HCEO WORKING PAPER SERIES

Working Paper

HUMAN CAPITAL AND
ECONOMIC OPPORTUNITY
GLOBAL WORKING GROUP

The University of Chicago
1126 E. 59th Street Box 107
Chicago IL 60637

www.hceconomics.org

Neighborhood Effects and Housing Vouchers*

Morris A. Davis

Rutgers University

morris.a.davis@rutgers.edu

Jesse Gregory

University of Wisconsin - Madison

jmgregory@ssc.wisc.edu

Daniel A. Hartley

Federal Reserve Bank of Chicago

Daniel.A.Hartley@chi.frb.org

Kegon T. K. Tan

University of Rochester

ttan8@ur.rochester.edu

December 14, 2019

Abstract

Researchers and policy-makers have explored the possibility of restricting the use of housing vouchers to neighborhoods that may positively affect the outcomes of children. Using the framework of a dynamic model of optimal location choice, we estimate preferences over neighborhoods of likely recipients of housing vouchers in Los Angeles. We combine simulations of the model with estimates of how locations affect adult earnings of children to understand how a voucher policy that restricts neighborhoods in which voucher-recipients may live affects both the location decisions of households and the adult earnings of children. We show the model can replicate the impact of the Moving to Opportunity experiment on the adult wages of children. Simulations suggest a policy that restricts housing vouchers to the top 20% of neighborhoods maximizes expected aggregate adult earnings of children of households offered these vouchers.

JEL Classification Numbers: I240, I31, I38, J13, R23, R38

Keywords: Neighborhood Choice, Housing Vouchers

*We thank numerous discussants and seminar participants for helpful comments. The views expressed herein are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

1 Introduction

We study if housing policy that was enacted to reduce housing costs of low-income households can also affect intergenerational mobility. Specifically, we consider an environment in which policy-makers restrict the use of housing vouchers to a set of neighborhoods that may positively impact the earnings of children once they are adults. We investigate the extent that this change in policy affects both the willingness of low-income households to use housing vouchers and the adult earnings of children of those households.

So, why is this interesting? A large body of evidence suggests that neighborhoods can directly affect many child outcomes, including the income of children once they are adults. An older empirical literature using observational data often finds strong associations between neighborhood quality, broadly defined, and positive child-level outcomes: See [Leventhal and Brooks-Gunn \(2000\)](#), [Durlauf \(2004\)](#) and [Ross \(2011\)](#) for surveys. While the researchers of these studies typically attempt to account for selection issues, the fact that individuals endogenously sort into neighborhoods leaves open the possibility of non-causal explanations for documented patterns.

A recent set of papers using experimental or quasi-experimental evidence also finds strong effects of neighborhoods on child outcomes. [Chyn \(2018\)](#) shows that children of families forced to relocate out of demolished public housing projects in Chicago are more likely to be employed and earn more in young adulthood than peer children of nearby public housing that was not demolished. [Chetty, Hendren, and Katz \(2016\)](#) evaluate the impact of the Moving To Opportunity (MTO) program on adult earnings of children. MTO was an experiment undertaken in the 1990s that randomly assigned a group of households with children eligible to live in low income housing projects in five U.S. cities to three different groups: (i) a treatment group that received a Section 8 housing voucher that in the first year could be applied only in Census tracts with a poverty rate under 10% and could be applied unconditionally thereafter, (ii) a second treatment group that received a comparable Section 8 housing voucher with no location restriction attached, and (iii) a control group that received no voucher. [Chetty, Hendren, and Katz \(2016\)](#), hereafter CHK, show that children under the age of 13 from the group that received the location-restricted voucher experienced a \$3,477 annual increase in adult earnings relative to the control group.

Given this evidence, it may seem reasonable to ask if public policy should steer low-income households away from neighborhoods that might be detrimental to child outcomes and towards neighborhoods that might improve child outcomes. A public policy that achieves this goal may be implemented, in part, by restricting the locations in which housing vouchers may be applied. Low-income households that receive a location-restricted housing voucher

would only be able to use the voucher to pay rent in a pre-determined set of neighborhoods that are expected to improve child outcomes. [Bergman, Chetty, DeLuca, Hendren, Katz, and Palmer \(2019\)](#) are running a large experiment in Seattle in which randomly selected households receive location-restricted housing vouchers. The voucher-eligible locations consist of the top third of Opportunity Atlas Census tracts, which (loosely speaking) are the tracts where child income as an adult is expected to be largest conditional on parental income.¹

The [Bergman, Chetty, DeLuca, Hendren, Katz, and Palmer \(2019\)](#) paper uses an experimental design to understand barriers households face in accepting a location-restricted housing voucher. We take a different, structural approach in addressing similar issues. A brief summary of our paper is as follows: We use panel data from Los Angeles to estimate preferences for locations, consumption, housing and amenities for many different types of renting households in Los Angeles. Given estimated preferences and the structure of our location-choice model, we solve for the steady-state equilibrium of the model under various location-restricted housing voucher policies. These simulations show the extent to which the expected earnings of children of voucher recipients rise, due to their locating in neighborhoods that positively influence their earnings. The simulations also enable us to track the expected earnings of children of households not receiving vouchers. These earnings decline as some households relocate to relatively worse neighborhoods in response to an increase in rental prices in the relatively good locations that occurs as a result of policy. Additionally, we study the extent to which the location restrictions impact various households' willingness to accept a housing voucher and we discuss in some detail the distributional consequences of location-restricted voucher policies. While some of our conclusions are specific to renting households in Los Angeles, our methods can be used to study any area to inform policy design.

In the paper that is closest to ours, [Galiani, Murphy, and Pantano \(2015\)](#) use data on the location choices of the Boston participants in the Moving to Opportunity experiment to help identify the structural parameters of a location-choice model. Their approach exploits the randomization of MTO participants into the control group (no voucher) and Section 8 group (voucher with no location restriction) along with Census data on tract demographics to estimate the preference weights that households place on consumption, housing, amenities and various neighborhood characteristics. The randomization of households into the different MTO treatment and control groups allows the preference weight on consumption and housing to be identified without an instrument for rent. The model successfully matches a number of moments summarizing the location choices of MTO participants offered location-restricted vouchers, providing out-of-sample model validation. The spirit of our paper and [Galiani,](#)

¹We discuss the Opportunity Atlas in great detail later in the paper.

Murphy, and Pantano (2015) are similar, however there are a few key differences. Specifically, we focus on the location decisions of households and the implications for adult earnings of children. Additionally, we model and estimate the choices of all renters in Los Angeles (both voucher recipients and non-recipients), enabling us to study metro-wide implications of large-scale hypothetical changes to housing-voucher policies in a general-equilibrium framework.

The rest of this introduction highlights our methods, details and results.

We start by estimating the parameters of a discrete-choice, dynamic model of location choice for renters in Los Angeles. The model is in the spirit of Kennan and Walker (2011) and Bayer, McMillan, Murphy, and Timmins (2016). We use panel data on renting households from the Federal Reserve Bank of New York Consumer Credit Panel / Equifax (CCP) data to estimate optimized indirect utility for each neighborhood (Census tract) in Los Angeles and the cost of moving. These data are a 5% random sample of U.S. adults conditional on having an active credit file and any individuals residing in the same household. Our estimation sample includes more than 1.75 million person-year observations. We divide the sample into 144 types of households based on observable characteristics in the first period in which we observe the household.

Next, we specify that conditional on a choice of neighborhood, each household has Cobb-Douglas preferences for consumption and housing in that neighborhood. With Cobb-Douglas preferences, the ratio of expenditures on housing to expenditures on consumption is fixed, implying that households rent smaller units in neighborhoods where the rental price-per-unit of housing is high. We find the average expenditure share across types of households is 27%, consistent with the results of Davis and Ortalo-Magne (2011) and others.

Our specification requires estimation of one additional parameter that scales the deterministic portion of utility relative to the variance of utility shocks that are embedded in the dynamic location-choice model. This scale parameter determines how households respond to shocks that affect utility after controlling for consumption, housing and fixed location-specific amenities. To estimate this parameter we use the instrumental variables approach of Bayer, Ferreira, and McMillan (2007).

Finally, we determine the types of households that are eligible to receive a housing voucher and have at least one child. We use tract-level data from the 2000 Census to estimate average income and average number of children per household for each type. We identify 24 types of voucher-eligible households in our sample with children that accept a housing voucher if offered. These households are 1/4 African American and 3/4 Hispanic, have on average 2.1 children per household, an annual income of \$18.7 thousand and spend 36 percent of their income on rents.

In the final sections of the paper, we combine the predictions of the estimated model

with Data from the Opportunity Atlas to study how various housing-voucher policies affect optimal location choices of households and the earnings of children when they become adults. The Opportunity Atlas is a data set created by [Chetty, Friedman, Hendren, Jones, and Porter \(2018\)](#) that, for each Census tract in the United States, predicts the percentile of a child’s adult earnings in the age-26 income distribution given the percentile of the household’s income in the income distribution.² We begin the analysis by asking if our model can replicate the estimate of CHK that the MTO voucher program increased annual adult earnings of children under the age of 13 at the time the voucher was received by \$3,477. We show that the model can nearly exactly replicate this result; our model-based estimate of their statistic is \$3,507.

Interestingly, holding the poverty rate constant of the chosen neighborhood, our simulations show that if MTO voucher recipients had selected neighborhoods *randomly* then expected average adult earnings of children of voucher recipients would have increased by \$6,651, nearly double the estimate of CHK. In other words, we find that MTO voucher recipients selected into neighborhoods that yield relatively low adult earnings for children. This occurs for two reasons. For neighborhoods with a poverty rate less than 10%, households accepting an MTO voucher prefer the amenities of low Opportunity Atlas score neighborhoods to high Opportunity Atlas score neighborhoods. Additionally, rental prices tend to increase with Opportunity Atlas scores across neighborhoods.

In the final part of the paper we simulate our model under a plethora of policy scenarios to understand the extent to which a city-wide voucher program that restricts the neighborhoods in which housing vouchers can be used can increase the adult earnings of children of voucher-eligible households. In all simulations, we allow for rental prices to adjust in equilibrium in response to changes in tract-level housing demand. We consider two sets of simulations. At first we analyze results assuming the Opportunity Atlas score of all neighborhoods is fixed at its estimated value. After that, we allow the Opportunity Atlas score of a neighborhood to adjust based on changes in the racial composition and average income of that neighborhood.

We search for a cutoff Opportunity Atlas score, such that neighborhoods with higher Opportunity Atlas score are included in the set of acceptable locations of restricted-voucher holders, that (a) maximizes the aggregate adult earnings of all children and (b) maximizes aggregate adult earnings of the children in voucher-eligible households. In the analysis, we highlight essential trade-offs of a location-restricted voucher program: Some households decline the voucher because the set of acceptable neighborhoods is too restrictive but households that accept the voucher experience significant gains in the adult earnings of their

²We will sometimes refer to the expected percentile of the child’s adult earnings in the age-26 income distribution as the Opportunity Atlas “score” of the neighborhood.

children. We find that a voucher program that limits locations to the top 10 percent of Opportunity Atlas neighborhoods maximizes the aggregate annual earnings of all children of renting households in Los Angeles; and, a policy that limits location to the top 20 percent of Opportunity Atlas neighborhoods maximizes the aggregate earnings of children of renting households eligible to receive vouchers. In either case, many children of households accepting vouchers experience enormous gains to income and children of other households experience, on average, small losses. On net, the gains outweigh the losses. We conclude policymakers can implement a location-restricted voucher program that yields aggregate gains to adult earnings of children and significantly impacts intergenerational mobility for low-income households eligible to receive housing vouchers.

2 Location Choice Model and Estimates

2.1 Model

The first step in our analysis is to understand how household utility changes with location. To do this, we estimate the parameters of an optimal forward-looking location-choice model. The basic intuition of estimation is as follows: If we notice households moving to certain clusters of neighborhoods more frequently than others, then, on average, those neighborhoods must provide higher levels of utility. In other words, viewed from the lens of the model, probabilities over location choices are directly informative of net utility of locations.

We consider the decision problem of a household head deciding where his or her family should live using a dynamic discrete choice setting. Our basic framework is somewhat standard and similar models have been studied by [Kennan and Walker \(2011\)](#), [Bishop and Murphy \(2011\)](#) and [Bayer, McMillan, Murphy, and Timmins \(2016\)](#). For purposes of exposition, we write down the model describing the optimal decision problem of a single family which enables us to keep notation relatively clean. When we estimate the parameters of this model, we will allow for the existence of many different “types” of people in the data. Each type of person will face the same decision problem, but the vector of parameters that determines payoffs and choice probabilities will be allowed to vary across types of people.

The household can choose to live in one of J locations. Denote j as the household’s current location. We write the value to the household of moving to location ℓ given a current location of j and current value of a shock ϵ_ℓ (to be explained later) as

$$V(\ell \mid j, \epsilon_\ell) = u(\ell \mid j, \epsilon_\ell) + \beta EV(\ell)$$

In the above equation $EV(\ell)$ is the expected future value of having chosen to live in ℓ today and β is the factor by which future utility is discounted. Note that the expected future value of choosing to live in ℓ today does not depend on the value of ϵ_ℓ , as in Rust (1987). We assume the household problem does not change over time, explaining the lack of time subscripts.

u is the flow utility the agent receives today from choosing to live in ℓ given a current location of j and a value for ϵ_ℓ . We assume u is the simple function

$$u(\ell \mid j, \epsilon_\ell) = \delta_\ell - \kappa_{\ell j} + \epsilon_\ell$$

where δ_ℓ is the flow utility the household receives this period from living in neighborhood ℓ , net of rents and other costs. In section 2.4, we parse δ_ℓ into utility from consumption, housing and fixed neighborhood amenities, but for now just know that δ_ℓ has the interpretation of maximized indirect utility. $\kappa_{\ell j} = [\kappa_0 + \kappa_1 * \mathcal{D}_{\ell j}] \cdot 1_{\ell \neq j}$ are all costs (utility and financial) a household pays when it moves to neighborhood ℓ from neighborhood j , which we specify as the sum of a fixed cost κ_0 and a cost that increases at rate κ_1 with distance in miles between the centroid of tracts ℓ and j denoted $\mathcal{D}_{\ell j}$; $1_{\ell \neq j}$ is an indicator function that is equal to 1 if location $\ell \neq j$ and 0 otherwise, i.e. the household pays zero moving costs if it does not move; and ϵ_ℓ is a random shock that is known at the time of the location choice. ϵ_ℓ is assumed to be iid across locations, time and people. The parameters δ_ℓ , κ_0 and κ_1 may vary across households, but for any given household these parameters are assumed fixed over time. ϵ_ℓ induces otherwise identical households living at the same location to optimally choose different future locations. Dynamics in the model driven by moving costs and the ϵ_ℓ shocks. The model would be static if either the idiosyncratic shocks were time-invariant or moving costs were zero.

Denote ϵ_1 as the shock associated with location 1, ϵ_2 as the shock with location 2, and so on. After the vector of ϵ are revealed (one for each location), in each period households choose the location that yields the maximal value

$$V(j \mid \epsilon_1, \epsilon_2, \dots, \epsilon_J) = \max_{\ell \in 1, \dots, J} V(\ell \mid j, \epsilon_\ell) \quad (1)$$

$EV(j)$ is the expected value of (1), where the expectation is taken with respect to the vector of ϵ . We assume each period is one year.

When the ϵ are assumed to be drawn i.i.d. from the Type 1 Extreme Value Distribution,

the expected value function $EV(j)$ has the functional form

$$EV(j) = \log \left\{ \sum_{\ell=1}^J \exp \tilde{V}(\ell | j) \right\} + \zeta \quad (2)$$

where ζ is equal to Euler's constant,

$$\tilde{V}(\ell | j) = \delta_\ell - \kappa_{\ell j} + \beta EV(\ell) \quad (3)$$

and the tilde symbol signifies that the shock ϵ_ℓ has been omitted. Additionally, it can be shown that the log of the probability that location ℓ is chosen given a current location of j , call it $p(\ell | j)$, has the solution

$$p(\ell | j) = \tilde{V}(\ell | j) - \log \left\{ \sum_{\ell'=1}^J \exp [\tilde{V}(\ell' | j)] \right\} \quad (4)$$

Subtract and add $\tilde{V}(k | j)$ to the right-hand side of the above to derive

$$p(\ell | j) = \tilde{V}(\ell | j) - \tilde{V}(k | j) - \log \left\{ \sum_{\ell'=1}^J \exp [\tilde{V}(\ell' | j) - \tilde{V}(k | j)] \right\} \quad (5)$$

One approach to estimating model parameters such as [Rust \(1987\)](#) is to solve for the value functions at a given set of parameters, apply equation (5) directly to generate a likelihood over the observed choice probabilities, and then search for the set of parameters that maximizes the likelihood. This approach is computationally intensive because it requires solving for the value functions at each step of the likelihood, which involves backwards recursions using equations (2) and (3). In cases such as ours, involving many parameters to be estimated, this approach is computationally infeasible.

Instead, we use the approach of [Hotz and Miller \(1993\)](#) and employed by [Bishop \(2012\)](#) in similar work. This approach does not require that we solve for the value functions. Note that equation (3) implies

$$\tilde{V}(\ell | j) - \tilde{V}(k | j) = \delta_\ell - \delta_k - [\kappa_{\ell j} - \kappa_{kj}] + \beta [EV(\ell) - EV(k)] \quad (6)$$

But from equation (2),

$$EV(\ell) - EV(k) = \log \left\{ \sum_{\ell'=1}^J \exp \tilde{V}(\ell' | \ell) \right\} - \log \left\{ \sum_{\ell'=1}^J \exp \tilde{V}(\ell' | k) \right\}$$

Now note that equation (4) implies

$$\begin{aligned} p(k | \ell) &= \tilde{V}(k | \ell) - \log \left\{ \sum_{\ell'=1}^J \exp [\tilde{V}(\ell' | \ell)] \right\} \\ p(k | k) &= \tilde{V}(k | k) - \log \left\{ \sum_{\ell'=1}^K \exp [\tilde{V}(\ell' | k)] \right\} \end{aligned}$$

and thus

$$\log \left\{ \sum_{\ell'=1}^J \exp [\tilde{V}(\ell' | \ell)] \right\} - \log \left\{ \sum_{\ell'=1}^K \exp [\tilde{V}(\ell' | k)] \right\}$$

is equal to

$$\begin{aligned} &\tilde{V}(k | \ell) - \tilde{V}(k | k) - [p(k | \ell) - p(k | k)] \\ = &\quad -\kappa_{k\ell} \quad - [p(k | \ell) - p(k | k)] \end{aligned}$$

The last line is quickly derived from equation (3). Therefore,

$$EV(\ell) - EV(k) = -[p(k | \ell) - p(k | k) + \kappa_{k\ell}]$$

and equation (6) has the expression

$$\begin{aligned} &\tilde{V}(\ell | j) - \tilde{V}(k | j) \\ &= \delta_\ell - \delta_k - [\kappa_{\ell j} - \kappa_{kj}] - \beta [p(k | \ell) - p(k | k) + \kappa_{k\ell}] \end{aligned} \tag{7}$$

Combined, equations (5) and (7) show that the log probabilities that choices are observed are simple functions of model parameters $\delta_1, \dots, \delta_J$, κ_0 , κ_1 and β and of observed choice probabilities. In other words, a likelihood over choice probabilities observed in data can be generated without solving for value functions. Our estimation approach also relies on the fact that the expected value of choosing any neighborhood in the next period does not change over time. In other words, decisions today do not affect future expected values (net of moving costs). This allows us to estimate the model with a short panel, an insight from [Arcidiacono and Miller \(2011\)](#).

2.2 Data and Likelihood

We estimate the model using panel data from the Federal Reserve Bank of New York Consumer Credit Panel / Equifax (CCP). The panel is comprised of a 5% random sample of U.S. adults with a social security number, conditional on having an active credit file, and any individuals residing in the same household as an individual from that initial 5% sample.³ For years 1999 to the present, the database provides a quarterly record of variables related to debt: Mortgage and consumer loan balances, payments and delinquencies and some other variables we discuss later. The data does not contain information on race, education, or number of children and it does not contain information on income or assets although it does include the Equifax Risk ScoreTM which provides some information on the financial wherewithal of the household as demonstrated in [Board of Governors of the Federal Reserve System \(2007\)](#). Most important for our application, the panel data includes in each period the current Census block of residence. To match the annual frequency of our location choice model, we use location data from the first quarter of each calendar year. Other authors have used the CCP data to study the relationship of interest rates, house prices and credit (see [Bhutta and Keys \(2016\)](#) and [Brown, Stein, and Zafar \(2015\)](#)) and the impact of natural disasters on household finances ([Gallagher and Hartley, 2017](#)), but we are the first to use this data to estimate an optimal location-choice model.

We restrict our sample to individuals who (a) lived in Los Angeles County in the first quarter of any year from 1999 through 2013, (b) were observed in Los Angeles in the first quarter of the following year, and (c) never had a home mortgage, yielding 1,787,558 person-year observations. An advantage of the size of our data is that we can estimate a full set of model parameters for many “types” of people, where we define a type of person based on observable demographic and economic characteristics. We study renters to mitigate any problems of changing credit conditions and availability of mortgages during the sample window.⁴ We exclude from our estimation Census tracts with fewer than 150 rental units and tracts that are sparsely populated in the northern part of the county.⁵ The panel is not balanced, as some individuals’ credit records first become active after 1999.

³The data include all individuals with 5 out of the 100 possible terminal 2-digit social security number (SSN) combinations. While the leading SSN digits are based on the birth year/location, the terminal SSN digits are essentially randomly assigned. A SSN is required to be included in the data and we do not capture the experiences of illegal immigrants. Note that a SSN is also required to receive a housing voucher.

⁴In the CCP data, renters and homeowners without a mortgage are observationally equivalent. According to data from the 2000 Census, 85% percent of the units without a home mortgage are renter-occupied for the 1,748 Census tracts of our study. Since we drop households that ever had a mortgage and follow most households over multiple years, this implies the upper bound for the percentage of homeowners in our sample is 15%.

⁵On average, each Census tract in Los Angeles has about 4,000 people.

Table 1: Comparison of Equifax and Census Data

Poverty Rate (%)	Avg. Population 2000-2010		Equifax	Pct. w/ Mortgage 2008-2012	
	Equifax ^a	Census ^b	Share	Equifax ^c	ACS ^d
(1)	(2)	(3)	(4)	(5)	(6)
0-5	610,336	654,004	93.3%	61.6%	62.6%
5-10	1,395,831	1,478,114	94.4%	50.0%	50.2%
10-15	1,033,076	1,135,194	91.0%	40.5%	39.2%
15-20	751,098	870,869	86.2%	37.3%	34.9%
20-25	630,830	761,841	82.8%	30.7%	26.9%
>25	1,085,466	1,497,545	72.5%	23.9%	19.0%
Public Housing ^e	24,988	31,400	79.6%	19.1%	16.5%

Notes: This table compares population in the Census (column 3) and ACS (column 6) with the implied equivalent population in the Equifax data (columns 2 and 5). Column (4) is the share of the Census population accounted for by the Equifax data, computed as column (2) divided by column (3).

a Data are computed as 20 times the average (1999-2014) number of Equifax primary individuals ages 18-64.

b Data shown are the average (2000 and 2010) of the Census tract population ages 18-64.

c Data are the average share of households in Equifax with a mortgage, 2008-2012.

d Data are the average share of households in the American Community Survey tract-level tabulations with a mortgage, 2008-2012.

e Data shown are for 15 tracts with large public housing developments (250+ occupied, non-senior public housing units in 2000).

Table 1 compares sample statistics from the CCP data to Census data for the tracts in Los Angeles County. This table includes data for both owners and renters. Column (2) shows the implied total population of adults ages 18-64 in the CCP data, computed as twenty times the total number of primary individuals, and (3) shows the average population counts of adults from the 2000 and 2010 Census. The table shows that coverage in the low poverty tracts is very high, above 90%. Coverage remains high but falls for the higher-poverty tracts, either because many individuals lack credit history or do not have a social security number. Columns (5) and (6) compare the percentage of households with a mortgage in the two data sets. Not surprisingly, the percentages fall quite dramatically with the poverty rate, and generally speaking the percentages reported in the two data sets are close. The final row of Table 1 compares the CCP and Census data for 15 tracts containing large public housing developments, the residents of which will be the focus of some of our analysis later on.⁶ That row shows the two data sets closely align for these tracts.

We stratify households into types using an 8-step stratification procedure. We begin with

⁶We define large developments as those with at least 250 occupied, non-senior public housing units in 2000. We also include the Census tracts containing Avalon Gardens and Hacienda Village which are below the 250 unit threshold but are proximate to several large developments.

the full sample, and subdivide the sample into smaller “cells” based on (in this order): The racial plurality, as measured by the 2000 Census, of the 2000 Census block of residence (4 bins),⁷ 5 age categories (cutoffs at 30, 45, 55, and 65),⁸ number of adults age 18 and older in the household (1, 2, 3, 4+), and then the presence of an auto loan, credit card, student loan and consumer finance loan. We do not subdivide cells in cases where doing so would result in at least one new smaller cell with fewer than 20,000 observations. In a final step applied to all bins, we split each bin into three equally-populated types based on within-bin credit-score terciles. After all the dust settles, this procedure yields 144 types of households.

The following figures from our data are instructive. The top panel of Figure 1 shows the typical location choices made by type 133 in our sample: A household earning \$12,000 per year with an Equifax Risk ScoreTM below 580 and first observed living in a Census block that is predominantly African American.⁹ The light blue areas show all Census tracts with poverty rates less than 10% and the tan areas show all Census tracts with higher poverty rates. The areas in dark blue show the most chosen low-poverty Census tracts for this type and the areas in black show the most chosen high-poverty tracts. Panel (a) shows this type predominantly clusters its location choices in one crescent-shaped area in the south-central part of the county. The bottom panel of this figure shows the same set of location choices for type 28 in our sample, a household earning \$12,000 per year with an Equifax Risk ScoreTM below 600 and first observed in a predominantly hispanic Census block. Comparing the top and bottom panels, not many neighborhood choices overlap between the two types. If, counterfactually, we assumed that the vector of δ_j of the two types were the same, the model would attribute the systematic variation in optimal neighborhood choices entirely to differences in the i.i.d. utility shocks.

Households in our sample can choose to locate in one of 1,748 Census tracts in Los Angeles. Allowing a separate value of δ for each tract and for each type would require estimating more than 250,000 parameters. Conceptually, with a large enough sample we could separately estimate every δ for each type. Currently, we have data on approximately 2,000 households followed over 10 years for each type of household in our sample. For parsimony, and to exploit the fact that geographically nearby tracts likely provide similar utility, for each type we specify that the utility of location j , δ_j , is a function of latitude

⁷We assign race based on the racial plurality of all persons in the Census block, owners and renters, where they are first observed. The mean number of households and residents at the Census-block level in our sample of 1,748 tracts is 41 and 118, respectively.

⁸This refers to the age of the person in the household in the initial random sample.

⁹We discuss later how we generate the estimate of household income.

Figure 1: Location Choices by Type for Tracts Below and Above 10% Poverty Rate

(a) Type 133: African American households earning \$12,000 per year w/ <580 Equifax Risk ScoreTM

(b) Type 28: Hispanic households earning \$12,000 per year w/ <600 Equifax Risk ScoreTM

Notes: These graphs show the most frequent location choices for type 133 (top panel), an African American household with annual income of \$12 thousand per year and a Risk Score less than 580 and type 28 (bottom panel), a Hispanic household with annual income of \$12 thousand per year and a Equifax Risk ScoreTM less than 600. The light blue areas of the map indicate tracts with a poverty rate below 10% and the light brown areas are tracts with a poverty rate above 10%. The dark blue and black areas show the most frequently chosen tracts for each type. Dark blue are tracts with less than 10% poverty and black are tracts with greater than 10% poverty.

(lat_j) and longitude (lon_j) of that location according to the formula

$$\delta_j = \sum_{k=1}^K a_k B_k(lat_j, lon_j)$$

The B_k are parameter-less basis functions.¹⁰ For each type, we use $K = 89$ basis functions. Additionally, we allow the values of a_k to vary for tracts above and below the 10% poverty threshold. Inclusive of the two moving cost parameters, we estimate $2 \times 89 + 2 = 180$ parameters per type. With 144 types, we estimate a total of 25,920 parameters.

To define the log likelihood that we maximize we need to introduce more notation. Let i denote a given household, t a given year in the sample, j_{it} as person i 's starting location in year t and ℓ_{it} as person i 's observed choice of location in year t . Denote τ as type and the vector of parameters to be estimated for each type as θ_τ . The log likelihood of the sample is

$$\sum_{\tau} \sum_{i \in \tau} \sum_t p(\ell_{it} \mid j_{it}; \theta_\tau) \quad (8)$$

$p(\cdot)$ is the model predicted log-probability of choosing ℓ_{it} given j_{it} . For each τ we use the quasi-Newton BFGS procedure to find the vector θ_τ that maximizes the sample log likelihood.

Before moving on, note that the model assumes that all households have the ability to live in any neighborhood. Of course, some landlords may be racist or discriminate against households with low income¹¹ but households in our model do not need to be able to rent from every landlord in every location; they only need to be able to rent one unit of their desired size and quality in each Census tract. In the event racism or discrimination is systemic in certain tracts, the probability that certain types of households will live in those tracts will be low and this will affect estimates of δ for those types in those tracts.¹² If discrimination in certain tracks is significant, we conjecture our framework will be still be useful in predicting location choice for those tracts – and the policy experiments we discuss later will continue to be informative – as long as the degree to which landlords are discriminatory does not systematically change as a result of any policies we consider.

¹⁰Note that even though two adjacent tracts may have similar values of δ (possibly due to smoothing), the shocks for each tract do not have to be similar and there may be a discrete jump at the tract border in the value of the shock.

¹¹For evidence on discrimination in rental markets, see studies by [Yinger \(1986\)](#) and [Ewens, Tomlin, and Wang \(2014\)](#). [Popkin, Cunningham, Godfrey, Bednarz, and Lewis \(2002\)](#) and [Phillips \(2017\)](#) also demonstrate that landlords discriminate against rental applicants that wish to use housing vouchers.

¹²This may also affect estimates of κ , depending on the relative location of the tracts with these racial issues.

2.3 Estimates and Model Fit

Our procedure ultimately yields estimates of δ_j , κ_0 and κ_1 for each type to match model-predicted moving probabilities to those in the data.¹³ The top and bottom panels of Figure 2 show the surface of indirect utilities across Los Angeles County that we estimate for types 133 and 28, respectively, such that the model can replicate as best as possible the location choices shown in Figure 1. These figures illustrate the flexibility of our specification. The surfaces are quite different, reflecting the very different optimal location choices of these two types.

Due to our large number of types and tracts, it is impossible to report all parameter estimates. Instead, we summarize the estimates by examining the model’s in-sample fit along a number of dimensions. By design our model can nearly exactly match the average moving rate in the data for each type; a regression of the model-predicted average moving rate on the moving rate in the data for our 144 types has an R2 value of 0.9996. Figure 3 compares the distribution of distances moved (measured as the straight line distance between tract centroids) for all movers in the data and as predicted by our model.¹⁴ This figure shows that the model replicates the hump-shaped distribution of distances moved, with the most frequent moves around 4 miles. The model slightly overpredicts moves between 4 and 10 miles in length and slightly underpredicts moves less than 4 miles.

Figure 4 shows a detailed comparison of model-predicted and actual annual migration rates for households that choose to move by poverty rate of Census tracts. The tracts from which people are moving are split into six groupings based on the poverty rate of the originating tract: 0-5, 5-10, 10-15, 15-20, 20-25 and >25. For each of these groupings, the probability of choosing a destination tract of a given poverty rate is plotted for the data (dark blue solid line) and as predicted by the model (light blue dotted line). Figure 4 shows model fit for some very low-probability moves.¹⁵ The model tends to underpredict the probability that households living in low-poverty tracts move to a low-poverty tract, conditional on a move occurring. Aside from that, in our view the model fits the data well.

2.4 Preferences for Amenities, Housing and Consumption

We specify that δ_ℓ is the following function of consumption enjoyed in tract ℓ (c_ℓ), the quantity of housing rented in tract ℓ (h_ℓ), and type-specific amenities associated with location

¹³We fix $\beta = 0.95$ for all types.

¹⁴In the data we know the Census block of residence for each household. We eliminate any within-tract moves and for the remaining moves, we define distance moved as the distance between tract centroids of the sending and receiving tracts.

¹⁵For perspective, the unconditional probability of any move is less than ten percent.

Figure 2: Surfaces of Indirect Utility

(a) Type 133

(b) Type 28

Notes: These graphs show the maximum likelihood estimates of δ_j for all tracts in Los Angeles for household types 133 and 28.

Figure 3: Model Fit: Density of Moving Distance

Notes: This graph shows the predicted and actual density of distance moved in our data, conditional on a move to a different Census tract.

ℓ that are fixed over time (A_ℓ)

$$\delta_\ell = \left(\frac{1}{\sigma_\epsilon} \right) \ln A_\ell + \left(\frac{1-\alpha}{\sigma_\epsilon} \right) \ln c_\ell + \left(\frac{\alpha}{\sigma_\epsilon} \right) \ln h_\ell$$

σ_ϵ is a parameter that effectively rescales the variance of the draws of the ϵ utility shocks and α is a parameter that determines preferences for housing relative to consumption. We will specify that σ_ϵ is identical for all households but will allow α to vary across types of households. For the purposes of exposition, we temporarily suppress all type-specific subscripts.

We assume the renting households in our sample have no savings. Households choosing to live in location ℓ have the following budget constraint

$$w = c_\ell + r_\ell \cdot h_\ell$$

where w is type-specific income and r_ℓ is the quality-adjusted price-per-unit of housing in location ℓ . Given preferences and constraints, households choosing location ℓ choose optimal

Figure 4: Poverty Category Transitions $t-1$ to t , Conditional on Moving

Notes: These graphs show the predicted and actual density of moves in our data. Panels show outcomes for different poverty rates at the initial location. The x-axis of each panel is the poverty rate at the destination location and the y-axis is the frequency that poverty rate is chosen.

consumption and housing to satisfy

$$\begin{aligned} c_\ell &= (1 - \alpha) w \\ r_\ell \cdot h_\ell &= \alpha w \end{aligned}$$

Since income is fixed for each type in our sample, consumption is independent of location choice. The quantity of housing consumed varies across locations in a deterministic way determined by the rental price per unit of housing. As indicated by the first-order conditions, as households move from (say) cheaper to more expensive locations, their optimal quantity of housing falls such that the expenditure share on housing stays constant.

2.5 Indirect Utility

Continuing to suppress type-specific subscripts, when households do not receive a voucher we can derive optimized indirect utility in location ℓ by inserting the first-order conditions for consumption and housing into utility

$$\delta_\ell = \left(\frac{1}{\sigma_\epsilon} \right) \ln A_\ell + \left(\frac{1 - \alpha}{\sigma_\epsilon} \right) \ln [(1 - \alpha) w] + \left(\frac{\alpha}{\sigma_\epsilon} \right) \ln (\alpha w / r_\ell) \quad (9)$$

Households receiving a voucher rent a unit with market rent equal to the *payment standard* of an area, an amount we denote as *std*. Given market rent per unit of housing in tract ℓ is r_ℓ , the amount of housing voucher-receiving households rent in tract ℓ is std / r_ℓ . Government agencies set the dollar amount of the voucher households receive such that households contribute 30% of their income to rent, leaving 70% of income for consumption. This implies the following quantities of consumption and housing expenditures for voucher-receiving households choosing to live in location ℓ :

$$\begin{aligned} c_\ell &= 0.7w \\ r_\ell \cdot h_\ell &= std \end{aligned}$$

and indirect utility for voucher recipients living in location ℓ , call it $\delta_{\ell,\nu}$, is

$$\delta_{\ell,\nu} = \left(\frac{1}{\sigma_\epsilon} \right) \ln A_\ell + \left(\frac{1 - \alpha}{\sigma_\epsilon} \right) \ln [0.7w] + \left(\frac{\alpha}{\sigma_\epsilon} \right) \ln (std / r_\ell) \quad (10)$$

The difference in utility of living in tract ℓ with and without a voucher is

$$\delta_{\ell,\nu} - \delta_\ell = \left(\frac{1-\alpha}{\sigma_\epsilon} \right) \ln \left(\frac{0.7}{1-\alpha} \right) + \left(\frac{\alpha}{\sigma_\epsilon} \right) \ln \left(\frac{std}{\alpha w} \right) \quad (11)$$

Notice that conditional on α and w , the above expression does not vary across tracts. This implies the probability that any particular tract is chosen does not depend on whether or not the household is receiving a voucher, and thus our likelihood calculations are not affected by the presence of voucher recipients.

2.5.1 Estimating Type-Specific w and α

We wish to estimate w and α for each type, but the CCP data does not contain data on income or rental expenditures. Instead, we estimate w and α by type using data from the Census. Starting with w , for any given Census tract ℓ we compute the average income of renters in the tract in the 2000 Census, call it $\bar{w}_\ell$. We restrict our sample of tracts to tracts with at least 250 rental units, 1,642 tracts in total. Denote the share of type τ renters in tract ℓ according to the CCP data in the year 2000 as η_ℓ^τ . Ideally, we could estimate type-specific income, w^τ , by regressing average tract income $\bar{w}_\ell$ on the full set of type shares in each tract η_ℓ^τ , that is

$$\bar{w}_\ell = \sum_{\tau} w^\tau \eta_\ell^\tau + \text{error}_\ell^w$$

where by construction $\sum_{\tau} \eta_\ell^\tau = 1$. That said, we wish to enforce that estimates of annual income are at least $\underline{w} = \$12,000$ for every type. To do this, we run the regression

$$\bar{w}_\ell - \underline{w} = \sum_{\tau} (w^\tau - \underline{w}) \eta_\ell^\tau + \text{error}_\ell^w$$

and impose that $w^\tau \geq \underline{w}$ in estimation. We estimate that 13% of our types (19 of 144 types) have income at our lower bound of \$12 thousand per year. The average income of the other 125 types is \$47 thousand per year, with a standard deviation of \$31 thousand. The largest type-specific income we estimate is \$173 thousand per year.

Our next step is to estimate a value of α for each type. Denote our estimates of w^τ as $\hat{w}^\tau$ and denote the average level of rental expenditures (paid by renters) measured in the 2000 Census in tract ℓ as $\overline{r\bar{h}}_\ell$. The first-order condition of households implies

$$\overline{r\bar{h}}_\ell = \sum_{\tau} \alpha^\tau \hat{w}^\tau \eta_\ell^\tau \quad (12)$$

where α^τ is the type-specific expenditure share on rents. We transform this equation so regressions do not place disproportionate weight fitting tracts with relatively high average rents. Define predicted average income in tract ℓ as

$$\hat{w}_\ell = \sum_\tau \hat{w}^\tau \eta_\ell^\tau$$

Divide equation (12) by $\hat{w}_\ell$ and substitute our estimate of annual income $\hat{w}^\tau$ for w^τ to yield

$$\frac{\overline{r h}_\ell}{\hat{w}_\ell} = \sum_\tau \alpha^\tau \left(\frac{\hat{w}^\tau \eta_\ell^\tau}{\hat{w}_\ell} \right)$$

We run a regression of the form

$$\frac{\frac{\overline{r h}_\ell}{\hat{w}_\ell} - \underline{\alpha}_\tau}{\overline{\alpha}_\tau - \underline{\alpha}_\tau} = \sum_\tau \left[\frac{\alpha^\tau - \underline{\alpha}_\tau}{\overline{\alpha}_\tau - \underline{\alpha}_\tau} \right] \left(\frac{\hat{w}^\tau \eta_\ell^\tau}{\hat{w}_\ell} \right) + \text{error}_\ell^{rh} \quad (13)$$

This enables us to easily enforce in estimation that $\underline{\alpha}_\tau \leq \alpha_\tau \leq \overline{\alpha}_\tau$.¹⁶ We set $\underline{\alpha}_\tau = 0.1$ and $\overline{\alpha}_\tau = 0.7$. We find that that 3 types (2%) have an expenditure share of exactly 10 percent and 4 types have an expenditure share of exactly 70 percent. For the types with α^τ strictly between these bounds, we estimate the average value of α^τ is 27% with a standard deviation of 12.2%.

Denote our estimates of α^τ as $\hat{\alpha}^\tau$. Figure 5 presents a scatterplot of $\hat{\alpha}^\tau$ and $\hat{w}^\tau$. The small dots show the 144 type estimates and the larger diamonds show mean estimates of α^τ when we group $\hat{w}^\tau$ into 10 bins, one for each income decile, and compute average values in each bin for both $\hat{w}^\tau$ and $\hat{\alpha}^\tau$. Although the individual type data vary somewhat, on average the expenditure share on rent falls with income.

2.5.2 Estimating σ_ϵ

For convenience, rewrite the optimized indirect utility for each type of non-voucher households as

$$\begin{aligned} \delta_\ell &= \left(\frac{1}{\sigma_\epsilon} \right) \ln A_\ell + \left(\frac{1 - \alpha}{\sigma_\epsilon} \right) \ln [(1 - \alpha) w] + \left(\frac{\alpha}{\sigma_\epsilon} \right) \ln (\alpha w) - \left(\frac{\alpha}{\sigma_\epsilon} \right) \ln r_\ell \\ &= \text{const} + \left(\frac{1}{\sigma_\epsilon} \right) \ln A_\ell - \left(\frac{\alpha}{\sigma_\epsilon} \right) \ln r_\ell \end{aligned}$$

¹⁶When $\underline{\alpha}_\tau \leq \alpha_\tau \leq \overline{\alpha}_\tau$, the term in brackets is always between 0 and 1.

Figure 5: Estimates of α^τ and w^τ

Notes: This figure shows type-specific estimates of α^τ , the small blue dots, and the average estimated value of α^τ for each decile of predicted income, the larger red diamonds.

where *const* is a type-specific constant (and we have otherwise temporarily omitted type-specific notation). Assume that log amenities include both observed $\mathcal{O}_\ell$ and unobserved ξ_ℓ characteristics of tract ℓ such that the above can be rewritten as

$$\delta_\ell = \lambda \cdot \mathcal{O}_\ell - \left(\frac{1}{\sigma_\epsilon} \right) \cdot \alpha \ln r_\ell + \xi_\ell$$

The coefficient on α times log rent, $1/\sigma_\epsilon$, cannot be estimated using OLS because equilibrium rents will almost certainly be correlated with unobserved but valued characteristics of neighborhoods, ξ_ℓ . An instrument is required.

Given type-specific estimates of α from section 2.5.1, we use a three-step IV approach to estimate $1/\sigma_\epsilon$ that is similar to the procedure in [Bayer, Ferreira, and McMillan \(2007\)](#). As mentioned earlier, we impose in estimation that $1/\sigma_\epsilon$ is the same for all types. This means that after explicitly accounting for variation in how much people value housing relative to consumption and amenities, and abstracting from differences across types in moving costs, we impose that the importance of utility shocks in household decision making is constant across types. In the first step of our procedure, we estimate $1/\sigma_\epsilon$ using two-stage least squares. We include characteristics of the housing stock 0-5 miles from tract j in $\mathcal{O}_j$ as controls (number of rooms, number of units in the housing structure and age of structure)

and use characteristics of the housing stock 5-20 miles from the tract as instruments for rent. The first-stage F-statistic is 5.35: For more details, see the Appendix.

In the second step, we use estimates of $1/\sigma_\epsilon$ and type-specific estimates λ from the first step, call them $\widehat{\frac{1}{\sigma_\epsilon}}$ and $\widehat{\lambda}^\tau$, to construct predicted indirect utilities for each type that controls for unobserved amenities as

$$\widehat{\delta}_\ell^\tau = \widehat{\lambda}^\tau \cdot \mathcal{O}_j - \left(\widehat{\frac{1}{\sigma_\epsilon}} \right) \alpha^\tau \ln r_\ell$$

We simulate the model using this specification for indirect utility and adjust r_ℓ for all ℓ tracts until the simulated total housing demand in any tract is equal to the observed housing demand in the estimation sample for that tract.¹⁷ This procedure determines market-clearing rents in all tracts in the absence of unobserved amenities. We use these rents as instruments to estimate $1/\sigma_\epsilon$ in the third and final step with an F-statistic of 31.7. Intuitively, the F-statistic rises from 5 to 32 because the first step only uses information about the quality of substitutes for each tract individually whereas the third step uses similar information for all tracts. We estimate that $1/\sigma_\epsilon = 0.84$ with a standard error of 0.198.

2.6 Estimating Voucher-Eligible Households with Children

2.6.1 Number of Children

For our analysis of the MTO experiment and our alternative policy simulations, we wish to track the outcomes of households with children that are offered housing vouchers. This means we need estimates of which types have children, and for the types with children the number of children per household. In the 2000 Census, we know the average number of children by tract for *all* households, not just renting households. To estimate the average number of children by type for our sample of renting households, we invent a new type called “owner-occupiers.” We then run the regression

$$\bar{k}_\ell = \sum_{\tau} k^\tau \tilde{\eta}_\ell^\tau + \text{error}_\ell^k$$

where k^τ is the average number of children per household for type τ households and $\bar{k}_\ell$ is the average number of children per household in tract ℓ . $\tilde{\eta}_\ell^\tau$ is the percentage of type τ households in tract ℓ (which, relative to η_ℓ^τ , explicitly accounts for the fact that there is an additional type, homeowners). As before, $\sum_{\tau} \tilde{\eta}_\ell^\tau = 1$.

¹⁷Given Cobb-Douglas preferences, type-specific housing demand in tract ℓ is $\alpha^\tau w^\tau / r_\ell$.

Table 2: Differences between types with 0 and 3 children

	$k^\tau < 0.5$ (80 types)	$k^\tau > 2.9$ (17 types)
Average value of $\hat{w}^\tau$	\$47,321	\$34,883
Average value of $\hat{\alpha}^\tau$	0.286	0.309
Average Risk Score	686	615
African American	11.3%	5.9%
Hispanic	12.5%	94.1%
Other	13.8%	
White	62.5%	

Notes: This table shows various economic and demographic characteristics of the 80 types of households that we estimate have zero children and the 17 types of households that we estimate have 3 children.

To limit the influence that owner-occupiers have on our estimates of k^τ for renters, we restrict the estimation sample to tracts where at least 50% of the households rent.¹⁸ This restricts our sample to 1,052 tracts (from 1,642 tracts) with 250 or more renting households. We do not restrict k^τ to be an integer but we impose in estimation that $0 \leq k^\tau \leq 3$ for all types.

After discarding the owner-occupier type, we estimate that 80 types in our sample (56 percent of types) have less than 0.5 children on average and 17 types (12 percent) have more than 2.9 children. Table 2 shows how income, rental expenditure share, race and credit score vary across the types that have less than 0.5 children and the types that have 3 children.¹⁹ The types with 0 children have higher income (\$47 thousand as compared to \$35 thousand) and on average a better credit score than the types with 3 children. Additionally, the majority of renting households with 0 children are White and almost all the renting households with 3 children (16 of 17 types) are Hispanic.

2.6.2 Voucher Eligibility

In 2000, 2-person households with annual income less than \$25,020 and 3-person households with annual income less than \$28,140 were eligible to receive a housing voucher in Los Angeles County. Given these rules, we estimate 59 types out of 144 were eligible for a

¹⁸We experimented with setting the rental-share cutoff in 10 percentile increments, from 10% to 90%. Type-specific estimates seemed to stabilize at around a rental-share cutoff of 50%. Additionally, this cutoff minimized the number of types of households with k^τ exactly equal to 3, the upper bound on the number of children that we impose in estimation.

¹⁹When we set the cutoff for zero children to $k^\tau < 0.1$, the results are essentially identical as with our current cutoff of $k^\tau < 0.5$. For example, the number of types with $k^\tau < 0.1$ is 73, as compared to 80 types with $k^\tau < 0.5$.

voucher, 41 percent of households.²⁰ This implies that of the 1.634 million renting households in Los Angeles, 670 thousand households were eligible for a housing voucher in 2000. Only 62,487 households received a voucher, 9.3% of those eligible. Of the 59 types eligible to receive a voucher, 31 types have estimated income less than \$25,020 and 0 children ($k^\tau < 0.5$) and 28 types have estimated income less than \$28,140 and have at least one child ($k^\tau \geq 0.5$). Our estimate that 47% (28/59) of voucher-eligible households have at least one child in Los Angeles in 2000 is very close to the actual percentage of voucher households with children in Los Angeles that can be computed directly from public-housing-agency data, 52.8%.²¹

We use equation (11) to check if any of the 28 types of households with at least 0.5 children that are eligible for a housing voucher would choose to decline a housing voucher if offered. If the value of $\delta_{\ell,\nu} - \delta_\ell$ is less than zero for any type of household, that household would reject a voucher. Equation (11) shows that households with a low value of α and/or a high income relative to the payment standard should reject the voucher. For households with a low value of α , the voucher program forces them to forego some consumption, which they value greatly, in lieu of more housing which they do not value as much. We use the estimate of $\delta_{\ell,\nu} - \delta_\ell$ to determine the value of the voucher to households in terms of equivalent extra annual income. Specifically, we set the estimate of equivalent extra annual income from the voucher equal to annual income multiplied by $\exp[\sigma_\epsilon(\delta_{\ell,\nu} - \delta_\ell)] - 1$, which can be derived from equations (9) and (11).

We use the payment standard of \$9,192 per year (\$766 per month) as set by the U.S. Department of Housing and Urban Development (HUD) for a 2-bedroom apartment in Los Angeles in 2000. We find that 4 out of 28 types of households should reject the offer of a housing voucher. For these four types, accepting the voucher offer would be equivalent to losing \$367 per year of income, on average, as shown in Table 3. These four types are all Hispanic, have an average of 1.3 children, have an average income of \$27 thousand per year, and have an average expenditure share on rent of 18 percent.

Figure 6 shows the income-equivalent gain from the voucher (dots) and the housing subsidy from the voucher (solid line) for all 28 types. The figure also includes a dashed line at zero to highlight the four types that optimally do not take up the voucher. For the 24 types that accept the voucher, on average the voucher is equivalent to an increase in annual income of \$2,922, shown in Table 3. The average rent subsidy, defined as the payment

²⁰The average expenditure share on rents of these 59 types is 37.4 percent.

²¹The Housing Authorities of the City of Los Angeles, Los Angeles County and the City of Long Beach report 40,344, 16,583 and 5,372 voucher units, respectively. This total of 62,299 represents almost all of the 62,487 vouchers in the county. The share of voucher units with children is 52%, 54%, and 55% in the City of Los Angeles, the County of Los Angeles and City of Long Beach, respectively. The voucher-weighted average is 52.8% of voucher units have children, 32,993 units.

Table 3: Differences between voucher-eligible types with at least 0.5 children

	reject voucher (4 types)	accept voucher (24 types)
Average number of children	1.32	2.11
Average annual gain from accepting voucher	-\$367	\$2,922
Average value of $\hat{w}^\tau$	\$26,822	\$18,737
Average value of $\hat{\alpha}^\tau$	0.181	0.358
Average Risk Score	581	605
African American		25.0%
Hispanic	100.0%	75.0%

Notes: This table shows various economic and demographic characteristics of the 4 types of households that reject a housing voucher if offered and the 24 types of households that accept a voucher.

standard less 30 percent of household income, for these 24 types is equal to \$3,571 suggesting a “bang-for-the-buck” (dollars of income-equivalent utility per dollar of subsidy) of 82% on average. As Figure 6 shows, for many of the types the income-equivalent is nearly exactly equal to the housing subsidy yielding a bang-for-the-buck of 100%. The average income of the 24 types that take up the voucher is \$19 thousand, explaining why the benefits of the voucher are high. The types that accept a voucher are mixed racially, 6 African American and 18 Hispanic, and have more children (2.11) and have a higher expenditure share on rents (36 percent) than the types that do not take the voucher. When we use our model to simulate actual and counterfactual housing-voucher policies, we restrict the households that are offered housing vouchers to one of these 24 types.

3 Analysis of MTO

Moving to Opportunity (MTO) was a randomized control trial beginning in the 1990s that randomly assigned a group of households with children eligible to live in low-income housing projects in five U.S. cities to three different groups: (i) a treatment group that received a Section 8 housing voucher that in the first year could be applied only in Census tracts with a poverty rate under 10% and could be applied unconditionally thereafter, (ii) a second treatment group that received a comparable Section 8 housing voucher with no location restriction attached, and (iii) a control group that received no voucher. As described previously, voucher amounts were set such that after applying the voucher households spent no more than 30% of their income on rent.²² Summarizing the medium- to long-term impacts

²²Households that wanted to rent a more expensive unit were only allowed to contribute up to an additional 10% of their income. Note that our study focuses on the period after the Section 8 voucher and certificate

Figure 6: Income-Equivalent Value of Voucher and Amount of Housing Subsidy

Notes: This figure shows the income equivalent of the housing voucher, the blue dots, and the dollar amount of the housing subsidy provided by the housing voucher, the red line, for each of the 28 types of households with children that are eligible to receive housing vouchers based on their household income.

of MTO, [Sanbonmatsu, Kling, Duncan, and Brooks-Gunn \(2006\)](#), [Kling, Liebman, and Katz \(2007\)](#) and others show that on average the MTO treatment successfully reduced exposure to crime and poverty and improved the mental health of female children, but failed to improve child test scores, educational attainment or physical health. Later work by CHK demonstrated that MTO positively affected adult wages of children that were under the age of 13 at the time households first accepted an MTO voucher.

We ask if our estimated model can replicate the results of CHK. In a way we describe precisely in a moment, we use the Opportunity Atlas estimates of [Chetty, Friedman, Hendren, Jones, and Porter \(2018\)](#) to map the location choices of households receiving vouchers to the adult wages of children in those voucher-receiving households. The bottom line is that our model can nearly match the results of CHK. CHK estimate that the expected impact of accepting an MTO voucher on the annual adult earnings of each child under age 13 at the time the MTO voucher is accepted is \$3,477 with a standard error of \$1,418.²³ The equivalent estimate arising from our model simulations for Los Angeles is \$3,507.

We perform three sets of model simulations. For all simulations, we only consider the

programs were merged in 1999 and only vouchers were available.

²³See column (4) of Table 3 of their paper.

experiences of a small set of relevant households that we call “MTO Simulation Households,” the parameters of which we delineate soon. In the spirit of replicating the original, relatively small MTO experiment, we do not allow for any general-equilibrium effects on rental prices or any other variable. The three sets of simulations are:

1. Baseline: No household is offered a voucher. Household utility for type τ living in tract ℓ is δ_ℓ^τ , as estimated in section 2.
2. MTO: All households are offered MTO-style vouchers equal to the payment standard. Households receiving these vouchers must live in a Census tract with a poverty rate no greater than 10% in the first year. After the first year, they continue to receive the voucher and can live in any Census tract. Households that reject the initial offer of the MTO-style voucher are not offered a voucher in the future. All households understand the full set of program rules. We adjust δ_ℓ^τ for voucher-eligible tracts in the first year and all tracts in the second and subsequent years using equation (11).
3. MTO-R: We assign households to neighborhoods randomly according to the distribution of neighborhood poverty-rates that households are exposed to under the MTO simulations.²⁴

We define our MTO Simulation Households as those households that (a) are one of the 24 types of low-income agents with at least 0.5 children that are predicted to always accept an unrestricted housing voucher (see section 2.6.2) and (b) who reside at the start of the simulation in one of 15 Census tracts with at least 250 occupied non-senior-citizen public housing units.²⁵ While a few of the developments contain a small share of units set aside for senior citizens, these are predominately public housing developments for families with children. MTO Simulation Household types are represented in all simulations in proportion to their empirical distribution in the 15 public-housing tracts.

The assumptions we make about the age of children in households in the simulations in order to replicate the results of CHK are shown in Table 4 below. Column (2) of this table

²⁴Specifically, the procedure is as follows. (1) pool the set of simulated Census tract choices in MTO and the unconditional list of sample Census tracts. (2) Estimate a probit model predicting the probability that a record comes from the simulated data using only tract-poverty-rate categories as explanatory variables, and obtain the predicted probability p_j (propensity score) that a record from tract j comes from the simulated data. (3) Draw MTO-R simulated locations from the full set of Census tract with probability $Pr(j) = \frac{1}{J} \left(\frac{p_j}{1 - p_j} \right) \left(\frac{1 - \bar{p}}{\bar{p}} \right)$.

²⁵We also include the Census tracts containing Avalon Gardens and Hacienda Village which are below the 250 unit threshold but are proximate to several large developments. The MTO experiment also required the tracts to have a poverty rate of at least 40% in 1990. Of our 15 tracts, only 2 have a poverty rate of less than 40% in 2000; one tract has a poverty rate of 35.4% and the other tract poverty rate is 37.3%.

Table 4: Exposure by Age in Simulations

Percentage of Simulations	Years before Simulations Start	Years of Simulations
(1)	(2)	(3)
2.5%	3	15
5.0%	4	14
7.5%	5	13
10.0%	6	12
12.5%	7	11
12.5%	8	10
12.5%	9	9
12.5%	10	8
12.5%	11	7
12.5%	12	6

Notes: Column (2) shows the age of the child at which the MTO voucher is first offered. Column (3) shows the number of years we track the household, such that (2) and (3) sum to 18. Column (1) shows the percentage of times the particular row is included in the simulations.

shows the number of years we assume the child has lived in the initial location (one of the 15 tracts with public housing) prior to the simulation starting. Column (3) shows the number of years we simulate the model to determine optimal location choices in the baseline, MTO and MTO-R simulations. Column (1) shows the percentage of total simulations accounted for by the combinations shown in columns (2) and (3). We specify the distribution as shown in Table 4 to match three facts: First, the MTO experiment occurred between 1994 and 1998; second, CHK restrict their sample to children that are born on or before 1991; and third, the (significant) results of CHK are for children under the age of 13 at the time they are recruited for the MTO experiment. We assume households enter the MTO experiment uniformly between 1994 and 1998 and children of households in the MTO sample are born uniformly across years.

We keep track of the location of all households in all simulations and then map the sequence of locations to expected earnings of children using data from Opportunity Atlas. For each Census tract in the United States, [Chetty, Friedman, Hendren, Jones, and Porter \(2018\)](#) generate the Opportunity Atlas estimates by measuring the earnings of children given the earnings of parents using tax data from the IRS. For each tract, the Opportunity Atlas reports a child’s expected percentile in the nationwide income distribution at age 26 given household income of (a) the 25th percentile and (b) the 75th percentile of the nationwide income

distribution.²⁶ We map each type’s household income to the percentile of the nationwide household income distribution. Then, we use the two estimates from Opportunity Atlas to produce via linear interpolation (or extrapolation) an expected percentile in the nationwide age-26 income distribution for the children of that type of household in that tract. We interpret the Opportunity Atlas estimates as causal and for the analysis in this section, we assume the estimates for each tract are fixed. In a later section, we allow the Opportunity Atlas estimates to change in response to a possibly large policy intervention that alters the average income and racial composition of each neighborhood.

For all of the years before the simulation starts, column (2) of table 4, we assign the expected Opportunity Atlas percentile of the initial tract of residence (one of the 15 tracts described previously).²⁷ After the simulation starts, we assign the expected Opportunity Atlas percentile for each optimally chosen location in the simulation for the number of years shown in column (3). We average all 18 percentiles and then convert the resulting average percentile to a level of income using the nationwide age-26 income distribution for individual earners. The estimates we report are averages across simulated households of this level of income.

Table 5 reports simulation results. The table separately shows the outcomes of 8 household types that account for more than 90% of all MTO-voucher acceptances, the appropriately-weighted average outcome for the other 16 types, the overall average and the average of the top 8 rows. The top 8 rows sort types by household income (column 3) and then by housing expenditure share α (column 4). Column (2) shows the proportion of the type in the simulations, column (5) shows the race (B = African American and H = Hispanic) and column (6) shows number of children. The poverty rate and level of adult earnings of children in \$000s (per-child) from the Baseline simulations are shown in columns (7) and (8). Columns (9) - (11) show results from the MTO simulations. Column (9) shows the poverty rate for all households including those that do not accept the voucher; column (10) shows the percentage of households that accept the MTO-style voucher; and column (11) shows the per-child change in annual adult earnings, in \$000s, for the children of all households that accept the voucher. Column (12) shows results from the MTO-R simulations; this is the projected change in adult earnings in \$000s, per child, if households had randomly selected a tract with the same poverty rate as tracts actually chosen in the MTO simulation. The

²⁶The Opportunity Atlas data are for 2010 Census tracts. We use the Census Tract Relationship File from the U.S. Census Bureau to map the 2010 tracts to 2000 tracts. For the two cases in which this mapping fails, we assign an Opportunity Atlas number to those two tracts using spatial interpolation.

²⁷We assume that the household continuously resided in the initial tract of residence prior to the start of the simulation. Since this assumption is imposed for the baseline, MTO and MTO-R simulations, it will not affect comparisons across simulations.

averages reported are per-household except for the adult-earnings columns (8, 11 and 12), which are reported on a per-child basis such that these estimates are compatible with those of CHK.

Overall, three results are worth emphasizing. First, perhaps not surprisingly, the MTO experiment reduced exposure to poverty. The average poverty rate of the Census tract of residence falls from 41 percent in the baseline to 34 percent in the MTO simulations. The overall reduction in poverty for the top 8 types is more dramatic, from 43 percent to 27 percent. Second, our overall average simulated voucher take-up rate in the MTO simulations is only 26.2%. Recall that *all* of the Simulation Households are predicted to accept a location-unrestricted voucher. Our predicted take-up rate is much lower than the actual MTO take-up rate in Los Angeles of 62%. The additional counseling that MTO offered as noted by [Galiani, Murphy, and Pantano \(2015\)](#) likely played an important role in explaining the difference between simulated and actual voucher take-up rates. A different, complementary, story is that the distribution of types in the MTO experiment may be different than that in our simulations.

To see this more clearly, consider the experiences of the eight types of households that account for 90% of all households accepting a voucher in the simulations. These eight types of households are poor (average income of \$13 thousand) and mostly African American (6 of 8 types) and have a relatively high average expenditure share on rents of 45%. Each of these types has a voucher-acceptance rate in the MTO simulations of more than 20% such that the average voucher take-up rate of these types is 56%. It may be the case that these 8 types are over-represented in the MTO experimental data relative to the other 16 types of Simulation Households we consider. Since these 8 types account for almost all of the households accepting a voucher, a downweighting of the other 16 types would boost the simulated voucher take-up rate but would not affect our results on the impact on adult earnings of children conditional on households accepting a voucher.²⁸

Finally, our simulations nearly exactly match the reported CHK estimate on the impact of accepting the MTO voucher on the adult earnings of children under the age of thirteen. As mentioned, the CHK estimate is \$3,477 and our estimate is \$3,507. Additionally, the range across the 8 main types of voucher-recipients is relatively small, from \$2,093 (type 136) to \$3,942 (type 133). Interestingly, the results from the MTO-R simulations suggest the impact on adult earnings from MTO-style vouchers had the potential to have been much greater. Had the households that accepted a voucher selected a tract *randomly* with the

²⁸Of course, even if we strictly limit the simulations to only these 8 types, we would underestimate the overall take-up rate at 56% as compared to 62%. That gap may represent the impact of counseling; or it might reflect a weighting of the 8 types in the MTO experimental data that is different from the simulations as 5 of the 8 types have a take-up rate of 66% or greater.

Table 5: Simulations of MTO-Style Vouchers for MTO Simulation Households

Type	Sim. Share	Demographics				Baseline		MTO			MTO-R
		w^τ	α^τ	Race	k^τ	Pov. Rate	AE	Pov. Rate	Take-Up Rate	Treated Δ AE	Treated Δ AE
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
139	0.058	12.0	0.269	B	1.30	0.44	8.0	0.35	28.9%	2.75	6.61
142	0.031	12.0	0.346	B	1.29	0.39	9.4	0.23	66.2%	3.02	5.87
133	0.104	12.0	0.388	B	3.00	0.45	7.2	0.24	67.3%	3.94	8.13
143	0.023	12.0	0.530	B	0.53	0.40	10.0	0.18	83.3%	2.75	5.66
28	0.051	12.0	0.623	H	3.00	0.42	9.3	0.22	71.1%	3.50	6.25
136	0.053	12.0	0.657	B	0.96	0.43	8.3	0.23	72.1%	2.09	6.92
32	0.037	13.5	0.498	H	3.00	0.39	12.5	0.24	58.6%	3.02	4.01
137	0.065	18.4	0.380	B	1.82	0.43	10.3	0.37	21.4%	3.51	6.74
Other 16	0.578	21.7	0.303	H	2.23	0.40	13.9	0.39	4.4%	3.71	5.20
Avg. top 8		13.1	0.447		2.07	0.43	8.9	0.27	56.1%	3.45	6.80
Overall Avg.		18.1	0.364		2.16	0.41	11.9	0.34	26.2%	3.51	6.65

Notes: Column (1) is a type reference number and column (2) is the share of that type in the simulated samples. Column (3) is estimated household income in \$000s, (4) is estimated value of α , (5) is assigned race (B = African American and H = Hispanic) and (6) is estimated number of children. Column (7) is average poverty rate in the baseline simulations and column (8) is expected adult earnings in \$000s of each child, also in the baseline simulations. Columns (9) - (11) refer to the MTO simulations: (9) is the average poverty rate of everyone *offered* a voucher, (10) is the percentage of households that accept the MTO-style voucher and (11) is the change in the expected adult earnings in \$000s of each child relative to baseline *conditional on accepting a voucher*. Column (12) is the change in expected adult earnings in \$000s of each child for households that accept a voucher, relative to baseline, in the MTO-R simulation in which households randomly choose a tract.

same poverty rate as the tract they actually chose, the expected impact on per-child adult earnings from the MTO-experiment compared to baseline would have been \$6,651, nearly twice as large. In other words, conditional on the tract having a poverty rate of less than 10%, in the MTO experiment households negatively selected into tracts – a result that holds for every one of the 8 types we emphasize.

So, why did households select into relatively low Opportunity Atlas neighborhoods when offered an MTO-style voucher? Note that we can use equation (9) to estimate amenities for each type in each tract and then use equation (10) to estimate of the flow utility of the tract when using a voucher, $\delta_{\ell,\nu}$. Equation (10) can be rewritten as

$$\left(\frac{\sigma_\epsilon}{\alpha}\right) \delta_{\ell,\nu} - \left(\frac{1-\alpha}{\alpha}\right) \ln [0.7w] - \ln std = \left(\frac{1}{\alpha}\right) \ln A_\ell - \ln r_\ell \quad (14)$$

This equation shows that if households are reluctant to move into high Opportunity-Atlas neighborhoods, either the level of amenities is low or rental prices are too high given the level of amenities.

For each of the 8 types, we ran median regressions (least absolute deviation) of $\left(\frac{1}{\alpha}\right) \ln A_\ell$ on the Opportunity Atlas score for all 508 tracts in Los Angeles with a poverty rate less than 10%.²⁹ For all 8 types, the conditional median of log amenities is *decreasing* with Opportunity Atlas scores in these low-poverty neighborhoods. For six of the types, the estimated negative slope is statistically significant. Additionally, we ran a median regression of the Opportunity Atlas score on log rental prices and estimated a positive coefficient of 0.963 with a standard error of 0.20, implying that rental prices increase with Opportunity Atlas scores.³⁰ The bottom line is that MTO-voucher-receiving households negatively select into relatively low Opportunity Atlas score neighborhoods both because they prefer the amenities of these neighborhoods and because the rental prices are low.

4 Large Policy Experiments

In this section, we simulate our model to ask what would happen to the adult earnings of children of voucher recipients if the county of Los Angeles were to implement a policy like that in MTO, in which the location choices of voucher-recipients was restricted. Rather than directly condition feasible location choices on poverty rates, as was the case in the

²⁹The specific score we use is the child’s forecasted percentile in the age-26 earnings distribution given parent income in the 25th percentile of the earnings distribution.

³⁰This estimate implies that for each 10 percentage point increase in the neighborhood’s impact on the the child’s percentile in the earnings distribution according to the Opportunity Atlas data, log rental prices increase by 9.6 percent.

MTO experiment, we assume policy-makers in Los Angeles restrict the set of voucher-eligible neighborhoods based on the Opportunity Atlas scores of those neighborhoods.

In each policy experiment, we restrict the set of Census tracts where vouchers can be used based on the neighborhood’s Opportunity Atlas score, its forecasted percentile in the age-26 income distribution of a child’s adult earnings conditional on the parents earning the 25th percentile of the income distribution. We specify a cutoff value such that voucher-eligible neighborhoods are restricted to the top X^{th} percentile of Opportunity Atlas neighborhoods. We consider 10 possible cutoffs in total: $X = 10, 20, 30, \dots, 90, 100$. To illustrate, when $X = 10$, households receiving a voucher are only allowed to live in the top ten percent of neighborhoods based on the Opportunity Atlas score of that neighborhood. When $X = 100$, voucher recipients can live in any neighborhood. We call the results from the $X = 100$ experiment our baseline, since it essentially implements current policy.

We run each experiment exactly the same way: 11.18% of each of the 24 voucher-eligible types with children described earlier, currently living in any location, are offered a housing voucher exactly equal to the payment standard. The set of households that are offered vouchers is pre-determined and does not change; if a household ever declines the voucher in any given period, the household may accept the voucher in a later period. We choose 11.18% such that in the baseline simulation, the number of voucher-receiving households with children is equal to 2.02% of all renter households, the same as in the data for Los Angeles County in 2000 (32,993 voucher-receiving households with children and 1,634,030 rental households in total). In the experiments where we restrict the set of neighborhoods that are voucher-eligible, the percentage of households that accept housing vouchers falls, implying total expenditures on vouchers declines. Of course, policy-makers interested in maintaining constant expenditures on vouchers have the option of boosting the payment standard or increasing the number of households offered a voucher. We do not consider these alternatives as we wish to evaluate how restricting the feasible set of location choices of a fixed set of households, with no other policy parameters adjusted, changes the voucher take-up rates and adult earnings of children of those households.

In all simulations, we compute the optimal decisions of all households, including those that are not offered a voucher, to determine the steady state distribution of types across Census tracts. Rental prices in each simulation are determined in equilibrium such that total housing supply is equal to total housing demand in each tract.³¹ Additionally, we specify a tract-level housing supply elasticity of 0.25 such that the stock of housing can

³¹Housing demand in tract ℓ for a type τ household without a voucher is $\alpha^\tau w^\tau / r_\ell$ and is equal to std / r_ℓ for a type- τ household with a voucher.

expand or contract in the event rental prices change.³² Explaining, denote $\mathcal{H}_\ell^b$ as the total stock of housing in tract ℓ in the baseline and $\mathcal{H}_\ell^c$ as the stock of housing in a given policy experiment. Then the change in the housing stock in that tract is linked to the change in rental prices between baseline r_ℓ^b and experiment r_ℓ^c as follows:

$$\ln(\mathcal{H}_\ell^c/\mathcal{H}_\ell^b) = 0.25 \cdot \ln(r_\ell^c/r_\ell^b) \quad (15)$$

Before discussing our results, we mention a few nuances in the baseline simulation. Given estimated preferences, we adjust rental prices and the housing stock in each tract from what we observe in the 2000 Census to generate a stationary distribution of types in each tract. Similar to equation (15), we assume a tract-level housing supply elasticity of 0.25 in adjusting the stock of housing in the baseline relative to the data. Rental prices in the baseline and in current data are very similar. When we regress log rental prices in the baseline against log rental prices in the data for the 1,748 tracts in our sample, the R^2 of the regression is 0.78. The coefficient on log rental prices in the data is 1.04 with a standard error of 0.013, implying relatively expensive tracts in the data are even more expensive in the baseline. When we regress the Opportunity Atlas score for all 1,748 tracts on log rents in the data and then in the baseline, the coefficients are 1.50 and 1.68, respectively. These coefficients imply that for a household to increase its Opportunity Atlas score from the 37.1 percentile to the 52.1 percentile – this is a change from the bottom 10 percent of Opportunity Atlas tracts to the top 10 percent – the predicted change in log rents in the data is 0.226 and in the baseline is 0.253.

4.1 Fixed Opportunity Atlas

We consider two possibilities in our simulations. The first, which we discuss now, is that the Opportunity Atlas score for each tract does not change from the baseline. Later on, we allow each tract’s Opportunity Atlas score to depend deterministically on the steady-state mix of types occupying the tract.

The top panel of Figure 7 shows how the various alternate voucher policies affect the aggregate average annual adult earnings of *all* children of renting households in Los Angeles in millions of dollars.³³ The dashed blue line shows the positive impact to adult earnings of

³²Baum-Snow and Han (2019) estimate within-city tract-level housing supply elasticities ranging from 0.10 to 0.483.

³³For any given policy experiment, we know the cross-sectional steady state distribution of locations of households offered a voucher and households not offered a voucher. We use these distributions to compute the average Opportunity Atlas score of children of both sets of households. We then convert this averaged Opportunity Atlas score, which is a percentile of the age-26 income distribution, into an level of annual

children of households offered a voucher, relative to baseline; the dotted red line shows the negative impact of the policy on children of households not offered a voucher; and the solid black line shows the net impact for all children. At $X = 100$, there are no impacts at all since this experiment replicates current policy. The policy that maximizes the aggregate earnings of all children in Los Angeles is $X = 10$ which limits the voucher-eligible neighborhoods to the top 10% of all Opportunity Atlas neighborhoods. At this policy, the total net impact to adult annual earnings of children is \$28.7 million, about \$19 per year per child.³⁴ This net benefit reflects a positive benefit of \$43.1 million to all children of households offered a voucher and a loss of \$14.4 million to all children of households not offered a voucher. The policy that maximizes the benefit to only the children of households offered a voucher is $X = 20$. This policy yields an aggregate improvement in the adult earnings of children of households offered a voucher of \$43.2 million.

The bottom panel of Figure 7 shows how various policies affect the location decisions of the 24 types of households eligible to be offered a voucher. The y-axis indicates the percentage of these households that choose to locate in a voucher-eligible neighborhood. The red dots show the percentage for all households offered a voucher and the blue pluses show the percentage of households (given the same distribution over types) that are not offered a voucher. The gap between the red dots and blue pluses illustrates the impact of the voucher on location choices. The figure illustrates that the policy experiments can dramatically change where households live. For example, at $X = 20$, the 24-types of households we study that are not offered a voucher live in voucher-eligible tracts only about 10 percent of the time, whereas the households that are offered the voucher live in these tracts more than 60 percent of the time – a 50 percentage point increase.

At either $X = 10$ or $X = 20$, the policy creates enormous gains per-child for the relatively few children of households offered a location-restriction voucher and fairly small per-child losses for the large number of children of households not offered a voucher. To give an illustration of the gains, Table 6 shows outcomes for each of the 24 types offered a voucher. The types are sorted by household income, column (3), and then by housing expenditure share (not shown). Column (5) shows average adult earnings of children in \$000s in the steady state of the baseline simulations, where households with vouchers can live anywhere. Columns (6) and (7) of the table refer to results from the simulation that maximizes total adult earnings of all households offered a voucher, $X = 20$. Column (6) is the voucher take-up rate and (7) is the improvement in per-child adult earnings in \$000s of all households offered a voucher. Column (8) is the take-up rate at the type-specific value of X that maximizes

income.

³⁴The average number of children per renting household for the 1,634,030 renting households is 0.93.

Figure 7: Analysis of Various Voucher Policies: Fixed Opportunity Atlas

(a) Aggregate Impact, Adult Earnings of All Children

(b) Impact on Location Choices of Households Offered a Voucher

Notes: For various experiments restricting where voucher recipients can live ($X = 10, 20, \dots, 100$, with $X = 10$ corresponding to the top 10 percent of Opportunity Atlas tracts, and so forth), the top panel of this figure shows the aggregate impact to annual adult earnings of children of households offered a voucher (dashed blue line), children of households not offered a voucher (red dotted line), and all children of renting households in Los Angeles (solid black line). The bottom panel shows the frequency with which households offered a location-restricted voucher choose to live in one of the acceptable locations (red dots); the blue plusses show the frequency households not offered a voucher choose to live in one of the acceptable locations. In all experiments, each tract's Opportunity Atlas score is fixed.

per-child adult earnings of all households offered the voucher, i.e. the preferred “Pref” X for that type, and columns (9) and (10) are analogous to columns (6) and (7).

For convenience, we have divided the table into three bands of types. The top band of 9 types with the lowest income almost always accepts a voucher. For these types the average annual gain in adult earnings of children of households offered a voucher is an enormous \$11.44 thousand per child at $X = 20$, shown in column (7). This estimate includes experiences of children of households that do not accept the voucher. For reference, the average expected adult income of children of these households is \$15.4 thousand (column 5) implying the increase in income of \$11.44 thousand is equivalent to a 75% raise. At $X = 20$, the middle band of 8 types accepts a voucher with probability that ranges from 65 to 90 percent; for most of the types in this band, $X = 20$ is the value that maximizes adult earnings of children of households offered a voucher, shown in column (8). For these types, at $X = 20$, the average annual gain in adult earnings of children of households offered a voucher is \$7.67 thousand per child. The reduction in the benefit per-child relative to the 9 types of the top band reflects the lower take-up rate. Finally, at $X = 20$, the takeup rate of the bottom band of 9 types ranges from 15 to 50 percent. Since the takeup rate of these types is low, the impact of the the voucher on adult earnings of children is only \$2.56 thousand per year. Shown in column (8), these types would all prefer a voucher with fewer location restrictions. Even at the preferred value of X for these types, the takeup rates are still relatively low at 54%, on average. Overall, the impact on adult earnings of children of households of the 24 types offered a voucher is maximized at $X = 20$ because the take-up rate is high and the benefits of take-up are large for the 15 types of households in the top two bands of Table 6.

One might wonder why the average voucher takeup rate is 64% in the $X = 20$ policy experiment when the simulated MTO takeup rate with the same type mix of households is only 34 percent. After all, when $X = 20$ households with vouchers can live in only one of 350 tracts, but in the MTO experiment households could choose to live in one of 508 tracts with a poverty rate less than 10%. The difference can be explained by the nature of the policy simulations. The takeup rate in the $X = 20$ policy simulation is computed based on the steady-state. In contrast, we do not compute a steady state of the MTO policy experiment. Rather, any household that is offered a voucher must move to an eligible neighborhood in the first year and after that the household can live in any neighborhood. Had the MTO voucher been offered every year, presumably some households that refused the voucher in the first year may have eventually have taken it up, and given high moving costs, may have stayed in the neighborhood for quite some time.

Finally, it may seem surprising that there are any aggregate gains to adult earnings of

Table 6: Policy Experiments: Results by Type

Demographics				Base. AE	Results at $X = 20$		Results at Pref. X		
Type	Race	w^τ	k^τ		Take-Up Rate	Treated Δ AE	Pref. X	Take-Up Rate	Treated Δ AE
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
19	H	12.0	2.88	16.7	81%	8.24	20	81%	8.24
139	B	12.0	1.30	13.6	99%	14.65	10	99%	16.11
142	B	12.0	1.29	12.5	99%	15.12	10	99%	16.65
133	B	12.0	3.00	13.7	100%	12.88	10	100%	15.68
143	B	12.0	0.53	12.1	99%	14.95	10	100%	17.53
28	H	12.0	3.00	14.0	98%	13.03	10	98%	15.72
136	B	12.0	0.96	10.3	96%	17.05	10	95%	18.84
32	H	13.5	3.00	17.2	99%	9.59	10	99%	12.00
57	H	17.7	2.57	18.6	97%	10.14	10	96%	11.40
137	B	18.4	1.82	14.3	85%	11.21	20	85%	11.21
49	H	18.4	2.68	16.7	88%	9.83	10	76%	10.45
33	H	18.9	3.00	18.8	85%	7.36	20	85%	7.36
46	H	19.8	1.97	16.9	65%	7.26	20	65%	7.26
73	H	20.8	1.70	17.1	68%	7.51	20	68%	7.51
65	H	21.3	1.25	22.3	69%	3.26	20	69%	3.26
55	H	21.8	3.00	17.3	50%	5.12	30	64%	5.41
52	H	21.9	2.54	16.3	41%	4.73	30	55%	5.28
21	H	22.6	0.93	20.2	37%	2.08	30	50%	2.26
30	H	23.0	2.37	18.6	41%	3.62	30	55%	3.88
58	H	24.1	3.00	18.6	44%	3.75	30	54%	3.96
68	H	24.7	3.00	18.7	15%	0.80	50	53%	1.40
24	H	25.6	0.51	19.4	29%	1.92	40	54%	2.42
76	H	26.4	1.43	19.8	19%	0.53	50	53%	0.86
29	H	26.8	3.00	19.0	16%	0.75	50	53%	1.37
Average 1st 9 types				15.4	96%	11.44		96%	13.23
Average 2nd 6 types				17.8	75%	7.67		73%	7.81
Average 3rd 9 types				18.5	30%	2.56		54%	2.96
Average all types				17.3	64%	6.93		72%	7.70

Notes: Column (1) is a type reference number, column (2) is assigned race (B = African American and H = Hispanic), column (3) is estimated household income in \$000s and column (4) is estimated number of children. Column (5) is average adult earnings of children in the baseline ($X = 100$) simulations. Columns (6) and (7) refer to results from the simulation $X = 20$. Column (6) is the voucher take-up rate and (7) is the improvement in per-child adult earnings of all households offered a voucher. Column (8) is the takeup rate at the type-specific value of X that maximizes per-child adult earnings of all households offered the voucher, i.e the preferred “Pref” X for that type. Columns (9) and (10) are analogous to columns (6) and (7).

children from a voucher policy. Consider an environment, different from ours, where (a) all households have one child, (b) each household consumes one unit of housing, (c) the housing supply is fixed in all tracts and (d) no housing units are vacant. In this example, any voucher policy that encourages people to move out of a “bad” neighborhood and into some other “good” neighborhood displaces existing residents out of the good neighborhood. The voucher policy yields a reshuffling of the population but since all housing units in all neighborhoods are always occupied, the aggregate impacts of the voucher policy are zero.³⁵

There are three reasons why vouchers in our framework may yield positive aggregate impacts to the adult earnings of children. First, some households do not have children and a voucher policy that replaces childless households with households with children in high Opportunity Atlas score neighborhoods will yield improvements to aggregate earnings of children. Second, households consume differing quantities of housing. Even if the housing stock is fixed, a voucher policy may wind up generating increases in adult earnings in the aggregate by swapping one relatively rich household with children in a high Opportunity Atlas score area for two relatively poor households also with children. Finally, we allow for a relatively small housing supply elasticity, such that areas with increased rents also have more housing. We checked that the aggregate net gains we compute at $X = 10$ are not attributable to new housing by setting the housing supply elasticity to 0 and re-simulating the $X = 10$ experiment. When the housing stock is completely fixed, the aggregate net benefit at $X = 10$ falls very slightly, from \$28.7 million to \$27.9 million.

4.2 Varying Opportunity Atlas

A concern with the analysis of the previous section is that we hold the Opportunity Atlas scores fixed in every location while moving a possibly large number of households from one set of locations to a different set of locations. In other words, we assumed that the Opportunity Atlas score of a neighborhood does not depend on who lives in that neighborhood. In this section, we allow each neighborhood’s published Opportunity Atlas score to vary according to a simple function of the race and income of the residents of that neighborhood. Any voucher policy that changes neighborhood composition may also change the adult earnings of children of that neighborhood.

For each of the two published Opportunity Atlas scores we use in our analysis – the child’s expected percentile in the age-26 nationwide income distribution given household income of the 25th percentile and the 75th percentile of the nationwide income distribution – we regress

³⁵This result also requires that the effects of neighborhoods on adult earnings of children are independent of neighborhood composition; we return to this in a moment.

Table 7: Regressions of Opportunity Atlas Data on Income and Race

Regressor	Child Expected Income Percentile (x 100) if Household Income is at the:	
	25th Percentile	75th Percentile
(1)	(2)	(3)
Average Household Income (\$0000s)	0.740*** (0.239)	1.836*** (0.276)
African American share	-20.370*** (1.175)	-17.670*** (1.357)
Hispanic Share	-8.824*** (0.683)	-5.730*** (0.788)
Constant	47.060*** (1.333)	49.760*** (1.540)
Observations	1,748	1,748
R-squared	0.399	0.326

Notes: This table shows regressions of Census-tract-level Opportunity Atlas data on tract-level average income (in \$0000s) and share of African American and Hispanic households in the tract. The regressors (the Opportunity Atlas Data) are the expected child percentile in the income distribution as an adult (times 100) given household income in the 25th percentile and the 75th percentile. Standard errors are in parentheses.

*** $p < 0.01$

the published score multiplied by 100 on average household income (in tens of thousands of dollars) and the percentages of the neighborhood that are African American and Hispanic. The income and race regressors for each of the 1,748 tracts are generated using data from our 144 household types.

The regression results are shown in Table 7. Explaining the coefficients, using column (2) as an example: All else equal, if average income increases by \$10 thousand, then the predicted percentile in the income distribution of the child’s income at age 26 falls by 0.74; if the share of African American households increases by 10 percentage points then the predicted percentile falls by 2.037; and if the share of Hispanic households increases by 10 percentage points then the predicted percentile falls by 0.88. Remarkably, this simple regression can account for a large share of the variation of the Opportunity Atlas data, as the R2 values are 40 percent for the 25th percentile regression and 33 percent for the 75th percentile regression.

For each policy experiment, $X = 10, 20, \dots, 100$, we construct alternative Opportunity Atlas measures for each tract – for household income at both the 25th and 75th percentiles

in the income distribution – as follows. First, we use the regression coefficients reported in Table 7 to predict the Opportunity Atlas in each tract given the average household income and share of African American and Hispanic households resulting from the steady state of the policy experiment. Then, we add the residuals from the regression. This procedure only adjusts tract-level reported Opportunity Atlas scores if there is a change relative to the data of either household income or racial composition. Once we have revised estimates of the Opportunity Atlas scores in hand for household income in the 25th and 75th percentiles in the income distribution, we use the linear interpolation procedure described in section 3 to impute an Opportunity Atlas score to any household given the income of that household.

Before discussing our results, we wish to highlight important caveats. First, and obviously, correlation does not imply causation. Although racial shares and household income are highly correlated with the Opportunity Atlas scores, this does not imply that changing the racial composition or average income of a neighborhood will change the Opportunity Atlas score of that neighborhood. For example, high-income households may simply be more willing to pay higher rents that may be required to live in high Opportunity Atlas neighborhoods, thus inducing a correlation of the two series; this does not mean that moving lower-income households into a neighborhood will reduce the Opportunity Atlas score of that neighborhood.

Second, and equally importantly, we assume each type’s unobserved amenities from living in a neighborhood, the $\ln A_\ell$ term in equation (9), stays constant even if the racial composition or the average income of the neighborhood changes. Households may care about fixed neighborhood amenities that may be correlated with racial composition, average household income, and even Opportunity Atlas scores in the baseline. But, we assume households do not *directly care* about race or income of their neighbors or the Opportunity Atlas score. Related, if the racial or economic composition of a neighborhood changes we assume that the neighborhood amenities that households value do not change. This is a very strong assumption that we make for two reasons. First, in many models where households care about the composition of their neighbors, multiple equilibria may exist. It is unclear in our environment how to check for the presence of multiple equilibria and then how to select the appropriate equilibrium in the presence of multiplicity. Additionally, and related, households in these models need to form expectations about the composition of neighborhoods, and in rational-expectations equilibria these expectations must be consistent with outcomes, as documented by [Davis, Gregory, and Hartley \(2019\)](#). Computing an equilibrium is computationally very costly in our environment as it requires solving for expectation-consistent racial and economic composition of each of the 1,748 tracts in Los Angeles, in addition to market-clearing rents in those tracts.

Figure 8: Analysis of Various Voucher Policies: Varying Opportunity Atlas

Notes: For various experiments restricting where voucher recipients can live ($X = 10, 20, \dots, 100$, with $X = 10$ corresponding to the top 10 percent of Opportunity Atlas tracts, and so forth), this figure shows the aggregate impact to annual adult earnings of children of households offered a voucher (dashed blue line), children of households not offered a voucher (red dotted line), and all children of renting households in Los Angeles (solid black line). Each tract's Opportunity Atlas score is allowed to vary depending on the income and racial composition of the tract.

The results of the policy experiments when Opportunity Atlas scores are allowed to vary are shown in Figure 8, which plots the aggregate improvement in adult earnings of children of households offered a voucher (blue dashed line), the aggregate loss of adult earnings of children of households not offered a voucher (red dotted line), and the net aggregate gain (solid black line). Restricting vouchers to be used in the top 10% of Opportunity Atlas tracts ($X = 10$) maximizes the total gain in adult earnings of children in this environment.³⁶ At this policy, total net impact to adult annual earnings of children is \$33.6 million, which is *greater* than the equivalent estimate of \$28.7 million when we assumed Opportunity Atlas scores were invariant to policy. The net benefit of \$33.6 million can be decomposed into a positive benefit of \$39.2 million of all children of households offered a voucher and a loss of -\$5.6 million of all children of households not offered a voucher. Relative to the analysis in which Opportunity Atlas scores are fixed, the biggest change is that the aggregate loss in annual earnings to children of households not offered the voucher shrinks from -\$14.4 million to -\$5.6 million. This occurs because the racial and economic composition of relatively low Opportunity Atlas score neighborhoods changes such that Opportunity Atlas scores in those

³⁶In all policy experiments in this section, tracts are restricted or not based on Opportunity Atlas scores in the baseline.

neighborhoods improve.

5 Conclusion

In this paper we ask whether a policy that restricts the location choices of renting households in Los Angeles can improve the expected adult earnings of children of those households.

We answer this question by estimating household preferences over all Census tracts in Los Angeles using an infinite horizon, discrete-choice model that includes moving costs, and where households have preferences for consumption, housing and location-specific amenities. We allow preferences to vary across the population of renting households in Los Angeles by categorizing this population into one of 144 types. We estimate preferences separately for each type.

We simulate the behavior of 24 types of households that we estimate have children and have income sufficiently low to be eligible to receive a housing voucher to understand if the model can replicate results from the well-known MTO experiment. We offer to each of these types of households a restricted-location voucher that is in the style of the voucher offered in the MTO experiment. Our estimated model can nearly exactly replicate the results of CHK. These authors estimate that the expected impact of accepting an MTO voucher on the annual adult earnings of each child under age 13 at the time the MTO voucher is accepted is \$3,477. Our equivalent estimate is \$3,507.

We conclude our analysis by asking what would happen if Los Angeles were to convert its existing housing voucher program to one where vouchers can only be used in the top X percent of Opportunity Atlas neighborhoods. We set $X = 10, 20, \dots, 100$, where 100 means the voucher can be used in any neighborhood and $X = 10$ means the voucher can only be used in the top 10% of neighborhoods. In these experiments, we explicitly allow for general equilibrium effects in rents. As rents rise or fall, households may move in or out, and we explicitly keep track of all these changes.

We find that $X = 10$ maximizes the aggregate annual earnings of all children of renting households in Los Angeles and $X = 20$ maximizes the aggregate earnings of children of renting households offered location-restricted housing vouchers. The children of households accepting these vouchers experience substantial gains to annual income, in many cases as high as 75% of their baseline expected income. The children of households not offered vouchers experience small losses on average, as some households have to move from the locations that are most impactful on adult earnings to locations that are less impactful. On net, the gains to children of households offered vouchers outweigh the losses to other children and ultimately housing vouchers appear to have tremendous potential to improve intergenerational mobility.

References

- ARCIDIACONO, P., AND R. A. MILLER (2011): “Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity,” *Econometrica*, 79, 1823–1867. 9
- BAUM-SNOW, N., AND L. HAN (2019): “The Microgeography of Housing Supply,” Working Paper, University of Toronto. 35
- BAYER, P., F. FERREIRA, AND R. McMILLAN (2007): “A Unified Framework for Measuring Preferences for Schools and Neighborhoods,” *Journal of Political Economy*, 115(4), 588–638. 4, 22, 48, 49
- BAYER, P., R. McMILLAN, A. MURPHY, AND C. TIMMINS (2016): “A Dynamic Model of Demand for Houses and Neighborhoods,” *Econometrica*, 84(3), 893–942. 4, 6
- BERGMAN, P., R. CHETTY, S. DELUCA, N. HENDREN, L. F. KATZ, AND C. PALMER (2019): “Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice,” NBER Working Paper 26164. 3
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4), 841–890. 48
- BHUTTA, N., AND B. J. KEYS (2016): “Interest Rates and Equity Extraction During the Housing Boom,” *American Economic Review*, 106(7), 1742–1774. 10
- BISHOP, K. C. (2012): “A Dynamic Model of Location Choice and Hedonic Valuation,” Working Paper, Washington University in St. Louis. 8
- BISHOP, K. C., AND A. D. MURPHY (2011): “Estimating the Willingness to Pay to Avoid Violent Crime: A Dynamic Approach,” *American Economic Review*, 101(3), 625–629. 6
- Board of Governors of the Federal Reserve System (2007): “Report to Congress on Credit Scoring and Its Effects on the Availability and Affordability of Credit,” . 10
- BROWN, M., S. STEIN, AND B. ZAFAR (2015): “The Impact of Housing Markets on Consumer Debt: Credit Report Evidence from 1999 to 2012,” *Journal of Money, Credit, and Banking*, 47(S1), 175–213. 10
- CHETTY, R., J. FRIEDMAN, N. HENDREN, M. R. JONES, AND S. R. PORTER (2018): “The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility,” NBER Working Paper No. 25147. 5, 27, 29

- CHETTY, R., N. HENDREN, AND L. F. KATZ (2016): “The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment,” *American Economic Review*, 106(4), 855–902. 2
- CHYN, E. (2018): “Moved to Opportunity: The Long-Run Effects of Public Housing Demolition on Children,” *American Economic Review*, 108(10), 3028–3056. 2
- DAVIS, M. A., J. GREGORY, AND D. A. HARTLEY (2019): “The Long-Run Effects of Low-Income Housing on Neighborhood Composition,” Unpublished Paper. 42
- DAVIS, M. A., AND F. ORTALO-MAGNE (2011): “Household Expenditures, Wages, Rents,” *Review of Economic Dynamics*, 14(2), 248–261. 4
- DURLAUF, S. N. (2004): *Handbook of Regional and Urban Economics*chap. Neighborhood Effects, pp. 2174–2230. Elsevier B.V. 2
- EWENS, M., B. TOMLIN, AND L. C. WANG (2014): “Statistical Discrimination or Prejudice? A Large Sample Field Experiment,” *The Review of Economics and Statistics*, 96(1), 119–134. 14
- GALIANI, S., A. MURPHY, AND J. PANTANO (2015): “Estimating Neighborhood Choice Models: Lessons from a Housing Assistance Experiment,” *American Economic Review*, 105(11), 3385–3415. 3, 31
- GALLAGHER, J., AND D. HARTLEY (2017): “Household Finance after a Natural Disaster: The Case of Hurricane Katrina,” *American Economic Journal: Economic Policy*, 9(3), 199–228. 10
- HOTZ, V. J., AND R. A. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *The Review of Economic Studies*, 60(3), 497–529. 8
- KENNAN, J., AND J. R. WALKER (2011): “The Effect of Expected Income on Individual Migration Decisions,” *Econometrica*, 79(1), 211–251. 4, 6
- KLING, J. R., J. B. LIEBMAN, AND L. F. KATZ (2007): “Experimental Analysis of Neighborhood Effects,” *Econometrica*, 75(1), 83–119. 27
- LEVENTHAL, T., AND J. BROOKS-GUNN (2000): “The Neighborhoods They Live in: The Effects of Neighborhood Residence on Child and Adolescent Outcomes,” *Psychological Bulletin*, 126(2), 309–337. 2
- PHILLIPS, D. (2017): “Landlords Avoid Tenants who Pay with Vouchers,” *Economics Letters*, 151(C), 48–52. 14
- POPKIN, S. J., M. K. CUNNINGHAM, E. GODFREY, B. BEDNARZ, AND A. LEWIS (2002): “CHA Relocation Counseling Assessment,” The Urban Institute Final Report. 14

- ROSS, S. L. (2011): “Social Interactions within Cities: Neighborhood Environments and Peer Relationships,” in *The Oxford Handbook of Urban Economics and Planning*, ed. by N. Brooks, K. Donaghy, and G.-J. Knaap, chap. 9. Oxford Handbooks. [2](#)
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55(5), 999–1033. [7](#), [8](#)
- SANBONMATSU, L., J. R. KLING, G. J. DUNCAN, AND J. BROOKS-GUNN (2006): “Neighborhoods and Academic Achievement: Results from the Moving to Opportunity Experiment,” *Journal of Human Resources*, 41(4), 649–691. [27](#)
- YINGER, J. (1986): “Measuring Racial Discrimination with Fair Housing Audits: Caught in the Act,” *American Economic Review*, 76(5), 881–893. [14](#)

A Appendix

This appendix provides additional details on the instrumental variables method that we use to rescale preferences for consumption, housing and amenities (σ_ϵ) given the variance of the model’s i.i.d. preference shocks. Our maximum likelihood estimation of the location choice model identifies the indirect flow utility δ_ℓ provided by each tract for each household type. As shown in section [2.5.2](#), these estimated indirect utilities δ_ℓ for a given type can be related to tract rent and amenity levels by

$$\delta_\ell = \lambda \cdot \mathcal{O}_\ell - \left(\frac{1}{\sigma_\epsilon} \right) \cdot \alpha \ln r_\ell + \xi_\ell, \quad (\text{A1})$$

where r_ℓ is the tract rent level, and $\mathcal{O}_\ell$ and ξ_ℓ are observed and unobserved tract characteristics. The impact of log-rent on δ_ℓ depends on the budget share devoted to housing (α) and the scale of the ϵ -shocks (σ_ϵ) in consumption utility units. Having already estimated budget shares α (as described in section [2.5.1](#)), the parameter of interest ($\frac{1}{\sigma_\epsilon}$) can be thought of as the coefficient on α times log rent. Because equilibrium rents will almost certainly be correlated with unobserved but valued characteristics of neighborhoods, ξ_ℓ , consistent estimation of this coefficient requires an instrument Z that satisfies two conditions,

A1) *Instrument Relevance*: $\text{cov}(Z_\ell, \ln r_\ell) \neq 0$

A2) *Instrument Exogeneity*: $\text{cov}(Z_\ell, \xi_\ell) = 0$

The instruments must be predictive of tract rents but be uncorrelated with the unobservable component of tract amenity utility.^{[37](#)}

³⁷Note that the variation that must be instrumented is the log-rent variation exclusively, even though in

Our choice of instruments follows the approach proposed by [Bayer, Ferreira, and McMillan \(2007\)](#) for adapting the idea of “BLP instruments” ([Berry, Levinsohn, and Pakes, 1995](#)) from the IO literature to the urban setting where amenities are spatially correlated. The standard BLP instruments for the market-specific price of a product ℓ are (functions of) the *characteristics* of other products in the same market competing with ℓ . In differentiated consumer product markets, characteristics of competing products will affect the price mark-up that can be charged for ℓ in equilibrium, thus satisfying the instrument relevance condition, but will not directly affect the utility that a consumer derives from product ℓ conditional on choosing ℓ over the competing products, satisfying instrument exogeneity.

In our location choice framework, the city’s Census tracts are all “competing products” for one another whose characteristics are candidate instruments. As [Bayer, Ferreira, and McMillan \(2007\)](#) point out, however, the characteristics of tracts that are located very close to a tract ℓ are likely to be related to ξ_ℓ , because the quality of nearby housing can directly affect the utility one derives from living in a place. To address this concern, [Bayer, Ferreira, and McMillan \(2007\)](#) propose using as instruments the characteristics of the housing stock outside of a three mile buffer and including the characteristics of the housing stock inside of the three mile buffer as controls ($\mathcal{O}_\ell$).

The choice of the buffer distance involves a practical tradeoff for the researcher. With a larger buffer, the exclusion restriction is easier to believe, because the housing stock characteristics of neighborhoods farther from ℓ are less likely to provide a direct amenity value to residents of ℓ . However, instrument relevance may decline with the size of the buffer, because neighborhoods further from ℓ are likely be less substitutable with ℓ and therefore have a smaller influence on equilibrium rents in ℓ via competition. To be conservative, we chose a buffer of five miles instead of the three mile buffer used by [Bayer, Ferreira, and McMillan \(2007\)](#) after verifying doing so did not dramatically reduce the power of the first stage.

The rental housing stock characteristics that we include in the instrument list are as follows: The number of bedrooms (shares with 1, 2, 3, and 4 bedrooms), the number of rental units per building (share of units in buildings with 2, 3-4, 5-49, and 50+ units), and the vintage of the rental stock (share constructed pre-1939, shares by decade of construction from the 1940s to the 1980s, and two categories in the 1990s). Specifically, the instruments are the average of each tract-level share among tracts whose centroids are between 5 and 20 miles from the centroid of tract ℓ . We include in the list of control variables $\mathcal{O}_\ell$ each of these variables measured in tract ℓ itself and the average of each tract-level share among tracts

practice we estimate a coefficient $\alpha \times \ln r_\ell$ to yield a coefficient with the desired structural interpretation. That is because for any given type of household the parameter α does not vary across tracts.

with centroids within 5 miles of the centroid of tract ℓ .³⁸

We constrain the parameter $(\frac{1}{\sigma_\epsilon})$ in equation (A1) to be the same across types (τ) in the IV second stage regression by pooling the type-specific indirect utility measures into a single sample with one observation per type-tract pair ($144 \text{ types} \times 1,748 \text{ tracts} = 251,712$ observations). We allow for type-specific coefficients λ^τ on the tract-level controls $\mathcal{O}_\ell$ and estimate a single coefficient on $\alpha^\tau \times \ln r_\ell$, instrumenting for $\alpha^\tau \times \ln r_\ell$ with $\alpha^\tau \times Z_\ell$.

As described in section 2.5.2, the full IV procedure follows the three step approach used by Bayer, Ferreira, and McMillan (2007). In the first step, we recover an initial estimate of $(\frac{1}{\sigma_\epsilon})$ using the buffered list of housing stock characteristics as instruments. Then, in a second step, we use the estimates of $(\frac{1}{\sigma_\epsilon})$ and λ^τ from the first step, call them $\hat{\frac{1}{\sigma_\epsilon}}$ and $\hat{\lambda}^\tau$, to construct a new surface of indirect utilities for each type abstracting from unobservables as

$$\hat{\delta}_{\ell,\tau} = \hat{\lambda}^\tau \cdot \mathcal{O}_j - \left(\hat{\frac{1}{\sigma_\epsilon}} \right) \alpha^\tau \ln r_\ell$$

We simulate the model using this specification for indirect utility and adjust r_ℓ for all ℓ tracts until the simulated total housing demand in any tract is equal to the observed housing demand in the estimation sample for that tract.³⁹ This procedure determines market-clearing rents in all tracts in the absence of unobserved amenities. We use these rents as instruments to estimate $(\frac{1}{\sigma_\epsilon})$ in the third and final step.

Table A1 presents the first stage coefficients on the excluded instruments. For simplicity, the coefficients reported are from a single regression of $\ln r_\ell$ on the Z_ℓ and $\mathcal{O}_\ell$, as opposed to the coefficients from the pooled first-stage regression with $\alpha^\tau \times \ln r_\ell$ on the left hand side (which for each type equal the reported coefficients divided by each type's α^τ). Because each type contributes one observation per tract and α^τ is constant across each type's 1,748 observations, this single regression summarizes all of the exogenous variation extracted from the first stage and is sufficient for discussing the strength of the instruments. Column (1) presents first stage estimates from the “first step” that uses all buffered housing stock variables as instruments. The variables describing the distribution of units per building 5 to 20 miles from ℓ are most predictive of ℓ 's log-rent (joint p-value: 0.000), followed by the variables describing the age mix of the housing stock (joint p-value: 0.010). The F-stat for the joint significance of all excluded instruments is 5.35 (p-value: 0.000).

The second column reports first stage estimates from the “third step” where the iden-

³⁸Bayer, Ferreira, and McMillan (2007) also use variables related to land use in distant tracts as instruments.

³⁹Given Cobb-Douglas preferences, type-specific housing demand in tract ℓ is $\alpha^\tau w^\tau / r_\ell$, where w^τ is type-specific household income.

tifying variation is summarized in the single simulated rent instrument. The first stage F-statistic in column (2) is 31.7 (p-value: 0.000). Intuitively, the F-statistic rises because the first step only uses information about the quality of substitutes for each tract individually whereas the third step uses similar information for all tracts.

Table A1: Detail Table of IV Results

	(1)	(2)
Share of rental units with X bedrooms		
Share with 1 bedroom	-10.070 (8.002)	
Share with 2 bedrooms	-1.288 (4.188)	
Share with 3 bedrooms	-8.472 (9.604)	
Share with 4 bedrooms	-14.370 (16.31)	
Joint significance of bedrooms: p=	0.6340	
Share of renter-occ. units consisting of X units		
2 unit buildings	3.739 (16.780)	
3-4 unit buildings	3.033 (6.226)	
5-49 unit buildings	3.254* (1.706)	
50+ unit buildings	-15.910** (6.628)	
Joint significance of units per building: p=	0.0000	
Share of all rental units by vintage		
Share built 1995-1998	29.580 (38.720)	
Share built 1990-1994	-87.470** (40.220)	
Share built 1980-1989	35.700 (30.650)	
Share built 1970-1979	-3.465 (33.640)	
Share built 1960-1969	-7.330 (31.280)	
Share built 1950-1959	16.130 (32.490)	
Share built 1940-1959	-39.630 (34.040)	
Share built 1939 or earlier	0.781 (31.55)	
Joint significance of rental vintage: p=	0.0102	
Simulated ln(rent) instrument		0.226*** (0.0402)
Controls for own tract housing characteristics	X	X
Controls for housing characteristics w/in 5 miles	X	X
Observations	1,748	1,748
R-squared	0.669	0.732
All excluded instruments: F-statistic	5.35	31.69
All excluded instruments: p-value	0.000	0.000

Notes: This table shows the results of the 1st and 2nd stages of the IV. Robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.1