NYU Furman Center

Neighbors and Networks: The Role of Social Interactions on the Residential Choices of Housing Choice Voucher Holders

Ingrid Gould Ellen, Michael Suher, and Gerard Torrats-Espinosa

February 2017

furmancenter.org

Neighbors and Networks:
The Role of Social Interactions on the Residential Choices of Housing Choice Voucher Holders

Ingrid Gould Ellen
Michael Suher
Gerard Torrats-Espinosa

February 2017

Abstract:
The housing choice voucher program aims to reduce housing cost burdens as well as to enable recipients to move to a broader diversity of neighborhoods. Prior evidence shows voucher recipients still end up in neighborhoods with relatively high poverty rates and low performing schools. These constrained neighborhood choices can in part be attributed to landlord discrimination and the geographic concentration of units that rent below voucher caps. In this paper, we consider an additional explanation: the role of information and social influence in determining the effective set of potential housing choices. Using a strategy based on proximity of households in origin census tracts, we find evidence consistent with social influence effects being present in the neighborhood choices of voucher holders. Pairs of households living within the same or adjacent buildings are significantly more likely to relocate to the same neighborhood as each other than are more distant households within the same origin neighborhood. Further, we show that voucher holders who move to the same neighborhood as a nearby voucher holder end up on average in neighborhoods that have higher poverty rates, lower levels of labor market engagement, and higher exposure to environmental hazards --- in both absolute terms and relative to other voucher holders from their same origin tract.

Finding a home is not easy, especially for low-income households. This is all the more true in metropolitan areas with tight housing markets, where fewer units become available. Consider that a household with an income at the 25th percentile of the local renter income distribution in one of the nation's ten largest metropolitan areas in 2014 would have found fewer than 5 percent of available rental units to be affordable (Ellen and Karfunkel, 2014). Low-income households who are lucky enough to obtain a housing choice voucher should encounter many more affordable units, and yet they still typically live in neighborhoods that are only slightly less disadvantaged than the typical poor household (Pendall, 2000; Wood, Turnham and Mills, 2008; Galvez, 2011).

Recent studies have probed this puzzle of why voucher holders are so concentrated in high-poverty neighborhoods, a question that has taken on renewed urgency since Chetty, Hendren, and Katz (2015) published their experimental findings on the benefits of moving to lower poverty environments. Existing research has pointed to a number of potential contributors: the geographic concentration of units renting below voucher rent caps, limited search time, trade-offs between housing structure and neighborhood, and unwilling landlords. The existing literature has yet to pay much attention to the role of informal information networks. In this study, we explore whether voucher holders appear to use informal networks in their housing search and whether those networks tend to steer them to higher poverty neighborhoods.

We use geocoded longitudinal data on the universe of housing choice voucher holders in the U.S. between 2011 and 2014 to learn whether and how households with vouchers appear to use informal networks to aid their housing searches. While we do not have direct information on search methods, we can observe whether voucher holders tend to move to the same neighborhoods as other voucher holders who live in the same building or block, after controlling for the initial neighborhood and household attributes. We find robust evidence that pairs of voucher holders who live in the same building or in very close proximity to one another are more likely to move to the same neighborhood than other pairs of voucher holders who also live in the same neighborhood but further away from one another. We find that these effects are magnified in tight housing markets where searches are more challenging. They are also magnified in more racially segregated metropolitan areas, where voucher holders, about two thirds of whom are black or Hispanic, likely face more significant residential constraints. Finally, voucher holders who move to the same census tract of destination as other voucher holders who initially live in the same building or block are more likely to move to neighborhoods that have higher poverty rates, lower levels of labor market engagement, and higher exposure to environmental hazards, suggesting that social networks among voucher holders tend to guide them to more economically disadvantaged neighborhoods. Reliance on networks may be improving search outcomes along other dimensions such as housing unit quality. While we lack data on specific housing quality characteristics, evidence on gross rents does not support this possibility.

Background and Literature

Social Networks and Housing Search

Despite the fact that housing typically represents the largest expenditure that households make, there is surprisingly little research on housing search methods. Ford, Rutherford, and Yavas (2005) model how households may stop their housing search before finding the best match to reduce search costs, even when the distribution of properties available is fully known. Voucher holders in particular face an unusual search problem. Their past housing search experiences prior to receiving the voucher would, due to cost constraints, have precluded gaining knowledge of a set of units and neighborhoods which are now within their budget set. Marsh and Gibb (2011) note that the noisier the information available during housing market search, the greater the likelihood that the searcher will make a sub-optimal choice. In a classic geography paper, Brown and Moore (1970) argue that in the face of limited time and high search costs, households are likely to try to reduce housing search costs by turning to neighborhoods they know through social ties or day-to-day experience. The empirical evidence on these networks is limited, though the studies that exist generally find that informal social networks are an important source of information. Krysan (2008), for example, shows that homeseekers in Detroit frequently rely on conversations with friends and relatives in conducting their housing search. Specifically, more than a third of surveyed renters reported relying on conversations with friends, just below the share relying on newspaper and internet searches (43 and 37 percent respectively).

These patterns hold true across racial groups and for both renters and homeowners (Krysan 2008). It appears that informal social networks are an important source of information for households of all incomes engaging in housing searches. That said, these social networks are economically stratified, with lower income families relying on networks composed of lower income families (Lareau and Goyette 2014). They may be racially stratified too. For example, Krysan and Bader (2009) ask a random sample of Chicago residents to identify familiar neighborhoods and find strong racial differences, with white residents expressing little knowledge of integrated and largely minority areas, and black and Latino residents expressing little knowledge of the largely white areas. The role of social networks in driving these racial gaps in information (or "blind spots") is unclear.

Housing Choice Voucher Program

The federal government now spends about 18 billion dollars annually to provide assistance to over 5 million people in approximately 2.2 million households with housing choice vouchers (CBPP, Fact Sheet, 2016). On average, voucher holders receive an effective subsidy of approximately $8,000 per year (Collinson, Ellen and Ludwig, 2015).[1]

While slightly different variants of the voucher program have emerged over its more than four-decade existence, the basic structure has remained the same. Recipients use vouchers to help pay for housing units that they rent on the private market. They generally pay 30 percent of their income towards rent, while the federal government covers the difference between this tenant payment and the rent, as long as the rent is below a locally defined payment standard. The payment standard is set between 90 and 110 percent of the Fair Market Rent (FMR), which is defined as either the 40th or 50th percentile of rents in the metropolitan area, depending on market conditions. Voucher holders are allowed to rent units with rents above the payment standard, but

[1] The average monthly HUD subsidy was $647 in 2013 (HUD Congressional Justification FY 2015).

they must fully pay any amount that exceeds the payment standard. As a result, they can face rent burdens above 30 percent, and many do in tight markets, but when voucher holders first move into a new unit, their rent burden may not exceed 40 percent of income.

To receive a voucher, households apply to a local Public Housing Authority (PHA), which certifies that their income is below the eligibility threshold of 80 percent of the area median income (AMI). In practice, most voucher holders have far lower incomes, typically at or below the poverty line, as PHAs are required to set aside 75 percent of the vouchers they award each year for households with incomes at or below 30 percent of AMI. After receiving a voucher, households have a limited time period in which they must use it, which can be as short as 60 days. Households must find units that are not only affordable to them, but are of the appropriate size, meet federal housing quality standards, and charge a rent that the local housing authority deems reasonable given the local market. They must also find a willing landlord. Thirteen states and about 60 localities have now passed source of income discrimination laws that prohibit landlords from discriminating against voucher holders (Scott et al, 2013, updated 2016), but enforcement is weak.

Households receiving a voucher can remain in their same unit as long as that unit meets the voucher program's quality standards. While households "leasing in place" receive less than the maximum housing subsidy if they live in units with monthly rents below the voucher's payment standard, around 20 percent of newly issued voucher holders use the voucher to lease their previous unit (Finkel and Buron, 2001), perhaps reflecting inertia or the challenge of finding willing landlords and acceptable units in the specified time period.

Neighborhoods Reached by Housing Choice Voucher Holders

One of the original motivations for establishing the voucher program was its potential to help low-income families reach neighborhoods that offer better schools and greater opportunities for economic advancement. Research shows that vouchers have had some modest success in achieving this potential. On average, voucher holders live in less disadvantaged neighborhoods than the residents of public or other HUD-assisted housing (Hartung and Henig, 1997; Kingsley et al., 2003; Pendall 2000; Devine et al., 2003) and also in slightly less disadvantaged neighborhoods than the average poor household (Pendall, 2000; Wood, Turnham and Mills, 2008; Galvez, 2011). That said, voucher holders still live in very disadvantaged neighborhoods, with higher poverty rates than the neighborhoods surrounding developments subsidized through the Low Income Housing Tax Credit, now the largest federal vehicle for low-income housing production (McClure, 2006). Further, families with vouchers live near to lower-performing schools than other poor families (Ellen, Horn and Schwartz, 2014).

It is unclear why voucher holders don't get to better neighborhoods. One possible explanation is that voucher households, given their many pressing needs, do not prioritize neighborhoods. Given time pressure, they may instead use their subsidy to move out of crowded living situations (Wood et al., 2008), write down rent burdens (Mills et al., 2006), or find larger, higher quality homes (Rosenblatt & DeLuca, 2012). Further, they may prioritize certain aspects of neighborhoods, such as lower crime rates or proximity to family and friends (Lens, Ellen, & O'Regan, 2011; Desmond 2012). If people choose to locate near family and friends, and disadvantaged individuals tend to have disadvantaged social networks located in higher poverty neighborhoods, then this may restrict where voucher families look for housing.

Constraints may play a role as well. Most notably, voucher holders may find relatively few units with rents below the local, metropolitan area-wide payment standard in low-poverty neighborhoods, as these neighborhoods tend to be more expensive. That said, research suggests that voucher holders could afford many units in lower poverty areas. Ellen, Horn and Schwartz (2014) show that the average voucher holder lives in a neighborhood with a lower performing school than the average neighborhood where housing units renting below the FMR are located. McClure (2013) finds that one quarter of units that rent below the FMR are in neighborhoods with poverty rates below 10 percent.

Table 1 shows the average poverty rate of the neighborhoods where housing choice voucher holders live and compares it with the average poverty rate for neighborhoods affordable under voucher program payment standards. We use census tracts to proxy for neighborhoods. The table shows these figures averaged over all metropolitan areas and for the 20 metropolitan areas with the largest number of voucher holders. The poverty rate of the average neighborhood where voucher holders live (HCV poverty exposure) is consistently higher than the overall metropolitan area poverty rate. Voucher holders are also far less likely than other households to reach low poverty tracts, or tracts with a less than 10% poverty rate. Only 15% of voucher holder households in metropolitan areas nationwide live in such tracts, compared to 43% of the population overall. One potential driver of these discrepancies is the location of housing units that rent under the FMR. To test this possibility, we consider tracts as affordable to voucher holders if the median gross rent is below the FMR. The second-to-last column of the table shows that even if voucher holders chose neighborhoods at random among the affordable tracts and thus ended up in the average "affordable" tract, they would live in tracts with lower poverty rates in most cities. Moreover, the affordable tracts include many low-poverty neighborhoods. Across all metropolitan areas, the poverty rate of the neighborhood at the 25th percentile among affordable tracts is just 12 percent, which is actually lower than the average metropolitan area poverty rate. Thus, while affordability constrains choice, many tracts offer both relatively low rents and low poverty rates.

Landlords may also discriminate against voucher holders. Many landlords avoid the voucher program and refuse to house voucher holders (Rosen 2014). Owners of buildings in high-rent neighborhoods may be especially likely to refuse voucher holders because they have the option of charging the same (or higher) rent to other tenants (Rosen 2014; Collinson and Ganong, 2016). While it is difficult to test for discrimination directly, a few local studies have found evidence of such discrimination (Luna and Leopold, 2013). For example, a paired tester study in Washington, DC found significant discrimination against voucher holders, despite the fact that Washington, DC Human Rights Act has prohibited discrimination based on source of income since 1977 (Thabault and Platts-Mills, 2006).

Finally, voucher holders may have limited information about alternatives (Rosen 2014). And most relevant for us, they may rely on social networks and ties to learn about potential neighborhoods and homes. Rosenblatt, DeLuca, and Wood (2013) emphasize that low-income households facing limited options and time tend to turn to their networks to find information about available homes. These networks tend to be comprised of low-income individuals who refer them to homes that are near to their homes and typically located in high-poverty neighborhoods. Thus, they argue that residents end up living near to relatives and friends more as a result of expediency than preferences.

We focus our analysis on second and subsequent moves by voucher holders, which would be made under less immediate time pressure. Figure 1 demonstrates the relatively small number of neighborhoods to which voucher holders relocate in practice. The sample here is restricted to households who relocate across tracts within their metropolitan area between 2011 and 2014. The set of tracts available to moving voucher holders is taken to be those which have at least one voucher holding household residing in them. On average across all metropolitan areas, 60% of cross-tract moves are to just one fifth of available tracts (and just 15% of all tracts). A few studies suggest that voucher holders reach slightly better neighborhoods when making a subsequent move. For example, Eriksen and Ross (2013) find that voucher holders tend to move towards lower poverty neighborhoods after a few quarters in the program. Similarly, Feins and Patterson (2005) report that families who move after entering the voucher program choose neighborhoods that have slightly lower levels of poverty and higher homeownership rates than their original neighborhoods.

We explore the degree to which social networks among voucher holders appear to shape the destinations of these moves. In particular, we explore whether voucher holders who live very close to other voucher holder movers tend to move to the same neighborhoods as those other voucher holders. While this is not a definitive test of the role of social networks, finding that relocation to the same neighborhood is more likely among voucher holders who live in close proximity would suggest the presence of social connections through which information is shared.

Data and Analytical Strategy

We use the Housing Choice Voucher (HCV) Family Report records, a restricted household-level dataset from the Department of Housing and Urban Development (HUD) that contains information on the demographic attributes of households enrolled in the HCV program, their sources of income, and their rent payments in each year that they remain in the program. In addition, the HCV Family Report records provide us with the exact address of the home where the housing voucher was used in each year. We geocode each address to obtain its latitude and longitude as well as the Census block group and Census tract in which it is located.

Our analyses examine residential moves that took place between 2011 and 2014. We focus on the 1,458,423 households living in a Core Based Statistical Area (CBSA) who were enrolled in the HCV program in both 2011 and 2014.[2] We are able to geocode the addresses for 1,412,181 of them (96.8%). Among those, we keep the 1,369,401 households who remained in

[2] We are not restricting the enrollment in the program to be continuous from 2011 to 2014. To be included in our sample, a household must be in the HCV program in 2011 and in 2014.

the same CBSA as of 2014.[3] We then identify the subsample of movers, defined as voucher holders who moved to a different census tract between 2011 and 2014 (but remained in the same CBSA) and for whom we have complete demographic information. We further restrict the sample to those movers who lived in a tract with three or more voucher households in 2011 who also moved across tracts between 2011 and 2014. Our final sample of movers includes 272,329 voucher households initially living in 24,475 census tracts in 765 metropolitan and micropolitan areas.[4]

We supplement the HCV data with data from the American Community Survey that we use to compute poverty rates for Census tracts, and the rental vacancy rate and a measure of racial segregation in each CBSA in 2011. We also use tract-level data from HUD to measure the relative intensity of labor market engagement, the exposure to environmental hazards, and the accessibility to public transit of the Census tracts of origin and destination.

Estimation of Social Interactions in Relocation Decisions

The first part of our analysis examines the presence of social influence effects in relocation decisions of HCV households. Without an experimental design, identifying the presence and size of social interaction effects is empirically challenging.[5] The fundamental challenge in identifying the presence of social influence effects is that homophily, the common characteristics shared by close neighbors, can produce the same pattern of decisions as would arise from the exertion of influence. Similar households are more likely to make similar decisions as each other even if no influence is present. Therefore, observing that two households from the same originating neighborhood moved together could be attributed to a potential social influence effect, but it could also be a reflection of sorting into neighborhoods on the basis of unobserved household attributes. Research designs relying on observational data will be unable to account for all sources of selection and endogeneity that could drive any co-location decision observed in the data.

To partially address the non-random sorting of households into neighborhoods, we build on an estimation strategy proposed by Bayer, Ross, and Topa (2008). In their study of referral effects in the Boston metropolitan area, Bayer et al. (2008) estimate the propensity that two individuals living in the same neighborhood also work in the same location. To do so, they compare pairs of individuals who live in the same Census block to pairs of individuals who live in nearby Census blocks that are still part of the same Census block group.[6] The idea is to

[3] Three percent of voucher holders move to homes in a different CBSA. Because in some of the models we examine how conditions in the CBSA where the vouchers are used interact with our explanatory variables, we drop households who move across CBSAs.
[4] Approximately, 10% of voucher households move to another neighborhood every year. Therefore, the majority of households excluded from our sample are households that remained in the same census tract between 2011 and 2014.
[5] See Manski (1993) and Durlauf (2004) for a general discussion of identification of social interactions and Angrist (2014) and Sacerdote (2014) for recent reviews of the literature on estimation of peer effects.
[6] Bayer et al. (2008) use the term “reference group” to indicate the geographic area within which the comparisons of pairs from different Census blocks are made. They experiment with two reference groups for a given Census block: the Census block group where the Census block belongs and the set of 10 closest blocks (captured using physical distance between Census block centroids).

control for a reference group meant to proxy for the unobservable similarities of nearby households, while zeroing in on the decisions of even more proximate neighbors. Given the thinness of the market for owner occupied housing at the block level, the key assumption underlying their identification strategy is that selection into Census blocks within the same Census block group is ignorable. Their tests of that assumption indicate that sorting on the observable attributes that they measure is minimal. They find that living in the same versus nearby Census blocks increases the probability of working together by 33 percent.

Building on their design, we estimate whether voucher households living in close proximity to one another in 2011 are more likely to move to the same Census tract between 2011 and 2014 than other pairs of movers in the same initial census tract. Specifically, we ask whether a pair of households living within 50 feet of each other in 2011 is more likely to co-locate to the same Census tract in 2014 than another pair of mover households from the very same Census tract of origin but who lived further away from each other. If co-location is higher among pairs of movers who live within 50 feet from each other, this finding would suggest the presence of social interactions. The assumption underlying this research design is that residential sorting on the basis of unobserved attributes within Census tracts is minimal or ignorable. In a series of indirect tests of that assumption, we show that variation in observable characteristics such as race and income is driven largely by cross-tract rather than within-tract differences.

Formally, our statistical model to test social interactions in relocation decisions takes the following from:

$$M_{ij\,c} = \alpha_c + \delta_{50}\, B50_{ij\,c} + \delta_{100}\, B100_{ij\,c} + \delta_{250}\, B250_{ij\,c} + \delta_{500}\, B500_{ij\,c} + \delta_{1000}\, B1000_{ij\,c} + \varepsilon_{ij\,c} \quad (1)$$

where *i* and *j* denote mover households and *c* denotes a Census tract in 2011. $M_{ij\,c}$ is a binary indicator that takes on value 1 if the pair of households *i* and *j* originating from tract *c* in 2011 moved to the same tract in 2014, and 0 otherwise; α_c is a set of Census tract fixed effects; $B50_{ij\,c}$ is a binary indicator that takes on value 1 if the pair of households *i* and *j* lived within 50 feet or less from each other in tract *c* in 2011, and 0 otherwise; $B100_{ij\,c}$ is a binary indicator that takes on value 1 if the pair of households *i* and *j* lived between 50 and 100 feet from each other in tract *c* in 2011, and 0 otherwise; $B250_{ij\,c}$ is a binary indicator that takes on value 1 if the pair of households *i* and *j* lived between 100 and 250 feet from each other in tract *c* in 2011, and 0 otherwise; $B500_{ij\,c}$ is a binary indicator that takes on value 1 if the pair of households *i* and *j* lived between 250 and 500 feet from each other in tract *c* in 2011, and 0 otherwise; $B1000_{ij\,c}$ is a binary indicator that takes on value 1 if the pair of households *i* and *j* lived between 500 and 1,000 feet from each other in tract *c* in 2011, and 0 otherwise; and $\varepsilon_{ij\,c}$ is an idiosyncratic error term for the pair of households *i* and *j* originating from tract *c*.

Given the specification in Equation (1), the comparison group for pairs of movers in categories B50 to B1000 are all pairs of mover households in the same Census tract living more than 1,000 feet from each other in 2011. Therefore, the coefficients on each of the buffer indicators, δ_{50} to δ_{1000}, will estimate the increase in the probability of co-location with respect to that of pairs of mover households of the same tract living more than 1,000 feet from each other in 2011. Rejecting the null hypothesis that one or more of the δ coefficients are zero will suggest the presence of social interactions within the corresponding distance. Equation (1) is estimated using a linear probability model.

Our design differs from Bayer et al.'s (2008) in the way that reference groups are defined. They compare pairs of individuals living in the same Census block to pairs of individuals living in nearby blocks that are part of the same Census block group. By construction, their design ignores potential social interactions occurring between individuals living across the street from each other but who happen to live in different Census blocks. Our design with buffers of different sizes circumvents this problem by allowing for social interactions based on physical distance strictly, regardless of where the administrative boundaries of Census blocks are drawn. We also can test whether social interactions are stronger among voucher holders who live next door to one another or in the same building. Furthermore, by adding indicators for each of the buffer distances, we allow for the possibility that the strength of social interactions may decay as members of the pairs live farther away from each other.

To estimate our model, we generate a dataset that includes all unique combinations of paired mover voucher holders living in the same Census tract in 2011. As noted, our base sample includes 272,329 voucher households living in 24,475 Census tracts in 2011. We then generate all unique pairs of these households within each tract.[7] This yields a sample of 3,323,611 mover pairs. In Table 2, we show that, on average, there are 11 mover households and 136 unique combinations of mover pairs per tract.

We measure the physical distance between households in a pair using the latitude and longitude of their location in 2011.The distribution of the distance between pairs of voucher holders in the same tract in 2011 is shown in Table 3. The first row shows the distribution of distance between all pairs in the same census tract, and each of the rows below shows the distribution of distance between pairs within the different buffers. The mean of this distance is 2,189 feet, with a standard deviation of 2,533 feet. As for the distribution, 10 percent of pairs in the same tract live within less than 67 feet from each other and another 10 percent live within 4,805 feet or more from each other. For pairs living within 50 feet from each other, we find that the mean distance between them is 4 feet and that more than 75 percent of such pairs live in the same building.

In Table 4, we show the mean and standard deviation for our outcome. On average, 4.6 percent of pairs living in the same tract in 2011 moved to the same tract between 2011 and 2014. Table 4 also shows the share of pairs that fall in each of the buffer distances that we define. We find that 9.3 percent of pairs live within 50 feet or less from each other, 1.8 percent live between 50 and 100 feet, 3.3 percent live between 100 and 250 feet, 7.1 percent live between 250 and 500 feet, and 14.7 percent live between 500 and 1,000 feet from each other. Nearly two thirds of the mover pairs in our sample live more than 1,000 feet from each other.

[7] The resulting number of all possible unique pairs of households in a given tract obeys a sequence of triangular numbers of the form $\frac{(n-1)^2+(n-1)}{2}$, where n is the number of households in the Census tract.

Table 5 shows the characteristics of the mover pairs in our sample. For each household, we have data on race of the household head, number of dependents, gender of the household head, annual income, age, and building type. We construct all possible combinations of these attributes that can arise among pairs of households and show shares for the full sample and for pairs living in each of the buffer distances. We find that in 77.4 percent of all pairs, both households are racial minorities, in 54.1 percent of pairs, both households have at least one dependent, and in 75 percent of pairs, both are female-headed households.[8] The mean difference in annual income between the two households in a pair is $8,544. The average age difference between the household heads in a pair is 13.2 years. Finally, we find that in 33.5 percent of pairs, both households live in a single home, and that in 19.4 percent of the pairs both households live in an apartment building. Note that Table 5 shows that these means and shares appear to be quite similar across all buffer distances. While not conclusive, this suggests that pairs of households living in the 0-50ft buffer are not very different from pairs living more than 1,000ft from each other, at least in terms of this set of observed characteristics. In the next section, we test more rigorously for the presence of sorting on the basis of some of these attributes.

Testing for the presence of sorting on observables

To interpret any proximity effects as indicating the presence of social influence on household relocation decisions, we have to assume that sorting occurs across tracts but not within tracts. Bayer et al.'s (2008) justification of the latter assumptions relies on the thinness of the market for owner occupied housing at the blockgroup level, which is the geographic unit they use to proxy for neighborhoods. Specifically, they argue that because there are so few homes on the market within a blockgroup, homebuyers have little ability to choose which block they will live on within the blockgroup. Our study considers the market for rental housing, which, even at small geographies is not as thin, given higher turnover levels. Yet we focus on voucher holders, and the effective rental market for voucher holders is far thinner than for all renters. Only certain housing units within a given census tract will be available to voucher holders given the local payment standard and certification requirements, and thus voucher holders will be limited in their ability to choose individual blocks or buildings within a census tract. Beyond these institutional arguments, we can look for evidence of sorting based on observable characteristics. In Table 6, we decompose the variance of key household attributes into the contribution from within tract and between tract differences. The majority of variance in all the characteristics considered is due to across tract differences. This is particularly true for income, where 88 percent of the variance is due to cross-tract differences. This implies that tract fixed effects will be a strong control for any unobservables correlated with income that drive sorting among voucher households within census tracts. This does not speak directly to the level of sorting within tracts though so we turn to that next.

Within tract sorting would imply that household pairs become more alike the closer together they live. In Table 7, we test for whether this is apparent for observable characteristics at the sub-tract level. We examine whether pairs of households living in the 0-50ft buffer systematically differ from pairs living in other buffers of the same tract on the basis of differences in income, age, race, presence of dependents, and gender of the household head. To

[8] Racial minorities include all non-white households: non-Hispanic black, Hispanic, non-Hispanic Asian, and other non-white racial groups.

do so, we separately regress each of the following demographic attributes of the paired households on the set of buffer indicators described in Equation (1) and the set of tract fixed effects: the absolute difference in annual income between household heads in the pair (in thousands of dollars), the absolute difference in the age of the household heads in the pair, a binary indicator for whether the household heads in the pair are of the same race, a binary indicator for whether the two households in the pair have any dependents, and a binary indicator for whether the two household heads in the pair are female. We find that, relative to pairs in the same tract living more than 1,000 feet from each other, pairs living within 50 feet tend to be more similar in terms of annual income and age, and are more likely to be of the same race. In terms of presence of dependents and gender of the household head, we find that pairs living within 50 feet are no different from pairs living beyond 1,000 feet from each other. If we extend these comparisons to other buffers, we find that as the distance between the two households in the pair increases, these households tend to be more different in terms of income, age, and race.

While finding no statistically significant differences would be most convincing, the differences that Table 7 reveal are of small economic significance. At the bottom of the table, we report the mean of each of the demographic attributes that we examine for pairs living more than 1,000 feet from each other. If we compare the relative size of the point estimates for the 0-50ft buffer to those means, we find that none of these differences is larger than 10 percent. Focusing on income, voucher holding households in the sample have average incomes of about $11,500, but with significant variation: the average difference in income between a pair of voucher households is $8,700. The estimate in Table 7 shows that for near neighbors the typical difference in income falls, but only to $8,000. We conclude from this test that while there exists some sorting at the sub-tract level, the magnitude of this sorting is likely to have small implications for our identification strategy. Further, we will demonstrate the robustness of our results to directly controlling for these observable demographics. These tests on observables are encouraging that the social influence effects we seek to measure are not confounded by household sorting within tracts.

We demonstrate the robustness of our estimates in three additional model specifications. First, we estimate Equation (1) controlling for a rich set of attributes of the pairs. This strategy will account for the fact that, within the same census tract, pairs living within 50 feet from each other are more likely to be of the same race or have similar incomes than pairs of the same tract who live 1,000 feet or more from each other. Second, we estimate Equation (1) narrowing the reference group to households living in the same block group. By reducing the size of the reference group, the assumption of no sorting below the reference group level becomes even more plausible.[9] In our last robustness test, we take advantage of the fact that each household appears paired multiple times with other households in the same tract and estimate an individual fixed effects model. This model specification replaces the set of tract fixed effects with a set of fixed effects for the first member of the pair and a set of fixed effects for the second member of

[9] Changing the reference group from tracts to block groups changes the sample size. Because the combinations of unique pairs are constructed within the reference group, smaller reference groups will yield a smaller number of possible pairs. When using tracts as the reference group the sample includes 3,323,611 unique pairs. When using block groups as the reference group the sample includes 1,600,842 unique pairs. In both instances, we keep reference groups that have at least three households who moved to another tract between 2011 and 2014.

the pair.[10] That said, any sorting at the sub-tract level along unobservable characteristics not correlated with observables and relevant for housing search prowess or preference for specific neighborhoods is a challenge to identification of social influence effects from geographic proximity.

Results

In Column 1 in Table 8 we estimate Equation 1 to test for the potential presence of social influence effects in housing choice voucher holder relocation decisions. As noted, the unit of analysis is a mover pair, each of which is constructed from the set of voucher households in a tract who moved to a different tract in the same metropolitan area between 2011 and 2014. Each possible pair of voucher movers within each tract is a separate observation. The outcome variable is one if both households in a given pair moved to the same tract as each other and zero otherwise. Tract fixed effects are included to account for the average likelihood that such paired moves occur between any two households in the same tract. The results suggest that for mover households living within 50 feet of each other (in the same or adjacent buildings) in 2011, the likelihood of ending up in the same tract is about 2 percentage points higher than for other pairs of movers households who also both start out in the same tract but live 1,000 feet or more from each other. The average likelihood in the sample of such paired moves is 4.6 percent, so proximity increases this likelihood by about 40 percent. In other words, the odds of two mover households in the same tract in 2011 moving to the same tract increase from 1 in 22 to 1 in 15 when they live within 50 feet of each other. These proximity effects are still detectable though diminishing out to 250 feet.[11]

Column 2 in Table 8 shows results when we control for differences in the attributes of the pair including race, presence of dependents, income, age and gender of the household head and initial building type. Pairs of households where neither have dependents are significantly more likely to relocate to the same tract. Greater differences in income or age between households make paired moves to the same tract less likely. The proximity effect estimate is only slightly affected by the addition of these controls though, falling from 1.9% to 1.7%.[12] Column 3 In

---

[10] The individual fixed effects model has the following form: $M_{ij} = \alpha_i + \gamma_j + \delta_{50}\, B50_{ij} + \delta_{100}\, B100_{ij} + \delta_{250}\, B250_{ij} + \delta_{500}\, B500_{ij} + \delta_{1000}\, B1000_{ij} + \varepsilon_{ij}$, where α_i and γ_j are the two sets of individual fixed effects, M_{ij} is an indicator for whether the pair moved to the same tract in 2014, and the buffer indicators B50 to B1000 have the same interpretation as in Equation (1). Given the large number of dummy indicators included in α_i and γ_j, the estimation of such a model is very computationally demanding. To speed up the computation, we take advantage of the standard result in the Frisch–Waugh–Lovell (FWL) theorem, which enables the estimation of a fixed effects model by group-demeaning the left- and right-hand sides of the equation and using these demeaned variables in a regression without the fixed effects dummies. Greene (2003, pp. 291) shows how the FWL result can be used in the case of two-way fixed effects to recover algebraically equivalent estimates.

[11] We experimented with re-estimating our results with an even smaller proximity buffer of 10 feet (not shown). In practice, this means that the pair of moving households were in the same building in 2011. As shown in Table 3, in practice more than three quarters of household pairs within 50 feet of each other are in fact in the same building. Comparing the coefficients on the 10-foot and the 50-foot buffers, we see that the proximity effects seem to be predominantly driven by households residing within the same building, though the effect for neighbors in adjacent buildings is still detectable.

[12] Table 5 reveals that while demographic differences between household pairs are consistent for households living in close proximity and those more than 1,000 feet apart, this does not hold true for building type. Pairs both in single

Table 8 shows results when we narrow the reference group from pairs living in the census tract to pairs living in the same block group. If there is meaningful sub-tract level sorting by households, this specification should show a smaller proximity estimate as the reference group will be composed of households with even closer similarity. The 50ft buffer estimate does fall in this model, though only to 1.6%. Finally, column 4 in Table 8 shows results when we replace the set of tract fixed effects with the two sets of individual household fixed effects. These should account for any idiosyncratic unobservable differences in the propensity for paired moves by certain households who happen to live close to other voucher holders. The 50ft buffer estimate is reduced further, though only to 1.5%. We interpret this robustness to reference group and demographic controls as evidence that sub-tract level household sorting does not explain the proximity effect and that some social influence on relocation decisions by near neighbors appears to be present.

Social influence effects and local housing market conditions

If these estimated proximity effects in fact reflect social networks in home search then they should be stronger in cities where such search is more costly, or in tighter housing markets. Table 9 shows that there is substantial variation in vacancy rates in the 765 CBSAs in our sample, ranging from less than 6 percent of vacant housing units in some cities to more than 28 percent in others. In Table 10 we consider whether the estimated proximity effects are larger in metropolitan areas with lower vacancy rates, by interacting the buffer distance indicators with metropolitan area vacancy rates. For a point of comparison, we reproduce our baseline results in column 1. Column 2 shows the interaction between the buffer indicators and the vacancy rate in the CBSA. Focusing on the coefficients for the 50-foot buffer, we see that the proximity effects are largest in low vacancy cities. A one standard deviation increase in a city's vacancy rate decreases the overall proximity effect by about one third of the average effect estimated in column 1. This is consistent with the reliance on one's network on relocation choices being most pronounced when housing markets are tightest.

Neighborhood segregation can also play a role in restricting effective housing choices and making apartment search more costly, especially for black and Latino households. Thus, by the same logic as with vacancy rates, we expect the reliance on one's network in home search to be more pronounced in more segregated cities. To test this possibility, we compute a white/non-white dissimilarity index for each metro area.[13] As Table 9 shows, there is large variation in this measure of racial segregation across CBSAs. In column 3 in Table 10, we interact this measure with the distance buffers. As predicted, proximity effects for very near neighbors increase with metro area level neighborhood segregation. A one standard deviation increase in the dissimilarity index increases the proximity effect by about 40 percent of the average effect estimated in column 1. Some of the more distant buffers become statistically significant with these

---

family homes are more limited within the 50-foot buffer while pairs both in apartment buildings are rarely more than 1000 feet apart. We estimate our model separately for pairs that are both in single family homes, both in two-family homes, both in low-rise apartment buildings, both in high-rise apartment buildings, and both townhouses. We find that our point estimates for the 0-50ft dummy range from 0.10 to 0.21 across these subsamples of pairs. Results from this estimation are reported in Table A1 in the Appendix.

[13] This index measures the extent to which tract level racial composition deviates from metropolitan area level racial make-up, where zero represents a completely integrated metropolitan area and one represents a completely segregated one.

interactions included, but at the mean level of dissimilarity (.34) the overall effects are close to zero. These interaction results also provide an indirect test against confounding by sub-tract level sorting. There is little reason to expect sorting to be more prevalent in cities with tighter housing markets. That fact that proximity effects increase in predictable ways in such cities then further indicates that sorting is not the main driving of the results.

Exploring heterogeneity by demographic attributes of the pair

A priori, we expect households to be more likely to form social ties with neighbors who are demographically similar to themselves. In Table 11, we interact the distance buffers with a set of variables measuring differences between the characteristics of households in the pair. For simplicity, we only report estimates for the 0-50ft buffer indicator, the interactions between the 0-50ft buffer indicator and demographic attributes of the pair, and for the non-interacted demographic attributes. Column 1 examines heterogeneous effects by racial composition of the pair. We distinguish between pairs in which both households are of a racial minority, pairs in which one household is non-Hispanic white and the other is a racial minority, and pairs in which both households are non-Hispanic white. We find that the social influence effects are smaller in pairs that include at least one minority household, relative to pairs in which both households are non-Hispanic white.

Column 2 examines whether the presence of dependents in the household moderates the social influence effects. We distinguish between pairs in which none of the households have any dependents, pairs in which only one household has dependents, and pairs in which both households have dependents. We find that social influence effects are substantially larger among pairs in which neither of the households have any dependents, relative to pairs in which both households have dependents. While one may be inclined to think that the presence of children should lead to stronger social ties among neighbors, our findings may reflect the role of differing school choices for these households. Column 3 examines heterogeneous effects by gender of the household head. We find that for pairs in which both household heads are male the social influence effects are substantially larger. However, pairs in which both household heads are males represent only a 3 percent of all pairs in our sample, as shown in Table 5.

Column 4 examines heterogeneous effects by income (absolute difference in 1000s of USD). We find that social influence by near neighbors is less likely as income differences between households in a pair increases. For example, the effect for two households with income difference of $10 thousand is two thirds that of two households with the same income. Column 5 examines heterogeneous effects by age (absolute difference in years). Social influence also falls with age difference. For every 10 years of age difference, the social influence effect diminishes by 0.15 percentage points. If we add the two age difference coefficients, we find that the social influence effect within the 50-ft buffer would decrease 1.02 percentage points for pairs that are 35 years apart, as compared to pairs who are the same age. Column 6 examines heterogeneous effects by building type matches. Social influence effects are most apparent for pairs of households both in single homes. [14] This could in part reflect the smaller number of neighbors overall within 50 feet for single homes. In such single-family neighborhoods, each nearby

[14] Apartment includes low- and high-rise. Single home includes single-family homes, two-family homes, townhouses, and manufactured homes.

voucher holder represents a large fraction of a household's total stock of near neighbors and would thus be more likely to form a social tie with that household.

Finally, the juxtaposition of the large interaction coefficients with minimal main effects provides a particularly strong case against unobservable sub-tract level sorting of households. Demographic similarity, particularly for income and age, only helps explain neighborhood relocation choices if households originally live near one another. It is difficult to tell a story in which households sort into the same blocks or buildings on the basis of some unobservable characteristic which is uncorrelated with observable attributes overall but is then correlated with demographic differences in the same way within each buffer.

Relocation decisions and neighborhood quality

While the potential role of networks in housing search is of some interest on its own, our key interest lies in understanding whether these networks help to explain the concentration of voucher holders in high poverty neighborhoods. In particular, we would like to know whether network effects appear to be leading households to live in lower quality neighborhoods than they would otherwise. We examine the following indicators of neighborhood quality: poverty rate, labor market engagement, exposure to environmental hazards, and access to public transit. We measure the change in the poverty rate between the tracts of origin and destination using 5-year estimates from the American Community Survey from 2007-2011 and 2010-2014. We measure changes in labor market engagement, exposure to environmental hazards, and access to public transit using HUD's recently released Affirmatively Furthering Fair Housing (AFFH) data.[15] These indices range from 0 to 100 with larger values being more desirable. We provide further description of these measures in the Appendix.

In addition to examining changes in neighborhood quality, we also test whether the presence of network effects leads to changes in the quality of homes where households live, as it could be the case that households are trading off housing and neighborhood quality. For example, moves to lower quality neighborhoods might enable households to afford housing units of higher quality or of larger size. As suggested in Rosenblatt and DeLuca (2012), voucher holders may prioritize unit quality over neighborhood quality. We proxy for unit quality with the gross rent that landlords ask and by single-family homes, which are typically larger and offer more land than apartments in multifamily buildings.[16]

For this part of the analysis, we shift from comparing pairs of households to examining differences across individual households. As before, we focus on households who moved to a different tract between 2011 and 2014. Our sample includes 271,748 households for whom we

---

[15] The documentation for the AFFH data can be found here: https://www.hudexchange.info/resources/documents/AFFH-Data-Documentation.pdf. See also Ellen, Horn, and Kuai (2017) for more discussion of the usefulness of these metrics.

[16] Our measure of single-family homes includes two-family homes as well. In particular, it includes the following building types: single-family home, two-family home, townhouse, and manufactured home. We don't have data on the square footage of the units, which would be the best measure of unit size. We do know the number of bedrooms in the unit, but a change in number of bedrooms is likely to reflect a change in household composition, as the number of bedrooms a household qualifies for is determined by the housing authority according to the number of children in the household.

have data on all the outcomes that we examine.[17] For each household, we identify whether they moved from the same tract of origin in 2011 to the same tract of destination in 2014 as another voucher holder. In addition, for each household that moved to the same tract as another household from the same tract of origin, we identify whether the distance between them was smaller or larger than 1,000 feet. Given this distinction, each household in our sample of movers falls in one of these three mutually exclusive categories: the household moved to a tract to which no one else from the tract of origin moved, the household moved to another tract to which at least one other household living within 1,000 feet in the tract of origin also moved, and the household moved to another tract to which at least one household living beyond 1,000 feet in the tract of origin also moved. Table 12 shows that 41.5 percent of households moved to another tract to which someone else from the tract of origin also moved. For half of them, at least one of these households who moved to the same tract was living within 1,000 feet in the tract of origin.

Given this setup, we model change in neighborhood and housing unit quality between a moving household's origin and destination tract as a function of their distance to other movers from the same tract. Specifically, we model changes in neighborhood and housing unit quality as follows:

$$\Delta Y_{ic} = \alpha_c + \beta_1\, MovLT1000_{ic} + \beta_2\, MovGT1000_{ic} + \beta_3\, X_{ic} + \varepsilon_{ic} \quad (2)$$

where *i* denotes households and *c* denotes tracts. ΔY_{ic} is the change in neighborhood or housing unit quality, which we examine using the following measures in separate regressions: change in tract poverty, change in labor market engagement, change in exposure to environmental hazards, change in access to public transit, change in gross rent, and residence in a single home. $MovLT1000_{ic}$ is an indicator for whether at least one household from tract of origin *c* living within 1,000 feet from household *i* moved to the same tract in 2014. $MovGT1000_{ic}$ is an indicator for whether at least one household from tract of origin *c* living more than 1,000 feet away from household *i* moved to the same tract in 2014. X_{ic} includes a set of demographic controls for the household measured in 2011. We again control for tract fixed effects in all specifications to account for unobservable differences in household sorting across tracts as before, as well as to deal with the mechanical relationship that households who move from very high poverty tracts will tend to end up in lower poverty tracts.

In column 1 of Table 13, we consider the change in poverty exposure for any household who moved to the same tract as at least one other household from the origin tract. Poverty exposure for this group increased on average by 2.5 percentage points relative to households who moved away from the same tract but did not end up in the same neighborhood as any other voucher holder from their original neighborhood. In column 2, we add the set of household characteristics available in the data and find that our estimate remains unchanged. In columns 3 and 4, we specify the model represented in Equation 2 and find that these effects are pronounced for voucher households who initially live close to one another. Specifically, we find that the increase in tract poverty is larger for households who lived within 1,000 feet of another

[17] There are 272,343 households that meet the two criteria to be included in our sample: (1) the household moved to a different tract between 2011 and 2014, and (2) the tract of origin had at least three households moving out from the tract. Among these households, we have data on all outcomes for 271,748 of them. The missing households are those living in tracts for which HUD did not produce one or more of the indices that we use in our analyses of neighborhood quality.

household who moved to the same tract of destination. Those households experience an increase in tract poverty of 2.9 percentage points. Households living more than 1,000 feet from another household who moved to the same tract of destination experience an increase in tract poverty of 2.2 percentage points. As before, we find that these estimates are robust to adding household controls.

While these results suggest that network effects are potentially guiding households to worse neighborhoods, our findings could be driven by common shocks that households experienced in the neighborhoods of origin. A fire or a building collapse may force all households in the building to move out at the same time and possibly relocate together in the same tract of destination. If that was the case for a large share of households in our sample, we would observe a higher likelihood of moving to the same tract even in the absence of network effects. To account for this, we exclude households that moved out from buildings that lost all of their voucher holders between 2011 and 2014 and re-estimate models 3 and 4.[18] As shown in columns 5 and 6, our estimates remain the same after we exclude households that were potentially affected by a common shock.

Table 14 extends this analysis to other measures of neighborhood and unit quality. For each outcome, we re-estimate model 4 from Table 13. For a point of comparison, we reproduce the estimates for change in tract poverty in column 1. In line with our findings on poverty exposure, we find that households who were potentially affected by network effects end up in neighborhoods with lower levels of labor market engagement (column 2) and higher exposure to environmental hazards (column 3).[19] However, these households end up in tracts that have greater access to transit, suggesting that these moves lead them closer to the city center (column 4). In columns 5 and 6, we examine whether network effects potentially guide households to units of higher quality. We use change in the gross monthly rent as one possible indicator of unit quality. This specification also controls for the change in tract poverty rate between the origin and destination to isolate any increase in quality from a likely decrease in rents for comparable units in higher poverty neighborhoods. We find that households living closer to other households who moved to the same tract seem to rent in units that are of slightly better quality (column 5), but the magnitude is quite small (less than $10 per month).[20] We also look at the probability of renting a single home in 2014, as these types of units are typically larger (even conditional on number of bedrooms). We find that households that were potentially affected by network effects

[18] We identify buildings that had at least 3 voucher households living in them in 2011 but had no voucher households in 2014. We assume that households moving out from those buildings between 2011 and 2014 did so as a result of a common shock (a fire, a demolition, and so on). There are 7,048 of such households in our sample.
[19] The sample size in column 3 is smaller because the HUD index of exposure to environmental hazards is not available for all tracts.
[20] Model 5 controls for changes in the tract poverty between the neighborhood of origin and the neighborhood of destination. To evaluate whether the small effect on rents is a mechanical effect due to bunching of rents at the local payment standard, in Figures A1 and A2 in the Appendix we look at the distribution of the rent to payment standard ratio and how it changes with moves. There is no obvious bunching at the payment standard with most households renting at below 90% of the payment standard. Further, Figure A2 and the summary statistics on changes in gross rents in Table 12 show that in general, plenty of voucher holder moves do result in much higher rents implying at least in part an improvement in housing unit quality. This suggests that the finding that socially influenced moves do not tend to result in materially higher rents is not mechanical.

are less likely to rent this type of unit (column 6).[21] These results together show that potentially socially influenced moves lead households on average to lower opportunity neighborhoods along a number of dimensions. Such households do tend to experience better transit access and increased unit quality after relocation though the magnitude of these improvements is minimal.[22]

Conclusion

We test for the presence of social influence in the relocation decisions of housing choice voucher holders. Using a strategy based on proximity of neighbors while controlling for a broader reference group, we find evidence consistent with social influence effects being present in the neighborhood choices of voucher holders. Specifically, pairs of neighbors living within the same or adjacent buildings are about 40 percent more likely to relocate to the same neighborhood as each other, relative to pairs of neighbors from the same tract who live more than 1,000 feet away from each other. This tendency is greater in cities with tighter housing markets and more segregated neighborhoods, indicating that nearby neighbors are relying on their network connections to a greater degree when search costs are higher. Finally, we show that voucher holders whose relocation decisions are potentially socially influenced end up, on average, in higher poverty neighborhoods than other voucher holders moving from their same origin tract.

To be sure, we do not directly observe social interactions, and it is possible that pairs of voucher holders living close to one another also receive more similar information about available housing options through external channels. For example, landlords may share information with tenants about other buildings that they own or that others in their network own. We cannot distinguish between such potential information channels and social influence.

We also cannot definitively say that voucher holders are made worse by these social influence effects. Reduction in search costs is valuable as is maintenance of geographic proximity to members of one's network. We have motivated the social influence effects as operating through information exchange regarding neighborhoods made affordable with a voucher, there are other aspects of housing search that could plausibly be shared through these networks. Voucher holders may put more of a premium on unit quality and size rather than the surrounding neighborhood. Households guided by social networks may be led to larger, higher quality apartments. Unfortunately, we have no direct measures of housing quality or size, though evidence from rents does not support the idea that voucher holders that move to the same tracts as nearby neighbors move into nicer homes. Voucher holders may also gain information from near neighbors about which units and buildings meet quality standards and are owned by landlords willing to accept vouchers. This information could reduce search times, though it is not obvious why it would lead voucher holders to move to systematically higher poverty neighborhoods.

---

21 The dependent variable in Model 6 is not measured in changes. Instead, we regress an indicator for living in single family home in 2014 on the predictors shown in Equation 2 and an indicator for being in a single home in 2011.

22 We also estimated the specification in column 5 with interactions for the number of years in the voucher program (not shown) since earlier moves might be the most likely to be aimed at getting higher unit quality, but found no evidence of this.

At the very least, we believe these results suggest that providing more information and guidance to voucher recipients regarding the array of neighborhoods made newly affordable to them, and thereby improving the housing search information embedded in their social network, could broaden the set of destination neighborhoods and lessen the exposure to neighborhood poverty. Other approaches like shifting to Small Area Fair Market Rents, which peg voucher subsidies to rents at the ZIP Code level rather than the metropolitan area, might also open more options for voucher holders.

References

Angrist, J. D. (2014). The Perils of Peer Effects. *Labour Economics*, *30*, 98-108.

Bayer, P., Ross, S. L., & Topa, G. (2008). Place of work and place of residence: Informal hiring networks and labor market outcomes. *Journal of Political Economy*, *116*(6), 1150-1196.

Brown, L. A., & Moore, E. G. (1970). The intra-urban migration process: A perspective. Geografiska Annaler. Series B. Human Geography, 52(1), 1–13.

Chetty, R., Hendren, N., & Katz, L. (2015). The Effect of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment. NBER Working Paper.

Collinson, R., Ellen, I. G., and Ludwig, J. (2015). Low-Income Housing Policy. NBER Working Paper.

Collinson, R., & Ganong, P. (2016). Incidence and Price Discrimination: Evidence from Housing Vouchers. Working Paper.

DeLuca, S., Garboden, P. M., & Rosenblatt, P. (2013). Segregating shelter: How housing policies shape the residential locations of low-income minority families. *The ANNALS of the American Academy of Political and Social Science*, *647*(1), 268-299.

DeLuca, S., & Rosenblatt, P. (2010). Does Moving To Better Neighborhoods Lead to Better Schooling Opportunities? Parental School Choice in an Experimental Housing Voucher Program. Teachers College Record, 112, 1443–1491.

Deng, L. (2007). Comparing the Effects of Housing Vouchers and Low-Income Housing Tax Credits on Neighborhood Integration and School Quality. Journal of Planning Education and Research, 27, 20–35.

Desmond, M. (2012). Disposable Ties and the Urban Poor. *American Journal of Sociology*, *117*(5), 1295-1335.

Devine, D. J., Gray, R. W., Rubin, L., & Taghavi, L. B. (2003). Housing Choice Voucher Location Patterns: Implications for Participant and Neighborhood Welfare. Washington, DC: U.S. Department of Housing and Urban Development.

Durlauf, S. N. (2004). Neighborhood Effects. *Handbook of Regional and Urban Economics*, *4*, 2173-2242.

Ellen, I.G., & Karfunkel, B. (2016). Renting in America’s Largest Metropolitan Areas. New York, NY: NYU Furman Center for Real Estate and Urban Policy.

Ellen, I. G., & Horn, K. M. (2011). Creating a Metric of Educational Opportunity for Assisted Households. Report prepared for the Department of Housing and Urban Development.

Ellen, I. G., Horn, K. M., and Kuai, Y. (2017). Opportunity for what and for whom? Variation in neighborhood access among income housing tax credit tenants. Working paper.

Eriksen, M. D., & Ross, A. (2013). The Impact of Housing Vouchers on Mobility and Neighborhood Attributes. Real Estate Economics, 41, 255–277.

Feins, J. D., & Patterson, R. (2005). Geographic Mobility in the Housing Choice Voucher Program: A Study of Families Entering the Program, 1995–2002. Cityscape, 8, 21–47.

Finkel, M., & Buron, L. (2001). Study on Section 8 Voucher Success Rates. Volume I. Quantitative Study of Success Rates in Metropolitan Areas. *Prepared by Abt Associates for the US Department of Housing and Urban Development*, 2-3.

Ford, J.S., Rutherford, R.C., and Yavas, A. (2005). The Effects of the Internet on Marketing Residential Real Estate. *Journal of Housing Economics*, 14, 92-108.

Galvez, M. (2011). Defining "Choice" in the Housing Choice Voucher Program. Dissertation Manuscript.

Greene, W. H. (2003). *Econometric Analysis*. New Jersey: Prentice Hall.

Hartung, J. M., & Henig, J. (1997). Housing Vouchers and Certificates as a Vehicle for Deconcentrating the Poor: Evidence from the Washington, D.C., Metropolitan Area. Urban Affairs Review, 32, 403–419.

Horn, K. M., Ellen, I. G., & Schwartz, A. E. (2014). Do Housing Choice Voucher holders live near good schools? Journal of Housing Economics, 23, 28–40.

Jacob, B. A., Kapustin, M., & Ludwig, J. (2015). The impact of housing assistance on child outcomes: Evidence from a randomized housing lottery. Quarterly Journal of Economics, 130, 465–506.

Kan, K. (2007). Residential mobility and social capital. Journal of Urban Economics, 61, 436–457.

Kling, J.R., Liebman, J.B., and Katz, L.F. (2007). Experimental Analysis of Neighborhood Effects. *Econometrica*, 75(1), 83-119.

Kingsley, T. G., Johnson, J., & Pettit, K. L. S. (2003). Patterns of Section 8 Relocation in the HOPE VI Program. Journal of Urban Affairs, 25, 427–447.

Krysan, Maria. 2008. "Does Race Matter in the Search for Housing? An Exploratory Study of Search Strategies, Experiences, and Locations." Social Science Research 37(2): 581-603

Krysan, Maria and Michael Bader. 2009. “Racial Blind Spots: Black-White-Latino Differences in Community Knowledge.” Social Problems 56(4): 677-701.

Krysan, Maria, Kyle Crowder and Michael D.M. Bader. 2014. “Pathways to Segregation.” Pp. 27- 63 in Annette Lareau and Kimberly Goyette (eds), *Choosing Homes, Choosing Schools: Residential Segregation and the Search for a Good School*. New York: Russell Sage Foundation.

Lareau, Annette, and Kimberly Goyette. 2014. “Preface.” in Annette Lareau and Kimberly Goyette (eds), *Choosing Homes, Choosing Schools: Residential Segregation and the Search for a Good School*. New York: Russell Sage Foundation.

Lens, M., Ellen, I. G., & O’Regan, K. (2011). Do Vouchers Help Low-Income Households Live in Safer Neighborhoods? Evidence on the Housing Choice Voucher Program. Cityscape, 13, 135–159.

Luna, Jessica and Josh Leopold. (2013). Landlord Discrimination Restricts the Use of Rental Vouchers. Washington, DC: The Urban Institute. http://www.urban.org/urban-wire/landlord-discrimination-restricts-use-rental-vouchers

Manski, C. F. (1993). Identification of Endogenous Social Effects: The Reflection Problem. *The Review of Economic Studies*, *60*(3), 531-542.

Marsh, A. and Gibb, K. (2011). Uncertainty, Expectations and Behavioral Aspects of Housing Market Choices. *Housing, Theory and Society*, 28(3), 215-235.

McClure, K. (2006). The Low-Income Housing Tax Credit Program Goes Mainstream and Moves to the Suburbs. *Housing Policy Debate*, 17, 419–446.

McClure, K. (2013). Which metropolitan areas work best for poverty deconcentration with Housing Choice Vouchers? *Cityscape*, *15*(3), 209-236.

Mills, G., Gubits, D., Orr, L., Long, D., Feins, J., Kaul, B., Wood, M., Jones, A., Cloudburst Consulting Associates, & the QED group. (2006). The Effects of Housing Vouchers on Welfare Families. Washington, DC: U.S. Department of Housing and Urban Development, Office of Policy Development and Research.

Pendall, R. (2000). Why Voucher and Certificate Users Live in Distressed Neighborhoods. Housing Policy Debate, 11, 881–909.

Rosen, Eva (2014). Selection, Matching, and the Rules of the Game: Landlords and the Geographic Sorting of Low-Income Renters. Working Paper, W14-11. Joint Center for Housing Studies, Harvard University.

Rosenbaum, J. E. (1995). Changing the geography of opportunity by expanding residential choice: Lessons from the Gautreaux program. Housing Policy Debate, 6, 231–269.

Rosenblatt, P., & DeLuca, S. (2012). We Don’t Live Outside, We Live in Here: Neighborhoods and Residential Mobility Decisions Among Low-income Families. City and Community, 11, 254–284.

Rosenblatt, P., DeLuca, S., and Wood, H. (2013). Why Poor People Move (and Where They Go): Residential Mobility, Selection and Stratification. Paper Presented at NYU Law School. http://www.law.nyu.edu/sites/default/files/ECM_PRO_074751.pdf

Sacerdote, B. (2014). Experimental and quasi-experimental analysis of peer effects: two steps forward? *Annual Review of Economics*, *6*(1), 253-272.

Scott, Molly, Mary Cunningham, Jennifer Biess, Jennifer O’Neill, Phil Tegeler, Ebony Gayles, and Barbara Sard. 2013. *Expanding Choice: Practical Strategies for Building a Successful Housing Mobility Program*, Appendix B, Updated 2016. Washington DC: Urban Institute and Poverty and Race Research Action Council. http://www.prrac.org/pdf/ExpandingChoice.pdf

Susin, S. (2002). Rent Vouchers and the Price of Low-Income Housing. *Journal of Public Economics* 83(1): 109-152.

Thabault, Isabelle M. & Eliza T. Platts-Mills. 2006. Discrimination Against Participants in the Housing Choice Voucher Program: An Enforcement Strategy. Poverty and Race. Washington, DC: Poverty and Race Research Action Council.

Wood, M., Turnham, J., & Mills, G. (2008). Housing affordability and family well-being: Results from the housing voucher evaluation. Housing Policy Debate, 19, 367–412.

Appendix

We use a series of measures of neighborhood quality from the Department of Housing and Urban Development's (HUD) Affirmatively Furthering Fair Housing (AFFH) data tool. These data were created to help HUD grantees conduct an Assessment of Fair Housing (AFH) planning process.[23] Below, we show how the three indices that we use were constructed.

HUD labor market engagement index

HUD constructed a labor market engagement index from American Community Survey 2006-2010 data. The index incorporates the unemployment rate (*u*), labor-force participation rate (*l*), and share of the population with a bachelor's degree or higher (*b*). The index is constructed as follows:

$$LME_i = \left[\left(\frac{u_i - \mu_u}{\sigma_u}\right) * -1\right] + \left(\frac{l_i - \mu_i}{\sigma_l}\right) + \left(\frac{b_i - \mu_b}{\sigma_b}\right)$$

where the Census tract means of unemployment, labor-force participation, and share of the population with a bachelor's degree or higher (μ_u, μ_l, μ_b) and the corresponding standard errors (σ_u, σ_l, σ_b) are estimated over the national distribution. The value for unemployment rate is inverted, so that higher numbers correspond to lower unemployment. The values are percentile ranked nationally and range from 0 to 99.

HUD environmental hazard exposure index

The environmental health index is constructed using National Air Toxics Assessment (NATA) data from 2005. The index creates a measure of potential exposure to harmful toxins at the neighborhood level. It includes estimates of carcinogenic (*c*), respiratory (*r*) and neurological (*n*) air quality hazards. The index uses the following formula:

$$EnvHazard_i = \left[\left(\frac{c_i - \mu_c}{\sigma_c}\right) + \left(\frac{r_i - \mu_r}{\sigma_r}\right) + \left(\frac{n_i - \mu_n}{\sigma_n}\right)\right] * \ -1$$

where *i* indexes Census tracts. The index is calculated as a linear combination of *c*, *r* and *n*, where means of the three hazards (μ_c, μ_r, μ_n) and the corresponding standard errors (σ_c, σ_r, σ_n) are estimated over the national distribution. Values are percentile ranked from 0 to 99 and inverted (higher values correspond to lower exposure to environmental hazards).

HUD transit access index

23. For more information and documentation see https://www.hudexchange.info/resources/documents/AFFH-Data-Documentation.pdf

The transit access index is constructed using data from transit agencies that provide data through General Transit Feed Specification (GTFS) Exchange. The index assesses the relative accessibility to amenities via bus or train within a metropolitan area. Data from the Local Employment Dynamics dataset are used to obtain the number of jobs in retail (NAICS codes 44-45), arts entertainment and recreations (NAICS code 71), and food and accommodations (NAICS code 72) in block-group. Counts of these jobs are used as proxies for the prevalence of amenities at the block-group level. HUD identifies and sums the jobs in each of these sectors within ½ of a mile of each bus stop and ¾ of a mile of each rail transit stop. For each trip in the transit system, HUD calculates a stop-specific measures of the additional amenities accessed in each ensuing stop on that route, which is then divided by the additional travel time to each ensuing stop. Formally, the index is computed by first computing the accessibility of stop *i* on trip *j*, S_{ij}, as follows:

$$S_{ij} = \sum_{i}^{N} \frac{a_{i+1}}{T_{t+1}}$$

where *a* is the number of jobs defined above, and *T* is the marginal travel time with each stop. In other words, each stop of each trip is assigned a value equal to the sum of the amenities (proxied with the number of jobs in the NAICS categories 44-45, 71, and 72) of each ensuing stop divided by the time to that next stop for all stops on a trip.

These stop-journey specific values are summed over all journeys *j* made in 24 hours (where journeys in opposite directions are counted as two trips) as follows:

$$A_i = \sum_{j}^{k} S_{ij}$$

To translate the stop accessibility values (A_i) to block-groups, HUD calculates the distance between each stop and the population weighted centroid of each block-group. The three highest accessibility stops within ¾ of a mile are summed to generate a block-group value for accessibility. These values are places into decile buckets within metro areas and are scaled up by a factor of 10 to produce an index that ranges from 0 to 99. Larger values indicate higher accessibility to transit. To produce measures of accessibility at the tract level, we compute the average of the index across all block groups in the tract.