(\hat{P}_{ijt}\right) \tag{2.8}$$

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^*. Knowledge of the mechanism allows us to compare the patient's priority score, $s\left(t; x_i, z_j\right)$, with s_j^* directly. Further, our dataset contains rich information on donor proteins and patient immune system characteristics that allow us to accurately determine whether $c_{ij} = 1$.[24]

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 2.2) 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 under exactly the same circumstances 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 $\pi_{ij}\left(t\right)$. Second, our model of beliefs plays an important role in approximating the expectation in the expression above with the sample analog. Specifically, it relies on the assumption that the beliefs of agents do not depend on the history of observed offers. Therefore, the realized

[24]A patient must be both blood type and tissue type compatible for a transplant to take place (Danovitch, 2009). 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. However, before transplantation a crossmatch is conducted using blood from the donor and patient in case the virtual crossmatch yielded a false result. 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 instead of $1\left\{c_{ij} = 1\right\}$ in the expression above. Also, due to the possibility of false negatives in the "virtual crossmatch," we sometimes observe an acceptance that does not result in a transplant.

decisions of other agents are independent of i's presence and i's priority score. If our belief structure conditioned on a finer information set, then a simple sample analog would need to condition on the subset of histories where other agents' information were unchanged. This is because any changes in the choices of other agents could affect s_j^*.

Theorem 1 in the Appendix 2.8.2 shows that, for each x_i, 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 2.8.2. 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 $\Gamma\left(t, x_i, z_j, \eta_j\right)$. First, we estimate $V\left(t; x_i\right)$ for every time t in the dataset when patient i received an offer. We substitute the sample analog for $\pi_{ij}\left(t\right)$ in equation (2.8) into equation (2.6). Then, we use the estimated departures model for $p\left(\tau | t; x_i\right)$, and the observed donor arrival rate to estimate λ. Finally, the integral with respect to τ is calculated numerically by evaluating the integrand at a large number of points. Details of this procedure for computing $\hat{V}\left(t; x_i\right)$ are provided in Appendix 2.8.2. Once $\hat{V}\left(t; x_i\right)$ is calculated, we recover $\Gamma\left(\cdot\right)$ by inverting G :

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

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

2.4.2 Discussion

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 in terms of θ_Γ, recover the value function by solving the fixed point in equation (2.4), 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 this approach is to avoid directly parametrizing the (endogenous) difference in equation (2.7), $\Gamma(\cdot) - V(\cdot)$, 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 Arcidiacono and Miller, 2011; Hotz and Miller, 1993). 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 (2.8) 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 (2.4) 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 the optimization routine.[25]

[25] 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

The CCP approach avoids these complications and allows us to respect the details of the mechanism. The main cost is some loss in efficiency because we directly estimate the difference in equation (2.7), 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.

Unobserved Heterogeneity

Our current empirical specifications omit the patient unobserved heterogeneity terms α_i. The main complication in relaxing this assumption is that unobserved heterogeneity may affect both departure rates $\delta_i(t)$ as well as the conditional choice probabilities. Such dependence would require us to estimate the hazard rates and choice probabilities simultaneously. An approach by Arcidiacono and Miller (2011) could be adapted to do so in the case when α_i takes on discrete values. One may further argue that unobserved heterogeneity may be time varying if unobserved patient health is stochastic. Recent work by Connault (2016) provides a path forward in this case under certain assumptions. We believe that abstracting away from unobserved heterogeneity still yields useful results because our dataset contains a very rich set of patient characteristics.

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 2-1. The observable characteristics included in the model do not explain all of this sharp decline or the composition of offers, especially

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.

at lower positions on the waitlist. We will, however, compare model fit with and without such unobserved heterogeneity in the next section.

2.5 Parameter Estimates

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

2.5.1 Departure Rates

Table 2.6 presents estimates from hazard models of departures from the kidney waitlist prior to transplantation. We estimated several specifications that involve different parametric assumptions and sets of controls. Across specifications, we estimate an increasing baseline hazard of departure, consistent with patients progressively becoming less healthy over time. Column (1) presents a Gompertz model with no covariates. The estimated value of $\log(\delta_1)$ is -8.351 per day, and $\delta_2 = 3.6 \times 10^{-4}$. Column (2) presents a Weibull hazards model in which $\delta_0(t) = \delta t^{\delta-1}$, showing qualitatively similar patterns. In column (3), we examine how patient covariates correlate with departure rates when added to the baseline Gompertz model. Adding covariates does not change the pattern of increasing baseline departure rates, but the estimates reveal significant heterogeneity across patients. 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 live 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 (4) and (5) 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 non-parametric Cox proportional hazards model. Figure 2-7 in the appendix compares the estimated survival curves from columns (3) and (5) and shows that the Gompertz hazards model yields a survival curve that is very similar to the non-parametric Cox proportional hazards model. We therefore feel comfortable using the model in column (3) for our analysis.

2.5.2 Estimated CCPs

We estimated three specifications for the conditional choice probability of accepting an offer. Our choice sample consists of all offers made to NYRT patients between 2010 and 2013, including those screened out by pre-set criteria set by a patient. All specifications include the rich set of patient and donor observed characteristics summarized in tables 2.1 and 2.2. 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.[26] We also include any covariates that were part of the LYFT

[26]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

model of 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 with whether the donor was recovered in NYRT. In addition, we include patient-donor specific variables that capture match-specific heterogeneity, e.g. 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 2.7 presents select parameter estimates from the three specifications. Table 2.12 in the appendix shows the full set of estimated coefficients. 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. 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.

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 stabilizing after year three. We will interpret this time path later when we discuss the value function and values of transplantation.

Figure 2-3 describes the fit of these models. The first panel mimics the fit of

the procedure to request these modules.

acceptance rates by position described in figure 2-1, 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 first specification does not even accurately capture the average acceptance rate across the first 20 positions. Instead, it implies a steady decline as the composition of donor observables changes moving down the list.[27] The second panel of figure 2-3 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.

2.5.3 Estimated Value Functions V and NPV of Transplantation Γ

Table 2.8 presents estimates of the projection of V and Γ on $\chi(\cdot)$ for our preferred specification, which includes both donor unobserved heterogeneity and time waited. The projection coefficients for both Γ and V are intuitive. For example, younger donors are more valuable, and donors with antigen mismatches are less valuable. Similarly, donors that were recovered outside NYRT are less desirable. We also see that the value of a transplant decreases with waiting time. This rate of decline grows rapidly during the first few years, and then stabilizes in a steady decline after year three. One reason for this decline is that patient health deteriorates on dialysis, making an eventual transplant less useful. The table also shows that there is

[27] It appears that previous approaches, including the KPSAM model, have arbitrarily truncated the offer sequence for each donor by dropping all offers below a specific position on the list.

significant donor-level unobserved heterogeneity, equivalent to a standard deviation of 1.25.

The units in these estimates are in terms of standard deviations of ε_{ijt} in equation (2.1). Therefore, they are indicative of probabilities of accepting or rejecting an offer with given observed characteristics, but are not directly comparable across patients. Sequential assignment mechanisms can re-assign offers from some patients to others. Motivated by this fact, we convert the estimated values of the payoffs and value functions into a measure of equivalent number of offers. Specifically, for each patient, we calculate a multiplier that equals the marginal value of a one-time offer made to an agent at the time of registration:

$$\xi_i = \int \mathbb{P}\left(c_{ij} = 1 | z_j, x_i\right) \mathbb{E} \max\left\{0, \Gamma\left(0, x_i, z_j, \eta_j\right) + \varepsilon_{ij} - V\left(0; x_i\right)\right\} \mathrm{d}F.$$

We then transform the payoffs $\Gamma_{ij}(t)$ and $V_i(t)$ by dividing them by ξ_i.

We compute ξ_i for each patient by randomly sampling from the set of donors that were procured within NYRT. This is the set of donors for whom the mechanism could have been modified by NYRT's organ procurement organization through an application to the OPTN. Moreover, donors from NYRT are usually preferable to donors from outside NYRT, and are therefore more representative of an average donor whose organs a patient would seriously consider accepting.

This transformation to Equivalent Offers (EO) treats the marginal value of an average offer at the time of registration, given the current mechanism, equally for all agents. Therefore, EO expresses all payoffs and values as a multiple of the marginal value of an average offer at the time of registration. Differences in Γ across two kidneys or patients or differences in V across two mechanisms in EO units is a multiple of this marginal value. For small changes, the units approximate the number of average offers an agent is willing to forego in a trade-off between two kidneys or two mechanisms. This feature makes EO 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 2-4a and 2-4b present a more intuitive description of our estimates for Γ in terms of EO. The plots show how the value of transplant varies across specific patient and donor characteristics, holding all remaining characteristics fixed. Figure 2-4a shows that patients across all age groups prefer younger donors. However, the relative value of a young donor decreases with patient age. A patient under the age of 55 is willing to forego over 0.5 equivalent offers in order to obtain a kidney from a 30 year old donor as opposed to a 60 year old donor. This number is less than 0.4 for a patient over the age of 65. This result is intuitive because older patients typically have a lower remaining life expectancy after any transplant. Therefore, they place less value on receiving an organ that is likely to function for a very long period of time.

Similarly, Figure 2-4b shows that a kidney with a perfect tissue type match is especially valuable to patients. Patients are willing to forego between five and ten equivalent offers, depending on donor age, in order to obtain a kidney with no tissue-type mismatches. 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.

2.6 Evaluating Design Trade-Offs

This section begins by outlining an equilibrium concept for counterfactual analysis in which agents play optimal type-symmetric strategies and have consistent beliefs.

This concept leads to a tractable procedure to compute equilibrium outcomes under counterfactual allocation rules. We describe this procedure, and then present results from alternative mechanisms.

2.6.1 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 $z \in \zeta$. For computational reasons, we will treat $\mathcal{X}$ and ζ as finite sets, but the definition is agnostic regarding their cardinalities.

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. Consistent with our approach during estimation, these beliefs will be based on an equilibrium CDF of the priority score cutoff for each object type $z \in \zeta$, denoted $H_z : \mathbb{R} \to [0, 1]$.

We model the composition of the queue using a single steady state composition. A distribution over queue compositions would be extremely high dimensional, making the equilibrium concept intractable. Specifically, the queue composition will be governed by a probability density function, m, defined on the set $\mathcal{X} \times [0, T]$.[28] 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 on $x \times t$. The length of the queue is denoted by N.

[28] 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.

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 strategy σ_x^* after t.

2. Consistent beliefs: For each (t, x, z),

$$\pi^* (t; x, z) = H_z^* (s_{xz} (t)) \times \mathbb{P} (c_{ij} = 1 | \, x, z),$$

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^*.

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

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

   $$\dot{m}_x (t) = -m_x (t) \, \kappa_x (t) \text{ and } m_x (0) \propto \gamma_x,$$

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

   (b) N^* is the smallest value of $N \in \mathbb{N}$ such that $\sum_x \gamma_x \leq N \sum_{x,t} m_x^* (t) \, \kappa_x (t)$

The first condition assumes 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-Bellman-Jacobi equation defined in section 2.3.3. The second condition imposes that agents have correct beliefs for a large waitlist. It writes agent beliefs about future offers as the product of the steady state distribution of cutoffs and exogenous compatibility realizations. The cutoff distribution H_z^* is the distribution of cutoffs that arise when agents use strategies σ^* and N^* agents are drawn from a distribution governed by m^*.[29] 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 i 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. 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) determines the equilibrium queue length. The term $\sum_x \gamma_x$ is the total (exogenous) arrival rate of patients in the queue. This term must not be larger than the total equilibrium departure rate. The total equilibrium depature rate is the product which is the of the queue length and the departure rate for the average patient, $\sum_{x,t} m_x(t) \kappa_x(t)$. The condition imposes the constraint that N^* is an integer.

2.6.2 Computing Equilibria

We compute equilibria using an algorithm that iteratives between computing the value function and optimal deicisions, and the steady-state of the waitlist. A de-

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

tailed description with expressions for each step of the procedure and pseudocode is provided in Appendix 2.8.3. The following discussion provides a simplified description of the key steps.

In addition to the primitives, the algorithm uses a discrete time grid $t = t_0, \ldots, t_l, t_{l+1}, \ldots, T$, an 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 $V_x^k(t)$, given beliefs π^{k-1}, via backwards induction:

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

   The inner integral in the above expression is approximated by sampling a subset of donors. This calculation also yields patient departure rates $\kappa_x^k(t)$.

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

$$m_x^k(t_l) \propto \gamma_x \int_0^t \exp\left(-\int_\tau^t \kappa_x^k(\tau')\,\mathrm{d}\tau'\right) \mathrm{d}\tau.$$

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 $\sigma_x^k(\Gamma, t) = 1\left\{\Gamma \geq V_x^k(t)\right\}$.

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 right 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.

To keep the computational burden manageable, the results we present below are based on a type space given by a random sample of 500 patients and donors drawn from our dataset. Further, we discretize time into quarter-years.

2.6.3 Results on Alternative Mechanisms

As mentioned earlier, 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 three mechanisms aimed at achieving each of these goals. We also compare the mechanisms to the two mechanisms used prior to and following the 2014 re-design.

Following the design parameters considered by the kidney allocation committee, our exercises focus on sequential offer mechanisms in which patients and surgeons may refuse an offer.[30] One justification for this restriction is to respect patient and doctor discretion. Additionally, we restrict attention to mechanisms that use a scoring rule $s_{ij}(t)$ to order patients waiting for a transplant and break ties uniformly at random among patients with the same score. This is the class of mechanisms for which the algorithm described above can be used to compute an equilibrium.

We consider the following five mechanisms:

1. **Pre-2014:** This mechanism was used during our sample period and will be

[30]Our examination of the meetings of the kidney allocation committee indicate that they did not consider alternatives in which doctors were mandated to accept particular organs, or in which rejections were in any way penalized.

the point of reference for comparing other mechanisms.

2. **Post-2014:** In December 2014, the kidney allocation mechanism switched to a system that awards greater priority to patients that are extremely difficult to match. This change was motivated by the idea that these patients have few opportunities for transplantation and are likely to accept most organs. Therefore, giving them additional priority could reduce overall organ waste in addition to achiving equitable outcomes for sensitized patients. Additionally, patients with the top 20% of predicted post-transplant survival probabilities are given priority for organs in the lowest quintile of graft failure risk. The primary motivation is to offer high-quality kidneys to patients that are likely to benefit from them most.

3. **Random Order:** One concern of the kidney committee has been to maintain a transparent and procedurally fair offer system. A natural candidate for such a system is one that randomizes the order of patients each time a donor arrives. In this mechanism, we set $s_{ij}(t) = 0$ for all i, j, and t. When the waiting list is long, this system also has the potential advantage of incentivizing patients to accept more offers because they may be at the bottom of the queue when future organs arrive. Another natural candidate, which we are working on, resembles a first-come first-served system that only prioritizes patients according to how long they have waited.

4. **Hard-to-match:** Partly due to fairness considerations, the reforms prioritized patients that were very hard to match. An extreme version of this is to adopt a system in which priority is solely determined by immune sensitization. This system is likely to reduce waste because highly sensitized patients should accept

offers even with relatively high priority. These patients also have fewer options for receiving a living donor transplant. In this mechanism, we divide patients into 10 equally sized bins based on their CPRA, setting $s_{ij}(t)$ to 10 for those in the highest bin and 1 for those in the lowest bin.

5. **Match Value:** The final mechanism attempts to maximize efficiency by using a greedy waitlist procedure. For each donor, it divides patients into 10 equally sized bins based on our estimated $\Gamma_{ij}(t)$. $s_{ij}(t)$ is set to 10 for those in the highest bin and to 1 for those in the lowest bin.

Figure 2-5 and table 2.9 show our results. The figure shows that the 2014 reforms are predicted to do little in terms of improving patient welfare. However, it is predicted to slightly decrease kidney discard rates. The directions of these effects are theoretically ambiguous, since some patients received higher priority while others receive lower priority post-reform. In contrast, random order and mechanisms that prioritize patients that are hard to match substantially decrease kidney discard rates but at signifcant costs to patient welfare. The mechanism that prioritizes patients based on match value performs better than either the pre- or the post-2014 mechanism in terms of patient welfare and discard rates.

Table 2.9 describes the allocations in greater detail. Panel A shows the effects for all patients while panels B and C show the effects for adults younger than and older than age 55, respectively. The table shows that the change in overall value is correlated with the observed characteristics of donors allocated through the mechanism. For instance, Match Value results in the allocation of donors that are similar on the primary observed dimensions as the post-2014 priorities. However, panels B and C shows that the mechanism allocates more desirable donors as measured by age, head trauma, or hypertension to younger patients. The pre- and post-2014 mechanisms are very similar to each other in these dimensions. On the other hand, random order and priorities for hard-to-match patients are less likely find the highest-benefit patient for each donor, and lead to lower patient welfare.

2.7 Conclusion

Although the deceased donor kidney allocation system was reformed in 2014, there is much scope for improvement. The previous design process was assisted by a simple simulation model that did not account for the dependence of accept/reject rules on agents' incentives. This paper develops an empirically tractable method for estimating the value of various assignments in dynamic assignment systems and shows how to compute equilibria of counterfactual mechanisms. Our preliminary results show that there exist mechanisms that both improve match quality and reduce organ waste. Both the pre- and post-2014 reforms are worse on these dimensions.

These techniques and insights are more broadly applicable. Many resources, in addition to deceased donor organs of all types, are rationed via sequential offer mechanisms or waitlists. Previous theoretical approaches have not provided sharp guidance on how to organize these waitlists. The scope for empirical work on improving these systems is large and unexplored.

In ongoing work, we are exploring the comparison these mechanisms to a socially optimal assignment and finding mechanisms motivated by the resulting allocations. We are also working on comparing the predictions from our optimal stopping based approach to naive calculations in which the decision rules are not sensitive to the waitlist mechanism. Finally, we are working on comparing the predicted transplants from our approach to comparing the pre-2014 mechanism to the post-2014 mechanism with the observed outcomes in recent years.

Table 2.1: Patients 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 2.2: 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 2.3: Rates of Receiving and Accepting Offers

	Number of Patients	Offers to NYRT Patients: Offer & Acceptance Rates — All Donors		NYRT Donors		Perfect Tissue Type Match		Life Years From Transplantation (LYFT) — All Donors	
		Annual Rate	% Accepted	Annual Rate	% Accepted	Annual Rate	% Accepted	Mean Offered	Mean Accepted
Panel A: All Offers									
All	9917	218.4	0.15%	40.1	0.73%	0.094	10.6%	5.30	6.16
Peak CPRA < 0.80	8864	234.5	0.10%	42.8	0.49%	0.093	9.3%	5.29	6.07
Peak CPRA >= 0.80	1053	83.1	1.03%	17.3	4.39%	0.103	15.8%	5.60	6.31
Panel B: Offers that Met Screening Criteria									
All	9917	104.1	0.34%	23.5	1.35%	0.051	21.8%	4.89	6.16
Peak CPRA < 0.80	8864	112.3	0.22%	25.1	0.89%	0.049	19.7%	4.86	6.07
Peak CPRA >= 0.80	1053	35.8	2.51%	9.8	8.11%	0.068	29.0%	5.38	6.31
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%	7.32	8.32
Peak CPRA < 0.80	8864	0.8	15.79%	0.8	16.61%	0.019	51.1%	7.57	8.89
Peak CPRA >= 0.80	1053	0.9	44.10%	0.9	46.24%	0.038	43.6%	6.83	7.93

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 aggregated across patients. "Offered" and "Accepted" columns report average characteristics across offers or accepted offers. Peak CPRA is the highest level of Calculated Panel Reactive Antibodies (CPRA) recorded for each patient during the period of study.

Table 2.4: Evidence on Mismatch

	Match Characteristics				
	Patient Age 0-17	Patient Age 18-34	Patient Age 35-49	Patient Age 50-64	Patient Age 65+
Share of Patient Population	2.2%	11.1%	27.2%	41.5%	18.0%
Share of Deceased Donor Transplants	4.7%	9.6%	26.1%	41.9%	17.7%
Share of Ideal Donor Transplants	6.7%	9.8%	26.6%	39.8%	17.1%
	Panel A: Outcomes of all Patients				
Received Deceased Donor Transplant	47.9%	19.2%	21.3%	22.4%	21.9%
Received Live Donor Transplant	19.4%	18.6%	11.5%	8.9%	6.5%
Still Waiting	28.9%	48.6%	50.4%	48.8%	44.0%
Died or Too Sick to Transplant	1.9%	4.1%	6.8%	12.9%	20.0%
Departed for Another Reason	1.9%	9.5%	10.1%	7.0%	7.6%
	Panel B: Donor Age and Quality				
Donor is Ideal	97.0%	70.2%	69.7%	64.9%	66.2%
Donor Age 0-17	23.8%	14.6%	7.2%	4.3%	4.5%
Donor Age 18-35	73.3%	37.6%	26.7%	17.3%	12.9%
Donor Age 35-49	3.0%	30.2%	34.9%	28.7%	17.7%
Donor Age 50-64	0.0%	16.6%	29.2%	42.1%	52.0%
Donor Age >= 65	0.0%	1.0%	2.0%	7.6%	12.9%
	Panel C: Donor Age, Ideal Donors				
Donor Age 0-17	23.5%	17.4%	9.5%	5.8%	5.6%
Donor Age 18-35	73.5%	38.9%	23.7%	17.2%	11.2%
Donor Age 35-49	3.1%	26.4%	33.9%	22.5%	14.7%
Donor Age 50-64	0.0%	16.0%	30.6%	45.1%	51.8%
Donor Age >= 65	0.0%	1.4%	2.3%	9.4%	16.7%

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; had a non-cardiac death; has creatinine below 3; and is Hepatitis C negative.

Table 2.5: Acceptance Rate by Past Offer Rates

	Dependent Variable: Current Offer Accepted							
	All Offers						Ideal Donors	NYRT Donors
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
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)				
Time Since Last Offer, Including Inactive Periods					0.00421			
					(0.00514)			
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	X
Patient Characteristics		X	X	X	X	X	X	X
Donor and Match Characteristics		X	X	X	X	X	X	X
Listing Center Fixed Effects		X	X	X	X	X	X	X
Observations	2793098	2831262	2831262	2821641	2793098	2831262	1083686	526233
R-squared	0.098	0.097	0.097	0.099	0.098	0.099	0.133	0.116
Mean Acceptance Rate	0.15%	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	0.005	0.012	0.023
S.D. Time Since Last N Offers	0.016	0.016	0.011	0.008	0.019	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) - (8) control for current patient, donor, and match characteristics. Column (6) includes controls for donor characteristics of the patient's previous offer. Column (7) restricts to offers from ideal donors, and Column (8) restricts to NYRT donors. Controls are as described in the notes for Table 5. An ideal donor has no history of diabetes; had a non-cadiac death; has creatinine below 3; and is Hepatitis C negative.

Table 2.6: Survival Model Estimates

Patient Survival Model Coefficient Estimates					
	No Controls		With Controls		
	Gompertz (1)	Weibull (2)	Gompertz (3)	Weibull (4)	Cox (5)
Diabetic Patient			0.294 (0.0386)	0.290 (0.0385)	0.300 (0.0384)
Bloodtype A Patient			0.160 (0.0523)	0.157 (0.0523)	0.165 (0.0524)
Bloodtype O Patient			0.0976 (0.0457)	0.0872 (0.0457)	0.0889 (0.0457)
Zero cPRA			0.0933 (0.0507)	0.102 (0.0508)	0.103 (0.0507)
cPRA > 80			-0.260 (0.0775)	-0.214 (0.0769)	-0.212 (0.0770)
Age (at Registration)			-0.00476 (0.0255)	-0.0103 (0.0255)	-0.0145 (0.0255)
Age - 18 if Age>=18			0.00966 (0.0298)	0.0135 (0.0296)	0.0193 (0.0296)
Age - 35 if Age>=35			-0.00252 (0.0126)	-0.00376 (0.0126)	-0.00649 (0.0125)
Age - 50 if Age>=50			0.0193 (0.00865)	0.0219 (0.00864)	0.0246 (0.00864)
Age - 65 if Age>=65			0.0320 (0.0102)	0.0327 (0.0102)	0.0331 (0.0102)
Prior Transplant			-0.0216 (0.0641)	0.0308 (0.0636)	0.0463 (0.0637)
Body Mass Index (BMI)			-0.0109 (0.00757)	-0.0120 (0.00757)	
Missing BMI			0.339 (0.230)	0.275 (0.230)	
BMI >= 18.5			0.0888 (0.138)	0.0521 (0.138)	-0.185 (0.102)
BMI >= 18.5			-0.0587 (0.0575)	-0.0521 (0.0575)	-0.112 (0.0445)
BMI >= 30			0.0383 (0.0698)	0.0372 (0.0697)	-0.0538 (0.0453)
Total Serum Albumin			-0.242 (0.0650)	-0.248 (0.0650)	
Missing Total Serum Albumin			-0.849 (0.222)	-0.791 (0.222)	
Total Serum Albumin >= 3.7			-0.0623 (0.0696)	-0.0543 (0.0696)	-0.249 (0.0420)
Total Serum Albumin >= 4.4			0.121 (0.0602)	0.121 (0.0602)	-0.0240 (0.0456)
On Dialysis at Listing			-0.813 (0.0767)	-0.859 (0.0768)	-0.868 (0.0766)
Log(Days on Dialysis at Lising)			0.124 (0.0110)	0.126 (0.0111)	0.124 (0.0110)
Constant	-8.322 (0.0285)	-10.71 (0.159)	-7.546 (0.496)	-9.984 (0.517)	
Constant (delta(2))	0.000347 (0.0000188)		0.000418 (0.0000206)		
Constant (delta)		0.313 (0.0153)		0.350 (0.0157)	
Observations	9917	9917	9917	9917	9917

Table 2.7: Conditional Choice Probability Estimates (select co-efficients)

	Conditional Choice Probability of Accepting an Offer					
	Base Specification		Unobserved Heterog.		Waiting Time + UH	
	(1)		(2)		(3)	
Calculated Panel Reactive Antibody (CPRA)	0.32	(0.07)	0.24	(0.09)	0.09	(0.13)
Log Years on Dialysis at Registration	-0.08	(0.01)	-0.09	(0.01)	-0.09	(0.01)
Donor Age < 18	0.46	(0.13)	0.06	(0.21)	-0.04	(0.20)
Donor Age 18-35	0.76	(0.17)	0.09	(0.24)	-0.09	(0.27)
Donor Age 50+	-1.13	(0.35)	-0.64	(0.43)	-0.58	(0.47)
Expanded Criteria Donor (ECD)	-0.17	(0.02)	-0.50	(0.07)	-0.53	(0.08)
Donation from Cardiac Death (DCD)	-0.12	(0.03)	-0.44	(0.08)	-0.52	(0.07)
Perfect Tissue Type Match	2.45	(0.32)	2.96	(0.44)	2.95	(0.46)
2 A Mismatches	-0.02	(0.02)	-0.02	(0.02)	-0.03	(0.02)
2 B Mismatches	0.01	(0.02)	-0.03	(0.03)	-0.03	(0.03)
2 DR Mismatches	-0.09	(0.02)	-0.10	(0.02)	-0.10	(0.02)
Regional Offer	-1.00	(0.06)	-2.22	(0.17)	-2.42	(0.19)
National Offer	-1.15	(0.05)	-2.43	(0.11)	-2.64	(0.12)
Non-NYRT Donor, NYRT Match Run	0.94	(0.02)	1.78	(0.05)	1.94	(0.06)
Log Waiting Time (years)					-0.18	(0.06)
Log Waiting Time * Over 1 Year					-0.02	(0.07)
Log Waiting Time * Over 2 Years					-0.24	(0.13)
Log Waiting Time * Over 3 Years					0.17	(0.12)
NYRT Donor * Donor Age < 18	-0.01	(0.07)	0.35	(0.20)	0.37	(0.22)
NYRT Donor * Donor Age 18-35	0.12	(0.05)	0.28	(0.13)	0.27	(0.17)
NYRT Donor * Donor Age 50+	-0.26	(0.04)	-0.38	(0.11)	-0.45	(0.12)
Patient Age * Donor Age < 18	-0.01	(0.00)	-0.01	(0.00)	0.00	(0.00)
Patient Age * Donor Age 18-35	-0.02	(0.01)	0.00	(0.01)	0.01	(0.01)
Patient Age * Donor Age 50+	0.03	(0.01)	0.01	(0.01)	0.01	(0.01)
Patient Age - 35 if Age >= 35 * Donor Age 18-35	0.02	(0.01)	0.00	(0.01)	-0.01	(0.01)
Patient Age - 35 if Age >= 35 * Donor Age 50+	-0.01	(0.01)	0.01	(0.01)	0.01	(0.01)
Donor Unobservable Std. Dev.			1.03	(0.23)	1.25	(0.24)
Idiosyncratic Shock Std. Dev.	1.00		1.00		1.00	
Acceptance Rate	0.150%		0.150%		0.150%	
Number of Offers	2840937		2840937		2840937	

Table 2.8: Estimated Value of Waiting and Transplantation (select co-efficients)

	Value of a Transplant	Patient Value Function
Calculated Panel Reactive Antibody (CPRA)	4.02	0.74
Log Years on Dialysis at Registration	-0.30	-0.41
Donor Age < 18	0.49	
Donor Age 18-35	0.83	
Donor Age 50+	-1.35	
Expanded Criteria Donor (ECD)	-0.53	
Donation from Cardiac Death (DCD)	-0.51	
Perfect Tissue Type Match	2.73	
2 A Mismatches	0.02	
2 B Mismatches	0.05	
2 DR Mismatches	-0.05	
Regional Offer	-2.80	
National Offer	-3.02	
Non-NYRT Donor, NYRT Match Run	2.03	
Patient Age	0.18	0.22
Patient Age - 18 if Age >= 18	-0.29	-0.22
Patient Age - 35 if Age >= 35	0.15	0.04
Patient Age - 50 if Age >= 50	-0.06	-0.01
Patient Age - 65 if Age >= 65	-0.09	-0.21
Log Waiting Time (years)	-0.33	-0.74
Log Waiting Time * Over 1 Year	0.83	-0.45
Log Waiting Time * Over 2 Years	-0.01	-0.25
Log Waiting Time * Over 3 Years	4.54	0.01
NYRT Donor * Donor Age < 18	0.17	
NYRT Donor * Donor Age 18-35	0.18	
NYRT Donor * Donor Age 50+	-0.66	
Patient Age * Donor Age < 18	-0.01	
Patient Age * Donor Age 18-35	-0.02	
Patient Age * Donor Age 50+	0.03	
Patient Age - 35 if Age >= 35 * Donor Age 18-35	0.01	
Patient Age - 35 if Age >= 35 * Donor Age 50+	-0.02	
S.D. Idiosyncratic Shock	1.00	
S.D. Donor Unobserved Heterogeneity	1.25	
S.D.	2.62	1.13
S.D. Within Donor	1.63	
S.D. Between Donors	2.05	

Notes: coefficients and standard deviations are reported in reported in Standard Deviations of the Idiosyncratic Shock. Specification includes waiting time predictors and donor unobserved heterogeneity.

Table 2.9: Outcomes from Various Mechanisms

Mechanism	Steady State Queue Size	Reduction in Fraction Discarded	Average Life Years From Transplanta tion						
Pre-2014 Priorities	5953	0.0%	5.13	0.83	0.90	-0.07	1.82	0.11	0.56
Post-2014 Priorities	5984	-0.1%	5.17	0.55	0.61	-0.07	1.82	0.11	0.56
Random Order	4154	8.6%	7.57	-0.45	-0.25	-0.20	2.66	0.23	0.48
Hard-to-match	4461	6.9%	7.60	-0.43	-0.17	-0.25	2.47	0.20	0.49
Match Value	4898	5.2%	5.64	1.07	1.20	-0.13	2.33	0.17	0.51
Optimal	4911	3.2%	7.39	0.19	0.39	-0.19	2.14	0.16	0.53
Optimal (no fit)	4375	6.9%	7.83	-0.03	0.13	-0.16	2.49	0.21	0.49
Pre-2014 Priorities [Naive]	7032	-7.4%	6.24	0.36	0.07	0.28	1.14	0.06	0.64
Post-2014 Priorities [Naive]	7044	-7.4%	6.16	0.22	-0.06	0.29	1.14	0.06	0.64
Random Order [Naive]	7154	-8.6%	8.70	0.45	0.11	0.34	1.03	0.05	0.65
Hard-to-match [Naive]	6743	-8.3%	10.93	-0.03	-0.37	0.34	1.06	0.06	0.65
Match Value [Naive]	7363	-9.1%	5.69	2.18	1.82	0.36	0.98	0.05	0.65
Optimal [Naive]	6767	-7.8%	8.71	0.59	0.30	0.30	1.10	0.06	0.64
Optimal (no fit) [Naive]	6780	-8.1%	9.24	0.72	0.44	0.28	1.08	0.06	0.64
Pre-2014 Priorities	5953	0.0%	17.88	2.85	2.80	0.05	15.77	0.97	0.56
Post-2014 Priorities	5984	-0.1%	18.86	2.64	2.60	0.04	16.58	1.01	0.56
Random Order	4154	8.6%	14.76	0.54	0.53	0.01	6.42	0.56	0.48
Hard-to-match	4461	6.9%	14.77	0.48	0.48	0.00	5.76	0.47	0.49
Match Value	4898	5.2%	15.43	2.62	2.57	0.04	6.62	0.49	0.51
Optimal	4911	3.2%	16.87	1.91	1.88	0.04	13.25	0.99	0.53
Optimal (no fit)	4375	6.9%	16.85	2.19	2.15	0.04	8.37	0.70	0.49
Pre-2014 Priorities [Naive]	7032	0.0%	16.32	3.96	3.90	0.07	6.11	0.32	0.64
Post-2014 Priorities [Naive]	7044	0.0%	16.28	3.94	3.89	0.05	4.76	0.25	0.64

Figure 2-1: Acceptance Rate by Position

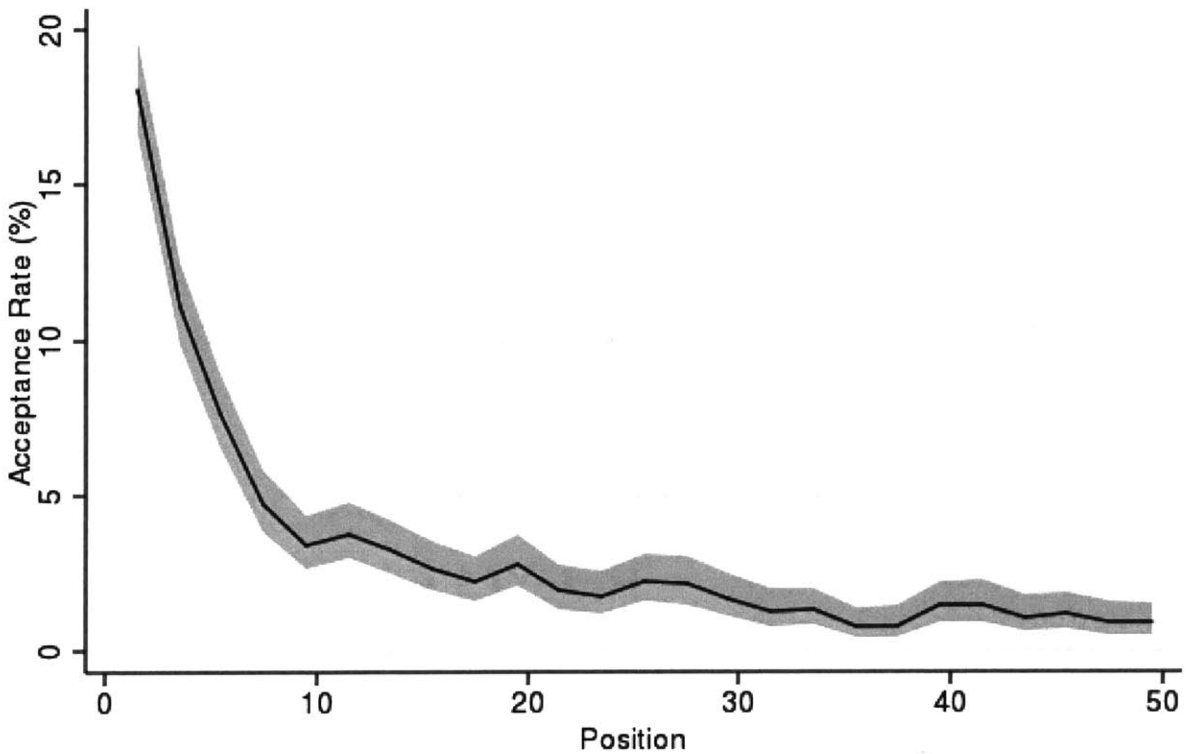

Note: Sample is offers made to NYRT patients between 2010 and 2013, excluding offers that did not meet a patient's pre-set donor screening criteria. Positive cross-matches are counted as acceptances. The shaded region represents pointwise 95% confidence intervals around the mean.

Figure 2-2: Waiting Time by Position

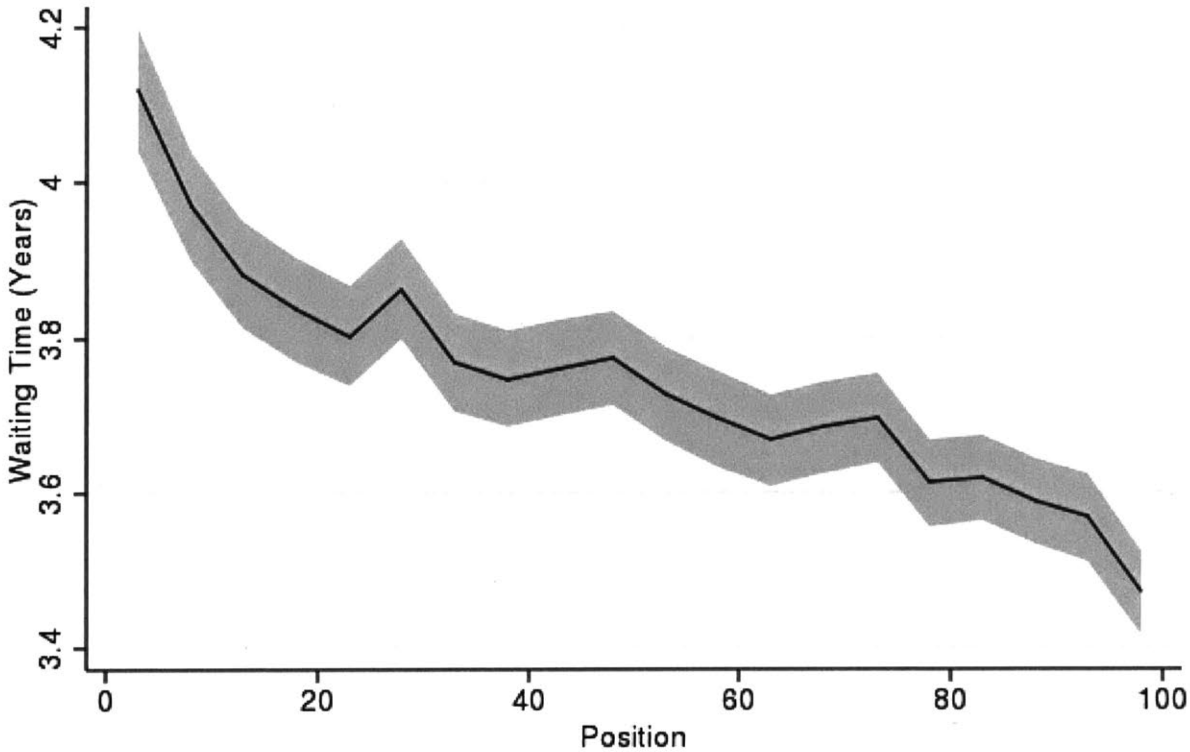

Note: Sample is offers made to NYRT patients between 2010 and 2013, excluding offers that did not meet a patient's pre-set donor screening criteria. The black line plots the mean waiting time among offered patients in each position group.

Figure 2-3: Model Fit

(a) Fit by position

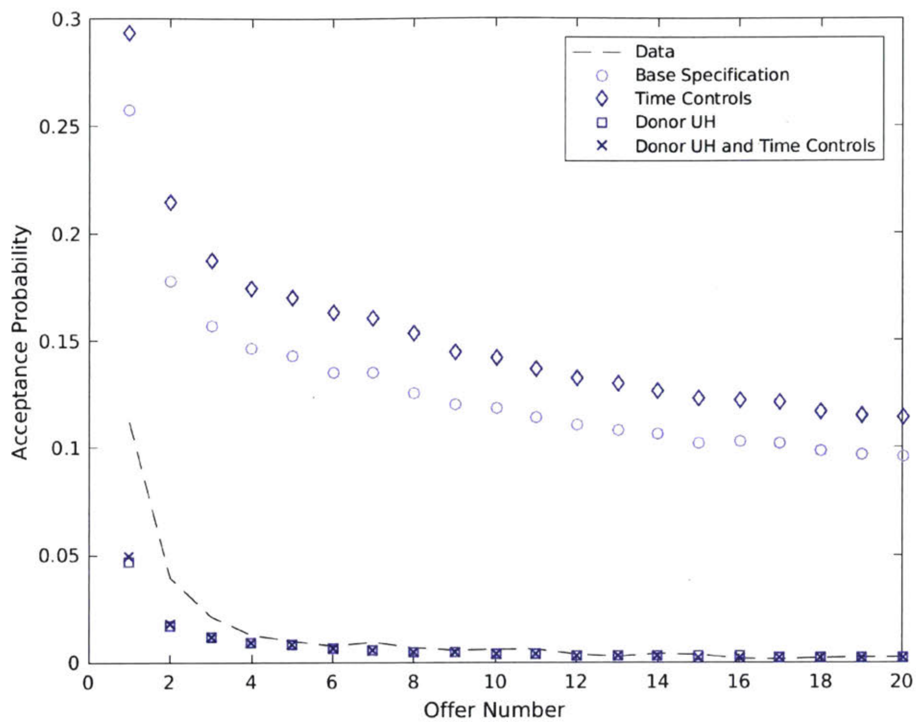

(b) Fit by time waited

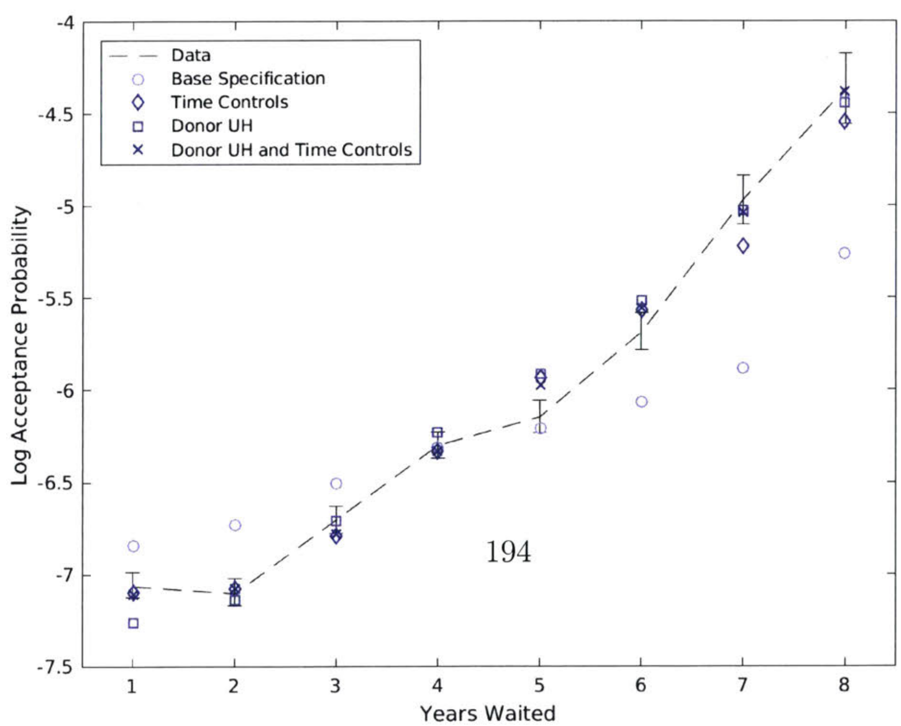

Figure 2-4: Value of Transplant by Donor Age

(a) Patient Age

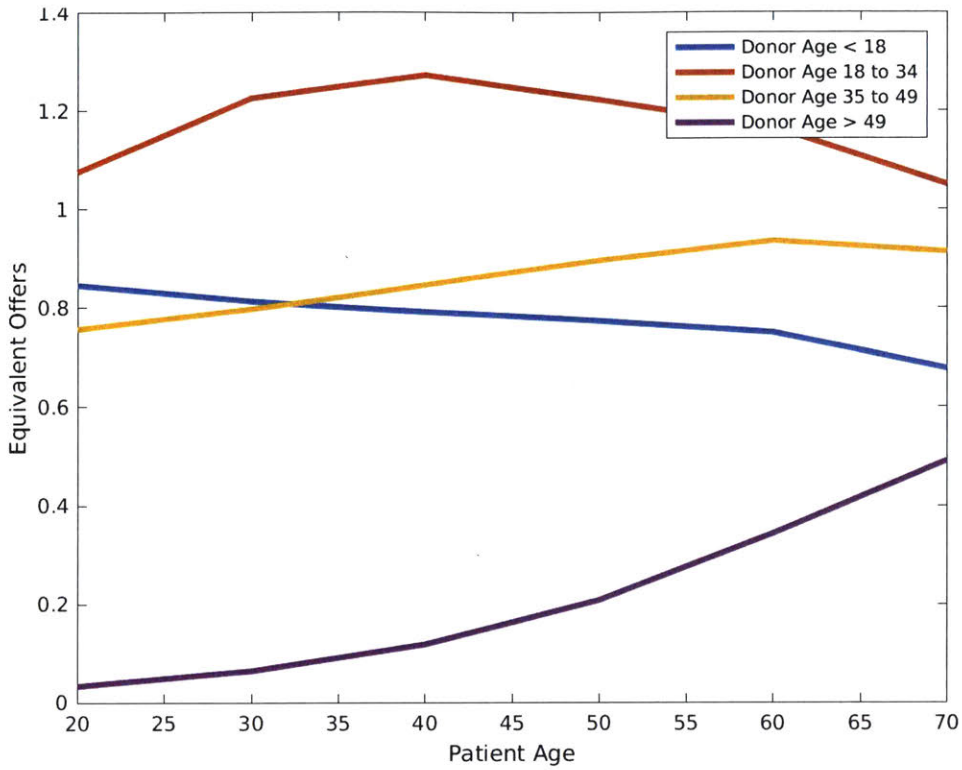

(b) Number of Tissue Type Mismatches

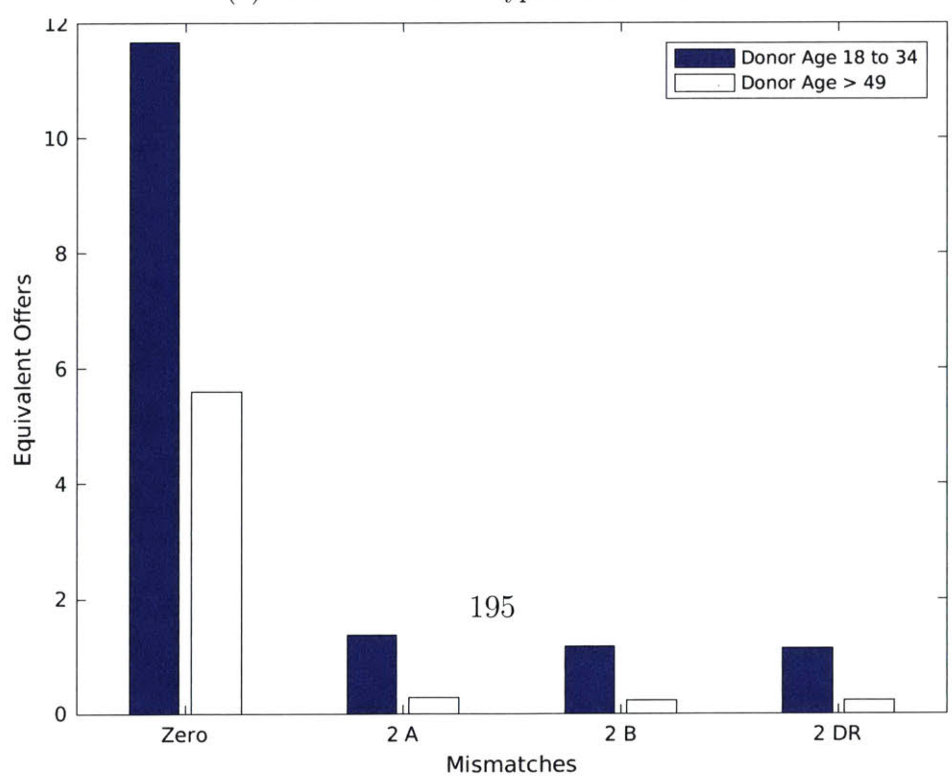

Figure 2-5: Mechanism Comparisons

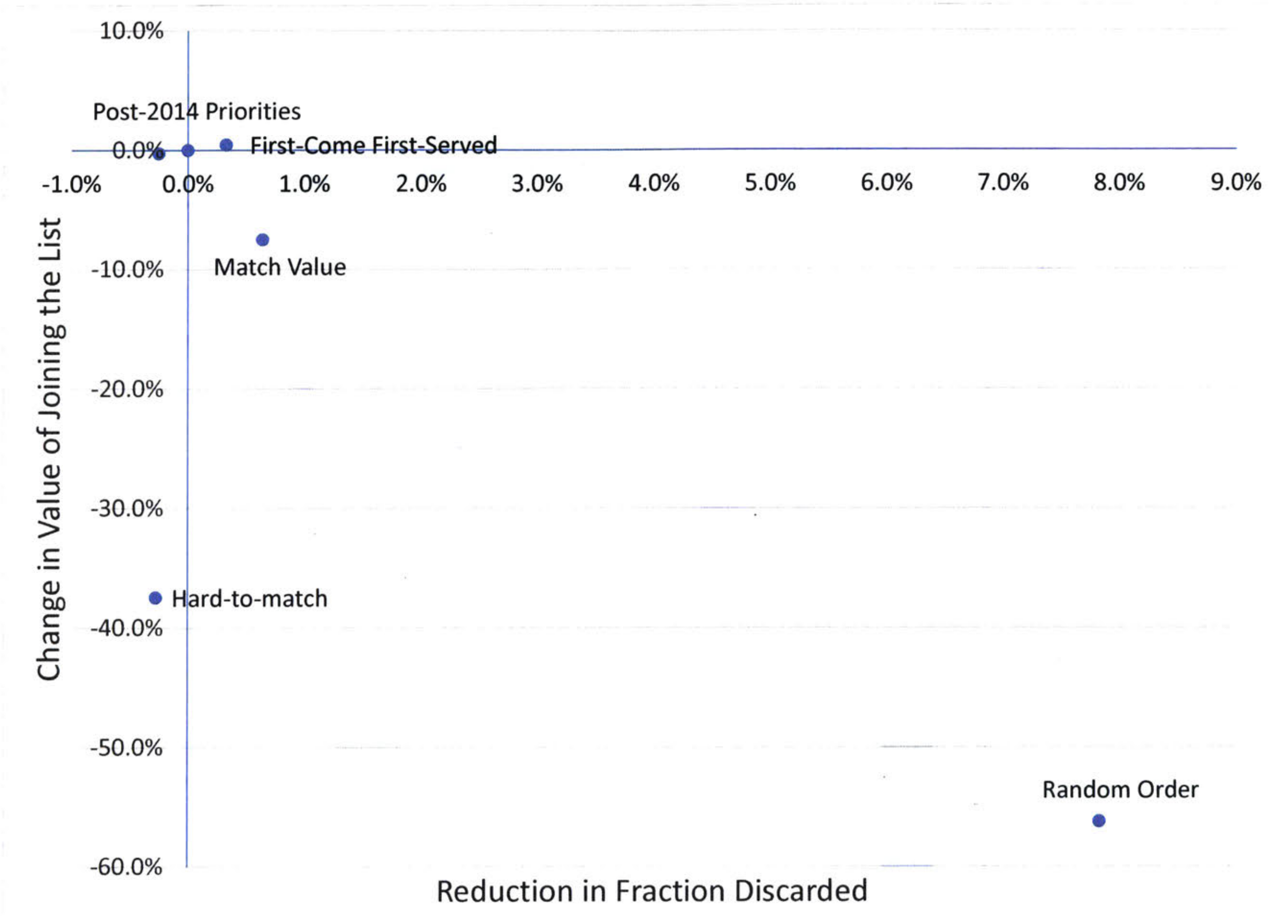

2.8 Appendix

2.8.1 Data Appendix

Data Description

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

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

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

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

Sample Selection

This section explains the selection of patients, donors, and offers used in our structural model. We consider patients who were registered in NYRT and actively waiting for a deceased donor kidney between January 1st, 2010 and December 31st, 2013. Our donor sample includes all U.S. deceased kidney donors whose organs were allocated according to the standard mechanism. Our offer sample includes valid offers from all deceased donors to NYRT patients during this period recorded in the PTR data, as well as offers that were not made because of pre-specified screening criteria. The next section discusses our replication of the offer mechanism, which we used to determine offers that were refused through these screening criteria, as well as the waiting time each patient would need to have access to each compatible donor. The remainder of this section discusses the details of how we selected our samples of donors, patients, and offers that met screening criteria.

Because NYRT patients may be offered donors from across the U.S., our procedure first constructs a nationwide sample of deceased donors, patients, and offers that meet our sample criteria. We then restrict the sample to NYRT and omit certain donors and patients who received non-standard treatment in the mechanism.

Our U.S. sample of deceased kidney donors comes from the intersection of donor identifiers in the PTR and STAR deceased donor files. Patients were active on the deceased donor kidney waitlist after 2010 and were not jointly registered for a pancreas transplant. Patient registration date and activity status were determined from the patient history file. We also exclude patients who departed the waitlist for reasons which indicate that they did not ultimately need a transplant. We exclude patients who were transplanted in another country, whose condition improved, or who could no longer be contacted. These departure reasons are recorded in the STAR patient and transplant outcome dataset.

We then determine which offers were valid and could have been accepted by and transplanted into the patient; patients' acceptance decisions; and the resulting priority score cutoffs in each match run.

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

Next, we created an algorithm to de-duplicate offers and acceptances within and across match runs, and to determine the true priority score cutoff for each donor in each match run. For some donors, multiple match runs are conducted, and these match runs can include offers to overlapping sets of patients. A specific kidney (e.g. the left kidney) may also be accepted in multiple match runs. Finally, a patient can have multiple offers recorded from the same donor, even in the same match run. Our algorithm assumes that later match runs take precedence over earlier ones (using the match run numbers provided by OPTN), and that the last observed match run in which an organ is placed takes that organ out of circulation for subsequent match runs. After de-duplication, there were 11,428,540 offers nationwide that met patient screening criteria. The U.S. sample contains 30,079 donors and 226,000 patients.

We then implement the sample restrictions for NYRT. We consider all patients who were registered in NYRT and had active status sometime between January 1st, 2010 and December 31st, 2013. At this stage, we exclude patients who received a transplant through non-standard allocation rules. This includes cases of medical urgency, an expedited placement attempt, a multi-organ transplant, or a military or directed donation. 68 patients were excluded because they received deceased donor kidney transplants for these reasons, leaving 9,917 patients in our NYRT sample. We also exclude the donors whose organs were placed according to these reasons, even if the organs were allocated to non-NYRT patients. The NYRT offer sample contains valid offers from the sample of deceased donors to the sample of NYRT patients. There are 1,281,024 such offers. These offers and patient acceptance decisions determine the priority score cutoff in each match run for each donor's available organs.

Replicating the Mechanism, Offer Dataset

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

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

From the simulation, we obtain a set of offers predicted by our simulation of the mechanism. These are pairs of donors and patients where the patient met the priority score cutoff and was blood and tissue type compatible with the donor. Some of these offers met the patient's screening criteria, while others did not. Those that did should appear in the PTR data. This provides a check on the performance of our mechanism code. Table 2.17 tabulates offers appearing in our filtered PTR data and those predicted by simulation. The vast majority of offers in the PTR data (93.1%) were predicted by our simulation. However, a substantial fraction of offers predicted by the simulation (29.6%) are not in our PTR offer sample.

To estimate the patient acceptance model, we take as our offer sample the union of the PTR offer dataset and the set of offers that the simulation predicts were filtered due to the patient's screening criteria but which would otherwise have appeared in the PTR data. In a final step, we de-duplicate offers at the patient level, since a patient registered at multiple centers will occasionally receive multiple offers from the same donor. The final offer sample contains 2,850,572 offers: 1,267,531 PTR offers, and 1,583,041 predicted offers that were "screened out." The PTR offers include the 88,366 offers that were not predicted by our simulation but which appear in the PTR data, and exclude the 502,239 offers predicted by the simulation but which did not appear in PTR. Offers that were screened out are interpreted as rejections since the patient deemed the donor's characteristics unacceptable.

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

Imputing Missing Donor DR Antigens

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

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

Second, for deceased donors who had no organs transplanted, we use the reported number of DR mismatches in the PTR offer dataset to impute the donor's DR antigens. Because we observe all patients' HLA antigens, the number of DR mismatches between a donor and patient is informative about the donor's antigens. For example, if a donor-patient pair has zero DR mismatches, the patient's tissue type limits the donor's antigens to a few possibilities.[31] A two DR mismatch pair also restricts the donor's DR antigens, though less so than a zero mismatch. Since deceased donors

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

whose organs are not transplanted are usually offered to many patients, we can combine information across all offers to make an educated guess of the donor's DR antigens.

We use the following imputation algorithm. For each donor without DR antigen information, we take all of the donor's offers in the PTR data. Based on these offers, the recorded number of DR mismatches, and the patient's DR antigens and listed unacceptable antigens, we calculate a score for each DR antigen that the donor might have. We dock one point from a DR antigen's score for each PTR offer it contradicts in the following cases:

- The offer has zero DR mismatches, and the antigen is not equivalent to one of the patient's DR antigens
- The offer has two DR mismatches, and the antigen is equivalent to one of the patient's DR antigens
- The antigen was listed as unacceptable by the patient

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

Life Years From Transplantation (LYFT)

This section describes the methodology used to estimate Life Years From Transplantation (LYFT), a measure of the medical benefits, in terms of quality-adjusted life years, provided by a deceased donor kidney transplant. We follow the methodology of Wolfe et al. (2008) and define LYFT by the following formula:

$$LYFT = 1.0 * (\text{Median Survival With Functioning Graft})$$

$$+0.8 * (\text{Median Survival After Graft Failure})$$

$$-0.8 * (\text{Median Survival Without Transplant})$$

LYFT compares a patient's predicted survival with and without a kidney transplant, weighing years with a functioning kidney more highly than years without one. The comparison is in terms of median predicted survival time in each possible state. See Wolfe et al. (2008) for a discussion of alternative definitions.

Following Wolfe et al. (2008), we estimate three survival models: (1) survival of a patient who remains on the deceased donor kidney waiting list; (2) survival of a patient after receiving a deceased donor kidney transplant; and (3) survival of the functionality of a transplanted deceased donor kidney. From each survival model, we can predict the median survival time for any patient or patient-donor pair. This allows us to construct LYFT for any patient-donor pair.

The rest of this section discusses our implementation of the LYFT measure. We discuss how to construct median survival times from a Cox proportional hazards model; describe the specification of the survival models, including sample and covariate selection; and describe how the survival curves are extrapolated when data are limited.

Estimating Median Survival Times

Our survival models use a Cox proportional hazards specification, allowing a flexible baseline hazard as a function of time and proportional effects of covariates on the hazard rate:

$$\lambda(x, t; \beta) = h_0(t) \exp x'\beta$$

where $\lambda(x, t; \beta)$ is the predicted hazard rate, x are covariates such as patient and donor characteristics, t is the number of years since the beginning of a spell (in our case, joining the waiting list or receiving a kidney transplant), and $h_0(.)$ is a flexible baseline hazard function. Given an estimate of β and characteristics x, median survival beginning at $t = 0$ is $t_{med}(x; \beta)$ such that

$$S_0(t_{med}(x; \beta)) = 0.5^{\frac{1}{\exp x'\beta}}$$

where $S_0(t) = \exp\left(-\int_0^t h_0(t)dt\right)$ is the baseline survival function. Therefore median survival time can be written as

$$t_{med}(x; \beta) = S_0^{-1}\left(0.5^{\frac{1}{\exp x'\beta}}\right)$$

Implementation of Cox Proportional Hazards Models

We estimated the Cox proportional hazards models in STATA using the *stcox* command. This section describes sample selection, covariate selection, and the definition of survival time for each model.

The sample for the waiting list survival model consists of all patients registered for a kidney transplant in the UNOS system between 1988 and 2014, for a total of

519,624 spells. The patient data contain mortality information for patients who died while still registered on the waiting list, and also from the Social Security Death Master File (SSDMF) for patients who departed the waiting list before death. The SSDMF records all patient deaths through 2011 but is incomplete thereafter. The beginning of a spell is defined as a patient's registration date, and the end of a spell occurs when a patient departs the waitlist. Waitlist survival time is censored if the patient departed the waitlist for a reason other than death, or at the end of 2011 (whichever occurred first). The sample for the graft and post-transplant survival models consists of all patients who received a deceased donor kidney transplant between 1988 and 2011, for a total of 153,479 (154,363) spells. Each spell begins at the time of transplantation, and ends when the graft fails or the patient dies. Survival times are censored at the end of 2011.

Covariates were chosen as in Wolfe et al. (2008). These variables are a subset of those used in our structural model and are specified in the same way here. Covariates in the waitlist survival model include patient characteristics at listing that predict health status: diabetes status, CPRA at registration, whether the patient received a previous transplant, dialysis time at registration, patient age at registration, body mass index, total serum albumin, and blood type. We also include indicators for the registration date in five-year intervals to account for technological changes over time that led to improved survival for dialysis patients.

For the graft and post-transplant survival models, we include patient and donor characteristics at the time of transplantation, as well as characteristics pertaining to the match. Patient characteristics at transplantation include patient age, waiting time, CPRA, diabetes status, body mass index, and total serum albumin. We also include indicators for a previous transplant, dialysis time at registration, blood type, and whether the patient was also on the pancreas waiting list. Donor characteristics include donor age, cause of death, gender, creatinine level, hypertensive status, and an expanded criteria donor (ECD) indicator. Match characteristics include the number of A, B, and DR mismatches; geographic match; and blood type compatibility; interactions between geography and donor age, and geography and number of HLA mismatches; interactions between zero HLA mismatches and geography, patient age, and CPRA; interactions between patient waiting time and geography, patient age, and CPRA; and interactions between patient and donor age. We also include indicators for the date of transplantation in five-year intervals to account for improved post-transplant survival over time.

Table 2.14 displays coefficient estimates from each of the three Cox proportional hazards models. Coefficient estimates generally go in the expected direction – for example, diabetic and highly sensitized patients have lower survival rates both on the waiting list and after transplantation. Younger donors, local donors, and perfect tissue type matches predict longer post-transplant survival. In addition, falling coefficient estimates on the time period indicators show that overall, survival rates have been improving over time both on the waiting list and after transplantation.

Survival Curve Extrapolation

The mapping from a patient-donor pair's hazard ratio, $\exp x'\beta$, to predicted median survival time depends on the baseline survival curve S_0. However, for some covariate values with low hazard ratios, predicted median survival times exceed the range of the baseline survival curve that can be estimated from the data. We therefore extrapolate the baseline survival curve for each model out to 40 years, parameterizing the log survival probability as a linear function of time:

$$log(S_0(t)) = \alpha t$$

Figure 2-8 compares the baseline survival curves from the model to the extrapolated curves. For each survival model, the parameter α is estimated by ordinary least squares using the baseline survival curve in the range of survival times observed in the data.

Construction of Median Survival Times and LYFT

For each estimated survival model, we construct a fine grid of hazard ratio values $r = \exp x'\beta$ and calculate the median survival time for each value r given the baseline survival function S_0. With this grid, we can map any covariates x to the closest hazard ratio r and the corresponding median survival time, $t_{med}(x; \beta)$. To predict median continued survival without a transplant at the time of a kidney offer, the offer date must be taken into account, since the baseline hazard function in the waitlist survival model depends on the patient's total time on the waitlist. In other words, we must condition on the fact that the patient has already survived for s years on the waitlist. Median survival on the waitlist, $t^{WL}_{med}(x, s; \beta)$ is therefore a function of s and given by the formula

$$\exp\{\log S_0(t^{WL}_{med}(x, s; \beta)) - \log S_0(s)\} = 0.5^{\frac{1}{\exp x'\beta}}$$

We construct an inverse map from r to median survival time at a grid of waiting times $s \in 0, 100, 200, ..., 5000$. When we calculate LYFT for a particular match x and patient waiting time s, we round the hazard ratio and waiting time to the nearest values in this grid. Letting $t^{GR}_{med}(x; \beta)$ denote graft survival and $t^{TX}_{med}(x; \beta)$ denote post-transplant survival, the final LYFT formula is

$$LYFT(x,s;\beta) = 1.0 * t_{med}^{TX}(x;\beta) - 0.2 * \max[0, t_{med}^{TX}(x;\beta) - t_{med}^{GR}(x;\beta)] - 0.8 * t_{med}^{WL}(x,s;\beta)$$

2.8.2 Estimation

Normalization

In this section, we show that the model described in Section 2.3.3 yields the same decision-rules as a model in which $d_i(t) + \delta_i(t) D_i(t) = 0$ and the net present value of a transplant is $\tilde{\Gamma}_{ij}(t) - O_i(t)$. To do this, the next proposition first derives implication on any assignment rule.

Proposition 1. *Let $p_{ij}(t)$ be the conditional probability that i is assigned object j given that j arrives in period t. Let $V_i(t;p)$ be i's value of the assignment rule p.*
1. For any assignment rule p, $\Gamma_{ij}(t) - V_i(t;p) = \tilde{\Gamma}_{ij}(t) - \tilde{V}_i(t;p)$ where $\tilde{\Gamma}_{ij}(t) = \Gamma_{ij}(t) - O_i(t)$ and

$$(\rho + \delta_i(t)) \tilde{V}_i(t;p) = \lambda \int p_{ij}(t) \left(\tilde{\Gamma}_{ij}(t) - \tilde{V}_i(t;p) \right) \mathrm{d}F + \dot{\tilde{V}}_i(t;p)$$

with boundary condition $\tilde{V}_i(T_i;p) = 0$.

2. For any two assignment rules p and p', $V_i(t;p) - V_i(t;p') = \tilde{V}_i(t;p) - \tilde{V}_i(t;p')$

Proof. Part 1: First, we verify that $\tilde{V}_i(t;p) = V_i(t;p) - O_i(t)$ satisfies the differential equation above. Note that

$$(\rho + \delta_i(t)) V_i(t;p) = d_i(t) + \delta_i(t) D_i(t) + \lambda \int p_{ij}(t) (\Gamma_{ij}(t) - V_i(t;p)) \mathrm{d}F + \dot{V}_i(t;p).$$

Therefore,

$$V_i(t;p) - O_i(t) = \lambda \int p_{ij}(t)\left(\tilde{\Gamma}_{ij}(t) - (V_i(t;p) - O_i(t))\right) \mathrm{d}F + \frac{\partial}{\partial t}(V_i(t;p) - O_i(t)).$$

Hence, $W_i(t;p)$ satisfies the necessary differential equation. It is straightforward to check that $V_i(T_i;p) = O_i(T_i) = D_i(T_i)$ showing that the solution with the boundary condition $\tilde{V}_i(T_i;p) = 0$ satisfies the requirements of the proposition.

Part 2: Observe that

$$\begin{aligned} V_i(t;p) - V_i(t;p') &= \lambda \int p_{ij}(t)(\Gamma_{ij}(t) - V_i(t;p)) \mathrm{d}F + \dot{V}_i(t;p) - \dot{V}_i(t;p') \\ &= \lambda \int p_{ij}(t)\left(\tilde{\Gamma}_{ij}(t) - W_i(t;p)\right) \mathrm{d}F + \dot{\tilde{V}}_i(t;p) - \dot{\tilde{V}}_i(t;p') \\ &= \tilde{V}_i(t;p) - \tilde{V}_i(t;p'). \end{aligned}$$

□

Refer to the model with $O_i(t) = 0$ for all t and the related value function $\tilde{V}_i(t;p)$ and the payoffs $\tilde{\Gamma}_{ij}(t)$ as the normalized model. Part 1 shows that the normalized model also yields $\Gamma_{ij}(t) - V_i(t;p)$ as the difference in the value of acceping j relative to the value waiting if one expects assignments according to p. In particular, the result holds for $p_{ij}(t) = \pi_{ij}(t) 1\{\Gamma_{ij}(t) - V_i(t) > 0\}$. Therefore, the normalized model yields an identical choice rule and value function relative to no assignment.

Part 2 shows that the normalized model yields an identical difference in value functions between any two assignment rules as the original model. This is useful for evaluating both alternative mechanisms as well as for solving for equilibria. To see this, consider any action space $\mathcal{A}_t$ and strategy $\sigma_i(t;j) \to \mathcal{A}_t$. The action space need not be binary and may depend on time. For example, to consider a counterfactual with multiple wait-lists we can define the action set as the choice of list at birth and an accept/reject decision thereafter. As the notation indicates, each action can depend on the currently offered object, if any. As long as the analyst can then evaluate the assignment rule $p_{ij}(t;\sigma)$ as a function of the strategy profile $\sigma = (\sigma_i, \sigma_{-i})$, the result says that the normalized model can be used to determine the difference in values. To solve for equilibria, we would need to evaluate deviations (σ_i', σ_{-i}) and compare

$$V_i(t; p(\sigma_i', \sigma_{-i})) - V_i(t; p(\sigma_i, \sigma_{-i})) = \tilde{V}_i(t; p(\sigma_i', \sigma_{-i})) - \tilde{V}_i(t; p(\sigma_i, \sigma_{-i})).$$

To identify the value function relative to the current mechanism, we would need to compute

$$V_i(t; p(\sigma^*)) - V_i(t; \hat{p}) = \tilde{V}_i(t; p(\sigma^*)) - \tilde{V}_i(t; \hat{p}),$$

where $\hat{p}$ denotes the assignment probabilities under the factual mechanism and $p(\sigma^*)$ denotes the equilibrium assignment probabilities in an equilibrium of the counterfactual mechanism.

Details on the Estimator

Gibbs' Sampler

The sampler is initialized at any value of θ^0, σ_η^0 and guesses for η_j^0 and y_{ijt}^0 corresponding to observed decisions such that $y_{ijt}^0 > 0$ if and only if agent i accepted object j in period t. We then sample from the conditional posteriors and draws of y

given the previous draws. The sampler iterates through the following sequence

$$y_{ijt}^{s+1}|\theta^s, \eta_j^s; a_{ijt}$$
$$\eta_j^{s+1}|y_j^{s+1}, \theta^s$$
$$\theta^{s+1}|y^{s+1}, \eta^{s+1}$$
$$\sigma_\eta^{s+1}|\eta^{s+1}$$

where the conditioning on the priors and the observables is implicit, y^s and η^s are vectors with components y_{ijt}^s and η_j^s, and y_j^{s+1} is a vector that stacks y_{ijt}^{s+1} across all it. The first two steps involve data augmentation to simplify the sampling problem of the key parameters in the next step. Each of these steps involves draws from a closed-form distribution if the prior distribution on σ_η is specified as an inverse-Gamma distribution and the prior for $\theta \sim N\left(\bar{\theta}, \Sigma_\theta\right)$. With these priors, the first step involved sampling from a truncated normal, the second and third steps involve sampling from a normal distribution and the final step involves sampling from an inverse-Gamma.

Computing the Value Function

Numerical integration using. 400 draws from exponential distribution. from t_0 to T. Re-use draw for any other t and reweight.

Given t, for each patient i, we compute the value of continuing is given by equation (2.6). Using equation (2.8), the sample analog of the value of continuing is given by

$$\hat{V}_i(t) = \lambda \int_0^{T_i - t} \exp(-\rho\tau)\, p(\tau + t|t; x_i)\, \hat{W}\left(x_i, \tau + t; \hat{\theta}\right) \mathrm{d}\tau.$$

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

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

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

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

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

This procedure ensures that there are B points of evaluation for each possible donor and patient-time pair. The numerical performance is superior to an alternative that approximates the integral in equation (2.6) as a sum over a fixed set of draws because some patient, donor, time combinations may have a very small window of availability, $\left[\underline{\tau}_{ijt}, T_i\right]$.

Auxiliary Models

Positive Crossmatch Probability

Not all accepted offers result in transplantation. Even after a provisional acceptance, additional testing may yield a positive crossmatch indicating that the patient is likely to reject the donor's kidney. These transplants are not executed, and if possible the organ is placed with another patient. To compute patient value functions and conduct counterfactual simulations, we must account for positive crossmatches. We therefore estimate a probit model to predict probability that a patient has a positive crossmatch with an organ they have accepted. We include as predictors interactions between the patient's CPRA and the number of HLA mismatches with the donor.[32] Coefficient estimates and standard errors are displayed in Table 2.13. Higher CPRA is associated with a higher positive crossmatch probability, as are more DR or HLA mismatches. In addition, CPRA and tissue type matches interact: high CPRA patients have an additional benefit, in terms of the normal index, from fewer HLA mismatches. This is intuitive: patients with more sensitized immune systems may be more likely to test positive against foreign antibodies, even if they have not tested positive in the past.

Maximum Number of Offers

Our data contain cases in which one of a donor's organs is discarded before being offered to all compatible patients. This usually occurs for two reasons. First, the organ may become unsuitable for transplantation if it remains outside donor's body for too long. Second, the organ may be accepted by a patient in another OPO. We call

[32]Using a predictive model in counterfactuals assume that test results are not manipulated, and therefore would follow the same distribution under counterfactual mechanisms. Conversations with transplant surgeons support this assumption.

these events "timeouts." Timeouts are driven by a combination of unobserved factors, including whether the organ remained in the donor's body during the offer process; the rate at which offers were made, which depends on patient/surgeon response times; the kidney's rate of physical deterioration once outside the body; and decisions of patients outside NYRT.

We model the maximum number of offers that can be made for a given donor using a censored exponential hazards model. Duration is the number of observed offers that met pre-specified screening criteria. Censoring occurs if a donor's organs are placed, or if they are discarded after being offered to all compatible patients. The hazard function is given by

$$\lambda_o(z) = \lambda_o \exp(z\beta) \tag{2.9}$$

where z are characteristics of the donor, β is a vector of coefficients, and λ_0 is the constant baseline hazard rate. We allow the timeout hazard to depend on geography and indicators of donor quality. Specifically, we control for whether the donor is an expanded criteria donor (ECD); the donor's cause of death (DCD); and whether the donor was recovered in NYRT, as well as interactions among these variables. The estimated timeout hazards are inputs to the structural model.

Consistency of $\hat{W}\left(x_i, t; \hat{\theta}\right)$

We now show that $\hat{W}\left(x_i, t; \hat{\theta}\right)$ is a consistently estimates the quantity

$$\begin{aligned} W(x_i, t; \theta) &= \int \pi_{ij}(t)\, \psi(P_{ijt})\, \mathrm{d}F \\ &= \int H_{z_j,\eta_j}(s(t; x_i, Z))\, \mathbb{P}(c_{ij} = 1 | z_j, x_i; \theta)\, \psi(x_i, Z, \eta, t; \theta)\, \mathrm{d}F_{Z,\eta}. \end{aligned}$$

To do this, we need to introduce some notation. Define

$$g_j(\theta) = \mathbb{P}(c_{ij} = 1 | z_j, x_i; \theta)\, \psi(x_i, z_j, \eta_j, t; \theta)\, 1\left\{s(t; x_i, z_j) > s_j^*\right\}.$$

For a vector $x = (x_1, \ldots, x_K)$, we write $|x| = (|x_1|, \ldots, |x_K|)$.

Finally, we index objects according to the order in which they arrive in our sample. Therefore, (z_j, η_j) denotes the observed and unobserved characteristics of the j-th donor that arrived. Therefore, the data on $\left(z_j, \eta_j, s_j^*\right)$ generates a sequence of object arrivals.

We make the following assumptions on $g_j(\theta)$:

Assumption 4. *(i) (z_j, η_j) is drawn i.i.d. with cdf F*

(ii) $g_j(\theta_0)$ is weakly stationary[33] *with $\sum_{k=-\infty}^{\infty} \gamma_k < \infty$, where $Cov(g_j(\theta_0), g_{j-k}(\theta_0)) = \gamma_k$.*

(iii) $\hat{\theta}_J$ is $\sqrt{J}$-consistent, i.e. $\left(\hat{\theta}_J - \theta_0\right) = O_p\left(J^{-1/2}\right)$

(iv) There exists an $m(z)$ with finite second moments, such that $|g(z_j; \theta) - g(z_j; \theta')| \leq m(z_j) \cdot |\theta - \theta'|$.

Part (i), in our empirical context, assumes that the characteristics of the donor are drawn independently each time a donor arrives. Part (ii) assumes that, at θ_0, the offers and their values for any given patient type x_i, at any given time t, follows a weakly stationary process. That is, the covariance in these values across any two donors falls as they are further apart in the sequence. Given part (i), the only potential source of dependence between $g_j(\theta_0)$ and $g_k(\theta_0)$ is that the characteristics of donor j may affect the state of the waitlist for donor k because of patient decisions. However, we expect that this dependence to fall as these donors become

[33]The process $\{g_j\}$ is weakly stationary if (i) $\mathbb{E}[g_j]$ does not depend on j, and (ii) $Cov(g_j, g_{j-k})$ exists, is finite and depends only on k, and not j.

further apart in their arrival sequence. Part (iii) assumes that $\hat{\theta}_J$ is consistently estimated at a rate that is at least as fast as the square-root of the number of donors. These parameters govern the conditional choice probabilities and the probability of a crossmatch failure at biological testing. Part (iv) is a regularity condition, assuming that $g(z_j;\theta)$ is Lipschitz continuous at each z, with a sufficiently bounded second moment. Proposition 2 shows that this property is satisfied under more primitive conditions stated in Assumption 5.

We now show that for each x_i, t,$\hat{W} = \frac{1}{J}\sum_{j=1}^{J} g_j\left(\hat{\theta}\right)$ is a $\sqrt{J}$-consistent estimator of $W(x_i, t; \theta_0)$ under this assumption.

Theorem 1. *Fix x_i, t. If Assumption 4 is satisfied, then*

$$\left| W(x_i, t; \theta_0) - \frac{1}{J}\sum_{j=1}^{J} g_j\left(\hat{\theta}\right)\right| = O_p\left(J^{-1/2}\right).$$

Proof. For each x_i, t, Assumption 4(i) implies that

$$\begin{aligned}
\mathbb{E}\left[g_j(\theta_0)\right] &= \mathbb{E}\left[p_c(z_j, x_i; \theta_0)\,\psi(x_i, z_j, \eta_j, t; \theta_0)\,1\left\{s(t; x_i, z_j) > s_j^*\right\}\right] \\
&= \mathbb{E}\left[p_c(z_j, x_i; \theta_0)\,\psi(x_i, z_j, \eta_j, t; \theta_0)\,\mathbb{E}\left[1\left\{s(t; x_i, z_j) > s_j^*\right\}|z_j, \eta_j\right]\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[p_c(z_j, x_i; \theta_0)\,\psi(x_i, z_j, \eta_j, t; \theta_0)\,H_{z_j,\eta_j}(s(t; x_i, z_j))\,|z_j, \eta_j\right]\right] \\
&= W(x_i, t; \theta_0),
\end{aligned}$$

where the equalities are a result of the law of iterated expectations, and the definitions of $H_{z_j,\eta_j}(s)$ and $W(x_i, t; \theta_0)$.

Because $W(x_i, t; \theta_0) = \mathbb{E}[g_j(\theta_0)]$, Chebychev's inequality implies that,

$$\mathbb{P}\left(J^{1/2}\left|W(x_i, t; \theta_0) - \frac{1}{J}\sum_{j=1}^{J} g_j(\theta_0)\right| > \varepsilon\right) \leq \frac{Var\left(\frac{1}{\sqrt{J}}\sum_{j=1}^{J} g_j(\theta_0)\right)}{\varepsilon^2}.$$

Assumption 4(i) and Proposition 6.8 in Hayashi (2001) implies that

$$\lim_{J\to\infty} Var\left(\frac{1}{\sqrt{J}}\sum_{j=1}^{J} g_j(\theta_0)\right) = K.$$

Therefore, we have that

$$W(x_i, t; \theta_0) - \frac{1}{J}\sum_{j=1}^{J} g_j(\theta_0) = O_p\left(J^{-1/2}\right). \tag{2.10}$$

Lemma 1 implies that

$$W(x_i, t; \theta_0) - \frac{1}{J}\sum_{j=1}^{J} g_j\left(\hat{\theta}\right) = O_p\left(J^{-1/2}\right),$$

as requirements (i) and (ii) of Lemma 1 are part of Assumption 4(iv), requirement (iii) is equivalent to Assumption 4(iii), and requirement (iv) is proved in equation (2.10). $\square$

Lemma 1. *Fix x_i, t. Suppose that (i) $g(z_j; \theta)$ is Lipschitz continuous for each z_j with the Lipschitz constant $m(z_j) \in \mathbb{R}^K$ i.e. $|g(z_j; \theta) - g(z_j; \theta_0)| \leq m(z_j) \cdot |\theta - \theta_0|$, (ii) $m(z_j)$ has finite second moments, (iii) $\left|\hat{\theta}_J - \theta_0\right| = O_p\left(J^{-1/2}\right)$, and (iv) $\frac{1}{J}\sum_j g(z_j; \theta_0) - E[g(z_j; \theta_0)] = O_p\left(J^{-1/2}\right)$, then*

$$\frac{1}{J}\sum_{j} g\left(z_j; \hat{\theta}_J\right) - E[g(z_j; \theta_0)] = O_p\left(J^{-1/2}\right).$$

Proof. Because $g(z_j;\theta)$ is Lipschitz continuous in θ, we have that

$$Var\left(\left|g(z_j;\theta) - g(z_j;\theta_0)\right|\right) \le \left|\theta - \theta_0\right|^T \mathbb{E}\left(m(z_j)\, m^T(z_j)\right) \left|\theta - \theta_0\right|,$$

where x^T is the transpose of the vector x. Therefore, because $|\theta - \theta_0| = O_p\left(J^{1/2}\right)$, with probability approaching 1,

$$Var\left(g(z_j;\theta) - g(z_j;\theta_0)\right) \le V_M K J^{-1},$$

for some finite constant $K > 0$, where $V_M = \mathbb{E}\left(\sum_{k,k'} m_k(z_j)\, m_{k'}(z_j)\right)$ and $m_k(z_j)$ is the k-th component of $m(z_j)$. By the covariance inequality, with probability approaching 1,

$$Cov\left(g\left(z_j;\hat{\theta}_J\right) - g(z_j;\theta_0), g\left(z_k;\hat{\theta}_J\right) - g(z_k;\theta_0)\right) \le V_M K J^{-1}.$$

Therefore, with probability approaching 1,

$$\begin{aligned}
&Var\left(\frac{1}{\sqrt{J}}\sum_j g\left(z_j;\hat{\theta}_J\right) - \frac{1}{\sqrt{J}}\sum_j g(z_j;\theta_0)\right) \\
=&\frac{1}{J}\sum_{k=1}^{J}\sum_{j=1}^{J} Cov\left(g\left(z_j;\hat{\theta}_J\right) - g(z_j;\theta_0), g\left(z_k;\hat{\theta}_J\right) - g(z_k;\theta_0)\right) \le V_M K.
\end{aligned}$$

By Chebychev's inequality,

$$\mathbb{P}\left(\sqrt{J}\left|\frac{1}{J}\sum_j g\left(z_j;\hat{\theta}_J\right) - E\left[g\left(z_j;\hat{\theta}_J\right)\right] - \frac{1}{J}\sum_j g(z_j;\theta_0) + E\left[g(z_j;\theta_0)\right]\right| > \varepsilon\right) \le \frac{V_M K}{\varepsilon^2}.$$

Therefore,

$$A_J = \frac{1}{J}\sum_j g\left(z_j;\hat{\theta}_J\right) - E\left[g\left(z_j;\hat{\theta}_J\right)\right] - \frac{1}{J}\sum_j g\left(z_j;\theta_0\right) + E\left[g\left(z_j;\theta_0\right)\right] = O_p\left(J^{-1/2}\right).$$

But, we know that $\frac{1}{J}\sum_j g\left(z_j;\theta_0\right) - E\left[g\left(z_j;\theta_0\right)\right] = O_p\left(J^{-1/2}\right)$. So, it must be that

$$\frac{1}{J}\sum_j g\left(z_j;\hat{\theta}_J\right) - E\left[g\left(z_j;\hat{\theta}_J\right)\right] = O_p\left(J^{-1/2}\right). \tag{2.11}$$

Lipschitz continuity of $g\left(z_j;\theta\right)$ and the Cauchy-Schwarz inequality imply that.

$$\begin{aligned} E\left[\left|g\left(z_j;\hat{\theta}_J\right) - g\left(z_j;\theta_0\right)\right|\right] &\leq E\left[m\left(z_j\right)\cdot\left|\hat{\theta}_J - \theta_0\right|\right] \\ &\leq \sqrt{V_M}\sqrt{E\left[\left|\hat{\theta}_J - \theta_0\right|\cdot\left|\hat{\theta}_J - \theta_0\right|\right]}. \end{aligned}$$

Because $\hat{\theta}_J$ belongs to the compact set Θ, and $\hat{\theta}_J - \theta_0 = O_p\left(J^{-1/2}\right)$, we have that $\sqrt{E\left[\left|\hat{\theta}_J - \theta_0\right|\cdot\left|\hat{\theta}_J - \theta_0\right|\right]} = O_p\left(J^{-1/2}\right)$. Together with the assumption that $V_M = \mathbb{E}\left(\sum_{k,k'} m_k\left(z_j\right) m_{k'}\left(z_j\right)\right)$ is finite, we have that

$$E\left[\left|g\left(z_j;\hat{\theta}_J\right) - g\left(z_j;\theta_0\right)\right|\right] = O_p\left(J^{-1/2}\right). \tag{2.12}$$

Finally, equation (2.11) and (2.12) together imply that

$$
\begin{aligned}
&\frac{1}{J}\sum_j g\left(z_j;\hat{\theta}_J\right)-E\left[g\left(z_j;\hat{\theta}_J\right)\right]\\
=&\frac{1}{J}\sum_j g\left(z_j;\hat{\theta}_J\right)-E\left[g\left(z_j;\theta_0\right)\right]+E\left[g\left(z_j;\theta_0\right)\right]-E\left[g\left(z_j;\hat{\theta}_J\right)\right]\\
\leq&\frac{1}{J}\sum_j g\left(z_j;\hat{\theta}_J\right)-E\left[g\left(z_j;\theta_0\right)\right]+E\left[\left|g\left(z_j;\theta_0\right)-g\left(z_j;\hat{\theta}_J\right)\right|\right]\\
=&O_p\left(J^{-1/2}\right).
\end{aligned}
$$

□

Lipschitz continuity of $g\,(z;\theta)$

We now show primitive regularity conditions under which Assumption 4(iv) is satisfied. Recall that $g_j\,(\theta)=\mathbb{P}\,(c_{ij}=1|z_j,x_i;\theta)\,\psi\,(x_i,z_j,\eta_j,t;\theta)\,1\left\{s\,(t;x_i,z_j)>s_j^*\right\}$. Fix x_i, t and omit it from the notation for simplicity.

Assumption 5. *(i) $\psi\,(z_j,\eta_j;\theta)=\mathbb{E}\,[\max\{0,\chi\,(z_j)\cdot\theta+\eta_j+\varepsilon\}]$ where ε has cdf F_ε*

(ii) There exists $m\,(z)\in\mathbb{R}^{k_\theta}$ with finite fourth moments such that for all θ, θ',

a. $|\mathbb{P}\,(c_{ij}=1|z_j;\theta)-\mathbb{P}\,(c_{ij}=1|z_j;\theta')|\leq m\,(z_j)\cdot|\theta-\theta'|$

b. $|\chi\,(z)|\leq m\,(z)\,.$

Part (i) follows from the definition of $\psi\,(z_j,\eta_j;\theta)$ and the parametrization of the conditional choice probabilities in our model. It is repeated simply to keep this exercise self-contained. Part (ii) is a regularity condition that assumes lipschitz continuity of primitives, with sufficiently small lipschitz constants.

Proposition 2. *If Assumption 5 is satisfied, then for all $\theta,\theta'\in\Theta$, $|g_j\,(\theta)-g_j\,(\theta)|\leq\tilde{m}\,(z_j)\cdot|\theta-\theta'|$, where $\tilde{m}\,(z_j)$ has finite second moments.*

Proof. First, we show that for all $\theta, \theta' \in \Theta$ $|\psi(z_j, \eta_j; \theta) - \psi(z_j, \eta_j; \theta')| \leq m(z_j) \cdot |\theta - \theta'|$. Define $\gamma(x) = \mathbb{E}[\max\{0, x + \eta + \varepsilon\}] = \int_{-x-\eta}^{\infty} (x + \eta + \varepsilon)\, \mathrm{d}F_{\varepsilon}$. Libniz's rule implies that

$$\gamma'(x) = \int_{-x-\eta}^{\infty} 1 \mathrm{d}F_{\varepsilon} = 1 - F(-x - \eta) \leq 1.$$

Therefore, $\gamma(x)$ is Lipschitz continuous with constant 1. Hence, Assumption 5(i) and (ii)b., and Lemma 3 imply that

$$|\psi(z_j, \eta_j; \theta) - \psi(z_j, \eta_j; \theta')| \leq \chi(z_j) \cdot |\theta - \theta'| \leq m(z_j) \cdot |\theta - \theta'|. \tag{2.13}$$

Equation (2.13), Assumption 5(ii)a. and Lemma 2 imply that $|g_j(\theta) - g_j(\theta)| \leq 2(m(z) * m(z)) \cdot |\theta - \theta'|$. Assumption 5(ii) implies that $\tilde{m}(z) = 2(m(z) * m(z))$, where $*$ is the hadamard (or component-wise) product, has finite second moments. □

Lemma 2. *Suppose that there exists a function $m(z)$ such that (i) $\tilde{\chi}(z) \leq m(z)$ and $|\chi(z, \theta) - \chi(z, \theta')| \leq m(z) \cdot |\theta - \theta'|$, and (ii) $\sup_{\theta} |\tilde{\chi}(z, \theta)| \leq m(z)$ and $\sup_{\theta} |\chi(z, \theta)| \leq m(z)$. Then,*

$$|\tilde{\chi}(z, \theta)\, \chi(z, \theta) - \tilde{\chi}(z, \theta')\, \chi(z, \theta')| \leq 2(m(z) * m(z)) \cdot |\theta - \theta'|,$$

where $*$ is the hadamard (or component-wise) product.

Proof. Suppose that $|\tilde{\chi}(z, \theta) - \tilde{\chi}(z, \theta')| \leq \tilde{m}(z) \cdot |\theta - \theta'|$ and $|\chi(z, \theta) - \chi(z, \theta')| \leq$

$m(z) \cdot |\theta - \theta'|$. Consider

$$\begin{aligned}
&|\tilde{\chi}(z,\theta)\chi(z,\theta) - \tilde{\chi}(z,\theta')\chi(z,\theta')| \\
=&|\tilde{\chi}(z,\theta)\chi(z,\theta) - \tilde{\chi}(z,\theta)\chi(z,\theta') + \tilde{\chi}(z,\theta)\chi(z,\theta') - \tilde{\chi}(z,\theta')\chi(z,\theta')| \\
\leq&|\tilde{\chi}(z,\theta)(\chi(z,\theta) - \chi(z,\theta'))| + |(\tilde{\chi}(z,\theta) - \tilde{\chi}(z,\theta'))\chi(z,\theta')| \\
\leq&|\tilde{\chi}(z,\theta)|\, m(z) \cdot |\theta - \theta'| + |\chi(z,\theta')|\, \tilde{m}(z) \cdot |\theta - \theta'| \\
\leq&2m^2(z) \cdot |\theta - \theta'|.
\end{aligned}$$

Suppose that $|\tilde{\chi}(z,\theta) - \tilde{\chi}(z,\theta')| \leq \tilde{m}(z) \cdot |\theta - \theta'|$ and $|\chi(z,\theta) - \chi(z,\theta')| \leq m(z) \cdot |\theta - \theta'|$. Consider

$$\begin{aligned}
&|\tilde{\chi}(z,\theta)\chi(z,\theta) - \tilde{\chi}(z,\theta')\chi(z,\theta')| \\
=&|\tilde{\chi}(z,\theta)\chi(z,\theta) - \tilde{\chi}(z,\theta)\chi(z,\theta') + \tilde{\chi}(z,\theta)\chi(z,\theta') - \tilde{\chi}(z,\theta')\chi(z,\theta')| \\
\leq&|\tilde{\chi}(z,\theta)(\chi(z,\theta) - \chi(z,\theta'))| + |(\tilde{\chi}(z,\theta) - \tilde{\chi}(z,\theta'))\chi(z,\theta')| \\
\leq&|\tilde{\chi}(z,\theta)|\, m(z) \cdot |\theta - \theta'| + |\chi(z,\theta')|\, m(z) \cdot |\theta - \theta'| \\
\leq&2(m(z) * m(z)) \cdot |\theta - \theta'|.
\end{aligned}$$

□

Lemma 3. *Suppose that (i) $\psi : \mathbb{R} \to \mathbb{R}$ is Lipschitz continuous with constant $K < \infty$, and (ii) for each z, there exists $m(z) \in \mathbb{R}^{K_\theta}$ such that $|\tilde{\chi}(z,\theta) - \tilde{\chi}(z,\theta')| \leq m(z) \cdot |\theta - \theta'|$, then $|\psi(\tilde{\chi}(z,\theta)) - \psi(\tilde{\chi}(z,\theta'))| \leq Km(z) \cdot |\theta - \theta'|$. In particular, if $\tilde{\chi}(z,\theta) = \chi(z) \cdot \theta$ where $\chi(z)$, $\theta \in \mathbb{R}^{K_\theta}$, then $|\psi(\tilde{\chi}(z,\theta)) - \psi(\tilde{\chi}(z,\theta'))| \leq K|\chi(z)| \cdot |\theta - \theta'|$.*

Proof. The first part follows definitionally because Lipschitz continuity of ψ implies that

$$\left|\psi\left(\tilde{\chi}\left(z,\theta\right)\right)-\psi\left(\tilde{\chi}\left(z,\theta'\right)\right)\right|\leq K\left|\tilde{\chi}\left(z,\theta\right)-\tilde{\chi}\left(z,\theta'\right)\right|\leq Km\left(z\right)\cdot\left|\theta-\theta'\right|.$$

For the second part, note that

$$\chi\left(z\right)\cdot\left(\theta-\theta'\right)\leq\left|\chi\left(z\right)\right|\cdot\left|\theta-\theta'\right|.$$

□

2.8.3 Computational Details

We start with a pseudo-code for the algorithm that solves for a steady-state equilibrium. It has three key steps. These steps are described in detail after the pseudo-code.

Algorithm 1 Steady State Equilibrium

Inputs: Patient and donor characteristics, scoring rule s, parameters Γ, δ, ρ, and patient age grid $\{t_0, \ldots, t_L = T\}$. Let $t_{l_x^0}$ be the arrival time for patient of type x.
Outputs: V^*, π^*, m^*, N^*
Initialize $k = 0$ and beliefs $\pi_x^k(t)$ for all x and $t \in \{t_0, \ldots, T\}$
repeat
    $V^k \leftarrow$ Backwards Induction(π^k,κ^k)
    $\kappa_x^k(t_l) \leftarrow \delta_x(t_l) + \lambda \sum_z \pi_{x,z}^k(t_l)\mathbb{P}(\Gamma(t_l; x, z) + \varepsilon > V_x(t_l))$
    $m^k \leftarrow$ Forward Simulation(κ^k) ▷ Waitlist Composition
    $N^k \leftarrow \lceil \sum_x \lambda_x / \sum_{x,t_l} m_x^k(t_l)\kappa_x^k(t_l) \rceil$ ▷ Waitlist Size: Definition 1, part 3(b)
    $\pi^k \leftarrow$ Compute Offer Probabilites(V^k, m^k, N^k) ▷ Offer Probabilities
    $k \leftarrow k + 1$
until $k > 1$, $\|V^k - V^{k-1}\|_\infty < \epsilon$, $\|m^k - m^{k-1}\|_\infty < \epsilon$, and $N^k = N^{k-1}$ ▷ Convergence
$V^* \leftarrow V^k, m^* \leftarrow m^k, N^* \leftarrow N^k, \pi^* \leftarrow \pi^k$
function BACKWARDS INDUCTION(π,κ)
    for all x **do**
        Set $V_x(T) = 0$
        Compute $w_x(T)$ using equation (2.14)
        for all x and $t_l = T$ to $t_{l_x^0}$ **do**
            Compute $V_x^k(t_l)$ using equation (2.15)
            if $t_l < T$ **then**
                Compute $w_x(t_l)$ using equation (2.14)
            end if
        end for
    end for
    return $V_x(t_l)$ for all x and $t_l \in \{t_{l_x^0}, \ldots, T\}$
end function
function FORWARD SIMULATION(κ)
    for all x **do**
        $m_x(t_{l_x^0}) \leftarrow \lambda_x$
        for all $t_l = t_{l_x^0+1}$ to T **do**
            $m_x(t_{l+1}) \leftarrow m_x(t_l)\exp(-\kappa_x(t_l)(t_{l+1} - t_l))$
        end for
    end for
    $m_x(t_l) \leftarrow m_x(t_l) / \sum_{x,\tau \in \{t_0,\ldots,t_L\}} m_x(\tau)$ for all x and t_l
    return $m_x(t_l)$ for $t_l \in \{t_{l_x^0}, \ldots, T\}$
end function
function COMPUTE OFFER PROBABILITIES(m^k, V^k, N^k)
    $p^a(t_l; x, z) \leftarrow \mathbb{P}(\Gamma(t_l; x, z) + \varepsilon > V_x(t_l))$ for all x, t_l
    for all $s = \max s(t_l; x, z)$ to $\min s(t_l; x, z)$ **do**
        Compute π using equation (2.17)
    end for
    return π^k
end function

Value Function Computation (Backwards Induction):

To ease notation, define

$$w_x^k(t) = \int \pi_x^k(t; z)\, \mathbb{E} \max\{0, \Gamma(t; x, z) + \varepsilon - V_x(t)\}\, \mathrm{d}F. \tag{2.14}$$

Conditional on continuing at time t, the duration until the arrival of the next organ in the list is an exponential random variable with parameter λ. Therefore, for any $h > 0$, we have that

$$V_x(t) = \int_0^h \lambda e^{-\lambda\tau} \left\{ e^{-(\rho+\delta_x(t+\tau))\tau} \left[w_x(t+\tau) + V_x(t+\tau) \right] \right\} d\tau + V_x(t+h)\, e^{-\int_0^h (\rho+\delta_x(t+\tau)+\lambda) d\tau}.$$

The first term represents the expected value conditional on the donor arriving between t and $t+h$, and thesecond term represents the value conditional on no arrival. For a small h and $\tau \in (0, h)$, approximate $\delta_x(t+\tau) = \delta_x(t)$, $w_x(t+\tau) = w_x(t)$, and $V_x(t+\tau) = V_x(t+h)$. We get that

$$\begin{aligned}
V_x(t) &\approx \lambda \int_0^h e^{-(\rho+\delta_x(t)+\lambda)\tau} \left[w_x(t) + V_x(t+h) \right] d\tau + e^{-(\rho+\delta_x(t)+\lambda)h} V_x(t+h) \\
&= \frac{\lambda \left(1 - e^{-[\rho+\delta_x(t)+\lambda]h}\right)}{\rho + \delta_x(t) + \lambda} w_x(t) + \left[\frac{\lambda \left(1 - e^{-[\rho+\delta_x(t)+\lambda]h}\right)}{\rho + \delta_i^* + \lambda} + e^{-(\rho+\delta_x(t)+\lambda)h} \right] V_x(t+h) \\
\implies V_x(t_l) &\approx \frac{\lambda \left(1 - e^{-(\rho+\delta_x(t_l)+\lambda)(t_{l+1}-t_l)}\right)}{\rho + \delta_x(t_l) + \lambda} w_x(t_l) + \left[\frac{\lambda + (\rho + \delta_x(t_l))\, e^{-(\rho+\delta_x(t_l)+\lambda)(t_{l+1}-t_l)}}{\rho + \delta_x(t_l) + \lambda} \right] V_x(t_{l+1}).
\end{aligned} \tag{2.15}$$

Offer Probabilities, $\pi_{x,z}(t)$:

In what follows, we fix a particular agent i with priority score s. Ties are broken randomly, and it is therefore without loss of generality to consider each agent's tiebreaker to be drawn from a uniform distribution on the unit interval. Let $1 - \alpha$

be the tie-breaker for agent i.

An agent receives an offer if i) the total number of acceptances by agents that are ranked higher than agent i is strictly less than the number of copies of the object available, and ii) the total number of offers that can be made for the object is larger than the total number of agents that are ranked higher than agent i. The first of these is governed by the acceptance behavior of agents and the priority rule. The second is governed by the model specified in equation (2.9). This model is equivalent to an exponential hazards model with equal probability of failure before the next offer is made.

We compute the probability of accepting an offer by considering waitlists that are composed of N agents drawn randomly drawn with distribution governed by m. For each agent drawn from m, consider two mutually exclusive events:

1. Event F: The agent drawn is ordered above i and an offer cannot be made because the a failure occurs as in equation (2.9). The probability of this event is given by

$$p^F(s, \alpha) = (m_H(s) + \alpha m_L(s)) p_0,$$

   where p_0 is the probability of a failure ocurring, $m_H(s) = \sum_{t,x} m(t; x) 1\{s(t; x) > s\}$ is the probability an agent with a higher priority (group H) is drawn, $m_E(s)$ is the probability an agent with a priority equal to s (group E) is drawn, and α represents the probability that an agent with an equal score has a higher draw of the tie-breaker.

2. Event A: The agent drawn is ordered above i, an additional offer can be made because the maximum number of offers is not exceeded (there is no failure as in equation (2.9)), and the drawn agent accepts the offer. The probability of

this event is given by

$$p^A(s,\alpha) = m_H(s)\,p_H^A(s)\,(1-p_0) + m_E(s)\,\alpha p_E^A(s)\,(1-p_0)\,.$$

The first term represents the case when an agent with a higher priority (group H) is drawn. The probability of acceptance by such an agent is given by

$$p_H^A(s) = \frac{1}{m_H(s)} \sum_{t,x} m(t;x)\,1\{s(t;x) > s\}\,\mathbb{P}\left(\Gamma(t;x) + \varepsilon > V_x(t)\right).$$

The second term represents the probability that an agent with priority score s is drawn. The term $p_E^A(s)$ is defined analagously as $p_H^A(s)$.

If either of these two events occur, then agent i's chances of receiving an offer are reduced. In the first case, one less object is made available. In the second case, agent i can no longer receive an offer. In all other cases, the chances that i receives an offer are not diminished by the agent just drawn.

These two events define a multinomial random variable with parameter $\left(p^F(s,\alpha), p^a(s,\alpha)\right)$. Let $\boldsymbol{X} = \left(X^F, X^A\right)$ be a random variable that denotes the total number of events of type F and A from N trials.

In this notation, an object is available if $X^F = 0$ and $X^A < q$, where q is the total number of copies of the object. In addition, for i to receive an offer, it must be the case that there is no failure right when i is approached. Hence, the probability that i receives an offer is given by

$$p_0 \int_0^1 \mathbb{P}\left(X^F = 0, X^A < q \middle| s, \alpha\right) \mathrm{d}\alpha, \tag{2.16}$$

where we have integrated over the tie-breaker α, and explicit conditioning on N is subsumed for simplicity.

The term $\mathbb{P}\left(X^F = 0, X^A < q|s, \alpha\right)$ is cumbersome as it depends on the number of draws of each type from a multinomial distribution. We approximate this term for large N and small p^A and p^F using Theorem 1 in McDonald (1980). Specifically, define two independent Poisson random variables Y^F and Y^A with parameters $\left(Np^F(s, \alpha), Np^a(s, \alpha)\right)$. Theorem 1 in McDonald (1980) implies that:

$$\sum_{z \in \mathbb{N}^2} |\mathbb{P}(\boldsymbol{X} = z) - \mathbb{P}(\boldsymbol{Y} = z)| \leq 2N\left(p^F(s, \alpha) + p^A(s, \alpha)\right)^2.$$

Therefore, the expression in equation (2.16) yields $\pi_{x,z}(t)$:

$$\pi_{x,z}(t) = p_0 \int_0^1 \mathbb{P}\left(Y^F = 0|s(t; x, z), \alpha\right) \mathbb{P}\left(Y^A < q|s(t; x, z), \alpha\right) \mathrm{d}\alpha, \qquad (2.17)$$

where

$$\mathbb{P}\left(Y^A < q|\alpha, s\right) = \sum_{q' < q} \frac{e^{-Np^A(s,\alpha)}\left(Np^A(s, \alpha)\right)^{q'}}{q'!}$$

and

$$\mathbb{P}\left(Y^F = 0|\alpha, s\right) = e^{-Np^F(s,\alpha)}.$$

The convergence of the integral in equation (2.16) by equation (2.17) is guaranteed by the dominated convergence theorem.

The expression in equation (2.17) can be simplified and solved for analytically. We use that solution in our algorithm.

Waitlist Size/Composition (Forward Simulation), m, N**:**

We use $\kappa_x(t)$ and λ_x to update the queue composition. Solving the ODE in

Definition 1, part 3(a), we get that for any $h > 0$,

$$m_x(t+h) = m_x(t)\exp\left(-\int_0^h k_x(t+\tau)\,\mathrm{d}\tau\right),$$

where $m_x(0) \propto \lambda_x$. Appoximating $\kappa_x(t+\tau) = \kappa_x(t+h)$ for all $\tau \in (0, h)$, we have that

$$m_x(t_{l+1}) = m_x(t_l)\exp\left(-\kappa_x(t_{l+1})(t_{l+1} - t_l)\right).$$

Finally, we scale the output so that $m_x(t_l)$ is a probability measure.

The size of the waitlist, N, is determined by part 3(b) of Definition 1.

Table 2.10: Evidence of Response to Dynamic Incentives

	Dependent Variable: Offer Accepted						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Calculated Panel Reactive Antibodies (CPRA)	0.0155***	0.00879***	0.00838***	0.00773***	0.00783***	0.00815***	0.0709***
	(0.000769)	(0.000893)	(0.000884)	(0.000836)	(0.000838)	(0.000866)	(0.0131)
Variables Affecting Priority		X	X	X	X	X	X
Patient Characteristics			X	X	X	X	X
Donor and Match Characteristics				X	X	X	X
Interaction between CPRA and # HLA Mismatches					X		
Excluding Zero HLA Mismatch Offers						X	
Restricting to Top 10 Offers							X
Mean Acceptance Rate	0.150%	0.150%	0.150%	0.150%	0.150%	0.146%	4.962%
Observations	2840937	2840937	2840937	2840937	2840937	2840048	45723
R-squared	0.003	0.006	0.009	0.099	0.099	0.016	0.266

Notes: Estimates from a linear probability model of offer acceptance on patient Calculated Panel Reactive Antibodies (CPRA). CPRA is the recorded CPRA, on a [0,1] scale, at the time the offer was made. Column 1 controls only for a CPRA=0 indicator. Column 2 adds controls affecting patient priority: indicators for CPRA>=0.2 and CPRA>=0.8; waiting time indicators and linear controls for 1-3, 3-5, and >5 years; and an indicator for patient age < 18. Column 3 adds controls for patient characteristics, and Column 4 adds controls for donor and match characteristics (detailed below). Column 5 adds interactions between CPRA indicators and # HLA mismatches. Column 6 restricts to offers that are not perfect tissue type matches. Column 7 restricts to the first ten offers for each donor. Patient controls are indicators for age 18-35, 35-50, and 50-65; linear controls and indicators for dialysis time 1-3, 3-5, 5-10, and >10 years; blood type and diabetes indicators; and health status at listing. Donor controls are linear age; linear creatinine with indicators for 0.6-1.8 and >1.8; and indicators for diabetes, cardiac death (DCD), and expanded criteria donor (ECD). Other controls are predicted Life Years from Transplantation (LYFT); linear # HLA mismatches; indicators for zero HLA mismatch, 0 and 1 DR mismatch, identical blood type, offer year, and local donor; linear controls for (+) and (-) age difference; and interactions between local and zero-HLA mismatch, and local and donor age. Finally, separate coefficients on donor age are estimated for: donor over 40, pediatric patient; donor over 55, patient age 18-35; and donor over 60, patient age 35-50 and over 50. Standard errors, clustered by donor, are in parentheses. *** Significant at 0.1% ** Significant at 1% * Significant at 5%

Table 2.11: Cutoff Rank Autocorrelation Tests

Testing for Serial Correlation in Donor Cutoffs Donor Type	N	Autocorrelation Statistic	p-value
Panel A: Donor Category			
Standard Criteria, Age < 35	1065	2.01	0.55
Standard Criteria, Age >= 35	1598	1.99	0.46
Cardiac Death, Age < 35	178	2.02	0.56
Cardiac Death, Age >= 35	404	2.01	0.51
Expanded Criteria	2462	1.95	0.10
Expanded Criteria or Cardiac Death	156	2.12	0.78
Panel B: Donor Category and Origin			
Standard Criteria, Age < 35; NYRT	234	2.09	0.74
Standard Criteria, Age < 35; non-NYRT	831	1.91	0.08
Standard Criteria, Age >= 35; NYRT	252	1.99	0.49
Standard Criteria, Age >= 35; non-NYRT	1346	2.01	0.60
Cardiac Death, Age < 35; NYRT	15	1.91	0.43
Cardiac Death, Age < 35; non-NYRT	163	2.25	0.95
Cardiac Death, Age >= 35; NYRT	50	2.51	0.97
Cardiac Death, Age >= 35; non-NYRT	354	2.08	0.77
Expanded Criteria; NYRT	227	2.16	0.88
Expanded Criteria; non-NYRT	2235	1.99	0.41
Expanded Criteria or Cardiac Death; NYRT	6	2.63	0.83
Expanded Criteria or Cardiac Death; non-NYRT	150	2.16	0.83
Panel C: Standard Criteria NYRT Donors, by Age and Blood Type			
Standard Criteria, Age < 35; Blood Type O	117	1.98	0.45
Standard Criteria, Age < 35; Blood Type B	37	1.70	0.20
Standard Criteria, Age < 35; Blood Type AB	5	1.50	0.25
Standard Criteria, Age < 35; Blood Type A	75	1.74	0.11
Standard Criteria, Age >= 35; Blood Type O	131	1.87	0.24
Standard Criteria, Age >= 35; Blood Type B	32	2.00	0.49
Standard Criteria, Age >= 35; Blood Type AB	7	1.86	0.45
Standard Criteria, Age >= 35; Blood Type A	82	1.98	0.46

Notes: results from tests for autocorrelation of donor cutoffs in the NYRT sample based on the rank version of von Neumann's ratio statistic (Bartels, 1982). Each donor's cutoff is the priority score above which a patient would have received an offer from that donor, which is determined by the last patient in the donor's offer sequence. The rank of each donor's cutoff is its order statistic among the cutoffs of donors of the same type, with ties broken by a random number. The autocorrelation statistic for each donor type is computed for the observed sequence of donor cutoff ranks. Each p-value is the fraction of 1,000 randomly sampled permutations of donor arrival sequences for which the rank autocorrelation statistic is below that observed in sample.

Table 2.12: Conditional Choice Probability Estimates (Detailed)

	Conditional Choice Probability of Accepting an Offer					
	Base Specification		Unobserved Heterog.		Waiting Time + UH	
	(1)		(2)		(3)	
Constant	-3.65	(0.02)	-4.44	(0.04)	-4.52	(0.04)
Patient Diabetic	-0.05	(0.02)	-0.02	(0.02)	0.00	(0.02)
Calculated Panel Reactive Antibody (CPRA)	0.32	(0.07)	0.24	(0.09)	0.09	(0.13)
CPRA >= 80%	0.36	(0.05)	0.19	(0.07)	0.14	(0.11)
CPRA = 0%	0.01	(0.03)	0.12	(0.04)	0.10	(0.04)
Patient had Prior Transplant	0.08	(0.02)	-0.07	(0.03)	-0.13	(0.04)
Log Years on Dialysis at Registration	-0.08	(0.01)	-0.09	(0.01)	-0.09	(0.01)
Donor Age < 18	0.46	(0.13)	0.06	(0.21)	-0.04	(0.20)
Donor Age 18-35	0.76	(0.17)	0.09	(0.24)	-0.09	(0.27)
Donor Age 50+	-1.13	(0.35)	-0.64	(0.43)	-0.58	(0.47)
Donor Cause of Death Anoxia	-0.02	(0.02)	-0.09	(0.06)	-0.07	(0.07)
Donor Cause of Death Stroke	0.00	(0.02)	0.04	(0.07)	0.03	(0.06)
Donor Cause of Death CNS	0.11	(0.10)	-0.24	(0.29)	-0.20	(0.37)
Donor Creatinine 0.5-1.0	-0.02	(0.04)	0.05	(0.10)	0.04	(0.11)
Donor Creatinine 1.0-1.5	-0.01	(0.04)	0.04	(0.11)	0.03	(0.12)
Donor Creatinine >= 1.5	-0.09	(0.04)	-0.16	(0.10)	-0.17	(0.12)
Donor Pancreas Offered	0.30	(0.03)	0.44	(0.09)	0.48	(0.10)
Patient awaits Pancreas	-2.42	(1.03)	-10.35	(0.89)	-5.97	(1.52)
Expanded Criteria Donor (ECD)	-0.17	(0.02)	-0.50	(0.07)	-0.53	(0.08)
Donation from Cardiac Death (DCD)	-0.12	(0.03)	-0.44	(0.08)	-0.52	(0.07)
Donor Male	0.03	(0.02)	0.05	(0.04)	0.06	(0.04)
Donor History of Hypertension	-0.01	(0.02)	-0.02	(0.06)	-0.02	(0.06)
Perfect Tissue Type Match	2.45	(0.32)	2.96	(0.44)	2.95	(0.46)
2 A Mismatches	-0.02	(0.02)	-0.02	(0.02)	-0.03	(0.02)
2 B Mismatches	0.01	(0.02)	-0.03	(0.03)	-0.03	(0.03)
2 DR Mismatches	-0.09	(0.02)	-0.10	(0.02)	-0.10	(0.02)
ABO Compatible	-0.40	(0.06)	-0.41	(0.08)	-0.46	(0.09)
Regional Offer	-1.00	(0.06)	-2.22	(0.17)	-2.42	(0.19)
National Offer	-1.15	(0.05)	-2.43	(0.11)	-2.64	(0.12)
Non-NYRT Donor, NYRT Match Run	0.94	(0.02)	1.78	(0.05)	1.94	(0.06)
Patient Blood Type A	-0.22	(0.02)	-0.30	(0.06)	-0.34	(0.07)
Patient Blood Type O	-0.42	(0.02)	-0.44	(0.06)	-0.42	(0.06)
Patient on Dialysis at Registration	-0.50	(0.05)	-0.68	(0.06)	-0.68	(0.06)
Patient Age	0.02	(0.01)	0.05	(0.01)	0.05	(0.01)
Patient Age - 18 if Age >= 18	-0.02	(0.01)	-0.04	(0.01)	-0.04	(0.01)
Patient Age - 35 if Age >= 35	0.01	(0.01)	0.00	(0.01)	-0.01	(0.01)
Patient Age - 50 if Age >= 50	0.00	(0.00)	0.00	(0.00)	0.00	(0.00)
Patient Age - 65 if Age >= 65	-0.01	(0.00)	-0.01	(0.00)	-0.01	(0.00)
Log Waiting Time (years)					-0.18	(0.06)
Log Waiting Time * Over 1 Year					-0.02	(0.07)
Log Waiting Time * Over 2 Years					-0.24	(0.13)
Log Waiting Time * Over 3 Years					0.17	(0.12)
Patient BMI at Departure	-0.01	(0.03)	-0.03	(0.04)	-0.02	(0.04)
Patient BMI - 18.5 if BMI >= 18.5	0.01	(0.03)	0.04	(0.04)	0.02	(0.04)
Patient BMI - 25 if BMI >= 25	0.00	(0.01)	-0.01	(0.01)	-0.01	(0.01)
Patient BMI - 30 if BMI >= 30	0.00	(0.01)	0.00	(0.01)	0.00	(0.01)
Patient Serum Albumin	-0.06	(0.03)	-0.08	(0.04)	-0.08	(0.04)
Serum Albumin - 3.7 if >= 3.7	0.02	(0.05)	0.03	(0.06)	0.00	(0.07)
Serum Albumin - 4.4 if >= 4.4	0.04	(0.05)	0.09	(0.06)	0.10	(0.07)
Perfect Tissue Type Match * Prior Transplant	-0.16	(0.19)	0.03	(0.27)	0.05	(0.28)
Perfect Tissue Type Match * Diabetic Patient	0.01	(0.16)	0.05	(0.22)	0.06	(0.23)
Perfect Tissue Type Match * Patient Age	-0.01	(0.01)	-0.02	(0.01)	-0.02	(0.01)
Perfect Tissue Type Match * CPRA	1.25	(0.34)	1.78	(0.48)	2.01	(0.50)
Perfect Tissue Type Match * CPRA above 80%	-0.76	(0.29)	-0.54	(0.41)	-0.51	(0.43)
Perfect Tissue Type Match * ECD Donor	-0.52	(0.16)	-0.59	(0.24)	-0.63	(0.24)
Perfect Tissue Type Match * DCD Donor	-0.42	(0.32)	-1.08	(0.46)	-1.12	(0.47)
Perfect Tissue Type Match * NYRT Donor	0.64	(0.18)	0.17	(0.26)	0.17	(0.27)

Perfect Tissue Type Match * ABO Compatible	-0.02	(0.17)	0.05	(0.24)	0.09	(0.25)
Donor Pancreas Offered * Patient awaits Pancreas	2.48	(1.12)	10.53	(1.04)	6.12	(1.60)
NYRT Donor * 2 A Mismatches	0.06	(0.03)	0.03	(0.04)	0.03	(0.04)
NYRT Donor * 2 B Mismatches	0.02	(0.03)	0.00	(0.05)	-0.01	(0.05)
NYRT Donor * 2 DR Mismatches	-0.04	(0.03)	-0.02	(0.04)	-0.01	(0.04)
NYRT Donor * Donor Age < 18	-0.01	(0.07)	0.35	(0.20)	0.37	(0.22)
NYRT Donor * Donor Age 18-35	0.12	(0.05)	0.28	(0.13)	0.27	(0.17)
NYRT Donor * Donor Age 50+	-0.26	(0.04)	-0.38	(0.11)	-0.45	(0.12)
Patient Age * Donor Age < 18	-0.01	(0.00)	-0.01	(0.00)	0.00	(0.00)
Patient Age * Donor Age 18-35	-0.02	(0.01)	0.00	(0.01)	0.01	(0.01)
Patient Age * Donor Age 50+	0.03	(0.01)	0.01	(0.01)	0.01	(0.01)
Patient Age - 35 if Age >= 35 * Donor Age 18-35	0.02	(0.01)	0.00	(0.01)	-0.01	(0.01)
Patient Age - 35 if Age >= 35 * Donor Age 50+	-0.01	(0.01)	0.01	(0.01)	0.01	(0.01)
Log Waiting Time * Prior Transplant					0.03	(0.02)
Log Waiting Time * Patient Diabetic					-0.03	(0.02)
Log Waiting Time * Patient Age					0.00	(0.00)
Log Waiting Time * CPRA					0.13	(0.08)
Log Waiting Time * CPRA >= 80					0.02	(0.08)
Log Waiting Time * Patient Serum Albumin					0.01	(0.01)
Log Waiting Time * Patient BMI at Departure					0.01	(0.00)
Log Waiting Time * Patient Blood Type A					0.01	(0.03)
Log Waiting Time * Patient Blood Type O					-0.02	(0.03)
Patient BMI Missing	-0.35	(0.54)	-0.90	(0.68)	-0.50	(0.70)
Patient Serum Albumin Missing	-0.15	(0.11)	-0.31	(0.13)	-0.24	(0.14)
Donor Unobservable Std. Dev.			1.03	(0.23)	1.25	(0.24)
Idiosyncratic Shock Std. Dev.	1.00		1.00		1.00	
Acceptance Rate	0.150%		0.150%		0.150%	
Number of Offers	2840937		2840937		2840937	
Number of Donors	5863		5863		5863	
Number of Patients	9917		9917		9917	

Table 2.12: Value of Transplant Estimates (Detailed)

	Estimated Value of a Transplant		
	Base Specification	Unobserved Heterog.	Waiting Time + UH
	(1)	(2)	(3)
Constant	-3.24	-1.32	-1.45
Patient Diabetic	-0.30	-1.19	-1.01
Calculated Panel Reactive Antibody (CPRA)	2.37	5.45	4.02
CPRA >= 80%	0.65	-3.32	-1.71
CPRA = 0%	-0.28	-1.12	-0.82
Patient had Prior Transplant	1.51	3.63	4.77
Log Years on Dialysis at Registration	-0.14	-0.31	-0.30
Donor Age < 18	0.47	0.49	0.49
Donor Age 18-35	0.83	0.53	0.83
Donor Age 50+	-1.10	-1.10	-1.35
Donor Cause of Death Anoxia	-0.03	-0.09	-0.08
Donor Cause of Death Stroke	0.01	0.02	0.02
Donor Cause of Death CNS	0.18	-0.05	-0.09
Donor Creatinine 0.5-1.0	-0.06	-0.03	0.00
Donor Creatinine 1.0-1.5	0.02	-0.04	-0.03
Donor Creatinine >= 1.5	-0.13	-0.25	-0.23
Donor Pancreas Offered	0.38	0.76	0.53
Patient awaits Pancreas	-2.80	-11.77	-7.72
Expanded Criteria Donor (ECD)	-0.15	-0.51	-0.53
Donation from Cardiac Death (DCD)	-0.11	-0.47	-0.51
Donor Male	0.00	0.08	0.07
Donor History of Hypertension	0.01	-0.01	-0.01
Perfect Tissue Type Match	2.38	2.45	2.73
2 A Mismatches	-0.07	0.05	0.02
2 B Mismatches	0.07	0.07	0.05
2 DR Mismatches	-0.06	-0.01	-0.05
ABO Compatible	-0.32	1.50	1.91
Regional Offer	-1.36	-3.02	-2.80
National Offer	-1.52	-3.24	-3.02
Non-NYRT Donor, NYRT Match Run	1.26	2.64	2.03
Patient Blood Type A	-0.03	0.99	1.29
Patient Blood Type O	-0.23	0.24	0.15
Patient on Dialysis at Registration	-0.85	-2.19	-2.25
Patient Age	0.02	0.15	0.18
Patient Age - 18 if Age >= 18	-0.05	-0.20	-0.29
Patient Age - 35 if Age >= 35	0.04	0.08	0.15
Patient Age - 50 if Age >= 50	-0.01	-0.06	-0.06
Patient Age - 65 if Age >= 65	-0.03	-0.11	-0.09
Log Waiting Time (years)			-0.33
Log Waiting Time * Over 1 Year			0.83
Log Waiting Time * Over 2 Years			-0.01
Log Waiting Time * Over 3 Years			4.54
Patient BMI at Departure	-0.15	-0.63	-0.57
Patient BMI - 18.5 if BMI >= 18.5	0.11	0.54	0.54
Patient BMI - 25 if BMI >= 25	0.07	0.20	0.14
Patient BMI - 30 if BMI >= 30	-0.02	-0.12	-0.12
Patient Serum Albumin	0.01	0.54	0.55
Serum Albumin - 3.7 if >= 3.7	-0.05	-0.62	-0.59

Serum Albumin - 4.4 if >= 4.4	0.17	1.00	0.89
Perfect Tissue Type Match * Prior Transplant	-0.12	-0.22	-0.60
Perfect Tissue Type Match * Diabetic Patient	-0.02	0.35	0.10
Perfect Tissue Type Match * Patient Age	-0.01	-0.02	-0.02
Perfect Tissue Type Match * CPRA	0.35	-0.53	0.95
Perfect Tissue Type Match * CPRA above 80%	-1.40	-0.77	-2.91
Perfect Tissue Type Match * ECD Donor	-0.60	-0.66	-0.82
Perfect Tissue Type Match * DCD Donor	-0.61	-1.31	-1.40
Perfect Tissue Type Match * NYRT Donor	0.56	0.18	0.23
Perfect Tissue Type Match * ABO Compatible	0.01	-1.48	-1.83
Donor Pancreas Offered * Patient awaits Pancreas	2.49	10.12	6.07
NYRT Donor * 2 A Mismatches	0.15	0.09	0.07
NYRT Donor * 2 B Mismatches	-0.02	-0.08	-0.06
NYRT Donor * 2 DR Mismatches	-0.02	0.00	0.02
NYRT Donor * Donor Age < 18	-0.07	0.02	0.17
NYRT Donor * Donor Age 18-35	0.10	0.17	0.18
NYRT Donor * Donor Age 50+	-0.42	-0.77	-0.66
Patient Age * Donor Age < 18	-0.01	-0.01	-0.01
Patient Age * Donor Age 18-35	-0.02	-0.01	-0.02
Patient Age * Donor Age 50+	0.03	0.02	0.03
Patient Age - 35 if Age >= 35 * Donor Age 18-35	0.02	0.00	0.01
Patient Age - 35 if Age >= 35 * Donor Age 50+	-0.01	-0.01	-0.02
Log Waiting Time * Prior Transplant			2.07
Log Waiting Time * Patient Diabetic			-0.30
Log Waiting Time * Patient Age			0.00
Log Waiting Time * CPRA			2.91
Log Waiting Time * CPRA >= 80			-3.23
Log Waiting Time * Patient Serum Albumin			-0.06
Log Waiting Time * Patient BMI at Departure			0.01
Log Waiting Time * Patient Blood Type A			0.28
Log Waiting Time * Patient Blood Type O			0.18
Patient BMI Missing	-3.01	-12.64	-12.28
Patient Serum Albumin Missing	0.43	2.81	1.97
Donor Unobservable Std. Dev.		1.03	1.25
Idiosyncratic Shock Std. Dev.	1.00	1.00	1.00

Table 2.12: Net Present Value Estimates (Detailed)

	Base Specification		Value Function Unobserved Heterog.		Waiting Time + UH	
	(1)		(2)		(3)	
Constant	0.46	(0.00)	3.14	(0.00)	3	(0.00)
Patient Diabetic	-0.19	(0.00)	-1.11	(0.00)	-0.94	(0.00)
Calculated Panel Reactive Antibody (CPRA)	1.76	(0.00)	4.75	(0.01)	3.46	(0.01)
CPRA >= 80%	0.42	(0.00)	-3.38	(0.01)	-1.81	(0.01)
CPRA = 0%	-0.17	(0.00)	-1.1	(0.00)	-0.79	(0.00)
Patient had Prior Transplant	1.08	(0.00)	3.26	(0.00)	4.59	(0.00)
Log Years on Dialysis at Registration	-0.06	(0.00)	-0.22	(0.00)	-0.21	(0.00)
Patient awaits Pancreas	-0.18	(0.01)	-1.39	(0.01)	-1.62	(0.01)
Patient Blood Type A	0.13	(0.00)	0.72	(0.00)	0.8	(0.00)
Patient Blood Type O	0.06	(0.00)	0.08	(0.00)	-0.24	(0.00)
Patient on Dialysis at Registration	-0.33	(0.00)	-1.53	(0.00)	-1.6	(0.00)
Patient Age	-0.01	(0.00)	0.06	(0.00)	0.09	(0.00)
Patient Age - 18 if Age >= 18	-0.01	(0.00)	-0.11	(0.00)	-0.2	(0.00)
Patient Age - 35 if Age >= 35	0.03	(0.00)	0.08	(0.00)	0.14	(0.00)
Patient Age - 50 if Age >= 50	-0.01	(0.00)	-0.06	(0.00)	-0.06	(0.00)
Patient Age - 65 if Age >= 65	-0.02	(0.00)	-0.1	(0.00)	-0.08	(0.00)
Log Waiting Time (years)					-0.51	(0.01)
Log Waiting Time * Over 1 Year					0.93	(0.00)
Log Waiting Time * Over 2 Years					0.05	(0.01)
Log Waiting Time * Over 3 Years					4.31	(0.01)
Patient BMI at Departure	-0.1	(0.00)	-0.54	(0.00)	-0.5	(0.00)
Patient BMI - 18.5 if BMI >= 18.5	0.07	(0.00)	0.46	(0.00)	0.48	(0.00)
Patient BMI - 25 if BMI >= 25	0.04	(0.00)	0.18	(0.00)	0.12	(0.00)
Patient BMI - 30 if BMI >= 30	-0.02	(0.00)	-0.1	(0.00)	-0.11	(0.00)
Patient Serum Albumin	0.04	(0.00)	0.56	(0.00)	0.57	(0.00)
Serum Albumin - 3.7 if >= 3.7	-0.05	(0.00)	-0.63	(0.01)	-0.59	(0.01)
Serum Albumin - 4.4 if >= 4.4	0.11	(0.00)	0.92	(0.01)	0.81	(0.01)
Log Waiting Time * Prior Transplant					1.88	(0.00)
Log Waiting Time * Patient Diabetic					-0.27	(0.00)
Log Waiting Time * Patient Age					0	(0.00)
Log Waiting Time * CPRA					2.81	(0.01)
Log Waiting Time * CPRA >= 80					-3.23	(0.01)
Log Waiting Time * Patient Serum Albumin					-0.06	(0.00)
Log Waiting Time * Patient BMI at Departure					0.01	(0.00)
Log Waiting Time * Patient Blood Type A					0.55	(0.00)
Log Waiting Time * Patient Blood Type O					0.44	(0.00)
Patient BMI Missing	-2.02	(0.03)	-10.99	(0.07)	-11	(0.07)
Patient Serum Albumin Missing	0.4	(0.01)	2.85	(0.01)	2.02	(0.01)
R-Squared	0.69		0.46		0.77	

Table 2.13: Positive Crossmatch Model

Dependent Variable: Positive Crossmatch	
CPRA	1.605
	(0.0667)
0 or 1 HLA Mismatches	-1.259
	(0.416)
2 or 3 HLA Mismatches	0.131
	(0.0798)
0 DR Mismatches	-0.512
	(0.0873)
CPRA * (0 or 1 HLA Mismatches)	-0.600
	(0.575)
CPRA * (2 or 3 HLA Mismatches)	-0.578
	(0.159)
Constant	-0.459
	(0.0267)
Observations	4283

Notes: coefficient estimates from a probit regression of positive crossmatch on patient CPRA and the number of HLA mismatches. The sample is all offers accepted by NYRT patients between 2010 and 2013. Positive crossmatches are identified by the appropriate refusal code in the PTR data. CPRA is measured on a [0,1] scale.

Table 2.14: Survival Model Estimates for LYFT Calculation

	Survival Model Coefficient Estimates					
	Waiting List		Functioning Graft		After Transplantation	
	(1)		(2)		(3)	
Patient Diabetic	0.5990	(0.009)	0.2910	(0.0102)	0.5220	(0.0118)
Calculated Panel Reactive Antibodies (CPRA)	0.0063	(0.0239)	0.2050	(0.0254)	0.1390	(0.0322)
CPRA = 0%	-0.0432	(0.0141)	-0.0184	(0.0114)	-0.0065	(0.0139)
Patient Age	-0.0156	(0.0066)	0.0289	(0.0041)	0.0522	(0.0089)
Patient Age - 18 if Age >= 18	0.0505	(0.0076)	-0.0578	(0.0048)	-0.0296	(0.0101)
Patient Age - 35 if Age >= 35	-0.0080	(0.0029)	0.0318	(0.0026)	0.0171	(0.0042)
Patient Age - 35 if Age >= 35	0.0007	(0.0019)	0.0199	(0.0019)	0.0024	(0.0024)
Patient Age - 35 if Age >= 35	-0.0016	(0.003)	0.0101	(0.0031)	0.0054	(0.0033)
Patient had Prior Transplant	0.1560	(0.0115)	0.1670	(0.0141)	0.1990	(0.0182)
Patient BMI at Departure	-0.0540	(0.0175)	-0.0031	(0.0121)	-0.0483	(0.0197)
Patient BMI missing	-0.6640	(0.318)	0.0396	(0.22)	-0.8500	(0.359)
Patient BMI - 18.5 if BMI >= 18.5	0.0050	(0.0188)	0.0074	(0.0133)	0.0402	(0.0212)
Patient BMI - 25 if BMI >= 25	0.0269	(0.0053)	0.0183	(0.0051)	0.0218	(0.0063)
Patient BMI - 30 if BMI >= 30	0.0251	(0.004)	-0.0080	(0.0043)	0.0021	(0.0052)
Patient Serum Albumin	-0.5240	(0.0138)	-0.1820	(0.0223)	-0.2560	(0.0275)
Patient Serum Albumin missing	-1.9060	(0.0471)	-0.6270	(0.0762)	-0.9080	(0.0937)
Serum Albumin - 3.7 if >= 3.7	-0.0550	(0.0328)	-0.0009	(0.0462)	-0.0326	(0.0588)
Serum Albumin - 4.4 if >= 4.4	0.6880	(0.0386)	0.2370	(0.0434)	0.3410	(0.0562)
On Dialysis at Registration	-0.3600	(0.032)	-0.3530	(0.0421)	-0.5120	(0.0553)
Log Years on Dialysis at Registration	0.1150	(0.0049)	0.0810	(0.0066)	0.1060	(0.0087)
Regional Donor			0.0154	(0.0239)	0.0284	(0.0291)
National Donor			0.0119	(0.0215)	0.0158	(0.0263)
Patient Blood Type A	0.0959	(0.0119)	0.0132	(0.0122)	0.0297	(0.015)
Patient Blood Type O	0.0151	(0.0108)	0.0181	(0.0122)	0.0232	(0.015)
ABO Compatible			0.0278	(0.021)	0.0042	(0.0256)
Expanded Criteria Donor (ECD)			0.0974	(0.0171)	0.0649	(0.0208)
Donation after Cardiac Death (DCD)			0.0512	(0.0214)	0.0162	(0.0273)
Male Donor			-0.0522	(0.0089)	-0.0167	(0.011)
Donor Pancreas Offered			-0.0635	(0.0103)	-0.0518	(0.0128)
Donor Creatinine 0.5-1.0			-0.0244	(0.0116)	-0.0131	(0.0143)
Donor Creatinine 1.0-1.5			0.0283	(0.0107)	0.0211	(0.0132)
Donor Creatinine >= 1.5			0.0466	(0.0146)	0.0134	(0.0181)
Donor History of Hypertension			0.1040	(0.0116)	0.0485	(0.0142)
Donor Age < 18			-0.3250	(0.0467)	-0.2020	(0.0705)
Donor Age 18-35			-0.1080	(0.0662)	-0.0039	(0.129)
Donor Age 50+			0.0149	(0.105)	0.0857	(0.195)
Donor Cause of Death Anoxia			0.0279	(0.0138)	0.0084	(0.0173)
Donor Cause of Death Stroke			0.0753	(0.0109)	0.0579	(0.0135)
Donor Cause of Death CNS			-0.0084	(0.0445)	0.0308	(0.0533)
NYRT Donor * Donor Age < 18			0.0474	(0.0275)	0.0400	(0.0339)
NYRT Donor * Donor Age 18-35			0.0331	(0.0235)	0.0362	(0.0289)
NYRT Donor * Donor Age 50+			0.0176	(0.0233)	0.0364	(0.0284)
NYRT Donor * 2 A Mismatches			-0.0399	(0.02)	-0.0248	(0.0245)
NYRT Donor * 2 B Mismatches			-0.0779	(0.0205)	-0.0671	(0.0251)
NYRT Donor * 2 DR Mismatches			-0.0111	(0.0208)	-0.0248	(0.0256)
Perfect Tissue Type Match			-0.3850	(0.06)	-0.2060	(0.0838)
Perfect Tissue Type Match * Diabetic Patient			0.0548	(0.0319)	0.0358	(0.0362)
Perfect Tissue Type Match * Prior Transplant			0.0441	(0.0401)	0.0534	(0.0499)
Perfect Tissue Type Match * Patient Age			0.0043	(0.0011)	0.0021	(0.0015)
Perfect Tissue Type Match * CPRA			0.0311	(0.0503)	0.0632	(0.0621)
Perfect Tissue Type Match * ECD Donor			0.0575	(0.0432)	0.0708	(0.0516)
Perfect Tissue Type Match * DCD Donor			0.1520	(0.097)	0.1770	(0.117)
Perfect Tissue Type Match * NYRT Donor			-0.0478	(0.0434)	-0.0592	(0.0516)
Log Waiting Time (years)			0.0017	(0.0151)	-0.0579	(0.0215)
Log Waiting Time * Over 1 Year			0.1910	(0.0255)	0.2470	(0.0319)
Log Waiting Time * Over 2 Years			-0.0076	(0.038)	0.0044	(0.0479)

Log Waiting Time * Over 3 Years			0.4810	(0.34)	-0.2530	(0.489)
Log Waiting Time * Patient Diabetic			-0.0037	(0.0071)	-0.0004	(0.0082)
Log Waiting Time * Prior Transplant			0.0113	(0.0097)	0.0124	(0.0124)
Log Waiting Time * Patient Age			0.0000	(0.0002)	0.0007	(0.0003)
Log Waiting Time * CPRA			0.0060	(0.0148)	0.0066	(0.0186)
Log Waiting Time * Patient Blood Type A			0.0074	(0.0088)	0.0143	(0.0108)
Log Waiting Time * Patient Blood Type O			-0.0123	(0.0086)	-0.0084	(0.0106)
Log Waiting Time * Patient Serum Albumin			-0.0054	(0.0018)	-0.0073	(0.0023)
Log Waiting Time * Patient BMI at Departure			0.0002	(0.0004)	0.0009	(0.0005)
Donor Pancreas Offered * Patient awaits Pancreas			0.1280	(0.136)	0.4470	(0.15)
Patient Age * Donor Age < 18			0.0037	(0.0009)	0.0021	(0.0013)
Patient Age * Donor Age 18-35			-0.0025	(0.002)	-0.0037	(0.0039)
Patient Age * Donor Age 50+			0.0052	(0.0032)	0.0002	(0.0058)
Patient Age - 35 if Age >= 35 * Donor Age 18-35			0.0038	(0.0026)	0.0036	(0.0044)
Patient Age - 35 if Age >= 35 * Donor Age 50+			-0.0087	(0.0037)	-0.0022	(0.0063)
2 A Mismatches			0.0876	(0.0173)	0.0380	(0.0212)
2 B Mismatches			0.1070	(0.0178)	0.0535	(0.0218)
2 DR Mismatches			0.0937	(0.018)	0.0801	(0.0221)
Patient Registered 1988-1989	--	(.)				
Patient Registered 1990-1994	1.0070	(0.0218)				
Patient Registered 1995-1999	0.8690	(0.0186)				
Patient Registered 2000-2004	0.5210	(0.013)				
Patient Registered 2005-2009	0.2430	(0.0115)				
Transplanted 1988-1989			--	(.)	--	(.)
Transplanted 1990-1994			1.0060	(0.0354)	1.0430	(0.0474)
Transplanted 1995-1999			0.8120	(0.0319)	0.8050	(0.0435)
Transplanted 2000-2004			0.6090	(0.0295)	0.5860	(0.0408)
Transplanted 2005-2009			0.3370	(0.0289)	0.3230	(0.0401)
Observations	519624		153479		154363	

Notes: coefficient estimates from Cox proportional hazard models of patient survival on the waiting list, survival of a functioning graft, and patient survival after transplantation. The sample for waiting list survival is all patients in the UNOS database registered between January 1st, 1988 and December 31st, 2014. The sample for post-transplant and graft survival is patients who received a transplant between January 1st, 1988 and December 31st, 2011. Survival on the waiting list is censored on December 31st, 2014 for patients who received a transplant before or were still alive on that date. Patient and graft survival after transplantation are censored on December 31st, 2011 for patients who received another transplant or if the patient was still alive or the graft remained functional on that date.

Table 2.14: Conversion to Life Years from Transplantation (LYFT) Units

Value of Characteristics for a Representative Patient and Donor			
	Life Years From Transplant (LYFT)	Kidney Utility	Kidney Utility in LYFT Units
Perfect Tissue Type Match	6.28	2.18	2.64
Zero DR Mismatches rather than Two	1.32	0.12	0.15
ABO Identical rather than Compatible	0.16	-1.27	-1.54
Local Donor rather than National	1.36	2.66	3.22
Donor Age 45 rather than 55	2.40	0.61	0.74
Donor not Hypertensive	1.20	0.02	0.02
Not a Donation after Cardiac Death (DCD)	1.88	1.05	1.26
Donor died of Head Trauma, not Anoxia	0.24	0.07	0.09
Donor Creatinine 0.5-1.0 rather than above 1.5	0.69	0.22	0.26

Notes: Reports the value of changing one offer characteristic for a representative patient-donor pair in terms of Life Years from Transplant (LYFT); kidney utility; and kidney utility mapped to LYFT units using the regression coefficient of LYFT on Gamma with patient by year fixed effects, for local offers only. The representative patient is 45 years old, non-diabetic, and unsensitized. The representative donor is 45 years old and from NYRT. The patient and donor are of identical blood type and have 5 tissue type mismatches.

Figure 2-6: Offer and Acceptance Rate by cPRA

(a) Offer Rate

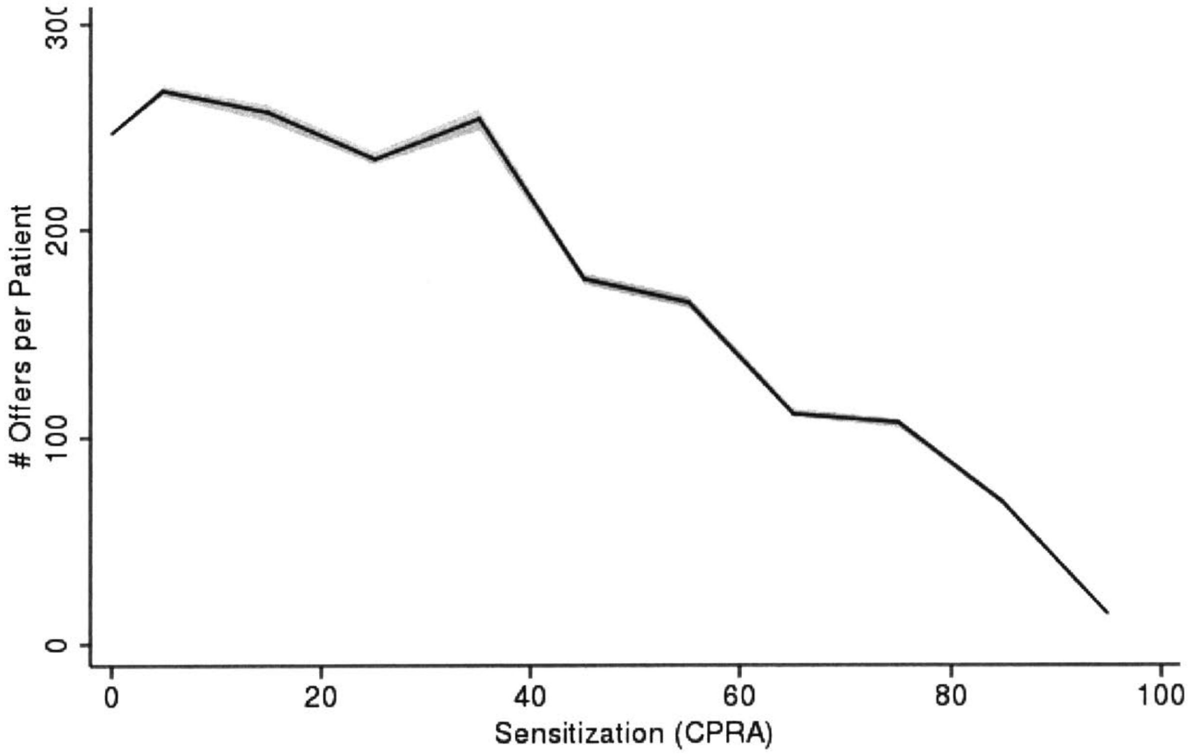

(b) Acceptance Rate

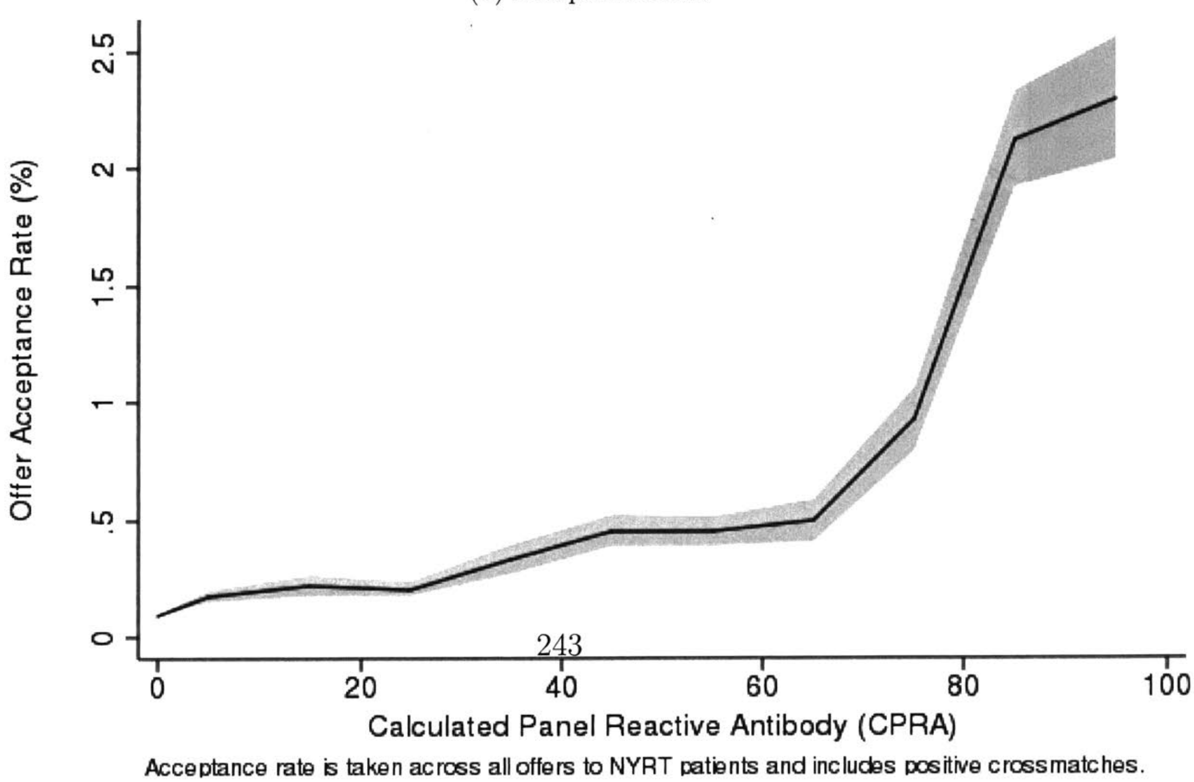

Note: Sample uses all offers, including offers that did not meet pre-set donor screening criteria.

Figure 2-7: Comparing Hazard Rates for Gompertz and Cox Proportional Hazards Models

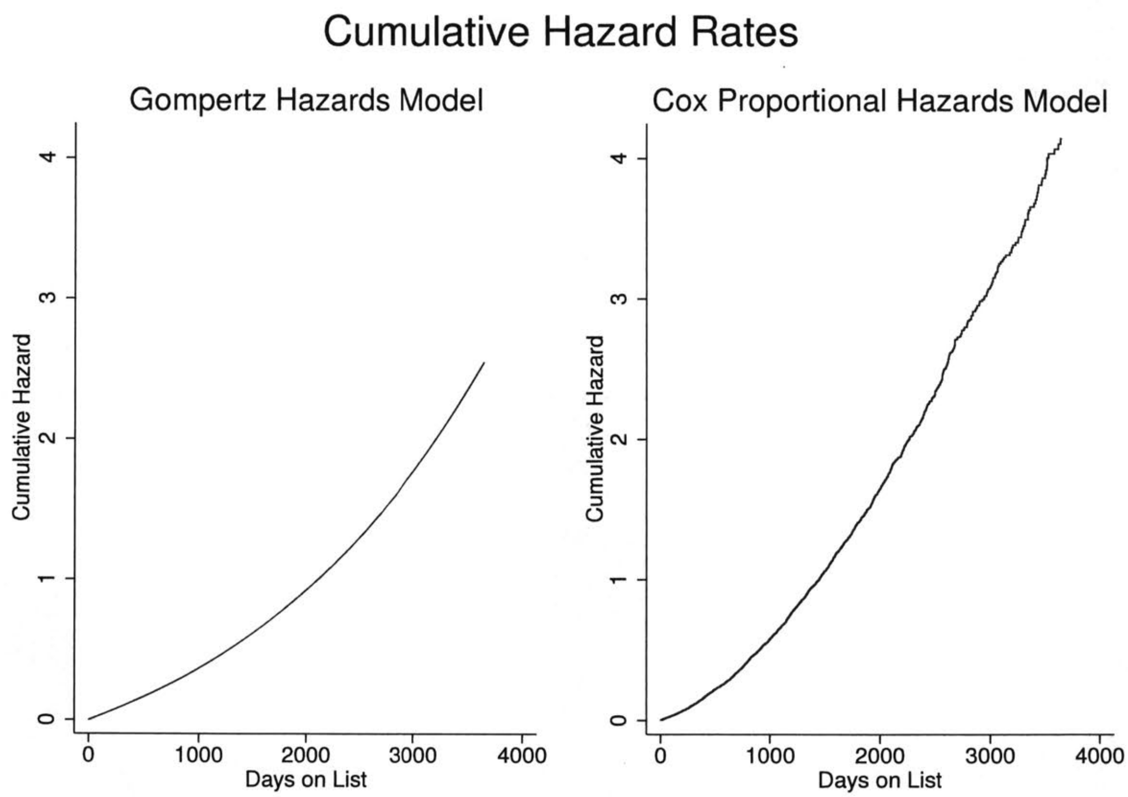

Figure 2-8: Extrapolation of Baseline Survival Curves

(a) Survival without a Kidney Transplant

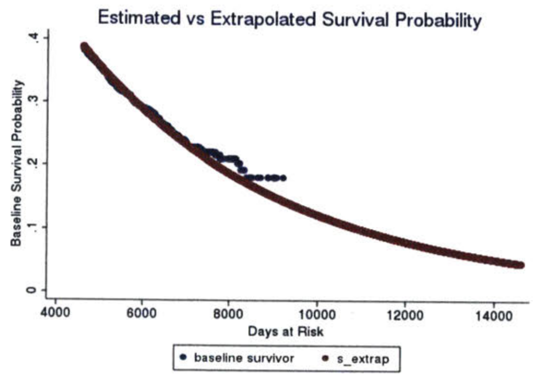

(b) Survival of Functioning Graft

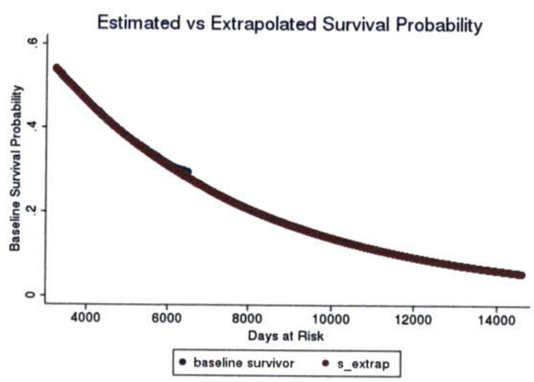

(d) Survival after a Kidney Transplant

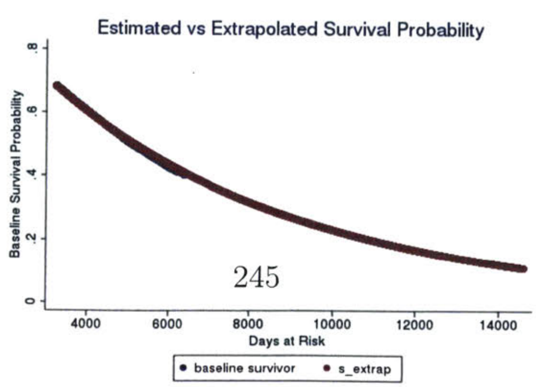

Table 2.15: Patient Sample Restrictions

	Number of Patients Registered
Patients Registered in NYRT Between 2010 and 2013	15063
Excluding Candidates for a Pancreas	14499
Excluding Candidates Who Did Not Need a Transplant	13950
Excluding Inactive Candidates	9985
Excluding Candidates Receiving Non-Standard Allocations	9917

Notes: candidates who did not need a transplant include patients who departed the waitlist because they refused transplantation, received a transplant in another country, could not be contacted, or had an improved condition. Candidates receiving non-standard allocations include placements from military or directed donations, expedited placement attempts, and medical emergencies.

Table 2.16: Offer Sample Restrictions

	Number of Offers
All PTR offers in U.S. sample	61038882
NYRT Patients	4516460
Offers made between January 1st, 2010 and December 31st, 2013	2883287
Excluding non-genuine refusals	1880672
Excluding offers after the donor's cutoff	1318961
Excluding patients and donors receiving non-standard allocations	1281024

Notes: non-genuine refusals include offers in the Potential Transplant Recipient (PTR) dataset which did not meet the patient's pre-set screening criteria; for which the patient or transplant surgeon was unavailable; or where the transplant could not occur for medical reasons. Non-standard allocations include placements from military or directed donations, expedited placement attempts, and medical emergencies.

Table 2.17: Fit of Mechanism Code: Predicted Offers

		In PTR Data		
		No	Yes	Total
Predicted by Simulation	No	14,907,971	88,366	14,996,337
	Yes	502,239	1,192,658	1,694,897
	Total	15,410,210	1,281,024	16,691,234

Notes: comparison between offers predicted by simulation of the mechanism and offers appearing in the PTR dataset. Offers not in either dataset set include compatible patient-donor pairs where the patient did not meet the donor's priority score cutoff, or where the donor did not meet the patient's pre-set screening criteria.

Chapter 3

Tax Refund Expectations and Financial Behavior

3.1 Introduction

Income uncertainty is thought to play a central role in household finances. While pre-tax income volatility is often emphasized as a source of this uncertainty, households may also have substantial uncertainty about their income *tax*. For low-income individuals, tax-linked transfer payments, including payments from the Earned Income Tax Credit (EITC), comprise a substantial portion of annual income.[1] Quantifying and understanding uncertainty about income taxes is therefore critical for understanding the role of transfer payments through the tax system in household finances, the potential consequences of changes to the tax system, and the effects of income uncertainty on consumption and financial decisions more broadly.

In this paper, we study what low- and moderate-income households do and do not

[1]For instance, in our sample, the mean refund totals nearly eight percent of annual income (roughly one month of earnings).

know about their income tax refunds before they file taxes. We then examine how financial behavior responds to expectations of future tax refunds, refund uncertainty, and surprises in realized tax refund amounts. We do so using a unique combination of (1) administrative tax records, (2) a linked panel of consumer credit reports, and surveys to measure both (3) expectations of tax refunds before tax filing and (4) consumption behavior after tax refund receipt. One key innovation in our setting is our direct measurement of taxpayers' beliefs about the probability distribution over their own future tax refund amounts. These expectations data allow us to study the amount of, and the effects of, income tax uncertainty on consumption without making strong assumptions about the sources of taxpayers' uncertainty.

We start by showing that taxpayers have correct mean expectations about their refund on average, but also face substantial uncertainty. This self-reported uncertainty accurately tracks "true" uncertainty, as measured by the difference ("surprise") between realized and mean expected tax refund amounts. We examine sources of refund uncertainty. Surprises are driven by changes in income and family structure in ways that are consistent with households misunderstanding how marginal tax rates change at different parts of the earned-income tax credit schedule. Nevertheless, we also find that much of this uncertainty is not explained by observables or by changes in household circumstances.

We then show that household consumption and borrowing behavior depends on expectations about tax refund amounts, refund uncertainty, and refund surprises. Households borrow a moderate amount out of their expected tax refunds: for each dollar of expected refund, households repay roughly 15 cents in debt shortly after tax refund receipt. Households also exhibit precautionary behavior in borrowing out of future tax refunds, as these borrowing and repayment patterns out of expected tax refunds are less pronounced for households that report being more uncertain about their refunds ex-ante. To our knowledge, this is some of the first evidence of precautionary behavior (prudence) among a low-income population in the US. This finding contrasts with prior work which has interpreted the combination of high income volatility and low savings rates as evidence against the existence of precautionary behavior among low-income households (Carroll et al., 2003).

Finally, we examine the link between tax refund surprises and debt, and find that surprises in tax refund amounts are not used to repay debt. In fact, we find that larger refund surprises lead to *increases* in overall debt, an effect that is entirely driven by higher balances on installment loans such as auto loans. This pattern implies a medium-run marginal propensity to consume (MPC) out of windfall income above one. One explanation for this stark finding is that refund surprises may be used to relax collateral constraints for newly financed durable purchases. We find suggestive evidence from a follow-up survey on durable consumption choices to corroborate this interpretation.

There are two primary advantages to our empirical approach. First, we obtain rich data on household balance sheets before and after the resolution of tax-related income uncertainty: administrative data on all reported income and nearly all financial liabilities, as well as survey-based measures of real and financial assets. Such data are particularly difficult to assemble for lower-income populations. Second, we directly elicit individuals' uncertainty about the component of future income risk driven by tax refund uncertainty. This stands in contrast to much of the existing literature that looks for evidence of precautionary behavior in response to uncertainty; we know of one notable exception (Jappelli and Pistaferri, 2000).

One substantial caveat to our approach is that we, like most researchers who use U.S.-based data, have relatively poor data on households' real and financial assets. All of our asset measures are survey-based, whereas we have administrative data on income and debt liabilities. Perhaps reassuringly, the low-income population in the United States has elsewhere been shown to hold low levels of financial and real assets, a finding that we corroborate in our survey measures. A second important caveat is that, while we analyze differences in financial behavior within groups that are at similar stages in the life-cycle (age, income, and family structure), there nevertheless may be important unobservable differences across individuals within these groups – for example, in unobservable labor income risk – that we cannot control for and that are correlated with tax-relevant uncertainty.

Related Literature This paper contributes to at least three distinct literatures.

First, we contribute to a large empirical literature in macroeconomics on household consumption, savings, and borrowing decisions. This work studies how households respond to income uncertainty ex-ante, and how households react to income surprises ex-post. A robust theoretical literature predicts that households will save precautionarily – maintaining a "buffer stock" – in the presence of future income uncertainty (Carroll, 1996; Deaton, 1991; Kimball, 1990), and calibration exercises suggest that the role of precautionary motives in saving over the life-cycle is substantial (Carroll and Samwick, 1998). However, other empirical work has found limited evidence for precautionary behavior (Dynan, 1993), especially among low-income households (Carroll et al., 2003). This latter result stems from substantial labor income uncertainty faced by low–income households coupled with low observed savings rates. Much of the empirical work testing for precautionary motives uses labor income uncertainty implied by income processes imputed using observables such as age and occupation (Carroll and Samwick, 1998; Dynan, 1993; Skinner, 1988). One notable exception uses self-reported uncertainty measured through a survey, as we do (Jappelli and Pistaferri, 2000). We believe we are the first paper to link such survey-based measures of uncertainty to administrative data on income and borrowing.

Another vein of empirical macroeconomics research studies how consumers respond to windfall income surprises. Most closely related to our study of tax refund surprises is a set of papers analyzing responses to tax rebates (Agarwal et al., 2007; Baugh et al., 2018; Broda and Parker, 2014; Johnson et al., 2006; Parker et al., 2013). These papers find, as we do, high marginal propensities to consume (MPCs) out of such windfall income. Of particular note is Parker et al. (2013), which finds that up to 60 percent of tax rebate payments are used to purchase durables, and especially vehicles, within 3 months of rebate receipt. These findings are consistent with our result that positive refund surprises are used to finance durable purchases.

Second, we contribute to a growing literature on the limits of taxpayers' understanding of the tax code and on the consequences of tax complexity for individuals and firms. Several recent papers have shown that individuals and firms fail to take full advantage of the credits and refunds for which they are eligible. Part of this is likely due to hassle costs: individuals may rationally choose not to invest the time or money required to optimize their tax benefits (Benzarti, 2017). Among low-income EITC filers, like those in our sample, part of this failure to optimize may be due to lack of information about the tax system (Aghion et al., 2017; Chetty et al., 2013; Zwick, 2018). Prior research has shown that many individuals are unaware of EITC program rules and that lack of information has real consequences for earnings behavior (Chetty and Saez, 2013; Chetty et al., 2013; Romich and Weisner, 2000; Smeeding et al., 2000). We contribute to this literature by directly quantifying the amount of uncertainty faced by our population of low-income tax filers, and by linking this uncertainty to actual consumption decisions.

Third, we contribute to a diverse literature on the measurement and reliability of subjective expectations data (Manski, 2004). Following the pioneering work of Engelberg et al. (2009), we elicit not just point forecasts or mean expectations, but individuals' subjective probability distributions over future events. These methods have previously shown success in measuring inflation expectations (Armantier et al., 2016), income expectations among college students (Zafar, 2011), and income expectations in a developing country (Delavande et al., 2011). Elicited expectations have been shown to affect financial behavior in lab settings (Wiswall and Zafar, 2014). Our contribution is to link probabilistic income expectations with a panel of administrative data to study how financial behavior responds to such expectations in a "real world" (non-lab) setting, and to demonstrate success of these survey questions even in a low-income, relatively low-education U.S. population.

The rest of the paper proceeds as follows. The next section describes the empirical setting and data. Section 3.3 describes how we translate our survey measures of beliefs into probabilistic distributions and compares these distributions to actual refund amounts. Section 3.4 shows how refund expectations, uncertainty, and surprises translate into consumption and borrowing decisions. Section 3.5 concludes.

3.2 Data and Empirical Setting

In this section we describe our data and empirical setting. We first provide institutional background on the setting, a clinic that provides free income tax preparation services in Boston. We then describe our administrative (tax and credit) and survey (expectations, assets, and consumption) data sources. We conclude by describing the characteristics of our sample.

3.2.1 Boston Tax Sites

Our data come from a Volunteer Income Tax Assistance (VITA) tax preparation center operated by the Boston Tax Help Coalition and the Boston Office for Financial Empowerment (OFE). The City of Boston runs over 30 free tax preparation centers, which annually serve more than 13,000 clients. Our data come from one of the largest of these centers, Dorchester House.

Boston residents are eligible to receive these free tax preparation services if they worked in the prior year, earned less than $54,000, and do not own their own business. Eligible individuals who come to the tax site ("clients") typically go through three separate stations. First, they complete an intake survey, which includes questions on demographics, use of city services, savings behavior, and credit usage. Next, clients are offered a free "financial check-up" from a trained volunteer referred to as a "financial guide." The financial guides offer the client a free credit report and provide information on other city services for which the individual may be eligible. Finally, the client is sent to a tax preparer who electronically prepares and submits the individual's tax return.

We partnered with the Boston OFE to field a survey of clients' expectations about their tax refund (detailed in Section 3.2.3) at the second of these three stations, together with the financial guide. This survey came before clients filed taxes, and so measures their pre-filing uncertainty about their tax refund. At this stage, clients also provided consent for their tax, credit, and survey information to be used for research purposes. Figure 3-1 describes the flow of clients through Dorchester House in more detail.

Two operational features of the financial check-up stage deserve mention. First, because of financial guide shortages and constrained tax site operating hours, many tax filers skipped the financial check-up during busy periods. As a result, we obtained consent from only 60 percent of tax filers. However, among clients who completed the financial counseling session our consent rate was 96 percent. Therefore, we do not believe that tax filer consent was a major source of selection into our research sample.

Second, the OFE implemented a separate randomized controlled trial as part of the financial check-up wherein clients were randomly assigned to a more or less detailed check-up. Those assigned the more detailed check-up were given an in-depth explanation of their credit report, as well as financial advice and referrals to a variety of services provided by the City of Boston and state and federal organizations. Those assigned to the less detailed check-up also received their credit report, but no detailed financial advice or referrals. In our analysis of consumption responses in Section 3.4, we control for treatment status at the financial check-up stage.[2]

3.2.2 Administrative Tax and Credit Data

We obtain administrative tax returns for consenting clients who filed their taxes at Dorchester House. These data include information on income, family structure (filing status and number of dependents), and refund amount. For individuals who previously used the city's tax preparation services, we are able to link these returns to those from earlier years. We have two years of returns for 69 percent of the tax filers in the credit and expectations survey sample.

[2]An OFE analysis of the randomized controlled trial finds balance across treatment assignment on a range of taxpayer characteristics. The report can be accessed at https://owd.boston.gov/wp-content/uploads/2017/07/DES-89-Financial-Check-up-Evaluation-2017-Web.pdf?utm_source=Office+of+Workforce+Development&utm_campaign=699487e955-EMAIL_CAMPAIGN_2017_07_24&utm_medium=email&utm_term=0_f071f9ca69-699487e955-226050949

We merge these administrative tax records with a short panel of consumer credit reports for clients who provided consent during the financial check-up. We have four reports for each individual in our sample: one that was pulled when they visited the tax site, and three that were pulled one, two, and six months later. The one and two month credit reports measure changes in debt levels soon after tax filing. For clients who receive their refund by direct deposit, both the first and second-month follow-up credit reports show loan balances after tax refund receipt; for clients who receive their tax refund by paper check, the first of these two credit reports likely show balances from prior to refund receipt. The six month report allows us to observe longer-run deleveraging and new loan originations (e.g. auto loans) that may not have been reported in time for the one and two month follow-ups.

3.2.3 Expectations and Consumption Surveys

We supplement these administrative data sets with three distinct surveys. These surveys provide information on taxpayer demographics and assets, refund expectations, and consumption before and after refund receipt. Our first source of survey data is the demographics and assets survey individuals completed when they arrived at Dorchester House. From this survey we obtain information on a client's gender and level of education (high school degree or some college), and on a client's savings behavior. The response rate for this survey was high: of the 1,186 individuals who filed taxes at Dorchester House during the spring 2016 season, 995 completed the survey. A copy of the survey is provided in Appendix 3.7.2.

We obtain information on tax filers' expectations and uncertainty about their refunds from a short four-question survey. Tax filers completed this survey after they had been paired with a financial guide, but before they filed their taxes and learned their actual refund amount. We elicited beliefs in two ways. First, we directly asked each filer how much they expected their refund to be, and their qualitative certainty that the refund would fall within $500 of this amount. Second, we provided individuals with a set of six bins, and asked them the probability that their refund would fall within each bin. A copy of the survey is provided in Appendix 3.7.3. We discuss how we translate the answers to these questions into probabilistic beliefs in Section 3.3.

Finally, we merge these expectations with data from a second consumption survey designed to measure saving and consumption behavior before and after refund receipt. While we obtain substantial information on consumption from the panel of credit reports – for example, the presence of new auto loans or pay-down of debt – these reports do not contain information on durable purchases or on the timing of purchases relative to refund receipt. They also contain no information on savings. The response rate to our consumption survey was 46 percent (291 out of 625 filers in our sample), which is high compared to similar phone-based surveys.[3] The consent language and questions are provided in Appendix 3.7.3.

3.2.4 Descriptive Statistics

Table 3.1 presents descriptive statistics from our sample of low-income Boston tax filers. Column 1 includes all 995 tax payers who visited Dorchester House and completed the in-take survey during the spring 2016 tax season. The average adjusted gross income (AGI) in this sample is $21,603, and the mean refund size is $1,765. Thirty-eight percent of filers receive the Earned Income Tax Credit (EITC). Most filers are unmarried, but only 28 percent file as a single head of household, and 34

[3]For example, Allcott and Kessler (2018) obtain an 18 percent response rate to a phone survey of energy usage, a higher rate than they expected.

percent have dependents (including married filers). Eighty percent of filers have a high school degree, but only 15 percent have attended college. The average age is 41 years old.

The remaining columns restrict the sample to the subsets for which we have credit reports, the expectations and consumption surveys, and tax returns from the prior year (2015). Column 2 reports statistics from the 714 filers in the credit report sample. Tax filers in this sample are highly leveraged with very low savings rates. The average filer has roughly $9000 in installment debt, $1700 in credit card debt, and $500 in savings. Average savings is less than one third of the average refund amount, and less than 5 percent of average debt. The mean FICO score for those with credit reports is 664, below the 2016 U.S. average of 700. The credit report sample is similar to the asset survey sample in terms of age, gender, education, family status, AGI, and refund amounts.

These economic and demographic variables remain stable across columns 3-5, suggesting that attrition across surveys is largely unrelated to tax status or demographic characteristics that could bias our results. Column 3 restricts to the 625 tax filers in the credit sample who completed the expectations survey. The vast majority (557) of these filers also completed the asset survey. Column 4 restricts to the much smaller sample of 291 filers who completed the follow-up consumption survey. Despite a 46 percent response rate, households that did and did not respond are nearly identical in terms of their average characteristics. Column 5 restricts to tax filers with expectations and credit reports who filed their taxes with the City of Boston in the previous year (2015).

Table 3.1 shows that refund amounts are large relative to income, savings, and debt levels. The mean refund of $1,776 is approximately eight percent of the average individual's adjusted gross income and is triple the average individual's savings at tax filing. In addition, we show in section 3.3 that tax filers face a large degree of uncertainty about their refunds. This suggests that tax refunds and uncertainty about them can have important implications for financial behavior in this population.

3.3 Tax Refund Expectations and Realizations

We surveyed Dorchester House tax filers to elicit their beliefs about the tax refund they would receive after filing. Since consumption responses to refund amounts may depend on both mean expectations and uncertainty, our survey elicited both aspects of filers' beliefs through a probabilistic survey question. This section describes the belief survey; explains how we converted survey responses to smooth belief distributions; and compares beliefs to realized refund amounts and surprises. Although tax filers reported substantial uncertainty about their refunds, their mean expectations were, on average, correct. In addition, filers reporting greater uncertainty saw larger refund surprises. This suggests that most tax filers had an accurate sense of the refund amount they could expect to receive, and also had an accurate sense of the uncertainty they faced. However, there is evidence of (mean) inaccurate expectations particularly among tax filers whose incomes or family status changed relative to past years, suggesting that tax filers are imperfectly aware of after-credit tax rates at different points in the income tax schedule.

3.3.1 Belief Elicitation Survey

The survey was administered at the beginning of the financial counseling session at Dorchester House, which took place prior to the tax preparation session. We view this as the ideal time to survey program participants on their refund expectations:

tax filers had not yet received any information about their refunds. However, tax filers had collected their tax documents, come to the tax site, and filled out a detailed economic and demographic survey for use during the tax preparation session.

The final question in this survey elicited probabilistic refund expectations. Respondents were asked the percent chance that their refund would fall in each of six bins: negative (they would have to pay taxes), $0-$500, $500-$1,000, $1,000-$2,500, $2,500-$5,000, over $5,000. We asked for points in a cumulative density function rather than moments such as the mean and standard deviation because subjective probabilities are easier to understand and calculate. In addition, probabilistic survey questions can provide richer information about beliefs. We would have ideally constructed bins around each filer's point estimate to obtain comparable uncertainty measures across households. The need to conduct the survey quickly made this approach too difficult to implement, so we used fixed intervals. Nevertheless, we show in the next section that the fitted distributions accurately capture both expected refund amounts and uncertainty.

Appendix Table 3.9 describes features of the elicited belief distributions. The first column presents statistics for all tax filers in our main analysis sample, and the remaining columns disaggregate those statistics into subgroups. Forty percent of respondents put nonzero probability on three or more bins, while 60 percent did so on only one or two bins.

3.3.2 Fitting Belief Distributions

To summarize beliefs and to quantify both mean expectations and uncertainty, we convert each probabilistic elicitation into a smooth probability distribution following Engelberg et al. (2009) (hereafter EMW). Our goal is to use all information available in respondents' subjective probabilities and to smooth between points of the cumulative density function in a reasonable way. We fit a distribution which depends on the number of bins on which the respondent placed positive probability. Single bin reports are fit with a scalene triangle; the support is the full bin, and the mode is the point estimate. In this case, we depart from EMW by using additional information from the respondent's point estimate, fitting a scalene triangle rather than an isosceles triangle. Meanwhile, two-bin reports are fit with an isosceles triangle with the widest possible support that is consistent with the probabilities for each bin. These sets of assumptions uniquely pin down a distribution for one- and two-bin responses. For three or more bins, we follow EMW in fitting a beta distribution to the reported quantiles. Triangle and beta distributions are appropriate for our setting because they have finite support, and because beta distributions can match a wide range of distributional shapes that might be implied by probabilistic survey questions.[4] The maximum refund amount was a little below \$20,000, and the lowest refund amount was about -\$500. We take these two values as the endpoints of the support of the highest (over \$5,000) and lowest (negative) bins.

The triangle distributions are exactly identified and fit using analytical formulas. To fit the beta distributions, we follow EMW and minimize the sum of squared differences between the reported cumulative probabilities at each point in the distribution's support and those of a beta distribution with the same support. Let $\mathcal{X}$ denote the support points of the response to the probabilistic survey question. Let Z denote a beta-distributed random variable governed by parameters (α, β) and normalized to have support on $\mathcal{X}$. Finally, let p_x denote the reported cumulative probability at each point $x \in \mathcal{X}$. We find the $(\hat{\alpha}_i, \hat{\beta}_i)$ for the elicited distribution from each individual i which solves

[4]We also depart from EMW by not constraining the estimated beta densities to be single-peaked.

$$\min_{\alpha,\beta} \sum_{x \in \mathcal{X}_i} [p_{x,i} - P(Z \leq x \mid \alpha, \beta)]^2$$

The fitted distributions reveal large variation in the expected refund amounts and uncertainty across tax filers. Appendix Table 3.9 shows that the average mean expectation is \$1,970, and the average standard deviation is \$740. The average coefficient of variation is 0.50 – so refund uncertainty is, on average, large relative to the expected amount. These averages mask an enormous amount of variation across tax filers in their refund expectations. The standard deviation across tax filers of their mean expectations is \$2,850, and the standard deviation of subjective uncertainty (where uncertainty is measured using the standard deviation of each tax filer's fitted distribution) is \$1,019.

It is illustrative to compare self-reported measures of qualitative and quantitative uncertainty as a validity check on these survey responses . Table 3.2 summarizes the coefficients of variation of respondents' belief distributions depending on whether they were "very sure," "somewhat sure," or "not sure at all" about whether their refund would be within \$500 of their point estimate. The most uncertain individuals have much larger coefficients of variation. Two-way t-tests of equal means strongly reject equal quantitative uncertainty for any two qualitative responses. The next subsection provides additional evidence that the quantitative measures of uncertainty meaningfully capture tax filers' subjective beliefs.

3.3.3 Beliefs and Realizations

Our unique institutional setting allows us to compare applicants' refund expectations to what they actually received. This comparison shows not only that applicants have correct mean expectations, but also that they understand the degree of uncertainty they face, at least on average.

Figure 3-2 compares mean expectations from survey responses to actual refund amounts. Mean expectations closely track realized amounts. The slope of the regression line is close to one, though beliefs are slightly attenuated: those with the most extreme realizations had slightly less extreme expectations. The strongly linear relationship between expected and actual refund amounts does not imply that tax filers faced little uncertainty, or that any individual had unbiased beliefs. Rather, it shows that beliefs tracked realized refund amounts on average, suggesting that the probabilistic survey question does contain meaningful quantitative information.

Figure 3-3 performs a similar exercise for the degree of self-reported uncertainty. It compares the magnitude of each tax filer's refund "surprise"–the difference between the realized and expected refund amounts–to the fitted standard deviation of their belief distribution. There is a clear linear relationship between subjective uncertainty and realized absolute errors.[5] Thus, tax filers face substantial refund uncertainty, and furthermore they seem to be aware of the degree of uncertainty that they face. The next section investigates the determinants of refund uncertainty and shows that some but not all of the variation in refund uncertainty across individuals can be explained by observed characteristics.

3.3.4 Predictors of Refund Uncertainty and Surprises

In this section we investigate the predictors of refund uncertainty and the magnitude and direction of refund surprises. We find that current income and family structure are highly predictive of refund uncertainty, while demographic variables such as age,

[5]Note that the slope of the line should not necessarily be one – a standard deviation is the square root of the expected squared error, not the expected absolute error – and the conditional expectation function need not be linear.

gender, and education are less predictive. Filers whose income or family structure changed from previous years are also more uncertain, but tax filers whose situation did not change still report substantial uncertainty. Furthermore, the characteristics that predict greater uncertainty are also associated with larger surprises, suggesting that these relationships reflect real differences in uncertainty across tax filers. Even after controlling for demographic characteristics and changes in tax situation there is substantial variation in both uncertainty and surprises.

To analyze determinants of subjective refund uncertainty, we first regress three of measures of refund uncertainty on a range of economic and demographic characteristics. Our specifications take the form

$$y_i = X_i\beta + \epsilon_i \tag{3.1}$$

where y_i is a measure of uncertainty, and X_i includes sociodemographic variables capturing age, education, and gender; demographic variables including marital status and number of dependents; and dummy variables for each quartile of adjusted gross income (AGI). Results from a series of these regressions are shown in Table 3.3. Columns 1-3 use the standard deviation of the household's parametric belief distribution as the measure of uncertainty, y_i; columns 4-6 use the absolute value of the refund surprise (refund amount - mean expectation); and columns 7-9 use the size of the refund surprise.

Column 1 shows that the number of dependents and income quartile are quite predictive of refund uncertainty, as measured by the standard deviation of beliefs. Tax filers in the third and fourth income quartiles report greater uncertainty, as do households with more dependents. These differences are large: for example, an additional dependent is associated with \$449 more in refund uncertainty (as measured by the standard deviation of an individual's fitted belief distribution), and filers in the third AGI quartile report over $600 greater uncertainty than filers in the first quartile. In contrast, the demographic variables capturing age, education, and gender are less predictive of uncertainty.[6] These patterns are consistent with a model in which cognitive limitations and total experience with the tax system are less important determinants of refund uncertainty than economic characteristics that directly determine tax liabilities.

Columns 2 and 3 add several variables related to changes in financial and family status: whether the filer received unemployment insurance payments in the past year; whether their filing status changed, e.g. from single to married; the (absolute or level) change in their AGI; and the (absolute or level) change in the number of dependents. The sample size in these columns is lower, reflecting the fact that we only observe these changes for filers who filed at a Boston tax site in previous years. Column 2 controls for indicators for, and magnitudes of, these year-to-year changes, while column 3 replaces absolute changes with level (signed) changes.

Our results in these two columns provide mixed evidence on whether changes in financial and family status contribute to refund uncertainty. In column 2, the magnitude of change in AGI is positively related to uncertainty but not statistically significant. Households that experienced an increase in the number of dependents actually report significantly lower refund uncertainty, but there is no significant correlation between uncertainty and the *absolute* change in number of dependents. These results however are noisy enough to be consistent with taxpayer uncertainty being partially driven by changes in financial or family situations, which may result from how after-credit tax rates depend directly on family size and structure as well as

[6]The coefficient on the indicator for age > 50 is negative and marginally significant, but this pattern disappears after adding controls for change in filing status.

income. In both columns 2 and 3, the coefficient estimates on third quartile of AGI and number of dependents remain statistically significant and of similar magnitudes as seen previously in column 1.

A natural question is whether the tax filers who reported greater uncertainty actually saw higher variance in their refund surprises. Columns 4-6 of Table 3.3 repeat the regressions in columns 1-3 with the absolute value of the tax filer's refund surprise as the dependent variable. Almost all variables which significantly predict refund uncertainty in columns 1-3 significantly predict absolute errors in the corresponding specification in columns 4-6, with the same sign and similar magnitudes. Number of dependents and income quartile remain the main predictors of refund surprise magnitudes, while demographic variables are less predictive. In addition, changes in AGI and changes in the number of dependents are significantly predictive of surprise sizes. Column 5 shows that tax filers with larger AGI changes saw larger surprises (30 dollars of uncertainty per 1,000 dollar change in AGI), as did filers with changes in number of dependents (1,000 dollars of additional uncertainty). In addition, surprise size is predicted by the direction of these changes. Households with increases in AGI and dependents actually saw smaller surprises. Taken together, these results suggest that different household types accurately assess the uncertainty they face as a result of unchanging family characteristics, but they may not fully update their beliefs or their subjective uncertainty based on changes in income and family structure.

Finally, columns 7-9 of Table 3.3 regress the surprise amount, rather than the magnitude, on the same sets of predictors to investigate whether certain types of households systematically over- or under-estimate the size of their tax refunds beliefs. Consistent with the finding in section 3.3.3 that expected refund amounts track realizations, most of the coefficient estimates on current tax filer characteristics are statistically insignificant. In particular, current AGI and number of dependents, which were predictive of refund uncertainty and surprise magnitudes, do not systematically predict the direction of mistakes. Though the coefficients on the second and third income quartiles are marginally significant in column 7, they become insignificant after controlling for the change covariates.

That said, there is one demographic variable which significantly predicts bias: married tax filers systematically overestimate their refund amounts by $1,000 relative to unmarried filers. Additionally, we find in columns 9 that changes in financial or family status are predictive of under- or over-estimates: filers whose AGI rose had lower surprises, while filers with an increase in their number of dependents saw higher surprises. In particular, tax filers whose incomes rose overestimated their refund amounts by 3 cents per each dollar of change in AGI, whereas filers underestimated their tax refund by more than 700 dollars for each additional dependent relative to the previous year. This is consistent with tax filers underestimating the slope of their refund with respect to characteristics: for example, the EITC claw-back rate as incomes rise and the generosity of EITC or child-tax credit benefits for additional dependents.

The above discussion yields several takeaways regarding tax refund expectations. First, the relationship between uncertainty and variables directly relevant to a household's tax liability, but not sociodemographic variables, suggests that uncertainty about how financial characteristics map to tax liabilities is common across a range of households with varied levels of sophistication and experience. Second, the types of households who report greater uncertainty also see larger refund surprises. Third, current taxpayer characteristics do not predict the direction of mistakes, with the exception of marital status; however, tax filers do not fully update about how changes in income and family structure affect their tax liabilities. A final observation is that even this broad set of tax filer attributes fails to explain all of the variation in uncertainty. The R-squared in column 3 – the highest across specifications – is 0.358, leaving substantial unexplained variation.

3.4 Borrowing and Consumption Responses to Tax Refunds

In this section we study how individuals' borrowing and consumption behavior around the time of tax filing responds to their expectations about, and actual realizations of, their tax refunds. In our sample of low-income tax filers, we find that roughly 15 cents per dollar of expected tax refund is used to repay revolving debt after tax refund receipt. In contrast, the unexpected ("surprise") component of tax refunds has a precisely estimated near-zero effect on revolving debt repayment. These results are consistent with individuals borrowing out of their expected tax refunds to smooth consumption over the course of the year, while also having a high propensity to consume out of windfall income in the form of tax refund surprises. Furthermore, we find that individuals exhibit precautionary behavior in their willingness to borrow out of expected tax refunds, as post-refund debt repayment is significantly more pronounced for individuals who report being more certain of their refund amount ex-ante.

Further exploring the effects of tax refund surprises, we find that surprises have a significantly positive effect on installment debt balances: unexpectedly high refunds lead to higher installment debt levels. This result is consistent with tax refunds partly being used to fund down payments for newly financed durable purchases, such as new auto loans. We use our follow-up survey of consumption behavior to corroborate this possible mechanism, finding that individuals with higher tax refund surprises indeed more frequently report that they bought a new car or initiated home repairs after tax filing. Summing the effects of refund surprises across both installment balances and revolving balances, we estimate that an additional dollar of refund surprise leads to an additional 40 cents of debt. While noisy, this estimate is significantly greater than zero and suggests a mechanism whereby medium-run MPCs can lie above 1 when windfall income is used to relax collateral constraints for new borrowing.

Finally, we test whether ex-ante uncertainty about tax refunds affects individuals' propensity to consume out of tax refund surprises. If individuals behave precautionarily and if the consumption function is concave in cash on hand, as predicted under a broad set of conditions (Carroll and Kimball, 1996; Zeldes, 1989), then individuals with more ex-ante uncertainty should have lower propensities to consume out of tax refund surprises. Our results are consistent with this prediction in sign, but underpowered, and we fail to reject the null that MPCs out of surprises are the same for different levels of ex-ante uncertainty.

3.4.1 Revolving Debt Repayment

We begin by examining financial behavior around tax filing by studying revolving debt balances.[7] These loans are sensible to examine first, as their balances are most readily adjustable over a short time horizon; we defer until section 3.4.2 a discussion of less easily adjustable installment debt.[8]

Using the linked panel of credit report data, we calculate the change in each

[7]Revolving debt includes all loans with a flexible repayment schedule and an open line of credit that can be used flexibly over time and over purchases, including credit cards, retail store cards, and home equity lines of credit (HELOCs).

[8]Installment debt includes all loans with a fixed repayment schedule. These loans are often used to fund one-time purchases, including car loans, student loans, and mortgages.

tax filer's revolving debt balance between the credit report drawn just prior to tax filing and credit reports at subsequent two-month and six-month horizons. These provide short- and medium-run measures of responses to tax refunds. We then regress these two-month and six-month changes on three features of tax filers' beliefs and realizations of tax refunds: (1) their expectation of their tax refund amount, (2) their uncertainty about their tax refund amount as measured by the standard deviation of their elicited subjective probability distribution over refund amounts, and (3) the surprise in their realized tax refund relative to their expected refund. Debt changes are signed so that a negative change is a decrease in debt levels. Refund surprise is defined as realized tax refund minus expected tax refund; thus, a positive surprise is "good news" for the tax filer. We estimate regressions of the form

$$\Delta b_i = \alpha + \beta_1 \mu_i + \beta_2 \text{surprise}_i + \beta_3 \sigma_i + \gamma Z_i + \eta_i \tag{3.2}$$

where i indexes a tax filer; Δb_i denotes change in balances; μ_i is i's mean expected refund; σ_i is their subjective standard deviation; and surprise$_i$ is their refund surprise. The vector Z_i controls for a range of interacted tax filer characteristics because household debt paths may differ at different stages of the lifecycle. We include fully interacted fixed effects for age group, income quartile, marital status, and whether an individual has dependents.[9] These interaction terms aim to absorb differences in levering or deleveraging over time that are due to differences between, for example, a young unmarried parent in the middle of our sample's income distribution, and a married elder at the bottom of our sample's income distribution. All residual variation is within a set of individuals who have similar lifecycle circumstances.[10] We

[9]We group individuals into three age bins based on whether they are younger than 25, between 26 and 50, or over 50.

[10]The most notable omission from these lifecycle controls is arguably the variability of individuals'

also add controls for whether an individual received their refund by direct deposit or by paper check, and for an individual's treatment status in the randomized trial being conducted simultaneously at the tax site as discussed in section 3.2.1.

Table 3.4 reports estimates from specifications without and with lifecycle controls at a two-month horizon (columns 1 and 2) and a six-month horizon (columns 3 and 4). Column 1 indicates that for every dollar of expected tax refund, our sample repays roughly 15 cents in revolving balances after refund receipt. These estimates remain stable and significantly different from zero across both horizons and with the inclusion of lifecycle controls. These results quantify how revolving lines of credit are used to transfer a moderate share of expected tax refunds forward in time to fund earlier consumption.

Turning to the second row of the table, we see that surprises in tax refunds are *not* used to repay revolving debt. After including lifecycle controls (columns 2 and 4), we can reject more than 13 cents of refund surprise being put toward revolving debt repayment at a two-month horizon, and 21 cents at a six-month horizon. Considering the low savings levels in our sample (see section 3.2.4) this suggests that households have a marginal propensity to consume (MPC) of close to one out of cash on hand.[11] This is consistent with existing evidence of high MPCs from windfall income among low-income consumers (Jappelli and Pistaferri, 2014). We further explore such propensities to consume in section 3.4.2.

In the final row of the table we test for the presence of precautionary behavior in

labor income. In future work, additional data cleaning of individuals' self-reported industry and/or occupation, together with our data on age, income level, and history of unemployment insurance receipt, will make it possible to impute a measure of labor income uncertainty using data such as the Panel Study of Income Dynamics (PSID) or Current Population Survey (CPS) for long-run or short-run income risk, respectively.

[11]Note that a coefficient of -1 on refund surprise would indicate an MPC of zero; tax filers would then be spending their entire refund surprise to pay down debt rather than changing consumption.

borrowing out of tax refunds. If tax filers are less certain of their tax refund amount ex-ante, they may borrow less of their expected refund before filing. We find that revolving balances are repaid less after refund receipt for more uncertain tax filers; for every dollar of standard deviation in a tax filer's subjective beliefs about their tax refund amount ex-ante, we estimate 35 to 40 cents less is used to repay debt ex-post. This pattern is consistent with uncertain tax filers precautionarily taking on less debt prior to filing their taxes. The estimated relationship between uncertainty and debt changes remains stable across both horizons and with the inclusion of lifecycle controls, although standard errors become larger at the six-month horizon. In dollar terms (results not shown), we estimate that having above-median refund uncertainty predicts roughly $275 less repaid toward revolving debt after refund receipt.

3.4.2 Installment Debt Repayment and Durable Consumption

We now turn our attention from revolving debt balances, such as credit card borrowing, to non-mortgage installment debt such as auto loans, retail loans, and student loans. We conduct the same analyses as in table 3.4 for installment debt instead of revolving debt, again estimating equation 3.2 with and without taxpayer controls at two-month and six-month horizons. We present these results in table 3.5.

We find that higher (more positive) tax refund surprises lead to relative *increases* in installment debt. The effect of surprises on installment debt is near zero at a two-month horizon, but at a six-month horizon each additional dollar of tax refund surprise leads to an additional 65 cent increase in installment debt. This estimated effect remains stable as lifecycle controls are added in column 4. We reject a zero effect on installment debt changes in columns 3 and 4 with 95% and 90% confidence, respectively.

Other coefficient estimates for installment debt are noisier than for revolving debt. While we estimate that an even larger share of expected refunds are used for installment debt repayment than for revolving debt repayment, the estimates are not significantly different from zero or from the revolving debt estimates at reasonable confidence levels.[12] Similarly, the effect of uncertainty on borrowing behavior for installment debt cannot be distinguished from zero or from the estimated effects for revolving debt.

The result that positive refund surprises lead to rising installment balances is intriguing. We corroborate this relationship visually in figure 3-4, using a binned scatter plot to show conditional means of the dependent variable across bins of refund surprise after partialling out the controls in column 4 of table 3.5. The visual evidence strongly confirms the relationship between refund surprises and changes in installment debt. Across large and small, positive and negative surprises, an approximately linear relationship holds between surprises and installment debt changes, confirming that this pattern is not driven by non-linearity or outliers.

Given that installment debt is less adjustable over short horizons than revolving debt – it has a fixed repayment schedule, and taking out additional debt typically must coincide with a new purchase – this result suggests that positive tax refund surprises may be used to fund down payments on newly financed durable purchases. Conversely, a negative tax refund surprise may make an anticipated durable purchase no longer possible due to collateral constraints. This mechanism is illustrative of a case where medium-run MPCs out of windfall income can in fact be above 1, when

[12]Greater rates of borrowing out of tax refunds through installment rather than revolving debt would be consistent with optimal consumption smoothing behavior when installment debt interest rates are lower than revolving debt interest rates (and when installment borrowing is sufficiently fungible to substitute for revolving debt).

such income is used to relax collateral constraints for new durable financing.

To investigate the durable purchases explanation, table 3.7 examines the relationship between tax refund surprises and self-reported durable purchases in the follow-up durable consumption survey we conducted approximately two months after tax filing. See section 3.2 for a survey description. This table follows a similar format to tables 3.4 and 3.5. While the coefficient estimates are not statistically significant at conventional levels, they are consistent ($p = 0.263$) with the interpretationthat higher refund surprises lead to a greater likelihood of new durable purchases after refund receipt.

To conclude this section, we pool both revolving and installment debt balances together and estimate the effects of tax refund expectations, uncertainty, and surprises on overall non-mortgage debt balances. We estimate the same specifications as in tables 3.4 and 3.5, but with the total change in installment and revolving debt balances as the dependent variable.[13] Estimates for these specifications are shown in table 3.6. A moderate amount of each dollar of expected refund is used to repay debt shortly after refund receipt, suggesting that individuals use both revolving and installment debt to smooth consumption by borrowing out of refunds ex-ante. We also find evidence that individuals behave precautionarily in borrowing out of tax refunds, as this borrowing is less pronounced for individuals with more uncertainty ex-ante. However, unlike in our specification using only revolving debt, these differences are not statistically significant. Finally, we again find that more positive refund surprises lead to significantly higher total debt at a six-month horizon. Our preferred estimate when including lifecycle controls in column 4 is an additional 38 cents of total debt

[13]Individuals are included in this regression sample if they ever have either revolving loans or installment loans over our panel horizon, so the sample differs from that in our analyses of revolving debt and installment debt alone.

for each additional dollar of refund surprise. The modest amount of deleveraging on revolving debt does not overwhelm the larger, positive effect on installment debt balances. On net, positive surprises appear to lead to higher indebtedness in the medium term.

3.4.3 Further Tests of Precautionary Behavior

In this final subsection, we study the relationship between ex-ante uncertainty about tax refunds and individuals' later consumption out of tax refund surprises. This provides a further test of precautionary behavior. Here, we test a central prediction of the buffer stock consumption-savings model: the consumption function should be concave in cash on hand (Carroll and Kimball, 1996; Zeldes, 1989).

The logic of our test is that conditional on other characteristics, refund uncertainty is an instrument for initial debt levels. Suppose two tax filers have identical characteristics, mean refund expectations, and refund surprises, but have different levels of refund uncertainty at filing. With a precautionary savings motive, the filer with greater uncertainty should enter tax season with lower debt or higher assets to maintain a buffer against lower-than-expected tax refund realizations. Since the two filers are identical after refund uncertainty is realized, they should have the same consumption function, but the measured MPC[14] will be evaluated at different wealth levels. As a result, we should estimate a negative interaction between refund surprise and uncertainty after controlling for mean expectations and mean tax filer characteristics. We implement this test by adding an interaction between refund surprise and the subjective standard deviation of ex-ante expectations to equation 3.2:

[14]Precisely, we measure changes in debt balances, and invoke the low levels of saving (in either real or financial assets) in our sample to translate from changes in debt balances to consumption. See also section 3.2.

$$\Delta b_i = \alpha + \beta_1 \mu_i + \beta_2 \text{surprise}_i + \beta_3 \sigma_i + \beta_4 \text{surprise}_i \times \sigma_i + \gamma Z_i + \eta_i \qquad (3.3)$$

Estimates from these specifications are shown in Table 3.8. Each pair of columns shows changes in debt at two-month and six-month horizons, and the three pairs of columns respectively show results for revolving debt, non-mortgage installment debt, and both debt categories together. The estimates are noisy, but we generally estimate negative signs on the interaction term that are consistent with precautionary behavior having induced lower MPCs. We conclude that this test for precautionary behavior is underpowered in our setting.

3.5 Conclusion

This paper uses a rich dataset linking administrative tax and credit data to surveys on taxpayer expectations and consumption behavior to shed new insight on low- and moderate-income households' choices to pay down debt, save, and consume. We showed that simple questions about an individual's expected tax refund can be used to generate rich probabilistic distributions that are informative about both mean expectations and uncertainty. We then showed that, in our sample of low-income filers, individuals face substantial uncertainty about the size of their tax refund. This is true despite the fact that annual refunds make up a substantial part of individuals' annual income and is true even for individuals whose tax situation has not changed since they last filed. Finally, we showed that refund expectations and surprises influence household financial decisions after tax filing. Filers use roughly 15 cents per dollar of expected tax refund to repay revolving debt after tax refund receipt. In contrast, refund surprises are not used to pay down debt, but rather, lead to higher borrowing through installment credit such as auto loans, consistent with MPCs potentially lying above one due to relaxed collateral constraints for financed durables. Post-refund debt repayment is most pronounced for less uncertain individuals, suggesting precautionary behavior.

There are two key limitations to our work. First, because of the small size of our sample, we are unable to generate precise estimates of the impact of uncertainty on consumption decisions. In particular, while our results are in line with buffer-stock consumption theory's predictions about how uncertainty should affect ex-ante borrowing out of expected refunds and ex-post propensities to consume out of surprises, our estimates remain somewhat noisy, especially in our tests for heterogeneous MPCs. Second, because we focus on individuals who take advantage of the city of Boston's free tax preparation services, our results may be specific to this population of low-income filers. The fact that these individuals sought out city services suggests that they may already be in a distressed financial situation. As a result they might be more responsive to income surprises than other individuals with similar incomes. However, the fact that they were aware of the city's services and were able to gather their paperwork and, in many cases, file their taxes months ahead of the deadline, suggests that they may be more conscientious on average. This would suggest that they may be less uncertain than the average taxpayer with their demographic characteristics.

One possible direction for future work would be to study how uncertainty and expectations evolve over time. While we were only able to collect one year of expectations data, it would be interesting to examine whether individuals' beliefs about their refunds become more precise if they file similar returns for several years. A second possibility would be to consider a broader sample of tax filers, who may have higher incomes or who may not not take advantage of free government-provided tax filing services. Finally, given the recent changes to the U.S. tax code, a complementary question is how taxpayers update their beliefs about their tax liabilities when the tax code itself – in addition to their own financial status – has changed.

3.6 Figures and Tables

Figure 3-1: Dorchester House Tax Site Flow

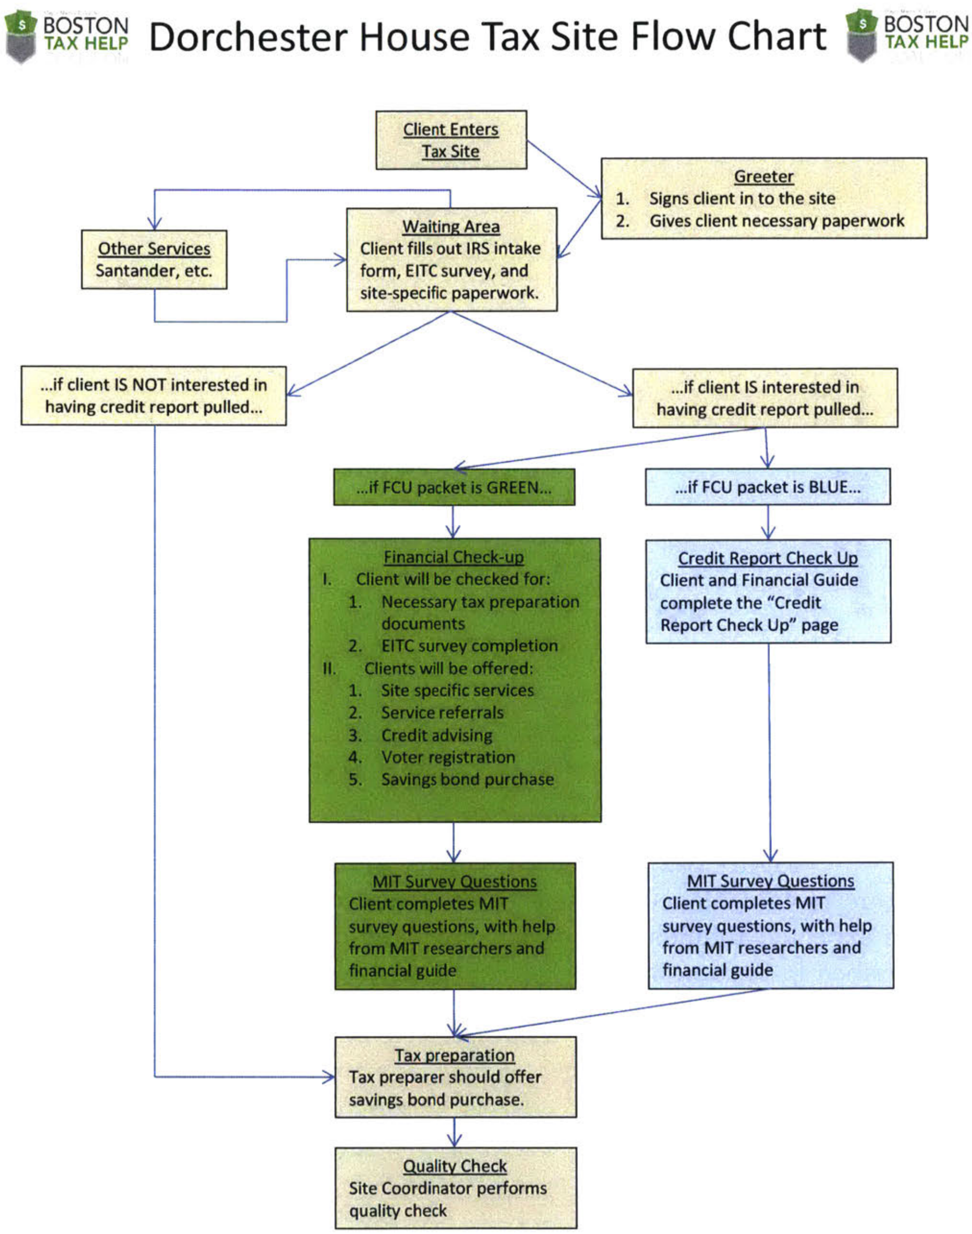

Note: This figure shows the steps a Dorchester House tax client would go through upon arriving at the center.

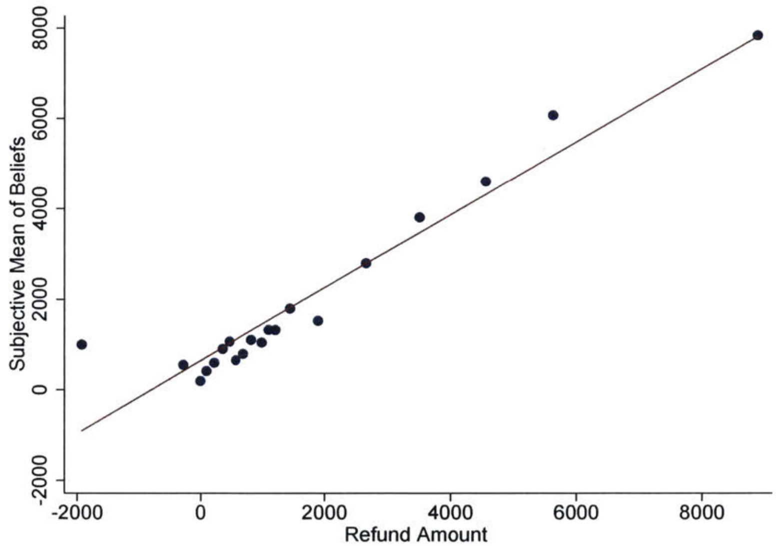

Figure 3-2: Expected Versus Actual Refunds

Note: This figure plots a binned scatterplot of mean expectations against actual refund amounts. The expected refunds are the means of the distributions calculated using the procedure described in Section 3.3.

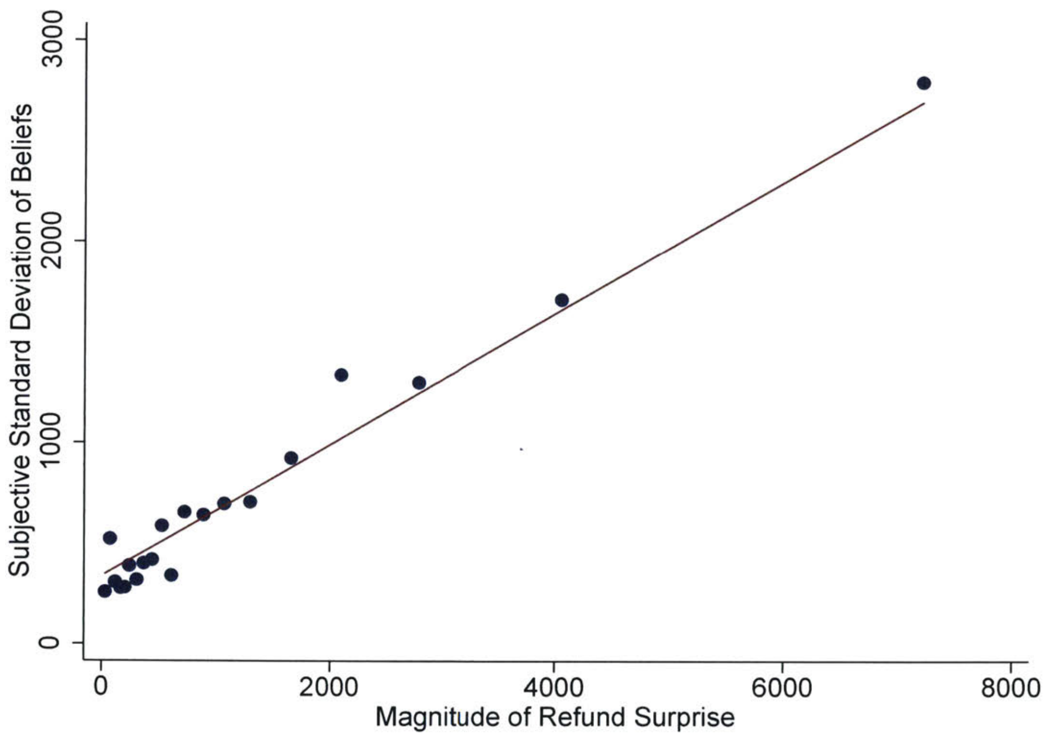

Figure 3-3: Refund Uncertainty and Refund Surprises
Note: This figure plots the size of the refund "surprise" (actual refund - mean expectation) against the standard deviation of beliefs.

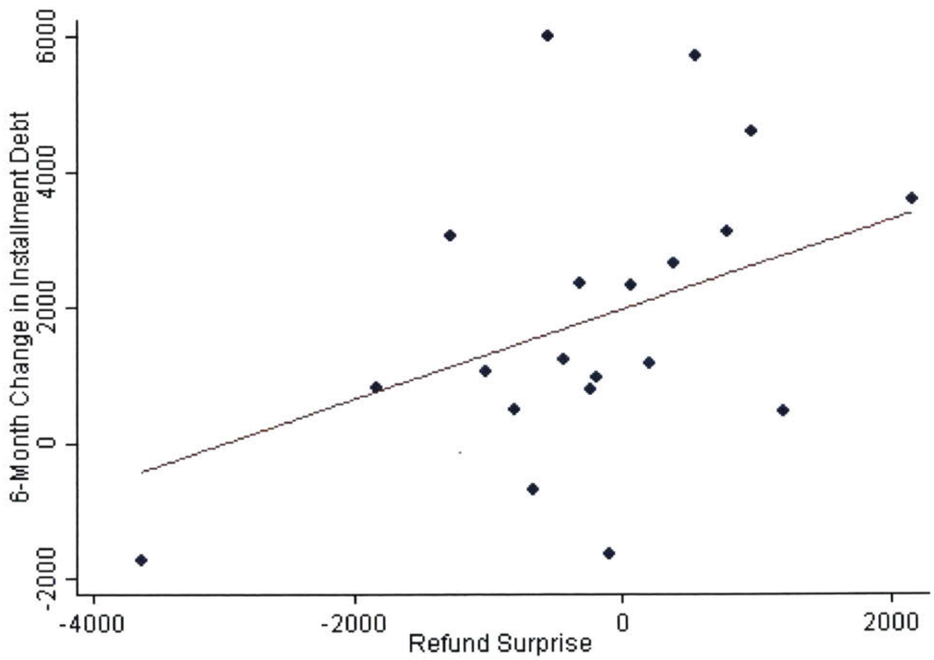

Figure 3-4: Refund Surprises and Changes in Installment Debt
Note: This figure plots a binned scatterplot of 6-month changes in installment balances against tax refund surprises. Surprises are defined as realization minus expectation, such that a positive surprise is "good news" for the tax filer. These data are plotted after partialling out the other controls included in column 4 of Table 3.5.

Table 3.1: Descriptive Statistics

	Asset Survey Sample	Credit Report Sample	Expectations and Credit Reports	Consumption, Expectations, and Credit Reports	Past Year Taxes, Expectations, and Credit Reports
	(1)	(2)	(3)	(4)	(5)
Female	0.64	0.62	0.62	0.61	0.63
Age	40.91	40.92	40.71	40.50	42.15
	(15.66)	(15.69)	(15.85)	(16.24)	(15.62)
Adjusted Gross Income ($)	$21,603	$21,535	$21,572	$21,092	$23,784
	($15,962)	($15,924)	($15,988)	($16,152)	($15,938)
Has Dependents	34%	33%	33%	32%	36%
Filing Status					
Married	10%	9%	8%	9%	6%
Single Head of Household	29%	28%	28%	28%	32%
Filed Schedule C	6%	6%	7%	8%	6%
Refund Size	$1,765	$1,677	$1,635	$1,605	$1,849
	($2,406)	($2,413)	($2,384)	($2,484)	($2,446)
Received EITC	38%	37%	35%	33%	36%
EITC Refund (If >0)	$1,769	$1,724	$1,760	$1,777	$1,940
	($1,686)	($1,644)	($1,671)	($1,749)	($1,724)
Chose Direct Deposit	59%	60%	59%	56%	62%
Total Savings Balance	$518	$525	$524	$603	$540
	($562)	($568)	($572)	($599)	($578)
High School or Above	80%	83%	83%	84%	85%
Some College or More	15%	16%	15%	15%	17%
FICO Score		664	664	671	668
		(86)	(87)	(87)	(87)
Credit Card Balances ($)		$1,732	$1,672	$1,548	$1,840
		($4,888)	($4,845)	($3,994)	($5,493)
Installment Balances ($) (non-mortgage)		$9,205	$9,171	$9,306	$10,396
		($22,500)	($22,046)	($22,264)	($24,202)
Has Mortgage		4%	4%	5%	5%
Observations	995	714	625	291	424
Obs. with Asset Survey	995	626	557	253	383

Note: This table provides descriptive statistics on our population of low-income tax filers. The first column includes all individuals who visited Dorchester House and responded to the demographics and asset survey. The second column restricts the sample to the population for whom we have both initial and follow-up credit reports. The third column includes individuals who have both credit reports and completed expectations surveys. The fourth column includes individuals with credit reports, completed expectations surveys, and consumption surveys. The fifth column includes individuals who additionally could be matched to the preceding year's tax return by virtue of being return clients. In each column, gender, savings balance, and education are provided for the subset of individuals in that column who also completed the asset survey.

Table 3.2: Comparison of Qualitative and Quantitative Uncertainty

Qualitative Uncertainty	Quantitative Uncertainty				
		Coefficient of Variation		p-Value from t-test for Equality of Means	
	N	Mean	S.D.	Not Sure	Somewhat Sure
Not Sure at All	149	0.79	3.73		
Somewhat Sure	254	0.45	0.65	0.00038	
Very Sure	215	0.38	0.66	0.00002	0.00000

Note: This table compares the coefficient of variation calculated for the parametric belief distributions to the qualitative uncertainty responses. The sample includes all individuals who responded to the expectations survey.

Table 3.3: Correlates of Refund Uncertainty and Surprises

	Elicited Standard Deviation of Refund Amount			Magnitude of Refund Surprise			Refund Surprise		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
25 or Younger	62.90	-93.53	-66.94	19.34	304.2	485.8	297.0	398.8	487.4
	(108.2)	(161.9)	(160.3)	(209.0)	(317.4)	(315.8)	(271.9)	(434.9)	(430.3)
Older than 50	-171.2*	-134.7	-130.1	-166.6	136.2	79.31	308.7	222.2	143.8
	(98.69)	(126.3)	(123.8)	(190.3)	(247.9)	(244.2)	(247.6)	(339.6)	(332.7)
Any College	-24.99	-33.90	-2.697	101.3	143.9	151.1	204.2	172.0	33.05
	(83.72)	(113.3)	(112.2)	(161.9)	(222.9)	(221.8)	(210.6)	(305.4)	(302.2)
Female	-79.07	-89.00	-77.56	-51.55	-149.0	-162.2	157.5	229.0	204.9
	(87.97)	(124.4)	(121.4)	(170.4)	(244.4)	(239.7)	(221.6)	(334.9)	(326.5)
Has Dependent	347.8**	545.5**	516.7**	676.8**	1162.6***	823.1**	-172.1	-466.4	-504.0
	(168.0)	(219.0)	(211.2)	(323.5)	(428.8)	(415.7)	(420.8)	(587.5)	(566.4)
Number of Dependents	448.7***	321.0**	377.9***	494.9***	296.4	663.2***	312.2	565.7*	455.8
	(90.97)	(126.7)	(118.1)	(175.1)	(248.0)	(232.4)	(227.7)	(339.8)	(316.7)
Married	143.1	309.1	390.6	184.2	198.3	227.0	-967.3**	-1383.6**	-1598.6**
	(159.8)	(261.7)	(254.7)	(312.1)	(512.2)	(500.9)	(406.0)	(701.7)	(682.5)
AGI in 2nd Quartile	88.16	67.49	53.96	330.5	281.1	407.0	558.0*	696.2	891.5*
	(113.1)	(171.1)	(170.2)	(218.5)	(334.9)	(334.8)	(284.3)	(458.9)	(456.2)
AGI in 3rd Quartile	608.8***	489.7***	480.2***	1000.8***	925.2***	1031.3***	-686.8**	-422.3	-112.2
	(115.2)	(168.4)	(169.8)	(222.6)	(330.4)	(334.8)	(289.5)	(452.7)	(456.2)
AGI in 4th Quartile	289.4**	174.5	187.6	501.4**	93.67	304.4	-460.5	154.8	485.7
	(118.3)	(174.4)	(178.1)	(228.2)	(341.2)	(350.3)	(296.8)	(467.5)	(477.3)
Received UI in Past Year		27.56	28.98		-336.2	-278.9		752.2	758.6
		(221.5)	(216.7)		(433.5)	(426.2)		(593.9)	(580.7)
Change in Filing Status		-121.1	-228.1		-723.6	54.14		-43.27	69.35
		(266.0)	(187.6)		(520.7)	(369.1)		(713.4)	(502.9)
Magnitude of Change in AGI ($1,000)		8.298			28.48**			-6.706	
		(6.555)			(12.83)			(17.58)	
Any Change in Number of Dependents		-130.7			1011.5**			173.3	
		(235.8)			(461.5)			(632.3)	
Change in AGI ($1,000)			-1.726			-22.53**			-28.01**
			(5.206)			(10.24)			(13.95)
Change in Number of Dependents			-296.3***			-517.3***			736.2***
			(92.90)			(182.8)			(249.0)
Constant	256.3**	292.9*	280.2	326.7	124.1	154.7	-592.7**	-873.7*	-899.1*
	(114.5)	(177.3)	(172.1)	(220.6)	(347.4)	(338.8)	(286.9)	(476.0)	(461.6)
N	463	268	268	460	267	267	460	267	267
R-squared	0.348	0.333	0.358	0.222	0.281	0.301	0.072	0.083	0.118

Note: This table reports estimates from ordinary least squares regressions of refund uncertainty and surprises on tax filer characteristics. The sample in columns 1, 4, and 7 is all Dorchester House tax filers who completed the assets and beliefs surveys. The sample in the remaining columns is the subset of those tax filers who could be linked to the previous year's tax return by virtue of being a repeat client. The Elicited Standard Deviation of Refund Amount (columns 1-3) is the standard deviation of the parametric belief distribution fit to each tax filer's probabilistic survey question response, described in Section 3.3.2. Magnitude of Refund Surprise (columns 4-6) is the absolute value of the difference between each tax filer's refund amount and the mean of their parametric belief distribution. Refund Surprise (columns 7-9) is the (signed) level of the same quantity. AGI is Adjusted Gross Income as reported on the tax return. Change in Filing Status is an indicator for whether a household's filing status changed from the previous year of filing. Standard errors are in parentheses. * $p < .1$ ** $p < 0.05$ *** $p < 0.01$

Table 3.4: Impact of Refund Surprise on Revolving Debt Balances

	Dependent Variable: Change in Revolving Debt ($)			
	2-Month Follow-Up		6-Month Follow-Up	
	(1)	(2)	(3)	(4)
Mean Expectation ($)	-0.150***	-0.103*	-0.145**	-0.142*
	(0.0457)	(0.0569)	(0.0671)	(0.0826)
Surprise ($)	-0.0677*	-0.0393	-0.0831	-0.0831
	(0.0370)	(0.0462)	(0.0546)	(0.0674)
S.D. of Beliefs	0.395**	0.378**	0.377*	0.358
	(0.154)	(0.168)	(0.226)	(0.243)
Controls for Taxpayer Characteristics		X		X
R-Squared	0.042	0.198	0.021	0.204
N	302	302	301	301

Note: This table reports estimates from ordinary least squares regressions of changes in revolving debt balances on tax filer expectations, refund surprises, and characteristics. The sample in all columns is individuals with non-missing data on demographics as measured in the asset survey, expectations, and an open revolving loan observed at any point during the sample period. Balances of zero are assigned to loans reported as closed with no balance. In columns 1 and 2, the dependent variable is the change in revolving debt between the week of tax filing and the two-month credit report follow-up. Column 1 controls for the mean and standard deviation of each tax filer's parametric belief distribution, fit as described in Section 3.3.2, as well as their refund surprise. Column 2 adds controls for tax filer characteristics. Columns 3 and 4 repeat these specifications for the six-month change in revolving debt. Tax filer characteristics are fully interacted bins of age less than 25, 25-50, and over 50; adjusted gross income (AGI) quartile; marital status; and an indicator for any dependents. Standard errors are in parentheses. * $p < .1$ ** $p < 0.05$ *** $p < 0.01$

Table 3.5: Impact of Refund Surprise on Installment Debt Balances

	Dependent Variable: Change in Installment Debt ($)			
	2-Month Follow-Up		6-Month Follow-Up	
	(1)	(2)	(3)	(4)
Mean Expectation ($)	-0.128	-0.233	-0.0231	-0.303
	(0.150)	(0.196)	(0.325)	(0.416)
Surprise ($)	0.0599	0.0318	0.676**	0.659*
	(0.125)	(0.158)	(0.272)	(0.338)
S.D. of Beliefs	-0.353	-0.146	-0.539	-0.374
	(0.478)	(0.527)	(1.036)	(1.115)
Controls for Taxpayer Characteristics		X		X
R-Squared	0.036	0.218	0.043	0.260
N	216	216	215	215

Note: This table reports estimates from ordinary least squares regressions of changes in installment debt balances on tax filer expectations, refund surprises, and characteristics. The sample in all columns is individuals with non-missing data on demographics as measured in the asset survey, expectations, and an open non-mortgage installment loan observed at any point during the sample period. Balances of zero are assigned to loans reported as closed with no balance.. In columns 1 and 2, the dependent variable is the change in installment debt between the week of tax filing and the two-month credit report follow-up. Column 1 controls for the mean and standard deviation of each tax filer's parametric belief distribution, fit as described in Section 3.3.2, as well as their refund surprise. Column 2 adds controls for tax filer characteristics. Columns 3 and 4 repeat these specifications for the six-month change in installment debt. Tax filer characteristics are fully interacted bins of age less than 25, 25-50, and over 50; adjusted gross income (AGI) quartile; marital status; and an indicator for any dependents. Standard errors are in parentheses. * $p < .1$ ** $p < 0.05$ *** $p < 0.01$

Table 3.6: Impact of Refund Surprise on Non-Mortgage Debt Balances

	Dependent Variable: Change in All Non-Mort. Debt ($)			
	2-Month Follow-Up		6-Month Follow-Up	
	(1)	(2)	(3)	(4)
Mean Expectation ($)	-0.163**	-0.106	-0.223	-0.348
	(0.0772)	(0.0960)	(0.206)	(0.255)
Surprise ($)	-0.0194	0.0303	0.384**	0.382*
	(0.0641)	(0.0813)	(0.172)	(0.217)
S.D. of Beliefs	0.183	0.165	0.673	0.789
	(0.256)	(0.272)	(0.683)	(0.723)
Controls for Taxpayer Characteristics		X		X
R-squared	0.024	0.158	0.021	0.167
N	352	352	351	351

Note: This table reports estimates from ordinary least squares regressions of changes in total non-mortgage debt balances on tax filer expectations, refund surprises, and characteristics. The sample in all columns is individuals with non-missing data on demographics as measured in the asset survey, expectations, and an open revolving or non-mortgage installment loan observed at any point during the sample period. Balances of zero are assigned to loans reported as closed with no balance.. In columns 1 and 2, the dependent variable is the change in debt between the week of tax filing and the two-month credit report follow-up. Column 1 controls for the mean and standard deviation of each tax filer's parametric belief distribution, fit as described in Section 3.3.2, as well as their refund surprise. Column 2 adds controls for tax filer characteristics. Columns 3 and 4 repeat these specifications for the six-month change in debt. Tax filer characteristics are fully interacted bins of age less than 25, 25-50, and over 50; adjusted gross income (AGI) quartile; marital status; and an indicator for any dependents. Standard errors are in parentheses. * $p < .1$ ** $p < 0.05$ *** $p < 0.01$

Table 3.7: Impact of Refund Surprise on Durable Purchases

	Dependent Variable: Survey-Reported Durables Purchase	
	2-Month Follow-Up	
	(1)	(2)
Mean Expectation ($)	0.0000143	0.0000197
	(0.0000138)	(0.0000176)
Surprise ($)	0.00000182	0.00000873
	(0.0000118)	(0.0000148)
S.D. of Beliefs	-0.0000228	-0.0000398
	(0.0000442)	(0.0000476)
Controls for Taxpayer Characteristics		X
R-squared	0.003	0.127
N	443	443

Note: This table reports estimates from ordinary least squares regressions of an indicator for durable purchases between tax filing and the two-month follow-up consumption survey. Controls are tax filer expectations, refund surprises, and characteristics. The sample in all columns includes individuals with non-missing data on demographics as measured in the asset survey, expectations, and a response to the follow-up consumption survey conducted approximately two months after tax refund receipt. Column 1 controls for the mean and standard deviation of each tax filer's parametric belief distribution, fit as described in Section 3.3.2, as well as their refund surprise. Column 2 adds controls for tax filer characteristics. Tax filer characteristics are fully interacted bins of age less than 25, 25-50, and over 50; adjusted gross income (AGI) quartile; marital status; and an indicator for any dependents. Standard errors are in parentheses. * $p < .1$ ** $p < 0.05$ *** $p < 0.01$

Table 3.8: Testing Concavity of the Consumption Function

Dependent Variable	Change in Revolving Debt		Change in Installment Debt		Change in All Non-Mort. Debt	
Horizon	2-Mo.	6-Mo.	2-Mo.	6-Mo.	2-Mo.	6-Mo.
	(1)	(2)	(3)	(4)	(5)	(6)
Surprise ($)	-0.0697	-0.112	0.0839	0.922**	0.0291	0.541**
	(0.0569)	(0.0833)	(0.197)	(0.422)	(0.102)	(0.272)
Surprise * S.D. of Beliefs	0.0000395	0.0000376	-0.0000639	-0.000319	0.00000151	-0.000190
	(0.0000431)	(0.0000628)	(0.000144)	(0.000305)	(0.0000735)	(0.000196)
S.D. of Beliefs	0.409**	0.387	-0.197	-0.627	0.166	0.622
	(0.172)	(0.249)	(0.540)	(1.141)	(0.280)	(0.743)
Controls for Taxpayer Characteristics	X	X	X	X	X	X
Controls for Mean Expectation	X	X	X	X	X	X
R-squared	0.200	0.205	0.219	0.264	0.158	0.170
N	302	301	216	215	352	351

Note: This table reports estimates from ordinary least squares regressions of changes in debt balances on tax filer expectations, refund surprises, and tax filer characteristics. The sample in all columns is individuals with non-missing data on demographics as measured in the asset survey, expectations, and an open loan (either revolving loan, installment loan, or either type of loan, depending on the column) observed at any point during the sample period. Balances of zero are assigned to loans reported as closed with no balance. In columns 1 and 2, the dependent variable is change in two- and six-month revolving debt balances, respectively. Both columns control for the mean and standard deviation of each tax filer's parametric belief distribution, fit as described in Section 3.3.2, as well as their refund surprise, an interaction between the surprise and belief standard deviation, and tax filer characteristics. Columns 3 and 4 repeat these specifications for changes in installment debt balances, and columns 5 and 6 do the same for all non-mortgage debt. Tax filer characteristics are fully interacted bins of age less than 25, 25-50, and over 50; adjusted gross income (AGI) quartile; marital status; and an indicator for any dependents. Standard errors are in parentheses. * $p < .1$ ** $p < 0.05$ *** $p < 0.01$

3.7 Appendix

3.7.1 Appendix Tables and Figures

Table 3.9: Elicited Beliefs by Tax Filer Group

	Features of Probabilistic Survey Question Responses								
	Full Sample	Has Dependents		Marital Status		Adjusted Gross Income (AGI)		Education	
		Yes	No	Married	Single	Above $20,000	Below $20,000	Some College	No College
Number of Bins with Positive Probability									
1 Bin	22.2%	24.0%	21.4%	28.3%	21.7%	24.7%	20.1%	21.2%	23.0%
2 Bin	38.8%	39.9%	38.3%	32.6%	39.3%	34.4%	42.5%	37.3%	39.9%
3 Bin	20.3%	14.2%	23.2%	13.0%	20.9%	18.1%	22.1%	18.7%	21.5%
4 Bin	11.8%	12.0%	11.7%	13.0%	11.7%	13.9%	10.1%	14.5%	9.8%
5 Bin	5.5%	8.2%	4.2%	8.7%	5.2%	6.9%	4.2%	6.6%	4.6%
6 Bin	1.4%	1.6%	1.3%	4.3%	1.2%	1.9%	1.0%	1.7%	1.2%
Qualitative Uncertainty									
Very Sure	34.4%	30.6%	36.3%	48.0%	33.2%	29.9%	38.4%	32.4%	35.8%
Somewhat Sure	40.6%	47.6%	37.2%	34.0%	41.2%	43.2%	38.4%	40.2%	41.0%
Not Sure at All	23.8%	21.4%	25.1%	18.0%	24.3%	25.5%	22.4%	25.9%	22.4%
Quantitative Responses									
Point Estimate	1,758	3,466	921	2,336	1,708	2,377	1,202	1,753	1,762
Minimum	-364	1,071	-1,048	-304	-369	-75	-607	-527	-244
Maximum	5,922	10,885	3,557	7,783	5,758	7,851	4,300	6,344	5,610
Features of Parametric Distribution									
Mean	1,970	4,211	902	2,891	1,889	2,817	1,258	1,995	1,952
Median	2,073	4,225	1,047	2,768	2,011	2,889	1,386	2,089	2,061
Std. Dev.	740	1,475	390	1,052	713	1,023	502	803	693
Coefficient of Variation	0.50	0.33	0.58	0.35	0.51	0.37	0.61	0.57	0.45

Notes: This table reports responses to the beliefs survey. All statistics are means within each group. The last panel contains statistics based on the parametric distributions fit to the probabilistic survey question described in Section 3.3.

Table 3.10: Parametric Belief Distributions

	Features of Elicitations under Alternative Parametric Assumptions and Sample Restrictions											
	Baseline		Uniform		Lower Bound		Upper Bound		Omitting Only Top/Bottom Bin		Omitting 50-50 Reports	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Mean	1,970	2,850	2,312	3,297	1,098	1,418	3,527	5,212	1,566	2,047	2,061	2,992
Median	2,073	3,369	2,066	3,171	996	1,456	3,179	5,413	1,713	2,874	2,179	3,554
Std. Dev.	740	1,019	2,312	2,976					631	885	799	1,075
Minimum	548	1,261	525	1,251					353	816	553	1,315
Maximum	5,247	6,765	5,922	7,333					4,641	6,111	5,698	7,093

Notes: This table reports features of parametric belief distributions under alternative assumptions. Statistics are aggregated across all tax filers in the main analysis sample. The first pair of columns contains statistics based on the parametric distributions fit to the probabilistic survey question described in Section 3.3. Uniform assumes a uniform distribution within each bin with nonzero probability. Lower (Upper) bound calculates the lowest (highest) value of each tax filer's subjective mean and median expectation that is consistent with their subjective probabilities. The last two pairs of columns implement sample restrictions using the baseline parametric assumptions.

3.7.2 Survey Appendix

On-Site Survey

PLEASE DO NOT WRITE IN THIS AREA

02403

Mayor Martin J. Walsh
BOSTON TAX HELP COALITION

FEDERAL RESERVE BANK OF BOSTON™

2015 BOSTON TAX HELP SURVEY

The overall purpose of this survey is to inform how the services and resources available in the community can be improved. Your answers will be added with the answers we receive from other surveys. We will never share your personal information with anyone without your permission.

• USE A NO. 2 PENCIL ONLY

- Make solid marks that fill the oval completely
- Erase cleanly any marks you wish to change

1. GENDER? ○ Female ○ Male

2. ARE YOU A VETERAN? ○ Yes ○ No

3. ARE YOU REGISTERED TO VOTE? ○ Yes ○ No

4. HOW WOULD YOU DESCRIBE YOUR RACE OR ETHNICITY? (please mark all that apply)
○ Asian/Pacific Islander ○ Hispanic/Latino ○ White ○ African
○ Black/African American ○ Native American ○ Caribbean ○ Other ______

5. HOW FLUENT ARE YOU IN ENGLISH?
Please rate from 0 (beginner) to 5 (native speaker): **Beginner** ⓪ ① ② ③ ④ ⑤ **Native Speaker**

6. WHAT IS THE HIGHEST LEVEL OF EDUCATION YOU HAVE COMPLETED? (please mark one)
○ 0–8 grades ○ High School grad or GED ○ Associate's Degree ○ Some graduate school
○ 9–12 grades (no diploma) ○ Some college ○ Bachelor's Degree ○ Graduate degree

7. DO YOU WANT TO CONTINUE YOUR EDUCATION OR VOCATIONAL TRAINING? (please mark all that apply)
○ Yes, but I don't have enough time ○ Yes, but I am not sure investing in additional education or training will pay off for me
○ Yes, but I don't have enough money to cover all the cost ○ No, I am all set

8. LAST YEAR, WHAT BENEFITS DID YOU OR YOUR FAMILY RECEIVE? (please mark all that apply)
○ WIC ○ SNAP/food stamps ○ TAFDC ○ Child Care Vouchers ○ SSI ○ SSDI ○ None

9. HOW MANY MONTHS WERE YOU EMPLOYED LAST YEAR?
Full-time: ○ 0 ○ 1–3 ○ 4–6 ○ 7–9 ○ 10–12
Part-time: ○ 0 ○ 1–3 ○ 4–6 ○ 7–9 ○ 10–12

10. WHICH OF THE FOLLOWING BEST DESCRIBES YOUR HOUSING STATUS? (please mark one)
○ Rent, no subsidy ○ Rent, public housing ○ Living with family or friends ○ Homeless, in shelter ○ Own a home with a mortgage
○ Rent, Section 8 ○ Rent, other subsidy ○ Assisted Living Facility ○ Homeless, no shelter ○ Own a home with no mortgage

11. DID YOU LOSE YOUR HOUSING THROUGH EVICTION OR FORECLOSURE IN THE LAST YEAR? (please mark all that apply)
○ Yes, foreclosure ○ Yes, eviction ○ No

12. DID YOU FILE TAXES LAST YEAR (2014)? ○ Yes ○ No **If No, Then Skip To 15.**

13. Did you receive a refund for the 2014 tax year?
○ Yes, and it was more than $1,000 ○ No **If No, Then Skip To 15.**
○ Yes, and it was less than $1,000 ○ Don't Remember

14. If you received a refund, were you able to use it as you intended?
○ Yes, pretty much as I had planned
○ No, I had planned to save more of my refund
○ No, I had planned to use it for something else but ended up spending most of it on bills
○ I did not have any specific plans for the refund
○ I don't remember

***15. If you get a tax refund this year, what do you plan to do with the money?* (please mark all that apply)**
○ Buy groceries ○ Pay bills ○ Home improvement
○ Pay child expenses (including K-12 school fees) ○ Pay medical bills ○ Pay for college
○ Buy a car ○ Pay moving expenses ○ Save for college
○ Car repairs/expenses ○ Pay down debt ○ Save for retirement
○ Save for emergencies ○ Go on vacation ○ Other ______

16. DO YOU HAVE A BANK ACCOUNT?
○ Checking & Saving ○ No, but I am interested in one **If No, Then Skip To 21.**
○ Just checking ○ No, not interested **If No, Then Skip To 21.**
○ Just savings

***17. If you have bank account(s), with whom do you bank?* (please mark all that apply)**
○ Bank of America ○ TD Bank ○ Capital One 360
○ Citibank ○ One United ○ Century Bank
○ Citizens ○ Santander ○ Webster First Financial
○ City of Boston Credit Union ○ Cathay ○ Other ______
○ East Boston Savings/Mt. Washington ○ Mass State Employees Credit Union
○ Eastern ○ Metro Credit Union

SCANTRON Mark Reflex® EM-257980-12:654321 AS99

PLEASE CONTINUE ON BACK

18. *Did you pay an overdraft fee in the past 12 months?*

- ○ Yes, once ○ Yes, more than once ○ No

19. *If you have bank account(s), how much money do you regularly keep in it (them) all together?*

- ○ $0 ○ $1 - $100 ○ $101 - $500 ○ $501 - $1,000 ○ More than $1,000

20. *On a monthly basis, how much money do you regularly put into savings?*

- ○ $0 ○ Between $1 - $50 ○ Between $51 - $100 ○ More than $100

If you have a bank account, Then Skip To 22.

21. *If you do not have bank account(s), why not?* (please mark all that apply)

- ○ Overdraft fees
- ○ Other fees
- ○ I don't trust them
- ○ Goes against my religious beliefs
- ○ I don't think I can get an account
- ○ They don't speak my language
- ○ Locations or hours are not good for me
- ○ Isn't really worth it
- ○ Other ____________________

22. IN THE EVENT OF A FINANCIAL HARDSHIP, DO YOU KNOW ANYONE LIKELY TO LOAN YOU $200 OR MORE IF YOU ASKED? (please mark one)

- ○ No, Nobody
- ○ Yes, Maybe 1 or 2 people
- ○ Yes, Maybe 3 or 4 people
- ○ Yes, Maybe 5 or more people

23. *In 2014, did you have a friend or relative that borrowed $200 or more from you because of a financial hardship?*

- ○ No
- ○ Yes, and I expect to get paid back
- ○ Yes, and I do not expect to get paid back

24. DO YOU HAVE ONE OR MORE CREDIT CARDS?

- ○ Yes, one credit card
- ○ Yes, more than one credit card
- ○ No, but I want a credit card **If no, Then Skip To 28.**
- ○ No, I don't want a credit card **If no, Then Skip To 28.**

25. *How much do you usually pay each month on your highest balance card?*

- ○ The full balance ○ More than the minimum payment ○ The minimum payment ○ Less than the minimum payment

26. *What is your current outstanding balance on all credit cards?*

- ○ Less than $1,000 ○ $1,000 - $5,000 ○ $5,000 - $10,000 ○ More than $10,000 ○ Unsure

27. *What is the interest rate you are charged on your credit card with the highest balance? If you are unsure, please approximate.*

- ○ Less than 10% ○ 10% - 14.9% ○ 15% - 19.9% ○ More than 20% ○ Don't Know (Can't even guess)

28. DID YOU EXPERIENCE ANY OF THE FOLLOWING IN THE LAST YEAR (2014) WITH CREDIT CARD(S)? (please mark all that apply)

- ○ Missed a credit card payment
- ○ Went over the credit card limit
- ○ Credit Limit decreased
- ○ Credit Limit increased
- ○ Interest Rate changed
- ○ Offered a credit card to pay a medical expense
- ○ I don't know
- ○ Did not have a credit card in the last year (2014)

29. *Did you use a credit card in the last year (2014) for any of the following reasons because you had no access to cash to pay for them?* (please mark all that apply)

- ○ Medical Bills ○ Prescriptions ○ Groceries ○ Utilities ○ Phone ○ Did not have a credit card in the last year (2014)

30. DO YOU KNOW WHAT A CREDIT SCORE IS?

- ○ No ○ Yes, I know but it is not important for me ○ Yes, I know and it is important for me

31. How would you rate your credit score?

- ○ Very Bad
- ○ Bad
- ○ Fair
- ○ Good
- ○ Very Good

32. How easy do you think it is to improve your credit score?

- ○ Very Easy
- ○ Easy
- ○ Neutral
- ○ Hard
- ○ Very Hard

33. How often do you get your credit report?

- ○ Never **If never, Then Skip To 35.**
- ○ Less than once per year
- ○ Once per year
- ○ More than once per year

34. How do you get your credit report?

- ○ Request it on-line through AnnualCreditReport.com
- ○ Request it by mail through AnnualCreditReport.com
- ○ Access it online for a fee
- ○ A different method
- ○ Through the tax sites

35. DURING 2014, DID YOU DO ANY OF THE FOLLOWING? (please mark all that apply)

- ○ Attend a credit counseling session
- ○ Experience a financial crisis
- ○ Filed a credit report complaint
- ○ Lose a job
- ○ Open a checking account
- ○ Open a retirement account
- ○ Open a savings account
- ○ Pay a bill over the internet
- ○ Purchase a Money Order
- ○ Purchase a US Savings Bond
- ○ Received a court order for debt collection
- ○ Received a notice for debt you don't owe
- ○ Received threatening phone calls for debt collection
- ○ Take a class on personal finances
- ○ Take a loan from a payday lender
- ○ Take out a loan from a bank
- ○ Use a Check Casher (Western Union, PL$, MoneyGram, etc.)
- ○ Use Direct Deposit (paycheck, gov't check, etc.)
- ○ Used a pre-paid debit card with fees
- ○ Used mobile banking on a smartphone

36. WOULD YOU LIKE HELP WITH ANY OF THE FOLLOWING? (please mark all that apply)

- ○ Emergency/short term loans
- ○ Foreclosure help/prevention
- ○ Small business loans
- ○ Credit advising
- ○ Low cost-prepaid cards
- ○ Helping me work with my bank
- ○ Budgeting help
- ○ Debt Management
- ○ Continuing education
- ○ Employment

Thank you for your participation

3.7.3 Expectations Survey

The expectations survey consisted of four questions, printed below. The survey was administered by the financial guides at Dorchester House. Along with the answers to these four questions, financial guides recorded each individual's tax client number so that the survey could be linked to the other data we collected.

1) If you get a tax refund this year, how much do you think it will be? Please choose an amount:

$________________

(Financial Guide volunteer: please write $500 above this number, and $500 below this number, in the two blank lines in the question below)

2) How sure are you that your refund will be between $________ and $________? Please circle one:

NOT SURE AT ALL　　SOMEWHAT SURE　　VERY SURE

3) Suppose you want to make some extra money by working more hours next week. Do you think you could you get your manager/supervisor to schedule you for more hours?

YES

NO

I AM NOT WORKING RIGHT NOW

I AM NOT PAID HOURLY

4) We have one final question about your tax refund. Below we show six possible amounts that your refund could be (for example, "between $1000 and $2500"). For each of the six possibilities, please say what is the "percent chance" that you think your refund could be that amount:

Could my refund be...	**(Please Enter % Chance for Each)**
Over $5000	%
Between $2500 and $5000	%
Between $1000 and $2500	%
Between $500 and $1000	%
Between $0 and $500	%
Negative: I will owe taxes	%

Follow-Up Survey

The follow-up survey was conducted via phone by a Dorchester House volunteer. After introducing herself and reading the consent statement, the volunteer went through a pre-specified script and coded the answers into a spreadsheet. Individuals who completed the survey were mailed a $10 gift card.

Consent Statement: The survey information will be stored securely at the City of Boston's Office of Financial Empowerment, will be kept confidential, and will only be accessed by OFE employees. The information will also used as part of ongoing research with researchers at MIT. All survey questions are voluntary and you can stop the survey at any time. Participation will not affect your eligibility for city services. The survey should take about 4 minutes. To thank you for your participation, you will be given a $10 gift card at the end of the survey.

Questions

1. Have you made any of the following large purchases in 2016?

    (a) Car or motorcycle
    (b) Large household appliance, for example a dishwasher, refrigerator, or clothes dryer
    (c) A major repair to your home or the place you live
    (d) Television or computer
    (e) Car repairs
    (f) Wedding, funeral, or party expenses
    (g) [Repeat for each of the items purchased:]

        i. About when was it that you purchased? How certain are you of this date?

ii. How much did it cost?
iii. How did you pay for it? (cash/check/credit...)

2. Have you faced any unexpected expensive life events, such as job loss, job change, or medical bills, in 2016?

   (a) [Repeat for each event:]

      i. About when did — happen? How certain are you of this date?
      ii. If applicable: how much did the expense cost, and how did you pay for it?

3. About what time did you receive your tax refund this year?

   (a) How long was it after you first came to Dorchester House to file taxes?

4. Did you use your tax refund to put more money in a savings or checking account?

5. OK, now I have just one last question about the things we've discussed so far.

   (a) [Repeat for each large purchase or life event in Questions 1 and 2:]

      i. Do you recall if — was before or after you got your tax refund?
      ii. Do you recall if that was before or after you came to Dorchester House to file taxes?

Bibliography

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 Assignment: The New York City High School Match," *American Economic Review*, 2017, *95* (2), 364–367.

— **, — , and —** , "The Welfare Effects of Coordinated School Assignment: Evidence from the NYC High School Match," *American Economic Review, forthcoming*, 2017.

Ackerberg, Daniel A., "A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation," *Quantitative Marketing and Economics*, 2009, *7* (4), 343–376.

Agarwal, Nikhil, "An Empirical Model of the Medical Match," *American Economic Review*, 2015, *105* (7), 1939–1978.

— **and Paulo Somaini**, "Demand Analysis Using Strategic Reports: An Application to a School Choice Mechanism," *NBER Working Paper*, 2015, *No. 20775.*

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

Agarwal, Sumit, Chunlin Liu, and Nicholas S Souleles, "The Reaction of Consumer Spending and Debt to Tax Rebates—Evidence from Consumer Credit Data," *Journal of Political Economy*, 2007, *115* (6), 986–1019.

Aghion, Philippe, Ufuk Akcigit, Matthieu Lequien, and Stefanie Stantcheva, "DP12471 Tax Simplicity and Heterogeneous Learning," 2017.

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.

Akerlof, George, "The Economics of "Tagging" as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning," *American Economic Review*, 1978, *68* (1), 8–19.

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Poi, "Self-Targeting : Evidence from a Field Experiment in Indonesia," *Journal of Political Economy*, 2016, *124* (2).

Allcott, Hunt and Judd B Kessler, "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons," *American Economic Journal: Applied Economics*, 2018.

Andersson, Fredrik, Mark J. Kutzbach, Giordano E. Palloni, Henry O. Pollakowski, and Daniel H. Weinberg, "Childhood Housing and Adult Earnings: A Between-Siblings Analysis of Housing Vouchers and Public Housing," *NBER Working Paper 22721*, 2016.

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.

Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert Van der Klaauw, and Basit Zafar, "The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment," *Review of Economics and Statistics*, 2016, *98* (3), 503–523.

Arnosti, Nick and Peng Shi, "Non-Monetary Dynamic Allocation to Agents with Unit Demand," 2017, pp. 1–46.

Atkinson, Anthony B, "On the Measurement of Inequality," *Journal of Economic Theory*, 1970, *2*, 244–263.

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

Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A Parker, "A Test of Consumption Smoothing and Liquidity Constraints: Spending Responses to Paying Taxes and Receiving Refunds," Technical Report, Massachusetts Institute of Technology March 2018.

Benzarti, Youssef, "How Taxing Is Tax Filing? Using Revealed Preferences to Estimate Compliance Costs.," Technical Report, National Bureau of Economic Research 2017.

Berry, Steven, Ariel Pakes, and James Levinsohn, "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," *Journal of Political Economy*, 2004, *112* (1), 68–105.

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*, 2017, *9* (1), 88–122.

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

Broda, Christian and Jonathan A Parker, "The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption," *Journal of Monetary Economics*, 2014, *68*, S20–S36.

Carroll, Christopher D, "Buffer Stock Saving: Some Theory," September 1996.

_ **and Andrew A Samwick**, "How Important Is Precautionary Saving?," *Review of Economics and Statistics*, 1998, *80* (3), 410–419.

_ **and Miles S Kimball**, "On the Concavity of the Consumption Function," *Econometrica: Journal of the Econometric Society*, 1996, pp. 981–992.

_ **, Karen E Dynan, and Spencer D Krane**, "Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets," *Review of Economics and Statistics*, 2003, *85* (3), 586–604.

Chade, Hector and Lones Smith, "Simultaneous Search," *Econometrica*, 2006, *74* (5), 1293–1307.

Chetty, Raj and Emmanuel Saez, "Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients," *American Economic Journal: Applied Economics*, 2013, *5* (1), 1–31.

_ **, John N Friedman, and Emmanuel Saez**, "Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings," *American Economic Review*, 2013, *103* (7), 2683–2721.

_ **, Nathaniel Hendren, and Lawrence F. Katz**, "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportuity Experiment," *American Economic Review*, 2015, *106* (4), 1–88.

Collinson, Robert, Ingrid Gould Ellen, and Jens Ludwig, "Low-Income Housing Policy," *NBER Working Paper 21071*, 2015.

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

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

Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica*, 1991, *59* (5), 1221–1248.

Delavande, Adeline, Xavier Giné, and David McKenzie, "Measuring Subjective Expectations in Developing Countries: A Critical Review and New Evidence," *Journal of Development Economics*, 2011, *94* (2), 151–163.

Deshpande, Manasi and Yue Li, "Who Is Screened Out? Application Costs and the Targeting of Disability Programs," *NBER Working Paper 23472*, 2017.

Desmond, Matthew and Weihua An, "Neighborhood and Network Disadvantage among Urban Renters," *Sociological Science*, 2015, *2*, 329–350.

Dynan, Karen E, "How Prudent are Consumers?," *Journal of Political Economy*, 1993, *101* (6), 1104–1113.

Engelberg, Joseph, Charles F Manski, and Jared Williams, "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters," *Journal of Business & Economic Statistics*, 2009, *27* (1), 30–41.

Fack, Gabrielle, Julien Grenet, and Yinghua He, "Beyond Truth-Telling: Preference Estimation with Centralized School Choice," 2015.

_ , _ , and _ , "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.

Geyer, Judy and Holger Sieg, "Estimating a Model of Excess Demand for Public Housing," *Quantitative Economics*, 2013, *4* (3), 483–513.

Glaeser, Edward L. and Erzo F.P. Luttmer, "The Misallocation of Housing Under Rent Control," *American Economic Review*, 2003, *93* (4), 1027–1046.

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.

Hastings, Justine S., Thomas J. Kane, and Douglas O. Staiger, "Heterogeneous Preferences and the Efficacy of Public School Choice," 2009.

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

He, Yinghua, "Gaming the Boston School Choice Mechanism in Beijing," 2017.

_ , Antonio Miralles, Marek Pycia, and Jianye Yan, "A Pseudo-market Approach to Allocation with Priorities," *American Economic Journal: Microeconomics*.

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.

Hylland, Aanund and Richard Zeckhauser, "The Efficient Allocation of Individuals to Positions," 1979.

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.

Jappelli, Tullio and Luigi Pistaferri, "Using Subjective Income Expectations to Test for Excess Sensitivity of Consumption to Predicted Income Growth," *European Economic Review*, 2000, *44* (2), 337–358.

_ **and** _ , "Fiscal Policy and MPC Heterogeneity," *American Economic Journal: Macroeconomics*, 2014, *6* (4), 107–36.

Johnson, David S, Jonathan A Parker, and Nicholas S Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 2006, *96* (5), 1589–1610.

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

Kimball, Miles S, "Precautionary Saving in the Small and in the Large," *Econometrica: Journal of the Econometric Society*, 1990, pp. 53–73.

Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz, "Experimental Analysis of Neighborhood Effects," *Econometrica*, 2007, *75* (1), 83–119.

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 D, "Dynamic Matching in Overloaded Waiting Lists," 2017, pp. 1–78.

Lieber, Ethan M.J. and Lee M. Lockwood, "Targeting with In-Kind Transfers: Evidence from Medicaid Home Care," *SSRN Electronic Journal*, 2017, (January).

Ludwig, Jens, Greg J. Duncan, Lisa A. Gennetian, Lawrence F. Katz, Ronald C. Kessler, Jeffrey R. Kling, and Lisa Sanbonmatsu, "Long-Term Neighborhood Effects on Low-Income Families: Evidence from Moving to Opportunity," *American Economic Review, Papers & Proceedings*, 2013, *103* (3), 226–231.

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

Manski, Charles F, "Measuring Expectations," *Econometrica*, 2004, *72* (5), 1329–1376.

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.

McDonald, David R., "On the poisson approximation to the multinomial distribution," *Canadian Journal of Statistics*, 1980, *8* (1), 115–118.

McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," in "Frontiers in Econometrics" 1973, pp. 105–142.

_ , "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration," *Econometrica*, 1989, *57* (5), 995.

Meyer, Bruce and Nikolas Mittag, "Using Linked Survey and Administrative Data to Better Measure Income: Implications for Poverty, Program Effectiveness and Holes in the Safety Net," *Upjohn Institute Working Paper 15-242*, 2015.

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

Narita, Yusuke, "Match or Mismatch: Leaning and Inertia in School Choice," 2016, pp. 1–85.

Nichols, Albert L. and Richard J. Zeckhauser, "Targeting Transfers through Restrictions on Recipients," *The American Economic Review Papers and Proceedings*, 1982, *72* (2), 372–377.

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 David Pollard**, "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 1989, *57* (5), 1027–1057.

— **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.

Parker, Jonathan A, Nicholas S Souleles, David S Johnson, and Robert McClelland, "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 2013, *103* (6), 2530–53.

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.

Romich, Jennifer L and Thomas Weisner, "How Families View and Use the EITC: Advance Payment Versus Lump Sum Delivery," *National Tax Journal*, 2000, pp. 1245–1265.

Rust, J, "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 1987, *55* (5), 999–1033.

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

_ and _ , "Waiting for Affordable Housing in New York," 2016.

Skinner, Jonathan, "Risky Income, Life Cycle Consumption, and Precautionary Savings," *Journal of Monetary Economics*, 1988, *22* (2), 237–255.

Smeeding, Timothy M, Katherin Ross Phillips, and Michael O'Connor, "The EITC: Expectation, Knowledge, Use, and Economic and Social Mobility," *National Tax Journal*, 2000, pp. 1187–1209.

Stack, Carol, *All Our Kin: Strategies for Survival in a Black Community.*, New York: Basic Books, 1974.

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," *Journal of the American Society of Nephrology*, 2004, *15* (6), 1656–1663.

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, Neil, "The Public-Housing Allocation Problem : Theory and Evidence from Pittsburgh," 2016, (August), 1–55.

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.

Van Ommeren, Jos N. and Arno J. Van der Vlist, "Households' Willingness to Pay for Public Housing," *Journal of Urban Economics*, 2016, *92*, 91–105.

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.

Wiswall, Matthew and Basit Zafar, "Determinants of College Major Choice: Identification Using an Information Experiment," *The Review of Economic Studies*, 2014, *82* (2), 791–824.

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.

Zafar, Basit, "How Do College Students Form Expectations?," *Journal of Labor Economics*, Apr 2011, *29* (2), 301–348.

Zeldes, Stephen P, "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *The Quarterly Journal of Economics*, 1989, *104* (2), 275–298.

Zenios, S A, *Models for Kidney Allocation*, operations ed., Vol. 70, International Series in Operations Research & Management Science, 2004.

_ **, G M Chertow, and L M Wein**, "Dynamic allocation of kidneys to candidates on the transplant waiting list," *Operations Research*, 2000, *48* (4), 549–569.

Zhang, J, "The sound of silence: Observational learning in the US kidney market," *Marketing Science*, 2010, *29* (2), 315–335.

Zwick, Eric, "The Costs of Corporate Tax Complexity," 2018.