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

Daniel Waldinger

January 22, 2021

*NYU Department of Economics. Email: danielwaldinger@nyu.edu. I am grateful to Nikhil Agarwal, Parag Pathak, Michael Whinston, and Amy Finkelstein for invaluable guidance and support. I also thank Nick Arnosti, Glenn Ellison, Ingrid Gould-Ellen, Tatiana Homonoff, John Lazarev, Jacob Leshno, Peng Shi, Paulo Somaini, Neil Thakral, and Heidi Williams; seminar participants at the MIT Industrial Organization and Public Finance workshops, and at several other universities; and my colleagues at NYU. Suggestions from the editor and four anonymous referees greatly improved the paper. 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. I acknowledge support from a National Science Foundation Graduate Research Fellowship. All mistakes are my own.

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 households who submitted an application after October 1st 2009, as well as households who applied earlier and were still on the waiting list on that date.

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, race, and 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.[1] 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 the CHA sent the letter. Further details on these and other data processing steps are provided in the replication archive.

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 would likely have been eligible to appear in the applicant sample, which contains households with priority for 2 and 3 bedroom apartments in Cambridge Family Public Housing. Beginning with the universe of

[1] To 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 their application date.

ACS households living in the state of Massachusetts, I first determine whether each household lived or worked in Cambridge.[2] 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 create separate households for the elderly members and the younger members.[3] 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.

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

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

B Robustness Checks for Final Choice Analysis

For the evidence of responsiveness to waiting time information in Section IIC to be valid, position information provided to applicants when they make their final choice should be uncorrelated with their preferences, conditional on first-stage decisions. While difficult to test directly, the data rule out two possibilities that would suggest this condition is violated. First, conditional on an applicant's initial choice, the position information they receive at final choice is not significantly correlated with demographic or economic characteristics. Any selection into the final choice stage based on preferences would have to be a function of unobservables. Second, in contrast to the final choice analysis, initial choices are not predicted by list lengths on the specific date a household applied. This suggests that applicants were not aware of short- or medium-term fluctuations in list lengths before they received their final choice letters.

B.1 Testing for Selection in Final Choice Analysis

This section tests whether final choice positions are correlated with observed applicant characteristics. The idea is analogous to a test for balance between treatment and control group characteristics in a randomized controlled trial. Here, the analysis faces two complications. First, because each applicant selected their final choice set in the first stage, the test must condition on initial choices. Second, the "treatment" is multi-dimensional because each applicant learns up to three list positions at the final choice stage.

I therefore test whether each pair of list positions in the applicant's final choice letter predicts the applicant's characteristics. Let C_i denote applicant i's initial choice, p a pair of developments $j, k \in C_i$, and Z_i an applicant characteristic. I run an ordinary least squares regression with one observation for each applicant and pair of developments in a final choice letter:

$$Z_i = \alpha_{p,C_i} + \beta_{p,C_i} \frac{x_{ij}}{x_{ik}} + \epsilon_{ip} \tag{17}$$

The predictors include choice set, pair interaction dummies and the ratio of applicant i's positions on lists j and k. This ratio can have a different relationship with the dependent variable for each choice set, pair interaction. One can interpret the ratio as a relative price; a higher value means that the applicant faces a longer continued wait for development j relative to development k. To test whether certain types of applicants systematically receive different position information at final choice, captured by these relative prices, I perform a joint F-test of the hypothesis $\beta_{p,C_i} = 0 \ \forall p, C_i$.

Table 9 shows that final choice list positions are not significantly predictive of most applicant characteristics. Panel A constructs the relative price using list position, while Panel B uses expected continued waiting time. The first two characteristics – application date and final choice letter date – are strongly predicted by final choice positions. These relationships are to be expected, and they demonstrate that the regressions have sufficient power to reject the null hypothesis. The correlation between application date and position information is consistent with the fact that some lists are becoming longer relative to others over the sample period. The correlation between final choice date and position information would occur for purely mechanical reasons, even without trends in relative list lengths over time. If an applicant receives a final choice letter early compared to others who made the same initial choice, one of the lists must be unusually short.

Among the other characteristics, only number of children and number of household members, which are highly (and mechanically) correlated, have F-test p-values below .05. Importantly, characteristics that are important predictors of applicant behavior in the structural model, including household race/ethnicity and annual income, are not correlated with position information conditional on initial choices. If applicants were selecting into the final choice stage based on their development preferences, one would expect selection to be correlated with these characteristics. The absence of a correlation supports interpreting the final choice regressions in Section IIC as revealing a causal response of final choices to waiting time information.

B.2 Testing for Responsiveness of Initial Choices to List Position

To test for responsiveness of initial choices to list position, I construct a dataset similar to the one used for the final choice analysis in Section IIC. Specifically, for each applicant, I determine the position they would be on the waiting list for each development if they included that development in their initial choice. I then test whether each development is more likely to be selected on dates when that development's waiting list is short relative to those of other developments. Conducting this analysis separately for each development deals with the fact that an applicant may select multiple developments in their initial choice.

Define $y_{ij} = 1\{j \in C_i\}$ to be an indicator for whether applicant i selected development j as part of their initial choice, and let x_{ij} be the position number the applicant would have had on list j if they selected it (regardless of whether they did). Each applicant has characteristics $\mathbf{Z}_i$. I estimate the following regression equation for each development j:

$$y_{ij} = \alpha_{ij} + \mathbf{Z}_i' \boldsymbol{\delta} + \sum_{k \neq j} \beta_{jk} \frac{x_{ik}}{x_{ij}} + \epsilon_{ij} \tag{18}$$

Equation 18 allows the probability development j is selected to depend on the ratio between the length of list j and list k for every other development k. This specification captures the idea that *relative* list lengths should matter for applicants' decisions, and also allows applicant characteristics to predict their choices.

Table 10 presents F-statistics and p-values from a test of the joint hypothesis $\beta_{jk} = 0 \ \forall k$ for each development j. In Panel A, which controls for list lengths but not applicant characteristics, three developments have p-values below 0.05. Only one development has a p-value below 0.05 once applicant characteristics are included in Panel B. There is therefore little evidence of responsiveness to list position at the initial choice stage. These patterns contrast starkly with the clear response to position information at final choice. They are also consistent with institutional facts about the Cambridge Mechanism. The CHA did not make list length information readily available to new applicants, and although an applicant could call the CHA and ask for its position number on each list after it applied, few did so.

In addition to validating the final choice analysis by ruling out a particular source of selection into the final choice stage, Table 10 also motivates the information structure in the development choice model. Applicants do not behave as though they know the length of each list when they apply; instead, their initial choices are consistent with a common prior based on a steady state distribution of waiting times, while their final choices show updating based on the specific position information in their final choice letters.

C Details of Estimation Procedure and Counterfactual Simulations

C.1 Distribution of Eligible Households

The first decision an eligible household makes is whether to apply for public housing at all. Estimating application rates requires the distribution of characteristics among all households that could have applied for Cambridge public housing during the sample period. This includes households that did apply and also *eligible non-applicants* – eligible households that did not apply and were not already Cambridge public housing applicants or tenants at the beginning of 2010. The CHA dataset contains information on households who applied during the sample period, but not on households which 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 waitlist 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.

The ACS publishes a 5 percent sample of U.S. households covering 2010 through 2014, the same period covered by the CHA applicant dataset. It contains information on household structure and economic and demographic characteristics that determine eligibility and priority for Cambridge public housing – in particular, 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 the probability that each eligible household surveyed in the ACS appears in the CHA dataset, either as a tenant or as a past or current applicant, as a parametric function of household characteristics. The parameters are estimated by minimum distance using a probit link function and moments based on the characteristics of the households in the CHA dataset. 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 waitlist 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.

To formally describe the minimum distance estimator, ACS households are indexed by $b = 1, ..., B$ and have observed household characteristics $\mathbf{Z}_b$. The ACS assigns each surveyed household a weight w_b based on household b's inverse probability of being sampled. In other words, w_b is the expected number of households that b represents. The estimator chooses a parameter vector $\boldsymbol{\theta}_{ACS}$. Denote statistics from the Cambridge dataset by $\boldsymbol{m}_{data}$, and denote the contribution of each ACS household to the same statistics by $\boldsymbol{m}_b$. The minimum distance estimator solves

$$\min_{\boldsymbol{\theta}_{ACS}} \; (\boldsymbol{m}_{ACS}(\boldsymbol{\theta}_{ACS}) - \boldsymbol{m}_{data})'(\boldsymbol{m}_{ACS}(\boldsymbol{\theta}_{ACS}) - \boldsymbol{m}_{data})$$

where

$$\boldsymbol{m}_{ACS}(\boldsymbol{\theta}_{ACS}) \equiv \sum_{b=1}^{B} p(\mathbf{Z}_b, \boldsymbol{\theta}_{ACS}) w_b \boldsymbol{m}_b$$

$$p(\mathbf{Z}, \boldsymbol{\theta}) = \Phi(\mathbf{Z}'\boldsymbol{\theta})$$

C.2 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.

C.2.1 Cambridge Mechanism

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

The simulation of the Cambridge Mechanism 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. Applicant i is told their positions $\{p_{ij}\}_{j \in C_i}$ on each list in their initial choice set.

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

(c) If i responds, they make a final choice $f \in C_i$ which may depend on their list positions.

(d) If i does not respond, they are 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, $(L_{j,k}, K_{j,k})$ is held for the next period.

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

C.2.2 Structure of Simulation Inputs

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 waitlist 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 waitlist.

- **Apartment Vacancies:** 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 approximately once every 10 years. The vacancy rate of list j is set to $0.1 * S_j$.
- **Applicant Arrivals and Exogenous Departures:** applicants arrive according to a poisson process with rate α. Each applicant becomes unresponsive immediately with probability a_0, and departs at an exponential annual rate a_1 thereafter. I estimate these parameters by non-linear least squares.
- **Initial Choices:** applicant i makes an initial choice $C_i \subset \{1, ..., J\}, |C_i| \leq 3$ upon arrival. Since applicants do not know the state of the waitlist when they apply, their initial choices are independent of the current state. Each C_i is therefore drawn independently from the empirical distribution in the CHA dataset.
- **Final Choice Letters:** 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 $(L_{j,l}, K_{j,l})$ is drawn independently from the empirical distribution for list j in the Cambridge dataset. After batch l of final choice letters is sent for list j, $(L_{j,l+1}, K_{j,l+1})$ becomes the next trigger and batch policy.
- **Final Choices:** I use a reduced form model to capture the sensitivity of the final choice to position information. Given initial choice C_i, 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.

Given estimated parameters, I draw sequences of inputs and run the Cambridge Mechanism for 500 years. 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 otherwise follows an exponential distribution with mean $\frac{1}{a_1}$ years after the arrival date. 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. Waiting times converged after about 10 years. I used the last 480 years of the simulation to construct beliefs.

C.2.3 Constructing Beliefs from Simulation Outputs

The simulation produces the state of all Cambridge waiting lists every day for 480 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 choices the applicant could have made.

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

Constructing beliefs $\{F_{j,C}(. \mid \mathbf{p}_C)\}_{j,C,\mathbf{p}_C}$ 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 $\mathbf{p}_C$ induces a different continued waiting time distribution. Even using the simulation results, there is a limit to how flexibly these distributions can be estimated. My approach is to specify a hierarchical parametric model for the continued waiting time distribution. I assume that continued waiting time follows a beta distribution

$$T_j \mid j, C, \mathbf{p}_C \sim Beta(\alpha_{j,C}(\mathbf{p}_C), \beta_{j,C}(\mathbf{p}_C))$$

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

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

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

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

C.3 Development Preferences

C.3.1 Moments

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

- Application rates by income and demographics. I use the following indicator variables for $\mathbf{Z}_i$: a dummy equal to 1 for all households; annual household income in the ranges of $[X, X+10,000]$ for X in \$10,000 intervals from \$0 to \$50,000; household head is black, and hispanic; the household currently lives in Cambridge; the youngest household member is under age 10; the household qualifies for three bedrooms; and household income is below \$20,000. I also match the rate at which all households and households earning \$0-\$20,000 and \$20,000-\$40,000 select three developments in their initial choice.
- Development shares: there is one moment for the initial choice share 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, and the covariance between chosen development size and whether the household head is hispanic, the household requires three bedrooms, and the household's youngest member is less than 10 years old.
- 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 and variances of chosen waiting times within and between applicants, by income and demographics. The first and second waiting time moments are interacted with household income bins for \$0-\$20,000, \$20,000-40,000, and \$40,000+.
- The final choice moments used are:
  - The fraction of eligible households who made a final choice:

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

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

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

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

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

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

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

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

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

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

C.3.2 Simulation Procedure

The method of simulated moments estimates $E(\mathbf{m}_i \mid \mathbf{Z}_i, \boldsymbol{\theta})$ in the following steps:

(i) For each sampled household i and simulation draws $s = 1, ..., S$,

  (a) Draw preference shocks $\{\eta_{is}, \boldsymbol{\nu}_{is}, \boldsymbol{\epsilon}_{is}\}$.

  (b) For each possible initial choice C, draw the date and position information of the final choice $(s_{is}^C, \mathbf{p}_{is}^C)$, drawn from the distribution $G_C(S_C, \mathbf{P}_C)$.

  (c) Draw an exogenous departure time using the attrition model. This determines whether the simulated applicant makes a final choice for a given final choice date s_{is}^C.

(ii) For each proposed value of $\boldsymbol{\theta}$ and each (i, s),

  (a) Compute $\mathbf{v}_{is}$ according to equation 10 given $\mathbf{z}_i$, $\boldsymbol{\theta}$, and the simulation draws.

  (b) Compute the optimal initial choice C_{is} according to equation 4 given $\mathbf{v}_{is}$, the discount factor ρ, and waiting time beliefs.

  (c) If the exogenous departure date is after the final choice date, compute the applicant's final choice according to equation 3 given preferences and beliefs.

(d) Construct the conditional expectations

$$\hat{E}(\mathbf{m}_i \mid \mathbf{z}_i, \boldsymbol{\theta}) = \frac{1}{S} \sum_{s=1}^{S} \mathbf{m}_{is}(\boldsymbol{\theta})$$

and form moment conditions.

The one non-standard component of the simulation comes from the applicant's two-stage decision problem. Different parameter values $\boldsymbol{\theta}$ will lead a simulated applicant to make different initial choices, inducing a different distribution over final choice states. I draw one final choice state for each possible initial choice and hold these draws fixed across candidate parameter values.

C.3.3 Objective Function and Optimization

Because the moments used in estimation are highly correlated, the optimal weight matrix performed poorly. The estimator 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 use a diagonal weight matrix with elements inversely proportional to the sampling variance of the corresponding moment functions. I also placed additional weight on application rates, means and variances of chosen waiting times, and final choice moments.

Minimizing the objective function was challenging because the objective function is discontinuous and not guaranteed to be convex. Fortunately, Monte-Carlo simulations suggested that a combination of global and local search consistently found a global minimum close to the true parameters. I used the following procedure: I first used MATLAB's *fmincon* function, approximating the gradient by finite differences. I found that iteratively decreasing the finite difference minimum step size, using the previous solution as a starting value, helped to ensure that the estimator searched widely while also finding a local minimum. At each local minimum, I used MATLAB's *patternsearch* algorithm to ensure that an exact local minimum was attained and to search for other local minima. I used several starting values covering a range of parameters. To limit numerical instability, the variance of each random coefficient was constrained to be less than one million.

C.3.4 Inference

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

The asymptotic variance of the method of simulated moments estimator is

$$(\mathbf{G}'\mathbf{A}\mathbf{G})^{-1}\mathbf{G}'\mathbf{A}\boldsymbol{\Omega}\mathbf{A}\mathbf{G}(\mathbf{G}'\mathbf{A}\mathbf{G})^{-1},$$

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

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

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

$$\boldsymbol{\Omega} = \boldsymbol{\Omega}_m + \frac{1}{S}\boldsymbol{\Omega}_s \,.$$

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

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

$\boldsymbol{\Omega}_s$ can be estimated consistently by

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

where

$$\hat{\mathbf{m}}_i(\hat{\boldsymbol{\theta}}) = \frac{1}{S}\sum_{s=1}^{S} \mathbf{m}_{is}(\hat{\boldsymbol{\theta}}) \,.$$

The variance estimate is

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

C.4 Counterfactual Simulations

C.4.1 Computational Details

To compute counterfactual equilibria, I draw 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 IVA, and then draw flow payoffs given those characteristics using the estimates from Specification (3) of the structural model. Apartment vacancies and exogenous departure dates are drawn from the same distributions used to construct beliefs in Section IVB.

These sequences are used to compute counterfactual allocations under all mechanisms. In computing features of the equilibrium and allocation, the first 20 years are discarded to allow the waiting list to approach steady state. All applicants are eligible for all 13 public housing developments, and all waiting lists remain open during the entire simulation. This abstracts from temporary list closures (which do occur 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 search for a fixed point between applicants' choices and the implied weights $\{w_{j,C}(\psi_\varphi(y_i))\}_{C \in \mathcal{C}_\varphi}^{j=1,\ldots,J}$. The algorithm works as follows. Iteration q begins with a vector of proposed weights $\mathbf{w}^{(q)}$. The following steps then occur:

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

2. The waiting list is run, yielding predicted weights $\mathbf{w}^{(q)'}$ with distance $D^{(q)} = \|\mathbf{w}^{(q)'} - \mathbf{w}^{(q)}\|$. To calculate the predicted weights, I consider the experience of one additional applicant on each possible application date in the simulation (after the 20-year burn-in period). For each possible choice C, list j, and priority group ψ, I calculate the waiting time for list j, T_j, and the indicator $1\{T_j \leq \min_{k \in C} T_k\}$ for whether a development j apartment would arrive before any other development in C if the applicant arrived on that particular date. The weights are then calculated using the sample analog of equation 14:

$$w_{j,C}^{(q)'}(\psi) \equiv \frac{1}{\rho} \sum_{d \in D} e^{-\rho\, t_{j,d}(\psi)} * 1\left\{t_{j,d}(\psi) = \min_{k \in C} t_{k,d}(\psi)\right\},$$

   where $t_{j,d}(\psi)$ is the waiting time for an applicant in priority group ψ arriving on date d, and D is the set of application dates averaged over in the simulation. Thus, as in section IVB, expectations are constructed from the empirical distribution of waiting times generated by the simulation. This method has the benefit of fully accounting for transition dynamics as the number and types of applicants in the queue fluctuate over time.

3. Weights are updated as a convex combination of the proposed and implied weights:

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

The factor λ determines how aggressively the weights are updated. If $\lambda = 1$, then the weights implied by applicant choices ($\mathbf{w}^{(q)'}$) 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 percent 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. While multiple equilibria are theoretically possible under some of the development choice systems considered in the paper, I did not find evidence of multiplicity.

C.4.2 Additional Development Choice Systems

The PHAs surveyed in Table 8 use development choice systems that fall into four additional categories other than Limited Choice and No Choice:

- **Choose Any Subset**: $\mathcal{C} = 2^{\{1,...,J\}}$. Applicants may choose any subset of developments, as in Boston and San Antonio.
- **Choose All or One**: $\mathcal{C} = \{\{1\}, ..., \{J\}, \{1, ..., J\}\}$. Applicants may either wait for their preferred development or take the first available offer from any development. This choice system approximates the policies used in Philadelphia, Baltimore, and Newark.
- **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.
- **Choose All or Neighborhood**: $\mathcal{C} = \{C_{\text{north}}, C_{\text{east}}, C_{\text{central}}, \{1, ..., J\}\}$. Applicants may either choose a neighborhood or receive the first offer city-wide. Chicago uses this development choice system for Family Public Housing.

In all of these systems, the applicant is offered the first available apartment from any development in their chosen set C.

Table 16 presents counterfactual simulations for these four choice systems, along with Choose One and No Choice, under Equal Priority. As expected, these systems produce allocations that are between Choose One and No Choice in terms of efficiency and redistribution. Choose Any Subset and Choose All or One, which allow applicants to select several developments as a hedge against waiting time uncertainty, have modest 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. Each neighborhood contains at least three developments, so some applicants reject offers, lowering match quality and improving targeting relative to Choose One. However, these mechanisms achieve higher match quality than No Choice by giving applicants some choice over where they might be assigned.

D Robustness to Alternative Modeling Assumptions

The estimation procedure outlined in sections III and IV and implemented in section V relies on a particular model of applicants' decision rule, beliefs, and underlying utility functions. This section checks whether the main findings in the paper are sensitive to the specific assumptions made about these objects. It repeats the analysis in the paper under four departures from the baseline model. In the first, applicants use a simpler heuristic to make their first-stage development choices, relaxing the degree of applicant sophistication in the main analysis. The second departure explores belief heterogeneity by assuming that half of applicants are completely naive, simply choosing their preferred developments without considering waiting time, while the remainder are sophisticated. The third departure assumes that applicants' beliefs follow the empirical distribution of waiting times in the data. Finally, the fourth alternative model assumes a minimum level of housing expenditure; this is motivated by the empirical fact that higher-income households spend a lower fraction of their incomes on rent.

The rest of this section describes the four alternative models in more detail and then presents the main results. While the quantitative results differ somewhat across these specifications, the trade-off between efficiency and redistribution is qualitatively robust. The full set of parameter estimates and counterfactual predictions are omitted for brevity, but the author is happy to provide additional details upon request.

D.1 Alternative Models

This section details how each model is implemented in estimation and counterfactuals.

D.1.1 A Simpler First-Stage Decision Rule

In the benchmark development choice model, applicants anticipate that they will receive new information in the final choice stage, and understand that this generates a portfolio choice problem at the initial choice stage. This decision rule and belief structure entail a high degree of sophistication in a socioeconomically disadvantaged population.

This section repeats the analysis in the paper assuming that applicants do not consider the full complexity of the portfolio choice problem generated by the Cambridge Mechanism. Instead, in the initial choice stage applicants use a heuristic: they consider the value of applying for each development on its own and select the developments with the highest expected value according to this criterion. This heuristic choice rule rules out certain types of sophisticated behavior. For example, in a portfolio choice problem it can be optimal to select a development for its option value – even if it has a longer expected waiting time and is therefore much less likely to be eventually chosen than another slightly less desirable development, it may yield a greater increase in the value of the applicant's portfolio. It may also be optimal to omit a development in order to delay the timing of the final choice stage and obtain a more precise measure of continued waiting time.

Formally, at the initial choice stage applicants form beliefs about the marginal distribution of waiting times for each development j. Let $G_j(t)$ denote the believed probability that the waiting time for development j is less than t years. At the final choice stage, applicants form beliefs in the same way as in the sophisticated model, taking all list positions p into account when predicting the continued waiting time for development j. Let $F_{j,C}(t \mid p)$ denote the probability that continued waiting time for development j is less than t years given current list positions p. At final choice, the applicant solves the problem defined in equation 3, just as in the sophisticated model. At initial choice, the applicant solves a different problem than the one defined in equation 4:

$$\max_{C \in \{0,1,\ldots,J\}^3} \sum_{j \in C} E\left[e^{-\rho T}\right](v_{ij} - v_{i0}) \tag{19}$$

$$= \max_{C \in \{0,1,\ldots,J\}^3} \sum_{j \in C} \int \frac{1}{\rho} e^{-\rho T}(v_{ij} - v_{i0}) dG_j(T) \ . \tag{20}$$

In estimation, the distributions G_j and $F_{j,C}$ come from the same simulation that generated the belief distributions used for the main estimates. Therefore, beliefs are consistent across the two stages of choice in the sense that they are generated by a simulation that respects the structure of the Cambridge Mechanism. However, those beliefs are now inputs to a decision rule that is suboptimal because of the heuristic employed in the first stage.

I estimate the specifications in Section V under the simpler decision rule using the same procedure as in section IV. The only difference is the decision rule and belief objects used to predict a simulated applicant's development choices in the method of simulated moments procedure. Then, I re-solve for counterfactual equilibria under alternative mechanisms using estimates obtained under specification (3) with the simpler decision rule.

D.1.2 Applicant Confusion

It is also possible that some applicants are confused about the application process and do not strategize at all. The estimates in the paper might interpret applicants' limited responsiveness to waiting time as evidence that applicants are patient and have very heterogeneous preferences, when in fact applicants are impatient but confused.[4] With the data available, one cannot distinguish preference heterogeneity from arbitrary confusion or belief heterogeneity. However, one natural type of heterogeneity to explore is that some fraction of applicants are sophisticated, while the remainder are non-strategic and simply choose their preferred development without considering waiting time. Non-strategic behavior would be a reasonable response to confusion about waiting times and the structure of the Cambridge Mechanism, and it would also rationalize the observed responsiveness to waiting time information with greater

[4] I thank an anonymous referee for pointing out this form of potential bias.

impatience. I therefore explore how the fraction of non-strategic applicants would affect the trade-off between efficiency and redistribution in waitlist design.

Formally, I assume that a fraction NS of applicants are non-strategic, while the remaining applicants make their decisions as in the main analysis according to equations 3 and 4. Whether an applicant is strategic is uncorrelated with their preferences. A non-strategic applicant simply chooses their preferred development(s) at each stage in the application process without considering waiting time. At the final choice stage, given choice set C, a non-strategic applicant solves

$$\max_{j \in C} v_{ij} - v_{i0} \ . \tag{21}$$

In the first stage, a non-strategic applicant chooses the three most preferred developments, or the developments preferred to their outside option:

$$\max_{C \in \mathcal{C}} \sum_{j \in C} v_{ij} - v_{i0} \ . \tag{22}$$

As in the main analysis, applicants apply if and only if some development is preferred to their outside option.

This analysis also requires an assumption about the information structure and decision rules in counterfactuals. To hold the information environment constant, I assume that the same fraction of applicants are non-strategic in estimation and counterfactual exercises. Under Choose One, non-strategic applicants simply list their preferred development. As in the main analysis, equilibrium waiting time distributions are determined by the decisions of all applicants, including non-strategic ones. The belief distributions used for estimation are the same as in the main analysis; because the simulation of the Cambridge Mechanism relies on reduced-form policy functions, different information assumptions will not affect the simulation inputs, only how the structural model interprets those choice patterns.

Section D.2 reports counterfactual results assuming half of applicants are non-strategic ($NS = 0.5$). I have also run specifications with $NS = 0.25$ and $NS = 0.75$, though the results are omitted for brevity. Consistent with the above intuition, the model estimates a higher discount rate (more impatience) when a greater fraction of applicants are non-strategic. However, the estimated heterogeneity in match values and values of assistance remains similar, as does the trade-off between efficiency and redistribution.

D.1.3 Beliefs Matching Empirical Waiting Time Distribution

As noted in Section VB, while the simulation of the Cambridge Mechanism broadly captures differences in waiting times across developments, the waiting times generated by the simulation do not perfectly match those in the data. The estimates in the main paper assume that applicants' beliefs are governed by the simulation outputs in order to accommodate limited data on realized waiting times relative to the dimensionality of sophisticated applicants' beliefs. This section considers an alternative model of belief formation in which applicants' beliefs are initially governed by the empirical waiting time distributions observed in the data. This alternative assumption is attractive because it relies less on a specific model of the CHA waitlist; instead, it assumes that applicants form beliefs based on experience. Realized waiting times are something applicants could inquire about, if not directly observe.

In order to use the empirical initial waiting time distributions, the analysis makes two simplifications. First, applicants use the simpler first-stage decision rule defined in Section D.1.1, taking the empirical marginal distribution of waiting times as the initial waiting time distribution $\tilde{G}_j(.)$ for each development. This avoids estimating the joint distribution of waiting times across developments with only a couple hundred observations. Second, the data are too sparse to estimate final choice distributions $\tilde{F}_{j,C}(p)$ directly, so I still use the waiting time simulation to construct the continued waiting time distributions. To make the second-stage distributions more consistent with the empirical waiting time distributions in the first stage, I adjust vacancy rates so that the *initial* waiting time distributions generated by the simulation are close to the empirical average for each development, while keeping the other simulation inputs as in the main analysis. Since the goal of this exercise is to assess robustness of the paper's conclusions to alternative models of belief formation, the loss of internal consistency in the beliefs model is less concerning.

The preference estimates obtained using beliefs $\{\{\tilde{G}_j\}, \{\tilde{F}_{j,C}(T_j \mid p)\}_{j,p}\}_{C \in \mathcal{C}}$ are then used as inputs to the counterfactual simulations.

D.1.4 Relaxing Homotheticity

The first three alternative models explore robustness to specific assumptions about applicants' beliefs and decision rules. The paper also makes assumptions about the structure of applicants' underlying utility functions, which is important for welfare analysis. The Cobb-Douglas structure assumed in Section IVC implies that applicants spend a constant fraction of their income on housing, regardless of their incomes. This is at odds with survey data, which consistently shows lower-income households spending a larger fraction of their income on rent and other housing costs.

A natural way to accommodate non-homotheticity is to assume that households require a minimum level of housing expenditure $\underline{h}$. They then choose (c, h) outside of public housing to maximize

$$\tilde{u}(c, h; \underline{h}) \equiv \gamma \log c + (1 - \gamma) \log(h - \underline{h}) \ . \tag{23}$$

This utility function predicts that households in private market housing spend a lower fraction of their incomes on rent as their income rises. For this reason, it has been used in recent structural empirical work on housing (e.g. Sieg and Yoon (2020)). To accommodate a minimum required level of housing consumption, the minimum consumption level is also adjusted to $\underline{c} + \underline{h}$, and the distribution of η in equation 11 is modified accordingly:

$$\eta_i \overset{iid}{\sim} \begin{cases} TN(0, \sigma_\eta^2, \underline{c} + \underline{h} - y_i, \infty) & w.p. \;\; 1 - \Phi\left(\frac{\underline{c}+\underline{h}-y_i}{\sigma_\eta}\right) \\ \underline{c} + \underline{h} - y_i & w.p. \;\; \Phi\left(\frac{\underline{c}+\underline{h}-y_i}{\sigma_\eta}\right) \end{cases} \tag{24}$$

The remaining assumptions governing applicants' beliefs and decision rules, as well as the equivalent variation formula, are unchanged. The resulting preference estimates are inputs to counterfactual simulations. The next section presents results for $\underline{h} \equiv \$10,000$, which is close to minimum level of housing expenditure for any income group in the ACS (among likely eligible households). Similar conclusions hold for $\underline{h} \equiv \$5,000$, and are omitted for brevity.

D.2 Results

Appendix Table 5 compares the equilibrium allocations under Choose One and No Choice for each of the four alternative models. The qualitative trade-offs between efficiency and redistribution are similar across specifications. Moving from Choose One to No Choice lowers efficiency by 28-39 percent (compared to 35 percent in the baseline model) and leads to a loss of several thousand dollars of equivalent variation per unit allocated. Doing so significantly increases redistribution; tenant incomes fall by 13-17 percent (16 percent in the baseline model), and tenants' outside options fall by 16-22 percent (19 percent in the baseline model). The fraction of extremely high-need tenants also significantly increases under all four alternative models. The other findings discussed in section VIB.1 also hold qualitatively. Moving to No Choice, the fraction of tenants assigned their first choice falls dramatically; mean waiting times are shorter; and the degree of segregation across developments falls as applicants exercise less choice over where they are assigned. The predicted characteristics of tenants are similar to the baseline model in terms of observed income, race, and demographics.

There are some quantitative differences in the results across models. Applicants are estimated to be even more patient in when beliefs match the empirical waiting time distribution (columns 5-6). Because of this, most applicants are assigned their first choice development under Choose One, and almost all are assigned one of their top three choices. The loss in match quality, and therefore the overall efficiency loss from eliminating choice, is especially high. The welfare estimates also depend somewhat on the structure of underlying utility. With a \$10,000 minimum housing expenditure (columns 7-8), many more applicants are predicted to be at their consumption minimum. Tenants' outside options are much worse than in other specifications, and the fraction of extremely high-need tenants is higher. Tenants require smaller cash transfers to generate utility increases equivalent to those generated by their assignments, and so equivalent variation per unit assigned and overall efficiency are also significantly lower than in other models. Nevertheless, under all four alternative models, the trade-off between efficiency and redistribution exists and is quantitatively important.

Finally, Appendix Figure 5 compares social welfare under Choose One and No Choice for different levels of inequality aversion under each of the four alternative models. In all cases, a moderate to high degree of inequality aversion is needed to justify eliminating choice under the CHA's current priority system, ranging from $\lambda \approx 1.7$ when beliefs match the empirical waiting time distribution to $\lambda \approx 4$ with $\$10,000$ minimum housing expenditure. In the former case, a social planner indifferent between Choose One and No Choice would be willing to take about \$3.20 from a \$20,000 income household in order to transfer \$1 to a \$10,000 income household (burning \$2.20 in the process). No Choice is especially effective at increasing redistribution in this model because of applicants' high degree of patience; there is minimal targeting gain through applicants' choice of waiting time under Choose One. In the latter case, the planner would be willing to burn \$15 to make the same transfer, arguably an extreme preference for redistribution. The proportional increase in redistribution is relatively small with a minimum housing expenditure because so many tenants are near their consumption minimum under Choose One.

Table 8: Allocation Policies Used by Public Housing Authorities

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

Notes: Features of allocation mechanisms used by PHAs in 25 cities. PHAs were chosen based on city population and/or the size of their public housing stocks. * indicates that the PHA's administrative plan was not available online. In these cases, information was gleaned from the PHA website and application forms. A High SES priority system favors households above 30 percent of AMI, or which are economically self-sufficient or have a working member. A Low SES priority system prioritizes households below 30 percent of 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 9: Final Choice Balance Tests

	p-value	F-statistic
Panel A: List Position		
Date of Application	0.000	2.32
Date of Final Choice Letter	0.000	3.00
Annual Income	0.796	0.92
Has Labor Income	0.991	0.78
Has Public Assistance Income	0.844	0.90
Lives in Cambridge	0.623	0.97
Works in Cambridge	0.102	1.13
# Household Members	0.000	1.45
# Earners	0.982	0.80
# Adults	0.138	1.11
# Children	0.002	1.33
White Household Head	0.891	0.88
Black Household Head	0.334	1.04
Hispanic Household Head	0.989	0.79
Age of Household Head	0.137	1.11
Male Household Head	0.132	1.11
# Children under Age 10	0.109	1.13
Panel B: Continued Waiting Time		
Date of Application	0.000	3.60
Date of Final Choice Letter	0.000	2.41
Annual Income	0.968	0.82
Has Labor Income	0.998	0.73
Has Public Assistance Income	0.892	0.88
Lives in Cambridge	0.540	0.99
Works in Cambridge	0.176	1.09
# Household Members	0.001	1.37
# Earners	0.993	0.78
# Adults	0.054	1.17
# Children	0.011	1.25
White Household Head	0.893	0.88
Black Household Head	0.858	0.89
Hispanic Household Head	0.984	0.80
Age of Household Head	0.064	1.16
Male Household Head	0.574	0.98
# Children under Age 10	0.070	1.15

Notes: F-statistics and p-values from a joint test of significance for regression coefficients predicting applicant characteristics as a function of relative list lengths at final choice. Panel A measures list length using list position number, while Panel B uses expected continued waiting time. A different applicant characteristic is the dependent variable in each row.

Table 10: Testing for Responsiveness to List Position at Initial Choice

Development	p-value	F-statistic	DF(1)	DF(2)
Panel A: List Positions Only				
Corcoran Park	0.517	0.981	57	1661
East Cambridge	0.033	1.373	59	1660
Jackson Gardens	0.679	0.905	59	1660
Jefferson Park	0.114	1.235	57	1661
Lincoln Way	0.440	1.018	59	1660
Mid Cambridge	0.002	1.647	59	1660
Newtowne Court	0.090	1.266	57	1661
Putnam Gardens	0.458	1.009	57	1661
River Howard Homes	0.327	1.075	59	1660
Roosevelt Low-Rise	0.114	1.236	57	1661
Washington Elms	0.041	1.358	57	1661
Woodrow Wilson	0.084	1.269	59	1660
Roosevelt Mid-Rise	0.494	0.982	30	1144
Panel B: Applicant Covariates				
Corcoran Park	0.635	0.925	57	1641
East Cambridge	0.189	1.163	59	1640
Jackson Gardens	0.728	0.881	59	1640
Jefferson Park	0.175	1.177	57	1641
Lincoln Way	0.646	0.921	59	1640
Mid Cambridge	0.022	1.415	59	1640
Newtowne Court	0.243	1.127	57	1641
Putnam Gardens	0.262	1.115	57	1641
River Howard Homes	0.700	0.895	59	1640
Roosevelt Low-Rise	0.142	1.206	57	1641
Washington Elms	0.083	1.276	57	1641
Woodrow Wilson	0.101	1.246	59	1640
Roosevelt Mid-Rise	0.408	1.040	30	1125

Notes: F-statistics and p-values from tests for whether list positions predict applicants' choices at initial application. The sample is applicants in the structural estimation sample. For each development, the probability that each applicant chose that development initially is predicted as a function of the length of each list in its choice set. The F-statistic jointly tests for the significance of all coefficients on list position. Panel B adds controls for household income, race/ethnicity, and neighborhood of current residence if the household already lives in Cambridge.

Table 11: Sample Size by Stage of Application Process

Households Who	N
Are Predicted Eligible	6818
Made an Initial Choice	1725
Made a Final Choice	573
Received an Offer of Housing	163

Notes: Number of eligible households who made it to each stage of the application process. Households are restricted to the sample used in estimation. Final choices and housing offers occurred during the 2010-2014 sample period.

Table 12: Inputs to Waiting Time Simulation

Parameter	Value
Panel A: Apartment Vacancies	
Annual Vacancy Rate per Unit	0.10
Annual Vacancy Rate Total	108
Panel B: Applicant Arrivals and Departures	
Daily Applicant Arrival Rate	0.945
Annual Applicant Arrival Rate	345
Instant Departure Probability	0.239
Annual Departure Rate	0.222
Panel C: Final Choice Model	
List Position Coefficient	-0.021
Fixed Effects	
Corcoran Park	0.358
East Cambridge	-0.162
Jackson Gardens	0.304
Jefferson Park	-0.447
Lincoln Way	0.678
Mid Cambridge	0.241
Newtowne Court	0.073
Putnam Gardens	-0.303
River Howard Homes	0.000
Roosevelt Low-Rise	-0.597
Washington Elms	-0.335
Woodrow Wilson	-0.258
Roosevelt Mid-Rise	-0.878

Table 13: Coefficient Estimates Predicting Probability in CHA Dataset

	Point Estimate	90% Confidence Interval
Income $0-$8,000	0.77	[-0.56 , 14.75]
Income $8,000-$16,000	0.46	[-0.8 , 19.95]
Income $16,000-$32,000	-0.35	[-1.26 , 0.8]
Income $32,000-$48,000	-6.32	[-15.63 , -1.73]
Income Above $48,000	-7.16	[-19.92 , -2.52]
African American Household Head	4.38	[1.46 , 12.73]
Hispanic Household Head	-1.36	[-2.65 , 5.53]
Household lives in Cambridge	-1.17	[-2.24 , 3.38]

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 2010-2014 5 percent sample. The 90 percent confidence intervals are bootstrapped by re-sampling the ACS with replacement and re-running the estimation procedure.

Table 14: Simulated Waiting Times from Initial Application

	Simulation		Data	
Development	Mean	S.D.	Mean	S.D.
Corcoran Park	3.06	0.97	2.85	1.57
East Cambridge	5.60	1.99	3.91	1.05
Jackson Gardens	6.79	1.63	2.39	1.47
Jefferson Park	1.15	0.52	1.79	1.23
Lincoln Way	4.26	1.25	3.54	--
Mid Cambridge	5.98	2.20	3.91	1.05
Newtowne Court	2.36	0.77	2.18	1.65
Putnam Gardens	3.67	1.04	2.62	2.10
River Howard Homes	6.73	2.16	3.91	1.05
Roosevelt Low-Rise	2.55	0.78	3.04	1.33
Washington Elms	2.58	0.78	3.01	2.03
Woodrow Wilson	4.65	1.70	1.98	0.64
Roosevelt Mid-Rise	5.65	2.12	1.88	0.09

Notes: Realized waiting times are averaged across all applicants housed in each development during the simulation.

Table 15: Parameter Estimates, Full

	Baseline Specification		Richer Observed Heterogeneity		Unobserved Taste for Size and Location	
	(1)		(2)		(3)	
Annual Discount Rate	0.977	(0.011)	0.977	(0.008)	0.966	(0.009)
Development 1 Fixed Effect	-0.071	(0.018)	-0.079	(0.021)	-0.077	(0.025)
Development 2 Fixed Effect	-0.081	(0.024)	-0.076	(0.032)	-0.054	(0.034)
Development 3 Fixed Effect	-0.070	(0.037)	-0.058	(0.051)	-0.058	(0.044)
Development 4 Fixed Effect	-0.159	(0.023)	-0.226	(0.022)	-0.219	(0.048)
Development 5 Fixed Effect	0.056	(0.025)	0.069	(0.034)	0.067	(0.029)
Development 6 Fixed Effect	0.000	(0.021)	0.000	(0.036)	-0.094	(0.055)
Development 7 Fixed Effect	0.122	(0.013)	0.118	(0.023)	0.108	(0.033)
Development 8 Fixed Effect	-0.584	(0.093)	-0.622	(0.078)	-0.779	(0.079)
Development 9 Fixed Effect	0.008	(0.034)	0.040	(0.025)	0.084	(0.044)
Development 10 Fixed Effect	0.038	(0.01)	0.037	(0.02)	0.014	(0.022)
Development 11 Fixed Effect	-1.169	(0.084)	-1.399	(0.096)	-1.401	(0.098)
Development 12 Fixed Effect	0.001	(0.023)	0.003	(0.032)	0.003	(0.026)
Development 13 Fixed Effect	-0.023	(0.033)	0.002	(0.03)	0.046	(0.024)
S.D. Development Fixed Effects	0.352		0.415		0.432	
Panel A: Value of Assistance						
Head Is Black	0.933	(0.092)	0.838	(0.061)	0.839	(0.082)
Head Is Hispanic	0.032	(0.043)	0.138	(0.046)	0.083	(0.062)
Lives In Cambridge	0.528	(0.065)	0.384	(0.045)	0.381	(0.036)
Youngest Member < 10 Years			0.005	(0.041)	-0.018	(0.037)
3 Bedroom Household			0.258	(0.047)	0.259	(0.053)
Household Income < $20,000			0.321	(0.068)	0.320	(0.059)
Log Of Observed Income	0.164	(0.08)	0.158	(0.058)	0.166	(0.066)
Log Of Observed And Unobserved Income	-1.000	--	-1.000	--	-1.000	--
Scale of R.E. Unknown Income ($10,000)	1.115	(0.11)	1.115	(0.109)	1.090	(0.081)
Panel B: Match Values						
Applicant and Development Same Neighborhood	-0.137	(0.065)	-0.196	(0.031)	-0.193	(0.058)
Applicant Head Is Hispanic * Development Size			0.022	(0.033)	0.043	(0.043)
Youngest Member < 10 Years * Development Size			0.000	(0.016)	-0.006	(0.021)
Household Income < $20,000 * Development Size			0.000	(0.022)	-0.003	(0.021)
S.D. Unobserved Taste For Development Size					0.039	(0.011)
S.D. Unobserved Taste for North Cambridge					0.035	(0.019)
S.D. Unobserved Taste for East Cambridge					0.039	(0.013)
S.D. Idiosyncratic Shock	0.161	(0.013)	0.155	(0.01)	0.156	(0.015)

Table 16: Effects of Alternative Development Choice Systems under Equal Priority

	Choose One (1)	Choose Any Subset (2)	Choose All or One (3)	Choose Neighborhood (4)	Choose All or Neighborhood (5)	No Choice (6)
Panel A: Welfare Gain and Cost of Allocation						
Equivalent Variation ($)	22,318	22,349	22,318	16,617	16,612	14,617
Cost per Unit ($)	18,914	18,932	18,915	19,622	19,611	19,735
Equivalent Variation per $ Cost to Gvt.	1.18	1.18	1.18	0.85	0.85	0.74
Panel B: Targeting						
Observed Income ($)	16,953	16,894	16,952	14,595	14,629	14,217
Unobserved Income ($)	-2,548	-2,555	-2,545	-2,894	-2,926	-2,615
Observed and Unobserved Income ($)	14,405	14,338	14,407	11,701	11,703	11,602
% Extremely High-Need	29.7%	29.8%	29.7%	37.6%	37.5%	35.5%
Panel C: Match Quality						
% Assigned Top Choice	58.1%	56.7%	58.1%	16.3%	16.2%	8.9%
% Assigned Top 3	79.7%	79.4%	79.7%	35.6%	35.8%	24.4%
Panel D: Characteristics of Housed Applicants						
Waiting Time (days)	1688	1694	1688	805	807	643
% Black	52.4%	52.5%	52.4%	56.2%	56.1%	55.2%
% Hispanic	18.9%	19.0%	18.9%	17.3%	17.3%	17.0%
From Cambridge	61.4%	61.2%	61.4%	61.2%	61.1%	62.0%
Panel E: Distribution of Tenants across Developments						
S.D. Observed Income ($)	4,378	4,299	4,378	3,386	3,469	2,413
Range of Observed Income ($)	[10774 , 25107]	[10761 , 24922]	[10774 , 25107]	[9273 , 20486]	[9213 , 20676]	[8884 , 17163]
S.D. % Black	15.3%	15.2%	15.3%	6.7%	6.6%	3.5%
Range of % Black	[42.3% , 100.0%]	[42.6% , 100.0%]	[42.3% , 100.0%]	[48.0% , 70.6%]	[47.8% , 71.4%]	[51.6% , 64.0%]
S.D. % From Cambridge	8.8%	8.7%	8.8%	5.3%	5.5%	2.8%
Range of % From Cambridge	[47.8% , 72.1%]	[48.0% , 72.1%]	[47.8% , 72.1%]	[51.8% , 66.0%]	[50.9% , 66.8%]	[59.1% , 69.3%]

Notes: Statistics averaged across assigned apartments in each counterfactual simulation. Dollar amounts are annual. Cost per Unit is calculated based on market rate rental prices in Cambridge, MA during the sample period. Equivalent Variation is the equivalent cash transfer outside of public housing that would generate the same welfare change for a housed applicant as their assignment. Extremely High-Need applicants are at the minimum consumption level of $10/day outside of public housing.

Table 17: Effects of Development Choice under Alternative Modeling Assumptions

	Simple First-Stage Decision		50% of Applicants Naïve		Beliefs Match Empirical Waiting Time Distributions		$10,000 Minimum Housing Expenditure	
	Choose One	No Choice	Choose One	No Choice	Choose One	No Choice	Choose One	No Choice
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Welfare Gain and Cost of Allocation								
Equivalent Variation ($)	16,006	11,500	24,133	17,464	21,071	15,112	14,398	8,758
Cost per Unit ($)	18,597	19,503	18,646	19,318	18,365	19,231	19,087	19,763
Equivalent Variation per $ Cost to Gvt.	0.86	0.59	1.29	0.90	1.15	0.79	0.75	0.44
Panel B: Targeting								
Observed Income ($)	18,010	14,991	17,847	15,607	18,782	15,898	16,375	14,123
Unobserved Income ($)	-3,170	-3,224	-2,681	-2,939	-3,553	-4,048	-8,642	-8,025
Observed and Unobserved Income ($)	14,839	11,767	15,166	12,668	15,229	11,849	7,734	6,098
% Extremely High-Need	27.1%	31.6%	28.8%	32.3%	27.4%	34.7%	55.0%	63.9%
Panel C: Match Quality								
% Assigned Top Choice	56.5%	9.5%	76.9%	11.2%	90.0%	11.7%	61.5%	8.9%
% Assigned Top 3	89.0%	26.6%	91.8%	30.4%	99.4%	31.3%	85.1%	23.5%
Panel D: Characteristics of Housed Applicants								
Waiting Time (days)	1350	927	1416	1077	1135	739	1972	768
% Black	52.1%	54.9%	51.8%	54.0%	51.6%	53.8%	49.7%	53.8%
% Hispanic	22.9%	19.5%	21.0%	17.9%	18.1%	17.2%	18.2%	15.6%
From Cambridge	55.8%	56.7%	57.7%	58.6%	57.5%	58.7%	55.0%	55.6%
Panel E: Distribution of Tenants across Developments								
S.D. Observed Income ($)	3,267	907	3,387	1,037	2,095	352	3,416	1,910
Range of Observed Income ($)	[15349,27194]	[13375,16829]	[11496,24028]	[13653,16850]	[17058,22258]	[15423,16673]	[10490,23150]	[10074,16348]
S.D. % Black	3.7%	1.7%	5.9%	2.3%	7.6%	2.8%	8.7%	12.9%
Range of % Black	[45.6%,58.1%]	[52.6%,57.4%]	[42.7%,65.7%]	[49.8%,58.9%]	[38.3%,63.9%]	[50.2%,59.0%]	[41.1%,68.4%]	[42.7%,93.1%]
S.D. % From Cambridge	10.9%	2.2%	9.4%	1.9%	9.9%	2.2%	13.2%	3.3%
Range of % From Cambridge	[40.7%,71.2%]	[53.1%,60.9%]	[43.5%,68.0%]	[56.1%,63.1%]	[47.9%,72.1%]	[54.8%,63.0%]	[38.2%,71.9%]	[50.4%,61.6%]

Notes: Counterfactual estimates are obtained identically to those in Table 7 but under the alternative modeling assumptions described in Online Appendix Section D. Applicants receive Equal Priority in all simulations.

Figure 3: Locations of Cambridge Family Public Housing Developments

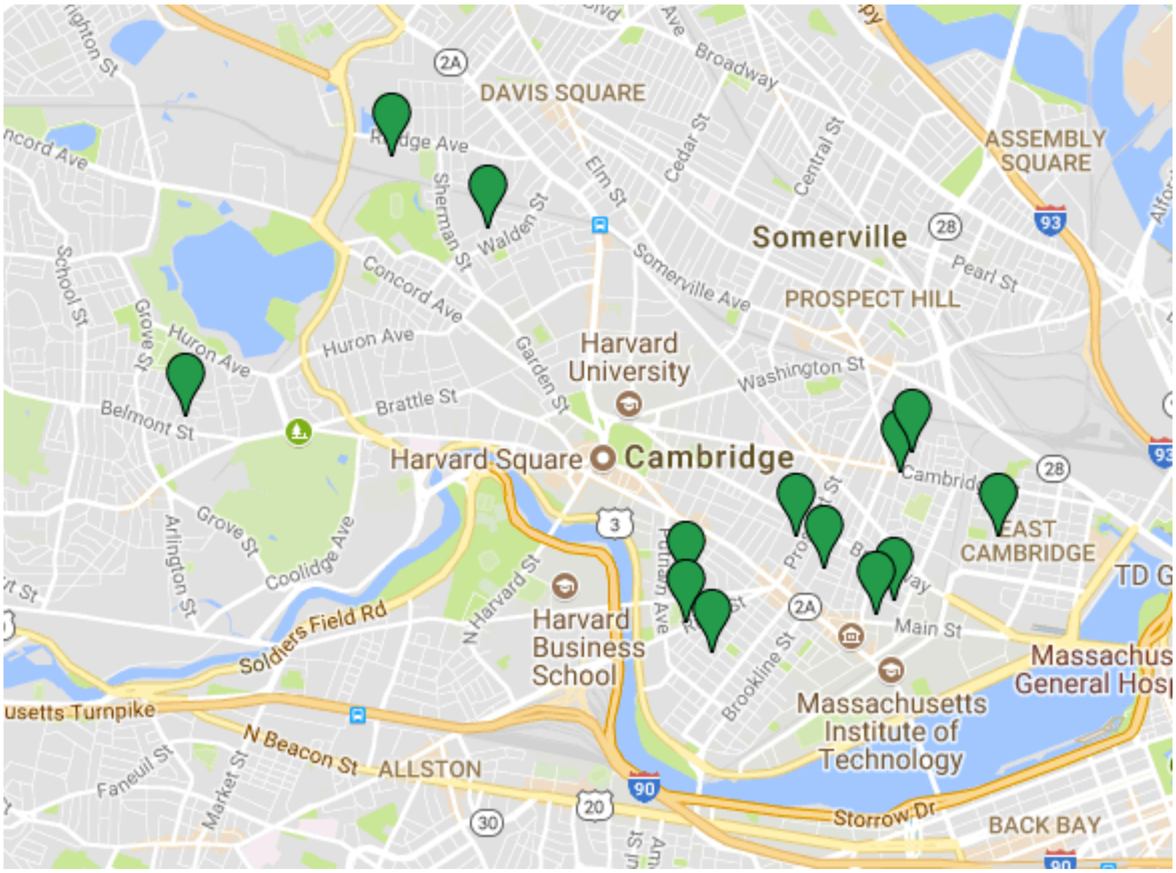

Figure 4: Application Rates by Income

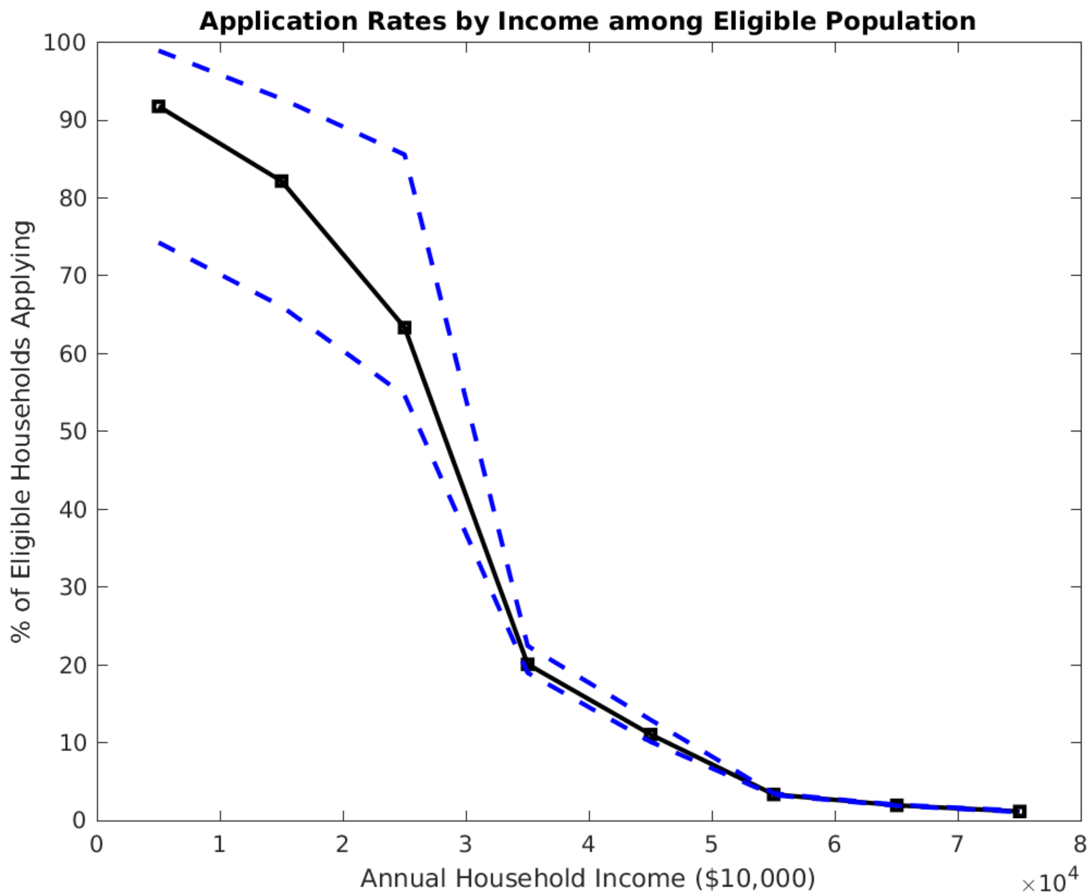

Notes: The estimated fraction of eligible households that applied for Family Public Housing in Cambridge between 2010 and 2014, by $10,000 income groups. For each group, the number of applicants is divided by the number of eligible households as estimated in Section VA. The dotted lines give point-wise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

Figure 5: Welfare Effect of Eliminating Choice under Alternative Modeling Assumptions

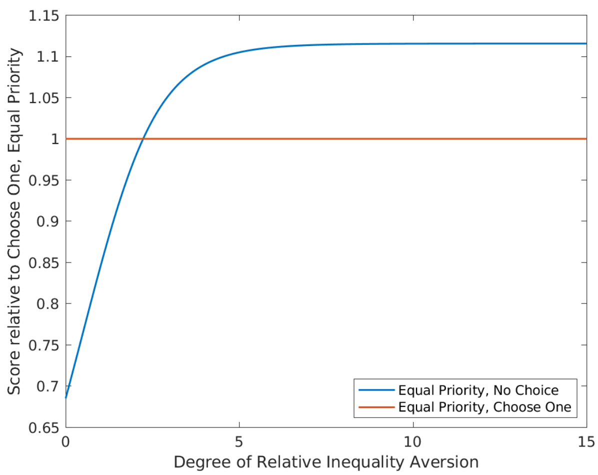

(a) Simple First-Stage Decision Rule

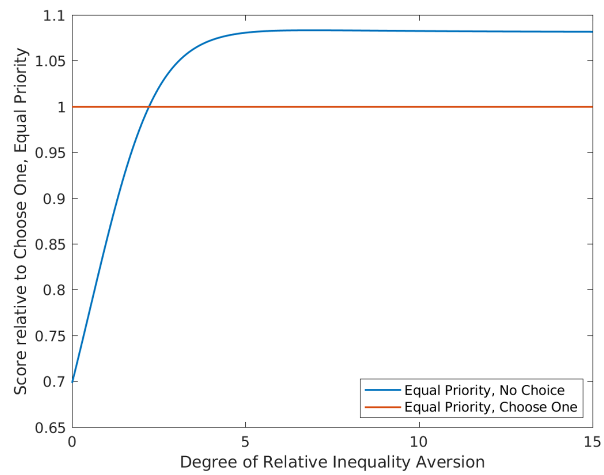

(b) 50% of Applicants Naive

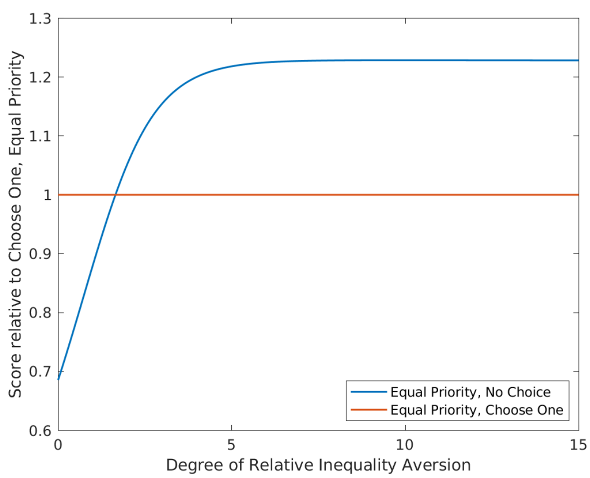

(c) Beliefs Matching Empirical Waiting Time Distributions

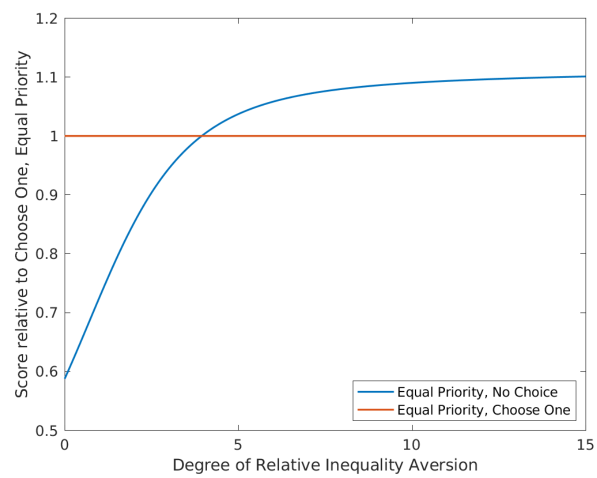

(d) $10,000 Minimum Housing Expenditure