

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

Daniel Waldinger, MIT[†]

Job Market Paper

January 2018

[Click here for the most recent version](#)

Abstract

In-kind transfer programs aim to provide valuable resources to beneficiaries while targeting those who most need assistance. This problem is particularly challenging for public housing authorities (PHAs), which allocate apartments to applicants who may differ in their outside options as well as their preferred apartment types. PHAs in the U.S. differ widely in the priority systems they use and how much choice they give potential tenants over where to live. This paper evaluates how these choice and priority systems affect two competing objectives: efficiency and redistribution. I use data on the submitted choices of public housing applicants to estimate a structural model of preferences for public housing in Cambridge, MA. I find substantial heterogeneity in applicants' preferred housing developments and in their values of obtaining assistance. Counterfactual simulations suggest that the range of mechanisms used by PHAs involves a large trade-off between efficiency and redistribution. When applicants are allowed to choose where they live, tenants enjoy welfare gains equivalent to cash transfers of more than \$6,500 per year. Removing choice would house applicants with worse outside options but provide low match quality, causing cost-adjusted welfare gains to fall by 30 percent. Prioritizing low-income applicants while allowing choice would improve targeting without lowering match quality. While several combinations of choice and priority are on the frontier of efficiency and redistribution, some commonly used mechanisms, such as prioritizing higher-income applicants without allowing choice, are never optimal.

*I am very grateful to Nikhil Agarwal, Parag Pathak, Michael Whinston, and Amy Finkelstein for invaluable guidance and support. This paper benefited from feedback from Zarek Brot-Goldberg, Alonso Bucarey, Sydnee Caldwell, Glenn Ellison, Chishio Furukawa, Andrey Fradkin, Colin Gray, Arda Gitmez, Ingrid Gould Ellen, Jonathan Gruber, Ray Kluender, Jacob Leshno, Benjamin Marx, Pooya Molavi, Anders Munk-Nielsen, Scott Nelson, Alan Olivi, Christina Patterson, Otis Reid, Nancy Rose, Peng Shi, Michael Stepner, Neil Thakral, Heidi Williams, and Ariel Zucker, as well as seminar participants at the MIT Industrial Organization and Public Finance lunches. The Cambridge Housing Authority generously provided the applicant and tenant data used in this paper, with special thanks to Tara Aubuchon, Tito Evora, Michael Johnston, Jay Leslie, Hannah Lodi, and John Ziniewicz. All analysis and views expressed in this paper are my own and do not represent the views of the Cambridge Housing Authority.

[†]MIT Economics. Email: dwalding@mit.edu. All mistakes are my own.

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, 2017; 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 appli-

cants 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 as-

sumptions 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 Akerberg (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 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, 2016; 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, 2017; 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

¹The 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.

implementation of the method of simulated moments uses a change of variables and importance sampling technique proposed by Akerberg (2009) to reduce the computational burden in estimation. This paper also joins a growing literature on revealed preference analysis in centralized matching markets (Abdulkadirolu et al., 2017; Agarwal, 2015; Fack et al., 2015; 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.

2 Institutional Background and Data

Section 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 2.2 describes the Cambridge Housing Authority and the mechanism it used to allocate public housing during the period of study. Section 2.3 describes the applicant dataset and sample criteria.

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 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

²Based 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_FMD.PDF. In 2013, the public housing program served a population with similar incomes.

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

applicants on voucher waiting lists.⁴

2.1.1 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.

The allocation policies of surveyed PHAs share several common features. Applicants 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

⁴Public 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>

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.

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

2.2.1 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 B.1 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 appointment, produce required documentation, or respond to mail from the CHA.

⁵Every 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>

⁶The 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.

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 A provides details of the CHA and ACS datasets, and Section 5.1 explains how they are used to estimate the distribution of potential applicants.

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.

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. There are 3 developments in East Cambridge, 3 in North Cambridge, and 7 near Central Square. Table 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 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

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

⁷The “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).

⁸There 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.

⁹The 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.

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 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 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.

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

¹⁰The 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.

¹¹Note 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.

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 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.

4 Model of Preferences and Development Choice

Section 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 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.

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.

4.1.1 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 $C \subset \{1, \dots, J\}$ 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 $\tilde{\alpha} \leq \alpha$.

4.1.2 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 $F_{j,C}(t | p)$ 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 B.1.1, 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 expected continued waiting time for each development $j \in C$.

4.1.3 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.

¹²Descriptive 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.

The present discounted value to i of being assigned to development j in t years is

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

4.1.4 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

$$\begin{aligned} & \max_{j \in C} \frac{1}{\rho} E [e^{-\rho T_j} | p] (v_{ij} - v_{i0}) \\ &= \max_{j \in C} \int \frac{1}{\rho} e^{-\rho T_j} (v_{ij} - v_{i0}) dF_{j,C}(T_j | p) . \end{aligned}$$

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

$$\begin{aligned} & \max_{C \in \{0,1,\dots,J\}^3} E \left[e^{-\rho S} \max_{j \in C} \frac{1}{\rho} E [e^{-\rho T_j} | P] (v_{ij} - v_{i0}) \right] \\ &= \max_{C \in \{0,1,\dots,J\}^3} \int e^{-\rho S} \max_{j \in C} \left[\int \frac{1}{\rho} e^{-\rho T_j} (v_{ij} - v_{i0}) dF_{j,C}(T_j | P) \right] dG_C(S, P) . \end{aligned}$$

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.

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 $v_i = (v_{i1} - v_{i0}, \dots, v_{iJ} - v_{i0})$. 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.

4.2.1 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 $h_i = (h_{i1}, \dots, h_{iJ})$. 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

$$v_{i0} \equiv \max_{c, h} u_1(c) + u_2(h) \quad \text{s.t.} \quad c + h \leq y_i + \eta_i \quad (1)$$

$$= v_0(y_i + \eta_i). \quad (2)$$

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

$$v_{ij} \equiv u_1((1 - \tau)y_i) + u_2(h_{ij}). \quad (3)$$

The difference in flow payoffs is given by

$$v_{ij} - v_{i0} = \underbrace{u_1((1 - \tau)y_i)}_{\text{value of assistance}} - \overbrace{v_0(y_i + \eta_i)}^{\text{outside option}} + \underbrace{u_2(h_{ij})}_{\text{match value}}. \quad (4)$$

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

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 Akerberg (2009) reduces the computational burden. The final subsection shows how estimates from the utility model can be interpreted in terms of equivalent cash transfers.

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

¹³One 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.

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 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 $b = 1, \dots, B$ 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,

¹⁴The 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.

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

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 characteristics 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

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

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

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 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.

5.2.1 Structure of Simulation Inputs

Appendix B.1.1 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 $t \in \{1, \dots, T\}$. Each list $j \in \{1, \dots, J\}$ represents a development and bedroom size. There are S_j apartments represented by list j .
- **Apartment Vacancies:** each vacancy $\nu \in \{1, \dots, V\}$ 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 $i \in \{1, \dots, N\}$ 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 $C_i \subset \{1, \dots, J\}, |C_i| \leq 3$ 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 .

5.2.2 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 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.

5.2.3 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

¹⁶Due 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.

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, $F_{j,C}(T_j | p)$, 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 B.1 provides details of how these distributions were constructed.

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.

5.3.1 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

$$v_{ij} - v_{i0} = \delta_j + \underbrace{\phi_1 \log y_i - \phi_2 \log(y_i + \eta_i)}_{\text{value of assistance}} + \underbrace{g(Z_i) + \sum_k X_{ijk} \beta_k^o + \sum_m X_{jkm} \nu_{im} \beta_m^u}_{\text{match value}} + \epsilon_{ij}, \quad (5)$$

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

$$\eta_i \stackrel{iid}{\sim} TN(0, \sigma_\eta^2, -y_i, \infty) \quad \nu_{im} \stackrel{iid}{\sim} N(0, 1) \quad \epsilon_{ij} \stackrel{iid}{\sim} N(0, 1) \quad (6)$$

In addition to placing parametric structure on the unobservables, this parameterization adds development fixed effects and demographic shifters to equation 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 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 6.3 I examine robustness of parameters governing the value of assistance to adding random coefficients for development size and location.

5.3.2 Moments and Objective Function

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

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

I estimate θ based on moment conditions

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

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 $q \in \{1, \dots, Q\}$, for specific choice features $m_i^{(q)}$ and household characteristics $Z_i^{(q)}$:

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

In estimation, the conditional expectation $\hat{E}(m_i | Z_i, \theta)$ 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 $\hat{\mathbf{g}}(\theta) \equiv (\hat{g}^{(1)}(\theta), \dots, \hat{g}^{(Q)}(\theta))'$ 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:

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

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

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

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

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

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

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

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:

$$m_i^{(q)} = \frac{1}{|C_i|} \sum_{j \in C_i} \bar{T}_j, \quad \left(\frac{1}{|C_i|} \sum_{j \in C_i} \bar{T}_j \right)^2, \quad \frac{1}{|C_i|} \sum_{j \in C_i} (\bar{T}_j)^2; \\ Z_i^{(q)} = 1\{(y_i, Z_i) \in \mathcal{Y}^{(q)} \times \mathcal{Z}^{(q)}\}$$

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

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

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:

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

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.

5.3.3 Change of Variables and Importance Sampling

Estimating the conditional expectation $E[m_i | Z_i, \theta]$ presents a computational challenge because the two-stage development choice problem is computationally burdensome to solve. A standard simulation procedure would draw unobserved characteristics $\{(\eta_{is}, \nu_{is}, \epsilon_{is})\}_{s=1, \dots, S}^{i=1, \dots, N}$ 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

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

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 Akerberg (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 $v_i = \{v_{i0}, v_{i1}, \dots, v_{iJ}\}$ and discount factor ρ . The technique draws flow payoffs $\{v_i^s\}_{s=1, \dots, S}^{i=1, \dots, N}$ from an initial (proposal) distribution $g(\cdot | Z_i)$; 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 :

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

Because flow payoffs were drawn from $g(\cdot | Z_i)$, 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 $p(\cdot | Z_i, \theta)$ 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 B.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 $\hat{E}[m_i | Z_i, \theta]$. They do not yet account for estimation error in the distribution of potential applicants or their beliefs.

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

$$v_{ij} - v_{i0} = v_0(y_i + \eta_i + EV_{ij}) - v_0(y_i + \eta_i), \quad (7)$$

where $v_0(\cdot)$ is the indirect utility function defined in equation 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:

$$EV_{ij} = (y_i + \eta_i) (\exp^{v_{ij} - v_{i0}} - 1). \quad (8)$$

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

$$v_{iA} - v_{iB} = u_1((1 - \tau)y_i + EV_{i,AB}) - u_1((1 - \tau)y_i), \quad (9)$$

where u_1 is utility from the numeraire as defined in equation 3. Equation 9 differs from equation 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 $u_1(\cdot)$ rather than the indirect utility function $v_0(\cdot)$. With homothetic preferences, the closed form expression is

$$EV_{i,AB} = (1 - \tau)y_i \left(\exp^{\frac{v_{iA} - v_{iB}}{\gamma}} - 1 \right). \quad (10)$$

6 Estimation Results

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

6.1 Eligible Population

Appendix Table 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 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 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.

6.2 Applicant Beliefs

Selected parameters governing inputs to the Cambridge Mechanism simulation are shown in Table 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 3.3, applicants are less likely to choose a development with a higher list position.

Table 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.

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.

6.3.1 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 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 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.

6.3.2 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 and 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 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 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 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.

7 Evaluating Design Trade-Offs

Using the estimates from Section 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 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 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.

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.

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

7.1.1 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}_φ is a subset of developments from which the applicant may receive apartment offers.
2. A **priority system** $\psi_\varphi : \mathcal{Z} \rightarrow \{1, \dots, B\}$ 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 $\nu \in \{1, \dots, V\}$ has a date t_ν and development j_ν . Each applicant $i \in \{1, \dots, N\}$ has arrival date t_i , departure date r_i , observed characteristics Z_i , and payoff vector $v_i = (v_{i0}, v_{i1}, \dots, v_{ij})$. 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 $j \in C_i$.

7.1.2 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, \dots, T_J} may depend on the applicant's priority $\psi_\varphi(Z_i)$. The applicant solves the following choice problem:

$$\max_{C \in \mathcal{C}_\varphi} \sum_{j \in C} w_j^C(\psi_\varphi(Z_i))(v_{ij} - v_{i0}) \quad (11)$$

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

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 $p(v_i | Z_i, \hat{\theta}_{MSM})$, and given that applicants choose developments according to equation 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 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 C 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.

7.1.3 Evaluating Allocations

Given sequences of inputs, a mechanism φ produces an eventual assignment $j_\varphi(i) \in \{0, 1, \dots, J\}$ 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:

$$W(\varphi) = \frac{1}{\sum_{i=1}^N 1\{j_\varphi(i) \neq 0\}} \sum_{i=1}^N EV_{i,j_\varphi(i)} \quad (12)$$

where $EV_{i,j_\varphi(i)}$ is as defined in equation 8. To summarize characteristics of housed applicants, one can do the same for transformations of applicant characteristics:

$$\frac{1}{\sum_{i=1}^N 1\{j_\varphi(i) \neq 0\}} \sum_{i=1}^N h(Z_i, v_i, j_\varphi(i)) \quad (13)$$

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:

$$W(\varphi; f) = \frac{1}{\sum_{i=1}^N 1\{j_\varphi(i) \neq 0\}} \sum_{i=1}^N f(Z_i, v_i, EV_{i,j_\varphi(i)}) \quad (14)$$

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

$$\tilde{W}(\varphi; f) = \frac{\sum_{i=1}^N f(Z_i, v_i, EV_{i, j_\varphi(i)})}{\sum_{i=1}^N 1\{j_\varphi(i) \neq 0\}(c - 0.3y_i)} \quad (15)$$

7.1.4 Simulated Mechanisms

The mechanisms used by the 24 surveyed PHAs in Section 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

1. **Choose One:** $\mathcal{C} = \{\{1\}, \dots, \{J\}\}$. 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:** $\mathcal{C} = 2^{\{1, \dots, J\}}$. Applicants may choose any subset of developments, as in Boston and San Antonio.
3. **Choose All or One:** $\mathcal{C} = \{\{1\}, \dots, \{J\}, \{1, \dots, J\}\}$. 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:** $\mathcal{C} = \{C_{\text{north}}, C_{\text{east}}, C_{\text{central}}, \{1, \dots, J\}\}$. 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:** $\mathcal{C} = \{\{1, \dots, J\}\}$. 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:

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

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.

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 12 - 15.

7.2.1 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 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 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.

7.2.2 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 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 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 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.

7.2.3 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 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):

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

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 15 to determine which mechanism should be used given a PHA’s degree of inequality aversion.

Figure 9 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 15 for each mechanism, normalized by welfare under “Equal Priority, Choose One” at a range of inequality aversion parameters. Consistent with Table 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 C 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.

7.2.4 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 10 summarizes the allocation that the Cambridge Mechanism would produce if applicants had the waiting time beliefs estimated in Section 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 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 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.

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-

¹⁸The 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.

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.

References

- Abdulkadirolu, Atila, 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.
- 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.
- 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).
- 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.
- 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.
- 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.
- Bloch, Francis and David Cantala**, “Dynamic Assignment of Objects to Queuing Agents,” *American Economic Journal: Microeconomics*, 2017, *9* (1), 88–122.
- Chade, Hector and Lones Smith**, “Simultaneous Search,” *Econometrica*, 2006, *74* (5), 1293–1307.

- Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz**, “The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity 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.
- 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.
- Fack, Gabrielle, Julien Grenet, and Yinghua He**, “Beyond Truth-Telling: Preference Estimation with Centralized School Choice,” 2015.
- 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.
- Hastings, Justine S., Thomas J. Kane, and Douglas O. Staiger**, “Heterogeneous Preferences and the Efficacy of Public School Choice,” 2009.
- He, Yinghua**, “Gaming the Boston School Choice Mechanism in Beijing,” 2017.
- Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz**, “Experimental Analysis of Neighborhood Effects,” *Econometrica*, 2007, 75 (1), 83–119.
- 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.

- 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.
- 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.
- Pakes, Ariel and David Pollard**, “Simulation and the Asymptotics of Optimization Estimators,” *Econometrica*, 1989, 57 (5), 1027–1057.
- Sieg, Holger and Chamna Yoon**, “Waiting for Affordable Housing in New York,” 2016.
- Stack, Carol**, *All Our Kin: Strategies for Survival in a Black Community.*, New York: Basic Books, 1974.
- Thakral, Neil**, “The Public-Housing Allocation Problem : Theory and Evidence from Pittsburgh,” 2016, (August), 1–55.
- 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.

9 Tables and Figures

Figure 1: Locations of Cambridge Family Public Housing Developments




Figure 2: Application Rates by Income




Figure 3: 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 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.

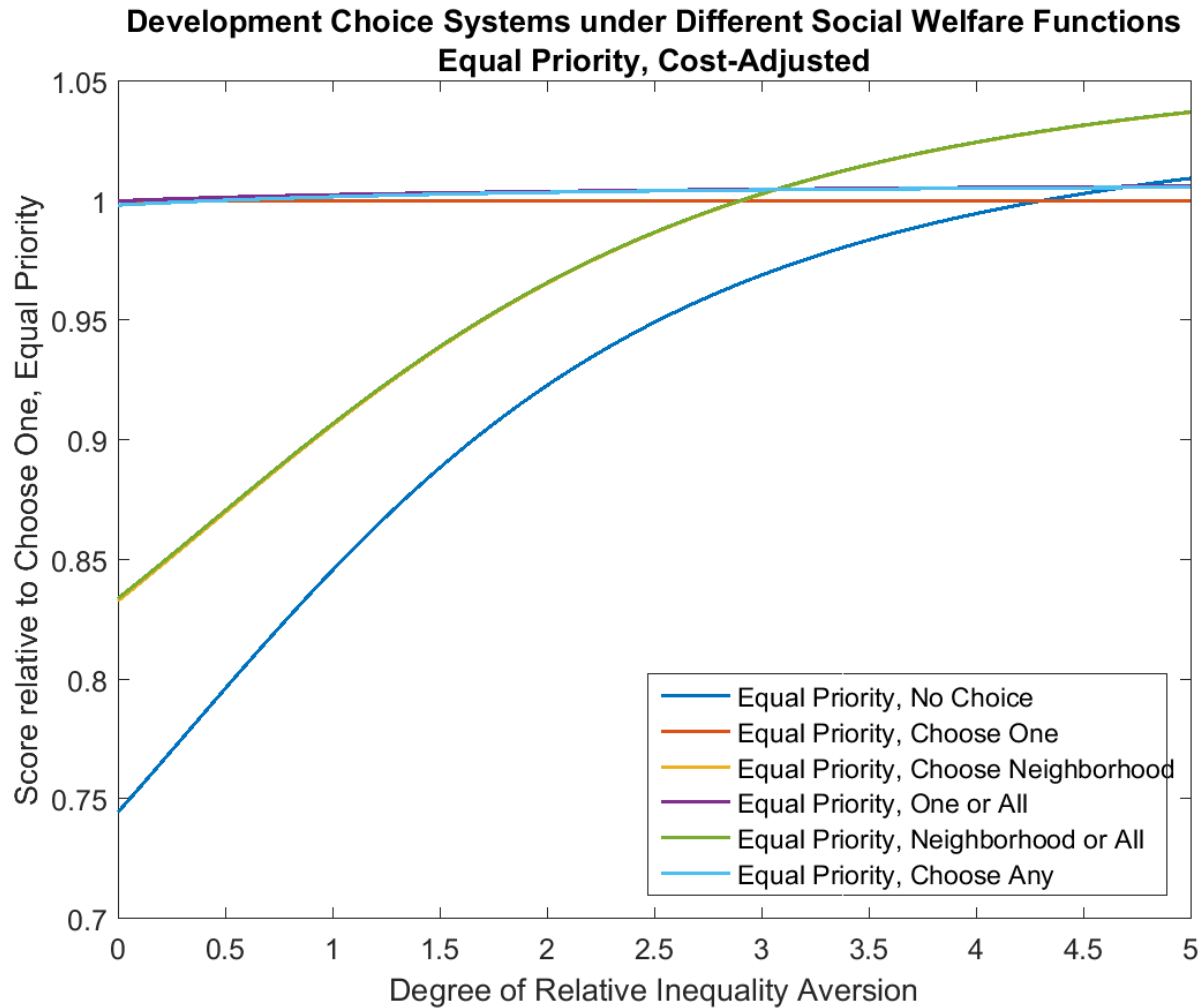


Figure 4: Comparison of cost-adjusted welfare gains produced by development choice systems used in practice, defined in Section 7.1. Applicants have Equal Priority in all mechanisms. Each point on the x-axis corresponds to a degree of relative inequality aversion. At each point, cost-adjusted welfare gains from each mechanism are normalized by the value for Equal Priority, Choose One.

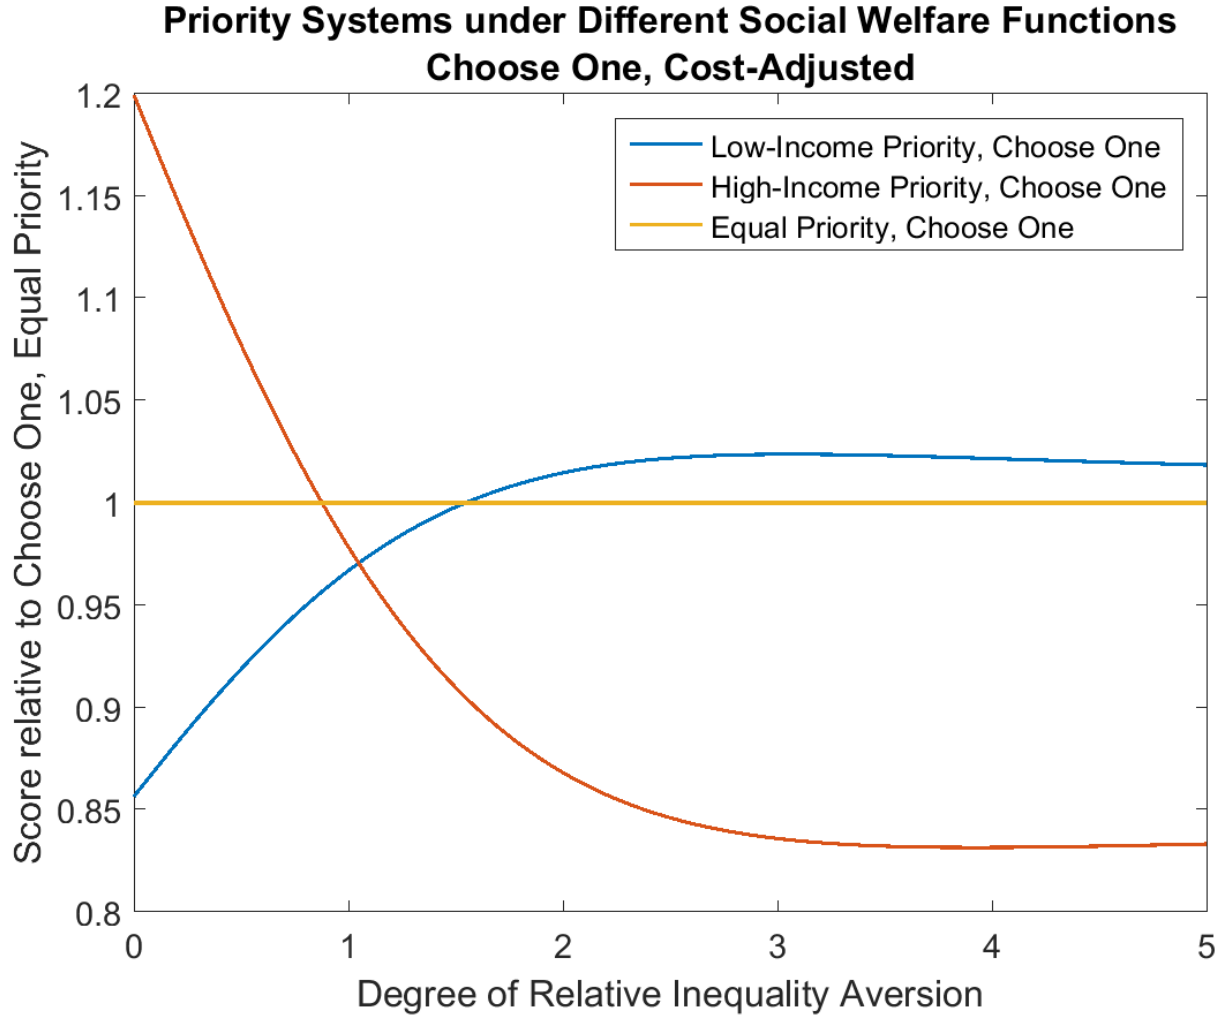


Figure 5: Comparison of cost-adjusted welfare gains produced by different priority systems used in practice. Low-Income Priority offers apartments to applicants below 30% AMI before other applicants, while High-Income Priority first offers apartments to applicants above 30% AMI. Applicants choose one 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.

Preferred Choice and Priority Systems under Different Social Welfare Functions
Cost-Adjusted

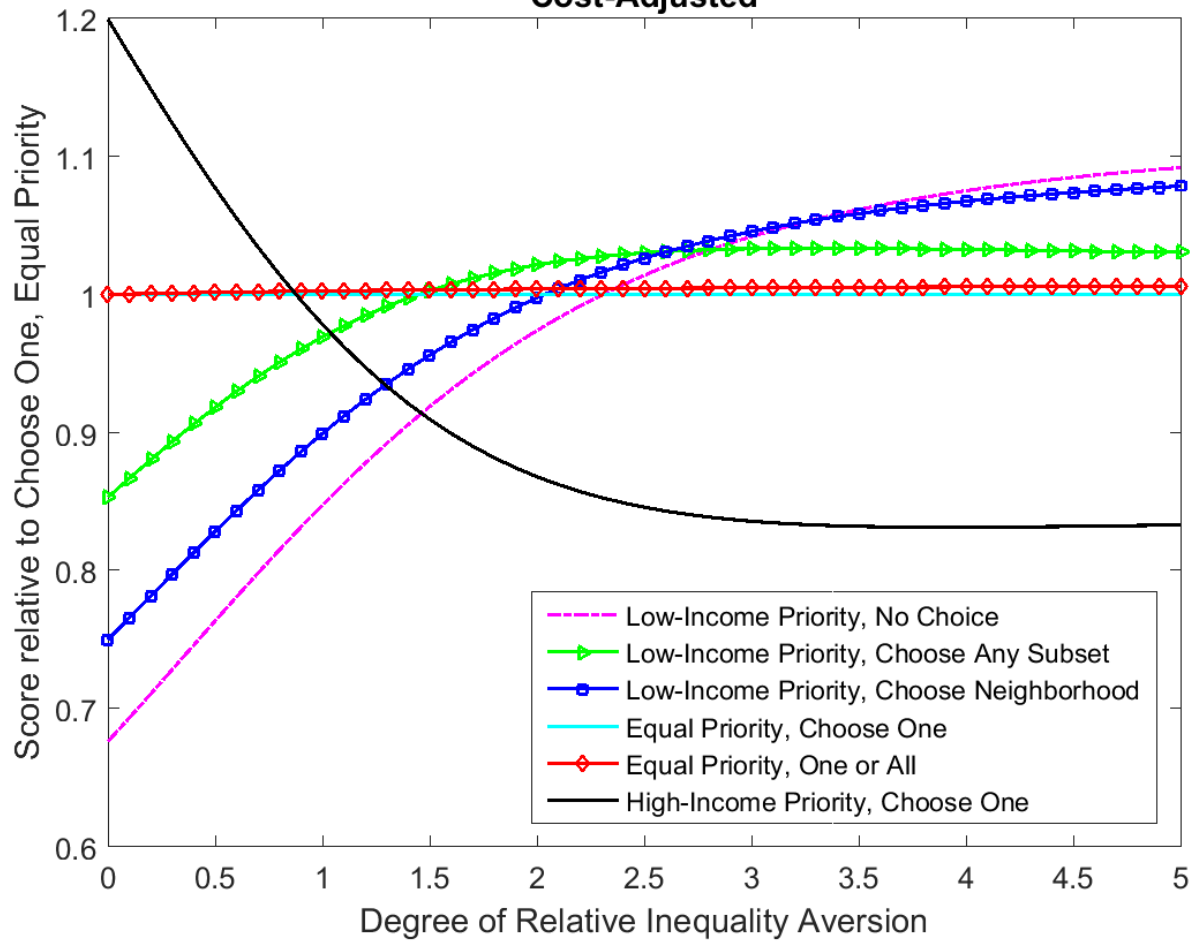


Figure 6: Cost-adjusted welfare gains from choice and priority systems that perform well for different degrees of relative inequality aversion. Each point on the x-axis corresponds to a degree of relative inequality aversion. At each point, cost-adjusted welfare gains from each mechanism are normalized by the value for Equal Priority, Choose One.

Table 1: 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
<i>Panel A: PHA's with Largest Public Housing Stock</i>				
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
<i>Panel B: PHA's comparable to Cambridge, MA</i> (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 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 3: Applicants

Characteristics of Eligible Population and Applicants		by Year of Initial Application					
		All	2010	2011	2012	2013	2014
		Eligible	Applied				
# Applicants		6828	1726	415	407	371	347
Income (\$)		42,219	18,477	17,971	18,718	18,191	19,835
2 Bedrooms		76.5%	69.8%	68.9%	69.8%	68.2%	72.6%
3 Bedrooms		23.4%	29.8%	30.8%	30.0%	31.8%	27.1%
Lives in Cambridge		49.3%	57.4%	55.2%	62.4%	52.6%	57.1%
Works in Cambridge		55.2%	39.7%	36.6%	39.8%	44.7%	44.1%
Age Youngest Member		10.5	8.5	8.2	8.2	8.7	9.0
Age Oldest Member		40.0	36.7	35.7	36.6	37.7	37.7
# Children		1.25	1.27	1.39	1.27	1.24	1.16
Child Under 10		60.8%	60.8%	56.6%	62.9%	62.0%	64.8%
Household Head Head White		55.2%	36.2%	32.3%	38.8%	38.8%	34.3%
Household Head Head Black		19.6%	50.3%	54.7%	47.7%	46.6%	49.3%
Household Head Head Hispanic		17.9%	19.2%	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 4: Initial Development Choices

Sub-Group	Number of Applicants	Initial Choices of Applicants					Location		
		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 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 (0.0031)		-0.0191 (0.0036)		-0.0259 (0.0063)	
Expected Waiting Time (Years)		-0.0639 (0.279)		-4.051 (0.755)		-4.992 (1.319)
Development FE's			X	X		
Development - Choice Set FE's					X	X
Implied Own-Price Elasticity	-0.657 (0.145)	-0.029 (0.128)	-0.747 (0.175)	-3.511 (0.669)	-1.125 (7.677)	-4.087 (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 6: Inputs to Waiting Time Simulation

Parameter	Value
<i>Apartment Vacancies</i>	
Annual Vacancy Rate per Unit	0.10
Annual Vacancy Rate Total	108
<i>Applicant Arrivals and Departures</i>	
Daily Applicant Arrival Rate	0.945
Annual Applicant Arrival Rate	345
Instant Departure Probability	0.243
Annual Departure Rate	0.245
<i>Final Choice Model</i>	
List Position Coefficient	-0.019
<i>Fixed Effects</i>	
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 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)
<i>Panel A: Value of Assistance</i>		
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)
<i>Panel B: Match Values</i>		
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 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 (\$/year)	8.6%	11.8%	8.2%	11.2%	8.0%	10.8%
	1,435	3,026	1,452	3,147	519	854
1st Choice instead of 3rd						
% Income (\$/year)	16.5%	19.5%	16.1%	18.7%	15.4%	18.0%
	2,957	4,941	3,062	5,195	1,023	1,435
1st Choice instead of Last						
% Income (\$/year)	100.6%	106.1%	99.1%	104.2%	97.3%	102.0%
	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 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
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 10: Effect of Development Choice

	Common Development Choice Systems						Full Information		
	Choose One (1)	Choose Any Subset (2)	Choose All or One (3)	Choose Neighborhood (4)	Choose All or Neighborhood (5)	No Choice (6)	Cambridge (7)	Equivalent Variation Maximizing (8)	Targeting Maximizing (9)
<i>Panel A: Welfare Gain and Cost of Allocation</i>									
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 Govt.	0.84	0.84	0.84	0.67	0.67	0.59	0.93	1.26	0.45
<i>Panel B: Targeting</i>									
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%
<i>Panel C: Match Quality</i>									
% 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%
<i>Panel D: Characteristics of Housed Applicants</i>									
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 11: Effect of Priority System

	Common Choice and Priority Systems						Cambridge		Full Information	
	Low-Income Priority		High-Income Priority		Equal Priority		Cambridge		Full Information	
	Choose One	No Choice	Choose One	No Choice	Choose One	No Choice	Equivalent Variation Maximizing	Targeting Maximizing		
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)		
<i>Panel A: Welfare Gain and Cost of Allocation</i>										
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 Govt.	0.68	0.50	1.11	0.65	0.84	0.59	0.93	1.26	0.45	
<i>Panel B: Targeting</i>										
Observed Income (\$)	13,581	11,855	27,950	20,580	20,509	17,555	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%	
<i>Panel C: Match Quality</i>										
% 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%	
<i>Panel D: Characteristics of Housed Applicants</i>										
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.

A Datasets

A.1 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 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

¹⁹To reduce waiting time uncertainty, CHA merged four small waiting lists with larger lists in 2013. As a result, an applicant’s initial choice set depended on its application date.

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.

A.2 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 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.

²⁰There 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.

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

B Estimation Details

B.1 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.

B.1.1 Cambridge Mechanism

Between 2010 and 2014, Cambridge ran its public housing waiting lists according to the following algorithm. Calendar time is indexed $t = 1, \dots, T$. Waiting lists are indexed by $j = 1, \dots, J$, where a list corresponds to a specific bedroom size apartment (2 or 3 bedrooms) in a specific development. Applicants are indexed $i = 1, \dots, N$, vacancies by $\nu = 1, \dots, V$. 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.

B.1.2 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.

B.1.3 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, \dots, 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 $s = 1, \dots, 1000$. These states – $\{(s_C^s, p_C^s)\}_{s=1, \dots, 1000}$ – form an empirical measure \hat{G}_C .

Constructing beliefs $\{F_{j,C}(\cdot | p_C)\}_{j,C,p_C}$ 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

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

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 $\{t_d\}_{d=1,\dots,D}$. 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 $\{p_{C,d}, t_{C,d}\}_{d=1,\dots,D}$, π and the upper bound of the support of the beta distribution for each $j \in C$ are estimated by maximum likelihood.

B.2 Development Preferences

B.2.1 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

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

where

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

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

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

B.2.2 Moments

To estimate the parameter vector $\theta = \{\rho, \delta, \beta, g(\cdot), \sigma_\eta\}$, 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.

B.2.3 Importance Sampling and Change of Variables

I estimate the parameter vector θ based on moment conditions

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

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 $\hat{E}(m_i | z_i, \theta)$ 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

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

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 Akerberg (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 $g(v, \eta | z_i)$ and compute the optimal choice for each v_{is} once for each value of ρ . Then, to estimate $E(m_i | z_i, \theta)$, I re-weight the simulation draws at new parameter vectors $\theta_{-\rho}$:

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

Since the flow payoffs and unknown income are drawn according to $g(\cdot | z_i)$, the above formula provides an unbiased estimate of $E(m_i | z_i, \theta)$. 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 .

B.2.4 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 $g(\cdot | z_i)$
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 $F_{j, C_{is}}(T_j | p_{is})$ 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.

B.2.5 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 $\beta \in \{1, 0.98, 0.96, \dots, 0.5\}$. To limit numerical instability in specifications with several random coefficients, the variance of each random coefficient was constrained to be less than one million.

B.2.6 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 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 $\hat{E}[m_i | z_i, \theta]$:

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

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

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

The variance estimate is

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

C 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 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 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 $\{w_j^C(\psi_\varphi(y_i))\}_{C \in \mathcal{C}_\varphi}^{j=1, \dots, J}$. 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 $w^{(q)'}$ with distance $D^{(q)} = \|w^{(q)' } - w^{(q)}\|$
3. Weights are updated as a convex combination of the proposed and implied weights:

$$w^{(q+1)} = \lambda^{(q)} w^{(q)'} + (1 - \lambda^{(q)}) w^{(q)}.$$

The factor λ determines how aggressively the weights are updated. If $\lambda = 1$, then the weights implied by applicant choices ($w^{(q)'}$) 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 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 13: Simulated Waiting Times from Initial Application

Simulated Waiting Time Realizations				
Development	Simulation		Data	
	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.

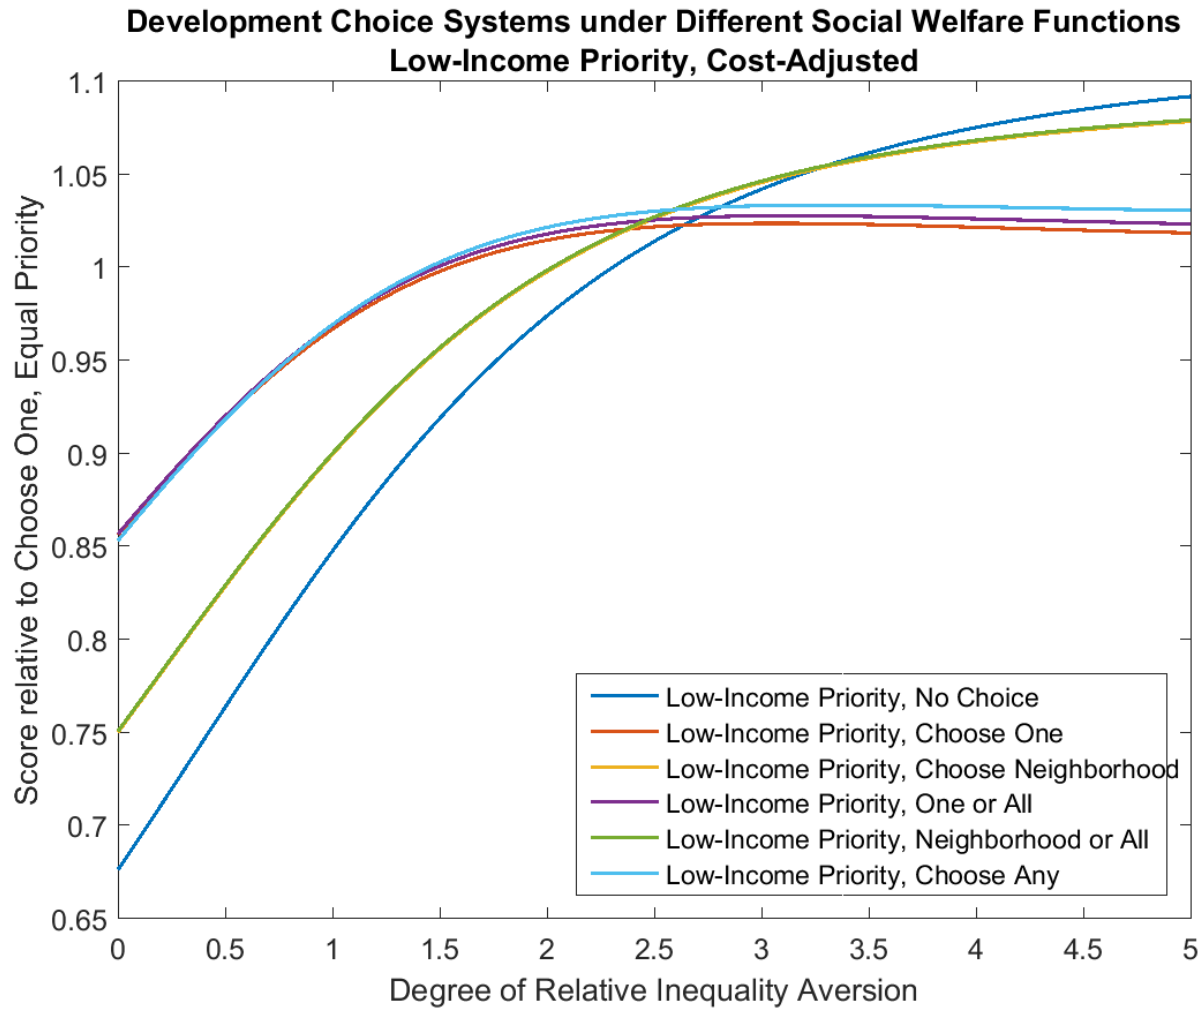


Figure 7: Comparison of cost-adjusted welfare gains produced by development choice systems used in practice, with priority for households with income below 30% AMI. Welfare gains are normalized by the value for Equal Priority, Choose One.

**Priority and Development Choice Systems under Different Social Welfare Functions
Not Cost-Adjusted**

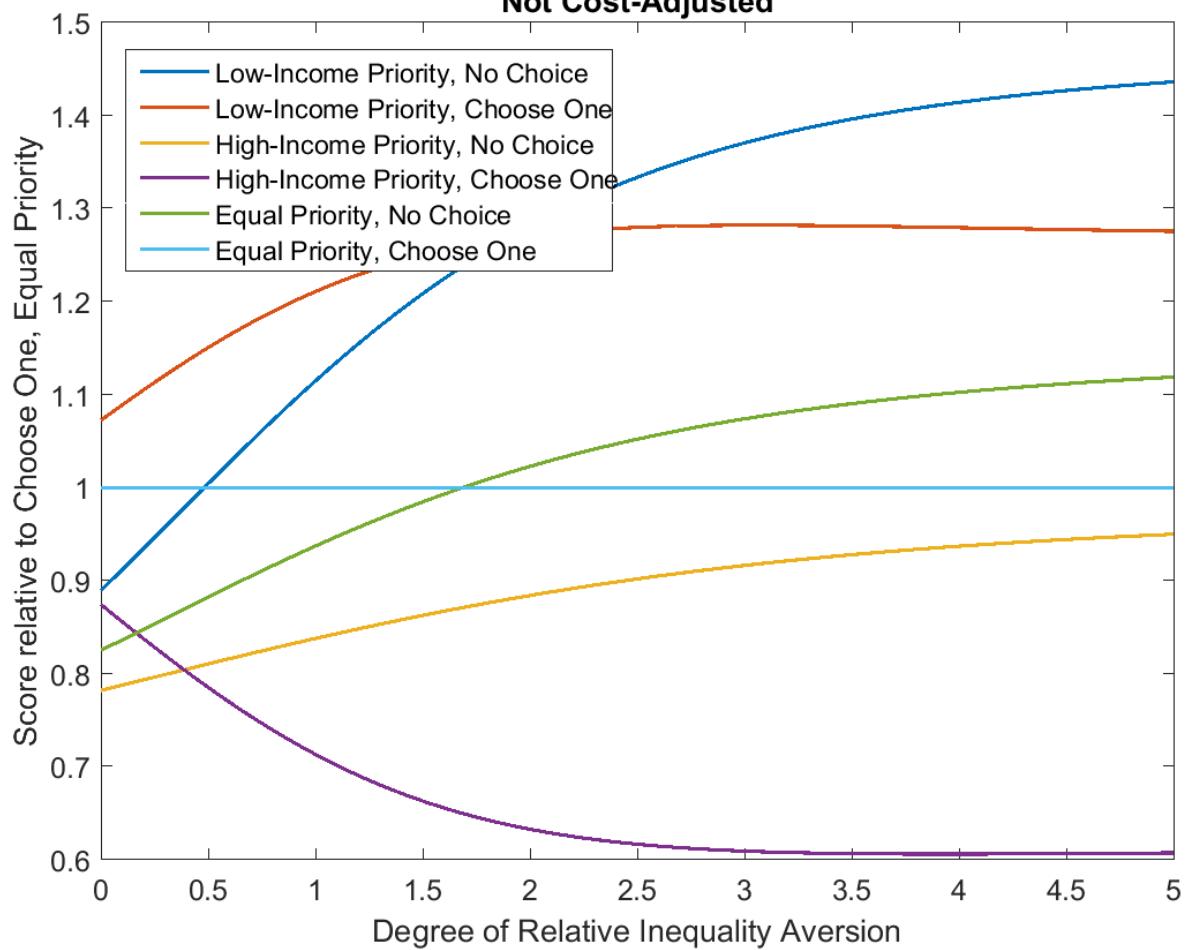


Figure 8: Comparison of welfare gains produced by development choice and priority systems used in practice. Welfare gains are not adjusted for cost, and are normalized by the gains from Equal Priority, Choose One.