working paper 15 06

Assessing the Evidence on Neighborhood Effects from Moving to Opportunity

Dionissi Aliprantis

FEDERAL RESERVE BANK OF CLEVELAND

Working Paper 15-06 March 2015*

Assessing the Evidence on Neighborhood Effects from Moving to Opportunity

Dionissi Aliprantis

The Moving to Opportunity (MTO) experiment randomly assigned housing vouchers that could be used in low-poverty neighborhoods. Consistent with the literature, I find that receiving an MTO voucher had no effect on outcomes like earnings, employment, and test scores. However, after studying the assumptions identifying neighborhood effects with MTO data, this paper reaches a very different interpretation of these results than found in the literature. I first specify a model in which the absence of effects from the MTO program implies an absence of neighborhood effects. I present theory and evidence against two key assumptions of this model: That poverty is the only determinant of neighborhood quality, and that outcomes only change across one threshold of neighborhood quality. I then show that in a more realistic model of neighborhood effects that relaxes these assumptions, the absence of effects from the MTO program is perfectly compatible with the presence of neighborhood effects. This analysis illustrates why the implicit identification strategies used in the literature on MTO can be misleading.

Keywords: Moving to Opportunity, Neighborhood Effect, Program Effect.

JEL Codes: C30, H50, I38, J10, R00.

Suggested citation: Aliprantis, Dionissi, 2015. "Assessing the Evidence on Neighborhood Effects from Moving to Opportunity," Federal Reserve Bank of Cleveland, working paper no. 15-06.

Dionissi Aliprantis is at the Federal Reserve Bank of Cleveland (dionissi.aliprantis@clev.frb.org). The author thanks Francisca G.-C. Richter for many helpful conversations and Jeffrey Kling, Joel Elvery, Becka Maynard, Juan Pantano, Ruby Mendenhall, Ramón García-Cobían, Jon James, Subhra Saha, Shawn Rohlin, Jason Seligman, Eugenio Peluso, Tim Dunne, Bruce Fallick, Daniel Carroll, Daniel Hartley, Nobuyuki Hanaki, Rick Mansfield, Susan Clampet-Lundquist, Douglas Massey, my Math Corps students, seminar participants at Aix Marseille School of Economics, the Cleveland Fed, the 2011 Federal Reserve System Applied Micro Conference, Ohio State (Glenn School), CSU, Akron, and several anonymous referees for contributing to this paper. Mary Zenker provided valuable research assistance, and Paul Joice at HUD has been extremely helpful. The research reported here was supported in part by the Institute of Education Sciences, U.S. Department of Education, through Grant R305C050041-05 to the University of Pennsylvania.

*Note: This paper first appeared as working paper 11-01 in January 2011. Later versions appeared as 11-22, 11-22R, 12-33, and 12-33R. The last revision of the previous version, 12-33R, was posted in September 2014.

1 Introduction

Understanding neighborhood effects is an imperative for public policy. For example, debates about the role of government in education cannot be resolved without understanding the nature of effects from localized differences in resources and social interactions (Friedman (1955), Manski (2013b)). Likewise, empirically characterizing neighborhood effects is crucial for understanding the persistence of racial inequality in the United States and for designing effective policy in response (Wilson (1987), Sampson (2012)).

Conclusive evidence on neighborhood effects is elusive, though, since spatial correlations in outcomes could reflect residential sorting as easily as neighborhood effects. To overcome this fundamental selection issue, researchers have studied housing mobility programs like Gautreaux, which relocated 7,100 public housing families throughout Chicago in a quasi-random manner between 1976 and 1998 (Polikoff (2006)). The results from Gautreaux have been interpreted as strong evidence of neighborhood effects: Those who moved to high-income, white-majority suburbs through Gautreaux had much better education and labor market outcomes than those who moved to segregated city neighborhoods (Rubinowitz and Rosenbaum (2000), Rosenbaum (1995), Mendenhall et al. (2006)).

The Moving to Opportunity (MTO) housing mobility program was designed to replicate the success of Gautreaux by randomly allocating housing vouchers to public housing residents in five US cities between 1994 and 1998. In a tremendous disappointment, the results from the MTO program were not as positive as the results from the Gautreaux program. Education and labor market outcomes did not improve (Sanbonmatsu et al. (2006), Kling et al. (2007a)), and the risky behavior of young males actually grew worse (Kling et al. (2005)).

The majority of the literature has interpreted the results from MTO as evidence against neighborhood effects. For example, Ludwig et al. (2013) interpret the results from the MTO program as being "Contrary to the widespread view that living in a disadvantaged inner-city neighborhood depresses labor market outcomes, ..." (p 228). Angrist and Pischke (2010)'s interpretation of MTO is that "The program has produced surprising and influential evidence weighing against the view that neighborhood effects are a primary determinant of low earnings by the residents of poor neighborhoods" (p 4).

Interpreting MTO as evidence against neighborhood effects has previously come under criticism for conflating program effects with neighborhood effects (Clampet-Lundquist and Massey (2008)). However, this critique has been dismissed as reflecting a misunderstanding of selection bias (Ludwig et al. (2008)). The literature continues to interpret MTO as an experiment that randomly allocated households to varying peer environments because housing vouchers were randomly assigned (Angrist (2014)).

This paper shows that the distinction made in Clampet-Lundquist and Massey (2008) between program effects and neighborhood effects is in fact critical to assessing the evidence

on neighborhood effects from MTO. To make the issues clear, one must first consider a standard joint model of potential outcomes and selection into treatment and note the following: Defining treatment as moving with an MTO voucher generates a model of program effects, while defining treatment as moving to a high-quality neighborhood generates a model of neighborhood effects.

I ask a question that follows from Clampet-Lundquist and Massey (2008)'s analysis: What model of neighborhood effects can be used to justify the view in the literature that "If neighborhood environments affect behavior ... then these neighborhood effects ought to be reflected in ITT and TOT impacts [of the program] on behavior" (Ludwig et al. (2008), pp 181-182)? By investigating this question, I not only distinguish between program and neighborhood effects, but also establish assumptions about models of neighborhood effects under which researchers can use program effects to learn about neighborhood effects. I find that these assumptions are strong, have led the literature to draw unwarranted conclusions from the MTO results, and can be relaxed by directly estimating a neighborhood effects model.

Put a bit more precisely, suppose that Y is an outcome variable like employment, D is neighborhood quality, Z is receipt of a housing voucher, and consider a model of neighborhood effects consisting of potential outcomes $Y(D)$ and $D(Z)$. Randomization of a housing voucher $Z \in \{0,1\}$ identifies a class of program effects, the potential outcomes $Y(Z)$ and $D(Z)$. A central contribution of Clampet-Lundquist and Massey (2008) was to make a distinction between program effects $Y(Z)$ and neighborhood effects $Y(D)$.

This paper asks the further question: What definition of D and resulting assumptions about $Y(D)$ allow us to draw conclusions about neighborhood effects from program effects? Different specifications of D, such as $D \in \{0,1\}$ or $D \in \{1,2,\ldots,J\}$, generate different models. These different models make distinct assumptions about how changing neighborhood characteristics affect outcomes. I show two sufficient assumptions for learning about neighborhood effects from program effects are that neighborhood quality is a binary variable, and that poverty is a proxy for quality. The resulting specification of potential outcomes $Y(D)$ imposes that the outcome variable Y changes only in response to crossing a single threshold of neighborhood poverty.

In more general models of neighborhood effects that relax these assumptions, it is entirely possible that neighborhood environments affect behavior but that these neighborhood effects are not reflected in the effects of the MTO program. I provide empirical evidence and theoretical arguments in favor of adopting a more general model of this type. I first show that outcomes in the model should be allowed to change across more than just one margin of quality. I then show that neighborhood quality should be defined as a function of other characteristics in addition to poverty. In order to conduct my empirical analysis, I use principal component analysis to construct a scalar measure of neighborhood quality that is a function of not only the neighborhood poverty rate, but also the percent with high school degrees, the percent

with BAs, the percent of single-headed households, the male employment-to-population ratio, and the female unemployment rate.[1]

I first provide theory and evidence in favor of adopting a model with more than two levels of quality. I show that MTO only induced transitions across low levels of neighborhood quality. As a result, MTO did not generate the variation in neighborhood quality necessary to learn whether changes to many types of neighborhood environments would alter outcomes. In other words, the neighborhood effects model with only two levels of quality implicitly used in the literature on MTO simply assumes that changes to many types of neighborhood environments would not alter outcomes.

I also provide empirical evidence against using poverty as a proxy for quality in the MTO experiment. I show that there are many low-poverty neighborhoods in MTO states that are still low-quality. Thus, even when focused on understanding effects from moves across low levels of quality, researchers must still be careful to identify the moves induced by MTO that changed neighborhood quality in conjunction with neighborhood poverty.

The paper proceeds as follows: Section 2 describes the MTO experiment. Section 3 characterizes the current literature on MTO in terms of the neighborhood effects model assumptions it implicitly imposes. Subsequent sections present theoretical reasoning and empirical evidence on these assumptions and how they might be relaxed. Section 4 presents a canonical joint model of potential outcomes and selection into treatment without any view of how such a model might be applied to MTO. Sections 5.1 and 5.2 then proceed, respectively, to discuss the program and neighborhood effects identified with the MTO data set under various assumptions. Section 6 concludes.

2 Moving To Opportunity (MTO)

Moving To Opportunity (MTO) was inspired by the promising results of the Gautreaux program. Following a class-action lawsuit led by Dorothy Gautreaux, in 1976 the Supreme Court ordered the Department of Housing and Urban Development (HUD) and the Chicago Housing Authority (CHA) to remedy the extreme racial segregation experienced by public-housing residents in Chicago. One of the resulting programs gave families awarded Section 8 public housing vouchers the ability to use them beyond the territory of CHA, giving families the option to be relocated either to suburbs that were less than 30 percent black or to black neighborhoods in the city that were forecast to undergo "revitalization" (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)),

[1] My measure of quality is a normalization of the first principal component of these variables, or the one-dimensional vector explaining the most variation in these variables.

and the children of suburban movers attended better schools, were more likely to complete high school, attend college, be employed, and had higher wages than city movers (Rosenbaum (1995)).[2]

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%, were current in their rent payment, and all families members were on the current lease and were without criminal records (Orr et al. (2003)).

Families were drawn from the MTO waiting list through a random lottery. After being drawn, families were randomly allocated into one of three treatment groups. The *experimental* group was offered Section 8 housing vouchers, but were restricted to using them in census tracts with 1990 poverty rates of less than 10 percent. However, after one year had passed, families in the *experimental* group were then unrestricted in where they used their Section 8 vouchers. Families in this group were also provided with counseling and education through a local non-profit. Families 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 families in the *control* group received project-based assistance.[3]

3 What Model of Neighborhood Effects Can Justify the Literature's Current Interpretation of MTO?

Program effects and neighborhood effects are different parameters defined in distinct models (Heckman (2010)). Yet ITT and TOT effects from receiving an MTO voucher have been interpreted as evidence on neighborhood effects in the literature on MTO. For example, Kling et al. (2007a) include ITT and TOT program effect estimates as "direct evidence on the existence, direction, and magnitude of neighborhood effects" (p. 84), and Ludwig et al. (2008) contend that "Both [ITT and TOT] estimators are informative about the existence of neighborhood effects on behavior" (p. 146).

What model of neighborhood effects can justify these statements? The current interpretation of the results from MTO does not equate program and neighborhood effects, but rather combines evidence on program effects from MTO together with logical arguments to indirectly

[2]It has also been found that suburban movers have much lower male youth mortality rates Votruba and Kling (2009) and tend to stay in high-income suburban neighborhoods many years after their initial placement (DeLuca and Rosenbaum (2003), Keels et al. (2005)).

[3]Section 8 vouchers pay part of a tenant's private market rent. Project-based assistance gives the option of a reduced-rent unit tied to a specific structure.

draw conclusions about neighborhood effects.[4] This Section shows that such an interpretation of MTO relies on an implicit, and therefore poorly-specified, model of neighborhood effects.

Suppose we were only focused on comparing the MTO experimental and control groups, and that for the sake of exposition we are focused on the single outcome of adult employment.[5] The following statement:

(†): "If neighborhood environments affect behavior... then these neighborhood effects ought to be reflected in ITT and TOT impacts [of the program] on behavior" (Ludwig et al. (2008), pp 181-182).

can be justified by a model of potential outcomes $D(Z)$, $Y(D)$, and $Y(Z)$ under the assumptions that D is a binary indicator of neighborhood quality, Z is a binary indicator of receiving an MTO voucher versus being in the control group, and Y is a binary indicator of employment:

M1: $D_i \equiv \mathbf{1}\{$ individual i lives in a high-quality neighborhood $\}$

M2: $Z_i \equiv \mathbf{1}\{$ individual i received an MTO voucher $\}$

M3: $Y_i \equiv \mathbf{1}\{$ individual i is employed $\}$

Note that treatment is defined here in terms of neighborhood quality, whereas most of the literature on MTO estimates models in which treatment is defined as moving with an MTO voucher.[6] It is important to distinguish between these definitions of treatment because they generate distinct models of potential outcomes and selection, with one being a model of program effects (D1), and the other being a model of neighborhood effects (D2):

D1 Treatment is moving with the aid of the program (ie, using an MTO voucher).

D2 Treatment is moving to a high-quality neighborhood.

Without any further empirical or theoretical restrictions on maintained assumptions M1-M3, these variables result in a neighborhood effects model that can generate any of $4^3 = 64$ possible counterfactual worlds displayed in Table 1.

To gain some intuition about the worlds in Table 1, consider Worlds 22 and 32. In both World 22 and 32 there are program effects on individual i's neighborhood quality. Columns 1 and 2 indicate that the individual would move to a "good" neighborhood when receiving a voucher, but would remain in a low-quality neighborhood without a voucher.

[4] This is the author's current interpretation of the literature, most prominently represented by Kling et al. (2007a) and Ludwig et al. (2008). However, the distinction between program and neighborhood effect parameters has not always been made clearly. Some studies do seem to equate program effects with neighborhood effects, even when using this indirect logic. Early examples where this distinction is unclear are Ludwig et al. (2001) and Kling et al. (2005), and more recent examples include Ludwig et al. (2013), Sanbonmatsu et al. (2012), and Gennetian et al. (2012).

[5] That is, we abstract from the Section 8 voucher group for the sake of exposition.

[6] See the Appendix of Ludwig et al. (2008) or Ludwig et al. (2013) for examples.

Worlds 22 and 32 differ according to the presence of program and neighborhood effects on individual i's employment. Columns 5 and 6 indicate that in World 22 the individual would have a job with a voucher, but not without a voucher. And Columns 3 and 4 indicate that in World 22 individual i would have a job when living in a "good" neighborhood, but not when in a "bad" neighborhood. In contrast, World 32 is characterized by no program effects on employment (Columns 5 and 6) and no neighborhood effects on employment (Columns 3 and 4).

Table 1: Counterfactual Worlds Possible in Unrestricted Nbd Effects Model with Binary Variables

Definitions:

Z ≡ "Individual i receives an MTO voucher."
D ≡ "Individual i lives in a 'good' neighborhood."
Y ≡ "Individual i is employed."

Truth Table

Column	1	2	3	4	5	6
Row	$D(Z=1)$	$D(Z=0)$	$Y(D=1)$	$Y(D=0)$	$Y(Z=1)$	$Y(Z=0)$
(World 1)	1	1	1	1	1	1
(World 2)					1	0
(World 3)					0	1
(World 4)					0	0
(World 5)			1	0	1	1
(World 6)					1	0
(World 7)					0	1
(World 8)					0	0
(World 9)			0	1	1	1
(World 10)					1	0
(World 11)					0	1
(World 12)					0	0
(World 13)			0	0	1	1
(World 14)					1	0
(World 15)					0	1
(World 16)					0	0
(World 17)	1	0	1	1	1	1
(World 18)					1	0
(World 19)					0	1
(World 20)					0	0
(World 21)			1	0	1	1
(World 22)					1	0
(World 23)					0	1
(World 24)					0	0
(World 25)			0	1	1	1
(World 26)					1	0
(World 27)					0	1
(World 28)					0	0
(World 29)			0	0	1	1
(World 30)					1	0
(World 31)					0	1
(World 32)					0	0
(World 33)	0	1	1	1	1	1
(World 34)					1	0
(World 35)					0	1
(World 36)					0	0
(World 37)			1	0	1	1
(World 38)					1	0
(World 39)					0	1
(World 40)					0	0
(World 41)			0	1	1	1
(World 42)					1	0
(World 43)					0	1
(World 44)					0	0
(World 45)			0	0	1	1
(World 46)					1	0
(World 47)					0	1
(World 48)					0	0
(World 49)	0	0	1	1	1	1
(World 50)					1	0
(World 51)					0	1
(World 52)					0	0
(World 53)			1	0	1	1
(World 54)					1	0
(World 55)					0	1
(World 56)					0	0
(World 57)			0	1	1	1
(World 58)					1	0
(World 59)					0	1
(World 60)					0	0
(World 61)			0	0	1	1
(World 62)					1	0
(World 63)					0	1
(World 64)					0	0

We could combine theory and empirical observations to rule out that our world as observed in the MTO data looked like some of the possible counterfactual worlds in Table 1. For example, based on empirical observations from MTO on the neighborhoods of residence of control group households as recorded at the time of the follow-up survey, it is likely to be uncontroversial that we can rule out $D(Z = 0) = 1$, or living in a "good" neighborhood without a voucher, in the real world. This would eliminate Worlds 1-16 or 33-48 from representing the real world, leaving the 32 counterfactual worlds displayed in Table 2 in consideration for accurately describing the world as observed in MTO.

Table 2: Counterfactual Worlds Possible in Empirically-Restricted Nbd Effects Model with Binary Variables

After Restrictions Imposed by Empirical Observations

Definitions:

Z $\equiv$ "Individual i receives an MTO voucher."
D $\equiv$ "Individual i lives in a 'good' neighborhood."
Y $\equiv$ "Individual i is employed."

Truth Table

Column	1	2	3	4	5	6
Row	$D(Z=1)$	$D(Z=0)$	$Y(D=1)$	$Y(D=0)$	$Y(Z=1)$	$Y(Z=0)$
(World 17)	1	0	1	1	1	1
(World 18)					1	0
(World 19)					0	1
(World 20)					0	0
(World 21)			1	0	1	1
(World 22)					1	0
(World 23)					0	1
(World 24)					0	0
(World 25)			0	1	1	1
(World 26)					1	0
(World 27)					0	1
(World 28)					0	0
(World 29)			0	0	1	1
(World 30)					1	0
(World 31)					0	1
(World 32)					0	0
(World 49)	0	0	1	1	1	1
(World 50)					1	0
(World 51)					0	1
(World 52)					0	0
(World 53)			1	0	1	1
(World 54)					1	0
(World 55)					0	1
(World 56)					0	0
(World 57)			0	1	1	1
(World 58)					1	0
(World 59)					0	1
(World 60)					0	0
(World 61)			0	0	1	1
(World 62)					1	0
(World 63)					0	1
(World 64)					0	0

So far our approach to relating program effects and neighborhood effects has only used empirical observations in addition to binary definitions of variables to rule out counterfactual worlds. It would be possible to further rule out of consideration some of the worlds from Table 2 solely on the basis of theory. One possibility would be to adopt the neighborhood effects model shown in Figure 1 below, along with the new model of neighborhood effects resulting from the MTO intervention.[7]

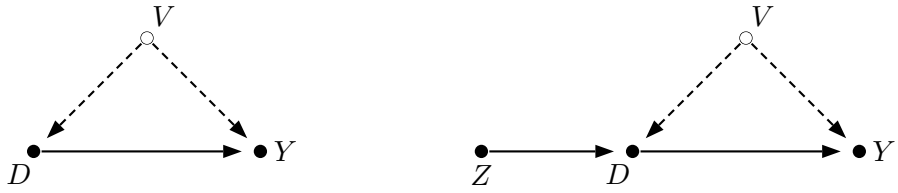

(a) The Neighborhood Effects Model (b) The MTO Intervention to the Nbd Effects Model

Figure 1: Directed Acyclic Graphs of the Neighborhood Effects Model

We could apply this neighborhood effects model to rule out particular counterfactual worlds from consideration. For example, we could rule out counterfactual Worlds 18, 19, and 20 as simply being inconsistent with the types of counterfactual worlds we believe are possibly similar to our own, as expressed by the restrictions on the Data Generating Process placed by our model.[8] We can proceed to eliminate counterfactual worlds from Table 2, with the worlds dropped from Table 3 all following the same pattern of elimination: They either contradict empirical observation, require that the MTO voucher affects outcomes through some pathway other than neighborhood quality, or else they would require some column to take different values in order to be consistent with our model.

[7]Figure 1 follows the convention from Pearl (2009) of communicating that a variable is observed by drawing a solid line to its descendants, and communicate that a variable is unobserved by drawing a dashed line to its descendants. These models correspond to the neighborhood effects model in Section 4 under assumptions A1-A6, definition of treatment D2, and V in the Figure defined to be (U_D, U_0, U_1).

[8]World 18 describes a world in which an individual will be employed regardless of the neighborhood in which they reside, yet receiving an MTO voucher will cause them to become employed. World 19 implies that an individual will be employed regardless of the neighborhood in which they reside, yet receiving an MTO voucher will cause them to become *unemployed*. Finally, World 20 describes a world in which the individual is both always employed (Columns 3 and 4) or else is never employed (Columns 5 and 6), which simply cannot happen in our model as structured.

Table 3: Counterfactual Worlds Possible in Empirically- and Theoretically-Restricted Nbd Effects Model

After Restrictions Imposed by Empirical Observations and Theory (ie, the Model)

Definitions:

Z ≡ "Individual i receives an MTO voucher."
D ≡ "Individual i lives in a 'good' neighborhood."
Y ≡ "Individual i is employed."

Truth Table

Column	1	2	3	4	5	6
Row	$D(Z=1)$	$D(Z=0)$	$Y(D=1)$	$Y(D=0)$	$Y(Z=1)$	$Y(Z=0)$
(World 17)	1	0	1	1	1	1
(World 22)			1	0	1	0
(World 27)			0	1	0	1
(World 32)			0	0	0	0
(World 49)	0	0	1	1	1	1
(World 56)			1	0	0	0
(World 57)			0	1	1	1
(World 64)			0	0	0	0

Suppose that Table 3 does in fact represent the counterfactual worlds that could possibly correspond with our own *under the assumptions* of our model (Figure 1 and D2) and M1-M3. Under these assumptions, and a few more, we can use evidence on the program effects pertaining to $D(Z)$ and $Y(Z)$ to draw conclusions about the neighborhood effects represented by $Y(D)$. To begin, since Z is randomized we can learn about $D(Z)$ and $Y(Z)$ from the values of $E[D|Z]$ and $E[Y|Z]$ observed in MTO.

If we also adopt the assumptions:

NQB Neighborhood quality D is a binary function of a latent index of neighborhood quality q: $D \equiv \mathbf{1}\{q \geq q^*\}$

NQP Neighborhood quality q is a one-dimensional vector that is a scalar function of neighborhood poverty p: $q = \alpha p$

then the reasoning proceeds that the changes in neighborhood poverty observed in MTO imply that we must be in one of Worlds 17, 22, 27, or 32. Within these worlds, only 22 and 27 "exhibit neighborhood effects" (See columns 3 and 4.), and in these worlds there are also program effects (See columns 5 and 6). Thus, *under the adopted modeling assumptions*, the empirical evidence can justify statement (†).

Once statement (†) is justified, conclusions about neighborhood effects follow quickly. The reasoning proceeds looking at Columns 5 and 6. The empirical evidence on program effects tells us that we are either in World 32, 56, or 64. Combined with the observed changes in neighborhood poverty rates implying we are in one of Worlds 17, 22, 27, or 32, we must be in World 32. Thus we conclude:

(⋆): "The evidence from MTO suggests neighborhood effects are not strong."

Because statement (†) is false in more general models of neighborhood effects relaxing assumptions NQB and NQP, conclusion (⋆) need not be true in such models.[9] We now consider theoretical and empirical reasons why we would want to relax Assumptions NQB and NQP.

4 The Definition of Causal Effects

4.1 A Joint Model of Potential Outcomes and Selection

We now define several treatment effect parameters within a standard model of potential outcomes and selection into treatment (Heckman and Vytlacil (2005), Rubin (1974), Holland (1986)), initially taking no stand on what effects the researcher aims to identify. Let $Y(1)$ and $Y(0)$ be random variables associated with the potential outcomes in the treated and untreated states, respectively, at the individual level. D is a random variable indicating receipt of a binary treatment, where

$$D \equiv \begin{cases} 1 & \text{if treatment is received;} \\ 0 & \text{if treatment is not received.} \end{cases} \tag{1}$$

The measured outcome variable Y is

$$Y = DY(1) + (1 - D)Y(0) \tag{2}$$

where potential outcomes are a function of observable characteristics X_D and some treatment level specific unobservable component U_j for $j \in \{0, 1\}$:

$$\begin{aligned} Y(0) &= \mu_0(X_0) + U_0 \\ Y(1) &= \mu_1(X_1) + U_1. \end{aligned} \tag{3}$$

[9] Aliprantis and Richter (2014) is one example of neighborhood effects estimated under weaker assumptions than NQB and NQP in which the estimated effects contradict conclusion (⋆).

Note that these are not structural equations under Definition 5.4.1 in Pearl (2009).[10] Thus, since unobserved factors U_0 and U_1 influence $Y(0)$ and $Y(1)$, respectively, exclusion restrictions will need to be made if particular variables are to be ruled out of being a part of U_0 or U_1.

In the case of social experiments, a researcher can typically control assignment but not receipt of treatment. Thus we define Z as an indicator for the treatment assigned to an individual:

$$Z \equiv \begin{cases} 1 & \text{if treatment is assigned;} \\ 0 & \text{if treatment is not assigned.} \end{cases} \tag{4}$$

Noting it need not be true that $D = Z$, we write $D(Z)$ to denote the treatment received when assigned treatment Z and we explicitly model how individuals select into treatment. We suppose there is a latent index D^* that depends on observable characteristics X, assigned treatment Z, and some unobserved component V as follows:

$$\begin{aligned} D^* &= \mu_D(X_0, Z) - V \\ &= \mu_X(X_0) + \gamma Z - V, \end{aligned} \tag{5}$$

and that individuals select into treatment status based on their latent index:

$$D = \begin{cases} 1 & \text{if } D^* \geq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Finally, define the propensity score conditional on Z to be $\pi^Z(X) \equiv F_V(\mu_D(X, Z)) \equiv Pr(D = 1|X, Z)$.

We adopt a simple version of Heckman and Vytlacil (2005) and Heckman et al. (2006) by assuming:

A1 $\gamma_i = \gamma$ for all i and $\gamma \neq 0$

A2 $\{U_j, V\} \quad \perp\!\!\!\perp \quad Z \mid X$ for $j = 0, 1$

A3 The distribution of V is absolutely continuous

A4 $E\big[\,|Y(0)|\,\big|\,X\big] < \infty$ and $E\big[\,|Y(1)|\,\big|\,X\big] < \infty$

A5 $0 < \Pr(D = 1|X) < 1$ for all X

A6 $X = X_1 = X_0$ almost everywhere

[10] See Aliprantis (2015), Aliprantis (2014), or Heckman and Vytlacil (2005) for further discussion.

Given this joint model of potential outcomes and selection into treatment, there are several treatment effect parameters we might be interested in investigating. We define Intent-to-Treat (ITT), Treatment-on-the-Treated (TOT), and Local Average Treatment Effect (LATE) parameters:

$$\begin{aligned}
\triangle^{ITT}(x, \pi^0(x), \pi^1(x)) &\equiv E[Y \,|\, x,\, Z = 1] - E[Y \,|\, x,\, Z = 0] && (7)\\
\triangle^{TOT}(x) &\equiv E[Y(1) - Y(0) \mid x,\, D = 1] && (8)\\
\triangle^{LATE}(x, \pi^0(x), \pi^1(x)) &\equiv E[Y(1) - Y(0) \mid x,\, D(1) - D(0) = 1], && (9)
\end{aligned}$$

Heckman and Vytlacil (2005) show that these and all of the remaining treatment effect parameters in the literature can be written as weighted averages of a parameter introduced by Björklund and Moffitt (1987), the Marginal Treatment Effect (MTE), which is defined as:

$$\triangle^{MTE}(x, v) \equiv E[Y(1) - Y(0) \,|\, x, v]. \tag{10}$$

We also define $U_D \equiv F_{V|X}(V|X)$, so we can refer interchangeably to $\triangle^{MTE}(x, u_D)$, the MTE at the conditional quantiles of V. It will be useful in some of the following discussion to alternatively define the MTE in terms of U_D:

$$\triangle^{MTE}(x, u_D) = E[Y(1)|x, u_D] - E[Y(0)|x, u_D].$$

Note that so far we have stated no assumption on the relationship between the unobservable components determining potential outcomes and selection into treatment. The treatment effects we have defined in Equations 7-10 exist regardless of the relationship between potential outcomes and V. However, the interpretation of the treatment effect parameters will be very different depending on the assumptions we make about the relationship between the unobservables in the model. Two mutually exclusive (but not exhaustive) assumptions often adopted in the literature are Strong Ignorability and Essential Heterogeneity:

SI $\{U_1, U_0\} \;\perp\!\!\!\perp\; V|\; X.$

EH $COV(U_1 - U_0, V) \;|\; X \;\neq\; 0.$

Under SI the MTE is the same for all V. Since the MTE is homogeneous,

$$\triangle^{MTE}(x, u_D) = \triangle^{TOT}(x) = \triangle^{LATE}(x, \pi^0(x), \pi^1(x)) \tag{11}$$

for all $u_D \in [0, 1]$, for all x in the support of X, and for all $\pi^0(x), \pi^1(x) \in [0, 1]$. Under EH the MTE is allowed to be heterogeneous across V. Further discussion can be found in Appendix A.

5 The Identification of Causal Effects

5.1 What Program Effects Are Identified by MTO?

Since the model defined in Section 4.1 is built around selection into treatment, it is not fully specified without first defining treatment. Unobservables will be different for different definitions of treatment, and thus our assumptions will change based on our definition of treatment. We now consider identifying assumptions under two definitions of treatment that correspond to effects we hope the MTO experiment will help us to understand.

One obvious definition of treatment we might wish to consider is:

D1 Treatment is moving with the aid of the program (ie, using an MTO voucher).

Under A4 we can identify the ITT parameter by comparing the expected value of the outcome for those assigned to different voucher groups:

$$E[Y \mid x,\, Z=1] - E[Y \mid x,\, Z=0] = \triangle^{ITT}(x, \pi^0(x), \pi^1(x)).$$

Consider an additional restriction placed on the choice model,

A5* $\Pr[D(1)=1|X] > 0$ and $\Pr[D(0)=1|X] = 0$ for all X.

Under A5*

$$D(1) - D(0) = 1 \iff D(1) = 1\ , \tag{12}$$

and thus under either assumptions (A1-A6, SI, D1) or assumptions (A1-A6, A5*, SI, D1) the Wald estimator allows us to identify the homogeneous program effect of MTO:

$$\frac{E[Y|x, Z=1] - E[Y|x, Z=0]}{E[D|x, Z=1] - E[D|x, Z=0]} = \triangle^{MTE}(x, \cdot) = \triangle^{TOT}(x) = \triangle^{LATE}(x, \cdot, \cdot) \tag{13}$$

If we relax SI by assuming EH, then under (A1-A6, EH, D1) MTO identifies the following program effect that is determined in part by selection into treatment:

$$\frac{E[Y|x, Z=1] - E[Y|x, Z=0]}{E[D|x, Z=1] - E[D|x, Z=0]} = \triangle^{LATE}(x, \pi^0(x), \pi^1(x)). \tag{14}$$

And under (A1-A6, A5*, EH, D1) MTO identifies the following program effect that is also dependent on selection into treatment:

$$\frac{E[Y|x, Z=1] - E[Y|x, Z=0]}{E[D|x, Z=1] - E[D|x, Z=0]} = \triangle^{TOT}(x) = \triangle^{LATE}(x, 0, \pi^1(x)). \tag{15}$$

Since assumptions (A1-A6, A5*, EH, D1) appear reasonable together, the program effect in Equation 15 is identified by MTO. Appendix A has a further discussion of assumptions

about the distribution of unobserved variables, and Appendix B a discussion of the external validity of this parameter.

Estimates of these program effects can be found in the literature on MTO. Some of the major findings are that there were no significant effects on earnings, welfare participation, or the amount of government assistance adults received 5-7 years after randomization (Kling et al. (2007a)). There were, however, positive program effects on measures of adult mental health such as distress and calmness (Tables III in Kling et al. (2007a) and F5 in Kling et al. (2007b)). Sanbonmatsu et al. (2006) find program effects on reading scores, math scores, behavior problems, and school engagement that are statistically indistinguishable from zero for MTO children who were 6-20 on December 31, 2001. And perhaps the most surprising result was that while the program improved outcomes for young females, MTO had negative TOT effects on some outcomes of young males (Kling et al. (2007a), Kling et al. (2005)).

5.2 What Neighborhood Effects Are Identified by MTO?

Another treatment whose effects we might be interested in understanding is defined as follows:

D2 Treatment is moving to a high-quality neighborhood.

Note that under alternative definitions of treatment the selection model in Equations 5 and 6 will be modeling fundamentally different choices. The choice in the selection model under D2 is whether to move to a neighborhood with particular characteristics, while under D1 the choice modeled is whether to move with an MTO voucher.[11] The corresponding change in effect parameters in the model is to effects from moving to neighborhoods of varying quality. In the literature evidence pertaining to parameters of the model under D1 has been presented in discussions on parameters under D2, and vice-versa, showing the importance of clearly stating which modeling assumptions are being made.

5.2.1 Defining Neighborhood Quality and Assumption A2

There are two key reasons unobservables might be correlated with the instrument, which violates assumption A2, and both reasons are related to how we choose to define neighborhood quality in D2. The first problem results from assuming neighborhood quality is a binary

[11] While using an MTO voucher did initially require moving to a neighborhood with particular poverty characteristics (<10 percent), this requirement only had to be met for one year. Since subsequent moves were frequent, often involuntary, and tended to be to low-quality neighborhoods (de Souza Briggs et al. (2010), Sampson (2008)), the initial MTO move does not to capture the entire sequence of neighborhood characteristics, even when measured by poverty alone. Here I measure mobility using residence at the time of the interim evaluation, but other ways of dealing with dynamics, whether within the static models discussed here or within an expanded dynamic model, could also be appropriate.

variable when it is in fact multi-valued or continuous. For the sake of implementation we might assume

NQB Neighborhood quality D is a binary function of a latent index of neighborhood quality q: $D \equiv \mathbf{1}\{q \geq q^*\}$

To see the problems resulting from dichotomizing neighborhood quality when it is truly multi-valued or continuous, consider an example in which treatment is defined as moving to a neighborhood at the 80th percentile of neighborhood quality or higher (ie, $q^* = 80$). A household that would move to a neighborhood with quality at the 82nd percentile when not assigned treatment would be an always-taker under this definition of treatment. It is possible that such a household would be induced to move into a neighborhood of higher quality, say at the 90th percentile, after being assigned treatment. If this instrument-induced move were to impact outcomes, then U_0 would be correlated with Z. Such a violation of A2 results from the fact that changes in treatment intensity across margins other than those defining the binary treatment affect outcomes.[12]

One way to resolve this issue is to generalize the model in Section 4.1 in terms of the ordered choice model developed in Heckman et al. (2006).[13] A generalized framework assumes

NQJ Neighborhood quality D is a multi-valued function of a latent index of neighborhood quality q: $D \equiv j \times \mathbf{1}\{C_{j-1} < q \leq C_j\}$ where $j \in \{1, \ldots, J\}$

Given J levels of treatment, there should be some J large enough so that a generalized version of A2 holds.

The second reason unobservables might be correlated with the instrument arises if neighborhood quality is assumed to be represented by one vector when it is in fact multivariate. In the models currently estimated in the literature this assumption is operationalized as:

NQP Neighborhood quality q is a one-dimensional vector that is a scalar function of neighborhood poverty p: $q = \alpha p$

For example, Kling et al. (2007a) estimate neighborhood effects from MTO using a model assuming D2, NQJ, and NQP where MTEs are constant across unobservables.[14]

If neighborhood quality is truly multivariate, then there might be some neighborhood characteristics affecting outcomes other than poverty. If these characteristics are not perfectly correlated with poverty, then the U_j might be correlated with the instrument Z. Consider

[12] A discussion related to Assumption NQB can also be found in Angrist and Imbens (1995).

[13] An alternative and complementary approach is to use an unordered choice model as in Pinto (2014).

[14] To be precise, the model in Kling et al. (2007a) is the limit of this model as $J \to \infty$. Ludwig and Kling (2007) estimate a similar model with poverty replaced by beat crime rate. MTEs in these analyses are constant in U under the specification in Equation 3 since they assume $U_j = U$ for all $j \in \{1, \ldots, J\}$, so $U_{j+1,i} - U_{j,i} = U_i - U_i = 0$.

an example in which the neighborhood unemployment rate impacts labor market outcomes, with $D \in \{1, \ldots, 10\}$, and $D = j$ if the poverty rate is in the interval $[100 - 10j, 100 - 10(j-1)]$. There is some distribution of unemployment rates for those living in high ($D = j - 1$) and low poverty ($D = j$) neighborhoods, (U_{j-1}, U_j). If the people induced to move into low poverty neighborhoods due to the instrument tend to move to neighborhoods with higher unemployment rates than those who move to low poverty neighborhoods without the instrument, then the distribution of U_j will be different for those with $Z = 0$ than for those with $Z = 1$.

Assumption NQP rules out this possibility. If poverty were perfectly correlated with the unemployment rate, then in this example moving to a low poverty neighborhood would imply moving to a neighborhood with a given unemployment rate regardless of the instrument value, ensuring the distribution of the U_j would not be correlated with Z. Empirical evidence related to NQP is presented in Section 5.2.2.

A generalization of NQP is:

NQK Neighborhood quality q is a one-dimensional vector that is a linear combination of K observable neighborhood characteristics: $q = \alpha_1 X_1 + \cdots + \alpha_K X_K$

Assumption A2 might be more plausible under NQK than NQP since it uses more information about a neighborhood to determine its quality than solely its poverty rate.

5.2.2 Empirical Evidence on Assumptions A5, NQP, and NQK

The first source of data used to examine the stated identifying assumptions is the MTO Interim Evaluation sample. The sample contains variables listing the census tracts in which households lived at both the baseline and in 2002, the time the interim evaluation was conducted. These census tracts are used to merge the MTO sample with decennial census data from the National Historical Geographic Information System (NHGIS, Minnesota Population Center (2004)), which provide measures of neighborhood characteristics. These measures are analyzed both as raw values and as the percentiles of the national NHGIS variables from the 2000 census. The variables created in this way include the poverty rate, the percent of adults who hold a high school diploma or a BA, the male Employed-to-Population Ratio (EPR), the share of households with own-children under the age of 18 that are single-headed, and the female unemployment rate.

This analysis focuses on the adults in the MTO Interim Evaluation sample. Weights are used in constructing all estimates.[15]

[15] Weights 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 neighborhood effects are not conflated with time trends. Second, sampling weights must be used to account for the subsampling procedures used during the interim evaluation data collection.

Consider the generalized model in which neighborhood quality is defined under assumptions D2, NQJ, and NQK with $j \in \{1, \ldots, 10\}$ and

$$D \equiv j \times \mathbf{1}\{10 \times (j-1) < q \leq 10 \times j\},$$

where q is the percentile of neighborhood quality. A key assumption that can be empirically tested under this definition is A5, which is an assumption about the observed treatment states. The generalized version of assumption A5 is that $0 < Pr(D = j|X) < 1$ for all X, or that there are some persons in each treatment state.

Given the difficulties related to assumption NQP discussed in Section 5.2.1, we adopt NQK by combining several measures of neighborhood quality into a single vector representing neighborhood quality. Principal components analysis is used to determine which single vector combines the most information about the national distribution of the poverty rate, the percent with high school degrees, the percent with BAs, the percent of single-headed households, the male EPR, and the female unemployment rate. Table 4 shows that the resulting univariate index explains 63 percent of the variance of these neighborhood characteristics, and that no additional eigenvector would explain more than 13 percent of the variance of these variables. Table 5 displays the coefficients relating each of these variables to the index vector. Relevant for deciding between assumptions NQP and NQK, the magnitudes of the coefficients for most variables are similar to the magnitude of the coefficient for poverty.

Table 4: Proportion of Variance Explained by Principal Components Eigenvectors

Eigenvector	Eigenvalue	Proportion of Variance
1	3.81	0.63
2	0.79	0.13
3	0.56	0.09
4	0.39	0.07
5	0.31	0.05
6	0.14	0.02

Table 5: Principal Components Analysis: First Eigenvector Coefficients

Variable	Coefficient
Poverty Rate	-0.46
HS Graduation Rate	0.44
BA Attainment Rate	0.35
Percent Single-Headed HHs	-0.38
Male EPR	0.41
Female Unemployment Rate	-0.40

Figure 2a shows the expected negative correlation between neighborhood quality and neighborhood poverty rate. We can see in Figure 2b that the US population distribution of neighborhood poverty rates in 2000 had a long right tail. Similarly, Figure 2c shows that the US population distribution of neighborhood quality had a long left tail in 2000. Figures 2d and 2e show how far in the tails of these national distributions much of the MTO sample typically resided.

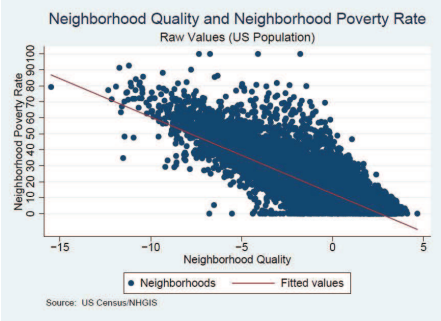

(a) Raw Measures of Neighborhood Quality and Poverty in 2000, US Population

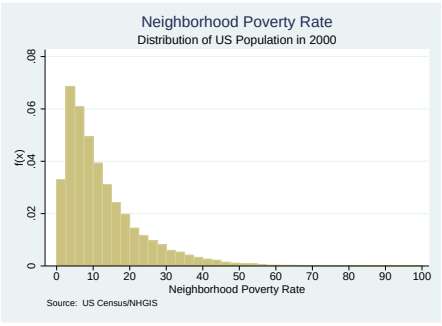

(b) Neighborhood Poverty Rate in 2000, US Population

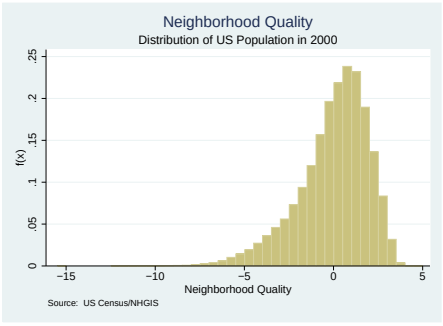

(c) Raw Measure of Neighborhood Quality in 2000, US Population

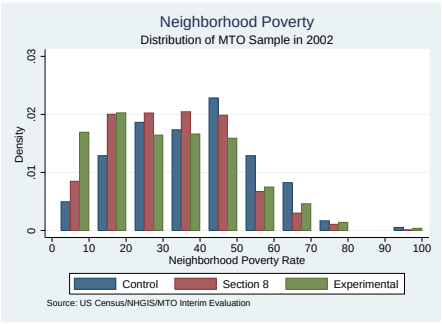

(d) Neighborhood Poverty Rate in 2002, MTO Sample

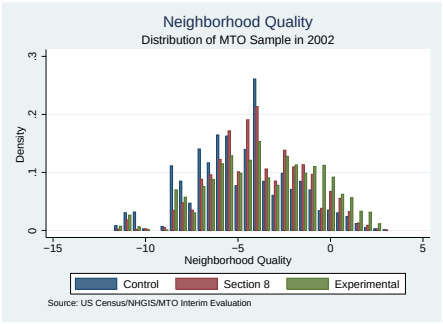

(e) Raw Measure of Neighborhood Quality in 2002, MTO Sample

Figure 2: Neighborhood Poverty Rate and Neighborhood Quality

Moving from a neighborhood with a poverty rate of 70 percent to a neighborhood with a 50 percent poverty rate might be a large change in the poverty rate, but Figure 2b suggests that we should also consider how big this change is relative to the national distribution of neighborhoods. An alternative way of measuring poverty and quality that addresses this question is to use the ranking of neighborhoods relative to those of the rest of the US population. These measures are shown for the entire US population in Figure 3. What we can see is that although the expected negative relationship still remains, there is now considerable variation in one variable conditional on the other. Consider, for example, that there are neighborhoods with the median poverty rate that are extremely low quality, and neighborhoods with the same poverty rate that are extremely high quality. This level of variation may not be surprising given the coefficients reported in Table 5, and can also be seen in Table 6, which presents evidence that in MTO states there were many low poverty neighborhoods that were also in the second and third deciles of the national distribution of quality. While the empirical evidence supports the adoption of assumption NQK over NQP if neighborhood characteristics other than poverty influence outcomes, simply comparing assumptions NQK and NQP in a theoretical way highlights that even defining neighborhood quality requires explicitly specifying which neighborhood characteristics influence outcomes.

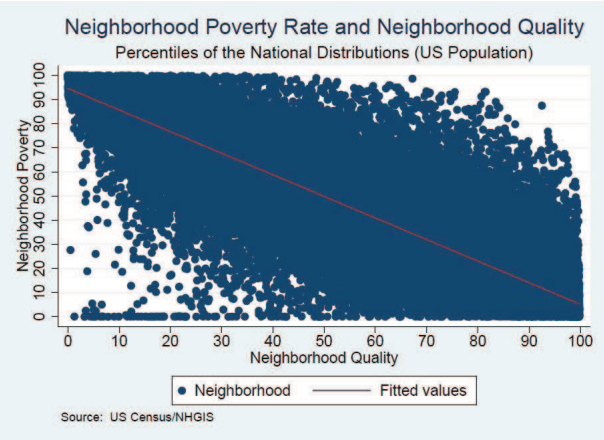

Figure 3: Neighborhood Poverty and Quality

Table 6: Low-Poverty ($\leq 10\%$), Low-Quality ($D \leq 3$) Neighborhoods in MTO States in 2000

Nbd Quality	Number of Residents
$D = 1$	6,362
$D = 2$	93,385
$D = 3$	751,738

Figure 4 shows that very few MTO adults were induced into high quality neighborhoods.[16] At the time of the interim evaluation less than 10 percent of the experimental group lived in neighborhoods whose quality was above the median of the national distribution. It is difficult to know for sure, but it appears reasonable to believe that the analogous distributions from Gautreaux would have had more mass in the right tail of the national distribution of neighborhood quality.[17]

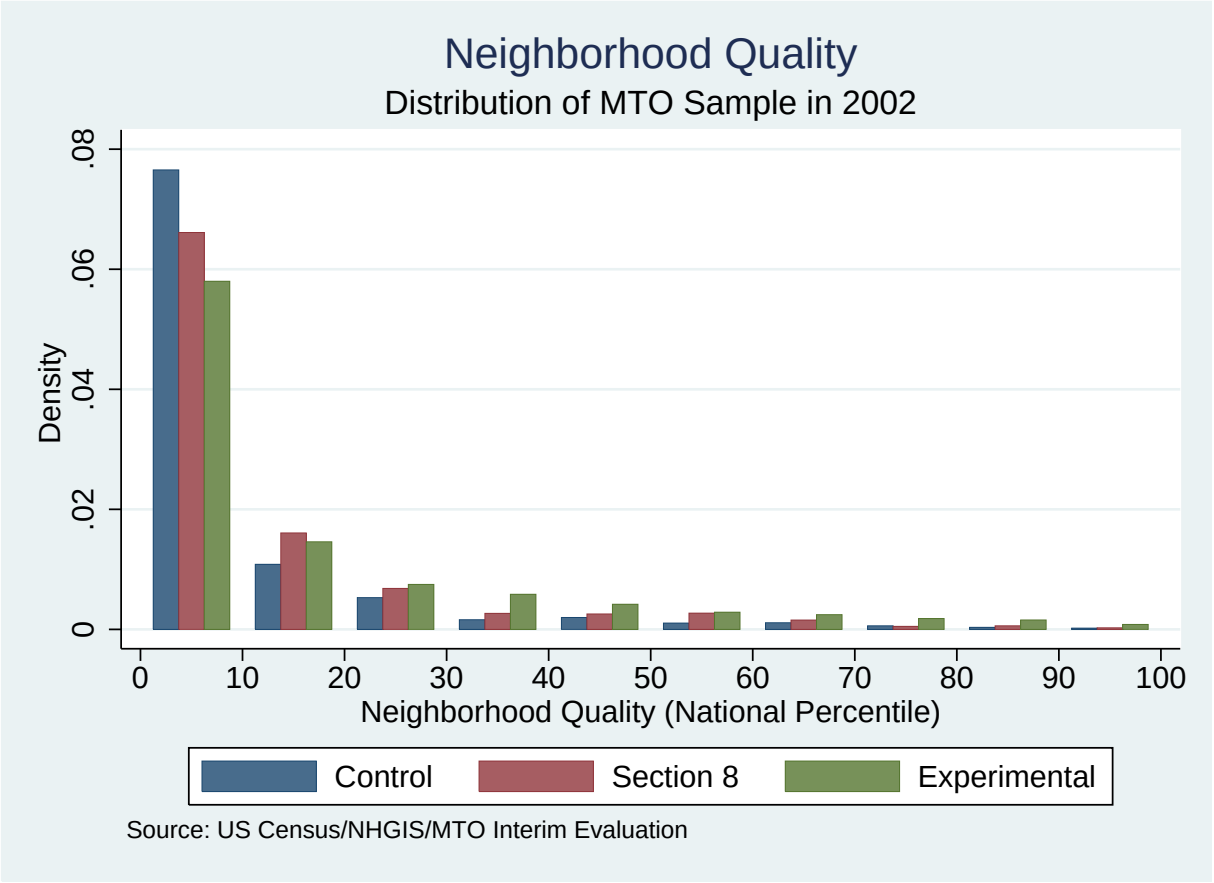

Figure 4: Neighborhood Quality of MTO Participants in 2002

[16]It is worth noting that the same general conclusion also holds in models assuming NQP. For example, Quigley and Raphael (2008) point out that "The effect of treatment under the MTO program was, on average, to move households in the five MTO metropolitan areas from neighborhoods at roughly the 96th percentile of the neighborhood poverty distribution to neighborhoods at the 88th percentile" (p 3).

[17]DeLuca and Rosenbaum (2003) find that 66 percent of the suburban group and 13 percent of the city group lived in the suburbs of Chicago 14 years after original placement through Gautreaux. DeLuca and Rosenbaum (2003) cite limited availability of housing, rather than the choice to not move through the program, as the reason only 20 percent of eligible applicants moved through Gautreaux. This claim is based on evidence that 95 percent of participating households accepted the first unit offered to them. Furthermore, it is likely that Gautreaux induced larger changes in school quality than MTO (Rubinowitz and Rosenbaum (2000), p 162). Taken together, this evidence is suggestive that Gautreaux induced more households into high quality neighborhoods than MTO.

The distributions in Figure 4 can be seen as a violation of the generalized version of assumption A5. While technically true for all j without conditioning on X, for the sake of estimation the generalized version of A5 is only likely to hold for $j \in \{1, \ldots, 5\}$ or $j \in \{1, \ldots, 6\}$. By the time of the interim evaluation less than 20 percent of the MTO experimental group lived in neighborhoods above the 30th percentile of the national distribution of quality, and less than 10 percent lived in neighborhoods above the median.

5.2.3 The Neighborhood Effects Identified by MTO

Effects from moving to high quality neighborhoods are not identified by MTO. Given the evidence in Section 5.2.2, any definition of treatment of the form D2 would have to restrict measures of quality to the lower half of the national distribution of neighborhood quality to satisfy assumption A5.

Once the focus on quality is restricted to accommodate A5, we can see that A5 appears more reasonable than A5*, as it is likely that some households will move to a relatively high quality neighborhood regardless of whether they receive a voucher through MTO or not. Under assumptions (A1-A6, EH, D2-NQB) the Wald estimator identifies the LATE:

$$\begin{aligned}\frac{E[Y|x, Z=1] - E[Y|x, Z=0]}{E[D|x, Z=1] - E[D|x, Z=0]} &= \triangle^{LATE}(x, \pi^0(x), \pi^1(x)) \\ &= \frac{1}{\pi^1(x) - \pi^0(x)} \int_{\pi^0(x)}^{\pi^1(x)} \triangle^{MTE}(x, u_D) du_D. \qquad (16)\end{aligned}$$

If we believe assumption A2 will fail to hold when treatment is defined under D2-NQB for the reasons discussed in Section 5.2.1, we could alternatively define treatment under D2-NQJ to generate level j specific analogues to 10 and 16:

$$\begin{aligned}\triangle^{MTE}_{j,j+1}(x, u_D) du_D &= E[Y(j+1) - Y(j) | X = x, U_D = u_D] \\ \triangle^{LATE}_{j,j+1}(x, \pi^0_j(x), \pi^1_j(x)) &= \frac{1}{\pi^1_j(x) - \pi^0_j(x)} \int_{\pi^0_j(x)}^{\pi^1_j(x)} \triangle^{MTE}_{j,j+1}(x, u_D) du_D.\end{aligned}$$

Versions of the model have been estimated in Kling et al. (2007a) and Ludwig et al. (2008) under (A1-A6, SI, and D2-NQJ-NQP). A dose-response analysis is used in Kling et al. (2007a) to determine if parameters are constant across all j to $j+1$ transitions in $\{1, \ldots, J\}$. Aliprantis and Richter (2014) estimate the model under (A1-A6, EH, D2-NQJ-NQK). That analysis makes A2 more plausible by relaxing D2-NQJ-NQP to D2-NQJ-NQK, and allows for the identification and estimation of LATEs that are heterogeneous over unobservables by relaxing SI to EH.[18]

[18] Note that NQK need not be adopted only in conjunction with NQJ. A version of Assumption NQB-NQK is adopted in Sampson et al. (2008) using a similar index of neighborhood quality to that used in this analysis.

6 Conclusion

This paper built on the Clampet-Lundquist and Massey (2008) critique to show that the literature draws unwarranted conclusions by using MTO program effects to learn about neighborhood effects. It was shown how the most popular interpretation of results from MTO in the literature uses ITT and TOT program effects from MTO to indirectly draw conclusions about neighborhood effects. The logic required to adopt this interpretation was expressed explicitly, in terms of the assumptions of a model of neighborhood effects.

The paper theoretically and empirically reviewed these assumptions that the literature on MTO implicitly imposes on the neighborhood effects model. An index of neighborhood quality was created that reflects a neighborhood's poverty rate as well as several other characteristics. Empirical evidence was presented that MTO did not induce participants into many levels of neighborhood quality, discrediting the assumption of a binary neighborhood quality assumption. Empirical evidence was also presented that there exist low-poverty neighborhoods that are also low-quality, casting doubt on the assumption that neighborhood poverty serves as a good proxy for neighborhood quality in MTO.

References

Aliprantis, D. (2014). Covariates and causal effects: The problem of context. *Federal Reserve Bank of Cleveland Working Paper 13-10r*.

Aliprantis, D. (2015). A distinction between causal effects in Structural and Rubin Causal Models. *Federal Reserve Bank of Cleveland Working Paper*.

Aliprantis, D. and F. G.-C. Richter (2014). Evidence of neighborhood effects from MTO: LATEs of neighborhood quality. *Federal Reserve Bank of Cleveland Working Paper 12-08r*.

Angrist, J. D. (2014, Forthcoming). The perils of peer effects. *Labour Economics*. http://dx.doi.org/10.1016/j.labeco.2014.05.008.

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.

Angrist, J. D. and J.-S. Pischke (2010). The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *Journal of Economic Perspectives 24*(2), 3–30.

Björklund, A. and R. Moffitt (1987). The estimation of wage gains and welfare gains in self-selection models. *The Review of Economics and Statistics 69*(1), pp. 42–49.

Brock, W. and S. Durlauf (2007). Identification of binary choice models with social interactions. *Journal of Econometrics 140*(1), 52–75.

Carneiro, P., J. J. Heckman, and E. J. Vytlacil (2011). Estimating marginal returns to education. *American Economic Review 101*(6), 2754–2781.

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.

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.

DeLuca, S. and J. E. Rosenbaum (2003). If low-income blacks are given a chance to live in white neighborhoods, will they stay? Examining mobility patterns in a quasi-experimental program with administrative data. *Housing Policy Debate 14*(3), 305–345.

Friedman, M. (1955). The role of government in education. In R. Solo (Ed.), *Economics and the Public Interest*. New Brunswick, NJ: Rutgers University Press.

Gennetian, L. A., M. Sciandra, L. Sanbonmatsu, J. Ludwig, L. F. Katz, G. J. Duncan, J. R. Kling, and R. C. Kessler (2012). The long-term effects of Moving to Opportunity on youth outcomes. *Cityscape 14*(2), 137–167.

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).

Heckman, J. J. (2010). Building bridges between structural and program evaluation approaches to evaluating policy. *Journal of the Economic Literature 48*(2), 356–398.

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.

Heckman, J. J. and E. Vytlacil (2005). Structural equations, treatment effects, and econometric policy evaluation. *Econometrica 73*(3), 669–738.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association 81*(396), 945–960.

Imbens, G. W. and J. D. Angrist (1994). Identification and estimation of Local Average Treatment Effects. *Econometrica 62*(2), 467–475.

Keels, M., G. J. Duncan, S. Deluca, R. Mendenhall, and J. Rosenbaum (2005). Fifteen years later: Can residential mobility programs provide a long-term escape from neighborhood segregation, crime, and poverty? *Demography 42*(1), pp. 51–73.

Kling, J. R., J. B. Liebman, and L. F. Katz (2007a). Experimental analysis of neighborhood effects. *Econometrica 75*(1), 83–119.

Kling, J. R., J. B. Liebman, and L. F. Katz (2007b). Supplement to "Experimental analysis of neighborhood effects": Web appendix. *Econometrica 75*(1), 83–119.

Kling, J. R., J. Ludwig, and L. F. Katz (2005). Neighborhood effects on crime for female and male youths: Evidence from a randomized housing voucher experiment. *The Quarterly Journal of Economics 120*(1), 87–130.

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., G. J. Duncan, and P. Hirschfield (2001). Urban poverty and juvenile crime: Evidence from a randomized housing-mobility experiment. *The Quarterly Journal of Economics 116*(2), 655–679.

Ludwig, J. and J. R. Kling (2007). Is crime contagious? *Journal of Law and Economics 50*(3), 491–518.

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. (2013a). Identification of treatment response with social interactions. *The Econometrics Journal 16*(1), S1–S23.

Manski, C. F. (2013b). *Public Policy in an Uncertain World: Analysis and Decisions*. Harvard University Press.

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. http://www.nhgis.org.

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.

Pearl, J. (2009). *Causality: Models, Reasoning and Inference* (2nd ed.). Cambridge University Press.

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.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology 66*(5), 688–701.

Rubin, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics 6*(1), 34–58.

Rubinowitz, L. S. and J. E. Rosenbaum (2000). *Crossing the Class and Color Lines: From Public Housing to White Suburbia*. University of Chicago Press.

Sampson, R. J. (2008). Moving to inequality: Neighborhood effects and experiments meet social structure. *American Journal of Sociology 114*(1), 189–231.

Sampson, R. J. (2012). *Great American City: Chicago and the Enduring Neighborhood Effect*. The University of Chicago Press.

Sampson, R. J., P. Sharkey, and S. W. Raudenbush (2008). Durable effects of concentrated disadvantage on verbal ability among African-American children. *Proceedings of the National Academy of Sciences of the United States of America 105*(3), 845–852.

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.

Sanbonmatsu, L., J. Marvakov, N. A. Potter, F. Yang, E. Adam, W. J. Congdon, G. J. Duncan, L. A. Gennetian, L. F. Katz, J. R. Kling, R. C. Kessler, S. T. Lindau, J. Ludwig, and T. W. McDade (2012). The long-term effects of Moving to Opportunity on adult health and economic self-sufficiency. *Cityscape 14*(2), 109–136.

Sobel, M. E. (2006). What do randomized studies of housing mobility demonstrate?: Causal inference in the face of interference. *Journal of the American Statistical Association 101*(476), 1398–1407.

Votruba, M. E. and J. R. Kling (2009). Effects of neighborhood characteristics on the mortality of black male youth: Evidence from Gautreaux, Chicago. *Social Science & Medicine 68*(5), 814–823.

Wilson, W. J. (1987). *The Truly Disadvantaged: The Inner City, the Underclass, and Public Policy*. University of Chicago.

A Assumptions about the Distribution of Unobservables

The interpretation of the treatment effect parameters will be very different depending on the assumptions we make about the relationship between the unobservables in the model. Strong ignorability is a standard assumption made in the statistics and econometrics literature about the relationship between the unobservable component determining selection into treatment and those determining potential outcomes. Strong ignorability is fundamentally an assumption about what the econometrician is able to observe; it is that the econometrician can observe all characteristics connecting selection into treatment with treatment effect heterogeneity. Although this assumption may be unrealistic in many applications, it is adopted frequently because it is helpful for identification for reasons that will be discussed shortly.

An implication of strong ignorability is that conditional on observables, selection into treatment is not related to treatment effect heterogeneity. Formally, strong ignorability can be written in our model as

SI $\{U_1, U_0\} \quad \perp\!\!\!\perp \quad V|\ X.$

Under SI the MTE is the same for all V. Since the MTE is homogeneous,

$$\triangle^{MTE}(x, u_D) = \triangle^{TOT}(x) = \triangle^{LATE}(x, \pi^0(x), \pi^1(x)) \tag{17}$$

for all $u_D \in [0,1]$, for all x in the support of X, and for all $\pi^0(x), \pi^1(x) \in [0,1]$.

Imbens and Angrist (1994) showed it is possible to identify an interpretable parameter, the LATE, even if strong ignorability fails. Recent work in Heckman and Vytlacil (2005), Heckman et al. (2006), and Carneiro et al. (2011) has further defined and estimated treatment effect parameters when relaxing the assumption of strong ignorability by assuming that unobservable treatment effect heterogeneity is related to the unobservable determinants of selection into treatment. Formally, the assumption of essential heterogeneity is that

EH $COV(U_1 - U_0, V)\ |\ X\ \neq\ 0.$

Figure 5 helps to illustrate the implications of SI and EH. The top panel in the figure shows that average treatment effects are allowed to vary across observable characteristics. SI and EH characterize different scenarios once we select a particular value of observable characteristics, x^*. In the middle panel of the figure we see a scenario of SI. The distributions of the potential outcomes must be independent of V given x^*, so the levels of the potential outcomes must be constant across V given x^*. The differences between these levels, the MTEs, are thus constant for all V given x^*.

The bottom panel of Figure 5 shows a contrasting scenario of EH. In this scenario the difference $U_1 - U_0$ is correlated with V, resulting in MTEs that vary across V. In the example displayed the effect of treatment is large for low levels of V, while for large values of V the

effect of treatment decreases. Given our latent index model, this implies that for the given observable characteristics x^*, treatment effects are large for those who would be most likely to select into treatment in the absence the program. Finally, Figure 6 shows that while SI and EH are mutually exclusive, they are not exhaustive since individuals might select on the level while not selecting on the gain.

The contrast in the role of instrumental variables under SI versus EH is shown clearly in Figure 5. Under SI it does not matter who is induced into treatment by the instrument since all variation from Z identifies the same homogeneous parameter. Unlike EH, one might assume SI and estimate parameters without the existence of an instrument, perhaps implemented with propensity score matching. In fact, it may appear to be superfluous to use an instrument in conjunction with the SI assumption. This is not necessarily the case, though, as adding a valid instrument Z to the latent index in Equation 5 can make SI a more plausible assumption.

In contrast to SI, under EH the selection into treatment induced by the instrument is of central interest for interpreting parameters. Since MTEs vary over the support of U_D, the subinterval induced into treatment by the instrument will determine the parameter(s) identified by the instrument. Different instruments that induce different intervals of U_D into treatment will identify different parameters.

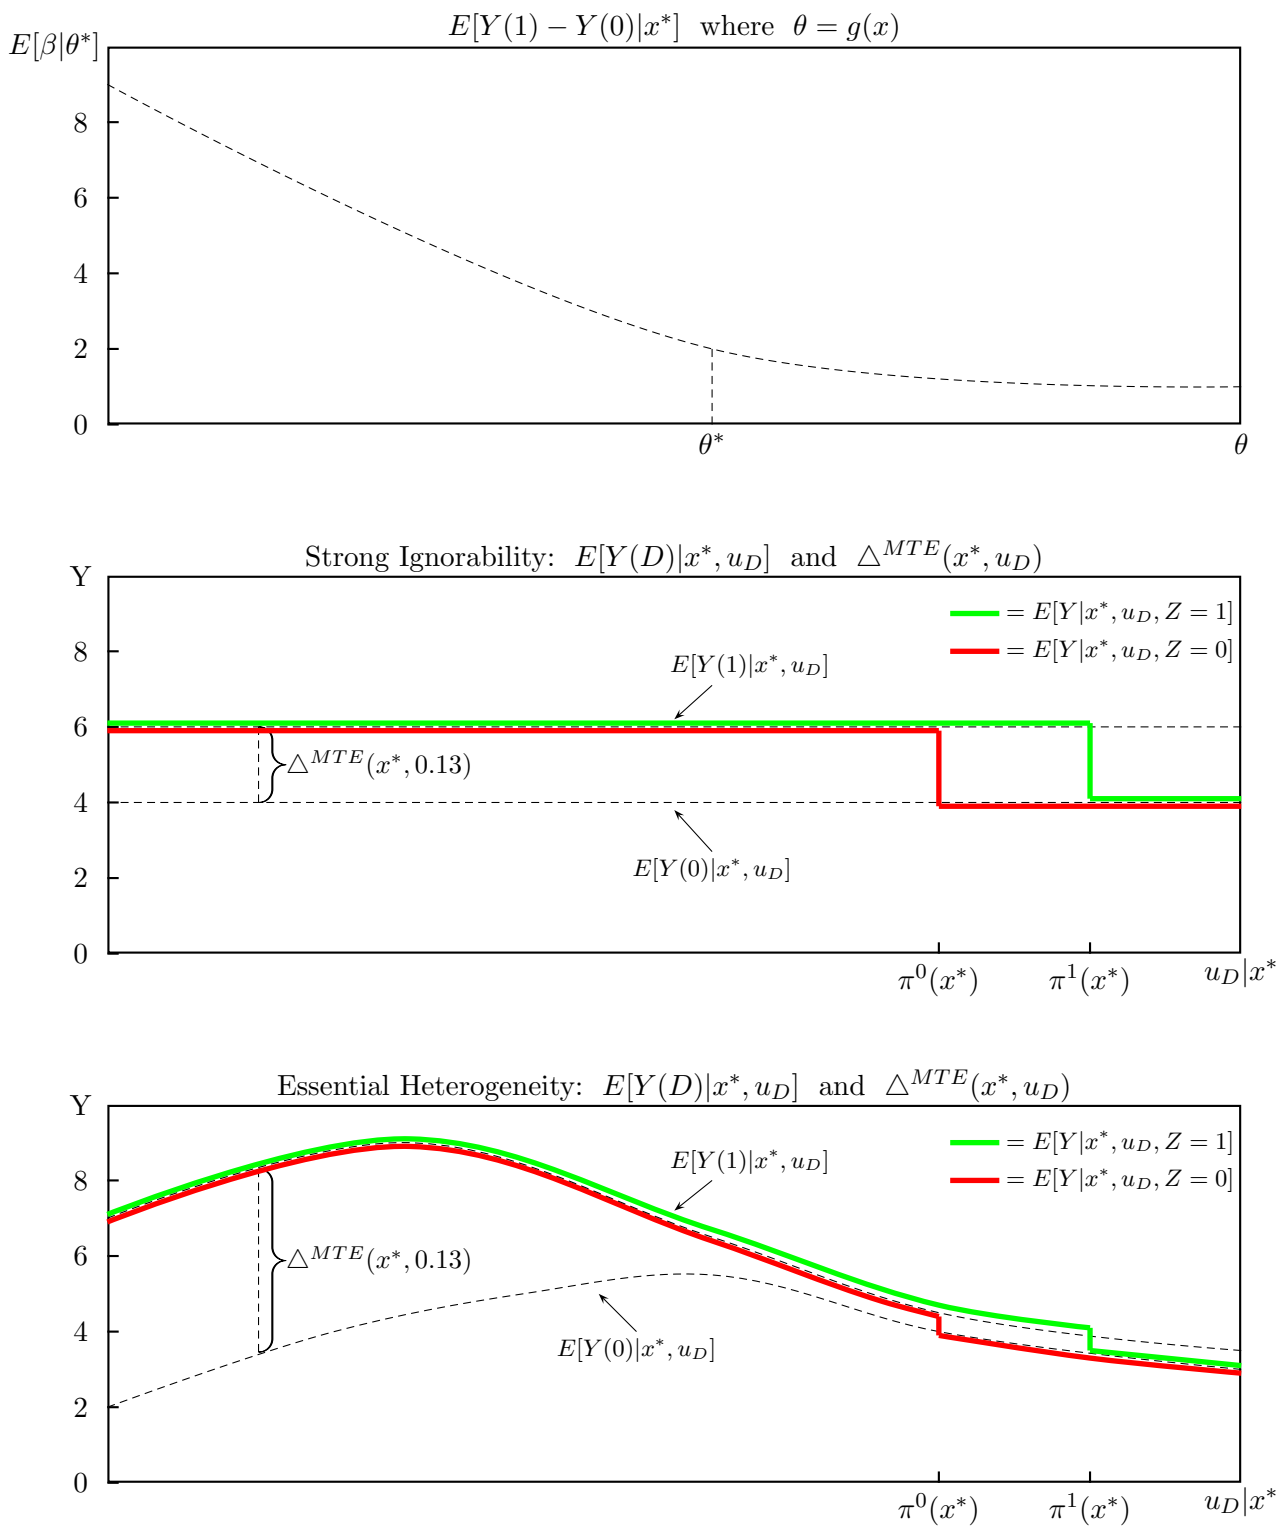

Figure 5: Examples of Strong Ignorability and Essential Heterogeneity

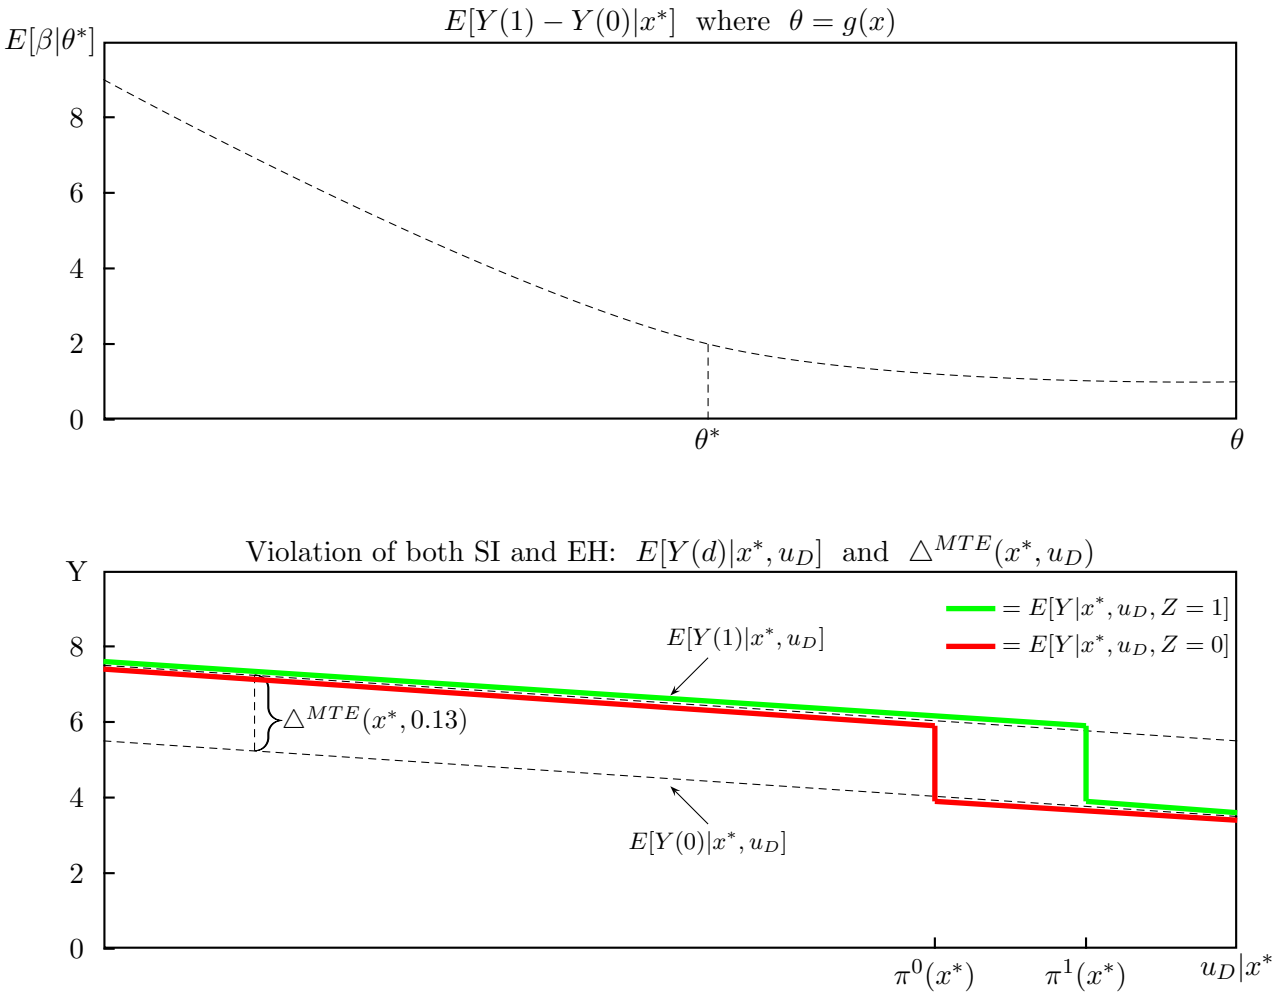

Figure 6: Example Violating Both Strong Ignorability and Essential Heterogeneity

B External Validity

Although external validity is the motivation for studying causal effects, and there is no clear reason for prioritizing internal validity over external validity (Manski (2013b)), the literature has focused most formal attention on internal validity (Aliprantis (2014)). The text has adopted these priorities for the sake of publication, but here we also consider why estimated parameters will not be experiment invariant unless an assumption also holds that restricts the permissible types of peer effects (Sobel (2006)). Interested readers are also directed to the careful discussions of these issues in Sobel (2006) and Ludwig et al. (2008).

B.1 Assumptions across and within Individuals

The parameters in Section 4.1 are all defined conditional on the joint distribution (U, V) where we define $U \equiv (U_0, U_1)$. Assumptions about how these random variables interact across individuals have implications for the joint distribution (U, V) and will change the interpretation of the parameters we have defined.

One possibility satisfying A6 is for X to be a bundle of individual level characteristics including baseline neighborhood characteristics, with one element captured in the unobservables V being peer effects on the selection decision.[19] We now take some terminology from Sobel (2006) to consider the implications of changes to the distribution of V. We suppose the MTO experiment involves N individuals, that there are k_1 people assigned to $Z = 1$, and that $k_0 = N - k_1$ are assigned to $Z = 0$, here again abstracting from the Section 8 group for the sake of exposition. Let $R(k_0, k_1)$ denote the set of possible realizations of such a randomization, with $r \in R(k_0, k_1)$ denoting one possible realization. If peer effects determining selection into treatment are a part of V, then different realizations r may result in different distributions of V, which we write as $F_{V|r}$. Returning to the fact that all of the parameters defined in Section 4.1 are defined assuming some distribution of (U, V), this implies that these parameters might be very different for some realization r compared to another realization r' (Sobel (2006)).

A standard assumption on the nature of peer effects resolves this problem by ensuring the effects defined in Section 4.1 are the same for all realized random assignments r. This assumption simply assumes there are no peer effects at all. In the context of our model, Angrist and Imbens (1995) state the Stable Unit Treatment Value Assumption (SUTVA) from Rubin (1978) as

SUTVA (a) $V_i \perp\!\!\!\perp Z_j$ for all $j \neq i$

SUTVA (b) $(U_{0i}, U_{1i}) \perp\!\!\!\perp Z_j$ and $(U_{0i}, U_{1i}) \perp\!\!\!\perp D_j$ for all $j \neq i$

[19] See page 677 of Heckman and Vytlacil (2005) for a relevant discussion of A6, and see Brock and Durlauf (2007) for a related model of peer effects on the selection decision.

Note that SUTVA is an assumption across different individuals. Under SUTVA, SI and EH are primarily assumptions within individuals. In this case, unobservables are primarily thought to represent individual-level causal variables. Although (U, V) can represent social interactions under SUTVA, these social interactions cannot be related to treatment or assigned treatment.[20] When SUTVA is relaxed, however, SI and EH become assumptions not only about individual-level causal variables, but also about social interactions.

A less restrictive assumption on peer effects that still keeps the effects in Section 4.1 identical across realizations of the randomization is that the distribution of peer effects will be identical under all realizations r. I label this as the Stable Peer Effects Assumption (SPEA):

SPEA $(U, V) \perp\!\!\!\perp R$

Note that neither SUTVA nor SPEA is necessary to define and estimate the parameters in Section 4.1. However, the model illustrates how the lack of such an assumption dramatically changes their interpretation. Since the distribution of peer effects included in V might change in different contexts, this could have very important consequences, both in terms of whether the parameters in the model are invariant to the realization of randomized voucher assignment (Sobel (2006)), as well as in terms of parameter invariance to classes of policy interventions. Importantly, this discussion illustrates that, just like SI or EH, parameter invariance is an assumption about the unobserved variables in the model.

[20] Although this model of neighborhood effects has additional mechanisms relative to those typically included in models of social interaction, such models are still useful to consider in this context. For example, Manski (1993) and Brock and Durlauf (2007) specify models relaxing SUTVA (a) and Manski (2013a) specifies a model relaxing SUTVA (b).

C List of Assumptions

Given the joint model of potential and outcomes and selection into treatment:

$$Y(D) = \mu_D(X_D) + U_D,$$

$$D^* = \mu_X(X_0) + \gamma Z - V,$$

with

$$D = j \quad \text{if} \quad D^* \in (C_{j-1}, C_j],$$

there are several assumptions about the model considered throughout the paper. I list them here for the reader's reference:

A1 $\gamma_i = \gamma$ for all i and $\gamma \neq 0$

A2 $\{U_j, V\} \perp\!\!\!\perp Z \mid X$

A3 The distribution of V is absolutely continuous

A4 $E\big[\,|Y(j)|\,\big|\,X\big] < \infty$ for all j

A5 $0 < \Pr(D = j|X) < 1$ for all X, j

A6 $X = X_j = X_k$ almost everywhere for all $j \neq k$

D1 Treatment is moving with the aid of the program (ie, using an MTO voucher).

D2 Treatment is moving to a high-quality neighborhood.

M1 $D_i \equiv \mathbf{1}\{$ individual i lives in a high-quality neighborhood $\}$

M2 $Z_i \equiv \mathbf{1}\{$ individual i received an MTO voucher $\}$

M3 $Y_i \equiv \mathbf{1}\{$ individual i is employed $\}$

NQB Neighborhood quality D is a binary function of a latent index of neighborhood quality q: $D \equiv \mathbf{1}\{q \geq q^*\}$

NQJ Neighborhood quality D is a multi-valued function of a latent index of neighborhood quality q: $D \equiv j \times \mathbf{1}\{C_{j-1} < q \leq C_j\}$

NQP Neighborhood quality q is a one-dimensional vector that is a scalar function of neighborhood poverty p: $q = \alpha p$

NQK Neighborhood quality q is a one-dimensional vector that is a linear combination of K observable neighborhood characteristics: $q = \alpha_1 X_1 + \cdots + \alpha_K X_K$

SUTVA (a) $V_i \perp\!\!\!\perp Z_j$ for all $j \neq i$

SUTVA (b) $(U_{0i}, U_{1i}) \perp\!\!\!\perp Z_j$ and $(U_{0i}, U_{1i}) \perp\!\!\!\perp D_j$ for all $j \neq i$

SPEA $(U, V) \perp\!\!\!\perp R$ for randomization R