

Evidence of Neighborhood Effects from Moving to Opportunity: LATEs of Neighborhood Quality

Dionissi Aliprantis* Francisca G.-C. Richter*

February 26, 2018

Abstract: This paper estimates neighborhood effects on adult labor market outcomes using the Moving to Opportunity (MTO) housing mobility experiment. We propose and implement a new strategy for identifying transition-specific effects that exploits identification of the unobserved component of a neighborhood choice model. Estimated Local Average Treatment Effects (LATEs) are large, result from moves between the first and second deciles of the national distribution of neighborhood quality, and pertain to a subpopulation of nine percent of program participants.

Keywords: Neighborhood Effect, Local Average Treatment Effect, Moving to Opportunity
JEL Classification Numbers: C31, C36, C50, D04, I38, R23

*: Federal Reserve Bank of Cleveland, +1(216)579-3021, dionissi.aliprantis@clev.frb.org

*: Case Western Reserve University, +1(216)368-8686, francisca.richter@case.edu

Acknowledgments: The authors thank Nathaniel Baum-Snow, Greg Caetano, Karim Chalak, Ben Craig, Joel Elvery, Kyle Hood, Jon James, Jeffrey Kling, Carlos Lamarche, Fabian Lange, Alvin Murphy, Stephen Ross, anonymous referees, the editor Bryan Graham, and seminar and conference participants for comments that have improved this paper. We also thank Emily Burgen and Nelson Oliver for valuable research assistance, and Paul Joice at HUD for his assistance with the data set. The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

1 Introduction

Moving to Opportunity (MTO) was a once-in-a-generation housing mobility experiment designed to identify neighborhood effects (Shroder and Orr (2012)). Through the program, households living in high-poverty neighborhoods were randomly assigned housing vouchers for use in low-poverty neighborhoods. Receiving an MTO voucher had no effect on labor market or education outcomes four to seven years after assignment (Kling et al. (2007), Sanbonmatsu et al. (2006)).

The lack of effects from receiving an MTO voucher has been interpreted as evidence against the existence of neighborhood effects (Ludwig et al. (2013), Ludwig et al. (2008)). One of the strongest assumptions required for effects from MTO to be informative about effects from neighborhood environments is that neighborhood poverty measures neighborhood quality (Aliprantis (2017a)). This assumption is strong because MTO voucher recipients tended to move to predominantly black low-poverty neighborhoods (Clampet-Lundquist and Massey (2008), Sampson (2008)), and black low-poverty neighborhoods are comparable to white high-poverty neighborhoods along dimensions other than poverty (Aliprantis and Kolliner (2015)).

This paper identifies neighborhood effects with the MTO data. We propose a methodological innovation through which transition-specific effects in an ordered treatment model can be identified with a single discrete instrument. We also use a measure of neighborhood quality that includes characteristics other than poverty.

We find that moving to a higher quality neighborhood had large, positive effects on employment, labor force participation, household income, and welfare receipt. Although effects on Body Mass Index (BMI) and individual earnings are not statistically significant, these effects are also estimated to be large and positive. And while our estimates are noisy, they are strongly suggestive that increasing neighborhood quality improves adult outcomes.

Identifying Local Average Treatment Effects (LATEs) in a model with multiple treatment levels is difficult (Athey and Imbens (2017)). Two prominent approaches either identify weighted averages of transition-specific effects (Angrist and Imbens (1995)) or else require transition-specific instruments for each level of treatment (Heckman et al. (2006)). Our identification strategy combines the strengths of each of these approaches, in that it is both empirically feasible and economically interesting. The key insight of the identification strategy is that by observing the continuous level of treatment associated with each individual, it is possible to identify the unobserved, idiosyncratic component of their latent index in the treatment choice model. Identification is obtained by comparing individuals who are similar along the observed and unobserved dimensions of the ordered choice model, and

who therefore select into different treatment levels just because of the randomly assigned instrument.

We think there are two key reasons why, in contrast to previous studies, we find evidence of neighborhood effects in the interim MTO data. The first is that previous studies have estimated effects for broader populations that on average did not experience large changes in neighborhood quality. For example, estimates of neighborhood effects on adult labor market outcomes are absent from the most prominent analysis of MTO because the program was found to have little effect on such outcomes for larger subpopulations than we study (Kling et al. (2007)).¹ To find movers from the first to the second decile of neighborhood quality, we focus on a subpopulation so restrictive that our LATE estimates pertain to nine percent of program participants.

The second reason is that while poverty may be the best single variable for characterizing neighborhood quality, other characteristics may be important drivers of outcomes. Consider that quasi-experimental studies show that neighborhood effects on employment could result from neighborhood referrals (Bayer et al. (2008)) or proximity to jobs (Andersson et al. (2018)). These mechanisms are likely to be captured by our neighborhood quality index because the index includes unemployment and the employment to population ratio. Important mechanisms may be missed when considering poverty alone if low-poverty neighborhoods in MTO were negatively selected along other dimensions (Davis et al. (2016)).

We do not conclude that MTO was a failed policy because few households were living in better neighborhoods four to seven years after randomization. While this limited mobility does restrict the evidence provided by MTO about neighborhood effects (Wilson (1987)), we think this mobility provides clear evidence for policy. Providing access to opportunity neighborhoods will require trying more interventions, whether that means making subsidies tied to smaller geographic areas (Collinson and Ganong (2017)), supporting landlords with capital-needs financing and program streamlining (Galvez et al. (2010)), or experimenting with other models of assistance.

The paper proceeds as follows: Section 2 specifies our joint model of neighborhood choice and potential outcomes, as well as our strategy for identifying treatment effect parameters from this model. Section 3 describes the MTO housing mobility program, the data used in estimation, and descriptive statistics of those data. Section 4 presents our empirical specification and estimation results, and Section 5 concludes. An Appendix presents a formal justification for our identification strategy, the algorithm we use to estimate the model, and the full specification of the estimated model.

¹Kling et al. (2007) focus their attention on “outcomes that exhibit significant treatment effects” (p 83).

2 Model

We present a new approach to estimating Local Average Treatment Effects (LATEs) where a multi-level treatment is modeled through an ordered choice model with a restricted class of random thresholds and a binary instrument. This threshold restriction allows us to estimate an unobserved component of choice that then determines the specific treatment transitions for which LATEs can be recovered.

We are interested in estimating transition-specific LATEs for an outcome of interest Y in response to an ordered treatment D , with instrument Z influencing the level of treatment selection. Our application is the effect of neighborhood quality level on individual economic outcomes, using the random assignment of a housing mobility voucher as an instrument.

Let Ω be the sample space of events representing all individual types i in the experiment, with each individual possibly defining its own type.² We define the following random variables on Ω :

$D(i) : \Omega \rightarrow \{1, \dots, K\}$ is neighborhood quality - the treatment;

$Y_k(i) : \Omega \rightarrow \mathbb{R}$ is the potential outcome for individual i associated with treatment level k ;

$Y(i) = (Y_1(i), \dots, Y_K(i))$ is a random vector of potential outcomes for i ;

$Z(i) : \Omega \rightarrow \{0, 1\}$ is a Binomial random variable representing voucher assignment;

$W(i) : \Omega \rightarrow \{0, 1\}$ is a Binomial random variable representing individual-specific voucher take up;

$V(i) : \Omega \rightarrow \mathbb{R}$ is a random variable representing an individual-specific component of choice unobserved to the researcher;

$U_k(i) : \Omega \rightarrow \mathbb{R}$ is an individual-treatment level specific component of outcome $Y_k(i)$ unobserved to the researcher;

$X(i) : \Omega \rightarrow \mathbb{R}$ is an individual-specific component of choice observed to the researcher.

We will use lowercases and subindexes to denote realized values of these variables. So $X(i)$ is realized as x_i . We will refer to $D_z(i)$ as the treatment level of individual i when the value of instrument z_i is counterfactually set to z , and to $W_z(i)$ analogously.

²While it is possible that more than two individuals will correspond to the same type i , we will refer to i as individual rather than individual-type for ease of notation. Because individuals are assumed to be household heads, we sometimes also refer to household i .

Treatment selection among each one of K possible treatment levels, $k = 1, \dots, K$ is an ordered choice model with random thresholds defined by:

$$D(i) = k \iff C_{k-1} - Z(i)W(i)\gamma_{k-1} \leq \mu(x_i) - V(i) < C_k - Z(i)W(i)\gamma_k, \quad (1)$$

where

$\{C_k\}_{k=1}^K$ are positive and monotonically increasing in k ,

$\{\gamma_k\}_{k=1}^K$ are non-negative treatment-level constants, and

$\mu(x_i)$ is a linear function of x_i .

Potential outcomes are defined by

$$Y_k(i) = \mu_k(x_i) + U_k(i) \quad (2)$$

where $\{\mu_k(x_i)\}_{k=1\dots K}$ are treatment level specific functions of observed individual characteristics influencing outcome.

For Z to satisfy standard instrument conditions, we require that Z be randomly assigned

$$Z(i) \perp (X(i), U_1(i), \dots, U_K(i), V(i), W_0(i), W_1(i)). \quad (3)$$

We note that the non-negativity of the $\{\gamma_k\}$ implies monotonicity, but need not imply that the instrument is relevant (Imbens and Angrist (1994)'s Condition 1 (ii)) without further assuming this to be the case.

Our model is different than the ordered choice models in Vytlačil (2006), where a function of the instrument shifts each threshold by the same amount. In our model, the instrument $Z(i)$ can influence each threshold differently through the $\{\gamma_k\}$, and the influence on each threshold can vary across individuals through the random indicator $W(i)$.

We are able to estimate the parameters of a random (individual-varying) thresholds choice model under the restriction that each threshold only has a binary support. Because we observe all of the random variables in Equation 1 needed to determine choice, once $Z(i)$ and $W(i)$ are realized and observed, we are left with a choice model with fixed thresholds whose parameters can be estimated via Maximum Likelihood optimization.³

³When $z_i = 0$, cutoff values are $\{C_k\}_{k=1}^K$. When individual i is assigned the instrument but does not take it up, we observe $w_i = 0$, so that cutoff values remain the same as when the instrument is not assigned, $\{C_k\}_{k=1}^K$. However, when the individual is assigned the instrument and takes it up, we observe $w_i = 1$, and cutoff values are $\{C_k - \gamma_k\}_{k=1}^K$.

This setup is consistent with the way in which the MTO experiment was implemented. We can think of the counterfactual choice facing experiment participants as a two-step process. The first step determines whether an offered voucher is taken up. In the second step, the individual chooses the level of quality of her neighborhood, which is measured four to seven years after voucher assignment.

In the first step, we interpret willingness to participate in the MTO program as the choice to use a voucher if feasible.⁴ If an individual is offered a voucher but does not take it up ($w_1(i) = 0$), we interpret this outcome as an indication that it was not feasible for the household to improve its housing using a Section 8 voucher under the program rules for reasons unknown to the household at the time it volunteered for the program.⁵ The first step is complete once the voucher is randomly assigned and household take up is determined, so that $Z(i)$ and $W(i)$ are realized.

Identifying transition-specific LATEs by allowing thresholds to respond to instrument assignment is attractive relative to the alternatives, even if we must restrict the heterogeneity in response to the instrument to be binary. While a single instrument can more easily identify the Average Causal Response (ACR), this weighted average of transition-specific LATEs (Angrist and Imbens (1995)) would be difficult to interpret in our application, where transitions are rare. And while transition-specific instruments might impose weaker assumptions for identifying transition-specific LATEs (Heckman et al. (2006)), in many applications only a single instrument will be available.

How do we estimate transition-specific LATEs based on the estimated ordered choice model? A necessary condition is that we observe a continuous measure of treatment even when the ordered choice model assumes a discretized version of treatment. Indeed, we observe a continuous measure of the treatment variable - neighborhood quality. This allows us to set up a continuous version of the treatment choice model in Equation 1, in which q denotes the observed continuous level of neighborhood quality. In this case, the optimal quality level of neighborhood q^* satisfies the following First Order Condition (FOC):

$$\mu(x_i) - V(i) - C(q) + Z(i)W(i)\gamma(q^*) = 0 \tag{4}$$

By interpolating between estimates of $\{\widehat{C}_k\}$ and $\{\widehat{\gamma}_k\}$ from the estimated model in Equation 1, we obtain estimates of the continuous functions $\widehat{C}(q)$ and $\widehat{\gamma}(q)$ that can be substituted

⁴Chyn (2015) analyzes a related program in which moving was required.

⁵We would generally interpret this as time restrictions on initial take up combined with supply constraints, where supply constraints could be driven at least in part by landlords avoiding voucher holders (Phillips (2017)). This outcome could also result from the realization of idiosyncratic shocks after volunteering for the program.

into the FOC:

$$\hat{\mu}(x_i) - V(i) - \hat{C}(q) + z_i w_i \hat{\gamma}(q^*) = 0, \quad (5)$$

from which we derive $\hat{V}(i)$.

Having estimated this individual unobserved component of choice for all participants, we are able to identify transition-specific LATEs by estimating average outcome differences among those with similar $V(i)$'s and x_i 's but different instrument assignments. As noted in Heckman et al. (2006), the LATE is remarkable for the patterns of selection under which the parameter is interpretable and potentially identified. In our application, this includes sorting on unobservables and matching on neighborhood amenities (Graham (2018)).

2.1 Identification

The causal effects identified with any data set are determined by selection into treatment and response to the instrument. In our case, we will only be able to identify j to $j + 1$ transition-specific LATEs (for $j \in \{1, \dots, J - 1\}$) that pertain to a specific region of x_i 's and v_i 's where Z induces changes in treatment.

In the empirical analysis it will be convenient to refer interchangeably to parameters and sets defined in terms of $V(i)$ and $U_D(i)$, where $U_D(i) \equiv F_V(v_i)$, the cumulative distribution function of $V(i)$. We are interested in the regions of observed and unobserved characteristics, $(\mu(X_i), U_D(i))$, corresponding to households that would select into treatment level j in the absence of a voucher, but could possibly select into treatment level $j + 1$ if given an MTO voucher. To focus on this group, we define

$$\Omega_j \equiv \left\{ i \in \Omega \mid D_0(i) = j, D_1(i) \in \{j, j + 1\}, \Pr(D_1(i) = j + 1) > 0 \right\},$$

and the identification support set \mathcal{S}_j as

$$\mathcal{S}_j \equiv \left\{ (\mu(x_i), U_D(i)) \mid i \in \Omega_j \right\}.$$

We stress that membership in the identification support set \mathcal{S}_j can only be determined after identifying each household's $V(i)$ as established in the previous section.

Determining \mathcal{S}_j gives us the ability to identify transition-specific LATEs with a single discrete instrument. In Appendix A we prove

Proposition 1. *The Wald estimator applied to the subsample of experimental and control*

households in \mathcal{S}_j identifies the j to $j + 1$ transition-specific LATE:

$$\begin{aligned} \Delta_{j,j+1}^{LATE}(Z^M) &\equiv \mathbb{E}[Y_{j+1} - Y_j \mid i \in \mathcal{C}_j] \\ &= \frac{\mathbb{E}[Y \mid Z^M = 1, (\mu(X), U_D) \in \mathcal{S}_j^M] - \mathbb{E}[Y \mid Z^M = 0, (\mu(X), U_D) \in \mathcal{S}_j^M]}{\mathbb{E}[D \mid Z^M = 1, (\mu(X), U_D) \in \mathcal{S}_j^M] - \mathbb{E}[D \mid Z^M = 0, (\mu(X), U_D) \in \mathcal{S}_j^M]}, \end{aligned}$$

where \mathcal{C}_j is the subset of \mathcal{S}_j for which $D_1(i) = j + 1$.

3 Moving to Opportunity (MTO)

3.1 Program Description

Moving to Opportunity (MTO) was inspired by the promising results of the Gautreaux housing mobility program (Polikoff (2006)). The initial relocation process of the Gautreaux program created a quasi-experiment, and its results indicated housing mobility could be an effective policy. Relative to city movers, suburban movers from Gautreaux were more likely to be employed (Mendenhall et al. (2006)), and the children of suburban movers attended better schools; were more likely to complete high school, attend college, and be employed; and had higher wages than city movers (Rosenbaum (1995)).

MTO was designed to replicate these beneficial effects, offering housing vouchers to eligible households between September 1994 and July 1998 in Baltimore, Boston, Chicago, Los Angeles, and New York (Goering (2003)). Households were eligible to participate in MTO if they were low income, had at least one child under 18, were residing in either public housing or Section 8 project-based housing located in a census tract with a poverty rate of at least 40 percent, were current in their rent payment, and all household members were on the current lease and were without criminal records (Orr et al. (2003)).

Households were drawn from the MTO waiting list through a random lottery. After being drawn, households were randomly allocated into one of three treatment groups. Households in the *experimental* group were offered Section 8 housing vouchers, but were restricted to using them in census tracts with 1990 poverty rates of less than 10 percent. Households in this group were also provided with counseling and education through a local non-profit. After one year had passed, households in the *experimental* group were unrestricted in where they used their Section 8 vouchers. Households in the *Section-8 only* comparison group were provided with no counseling, and were offered Section 8 housing vouchers without any restriction on their place of use. And households in the *control* group continued receiving project-based assistance.⁶

⁶Section 8 vouchers pay part of a tenant's private-market rent. Project-based assistance gives the option

3.2 Data

The first source of data we use in our analysis is the MTO Interim Evaluation. The MTO Interim Evaluation contains variables listing the census tracts in which households lived at both the baseline and in 2002, the time the interim evaluation was conducted. Neighborhood characteristics are measured at the tract level after merging the MTO sample with decennial census data from the National Historical Geographic Information System (NHGIS, Minnesota Population Center (2004)).

3.2.1 Variables

We create an index of neighborhood quality using a linear combination of several neighborhood characteristics. Neighborhood characteristics measured by NHGIS variables are first transformed into percentiles of the national distribution from the 2000 census. Principal components analysis is then used to determine which single vector accounts for the most variation in the national distribution of the poverty rate, the percent with a high school diploma, the percent with a BA or higher, the percent of single-headed households, the male Employment-to-Population Ratio (EPR), and the female unemployment rate.⁷

With respect to interpretation, we think of effects from this index of neighborhood quality as resulting from the neighborhood-level determinants of social interactions, access to resources, and exposure to institutions. In terms of the categories from Manski (1993) and Graham (2018), such effects include correlated effects from common exposure to neighborhood-level factors influencing the labor market and health outcomes of interest. These factors span anything from police interactions and strategies to schools, safety, lead paint, pollution, access to fresh food, transportation, etc. We note that our neighborhood effects will not include all correlated effects due to racial discrimination, such as the stress from exposure to racism in labor markets, the criminal justice system, the health care system, and interpersonal encounters. Another issue is that we cannot distinguish between endogenous effects like those from peers engaging in violence versus correlated effects due to exposure to violence resulting from weak state institutions (Aliprantis (2017b)).

With respect to identification, we note that in our model, each level of neighborhood quality represents homogeneous types of social interactions and resources available at the neighborhood level. We assume agents experience neighborhood quality but cannot influence it because we empirically define neighborhoods as census tracts, which contain about 4,000

of a reduced-rent unit tied to a specific structure.

⁷The coefficients relating each component to the index vector all have magnitudes similar to that of the coefficient for poverty. While poverty is, as expected, negatively correlated with quality, Aliprantis (2017a) shows that there exist neighborhoods that are both low poverty and low quality.

residents on average. If we were thinking about smaller reference groups, such as social groups in classrooms, we would be more interested in incorporating the endogenous formation of reference groups into the model and directly modeling interference (Brock and Durlauf (2007), Manski (2013)). Since our treatment is determined by large reference groups outside of our MTO sample, we believe our partial equilibrium abstraction is appropriate.⁸

Baseline characteristics of the MTO households used in the model include baseline neighborhood quality, whether the respondent had family living in his or her neighborhood of residence, whether a member of the household was a victim of a crime in the previous six months, and whether there were teenage children in the household. Site of residence is the only other observed characteristic included in X ; when models were estimated with additional variables such as the number of children or residence in an early HOPE VI building, the coefficients on the other observables were all statistically insignificant.

Outcome variables for adults from the MTO Interim Evaluation include the labor market status of the adult at the time of the interim survey (ie, two binary variables, one indicating labor force participation, the other indicating whether the adult was employed), the self-reported total household income (all sources summed), the individual earnings in 2001 of the sample adult, receipt of Temporary Assistance for Needy Families (TANF) benefits, and the respondent's BMI. Weights are used in constructing all estimates.⁹

3.2.2 Sample and Descriptive Statistics

The focus of this analysis is adults in the MTO Interim Evaluation sample satisfying two conditions. The first restriction ensures that we are focusing on a relatively homogeneous population. To satisfy this restriction we drop all households living at baseline in a neighborhood above the tenth percentile of the national distribution of quality. These are exceptional households in MTO: The median baseline neighborhood quality for MTO participants was below the first percentile of the national distribution (Figure 1a). For Chicago, Los Angeles, and New York City, nearly all participants lived at baseline in neighborhoods below the 10th percentile of the national distribution. In Baltimore and Boston, however, at


⁸Following the discussion on page 106 of Angrist (2014), our model and definitions of variables serve to distinguish in several ways between the subjects of investigation and the peers who might causally affect them. Each outcome variable is a single raw variable pertaining to the sample of MTO participants. Our treatment variable is an index of multiple variables, each measured in terms of percentiles, and pertains to samples of census tract residents. Discretizing treatment serves to break whatever mechanical relationship might remain between the outcome and treatment variables.

⁹Weights are used for two reasons. First, random assignment ratios varied both from site to site and over different time periods of sample recruitment. Randomization ratio weights are used to create samples representing the same number of people across groups within each site-period. This ensures that neighborhood effects are not conflated with time trends. Second, sampling weights must be used to account for the sub-sampling procedures used during the collection of the interim evaluation data.

baseline a non-trivial share of program participants lived in higher-quality neighborhoods, driven mainly by the male EPR and the share of adults holding a BA in their neighborhoods. These individuals represent a little under 15 percent of the interim evaluation sample and are dropped from our estimation sample.

The second sample restriction facilitates the estimation of the ordered choice model. To satisfy this restriction we top-code neighborhood quality at the median of the national distribution of quality in 2000. Figure 1b shows the final results of these restrictions.

The final estimation sample used in our analysis has approximately 3,100 adults (a little over 85 percent of the interim sample and a little under 75 percent of the original adult sample). Our sample represents “the other one percent.” At baseline, 67 percent of the estimation sample lived in a neighborhood whose quality was below the 1st percentile of the national distribution of neighborhood quality. There was enough mobility in the control group so that by 2002 only 39 percent were living in first percentile neighborhoods. On the other hand, though, the mobility of the control group was not extraordinary. By 2002, about 80 percent of the sample in the control group lived in a neighborhood whose quality was less than the 10th percentile of the national distribution of neighborhood quality.



(a) Baseline Distributions by Site

(b) Final Sample Used in Estimation

Figure 1: Neighborhood Quality in MTO

Note: This figure shows the neighborhood quality distribution of MTO adults in various subsamples.

4 Empirical Specification and Estimation Results

4.1 Ordered Choice Model Specification

When estimating the parameters of the ordered choice model, we define the discrete treatment levels by dividing the estimation sample into its deciles at the time of the interim survey. So $(\underline{q}_k, \bar{q}_k]$ are discrete levels of treatment with values for \underline{q}_k and \bar{q}_k in

$\{0, 0.2, 0.6, 1.1, 1.9, 3.6, 6.1, 11, 19, 33\}$, with $k = 1$ corresponding to $(0, 0.2]$, and $K=10$ corresponding to $(33, 50]$.

The marginal benefit of choosing to move from treatment level k to $k + 1$ in the ordered choice model is specified to be

$$MB_k(i) = \mu(X(i)) + Z(i)W(i)\gamma_k - C_k - V(i)$$

with components

$$\begin{aligned}\mu(X(i)) &= \beta_1 X_1(i) + \dots + \beta_8 X_8(i); \\ \gamma_k &= \Gamma_0 + \Gamma_k; \\ C_k &= \delta_0 + \delta_k.\end{aligned}$$

We can interpret $V(i)$ as the unobserved cost for household i of moving up in the absence of a voucher program, and $W(i)$ as the voucher take up outcome. We allow for these variables to be correlated in an arbitrary way, and assume for the sake of identification that $V(i) \sim \text{iid } N(0, 1)$.¹⁰ While assuming that $V(i)$ follows a normal distribution does impose a functional form, it might be best viewed as a normalization, since the specification of $C(q)$ is flexible. The full empirical specification and likelihood function are specified in Appendix C.

The first four variables in $X(i)$ are baseline neighborhood quality, whether the respondent had family living in the neighborhood of residence, whether a member of the household was a victim of a crime in the previous six months, and whether there were teenage children in the household at baseline. The final four variables in $X(i)$ are site indicators.

Γ_0 's are site-specific fixed effects that capture differences in factors such as local labor and housing markets across sites. We do not attempt to explicitly model housing market prices since these are largely offset by the nature of the payment structure of the rental vouchers and project-based programs. Individuals pay 30 percent of their income toward rent in project-based units, and pay this same rent when using vouchers as long as the price of rent is not above the FMR. Thus, the ability to lease up with a Section 8 voucher is more salient than price.

Like Galiani et al. (2015), we interpret Z as the random assignment of a potential reduction in the cost to accessing a higher-quality neighborhood relative to staying in the baseline neighborhood. It is important to note that secular trends outside the control of the program might swamp this cost reduction. One can imagine changing costs to accessing

¹⁰In the full model we adopt a functional form to learn about the relationship between $V(i)$ and the unobserved components of $W(i)$, but stress that identification of LATEs comes from the data on $W(i)$ and does not require this functional form assumption.

higher-quality neighborhoods due to changes in the local labor or housing markets, changes in school quality due to the provision of magnet schools, or simply an improvement in the quality of the baseline neighborhood.

We estimate the parameters of this ordered choice model via Maximum Likelihood using the log-likelihood function in Appendix C. We then interpolate to identify the continuous functions, and estimate the $V(i)$, as described in Appendix B.

4.1.1 Potential Outcomes Specification

When estimating treatment effect parameters, we define the discrete treatment levels in the intervals $(\underline{q}_j, \bar{q}_j]$ in terms of deciles of the national distribution so that $\underline{q}_j \in \{0, 10, 20, 30, 40\}$, with $j = 1$ corresponding to $(0, 10]$, and $J=5$ corresponding to $(40, 50]$.¹¹

We choose deciles to discretize neighborhood quality when investigating treatment effects not because we believe treatment should have an effect when crossing the particular thresholds of neighborhood quality used in this definition, but because we believe it offers the best balance between theoretical ideal and practical necessity. The model assumes that moves within a given level of treatment will not have effects on outcomes. Even if they do, it is enough to assume that individuals do not select *within* treatment levels based on rich information regarding neighborhood quality.¹² If these assumptions do not hold within entire deciles of quality, the effects from such moves will likely enter the estimation results through the $U_j(i)$. Theoretically, one way to handle this issue would be to increase the number of bins until moves within a given level do not have effects on outcomes. Another way to handle this problem would be to reformulate the model to accommodate a continuous treatment (Florens et al. (2008)).

Because of the limited mobility induced by MTO, we believe deciles of quality are likely to offer the smallest window on which it is feasible to estimate neighborhood effects using the MTO interim survey data. As the next section shows, even this discretization leaves us with an undesirably small sample size of compliers, resulting in noisy estimates providing little power to detect even very large effects. As a result, the only LATEs we attempt to estimate are of moves between $j = 1$ ($q \in (0, 10]$) and $j = 2$ ($q \in (10, 20]$).

¹¹Results in Appendix A allow us to move freely between ordered choice models with different discretizations of quality. This allows us to estimate the choice model on the discretization of quality $\{q_k\}$ that is best for predicting choice, and then, using the inferred continuous model, to translate those estimates to the discretization $\{\underline{q}_j\}$ that is most appropriate for potential outcomes.

¹²This can be seen as a stronger version of the central identifying assumption in Bayer et al. (2008).

4.2 Estimation Results: Ordered Choice Model

The full model we estimate, specified in Appendix C, takes into account the two types of vouchers used in MTO. Parameters corresponding to the experimental group voucher will be super-indexed with M , while those corresponding to the standard Section 8 voucher will be super-indexed by S . The estimated cost function $\widehat{C}(q)$ is shown in Figure 2 together with the cutoffs for treatment in the ordered choice model ($\{\underline{q}_k, \bar{q}_k\}$) and for effects on outcomes ($\{\underline{q}_j, \bar{q}_j\}$). The cost function is estimated to take the expected shape.

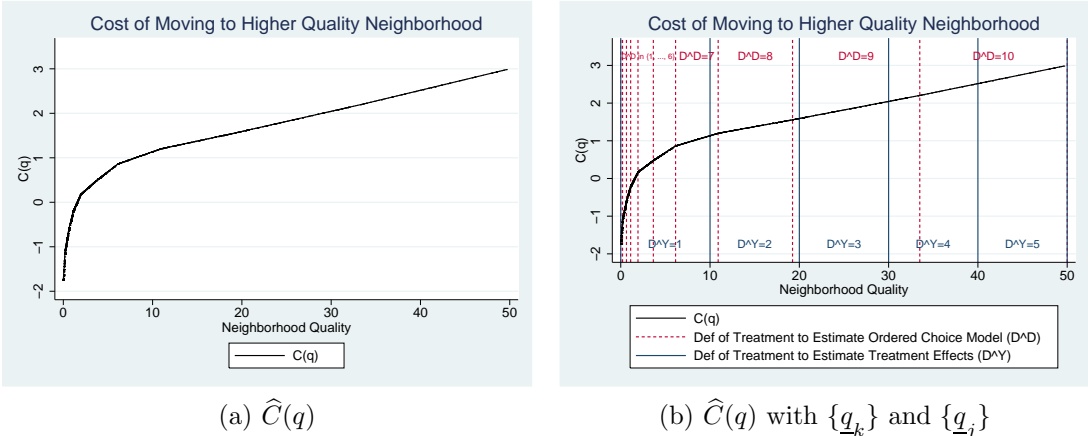


Figure 2: Cost Function

Note: (a) shows the estimated cost function $\widehat{C}(q)$. The dashed red vertical lines in (b) show the cutpoints of quality used to estimate the ordered choice model, $\{\underline{q}_k\}$, which generate the discrete treatment D^D . The cost function $C(q)$ is constructed by linearly interpolating between the cutpoints estimated in an ordered probit model. The solid blue vertical lines in (b) show the cutpoints of quality used in the model of potential outcomes, $\{\underline{q}_j\}$, which generate the discrete treatment D^Y .

Fewer people moved with the MTO voucher than with the unrestricted Section 8 voucher. Fifty-nine percent of our sample moved with the voucher when offered a Section 8 voucher, but only 43 percent of MTO voucher recipients moved with their voucher. However, the MTO voucher was much more effective than the standard Section 8 voucher in getting moving households to access higher-quality neighborhoods. We can see this from the estimated cost reductions $\widehat{\gamma}^S(q)$ and $\widehat{\gamma}^M(q)$ displayed in Figure 3. Furthermore, the effectiveness of both types of vouchers varied considerably by site. Vouchers represented the largest cost reductions in Los Angeles (LA) and New York City (NYC), and represented the smallest cost reductions in Baltimore and Boston. Chicago displays the largest gap in cost reduction between programs.




Figure 3: Cost Reduction from Voucher Use

Note: (a) shows the estimated cost reduction when moving with an experimental MTO voucher, and (b) shows the estimated cost reduction when moving with a standard Section 8 voucher. A higher cost reduction implies that the voucher is more effective at the site in helping residents to move to a higher-quality neighborhood.

Figure 4 shows program participants, color-coded by whether they lived in a neighborhood at the time of the interim evaluation ranked in the first, second, third, or fourth decile of the national distribution of neighborhood quality. On the x -axis is the $\hat{\mu}(x_i)$ of each household, and on the y -axis is the percentile of the household's unobserved determinant of selection in the absence of a program, $\hat{U}_D(i) \equiv F_V(\hat{V}(i)) = \Phi(\hat{V}(i))$. Since vouchers were randomly assigned in MTO, Figures 4a and 4b illustrate counterfactual distributions of neighborhood quality under external manipulations to voucher type. For each household, given observed variables summarized by $\mu(x_i)$ and unobserved variables u_{Di} , these figures show the neighborhood quality households would select into under each setting of the vouchers.

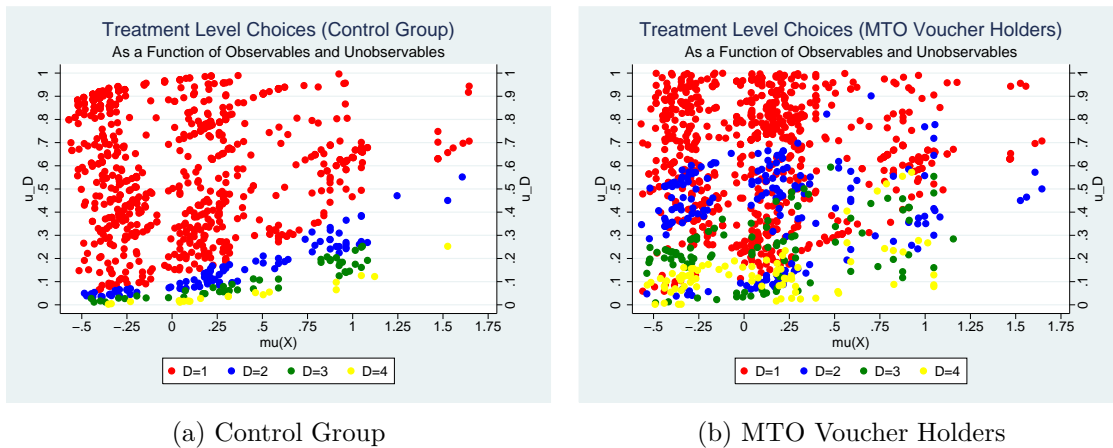


Figure 4: Selection into Treatment

Note: This figure shows how MTO adults sorted into the discrete quality levels $D(i) \in \{1, 2, 3, 4\}$ as a function of their estimated choice model observed characteristics ($\hat{\mu}(x_i)$ - on the x -axis) and unobserved characteristics ($\hat{U}_D(i)$ - on the y -axis). Since these characteristics determine how households would select into $D(i)$ in the absence of any voucher, and vouchers were randomly assigned to households, (b)-(d) characterize counterfactual choices under ideal interventions to voucher status.

In Figure 4a we can see that almost all of the control group remained in low-quality neighborhoods, most remaining in the first decile of neighborhood quality. Only households with very high observed factors $\hat{\mu}(x_i)$ and very low unobserved cost factors $\hat{U}_D(i)$ managed to move to higher quality neighborhoods, even when defined as moving only to the second, third, or fourth deciles of the national distribution.

Figure 4b shows that the MTO voucher induced some households to move to higher quality neighborhoods. Most of this mobility is from the first to the second decile of the national distribution of quality. Although it was still relatively rare, the MTO voucher did induce some households to move into the third and fourth deciles of neighborhood quality.

Table 1 reports the estimated coefficients in $\mu(X)$ of the ordered choice model. There are large differences between sites: In the absence of any voucher program, program participants in Boston tended to live in much higher-quality neighborhoods than their counterparts in the other MTO sites at the time of the interim evaluation. Baltimore and Chicago tended to be similar, while residents in LA and NYC were in the lowest-quality neighborhoods.

Table 1: Ordered Choice Model Parameter Estimates

Baseline Characteristics	$\hat{\beta}_k$	Site Fixed Effects	$\hat{\beta}_k$
Teens in HH	-0.08 (0.05)	Baltimore	0 -
Family in Nbd	-0.14 (0.05)	Boston	0.31 (0.10)
HH Member Victim	0.03 (0.05)	Chicago	-0.04 (0.09)
Baseline Nbd Quality	0.13 (0.01)	Los Angeles	-0.52 (0.10)
		New York City	-0.58 (0.09)

Note: The reported parameters of $\mu(X)$ correspond to the choice model in the absence of any program. The full likelihood is specified and the full set of estimated parameters is reported in Appendix C.

Table 1 in Appendix C shows the parameters of the full model. At the household level, having teens in the household reduces the likelihood of moving with and without a voucher. It is possible that room occupancy restrictions according to age and gender of children may have made it harder for households with older children to find housing. Having no family in the baseline neighborhood makes households more likely to move. Living in a higher-quality neighborhood at baseline increases the likelihood of moving without a voucher, but decreases moving prospects post MTO voucher assignment and counseling. The unobserved

components of V and W^M or V and W^S are only weakly correlated.

4.3 Estimation Results: LATEs of Neighborhood Quality

4.3.1 What Effects Are Identified?

Recalling the counterfactual distributions displayed in Figure 4, there is a range of values of $(\mu(X_i), U_D(i))$ for which households could be induced by receiving an MTO voucher to move from a $D = 1$ quality neighborhood to quality $D = 2$. Since D is not only a function of $\mu(x_i)$ and $U_D(i)$, but also of $W(i)$, households with the same $\mu(x_i)$ and $U_D(i)$ could be in different levels of D depending on whether or not they take up the voucher. There is another range for which households would not move from $D = 1$ (those with high $U_D(i)$), and there are also ranges for which households could be induced to make other moves, such as from $D = 2$ to $D = 3$.

Table 2: Average Neighborhood Characteristics in 2000

Nbd Characteristic	Mean $ D = 1$	Mean $ D = 2$	Unconditional	
			Median	Mean $ D = 10$
Poverty Rate (%)	33	22	9	3
HS Diploma (%)	55	65	83	95
BA (%)	7	11	19	52
Single-Headed HHs (%)	52	38	24	11
Female Unemployment Rate (%)	16	10	5	2
Male EPR ($\times 100$)	55	65	79	89

Owing to the observed patterns of neighborhood selection displayed in Figure 4, we focus on identifying the effects of moving from the first to the second decile of neighborhood quality.¹³ Table 2 characterizes some of the changes in neighborhood characteristics that would typically accompany a move from $D = 1$ to $D = 2$. On average, the poverty rate would decline from 33 to 22 percent, BA attainment would go from 7 to 11 percent, the share of single-headed households would drop from 52 to 38 percent, and the female unemployment rate would drop from 16 to 10 percent. While these changes in neighborhood characteristics are non-trivial, it is worth pointing out that they are still far worse than the unconditional median neighborhood in the US in 2000, and changes of these magnitudes would have to occur several times to achieve the characteristics of the highest-quality neighborhoods. As

¹³We have also estimated Average Causal Responses (ACRs) for subsets in which many possible moves are induced. The results are broadly consistent with our LATE estimates.

discussed in Aliprantis (2017a) and elsewhere, these are moves from the most extreme areas of the left tail of the distribution of quality to neighborhoods that are still within the left tail of quality.¹⁴

4.3.2 For Whom Are Effects Identified?

Appendix B shows how we determine the support of $(\mu(x_i), U_D(i))$ for which LATEs of moving from $D = 1$ to $D = 2$ are identified. The identification support set is

$$\mathcal{S}_{1,2}^M \equiv \left\{ (\mu(x_i), U_D(i)) \mid \mu(x_i) \in [-0.6, 0.4], \quad U_D(i) \in [0.43 + 0.30\mu(x_i), 0.68 + 0.15\mu(x_i)] \right\}.$$

Families in this group select into neighborhood quality $D = 1$ without any voucher, and into neighborhood quality $D = 1$ or $D = 2$ with an MTO voucher assigned. We now include the superscript on \mathcal{S} to indicate that the set is associated with the MTO voucher, and subscripts for both j and $j + 1$ to be clear about the transition with which the set is associated.

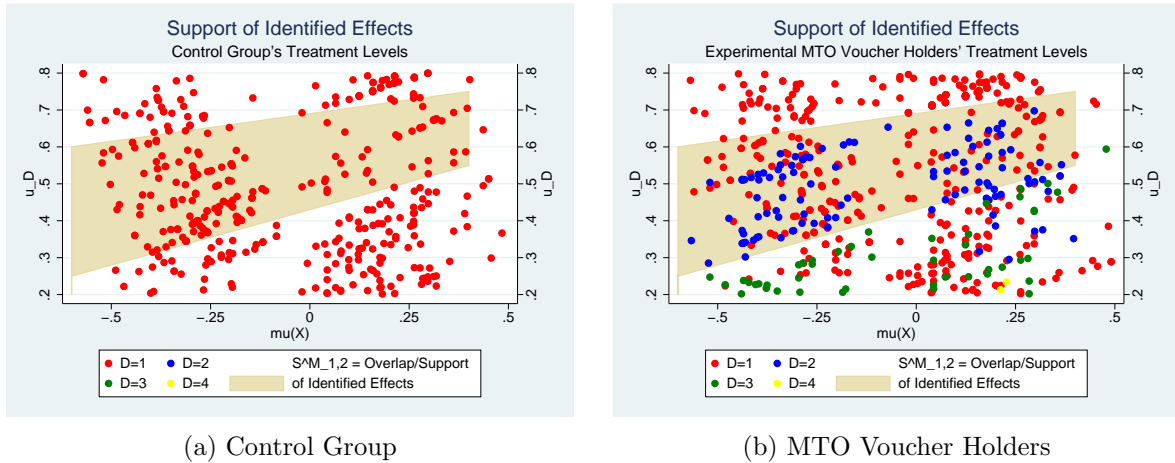


Figure 5: Selection into Treatment and Identification Support Set $\mathcal{S}_{1,2}^M$

Note: This figure shows how MTO adults sorted into the discrete quality levels $D(i) \in \{1, 2, 3, 4\}$ as a function of their estimated choice model observed characteristics $(\hat{\mu}(x_i))$ - on the x -axis and unobserved characteristics $(\hat{U}_D(i))$ - on the y -axis. We now also show the identification of the support set $\mathcal{S}_{1,2}^M$.

4.3.3 LATEs of Neighborhood Quality Estimation Results

Estimates of LATEs of neighborhood quality are reported for the subpopulation in $\mathcal{S}_{1,2}^M$ in Table 3. All of these effects conform with the theory that living in higher-quality neighborhoods improves adult labor market and health outcomes while decreasing receipt of welfare

¹⁴See Clampet-Lundquist and Massey (2008) or Quigley and Raphael (2008) for related discussions. Our effects are from changes in neighborhoods that are comparable to those in Pinto (2014), smaller than those considered in Altonji and Mansfield (2014), and in a different part of the distribution than in Galster et al. (2016).

benefits. We interpret the large standard errors as a result of the small sample size used in estimation, and note that the mechanical relationship generating bias in linear-in-means models does not hold in our model (see Footnote 8.). Since the estimates pertain to the select subgroup induced to move to higher-quality neighborhoods through MTO, they are not generalizable.

Table 3: Adult LATE Estimates

Outcome	$\Delta_{1,2}^{LATE} (Z^M)$	Control Mean in $\mathcal{S}_{1,2}^M$
Labor Market		
In Labor Force (%)	25.8** (18.3)	53.2
Employed (%)	31.2** (20.1)	41.7
Household Income (\$)	5,616* (3,914)	13,506
Earnings (\$)	1,970 (4,066)	7,642
Welfare Benefits		
Received TANF (%)	-40.0*** (19.2)	39.9
Health		
BMI (Raw)	-3.1 (2.8)	30.9

Note: $\Delta_{1,2}^{LATE} (Z^M)$ estimates pertain to individuals with observed and unobserved choice model components in $\mathcal{S}_{1,2}^M \equiv \{(\mu(x_i), U_D(i)) | \mu(x_i) \in [-0.6, 0.4], u_D(i) \in [0.43 + 0.30\mu(x_i), 0.68 + 0.15\mu(x_i)]\}$. Control means are also computed for the subsample in this region and outside of this region (both conditional on $D(i) = 1$). Standard errors are computed using 200 bootstrap replications, with * denoting statistical significance at the (one-tailed) 10% level determined using either the 10th or 90th percentile of the 200 bootstrapped replications, and ** and *** denoting significance at the 5% and 1% levels, respectively.

Point estimates are large, and moving from the first to second levels of neighborhood quality alone is estimated to have statistically significant effects at the 10 percent level or smaller on adult labor market outcomes like labor force participation, employment, and

household income. A statistically significant effect is also found on welfare receipt, which is estimated to be entirely eliminated. Although the estimated effect on BMI is large in magnitude, it is not statistically significant.

MTO has been interpreted as evidence against the theory that living in higher-quality neighborhoods improves adult labor market outcomes. The LATEs in Table 3 stand in contrast to this interpretation.

4.3.4 Falsification Test: Outcomes for Non-Complier Households

As a check on our selection model, and to highlight the difference between our analysis and the program effect approach adopted in most of the literature on MTO (Ludwig et al. (2008), Ludwig et al. (2013), Chetty et al. (2016)), we now use Figure 6 to define a falsification set $\mathcal{F}_{1,2}^M$ for which households would remain in neighborhoods of quality $D = 1$ even if they were assigned an MTO voucher:

$$\mathcal{F}_{1,2}^M \equiv \left\{ (\mu(x_i), U_D(i)) \mid D_0(i) = D_1(i) = 1 \right\}.$$

Household membership in the falsification set is determined not only by the values of $(\mu(x_i), U_D(i))$, but also by the general cost function $C(q)$, the value of the cost reduction function $\gamma^M(q)$ at various levels of quality q , and instrument take up $W^M(i)$. We can see from Figure 6 that restricting households to those with $U_D(i) \in [0.7, 1]$ identifies the falsification set $\mathcal{F}_{1,2}^M$.

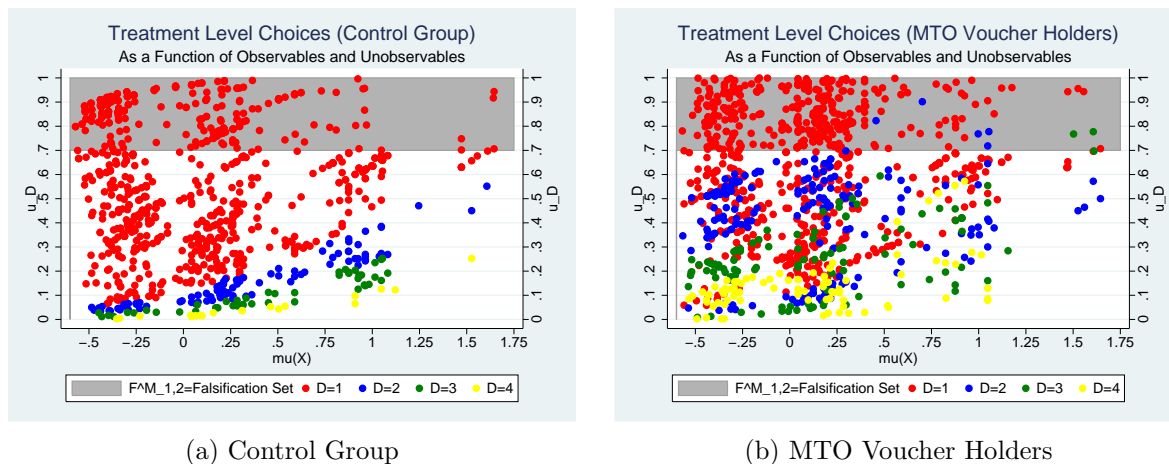


Figure 6: Falsification Set $\mathcal{F}_{1,2}^M$

Note: This figure shows how MTO adults sorted into the discrete quality levels $D(i) \in \{1, 2, 3, 4\}$ as a function of their estimated choice model observed characteristics $(\hat{\mu}(x_i))$ - on the x -axis and unobserved characteristics $(\hat{U}_D(i))$ - on the y -axis. The falsification set $\mathcal{F}_{1,2}^M$ is the group of households identified by the choice model to counterfactually all have $D_0(i) = D_1(i) = 1$.

Neighborhood selection for the subpopulations in $\mathcal{S}_{1,2}^M$ and $\mathcal{F}_{1,2}^M$ is shown in the CDFs in

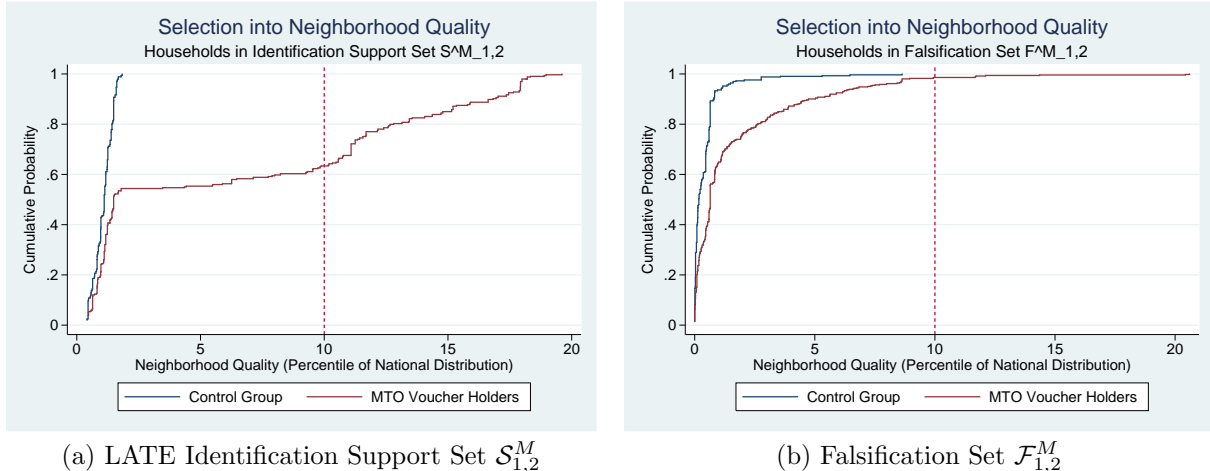


Figure 7: Selection into Neighborhood Quality for Various Subpopulations

Note: This figure shows the distributions of $q_Z(i)$ in (a) for the identification support set $\mathcal{S}_{1,2}^M$ and in (b) for the falsification set $\mathcal{F}_{1,2}^M$. We see that the model does appear able to identify groups whose counterfactual neighborhood selection is quite different, and that selection into the continuous measure of neighborhood quality is consistent with the selection into the discrete measure of neighborhood quality shown earlier.

Figure 7. Figure 7a shows that there is considerable variation in the neighborhood quality selected by the control and MTO voucher holders in the identification support set $\mathcal{S}_{1,2}^M$: No households in the control group selected into $D = 2$, while 37 percent of MTO voucher holders did. Neighborhood selection for households in the falsification set $\mathcal{F}_{1,2}^M$ was quite different: No households in the control group selected into $D = 2$, and only 2 percent of MTO voucher holders selected into $D = 2$. The effects of the MTO program are compared in Table 4 for households in the LATE identification support set $\mathcal{S}_{1,2}^M$ and for households in the falsification set $\mathcal{F}_{1,2}^M$. While receiving an MTO voucher resulted in large improvements to labor force participation rates for households in the group with improvements in neighborhood quality (ie, those in $\mathcal{S}_{1,2}^M$), there was no effect on labor force participation for households in the group with no improvement in neighborhood quality. Employment actually went down for those whose neighborhood quality did not improve, perhaps due to the disruptiveness of moving without the benefits of moving closer to jobs (Weinberg (2000), Andersson et al. (2018)). And while welfare (TANF) receipt and BMI decreased for voucher recipients who did not move to higher quality neighborhoods, this effect was much larger for those who did move to a higher quality neighborhood.

This falsification test helps to illustrate that the effects of the MTO program are not interchangeable with the effects from neighborhood quality. A list of assumptions must be made before translating effects of variation in MTO voucher assignment into effects of variation in neighborhood quality. Our assumptions have been stated explicitly in Sections 2-4; most assumptions in the MTO literature have been made implicitly (Aliprantis (2017a)).

Table 4: Adult Program Effects Estimates by Neighborhood Selection Groups

	Falsification Set $(\mu(x_i), U_D(i)) \in \mathcal{F}_{1,2}^M$ (No Change in Nbd Quality due to MTO)			Identification Support Set $(\mu(x_i), U_D(i)) \in \mathcal{S}_{1,2}^M$ (Some Improvement in Nbd Quality due to MTO)		
	$E[Y Z^M = 1]$	$E[Y Z^M = 0]$	$E[Y Z^M = 1] - E[Y Z^M = 0]$	$E[Y Z^M = 1]$	$E[Y Z^M = 0]$	$E[Y Z^M = 1] - E[Y Z^M = 0]$
Neighborhood Selection						
Nbd Quality ($D \in \{1, 2, 3, 4, 5\}$)	1.02 (0.07)	1.00 (0.00)	0.02 (0.07)	1.37 (0.10)	1.00 (0.09)	0.37 (0.13)
Nbd Quality ($q \in [0, 50]$)	1.7 (0.8)	0.4 (0.1)	1.2 (0.8)	6.4 (1.0)	1.1 (0.9)	5.3 (1.3)
Labor Market						
In Labor Force (%)	63.6 (3.9)	63.6 (5.4)	0.0 (6.4)	63.0 (3.2)	53.2 (3.8)	9.8 (4.7)
Employed (%)	47.1 (4.2)	53.6 (5.4)	-6.5 (6.8)	53.5 (3.3)	41.7 (3.9)	11.8 (5.0)
Household Income (\$)	14,252 (924)	14,134 (998)	119 (1,366)	15,629 (847)	13,506 (883)	2,123 (1,175)
Earnings (\$)	7,583 (914)	8,554 (992)	-971 (1,375)	8,364 (611)	7,642 (767)	722 (917)
Welfare Benefits						
Received TANF (%)	32.2 (3.7)	33.7 (5.0)	-1.5 (6.6)	24.9 (3.0)	39.9 (3.4)	-15.0 (4.6)
Health						
BMI (Raw)	30.0 (0.5)	30.4 (0.8)	-0.3 (1.0)	29.7 (0.5)	30.9 (0.5)	-1.2 (0.7)

Note: The first three columns of this table report the effects of receiving an experimental MTO voucher for households predicted by the estimated choice model to reside in a low-quality neighborhood even when receiving an MTO voucher. The last three columns report the effects of receiving an experimental MTO voucher for households predicted by the estimated choice model to potentially move to a higher quality neighborhood when receiving an MTO voucher.

5 Conclusion

Because households endogenously sort into neighborhoods, identifying the causal effects of neighborhood environments has proven to be a substantial challenge. The Moving to Opportunity (MTO) housing mobility experiment gave households living in high-poverty neighborhoods in five US cities the ability to enter a lottery for housing vouchers to be used in low-poverty neighborhoods. The results from MTO have led to surprising and controversial inferences about neighborhood effects.

This paper identified neighborhood effects in MTO using a methodological innovation allowing for the identification of transition-specific effects with a binary instrument. We found that moving to a higher-quality neighborhood had large, positive effects on employment, labor force participation, household income, and welfare receipt. Although effects on individual earnings and BMI were not statistically significant, these effects were also estimated to be large and positive. Due to the limited changes in neighborhood quality induced by MTO, these LATE estimates pertain to about nine percent of our estimation sample. We found no evidence from MTO against the theory that increasing neighborhood quality improves adult outcomes.

We think it is important to interpret our results while being conscious of the difficulty of conducting controlled, and not just randomized, experiments in social settings (Deaton (2010)). Programs and neighborhood changes differentially impacting the treatment and control groups will interfere with the ability to identify neighborhood effects using voucher assignment as an instrument. Ideally, the voucher assignment would induce a change in the cost of moving, holding all else equal. However, households may have responded to their group assignment, and baseline neighborhood conditions may well have changed during the multiple-year period between the decision to move and the time of the interim evaluation when outcomes were measured. For example, households assigned to the control group might have responded by applying for Section 8 vouchers on their own outside of the MTO program (Orr et al. (2003)). And according to de Souza Briggs et al. (2010), during the implementation of MTO, the Jobs-Plus program saturated public housing developments with state-of-the-art employment, training, and child care services, while providing rent incentives to encourage employment. In addition, the US enacted major welfare reform legislation in August 1996, precisely while MTO vouchers were being assigned (Blank (2002)).

We also think that our results support MTO-like policy innovations. The fact that households were more likely to move with Section 8 vouchers than MTO vouchers does not imply that Section 8 vouchers are preferable to MTO vouchers. Changes in neighborhood quality were much smaller for Section 8 movers than for MTO movers, and variation by site was

large. Since only about a quarter of eligible households are currently able to obtain a Section 8 housing voucher (Sard and Fischer (2012)), an area for future research is understanding what changes in voucher policy might optimize the extent to which households are able to realize positive neighborhood effects through the subsidy, and which of these policies might be feasible to implement (McClure (2010), Collinson and Ganong (2017)).

References

- Aliprantis, D. (2017a). Assessing the evidence on neighborhood effects from Moving to Opportunity. *Empirical Economics* 52(3), 925–954.
- Aliprantis, D. (2017b). Human capital in the inner city. *Empirical Economics* 53(3), 1125–1169.
- Aliprantis, D. and D. Kolliner (2015). Neighborhood poverty and quality in the Moving to Opportunity experiment. *Federal Reserve Bank of Cleveland Economic Commentary*.
- Altonji, J. and R. Mansfield (2014). Estimating school and neighborhood effects using sorting on observables to control for sorting on unobservables. *Mimeo., Cornell University*.
- Andersson, F., J. C. Haltiwanger, M. J. Kutzbach, H. O. Pollakowski, and D. H. Weinberg (2018). Job displacement and the duration of joblessness: The role of spatial mismatch. *The Review of Economics and Statistics*. Forthcoming.
- Angrist, J. D. (2014). The perils of peer effects. *Labour Economics* 30, 98–108.
- Angrist, J. D. and G. W. Imbens (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American Statistical Association* 90(430), 431–442.
- Athey, S. and G. W. Imbens (2017). The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Perspectives* 31(2), 3–32.
- Bayer, P., S. L. Ross, and G. Topa (2008). Place of work and place of residence: Informal hiring networks and labor market outcomes. *Journal of Political Economy* 116(6), 1150–1196.
- Blank, R. M. (2002). Evaluating welfare reform in the United States. *Journal of Economic Literature* 40(4), 1105–1166.

- Brock, W. and S. Durlauf (2007). Identification of binary choice models with social interactions. *Journal of Econometrics* 140(1), 52–75.
- 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.
- Chyn, E. (2015). Moved to Opportunity: The long-run effect of public housing demolition on labor market outcomes of children. *Mimeo., University of Michigan*.
- Clampet-Lundquist, S. and D. S. Massey (2008). Neighborhood effects on economic self-sufficiency: A reconsideration of the Moving to Opportunity experiment. *American Journal of Sociology* 114(1), 107–143.
- Collinson, R. and P. Ganong (2017). How do changes in housing voucher design affect rent and neighborhood quality? *American Economic Journal: Economic Policy*. Forthcoming.
- Davis, M. A., J. Gregory, D. A. Hartley, and K. Tan (2016). Neighborhood choices and neighborhood effects. *Federal Reserve Bank of Chicago*.
- de Souza Briggs, X., S. J. Popkin, and J. Goering (2010). *Moving to Opportunity: The Story of an American Experiment to Fight Ghetto Poverty*. Oxford University Press.
- Deaton, A. (2010). Instruments, randomization, and learning about development. *Journal of Economic Literature* 48(2), 424–455.
- Florens, J.-P., J. J. Heckman, C. Meghir, and E. Vytlacil (2008). Identification of treatment effects using control functions in models with continuous, endogenous treatment and heterogeneous effects. *Econometrica* 76(5), 1191–1206.
- Galiani, S., A. Murphy, and J. Pantano (2015). Estimating neighborhood choice models: Lessons from a housing assistance experiment. *The American Economic Review* 105(11), 3385–3415.
- Galster, G., A. Santiago, L. Stack, and J. Cutsinger (2016). Neighborhood effects on secondary school performance of Latino and African American youth: Evidence from a natural experiment in Denver. *Journal of Urban Economics* 93, 30 – 48.
- Galvez, M. M., J. Simington, and M. Treskon (2010). *Moving to Work and Neighborhood Opportunity: A Scan of Mobility Initiatives by Moving to Work Public Housing Authorities*. Washington, DC: What Works Collaborative.

- Goering, J. (2003). The impacts of new neighborhoods on poor families: Evaluating the policy implications of the Moving to Opportunity demonstration. *Economic Policy Review* 9(2).
- Graham, B. S. (2018). Identifying and estimating neighborhood effects. *Journal of Economic Literature*. Forthcoming.
- Heckman, J. J., S. Urzúa, and E. Vytlacil (2006). Understanding Instrumental Variables in models with essential heterogeneity. *The Review of Economics and Statistics* 88(3), 389–432.
- Imbens, G. W. and J. D. Angrist (1994). Identification and estimation of Local Average Treatment Effects. *Econometrica* 62(2), 467–475.
- Kling, J. R., J. B. Liebman, and L. F. Katz (2007). Experimental analysis of neighborhood effects. *Econometrica* 75(1), 83–119.
- Ludwig, J., G. J. Duncan, L. A. Gennetian, L. F. Katz, R. C. Kessler, J. R. Kling, and L. Sanbonmatsu (2013). Long-term neighborhood effects on low-income families: Evidence from Moving to Opportunity. *American Economic Review* 103(3), 226–231.
- Ludwig, J., J. B. Liebman, J. R. Kling, G. J. Duncan, L. F. Katz, R. C. Kessler, and L. Sanbonmatsu (2008). What can we learn about neighborhood effects from the Moving to Opportunity experiment? *American Journal of Sociology* 114(1), 144–188.
- Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies* 60(3), pp. 531–542.
- Manski, C. F. (2013). Identification of treatment response with social interactions. *The Econometrics Journal* 16(1), S1–S23.
- McClure, K. (2010). The prospects for guiding housing choice voucher households to high opportunity neighborhoods. *Cityscape* 12(3), 101–122.
- Mendenhall, R., S. DeLuca, and G. Duncan (2006). Neighborhood resources, racial segregation, and economic mobility: Results from the Gautreaux program. *Social Science Research* 35(4), 892–923.
- Minnesota Population Center (2004). *National Historical Geographic Information System* (Pre-release Version 0.1 ed.). Minneapolis, MN: University of Minnesota.

- Orr, L. L., J. D. Feins, R. Jacob, E. Beecroft, L. Sanbonmatsu, L. F. Katz, J. B. Liebman, and J. R. Kling (2003). *Moving to Opportunity: Interim Impacts Evaluation*. Washington, DC: US Department of Housing and Urban Development, Office of Policy Development and Research.
- Phillips, D. C. (2017). Landlords avoid tenants who pay with vouchers. *Economics Letters* 151, 48–52.
- Pinto, R. (2014). Selection bias in a controlled experiment: The case of Moving to Opportunity. *Mimeo., University of Chicago*.
- Polikoff, A. (2006). *Waiting for Gautreaux*. Northwestern University Press.
- Quigley, J. M. and S. Raphael (2008). Neighborhoods, economic self-sufficiency, and the MTO program. *Brookings-Wharton Papers on Urban Affairs* 8(1), 1–46.
- Rosenbaum, J. E. (1995). Changing the geography of opportunity by expanding residential choice: Lessons from the Gautreaux program. *Housing Policy Debate* 6(1), 231–269.
- Sampson, R. J. (2008). Moving to inequality: Neighborhood effects and experiments meet social structure. *American Journal of Sociology* 114(1), 189–231.
- Sanbonmatsu, L., J. R. Kling, G. J. Duncan, and J. Brooks-Gunn (2006). Neighborhoods and academic achievement: Results from the Moving to Opportunity experiment. *The Journal of Human Resources* 41(4), 649–691.
- Sard, B. and W. Fischer (2012, July). Renters’ tax credit would promote equity and advance balanced housing policy. *Center on Budget and Policy Priorities Report*.
- Shroder, M. D. and L. L. Orr (2012). Moving to Opportunity: Why, how, and what next? *Cityscape* 14(2), 31–56.
- Vytlacil, E. (2006). Ordered discrete-choice selection models and local average treatment effect assumptions: Equivalence, nonequivalence, and representation results. *The Review of Economics and Statistics* 88(3), 578–581.
- Weinberg, B. A. (2000). Black residential centralization and the spatial mismatch hypothesis. *Journal of Urban Economics* 48(1), 110–134.
- Wilson, W. J. (1987). *The Truly Disadvantaged: The Inner City, the Underclass, and Public Policy*. University of Chicago.