Parental Involvement and Neighborhood Quality: Evidence from Public Housing Closures in Chicago

Joel Kaiyuan Han*

This version: April 1, 2018
(most recent version)

Abstract

This paper studies whether, and how, parents respond to changing neighborhood quality through parental involvement. Empirical measurement of this response is complicated by neighborhood selection: through their residential location choices, parents exercise some control over neighborhood quality. To deal with neighborhood selection, this study makes use of the mass closures of public housing projects in Chicago. Displaced households relocated into other neighborhood in Chicago. Incoming relocations are associated with increased crime rates in receiving neighborhoods. By treating relocations received as unanticipated declines in neighborhood quality, this study estimates the causal effect of neighborhod quality on parental involvement. On average, parents compensate for decreased neighborhood quality by increasing parental involvement. However, the parental response is heterogeneous: parents with low initial involvement decrease parental involvement instead. These results suggest that existing estimates of neighborhood effects may underestimate the direct impact of neighborhoods on children, and provide evidence on an interesting spillover effect of public housing closures on other neighborhoods.

Keywords: Parental involvement, neighborhood effects, public housing
JEL Classification: I38, J13, R23

*Loyola University Chicago. Email: jhan12@luc.edu. This study was supported financially by the Institute for Research on Poverty Dissertation Research Fellowship, as well as the Department of Economics Fellowship at UW-Madison. This research benefited from the guidance of my dissertation committee - Steven Durlauf, Christopher Taber, and Jesse Gregory - as well as from seminar participants at UW-Madison, the National University of Singapore, and Singapore Management University. All remaining errors are mine.

1 Introduction

Differences in parenting behavior have been linked to differences in child behavior and development (Cunha et al., 2010; Pavan, 2013). Depending on the specific context, these differences may either reinforce or offset existing disparities between children. While parenting behavior has been found to be responsive to various changes in the external environment (Gelber and Isen, 2013; Avvisati et al., 2014; Fryer Jr et al., 2015; Bergman, 2015), little is known about how it responds to changes in the quality of the residential neighborhood.

In this paper, I consider how parental involvement responds to changes in neighborhood quality. Parental involvement refers to the time and effort a parent expends on the care, discipline and supervision of the child. Neighborhood quality is a catch-all term that refers to the overall influence of the neighborhood environment on child development. This paper attempts to answer the following questions: Do parents respond significantly to changing neighborhood quality through parental involvement? If so, what direction does this response take? A priori, it is unclear whether the changes will reinforce or offset the impact on children. The answer to this question depends on how neighborhood quality affects the marginal productivity of parental involvement, but also on how neighborhood quality affects the effort cost of parental involvement. Finally, do the answers to the above depend on parental characteristics, especially those associated with relative advantage/disadvantage? If different parents respond differently to the same change in neighborhood quality, this might imply an unequal impact of neighborhood change on children living in the same neighborhood.

When attempting to estimate the causal effect of neighborhood quality on parental involvement, the main difficulty arises from the parent's selection of residential neighborhood. A parent who chooses good neighborhoods for her child is likely to be one who exercises a high level of involvement. The neighborhood selection problem is further complicated by changing levels of parental motivation, which may also influence whether parents move out of or remain in a given neighborhood. For example, a parent who loses a job may become depressed, leading to lower parental involvement. Simultaneously, the inability to pay rent may force a move to a lower quality neighborhood. Along with job loss, other unobserved changes in the family's situation will generate a spurious correlation between parental involvement and neighborhood quality.

To address the neighborhood problem, I utilize a series of shocks to neighborhood composition, which were caused by the mandated closure and demolition of high-rise public housing projects in Chicago. Specifically, residents from closed housing projects were relocated into various other neighborhoods in Chicago, changing the composition of those destination neighborhoods. I study the effect of these changes on families whose neighborhoods received an influx of relocations. Henceforth, I address these families as **receiving families**, and their neighborhoods as **receiving neighborhoods**. These receiving families were by no means the only group affected by the public housing closures: families living close to a public housing closure also experienced a change in neighborhood quality, as did the displaced families themselves.[1] The analysis is restricted to receiving families for two reasons. First, the effect of relocations on receiving neighborhoods is simpler to understand: the only effect was to change neighborhood composition. In contrast, closure neighborhoods were affected by the eviction and demolition process, the change in neighborhood composition, as well as the subsequent redevelopment of the site. These disparate channels could have affected neighborhood quality in different ways, and the overall impact on neighborhood quality may be unclear. A second reason is that the data used in this study contains no families displaced by closures, and only a small number of families living in neighborhoods where a closure occurred.

This study uses data from the Project for Human Development in Chicago Neighborhoods (PHDCN), which contains longitudinal observations of parental involvement and residential neighborhoods. Using restricted neighborhood identifiers, PHDCN data is linked with neighborhood measures of crime rate and relocations received. Using the linked neighborhood data, I construct an instrument for neighborhood quality. This instrument is based on the spillover effects of public housing closures on receiving neighborhoods. Given the unpredictable timing of public housing closures, it is likely that these spillover effects represent unanticipated changes to receiving neighborhood quality. As a result, the instrument is unlikely to be correlated with the residential neighborhood selection of sampled parents.

To address the neighborhood selection problem as fully as possible, two additional features are implemented in the research design. First, the instrument is measured from the family's initial neighborhood, whether or not the family chose to remain there in subsequent survey waves. This feature exploits the fact that the initial neighborhood cannot be changed by the parent's subsequent residential location choices. By directly examining the residential location choices of families in the data, I verify that the instrumental variable appears to be uncorrelated with neighborhood selection. Second, instead of actual relocation flows, the instrument only measures likely relocation flows based on historical minority composition. This feature is motivated by a concern of reverse causation: that families displaced by closures may have chosen their destination neighborhoods based on observed parental involvement trends in those neighborhoods.

[1] Figure A1 provides a graphical description of the affected families.

I find that the overall intensity of parental involvement increases when neighborhood quality falls. Compared to zero relocations, parents in a neighborhood receiving the mean yearly level of relocations exhibited 0.1 standard deviations greater intensity of parental involvement.[2] Using crime rate as the measure of neighborhood quality, I find that a one standard deviation increase in the crime rate (decrease in neighborhood quality) is associated with a quarter standard deviation increase in parental involvement. The effect size is substantial - around half of the parental involvement gap between college educated and less educated parents. In contrast to the significant findings for parental involvement, the evidence for effects on child behavior is weaker and less consistent. Together, these two findings are consistent with the hypothesis that parents diminish the negative impact of neighborhood decline on their child by increasing parental involvement.

Additionally, I explore heterogeneity in the response of parental involvement. Most notably, the response of parental involvement to declining neighborhood quality depends strongly on the initial level of parental involvement. Parents with high initial involvement subsequently raise parental involvement in response to neighborhood quality declines. In contrast, parents with the lowest initial involvement instead make reinforcing decreases to parental involvement. I relate these findings to a model of parental involvement, where parents differ in their altruism levels. Under some assumptions on the nature of heterogeneity in parental involvement, this finding suggests that parental care and neighborhood quality are substitute inputs in the production of good child behavior. Consequently, if parental engagement policies are jointly implemented with place-based policies designed to improve the neighborhood, the effectiveness of each individual policy is likely to be reduced.

These findings provide more evidence that parents are important mediators of external influences on their children[3]. The magnitude of the response suggests that parental involvement is a major channel through which residential neighborhoods affect families. Additionally, these findings also relate to the study of neighborhood effects on child development, as a potential explanation for the surprisingly small impact of certain neighborhood interventions. For example, in the Moving to Opportunity experiment, Kling et al. (2007) use randomly-assigned housing vouchers to estimate the causal effects of moving to a better neighborhood. They find improvements in the mental health of female but not male children, and virtually no improvement to educational achievement. Studying the same experiment, Chetty et al. (2016) find more consistent positive impacts on subsequent adult-life incomes of exposed children, as does Chyn (2016). Furthermore, the duration and age of exposure to neighborhood influences also seems to play a major role (Wodtke et al.,

[2]The measure of parental involvement has been standardized by age group of the child, so the increase in parental involvement is measured relative to other parents with children of the same age.

[3]Cunha et al. (2010); Pavan (2013)

2011). In all of these studies mentioned, the indirect effect of parental behavior on child development is not accounted for. Assuming that parental involvement is productive in developing good outcomes in the child, then my results suggest that the direct impact of neighborhoods quality on child outcomes - holding parental involvement constant - could be more positive than previously thought.

Additionally, as a description of policy impact, the findings reveal an interesting spillover effect of public housing closures. This spillover effect has not perviously been highlighted: other studies have documented the effects of public housing closures on displaced families (Jacob, 2004; Chyn, 2016), and more generally, the effects of housing interventions on disadvantaged families (Andersson et al., 2016; Jacob et al., 2015). While Aliprantis and Hartley (2015), as well as by Sandler (2012), have documented spillover effects on receiving neighborhood crime rate, this study extends that work by showing that these neighborhood-level changes had significant impacts on behavior within the family. The findings suggest that, even if children in receiving neighborhoods were not severely affected by incoming relocations, this may have come at a significant cost to their parents in terms of effort and stress. If these parents are themselves a vulnerable population, then it may be important to devise policy responses to mitigate the impact of public housing closures on the receiving neighborhood.

This study is closely related to Patacchini and Zenou (2011) (henceforth referred to as PZ), who investigate whether neighborhood and parental inputs behave like complements or substitutes.[4] PZ do not observe longitudinal changes in neighborhood quality or parental involvement, instead, their identifying variation comes from random assignment of council housing location in the United Kingdom. Their overall findings are very different from those in this study: where I find that parental involvement increases in response to a neighborhood quality decline, they find the opposite effect. One major differences is that PZ study a sample that has selected into public housing.[5] To my knowledge, this study is the first to estimate neighborhood quality effects on parenting behavior, using a sample that has not selected into public housing.

The rest of the paper proceeds as follows. Section 2 provides some theoretical foundations for the response of parental involvement to neighborhood quality. Next, Section 3 describes the data used, paying special attention to the construction of the parental involvement variable. Section 4 describes the public housing closures and the resulting relocations: these are an important source of variation in neighborhood quality. Section 5 explains the empirical strategy for estimating neighborhood quality effects, and presents the main

[4]PZ also consider the contribution of both these factors to educational attainment of the child, while this current work is limited to concurrent measurements of behavior.

[5]Other potential reasons for the different are discussed in this paper, after the main results have been presented.

results. Section 6 then presents some evidence of heterogeneity in these neighborhood quality effects. Lastly, Section 7 concludes.

2 Expected Response of Parental Involvement to Neighborhood Quality

Neighborhood quality is likely to affect parental involvement through several channels. In this section, I summarize these channels using a simple model of parental utility. Abstracting from family structure, consider one-parent, one-child families living in a neighborhood with neighborhood quality q. This measure of neighborhood quality captures peer influences of other children and adults within the neighborhood, as well as the quality of amenities and institutions (e.g. schools) in the neighborhood.[6] Taking the level of neighborhood quality as given, parents choose the optimal level of parental involvement y. A high level of parental involvement contributes towards the development of good child behavior b, but exercising parental involvement requires costly effort c. Neighborhood quality affects parental decisions through child behavior, as well as parental effort cost.

$$U_i(b, y, q) = u(b(y, q)) - c(y, q) \tag{1}$$

The object of interest in this model is the response of optimal parental involvement to changes in neighborhood quality, which is given by $\partial y^*/\partial q$. Assuming an interior solution to the problem, this derivative can be expressed as (2). The denominator of (2) is the second order condition for utility maximization, hence, it is negative. As a result, the direction in which parental involvement changes is determined by the sum of three terms in the numerator. These terms correspond to channels proposed in the existing research literature.

$$\frac{\partial y^*}{\partial q} = -\frac{u'' b_y b_q + u' b_{yq} - c_{yq}}{u'' b_{yy} + u' b_{yq} - c_{yy}} \tag{2}$$

The first term, $u'' b_y b_q$, captures the effect of neighborhood quality on the marginal utility that parents derive from child behavior. The influence of neighborhood quality on child behavior has been extensively modeled: most models suggest a positive relationship between neighborhood quality and child behavior. This positive relationship can be transmitted through differences in institutional quality, changing adult influences, or changing peer influences from other children (Jencks and Mayer, 1990). Empirically, estimates

[6] Sacerdote et al. (2011) provides a review of the literature concerning school effects on child educational outcomes.

of peer effects support the hypothesis of a positive relationship, for a range of risky child/youth behaviors.[7] In turn, different levels of child behavior are likely to incentivize parental involvement to different extents. For example, the presence of parental expectations and aspirations (Fan and Chen, 2001) for child educational attainment suggests that that marginal utility of improved child behavior is decreasing: as it becomes more likely that the child achieves the target level of education, the marginal utility from improving child behavior goes down.

The second term, $u'b_{yq}$ captures the effect of neighborhood quality on the productivity of parental involvement. The sign of this term determines whether parental and neighborhood inputs are substitute or complement inputs in the production of child behavior. A priori, it is unclear which way neighborhood quality and parental involvement are related. Empirical evidence on this matter is also limited. By comparing families across schools with different resources, Hoffmann and Dufur (2008) find evidence that school resources serve as substitutes for parental involvement. However, their estimates are associative in the sense that selection into schools is not accounted for.

The final term, c_{yq}, captures the effect of neighborhood quality on parental effort cost. This term captures the effect of neighborhood quality on adult well-being, which in turn affects the parent's ability to exercise parental involvement. For example, the Moving to Opportuinty experiment found a reduction in adult stress level from an increase in neighborhood quality (Leventhal and Brooks-Gunn, 2003; Kling et al., 2007). A reduced stress level might, in turn, leave parents better able to care for their children. Parents may also have been affected through their expectations: Green et al. (2007) suggest that parental understanding of their role in child education, parental beliefs about the effectiveness of involvement, and parental understanding of external expectations are influential in driving different parental involvement decisions. Any changes in these would alter the marginal cost of parental involvement.[8]

3 Data Description

The data used in this study comes from the Project for Human Development in Chicago Neighborhoods (PHDCN). While this data has been widely studied within the social sciences (Morenoff et al., 2001; Sampson et al., 1997), it is relatively unfamiliar within the specific field of economics. Within this section, special

[7] These include smoking (Nakajima, 2007), alcohol, other drug use, dropping out of school (Gaviria and Raphael, 2001), absenteeism, disciplinary problems (Imberman et al., 2012), and crime Case and Katz (1991).

[8] Avvisati et al. (2011) present a thorough review of the literature regarding the effects and determinants of parental involvement, in the context of school quality.

attention is paid to the construction of the parental involvement variable, based on the available measures of parenting behavior contained in the PHDCN data.

The main subjects of the PHDCN are children/adolescents aged 0-18 years. In addition, the self-reported **Primary Caregivers** of these children were also surveyed. Henceforth, the terms parent and primary caregiver will be used interchangeably, and the term **family** will be used to denote a primary caregiver-child pair. Each family was surveyed up to three times, during the period from 1994 to 2002. At the time of each survey, the family's residential location was recorded. Information on residential location was recorded at the level of the **neighborhood cluster**, hence, this is the unit of analyis for neighborhoods adopted here. The term "neighborhood cluster" refers to an aggregation of adjacent, relatively similar census tracts: The average neighborhood cluster is about 2.5 census tracts in size, and contains around 8000 residents. Henceforth, the terms "neighborhood" and "neighborhood cluster" will be used interchangeably. Additionally, the neighborhood inhabited by the family at the time of the baseline (Wave 1) survey is referred to as the **initial neighborhood**.

The sample attrition rate is rather high: roughly 10 percent of the sample is lost in wave 2, and 20 percent in wave 3. However, the attrition is not strongly correlated with key family characteristics. Furthermore, attrition does not appear to be highly correlated with initial neighborhood. Figure 1 shows sample retention by mean income of the initial neighborhood. While respondents from poorer initial neighborhoods are slightly more likely to drop out of the sample, the difference in retention rate from the bottom to the top neighborhood quartile is at most 5 percentage points.

3.1 Measurement of Parental Involvement

The measure of parental involvement in this study is based on the features of the survey instrument: the Home Observation for Measurement of the Environment (HOME) inventory of the PHDCN. This battery of questions has been used in the National Longitudinal Study of Youth (NLSY) Child and Young Adult Survey, and is repeated in the PHDCN with slight modifications. The purpose of the HOME inventory is to measure the quality of the child's home and family life. To this end, the HOME inventory contains information about family practices regarding interaction with the child (e.g. conversations, outings) and discipline (e.g. TV usage, curfews).[9] This information was gathered through questions directed to the parent, as well as through direct observation by the interviewer. The number and type of parenting measures obtained differs by the age of the child: for example, parents of infants and toddlers were not asked about the presence of curfews.

[9]The exact behavior measures in each subset are listed in Table A1.

The available parenting measures from the HOME inventory are binary in nature. In order to aggregate these disparate measures into a single variable, I use the technique of multiple correspondence analysis (MCA). This procedure is a method for reducing the dimensionality of data. It is analogous to principal component analysis, but applied to discrete data. The first latent factor extracted from this procedure is used as the measure of **parental involvement**. This factor is a weighted sum of individual measure, which accounts for most of the observed differences in parenting behavior: more than 80 percent of the total variation is captured by this factor alone. At the same time, this factor has an intuitive interpretation: based on the signs of the estimated weights, this factor appears to capture the aggregate intensity of time and effort spent on child care. In other words, an increase in any individual parenting behavior will increase the measure of parental involvement.[10] Because the signs on the factor weights are unrestricted, this analysis also suggests that, in the PHDCN sample, the primary difference in parenting is one of overall intensity.

Because the exact parenting behaviors measured differ by child age, the multiple correspondence analysis is performed separately for children of different age ranges. Aside from the issue of non-measurement, another reason for doing so is that the importance of certain parenting behaviors might be age-dependent: for example, the interviewer-observed behavior of "encouraging child to contribute to the conversation" might be insignificant for the development of older children, and hence this measure would be uncorrelated with other measures of parental effort for older children. At the same time, for younger children, such behavior could be highly important in distinguishing more-involved from less-involved parenting.[11]

This constructed parental involvement measure bears many similarities to the measures of parental input in the NLSY, although it is not directly comparable. The NLSY measure of overall parental involvement (the HOME score) has been used as a measure of parental behavior: Todd and Wolpin (2007) use it as a measure of home inputs in the production of cognitive function in children. Relative to the HOME score, the parental involvement measure used here allows for each measure to be weighted differently, in a data-driven manner. In comparison, the HOME score is a simple additive index, weighting each individual parenting behavior equally.[12]

[10]The signs of these weights naturally depend on which direction of parenting behavior is defined as "more" parenting. For most behaviors, this choice is unambiguous. I have also chosen to treat avoidance of corporal punishment and harsh discipline as more intensive parenting. This choice reflects the notion that emotional restraint on the part of the parent is more costly than losing one's temper. Empirically, corporal punishment which has been associated with increased behavioral problems in children: such as aggressiveness and anger issues (Gershoff, 2002). In robustness checks, I show that the main results are not sensitive to the exclusion of all parental disciplinary behavior.

[11]In many studies, maternal language input appears to be positively associated with cognitive and vocabulary development in infants and toddlers: see, for example, Weisleder and Fernald (2013).

[12]Along with the main analysis, I will also present results using an additive index as the dependent variable. The main results are not changed by a different specification of parental involvement.

Common to the PHDCN and the NLSY, the measure of parental involvement covers a wide range of behaviors. In contrast, other research has focused on specific parental behaviors: Patacchini and Zenou (2011) use the frequency of reading to the child. On the other hand, both the PHDCN and NLSY measures of parental involvement share a drawback: they are insensitive to changes along the intensive margin of a particular parenting behavior. For example, both are able to capture the presence or absence of supervised homework time, but not an increase in the hours of supervised homework time.[13] This limitation should be kept in mind when interpreting the results.

3.2 Sample Restrictions

The following restrictions are placed on the sample: First, any families who were only observed once are excluded, since the identification of neighborhood effects comes from changes in parental involvement over time. Second, any families who left Chicago at any point during the survey are excluded. This group of families was not administered the full survey. As a result, there are several problems with missing information for these families. Finally, if the identity of the primary caregiver changed from one wave to the next, both waves are discarded from the sample.[14] Changes in the identity of the primary caregiver are likely to result in a change in residential neighborhood of the child; simultaneously, the different family structure is likely to result in changes in parental involvement as well. By excluding such families, any neighborhood quality effects can only be attributed to changes in parenting behavior, and not changes in the parent figure.[15]

After imposing all restrictions, the sample comprises roughly 2500 parent-child pairs (Table 1), each pair interviewed up to three times in the period 1994-2002. While some of these pairs are actually the same parent with multiple children, this information is not used in this study.

3.3 Descriptive Statistics - Families

Table 2 provides a basic description of the PHDCN sample. Among the self-reported primary caregivers, the black share is relatively high (0.34), but in line with the racial composition of urban Chicago. Roughly 40 percent of primary caregivers are employed, which is consistent with primary caregivers leaving the workforce to focus on child care. Table 2 also presents differences in mean baseline characteristics across different subsamples. Comparing families who moved out of the initial neighborhood during the study period

[13] In comparison, Guryan et al. (2008) have detailed measures of hours spent on each parenting activities.

[14] Suppose a child has the same primary caregiver reporting in waves 1 and 2, and a different one in wave 3. In that case, waves 1 and 2 are retained and wave 3 is discarded.

[15] This restriction may be concerning on the grounds that family composition could be endogenous to neighborhood quality: for example, deterioriating neighborhoods might lead parents to send children to live full-time with grandparents. For this reason, I also present results including changes in primary caregiver identity, and the effects are very similar.

with families who did not, those who moved appear to be more disadvantaged: the primary caregiver is more likely to be unemployed, earns less, and exercises less parental involvement. At the same time, children in such households perform worse on the Wide Range Achievement Test (WRAT), and are more likely to exhibit problematic behaviors related to internalizing and externalizing. [16] Comparing families from an affluent (above median income) initial neighborhood with those from poorer neighborhoods, the expected patterns of relative advantage appear. Finally, a comparison families with and without an educated (high school and above) parent reveals a similar pattern.

Table 3 describes the initial neighborhods of the families in the sample, in terms of neighborhood characteristics measured during the 1990 census. The sampled initial neighborhoods form a reasonably representative sample of neighborhoods in Chicago, in terms of average income, minority composition, and unemployment rate. As Column 5 of Table 3 shows, eight PHDCN neighborhoods contained a CHA public housing project. Five of those neighborhoods experienced a mandated closure. Since each initial neighborhood contains around 30 usable subjects, only about 150 respondents in the sample experienced the effects of a closure - too few for meaningful results on the effects of closures. For this reason, I focus on the effects of relocations instead of the effects of closures. [17]

4 Public Housing Closures and Relocations in Chicago

The closure of low-income public housing, and in particular the resulting relocations by displaced households, provide the key source of variation in neighborhood quality. In this section, I summarize the historical background of these events. Additionally, I show that the arrival of relocations in a particular neighborhood was associated with an increase in the crime rate. Finally, I show that incoming relocations were also associated with an increase in parental involvement, on the part of receiving parents.

4.1 Historical Background

In the 1990s, public housing in Chicago was notorious for a host of social ills, and the Chicago Housing Authority (CHA) was described as the "most troubled in the nation" (Terry, 1995). "Unemployment at

[16] The terms "internalizing" and "externalizing" are used widely in child psychology: internalizing refers to a group of withdrawn, inhibited, anxious or depressed behaviors. In contrast, externalizing refers to behavior that acts negatively on the external environment, including disruptive, hyperactive, and aggressive behaviors. This description of internalizing and externalizing behaviors is paraphrased from Liu (2004).

[17] In baseline analysis, I retain families in closure neighborhoods. These families were still affected by relocations to their neighborhood. In robustness checks, I show that the main results are unaffected by excluding all families who ever resided in a neighborhood containing a public housing project.

CHA developments ran as high as 90 percent... residents were at least twice as likely to be the victims of serious crime as other Chicagoans," (Chicago Tribune, 2002), and an estimated 95 percent of the resident population were on some other kind of public assistance (The Economist, 1998). Since the beginning of the 1990s, at least 20000 units have been demolished by the authorities, which constitutes over half of the 1995 stock of public housing. (Bennett et al., 2015). The first closures occurred in response to building-specific issues, rather than as part of a co-ordinated plan (Jacob, 2004). For example, four blocks of Cabrini-Green were closed and subsequently demolished due to a high-profile crime incident: a child was killed by a stray bullet fired from one of the blocks (Austen, 2012). Another reason for closing was poor/deteriorating building conditions, e.g. a building-wide loss of heat in the winter, linked to shoddy construction (Garza, 1999). Eventually, the scale of the problem prompted a federal takeover by the Department of Housing and Urban Development in 1995 (Terry, 1995). Shortly after, federal housing policy changes required viability testing of public housing blocks below a 90 percent occupancy, and around 18000 of Chicago's units failed this test (The Economist, 1998). At this point, plans were made for the widespread demolition of low-income housing in Chicago. This effort was aided by new federal legislation (HOPE VI), which funded the conversion of low-income housing to mixed income developments (Popkin et al., 2000). Ultimately, the CHA proposed the Plan for Transformation in order to carry out these objectives (CHA, 2011), and was returned from federal control to carry out this plan.

While the troubled history of Chicago public housing dates back to the 1970s (Popkin et al., 2000), the exact timing and volume of the studied closures was difficult to predict. The majority of closures during the study period (1994-2002) occurred before the year 2000: at this point, the closures were largely based on emergent issues and not an overall redevelopment plan. In addition, the federal takeover of CHA came largely as a surprise to the public (Terry, 1995), and the transition added to the uncertainty of the situation. While the last group of closures in the study period occur after the launch of the Plan for Transformation, they nonetheless occurred at the very start of its implementation.

A second difficulty in anticipating the timing of closures came from opposition by public housing residents and other members of the public. Even when official plans for demolition were made public, these legal and administrative challenges resulted in significant delays (Hunt, 2009), meaning that the timing of events did not coincide with stated plans. In the case of Cabrini-Green, a Local Advisory Council challenged parts of the proposed demolition plan, which resulted in a lawsuit (Bennett et al., 2015). The uncertainty generated by these challenges is sizeable: the Lathrop Homes were part of the Plan for Transformation, and again ordered demolished in 2006, but as of 2012, they were still standing, although almost completely vacant (Moser, 2012).

When mass closures eventually occcurred, residents of affected blocks were relocated in two ways. First, they were were offered Section 8 housing assistance and left to find housing on the private market. Alternatively, they were relocated to other public housing projects (Bennett et al., 2015). As a consequence, the destination neighborhoods of displaced households should be thought of as an intentional choice, made either by the displaced household or the public housing authority. At the same time, these choices are likely to have been constrained. For displaced househoulds, constraints include the spatial concentration of voucher housing (Oakley and Burchfield, 2009), discrimination from potential landlords, rushed relocations and insufficient information about destination neighborhoods (Thompson, 2006). Consistent with these claims, Desmond and Shollenberger (2015) find that involuntarily displaced renters in Milwaukee tend to settle in poorer-quality neighborhoods than those choosing to leave. When it came to resettlement organized by the public housing authority, the major constraint is likely to have been long wait-lists for housing at other low-income public housing developments.

4.2 Data on Closures and Relocations

Data on CHA closures is available at the address level,[18] and contains the yearly number of housing units (apartments) closed. The data is aggregated to the neighborhood cluster level. Hence, for each neighborhood cluster, I observe the precise number of housing units closed in a given year. The left panel of Figure 2 shows the variation in the number of housing units closed, over the two decade period between 1990 and 2010. Restricting attention to the time period of the PHDCN, over 10000 public housing units were closed between 1994 and 2002. During this period, the yearly volume of closures fluctuates considerably, and it is difficult to detect a consistent trend.

Table 3 describes PHDCN neighborhoods, CHA housing project neighborhoods, as well as CHA project closure neighborhoods, in terms of 1990 Census characteristics. For reference, Column 1 displays mean neighborhood characteristics for all neighborhoods in Chicago. Column 2 shows the same characteristics for PHDCN neighborhoods. Comparing Column 2 with Column 1, the neighborhoods sampled by the PHDCN appear to be fairly representative. Column 3 presents mean characteristics for neighborhoods containing a CHA housing project. As expected, neighborhoods containing public housing are - on average - poorer, higher in minority composition, and exhibit a higher unemployment rate, in comparison to the average Chicago

[18]This data was gathered partly by Brian Jacobs, and partly by Daniel Hartley.

neighborhood. Column 4 presents these same statistics for neighborhoods where a public housing closure occurred: these neighborhoods appear to be higher in minorities and more economically disadvantaged, even when compared to the average neighborhood containing public housing. This pattern suggests that closures occurred predominantly in public housing projects that were in the worst condition. Column 5 (resp. 6) describes the intersection of CHA (resp. closure) neighborhoods with the PHDCN neighborhoods. Column 5 shows that that only 8 neighborhoods sampled by the PHDCN also contained public housing, and Column 6 shows that only 5 of those neighborhoods experienced a closure during the PHDCN study period.

Data on yearly flows of relocations into receiving neighborhoods is not available. Instead, this study utilizes a CHA survey of ex-public housing residents, which contains the residential location of these households in 2008. This data is limited in two respects: first, this measure of relocations does not vary across time. Secondly, the geographical unit of measurement is a community area, which is a larger area than the neighborhood cluster.[19] While this level of geographical aggregation is large, the community areas correspond to well-known, distinct neighborhoods in Chicago.

Based on this data, I construct a shift-share measure of yearly relocations to a neighborhood j as follows: First, from the survey of ex-public housing residents, I calculate the share of total relocations received by each neighborhood cluster s_j. Since relocation shares are measured at the larger geographical area - the community area - I assign to each neighborhood cluster within the community area its population-weighted share of relocated households. This assumes that, relocations were uniformly distributed within each community area. The constructed shares are described in the right panel of Figure 2. The distribution contains many zeros: these mostly correspond to richer neighborhoods where rents are prohibitively high.

Next, I define the number of relocations to neighborhood j in year t as the product of yearly closures throughout Chicago $\sum_{j\in\mathcal{J}} c_{jt}$, and the neighborhood's share of relocations s_j (3). Finally, the raw number of relocations is divided by the neighborhood population, measured in 1994. Hence, the measure of relocations received by a neighborhood is a population-adjusted one.

$$r_{jt} = s_j \sum_{j\in\mathcal{J}} c_{jt}/n_{j,1994} \tag{3}$$

To describe the magnitude of the relocations, Figure 3 shows cumulative relocations received by each community area over the course of the sample period, as a fraction of year 1994 population.[20] Again, a fair

[19] There are 80 community areas covering Chicago, compared with 343 NCs.

[20] By construction, there is no variation in relocations received per resident, between different neighborhood clusters in the

number of neighborhoods receive no relocations at all, but there is also a long right tail to the distribution: conditional on receiving a non-zero share of relocations, the mean level of relocations per resident was around 3 percent. For the sake of comparison, across a 10 year time span (1990-2000), the mean absolute change in neighborhood population was roughly 12 percent of base year population. Hence, the population movements induced by the public housing closures appear to be non-negligible, relative to other factors affecting neighborhood population.

4.3 Effect of Relocations on Neighborhood Quality

To show how incoming relocations affected the receiving neighborhood, I focus on one particular aspect of neighborhood quality: the neighborhood crime rate. The neighborhood crime rate is the best available measure of neighborhood quality, because it is measured at a high (yearly) frequency and at a high level of geographical detail. The impact of incoming relocations on receiving neighborhood crime rate is well-documented. Using precise block-level information on relocation destinations for a subset of displaced households, Aliprantis and Hartley (2015) find significant increases on crime rates for most major types of crime. The average magnitude of such effects varies from 3 to 30 percent of base year crime rates, depending on the specific type of crime. Using tract level relocations data, Popkin et al. (2012) find similar increases ranging from 0.7 to 0.9 percent in violent/property/gun crime, associated with 1 relocated household received per 1000 existing residents.

Since the measure of relocations used in this paper is relatively less precise, I verify that the effects on neighborhood crime rate are qualitatively similar to those found in the previous studies. As shown in (4), neighborhood crime rate v_{jt} is regressed on lead, current, and lagged values of relocations received. Neighborhood fixed effects F_j and year effects G_t are included as controls.

$$v_{jt} = \sum_{k=-3}^{3} \rho_k r_{j,t-k} + F_j + G_t + \varepsilon_{jt} \tag{4}$$

Figure 4 presents the coefficients ρ_k from (4). There is little evidence of any association between neighborhood crime rate in a given year and subsequent relocations to that neighborhood. This negative finding is reassuring: a significant association would raise concerns of reverse causation, specifically that displaced public housing residents were choosing destination neighborhoods based on trends in neighborhood quality. In contrast, there is a significant, positive association between neighborhood crime rate and relocations re-

same community area.

ceived in the same year, as well as 1 and 2 years before. The latter finding is consistent with the notion that incoming relocations decrease neighborhood quality, but this effect appears to fade out after two years. Relative to Aliprantis and Hartley (2015), the magnitude of the coefficient is lower: one relocation per thousand is associated with an increase of roughly 2.5% in the neighborhood crime rate. Since the effects are being measured over a larger geographical area, the attenuated effect is not surprising.

Besides the crime rate, other aspects of neighborhood quality are also likely to have been changed by incoming relocations. These other effects are difficult to document: other aspects of neighborhood quality are not as well-measured as crime rate. Nevertheless, since it is unlikely that any of these dimensions were significantly improved by incoming relocations, it is likely that overall neighborhood quality was decreased by incoming relocations.

4.4 Effect of Relocations on Parental Involvement

In order to determine the spillover effects of public housing closures, I first analyze the effect of relocations on parental involvement. The regression model for this analysis is given by (5). The outcome variable is parental involvement, which is still denoted by y. Here, parental involvement is indexed by family (i) and year (t). The coefficient of interest is ρ, which measures the effect of relocations to family i's current (year t) neighborhood in the previous year (year $t-1$). Two versions of the relocations variable r_{jt} are used: the raw number of relocated households as well as the ratio of relocated households to neighborhood population. X'_{it} denotes a vector of controls: across all specifications, this vector includes parental employment status, marital status, and occupancy type. Additionally, all specifications include family level fixed effects F_i, to control for unobserved family characteristics that remain constant over time.

$$y_{it} = \rho r_{j(i,t),t-1} + X'_{it}\beta + F_i + \varepsilon_{it} \tag{5}$$

Table 4 presents the estimated effects from this regression model. Columns 1 through 3 present the effects of relocations, measured as raw numbers, while Columns 4 through 6 present the effects with relocations per resident. Within each set, the leftmost column (Column 1 and 4) shows the effects of relocations with the minimal set of controls. Relocations received by the neighborhood in the previous year are associated with an increase in parental involvement in the current year. Assuming that the overall effect of relocations on neighborhood quality is negative, this result suggests that parents compensate for lower neighborhood quality by increasing the intensity of parental involvement.

One alternative explanation for the positive coefficients is one of selective residential moves. Because relocations were predictably concentrated in relatively disadvantaged neighborhoods, parents with increasing motivation to take care of their children may have moved to neighborhoods that received less relocations. Alternatively, parents facing more difficulty in taking care of their children may have moved to neighborhoods receiving more relocations. To assess this explanation, the next pair of regressions (Columns 2 and 5) additionally controls for the cumulative sum of prior relocations to the neighborhood, as a measure of predictable neighborhood trends. Even conditional on this measure, the relocations received by the neighborhood in the previous year still has a significant effect on parental involvement. Finally, the last pair of regressions (Columns 4 and 6) also controls for year effects as well as neighborhood fixed effects. The final pair of coefficients is reduced in magnitude and barely significant. However, even after controlling for fixed differences between neighborhoods, common time trends across all neighborhoods, as well as a measure of neighborhood trends, relocations in the previous year still appear to increase parental involvement in the current year. Under the assumption that these controls are sufficient to control for expectations about neighborhood quality, last year's relocations to the neighborhood can be treated as an unanticipated shock to neighborhood quality, which has the observed effect on parental involvement.

5 The Effect of Neighborhood Quality on Parental Involvement

5.1 Estimating Equation

To arrive at an estimating equation for the effect of neighborhood quality on parental involvement, I specialize the model from Section 2 by assuming specific functional forms. Suppose that the production function for child behavior is given by (6). As before, b denotes child behavior and y denotes parental involvement, and both are now indexed by family i and year t. Neighborhood quality q is indexed by the family's chosen neighborhood $j(i,t)$ and year. The parameter α^p is the cross-partial derivative of child behavior, with respect to parental involvement and neighborhood quality. If this parameter is positive (negative), then neighborhood quality and parental involvement are complement (substitute) inputs, in the sense that increasing neighborhod quality has a positive (negative) effect on the marginal productivity of parental involvement. The next term, F_i^p, denotes unobserved, fixed characteristics influencing the marginal productivity of parental involvement for family i. The last term, ε_{it}^p, represents an idiosyncratic shock to productivity.

$$b_{it} = \left(\alpha^p q_{j(i,t),t} + F_i^p + \varepsilon_{it}^p\right) y_{it} \tag{6}$$

Parental effort costs take the quadratic structure given in (7). The notation in this equation is similar to that for child behavior. α^c now denotes the effect of neighborhood quality on the effort cost of parents, and analogous statements hold for F_i^c and ϵ_{it}^c.

$$c_{it} = \left(\alpha^c q_{j(i,t),t} + F_i^c + \varepsilon_{it}^c\right) y_{it} - y_{it}^2 \tag{7}$$

Finally, I assume that the utility of child behavior is linear: $u(b) = \omega_i b$. The parameter ω_i increases the marginal benefit of improving child behavior, relative to the cost. Hence, it captures the effect of parental altruism, for example. Differences in ω_i will become important when studying the heterogeneity of neighborhood quality effects. For now, I assume that $\omega_i = \omega$, which is common across all families.

With the chosen production, cost, and utility functions, the parental decision problem is given by (8). The resulting decision rule for parental involvement is linear (9). In (9), the parameter of interest is $\alpha = \omega\alpha^p - \alpha^c$. This parameter describes the behavioral response of parental involvement to changes in neighborhood quality, which is driven both by changing productivity as well as changing parental effort cost. Similarly, F_i and ε_{it} can also be interpreted as the aggregate of productivity and effort cost effects. Without resorting to an underlying utility model, (9) may be viewed as an approximation to the true decision rule. However, the derivation of (9) clarifies the interpretation of the coefficient α.

$$U_{ijt} = \Bigg(\underbrace{(\omega\alpha^p - \alpha^c)}_{\equiv\alpha} q_{j(i,t),t} + \underbrace{(\omega F_i^p - F_i^c)}_{\equiv F_i} + \underbrace{(\omega\varepsilon_{it}^p - \varepsilon_{it}^c)}_{\equiv\varepsilon_{it}}\Bigg) y_{it} - y_{it}^2 \tag{8}$$

$$y_{it} = \alpha q_{j(i,t),t} + F_i + \varepsilon_{it} \tag{9}$$

The empirical objective is simply to estimate (9) and recover the parameter of interest α. I now describe the difficulty posed by the neighborhood choice of parents.

5.2 Selection Problem

Naively, in the presence of panel data, (9) can be estimated using a fixed-effects regression, so that changes in parental involvement following changes in neighborhood quality identify α. To yield unbiased estimates of α, ε_{it} needs to be uncorrelated with $q_{j(i,t),t}$. The fixed-effects regression design addresses one specific problem of neighborhood sorting: highly-motivated parents tend to sort into better neighborhoods. Suppose that between period t and the next period t', parents move randomly between neighborhoods, or, alternatively, that no family moves. In that case, conditional on the fixed effect F_i, $q_{j(i,t'),t'}$ and $\varepsilon_{it'}$ will be uncorrelated.

A trickier selection problem arises when unobserved factors in $\varepsilon_{it'}$ influence parental incentives for living in neighborhoods of different quality. Let $j' = j(i, t')$ denote the parents choice of neighborhood in period t'. Since neighborhoods are not randomly assigned, $q_{j't'}$ is likely to be correlated with $\varepsilon_{it'}$, through the parent's choice of neighborhood j'. For example, consider a hidden factor - parental health - which changes over time. A parent who becomes ill is likely be less motivated to exercise parental involvement in the future. Simultaneously, due to medical expenses incurred, she may also be unable to afford rent in the current neighborhood, and hence be forced to move to a lower-quality neighborhood. Such patterns of neighborhood selection would lead to correlation between neighborhood quality and the regression error term, in turn causing biased estimates of α.[21]

Two details about the nature of this bias are worth mentioning. First, families who move and families who do not are both selected samples, and neither sample can be used in isolation to obtain consistent estimates of the neighborhood quality effect. Revisiting the example of a parent who falls ill, instead of being forced to move to a worse neighborhood, that parent might instead be forced to stay in a declining neighborhood. In the counterfactual scenario, the healthy parent might have found a house in a better neighborhood. The second detail concerns the interpretation of the parental involvement shock, $\varepsilon_{it'}$. Although parental health has been used as an example of a confounding factor, it can also be a channel through which neighborhood quality effects operate. For example, if an increase in secondhand smoke in the neighborhood affects the parents health, this corresponds to a direct effect on parental incentives. In other words, $\varepsilon_{it'}$ only contains idiosyncratic health shocks, not attributable to neighborhood quality, which affect where a family resides in the next period.

5.3 Empirical Strategy

In order to address neighborhood selection, the research design uses an instrumental variable for $q_{j(i,t),t}$, i.e., for neighborhood quality in the family's chosen neighborhood $j(i,t)$ in year t. The instrument is measured from the family's *initial* neighborhood, denoted $j(i, t_1)$. To reiterate, this is the family's neighborhood of residence during the first year of observation. The actual construction of the instrument follows a shift-share approach. The instrument is the product of : 1) closures occurring in all other neighborhoods during the

[21] Many models of neighborhood selection will generate the kind of correlation mentioned above. In Appendix B, I consider one specific example in order to formalize the intuition.

previous year $c_{-j(i,t_1),t-1}$, and 2) the black share of neighborhood population, measured in the year 1990 ($b_{j(i,t_1),1990}$). This instrument measures the shock to neighborhood quality caused by the relocations from closed public housing projects. The first component measures the number of public housing closures in a given year: neighborhood quality in receiving neighborhoods should be more affected in years when more public housing closures happened across Chicago. The second component is a historical characteristic of the receiving neighborhood, which influences how much neighborhood quality is affected by public housing closures elsewhere in Chicago. For a displaced public housing resident, the historical black share is likely to have been a major factor in the decision of where to relocate: Oakley and Burchfield (2009) record that, among other obstacles, racial discrimination placed severe constraints on the destination neighborhoods a displaced public housing resident could choose. As a result, for a given number of public housing closures, the historical black share of the neighborhoods is likely to be correlated with the fraction of displaced residents relocating into that neighborhood.[22]

$$E\left[c_{-j(i,t_1),t-1}b_{j(i,t_1),1990}\varepsilon_{it}|X_{it},\,F_i\right]=0 \tag{10}$$

This multiplicative interaction is a valid instrument for neighborhood quality, under the assumption that it is uncorrelated with idiosyncratic changes in parental motivation (10). Using this research design, of the neighborhood quality effect is identified by changes in parental involvement in response to fluctuations in the yearly flow of relocation shocks to the initial neighborhood. An alternative approach to identification would be to use changes in parental involvement in response to the cumulative stock of relocation shocks. While this alternative approach is intuitively appealing - cumulative relocations are easily interpreted as the total impact of public housing closures on receiving neighborhood quality - it is ultimately not used. This choice is motivated by the possibility that overall trend in relocation shocks might have been predicted by parents, and this information might have influenced the selection of initial neighborhoods. For example, a parent who expected to be less able to exercise parental involvement in the future may have chosen to live in a neighborhood that was less likely to be affected by public housing closures. In contrast, fluctuations in the flow of relocation shocks is less likely to have been well predicted by parents, and as a result are more likely to be uncorrelated with neighborhood selection behavior.

Two features of the research design must be highlighted. The first feature is that, for every family from a

[22]The historical black share of the neighborhood may also be correlated how much neighborhood quality was affected by a given number of relocated families.

given initial neighborhood, the instrument takes the same value whether that family subsequently moved out or not. To see how this feature addresses neighborhood selection, consider the following decomposition of neighborhood quality in the family's chosen neighborhood, measured in the year t: $q_{j(i,t),t}$ (11). The first term in this decomposition is the neighborhood quality in the initial neighborhood. The second term is the difference between chosen and initial neighborhood quality. If parents from the same initial neighborhood experience different parental motivation shocks, and consequently make different residential choices, the effect on chosen neighborhood quality will be isolated in the second term. In contrast, relocations to the family's initial neighborhood primarily affect $q_{j(i,t),t}$ through the first term.

$$q_{j(i,t),t} = q_{j(i,t_1),t} + \left(q_{j(i,t),t} - q_{j(i,t_1),t}\right) \tag{11}$$

Following the above intuition, I verify that relocations to the initial neighborhood appear to be uncorrelated with the second term in (11). These findings are presented in Table 5. Columns 4 to 6 present regression coefficients, for which the dependent variable is the crime rate difference, between chosen and initial neighborhood. The coefficients displayed correspond to the shift-share instrument, in the year prior as well as the year before $(t-2)$. All variables have been standardized for this particular table, hence, the coefficient sizes suggest a very small correlation: a one standard deviation change in the instrument is associated with roughly a 0.03 standard deviation change in the difference. The small magnitude of the coefficient suggests that the instrument is uncorrelated with the residential choices made by study families.

The second notable feature of the research design is that, in contrast to previous analysis (Section 4.4), it makes no use of the available information on relocation shares. One reason for this choice is that the data on relocation shares is only available at a coarser geographic scale and might, as a result, be imprecise. A second reason is that the actual relocation flows might be endogenous to trends in the receiving neighborhood, since relocation flows are the aggregated choices made by displaced public housing residents. For example, if displaced public housing residents favored destination neighborhoods where average parental involvement was increasing, neighborhoods were average parental involvement was increasing would receive higher levels of relocations. This form of reverse causation would generate a spurious positive correlation between neighborhood quality and parental involvement.[23] In the construction of the relocation shocks instrument, the historical minority composition in the neighborhoods is to be exogenous to any subsequent trends in average

[23] Alternatively, instead of displaced households doing the selecting, the public housing authority may have relocated residents to other public-housing projects in a manner correlated with parental involvement in those neighborhoods.

parental involvement, within the neighborhood. As a result, this instrument is much less likely to be subject to the problem of reverse causation.

5.4 Results

Following previous analysis, the neighborhood crime rate is used as a measure of quality. In application, the regression equation estimated is given by (12). Comparing (12) and (9), one difference is the addition of a vector of family characteristics X_{it} as controls. These controls are: marital status of the primary caregiver (as a set of dummy variables), employment of the primary caregiver, as well as occupancy type of the family (as a set of dummy variables). Since fixed effects are already included, the coefficients on these controls are identified by families who switch from one category to another, over time. The second difference is that the variable for neighborhood quality has been replaced with neighborhood crime rate v_{jt}. Since a higher crime rate is associated with lower neighborhood quality, the interpretation of the estimated coefficient α^v is changed: a positive coefficient means that *lower neighborhood quality* (higher neighborhood crime) is associated with greater parental involvement.

$$y_{it} = \alpha^v v_{j(i,t),t} + X_{it}'\beta + F_i + \varepsilon_{it} \tag{12}$$

The results of this regression are presented in Table 6. The first two columns show some preliminary estimates, which do not employ the empirical strategy described. Column 1 shows OLS estimates, which do not account for neighborhood selection by receiving families. While statistically significant, the magnitudes of these estimates are very small. Column 2 shows the estimated neighborhood quality effects when the relocations measure (from Section 4) is used as an instrument for neighborhood quality. The coefficient on neighborhood quality becomes more substantial. Column 3 presents the baseline estimates for the analysis: the magnitude of the coefficient is larger still. The greater magnitude of the IV estimates relative to the OLS estimates is consistent with expected neighborhood selection behavior: parents who experience increasing motivation to care for their children are also likely to move to better neighborhoods. To place the magnitude of the baseline estimate in context, I calculate the response of parental involvement to a one standard deviation increase in crime rate. Because crime rate varies widely between neighborhoods, and families rarely move to neighborhoods with drastically different crime rates, I use the within-neighborhood standard deviation of crime rate as a basis for comparison. According to the baseline coefficient in Column 3, a one standard deviation fall in neighborhood quality (17 crimes per 1000 residents) is associated with roughly a half standard deviation increase in parental involvement.

The baseline coefficients should be interpreted as the causal effect of neighborhood quality on parental involvement, and not the causal effect of neighborhood crime itself. Other aspects of neighborhood quality that were also altered by the relocations, and the effect of these other neighborhood characteristics is confounded in the estimated coefficients. I ignore the possibility that overall neighborhood quality might have been improved by incoming relocations: while certain isolated aspects of neighborhood quality might have been improved, it is highly likely that far more aspects were adversely affected, such that overall neighborhood quality decreases. In the preceding analysis, neighborhood crime rates are also assumed to accurately measure the *scale* of neighborhood quality: in other words, a one standard deviation increase in neighborhood crime rate is assumed to correspond to a one standard deviation decrease in neighborhood quality.[24]

The estimated neighborhood quality effect appears robust to different specifications of the neighborhood crime rate. Columns 4 to 6 of Table 6 vary the measure of crime used as an indicator of neighborhood quality. When only considering property crimes, the results are essentially unchanged. When only considering personal crimes, the estimated coefficient becomes much larger in magnitude, but this coefficient is likely unreliable due to a weak first-stage. Finally, when the raw number of crimes is used instead of the crime rate, the magnitude of the coefficient changes, but the standardized effect size is very similar to the baseline estimate.

The results suggest that increased parental care is a compensating response to decreased neighborhood quality. These results are very similar in spirit to those of Pop-Eleches and Urquiola (2013), who find that an increase in school quality results in a reduction in parental effort. Similarly, Katz et al. (2005) find that parents reduce monitoring of their children after relocating to a better neighborhood.[25] On the other hand, these findings sharply contrast with those of Patacchini and Zenou (2011) (PZ). They find positive effects of neighborhood quality on parental involvement, a relationship they refer to as "cultural complementarity". There are a few possible explanations for the different findings here. The first reason is the different populations being studied: the sample in PZ is restricted to residents of council housing in the UK during the 1970s - one possible reason for the discrepancy is the timing of variable measurement. In PZ, parental involvement variables are measured when the child was 7 years old, but the reference

[24]This assumption is made because neighborhood quality has no natural scale. As Appendix C describes in detail, this assumption is unlikely to overstate the magnitude of neighborhood quality effects: if neighborhood quality is additive in it's individual components, and each of these components are affected by the same amount (in standard deviation units), the estimated coefficient using neighborhood crime rates will actually understate the magnitude of the neighborhood quality effect.

[25]It should be noted that, in Katz et al. (2005), this result applies to residents of low income public housing - a different population to the one studied here.

neighborhood was measured when the child was 16 years old. During the intervening period, changes in the level of parental involvement made in response to changes in neighborhood quality (or even changes in neighborhood)[26] are not observed by PZ. In contrast, in this study, it is precisely those changes over time (e.g. parental involvement at age 7 vs. parental involvement at age 16) which identify neighborhood effects. Another issue with the timing is that neighborhood quality is measured after parental involvement, meaning that the interpretation of the effects is unclear. It is possible that, in PZ, initial levels of parental involvement are affecting neighborhood quality through selection into public housing, rather than the other way round.

Finally, these findings may appear to suggest that parent and neighborhoods are substitute inputs in producing good child behavior. If parental involvement and neighborhood quality are substitute inputs, then a reduction in neighborhood quality raises the marginal productivity of parental involvement, which in turn explains the observed increase in parental involvement. However, this explanation is not the only one: aside from marginal productivity, the marginal effort cost of parental involvement could have been been affected by reductions in neighborhood quality. For example, if deterioration to neighborhood amenities reduces the value of parental free time, parents would also spend more time around their children. I revisit the issue in Section 6, in light of additional empirical results.

5.5 Robustness

I further examine the sensitivity of the main findings in Table 7. These robustness specifications modify the baseline specification (Table 6, Columns 3). Uniformly, the coefficients remain positive and statistically significant. Column 1 estimates neighborhood quality effects using 2 lags of the relocation shocks instrument instead of 1. Column 2 changes the dependent variable to an additive index, instead of using the factor derived from multiple correlates analysis. Similar results are obtained for a more straightforward measure of parental involvement as well. Column 3 restricts the sample to single child families, since the choice of parents with multiple children is not explicitly modeled, and the effects obtained here are larger. Column 4 discards any family who ever lived in a neighborhood containing a public housing project: potentially, those family would not only be affected by incoming relocations, but also by closures or other changes occurring within public housing projects. Column 5 restricts the sample to families who never moved out of the initial neighborhood. If the instrument adequately addresses the neighborhood selection problem, the coefficients obtained should be identical for households who moved as it is for households who did not. The resulting

[26] It is not reported if the sample is restricted to individuals who lived in the same neighborhood at these two points, or if the information is even available.

estimates are slightly larger than the baseline, however, there is a large overlap in the confidence intervals.[27] Finally, Columns 6 and 7 alter the sample by respectively discarding any parents who changed marital status, and including any families where the primary caregiver changed. These last two columns show that the basic findings are robust to changes in the sample definition.

5.6 Subsets of Parenting Behavior

The next analysis describes, in more detail, the kinds of parenting behavior that respond the most to changes in neighborhood quality. Because individual measures are binary in nature, estimates of neighborhood quality effects on individual behavior measures are noisy and generally insignificant. Hence, for the next analysis, I divide the overall parental involvement measure into two groups of parenting behavior: interaction and discipline. Interaction contains behavior measures concerning parent-child interaction as part of a daily routine, as well as any activities that the parent organizes for the child.[28] Discipline covers the setting and enforcement of rules, both about homework, as well as chores and behavior with friends. Each subset of parenting behavior is collapsed into a single variable, using the same MCA procedure described in Section 3. In addition, a separate variable is created for housekeeping behavior, from behavior measures which are not included in the parental involvement measure. As a domestic behavior that is unrelated to child care, housekeeping behavior serves as a falsification test for the hypothesis that parents are responding to changing neighborhood influences on the child. If neighborhood quality primarily affects parental decisions through other channels, then other domestic behavior like housekeeping should also respond to changes in neighborhood quality. For example, one alternative explanation for the previous findings is that parents in a declining neighborhood face lowered incentives to leave the house, due to rising crime rate, deteriorating neighborhood amenities or other causes. This explanation would suggest that parental involvement would increase due to more time spent at home, but would also suggest that parents would do more housekeeping.

The effects of relocations on these subsets of parental behavior are presented on the left half of Table 8. The results suggest that most of the increase in parental involvement is driven by greater interaction between parent and child. Surprisingly, parenting behavior related to discipline appears to be little changed in response to neighborhood quality. While it might seem reasonable to expect that parents would exercise tighter disciplinary control over their children when neighborhoods worsen, there is little empirical evidence in support of this claim. Additionally, housekeeping behavior exhibits a small and statistically insignificant

[27]The sample for households who moved is smaller, and the effect sizes obtained are similar but insignificant.

[28]An example of the former is whether the parent engages in conversation daily with the child. An example of the latter is whether the child visits with family/friends regularly 2-3 times a month.

response to changes in neighborhood quality. As a result, it appears more likely that parents are strategically choosing behavioral responses which affect child development and behavior.

5.7 Effects on Child Behavior

There is extensive empirical research into the effects of neighborhood quality on child behavior. While this is not the main focus of the paper, I present estimates of the effect of neighborhood quality on three different measures of child behavior: the Wide Ranging Achievement Test (WRAT), internalizing behavior score, and externalizing behavior score. The first is a measure of cognitive achievement, and the other two are measures of potentially problematic, anti-social behaviors.[29] The analysis utilizes the same regression model described in (12), but with the dependent variable with the relevant child behavior. Because the same instrument is used for neighborhood quality, the estimated coefficients are not confounded with the residential neighborhood selection of the family. However, the estimated coefficients capture the sum of two conceptually distinct effects: the direct effect of neighborhoods on children and the mediating effect of altered parental involvement on children. This analysis does not attempt to separately identify both effects.

The estimated coefficients are shown in the right half of Table 8. Decreased neighborhood quality is associated with a modest *increase* in achievement test scores. At the same time, it is also associated with a very small increase in internalizing behavior, and a larger increase in externalizing behavior. Since the latter two variables describe problematic behavior, the overall story is conflicted: changing neighorhood quality appears to improve some aspects of child behavior but adversely affect others. However, bearing in mind the previous findings, compensating changes to parental involvment may be offsetting neighborhood effects, so that the net effect on the child behavior is very small.[30] These findings highlight the importance of considering parental mediation, if the research objective is to isolate the direct impact of neighborhoods on children. Furthermore, under the likely assumption that parental involvement positively affects child behavior, then, based on the findings in this paper, the direct effect of neighborhoods on children is likely to be larger (more positive) than existing estimates suggest.

[29] These variables are introduced in Section 3.3.

[30] This reasoning assumes that an increase in parental involvement is beneficial towards producing good child behavior. This assumption is not tested in this paper, since the causal effect of parental involvement on child behavior is not identified. However, there is a body of research which suggests the positive effect of increased parental involvement on child behavior. Fryer Jr et al. (2015) find that incentivizing parental involvement leads to increased child achievement in preschool, within an experimental setting. In a similar study, Avvisati et al. (2014) find positive effects on student behavior and achievement. These experimental findings is supported by a range of observational studies, which find a positive association between parental involvement and student achievement (Jeynes, 2007). Admittedly, the parental involvement variables featured in these other studies are restricted to academic matters (e.g. time spent in parent-teacher conferences), while the parental involvement variables measured here have a wider scope. However, Amato and Rivera (1999) finds that increased paternal involvement is associated with reduced behavioral problems in children, for a general definition of parental involvement that is similar to the one used here.

6 Heterogenous Effects of Neighborhood Quality

According to the results from the previous section, parents compensate for decreased neighborhood quality by increasing parental involvement. In this section, I consider whether the effect of relocations on parental involvement differ across different types of families. The main finding of this section is that parents with different initial parental involvement levels respond very differently to the same change in neighborhood quality.

In the basic model presented in Section 5, the coefficient of interest $\alpha = \omega\alpha^p - \alpha^c$ is uniform across all individuals. The basic model is extended to present an explanation for heterogenous neighborhood quality effects, by allowing the parameter ω is instead allowed to vary across families, then the effect of neighborhood quality on parental involvement will also vary across families (13). The parameter ω_i captures the utility benefits of improving child behavior, relative to the effort cost of parental involvement. Hence, this term captures any difference in parental altruism, as well as differences in the effort cost of parental care (due to factors like health or work stress). Henceforth, this parameter is referred to as *parental altruism*, with the understanding that this term actually captures benefits relative to costs.

$$\alpha_i = \omega_i \alpha^p - \alpha^c \tag{13}$$

While this parameter ω_i is unobserved at the family level, it is likely to be positively correlated with parental involvement levels. From (9), I obtain the partial derivative of parental involvement with respect to this parental altruism parameter. This partial derivative is equal to the marginal productivity of parental involvement in producing child behavior, which is positive by assumption (14). Hence, the model suggests a meaningful association between initial levels of parental involvement and the subsequent response to changes in neighborhood quality.

$$\frac{\partial y_{it}}{\partial \omega_i} = \alpha^p q_{j(i,t)t} + F_i^p + \varepsilon_{it}^p = \frac{\partial b_{it}}{\partial y_{it}} \tag{14}$$

In the empirical specification, initial parental involvement is broken into quartiles - denoted Q_k^y, $k \in \{1, 2, 3, 4\}$ - and the interaction between initial involvement quartile and neighborhood crime rate is estimated (15). The four coefficients α_k^v are potentially biased due to the residential neighborhood selection of the family: in order to identify these coefficients, the relocation shocks instrument from the previous section is interacted with initial parental involvement quartile.

$$y_{it} = \sum_{k=1}^{4} \alpha_k^v \mathbf{1}\left(y_{i,t_1} \in Q_k^y\right) v_{j(i,t),t} + X_{it}'\beta + F_i + \varepsilon_{it} \tag{15}$$

The results of this analysis are presented in the first column of Table 9, as well as in Figure 5. There is a positive association between initial parental involvement and the response to increased crime rate. Most parents exhibit a compensating response to changing neighborhood quality: when the crime rate increases, these parents increase their parental involvement. In contrast, parents from the lowest quartile actually respond by reducing involvement by a significant magnitude.

Since more involved parents tend to live in better neighborhoods, the the above findings could potentially be explained by parents in better neighborhoods responding differently to parents living in worse neighborhoods. Figure 6 shows estimates of the interaction between initial neighborhood income and neighborhood quality[31]. The results show that families from wealthier initial neighborhoods indeed compensate more heavily for an increase in crime rate. However, the coefficients do not vary as much, suggesting that a substantial part of the heterogeneity occurs across different families within the same neighborhood. Figure A2 shows another version of the interaction between initial parental involvement and neighborhood quality: this time, parental involvement is measured relative to other parents in the same initial neighborhood. The results from this analysis are very similar to Figure 5, suggesting that different parents within the same neighborhood respond differently to neighborhood quality change.

Another alternative explanation for these findings is that different families within the neighborhood are differentially exposed to negative neighborhood quality shocks. For example, within the same neighborhood, low-involvement parents might tend to live on the same block as incoming displaced households, while high-involvement parents tend to live farther away. This issue is essentially one of neighborhood mis-measurement: within the same "neighborhood", the effect of relocations on neighborhood quality may be non-uniform. Such issues would certainly limit the interpretation of the results. However, this kind of mis-measurement fails to explain the reversal of sign in the coefficients. As long as all families encountered a decrease in neighborhood quality, the reversal of sign still indicates a switch from compensating to reinforcing behavior.[32]

[31] The coefficients are also presented in Table 9, but are discussed in context of the figures.

[32] A switch from reinforcing to compensating parental behavior has previously been documented, in a different context by Hsin (2012). She finds that low-SES mothers invest less time in a low birth weight child relative to a sibling with normal birth weight; in contrast, high-SES parents invest more.

These findings have implications on the expected impact of place-based policies to improve child outcomes. Specifically, place-based policies to influence child development outcomes may have varied impacts, due to differences in parental mediation. Specifically, low-involvement parents might magnify the effect of such policies through increased parental involvement, while high involvement parents diminish these effects. Hence, when designing policies to address the impacts of worsening environmental influences on long-run child outcomes, policymakers might seek to target children receiving low levels of parental care. On the other hand, policies to improve the neighborhood environment are likely to be more effective on this particular group of children, because of parental mediation.

6.1 Are Parents and Neighborhoods Complements or Substitutes?

Based on the results presented in Section 5, parents on average tend to compensate (on average) for decreased neighborhood quality by increasing parental involvement. This finding has two potential explanations. The first explanation is that parental involvement and neighborhood quality are substitute inputs: when neighborhood quality decreases, the marginal productivity of parental involvement increases. The second explanation is that reduced neighborhood quality decreases the effort cost of parental involvement. Under certain assumptions, the heterogenous neighborhood quality effect found in this section can be used to rule out the second explanation. These assumptions are restrictive, and the purpose of this section is to clearly specify the assumptions under which such a claim can be made on the structure of the production function.

To conclusively determine that parental involvement and neighborhood quality are substitute inputs, three assumptions must be made: first, a higher parental altruism parameter is positively correlated with observed initial involvement levels. Second, the parameter α^p is constant across different families. This parameter is the cross-partial derivative of the child behavior production function, with respect to neighborhood quality and parental involvement. Third, the parameter α^c is also constant across different families. This parameter is the cross-partial derivative of the parental effort cost, again with respect to neighborhood quality and parental involvement. The interpretation of this assumption is analogous to the previous one. The second and third assumptions can be interpreted as saying the following: the marginal productivity of parental involvement is affected by neighborhood quality in exactly the same way across different parents, and the same holds for the marginal cost of parental involvement.

Proceeding under these assumptions, the heterogenous neighborhood quality effect indicates that parents with higher altruism (ω_i) compensate more in response to a neighborhood quality decline. Recalling (13), this pattern of heterogeneous responses suggests that $\alpha^p < 0$. In other words, parents and neighborhoods must be substitute inputs in the production of good child behavior. The intuition behind this result is as follows: greater parental altruism not only incentivizes a higher initial level of parental involvement, but also increases the parent's sensitivity to changes in the marginal productivity. Because parents with high altruism seem to have a more compensatory response to neighborhood quality, this implies that increased neighborhood quality decreases the marginal productivity of parental involvement.[33]

This result is relevant in anticipating the effects of comprehensive community-based policies targeting the early childhood development of disadvantaged children.[34] Such programs attempt to simultaneously improve the quality of the child's neighborhood environment, as well as foster greate engagement between parent and child. These results suggest that such programs will fact a problem of "crowding out": policies encouraging parental engagement may have smaller effectiveness than when implemented in conjunction with other policies improving the quality of the neighbohood environment, relative to being implemented alone.

The use of heterogeneous neighborhood quality effect allows a structural parameter (α^p) to be signed under more general assumptions: i.e. without assuming a particular sign for the paramete α^c. In contrast, when the average effect of neighborhood quality on parental involvement is considered in isolation, the identical claim would require the additional assumption that increasing neighborhood quality reduces the effort cost of parental involvement (i.e. $\alpha^c < 0$). Previous work on neighborhood quality effects has assumed parents and neighborhoods to be substitute inputs. (Patacchini and Zenou, 2011). Here, the relationship between parents and neighborhoods is delivered as a result, not an assumption. This improvement is made possible by the availability of panel data, which permits estimating the interaction between initial involvement and the subsequent response to changing neighborhood quality. [35]

At the same time, even the relatively general assumptions used here may be seen as overly restrictive. These assumptions severely restrict any heterogeneity in neighborhood quality effects: changes to neighborhood quality must affect the productivity and cost of parental involvement uniformly across all families. If either of the assumptions fail to hold, the current results are uninformative about the structural parameters of the model.

[33] For general functional forms for parental utility, the result holds under some restrictions on the magnitude of higher order utility derivatives. See Section D for a brief discussion.

[34] Examples of such programs include the Harlem Children's Zone, Promise neighborhoods, and the Even Start program.

[35] Patacchini and Zenou (2011) uses a single cross-section, and as a result the interaction effect cannot be observed.

6.2 Heterogeneity in Effects by other Characteristics

For the purposes of targeting poliy interventions, it may be important to determine which parents compensate for and which parents reinforce neighborhood quality effects. For example, to determine whih children are most adversely affected by neighborhood decline, a policymaker may wis to identify families who are unable to compensate through increased parental involvement. Since measures of parental involvement are less readily available, it is important to determine whether other observable characteristics of the family are correlated with neighborhood quality effects on parental involvement.

The following analysis proceeds by re-estimating (15). In each case, initial involvement is replaced with an variable describing the family. These descriptors were measured during the initial survey, and hence describe the characteristics of the family prior to the neighborhood quality changes induced by relocation shocks. Once again, the estimated coefficients show, for the relevant sub-group, the average response of parental involvement to a unit increase in the crime rate. Hence, a positive (negative) coefficient is interpreted as compensating (negative) behavior.

The first family characteristic considered is child cognitive achievement, measured using the Wide Ranging Achievement Test (WRAT). Figure 7 presents heterogenous neighborhood quality effects by child achievement test score quartile. Similar to the previous analysis, there is limited evidence of a switch from reinforcing to compensating parental involvement responses. Children from the lowest quartile of cognitive test scores (WRAT) tend to have parents who make reinforcing responses to neighborhood quality change, in contrast to children with higher scores. However, in this case, the differences are slight and statistically insignificant. Hence, child achievement test scores are of limited use in predicting how parental involvement will respond to changes in neighborhood quality.

Next, I consider whether basic demographic characteristics of the child - namely, child age and gender - are predictive of the parental involvement response to neighborhood quality changes. The estimated coefficients are presented in Figure 8. There appears to be a slight positive association between the age cohort of the child and response of parental involvement to neighborhood crime rate. Based on the child's age at the time of the initial survey, parents of older children appear to compensate more for neighborhood decline. Parents of children aged 12 and over appear to compensate especially strongly for neighborhood decline. In comparison, the parents of the youngest children (aged 0 to 3 years) do not record a significant response at all. One plausible explanation for this heterogeneity is that infants and toddlers are unlikely to be significantly affected by neighborhood influences. Hence, their parents would be unlikely to respond strongly to neighborhood change. When it comes to the gender of the child, parents of female children appear to compensate slightly more strongly for neighborhood decline. On the other hand, in this as well as in the previous case, there is no subgroup associated with a reinforcing response to neighborhood change.

Finally, Figure 9 shows that the response of parental involvement to neighborhood decline does not vary with parental education. Along with other measures of socioeconomic status, parental education has been found to be substantially correlated with the intensity of parental care (Guryan et al., 2008). Among the sampled households, this pattern is replicated as well (Table 2), for baseline levels of parental involvement. In contrast, the response to neighborhood decline is very similar across parents of different education levels: while highly-educated parents are more involved with their children, they respond similar to less-educated parents, increasing their involvement by roughly the same magnitude in response to neighborhood quality declines. Overall, none of the common family characteristics identify children who are most at-risk from neighborhood decline, which highlights the importance of measuring parental involvement.

7 Conclusion

In this paper, I have investigated how parental involvement changes in response to neighborhood quality. By using induced relocations as a natural experiment, I first verify that relocations induced by public housing closures cause neighborhood crime levels to increase in the subsequent year. This suggests the relocations had an adverse impact on neighborhood quality in the receiving neighborhoods. Following this, I find that a higher level of relocations causes parents to increase their level of involvement: a pattern which has not been previously documented or accounted for in studies of neighborhood effects on children. In addition, I find significant interaction effects between the relocations and basic demographic characteristics of the families. This suggests that the response of parental involvement to neighborhood quality differs meaningfully between families. The differing magnitude of parental response may result in different later-life outcomes in children, in response to the same neighborhood shock. Additionally, under some additional assumptions, the interaction between initial levels of parental involvement and its response to neighborhood is informative regarding the relationship between neighborhood quality and parental care is inputs to child development.

Further work is needed to incorporate child behavior into the model, in order to examine further the mechanisms affecting parental involvement, as well as its ultimate impact on children. One goal of such research would be to weight the relative contributions of neighborhood quality and parental involvement towards later-life outcomes in children. Patacchini and Zenou (2011) have made some steps in this direction, by estimating the effect of (randomly-assigned) neighborhood quality and parental investment towards child education, however, the analysis is limited by the single measure of neighborhood quality and of parental involvement. Information on the long-term outcomes of the PHDCN children would be invaluable in exploring such questions.

Another promising direction is to incorporate neighborhood selection by parents in more detail. While, the shock to neighborhood quality cause by incoming relocations is unlikely to have influenced neighborhood exit in the short run, it is likely to have had some influence in the long run, as the signs of neighborhood decline became more obvious. Another question about neighborhood selection is where parents facing a declining neighborhood choose to move to, and how the new neighborhood compares to the old one. Finally, analyzing the simultaneous choice of neighborhoods and parental involvement would answer a new set of interesting questions. For example, are the parents who move also the ones who change their level of involvement? Also, do the changes in neighborhood quality brought on by moving reinforce the changes (if any) in parental involvement? These are the questions to be pursued in the future.

Tables

Table 1: PHDCN Sample Size

	Overall			
	Wave 1	Wave 2	Wave 3	Total
All subjects	6226	5340	4855	16421
+ Consistent PC	4321	3507	3120	10948
+ >1 Interview	3369	3240	2845	9454
+ Never left Chicago	3088	2959	2564	8611
School Age Sample (6-18 years)	1760	2250	1761	5771
Toddler Sample (0-5 years)	1423	1337	466	3226
Young Adults (18+ years)	633	506	984	2123
	per Initial Neighborhood			
	Wave 1	Wave 2	Wave 3	Total
All subjects	75.0	48.8	34.8	52.9
+ Consistent PC	52.1	32.8	23.0	35.9
+ >1 Interview	40.6	32.8	23.0	32.1
+ Never left Chicago	37.2	30.1	22.8	30.0
School Age Sample (6-18 years)	21.2	23.2	16.0	20.1
Toddler Sample (0-5 years)	17.1	12.9	3.64	11.2
Young Adults (18+ years)	7.63	4.61	7.10	6.45

Table 2: Descriptive Statistics of PHDCN Families

	Mean (sd)	Differences (p)		
	Wave 1	Move - Stay	Rich - Poor Nbd.	PC HSG - Dropout
Primary Caregiver (PC):				
Employed	0.43	0.074	-0.15	-0.21
	(0.50)	(0.00)	(0.00)	(0.00)
Salary ($ 000)	21.1	-5.25	10.0	11.2
	(33.1)	(0.00)	(0.00)	(0.00)
Black	0.35	0.11	-0.049	0.18
	(0.48)	(0.00)	(0.00)	(0.00)
PC age	37.0	-3.82	1.83	1.00
	(8.86)	(0.00)	(0.00)	(0.00)
PC Female	0.93	0.027	-0.035	-0.026
	(0.25)	(0.00)	(0.00)	(0.00)
Parental Involvement:				
Sum, HOME Score Measures	12.6	-0.47	0.30	0.69
	(2.30)	(0.00)	(0.00)	(0.00)
Parenting Behavior Subsets:				
Interaction	5.75	-0.063	0.26	0.60
	(2.92)	(0.59)	(0.00)	(0.00)
Discipline	7.54	-0.59	0.092	0.12
	(3.00)	(0.00)	(0.23)	(0.12)
Parental Involvement Measure	0.024	-0.10	0.14	0.31
	(0.98)	(0.01)	(0.00)	(0.00)
Child:				
Age	9.34	-0.93	0.18	-0.055
	(4.74)	(0.00)	(0.09)	(0.59)
WRAT	96.0	-3.96	7.41	7.81
	(19.7)	(0.00)	(0.00)	(0.00)
Behavioral Scores:				
- Internalizing	8.58	1.02	-1.23	-1.72
	(7.40)	(0.00)	(0.00)	(0.00)
- Externalizing	9.64	2.14	-0.98	-0.63
	(8.23)	(0.00)	(0.00)	(0.00)
Observations	8711			

Sample excludes any subjects who left Chicago, changed PCs,
as well as those whose PC changed marital status.
Variables measured at Wave 1 of survey.
Rich - Poor Nbd.: whether wave 1 neighborhood was above/below median neigborhood income.
Move - Stay: families which changed / did not change neighborhood when next interviewed.
PC HSG - Dropout: families where the primary caregiver attained/did not attain a high school degree.
Salary: treated as missing if PC is unemployed.

Table 3: Neighborhood Characteristics

	(1) All Chicago mean/sd	(2) PHDCN Wave 1 mean/sd	(3) CHA project mean/sd	(4) CHA closure mean/sd	(5) (2) + (3) mean/sd	(6) (2) + (4) mean/sd
ln. family income	10.4	10.4	9.90	9.74	10.5	10.4
	(0.44)	(0.39)	(0.68)	(0.69)	(0.81)	(0.94)
Frac. minority	0.56	0.54	0.75	0.84	0.52	0.56
	(0.36)	(0.31)	(0.33)	(0.29)	(0.33)	(0.40)
Unemployment rate	0.14	0.12	0.27	0.32	0.15	0.18
	(0.095)	(0.073)	(0.16)	(0.16)	(0.13)	(0.16)
Population	8121.7	9081.6	8276	7950.4	10304.3	11909.8
	(2914.9)	(2966.0)	(3862.7)	(3912.2)	(4188.4)	(3781.6)
PHDCN observations		32.3			17.3	15
		(16.3)			(11.6)	(8.34)
Relocatees	5.36	4.54	5.62	4.00	2.45	2.66
	(6.85)	(5.61)	(9.70)	(4.76)	(3.36)	(3.70)
Observations	343	83	32	22	8	5

Table 4: Parental Involvement Responses to Relocations Received by Neighborhood

Effect of:	100 Relocations			10 Relocations / 1000 Residents		
	(1)	(2)	(3)	(4)	(5)	(6)
Relocations (t-1)	1.48***	1.42**	1.02+	1.03***	0.94**	0.61+
	(0.37)	(0.49)	(0.52)	(0.28)	(0.31)	(0.35)
Neighborhood Fixed Effects	No	No	Yes	No	No	Yes
Year Effects	No	No	Yes	No	No	Yes
Cumulative Prior Relocations	No	Yes	Yes	No	Yes	Yes
	Relocation Quantities, mean (s.d.)					
Number / Year	3.5 (7.7)			0.5 (1.1)		
If > 0	5.6 (9.0)			0.8 (1.3)		
Observations	7481	7481	7481	7481	7481	7481

Notes:
Dependent Variable: first axis of MCA on 15-23 parenting behaviors, standardized by child age group.
Neighborhood: refers to study-defined Neighborhood Cluster (343 Neighborhood Clusters in Chicago).
Sample excludes families that left Chicago, or changed parent (Primary Caregiver).
Controls: Parental employment, parental marital/co-habitation status, subject fixed effect.
Standard errors clustered by initial community area of residence (77 community areas in Chicago).

Table 5: Exclusion of Relocations from Neighborhood Selection

Dependent Variable:	Remaining in Initial Neighborhood			Difference, Crime Rate		
	(1) exp(coef.)	(2) exp(coef.)	(3) exp(coef.)	(4) coef.	(5) coef.	(6) coef.
Relocations per Resident, Initial Neighborhood:						
t-1	0.907	1.051	1.048	-0.023	-0.025*	-0.033*
	(0.058)	(0.056)	(0.053)	(0.012)	(0.012)	(0.014)
t-2					0.007	
					(0.006)	
x Parental Involvement			1.011			0.003
			(0.044)			(0.014)
Fixed Effect, Initial Neighborhood	Yes	Yes	Yes	No	No	No
Year Effect	No	Yes	Yes	No	No	No
Cumulative Prior Relocations	No	Yes	Yes	No	No	No
Observations	11799	11799	11799	6144	6144	4674

Notes:
Coefficients describe effect of 1 standard deviation in relocations received.
Columns 1 to 3: Exponentiated coefficients from parametric (exponential) survival time regression.
Columns 4 to 6: Coefficients from linear regression of difference between initial and subsequent neighborhood crime rate.
Neighborhood crime rates have been standardized.
All models control for parental involvement and parental employment/marital status/occupancy type, measured at initial survey. Column 3 also controls for interaction between cumulative prior relocations and parental involvement.
Regression sample excludes any child who changed parents (primary caregivers), as well as any child who left Chicago.
Robust standard errors in parentheses.

Table 6: IV Estimates: Effect of Neighborhood Quality on Parental Involvement

	Measure of Neighborhood Quality (default = Crimes / 10 Residents)					
	(1)	(2)	(3)	(4) Property	(5) Personal	(6) 100 Crimes
Neighborhood Quality Effect	0.44^{**}	2.57^{**}	3.18^{**}	3.80^{**}	7.94^{**}	0.38^{**}
	(0.11)	(0.64)	(0.89)	(1.06)	(2.54)	(0.11)
S.D. of Regressor, Within-nbd.	0.17	0.17	0.17	0.13	0.06	1.50
	Relocations Instrument for Neighborhood Quality					
	None	Observed	Frac. Black	Frac. Black	Frac. Black	Frac. Black
First Stage Coefficient		4.25^{**}	3.38^{**}	2.87^{**}	1.38^{*}	28.71^{*}
		(0.73)	(0.93)	(0.59)	(0.56)	(11.91)
F-statistic, Instrument		34.05	13.16	23.67	6.12	5.81
Observations	10963	7481	7481	7481	7481	7481

Notes:
Dependent Variable: first axis of MCA on 15-23 parenting behaviors, standardized by child age group.
Property/Personal: crime rate measures for specific categories of crime.
Neighborhood: refers to study-defined Neighborhood Cluster (343 Neighborhood Clusters in Chicago).
All instruments measured from initial neighborhood, neighborhood quality measured from current neighborhood.
Observed: Instrument for neighborhood quality is relocations per 1000 residents, to family's initial (Wave 1) neighborhood, constructed as product of yearly public housing closures in previous year, and observed relocation shares.
Frac. Black: Instrument is product of yearly public housing closures in previous year, and neighborhood black share in 1990.
Controls: Parental employment/marital status, occupancy type, subject fixed effect.
Standard errors clustered by initial neighborhood.

Table 7: Robustness Specifications: Effect of Neighborhood Quality on Parental Involvement

	(1) 2 Lags	(2) Raw Index	(3) Only Child	(4) No Project	(5) Never Moved	(6) Same Marital	(7) Change Parent
Crimes / 10 Residents	3.27^{**}	3.44^{**}	4.73^{**}	3.10^{**}	3.95^{**}	3.44^{*}	3.21^{**}
	(0.91)	(0.88)	(1.41)	(0.95)	(1.38)	(1.36)	(0.90)
Observations	7481	7481	4299	6872	4822	5404	7816

Notes:
Baseline specification is column 3 of Table 6.
2 Lags: First stage includes relocations/residents in year t-2 as well.
Raw Index: Dependent variable is additive index of parenting behaviors.
Only child: Drop observations where multiple children report same parent.
No Project: Drop initial neighborhoods containing CHA public housing.
Never Moved: Drop any family that changed neighborhoods.
Same Marital: Drop any family where parent changed marital status.
Change Parent: Include observations where primary caregiver changed.
Standard errors clustered by initial neighborhood.

Table 8: Effect of Neighborhood Quality on Other Behaviors

	Subset, Parenting			Child Behavior		
	(1) Interaction	(2) Discipline	(3) Housekeeping	(4) WRAT	(5) Internalizing	(6) Externalizing
Crimes / 10 Residents	2.06^{**}	0.57	0.26	1.55^{*}	0.86^{+}	3.98^{**}
	(0.80)	(0.46)	(0.66)	(0.62)	(0.49)	(0.67)
Observations	7481	7481	7261	4914	7271	7271

Notes:
Regression specification is column 3 of Table 6, with dependent variable changed.
All dependent variables standardized by age group of child.
Subset, Parenting: first axis, derived by MCA on subset of parenting behaviors.
WRAT: Wide Range Achievement Test.
Internalizing: Problematic behavior index, associated with depression, withdrawal, etc.
Externalizing: Problematic behavior index, associated with attention deficit, hyperaggression, etc.
Standard errors clustered by initial neighborhood.

Table 9: Heterogeneous Neighborhood Quality on Parental Involvement

	Interaction: (Crimes / 10 Residents) x (Family Characteristic at Initial Survey)								
	Parental Involvement	Neighborhood Income	Child WRAT	Age Cohort		Child Gender		Parent Education	
1st Quartile	-4.62***	2.21*	2.16	0-3	0.71	Male	3.98***	No HS	2.59*
	(0.92)	(0.91)	(1.99)		(1.25)		(1.19)		(1.19)
2nd Quartile	3.08**	2.59***	3.48***	3-6	2.87**	Female	2.52***	Some HS	3.28**
	(1.10)	(0.77)	(1.06)		(0.94)		(0.75)		(1.27)
3rd Quartile	5.68**	6.06	4.86*	6-9	2.26**		.	HS Degree	5.73
	(1.78)	(5.33)	(2.09)		(0.80)		(.)		(3.70)
4th Quartile	10.37***	19.11	3.36**	9-12	0.26		.	Post HS	2.92**
	(2.42)	(16.05)	(1.05)		(0.56)		(.)		(0.96)
	.	.	.	12+	6.36^{+}		.	College Degree	3.18**
	(.)	(.)	(.)		(3.31)		(.)		(1.22)
Continuous	6.78***	4.47	.	Years	0.68*		.	Years	0.02
	(1.13)	(3.28)	(.)		(0.28)		(.)		(0.16)
	0.00	0.65	0.57	Equal?	0.01		0.05		0.89

Notes:
Instruments for neighborhood quality: Relocations to family's initial neighborhood in the previous year, as well as its interaction with family characteristic.
Relocations is a constructed measure: the product of yearly closures in all other neighborhoods and 1990 neighborhood black share.
Controls: employment marital status, and occupancy type of parent, as well as subject fixed effect.
Standard errors clustered by initial neighborhood.
Equal?: shows p-value for test of joint equality of coefficients on all categories.
Years / Continuous: show coefficients from (separate) regression of continuous variable interacted with neighborhood quality.

Figures

Figure 1: Sample Retention by Neighborhood Income Quartile and Survey Wave

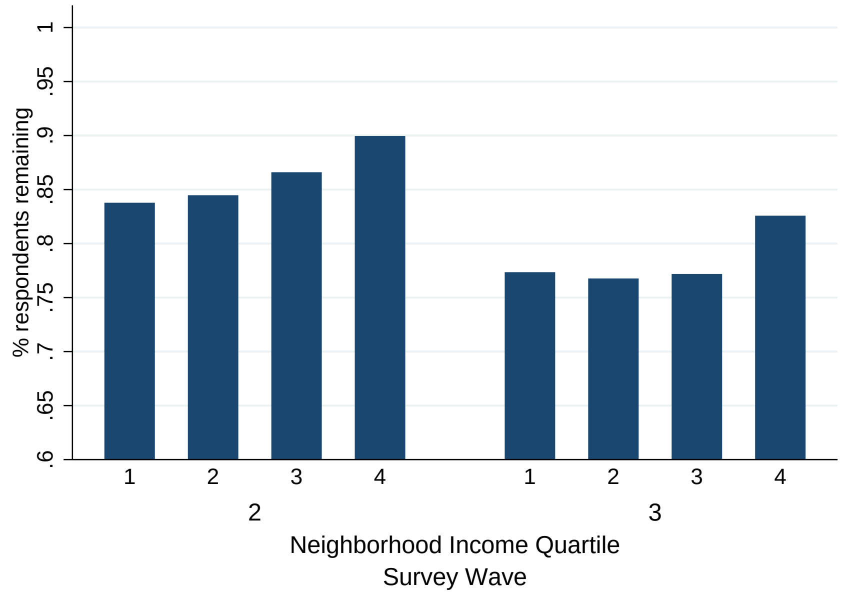

Notes:
Retention rates measured as a percentage of wave 1 respondents.
Neighborhood income: 1990 average family income of initial neighborhood.

Figure 2: Closures and Relocations in Chicago

(a) Total Number of CHA Public Housing Units Closed by Year

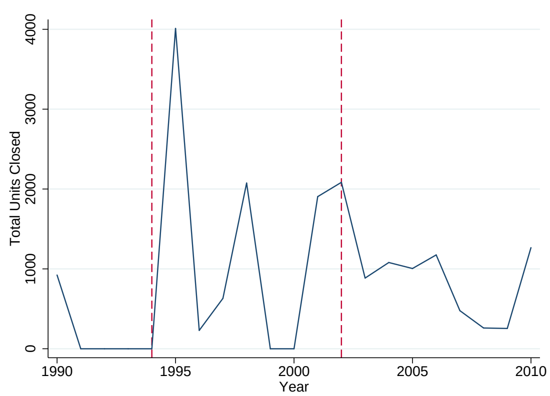

(b) Histogram of Community Area Relocation Shares

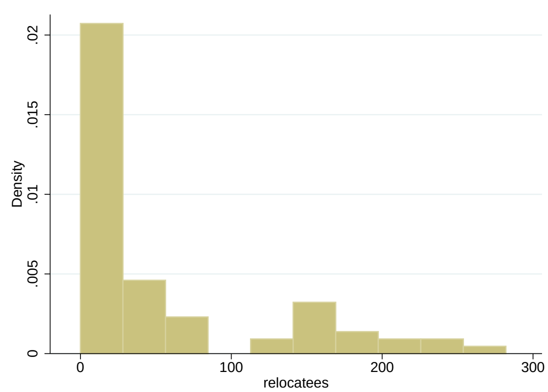

Figure 3: Histogram of Relocations Received per Capita

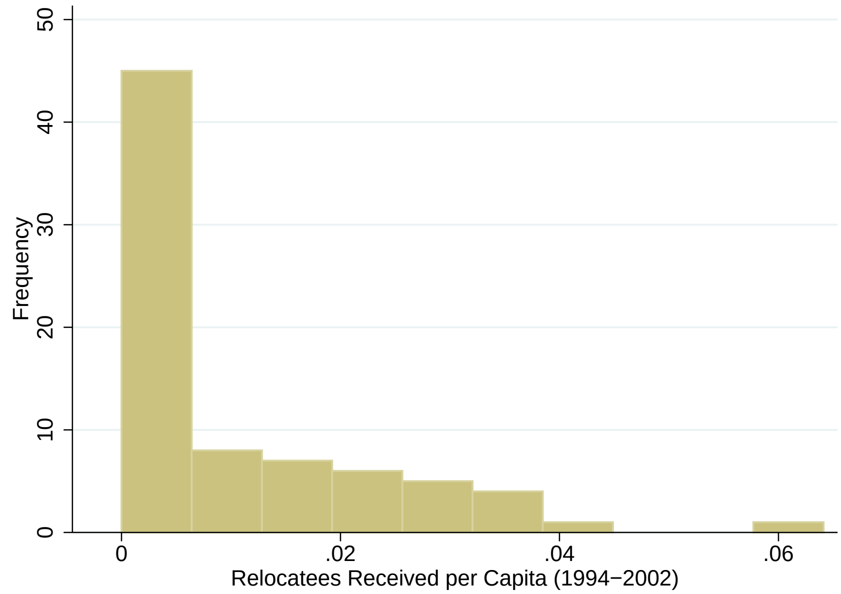

Notes:
Neighborhood: community area. Relocations calculated using shift-shares method.
Shifts: total public housing closures across all of Chicago (1994-2002).
Shares: fraction of relocations received by neighborhood, measured in 2008.
Each public housing unit is assumed to contain 2.2 individuals (mean).
Neighborhood population measured in 1990.

Figure 4: Effects of Relocation Leads/Lags on Neighborhood Crime Rate

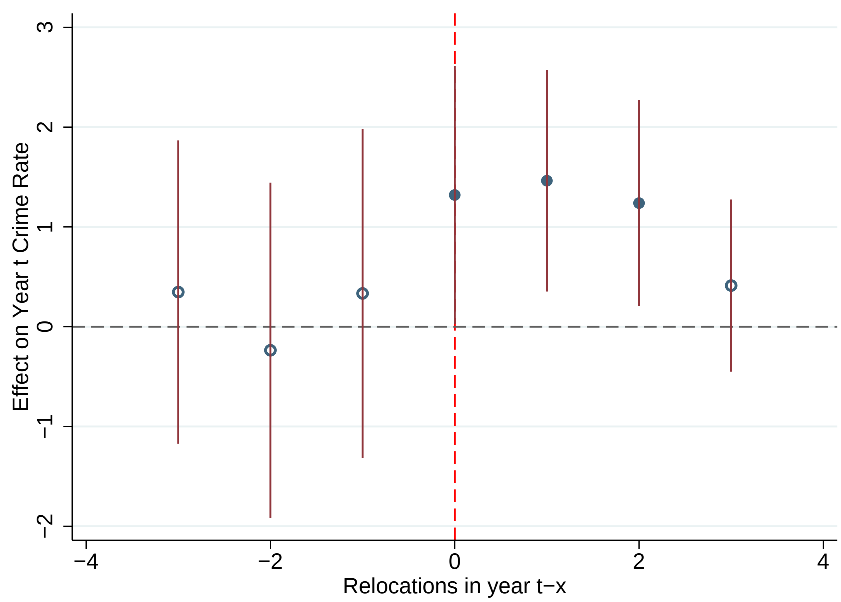

Notes:
Sample: All neighborhoods (NCs) in Chicago, years 1992 to 2004.
Controls: Neighborhood fixed effect, year effect.
90% confidence intervals shown, calculated using robust standard errors.

Figure 5: Heterogeneous Effects of Neighborhood Crime Rate on Parental Involvement, by Initial Involvement Quartile

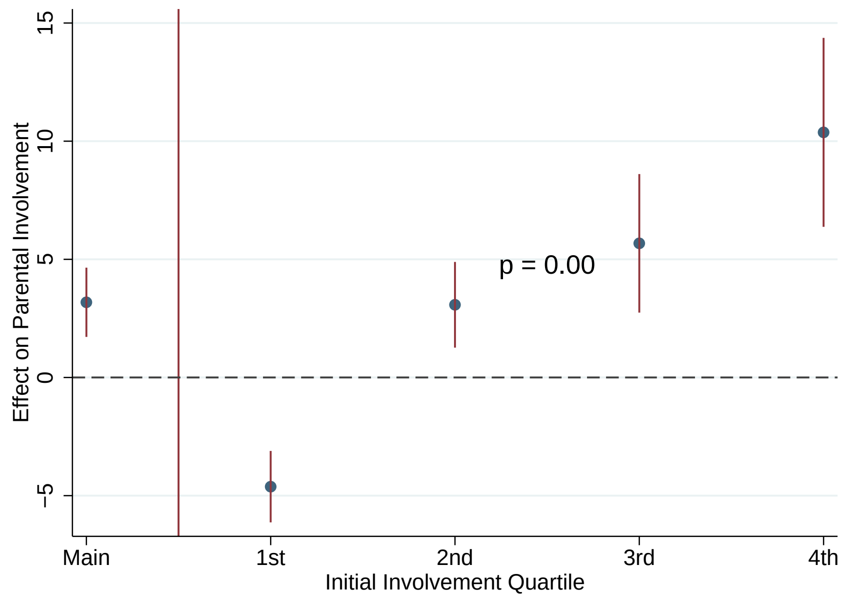

Notes:
Instrument for neighborhood crime rate: interaction between 1990 neighborhood black share, and closures in all other neighborhoods last year.
Neighborhood quality measured using family's current neighborhood. Instrument measured using family's initial (Wave 1) neighborhood.
Controls: parental employment/marital status, occupancy type, subject fixed effect.
Spikes indicate 90% confidence intervals.
P-value displayed for test of joint equality across all categories.

Figure 6: Heterogeneous Effects of Neighborhood Crime Rate on Parental Involvement, by Initial Neighborhood Income

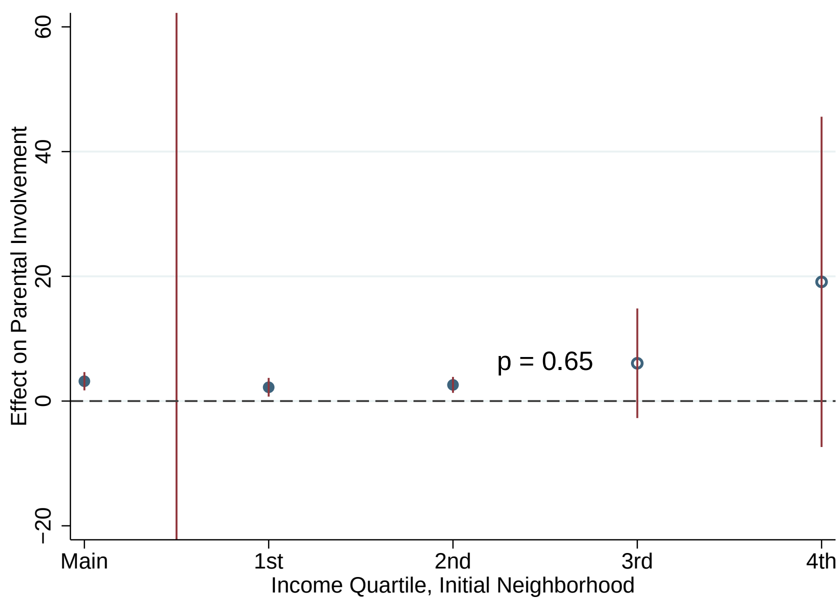

Notes:
Instrument for neighborhood crime rate: interaction between 1990 neighborhood black share, and closures in all other neighborhoods last year.
Neighborhood quality measured using family's current neighborhood. Instrument measured using family's initial (Wave 1) neighborhood.
Initial neighborhood income measured in year 1990.
Controls: parental employment/marital status, occupancy type, subject fixed effect.
Spikes indicate 90% confidence intervals.
P-value displayed for test of joint equality across all categories.

Figure 7: Heterogeneous Effects of Neighborhood Crime Rate on Parental Involvement, by Initial Child WRAT

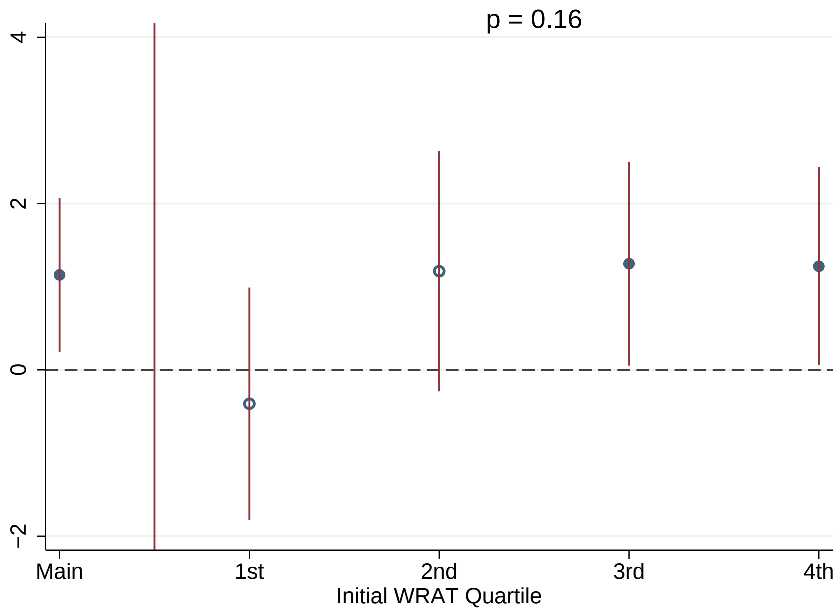

Notes:
Instrument for neighborhood crime rate: interaction between 1990 neighborhood black share, and closures in all other neighborhoods last year.
Neighborhood quality measured using family's current neighborhood. Instrument measured using family's initial (Wave 1) neighborhood.
Controls: parental employment/marital status, occupancy type, subject fixed effect.
Spikes indicate 90% confidence intervals.
P-value displayed for test of joint equality across all categories.

Figure 8: Heterogeneous Effects of Neighborhood Crime Rate on Parental Involvement, by Child Age / Gender

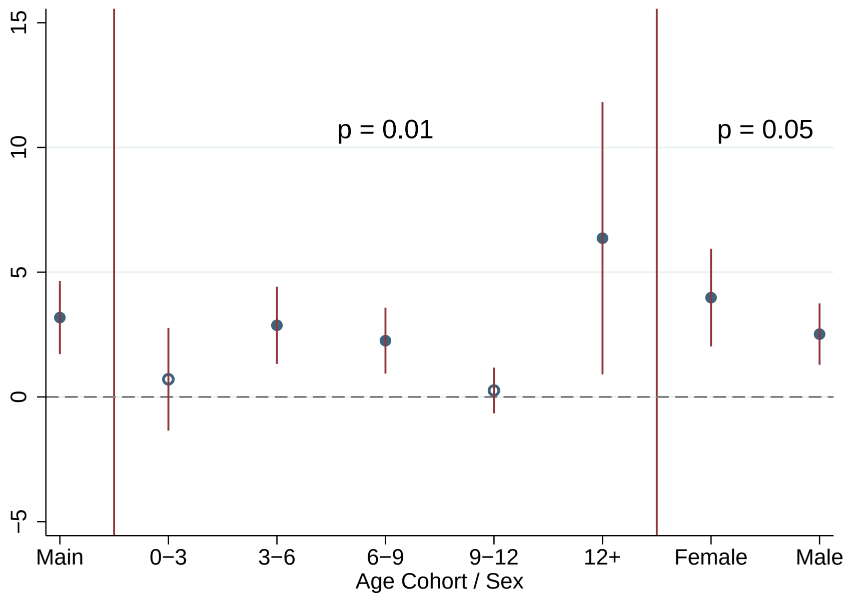

Notes:
Points denote neighborhood crime effect for population subgroup indicated on horizontal axis.
Age cohorts of child measured in Wave 1.
Instrument for neighborhood crime rate: interaction between 1990 neighborhood black share, and closures in all other neighborhoods last year.
Neighborhood quality measured using family's current neighborhood. Instrument measured using family's initial (Wave 1) neighborhood.
Controls: parental employment/marital status, occupancy type, subject fixed effect.
Spikes indicate 90% confidence intervals.
P-value displayed for test of joint equality across all categories.

Figure 9: Heterogeneous Effects of Neighborhood Crime Rate on Parental Involvement, by Parent Education

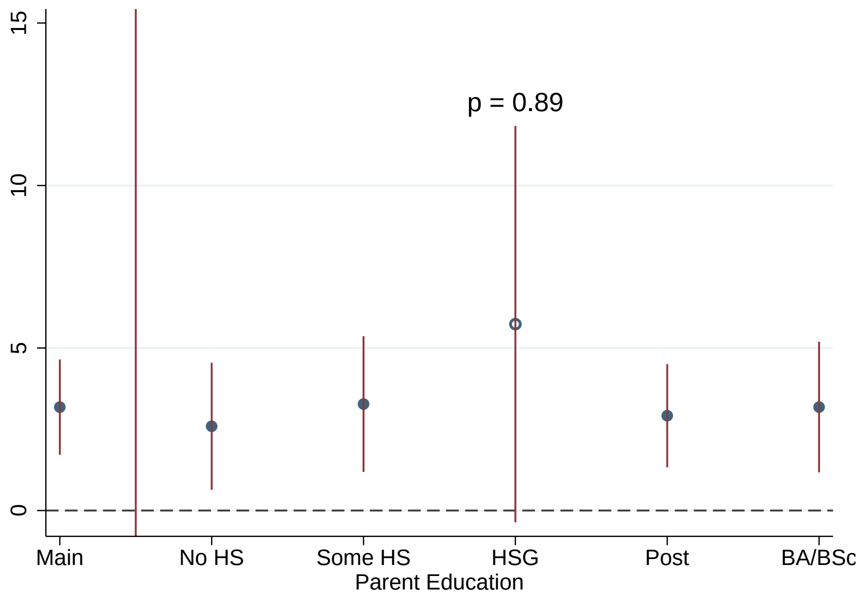

Notes:
Instrument for neighborhood crime rate: interaction between 1990 neighborhood black share, and closures in all other neighborhoods last year.
Neighborhood quality measured using family's current neighborhood. Instrument measured using family's initial (Wave 1) neighborhood.
Controls: parental employment/marital status, occupancy type, subject fixed effect.
Spikes indicate 90% confidence intervals.
P-value displayed for test of joint equality across all categories.

Appendix A Data Appendix

A.1 PHDCN Sampling

The panel data comes from the Longitudinal Cohort Survey of the PHDCN. Aside from this component, the PHDCN also contains a cross-sectional Community Survey and a Systematic Social Observation component, which characterized neighborhood block faces. These two components are not used in this paper.

For the PHDCN, 847 census tracts in Chicago were collapsed into 343 neighborhood clusters to form neighborhood clusters: the unit of analysis for residential neighborhoods in this study. The process of aggregation was guided by the following principles: first, the tracts to be aggregated were to be spatially contiguous. Second, these tracts were to be relatively homogeneous along a set of key census variables. In addition, the boundaries of neighborhood clusters were to respect physical boundaries (e.g. railroads and highways) and conform to local definitions of neighborhoods.

The sampling of respondents to the Longitudinal Cohort Study was performed as follows: first, neighborhood clusters were stratified by racial composition and income level, forming 21 cells. Next, 80 neighborhood clusters were selected from the 343 in Chicago, using with sampling weights designed to produce equal representation across each cell. Within each of these 80 initial neighborhoods, census block groups were selected randomly, and residents living in those block groups were contacted for participation in the survey. Within these block groups, children as well as young adults (aged 18 and above) were eligible for the survey. Because the young adults were mostly living independently of parental figures, they are not dropped from the data for this study.

A.2 Neighborhood Crime Data

Data on neighborhood crime rates was originally hand-collected by D. Garth Taylor, and has been used to study the effect of housinc choice vouchers on neighborhood crime rates Lens et al. (2011). This data was generously shared by Michael Lens. The data contains the total number of crimes reported within the census tract, per 1000 residents. These are available at a yearly frequency, and hence can capture the immediate change in neighborhood quality following incoming relocations. These crime rates are aggregated from the tract level to the neighborhood cluster level.

Appendix B Model of Neighborhood Selection

Suppose that, in each period t, parents choose j' according to (16). In selecting neighborhood j', parents value neighborhood influence on the child through U_{it}, and also incur some other benefits/costs from neighborhood residence, which are captured in the function $W(\cdot)$. This function contains the parental involvement shock ε_{it}, to reflect the idea that common factors shift parental involvement preferences and the parent's willingness or ability to reside in certain neighborhoods. Because there may be uncertainty about the t' values of neighborhood quality or the parental involvement shock, parents maximize expected utility, conditional on their information set $\mathcal{I}_{it}$.

$$j' \equiv j(i,t') = \underset{j \in \mathcal{J}}{\operatorname{argmax}} \, E\left[U_{ijt'}|\mathcal{I}_{it}\right] + E\left[W(F_i, q_{jt'}, \varepsilon_{it'})|\mathcal{I}_{it}\right] + \nu_{ijt} \tag{16}$$

Suppose that, when choosing neighborhoods, parents know $\varepsilon_{it'}$ and can perfectly forecast $q_{j't'}$ at the point of choosing neighborhood j'. Because of the quadratic form chosen, U_{ijt} is constant in neighborhood quality, so this term drops out of the decision problem.[36] As a result, the conditional expectation of neighborhood quality in the chosen neighborhood is given by (17). If $W_{q\varepsilon}$ is nonzero, for two arbitrary values of neigborhood quality $\overline{q}$ and $\underline{q}$, the difference $W(F_i, \overline{q}, \varepsilon_{it'}) - W(F_i, \underline{q}, \varepsilon_{it'})$ will depend on the value of $\varepsilon_{it'}$. As a result, the expected value of $q_{j't'}$ will be dependent on the value of $\varepsilon_{it'}$. For arbitrary distributions of shocks (ε, ν), it is likely that $q_{j't'}$ and $\varepsilon_{it'}$ will also be correlated.

$$E\left[q_{j't'} \middle| j' \text{ chosen}\right] = E\left[q_{j't'} \middle| W(F_i, q_{j't'}, \varepsilon_{it'}) - W(F_i, q_{kt'}, \varepsilon_{it'}) > \nu_{ij't} - \nu_{ikt}, \, \forall k \in \mathcal{J}\right] \tag{17}$$

To explicitly show the potential correlation between neighborhood quality and the parental involvement shock, consider the case where $W(F, q, \varepsilon) = (F + \varepsilon)q$. Then, because $\varepsilon_{it'}$ is assumed to be mean-zero, the covariance between $q_{j't'}$ and $\varepsilon_{it'}$ is given by (18). In general, the covariance is non-zero. Furthermore, if the shocks ε and ν are mutually independent, then the covariance is positive. This result is intuitive: by construction, a high parental involvement shock also increases the marginal benefit (relative to cost) that parents derive from living in a high-quality neighborhood. Hence, parental residential choices should reflect that preference.

[36] If U_{ijt} was chosen to be a different function, then this term would also have to be accounted for but would not change the main argument.

$$\text{Cov}\left[q_{j't'}, \varepsilon_{it'} | j' \text{ chosen}\right] = E\left[q_{j't'}\varepsilon_{it'} \middle| q_{j't'}\varepsilon_{it'} > q_{kt'}\varepsilon_{it'} + \nu_{ikt} - \nu_{ij't}, \forall k \in \mathcal{J}\right] \tag{18}$$

According (17) or (18), estimating (9) using a fixed effects regression is likely to result in biased estimates of α. Conditioning on fixed effects is not sufficient to address this problem: in (18) the covariance of $q_{j't'}$ and $\varepsilon_{it'}$ is independent of F_i, and non-zero in general.

Appendix C Crime Rate as a Measure of Neighborhood Quality

Suppose that true neighborhood quality is the sum of crime rate (with a coeffcient of -1) and an unobserved measure ξ_{jt}. (19).

$$q_{jt} = -v_{jt} + \xi_{jt} \tag{19}$$

For discussion purposes, suppose that ξ_{jt} represents school quality. If one unit of relocations raises crime rate and lowers school quality by θ_v and θ_ξ respectively, then the reduction in true neighborhood quality is the sum of these two: $\theta_v + \theta_\xi$. Hence, under the likely assumption that relocations change crime rate and school quality in different directions, the change in the crime rate alone will potentially understate the true change in neighborhood quality. In turn, estimates of the effect of neighborhood quality on parental involvement will be tend to be overstated in magnitude.

Suppose now, without loss of generality, that (19) is scaled such that v_{jt} has standard deviation of 1. Now, consider the effect of relocations on standardized neighborhood quality. This is given by $\theta_q = (\theta_v + \theta_\xi) / \sqrt{1 + \sigma_\xi - 2\rho_{v\xi}\sigma_\xi}$, where $\rho_{v\xi}$ denotes the correlation between crime and school quality. It is likely that the correlation is negative, so attention is restricted to this case. The relation between θ_q and θ_v is given by (20). In particular, in the case where $\rho_{v\xi}$ is negative, $\sqrt{1 + \sigma_\xi^2 - 2\rho_{v\xi}\sigma_\xi} - 1$ is less than σ_ξ. Then, under the additional assumption that θ_ξ/σ_ξ is approximately equal to θ_v, (20) shows that θ_q will be understated by θ_v. This last condition has the following interpretation: θ_ξ/σ_ξ is the effect of relocations on standard deviation units of school quality. As long as this effect is of roughly the same size as the effect of relocations on standard deviation units of crime rate, using crime rate in place of true neighborhood quality will not cause IV estimates to be overstated in magnitude. In fact, under all these assumptions, the estimates of the effect of neighborhood quality on parental involvement would be even more significant if the true measure of neighborhood quality were used. This reasoning provides some motivation for considering standard deviation units of crime rate, when estimating the effects on parental involvement.

$$\theta_q \geq \theta_v \iff \frac{\theta_\xi}{\sqrt{1+\sigma_\xi^2 - 2\rho_{v\xi}\sigma_\xi} - 1} \geq \theta_v \tag{20}$$

Appendix D Parental Involvement Choice under General Functional Forms

Consider a general utility function (21) for parental involvement choice. The assumptions The assumption that utility is linear in behavior is without loss of generality: if $b(y,q)$ were instead $u(b(y,q))$ for an increasing function y, the first order condition would be given by $\omega_i u_b(y,q) - c_y(y,q) = 0$. This can be rewritten as $\omega_i - \frac{c_y(y,q)}{u_b(y,q)} = 0$. Hence, by redefining the cost $c(y,q)$ as $\int \frac{c_y(y,q)}{u_b(y,q)} dy$, the maximization problem is unchanged.

$$U_i = \omega_i b(y,q) - l(1 - s(q)y) \tag{21}$$

The derivative of parental involvement with respect to neighborhood quality is given by (22). Because some parents respond in a compensating manner and others reinforce neighborhood change, this expression switches sign depending on ω_i.

$$\frac{\partial y}{\partial q} = \frac{\omega_i b_{yq} - c_{yq}}{\omega_i b_{yy} - c_{yy}} \tag{22}$$

The relevant cross-partial derivative describing this phenomenon is given by (23). It is the sum of two terms. The second term is related to the curvature of marginal productivity and marginal cost, and is difficult to sign. For now, I assume that this term is zero, and focus on the first term. Under standard assumptions ($b_{yy} < 0$ and $c_{yy} > 0$), the negative association between initial involvement and response to neighborhood quality implies that either $b_{yq} < 0$ and/or $c_{yq} < 0$. Assume that b_{yq} and c_{yq} have consistent sign across all families. To establish a contradiction, suppose that $b_{yq} > 0$: in that case, c_{yq} must be negative. However, from (22), this implies that the response of parental involvement to neighborhood quality is always negative: this is at odds with the empirical findings. Ultimately, this implies that b_{yq} must be negative.

$$\frac{\partial y}{\partial q \partial \omega_i} = \frac{b_{yq} c_{yy} - b_{yy} c_{yq}}{\left(\omega_i b_{yy} - c_{yy}\right)^2} + \frac{b_y \left[\omega \left(b_{yyq} - b_{yyy}\right) - \left(c_{yyq} - c_{yyy}\right)\right]}{\left(\omega_i b_{yy} - c_{yy}\right)^3} \tag{23}$$

If the second term in (23) is non-zero, then the above claim does not necessarily hold. Certain assumptions must be imposed before a conclusive statement can be made. One sufficient condition for the previous claim to hold is that $b_{yyy} = b_{yyq}$ and, at the same time $c_{yyy} = c_{yyq}$. This condition restricts the way neighborhood quality affects the curvature of the productivity (cost) of parental involvement.

Supplementary Tables

Table A1: List of Raw Variables in HOME Score

	Age group 1	Age group 2	Age group 3
Home Environment	Home is not dark	Home is not dark	Home is not dark
	Home is reasonably clean	Home is reasonably clean	Home is reasonably clean
	Home is minimally cluttered	Home is minimally cluttered	Home is minimally cluttered
	Play environment is safe (home or building for ages 36 mos +)	Play environment is safe (home or building for ages 36 mos +)	Play environment is safe (home or building for ages 36 mos +)
	home has at least 100 square feet space/person?	home has at least 100 square feet space/person?	home has at least 100 square feet space/person?
	home not too noisy from noise in house?	home not too noisy from noise in house?	home not too noisy from noise in house?
	home not too noisy from noise outside?	home not too noisy from noise outside?	home not too noisy from noise outside?
	no signs of alcohol/illegal drug use?	no signs of alcohol/illegal drug use?	no signs of alcohol/illegal drug use?
Family Activities	Child eats meal with both mother and father(-figure) once a day or more	Child eats meal with both mother and father(-figure) once a day or more	Child eats meal with both mother and father(-figure) once a day or more
	Child taken to museum in past year	Child taken to museum in past year	Child taken to museum in past year
	Mom showed physical affection to child	Family visits with family or friends 2-3 times a month	Family visits with family or friends 2-3 times a month
	Child receives lessons or belongs to sports/music/art/dance/ drama org	Child receives lessons or belongs to sports/music/art/dance/ drama org	Child receives lessons or belongs to sports/music/art/dance/ drama org
		Talk with child as works about house?	
		SP included in family activities	

	Age group 1	Age group 2	Age group 3
Discipline	Mom reports no more than 1 spank during past week	Mom reports no more than 1 spank during past week	Child expected to do routine chores
	Not more than 1 physical punishment / month	Not more than 1 physical punishment / month	Have rules for child's behavior with friends
	TV is on in home less than 5 hours per day	TV is on in home less than 5 hours per day	Saw child's friends, last week
			Mom reports no more than 1 spank during past week
			Not more than 1 physical punishment / month
			Child has curfew for weekend nights
			Child has curfew for school nights
			Fixed time for child to do homework
			Lost temper with child not more than once, past week
Other	Mom encouraged child to contribute to conversation	Mom encouraged child to contribute to conversation	Mom encouraged child to contribute to conversation
	Mom answered child's questions or requests verbally	Mom answered child's questions or requests verbally	Mom answered child's questions or requests verbally
	Mom spontaneously vocalize to/conversed with child at least twice	Mom spontaneously vocalize to/conversed with child at least twice	
	Mom's voice conveyed positive feeling about child	Child has 10 children's books	
		Child encouraged to read several times a week for enjoyment	

Notes:
Age group 1: Cohort 3 - Wave 1, Cohort 0 - Wave 2.
Age group 2: Cohort 6 - Wave 1, Cohort 3 - Wave 2, Cohort 0 - Wave 3.
Age group 3: Cohort 12, Cohort 9, Cohort 6 - Wave 2 & 3, Cohort 3 - Wave 3.
Shaded cells based on interviewer observation. Unshaded cells self-reported.
Exact text in questionnaire varies slightly across surveys.

Table A2: First Stage: Effect of Relocations Received on Neighborhood Crime

Share Component	Concurrent	Historical
	(1)	(2)
Relocations Received:		
t-1	8.989*	8.426
	(4.362)	(1007.6)
t-2	10.99+	9.033
	(5.775)	(1127.9)
F statistic	1.588	0.996
Observations	5831	5831

Notes:
Effects take units of standard deviations per 100 households received.
Neighborhood: refers to study-defined Neighborhood Cluster (343 Neighborhood Clusters in Chicago).
Relocations constructed as product of Shift Component and Share Component (Concurrent or Historical).
Shift Component: City-wide number of public housing apartments closed, excluding receiving neighborhood.
Concurrent Shares: Fraction of ex-public housing residents living in receiving neighborhood in 2008.
Historical Shares: Predicted value of Concurrent Share component, using 1990 neighborhood characteristics.

Table A3: Principal Component Weights of Parenting Behavior Measures in HOME Score

	Age Bracket		
	1	2	3
1+ family meal/day	-0.196	0.0551	0.213
Participates in child organization	0.0612	0.457	0.332
Brought child to museum, past year	-0.192		
2+ visits with family/friends per month		-0.570	0.213
Family Activities		0.0701	
Owns 3 children's books		0.505	
Encourages child to read at home		0.229	
3+ puzzles for sub		0.313	
Encourages child to contribute	0.425	0.596	0.620
Answers child	0.396	0.341	0.873
Talks to child twice during visit	0.467	0.596	
Introduces observer to child	0.290	0.477	0.745
Speaks positively to child	-0.0534		
Praises child twice	0.547		
Hugs child	1.054	0.278	
Daily conversation		0.405	-0.0351
Physical punishment less than once/week	-0.182	0.120	0.161
Does not slap/spank child	-1.931	-2.807	-1.074
Lost temper with child less than twice, last week			0.0225
Assigns chores to child			-1.586
Has rules for behavior with friends			0.141
Saw child's friends, last week			0.205
Sets and enforces rules about homework			0.340
Implements curfew, weekend nights			-0.162
Supervises child in public			0.227
TV used judiciously	0.00926		
Observations	1728	2251	5807

Notes:
Age bracket 1: Cohort 3 - Wave 1, Cohort 0 - Wave 2.
Age bracket 2: Cohort 6 - Wave 1, Cohort 3 - Wave 2, Cohort 0 - Wave 3.
Age bracket 3: Cohort 12, Cohort 9, Cohort 6 - Waves 2 & 3, Cohort 3 - Wave 3.

Supplementary Figures

Figure A1: Definition of Closure (C) and Receiving (R) Neighborhood Populations

(a) Before Closures

(b) After Closures

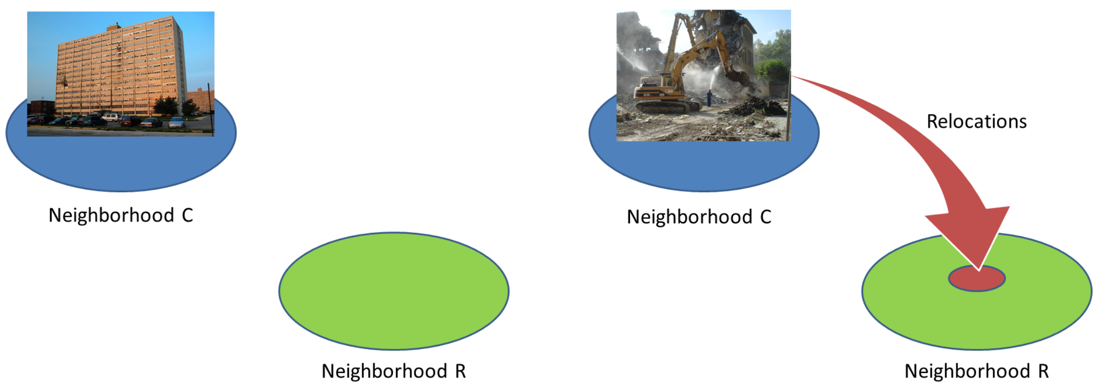

Figure A2: Heterogeneous Effects of Neighborhood Crime Rate on Parental Involvement, by Initial Involvement Quartile Relative to Initial Neighborhood

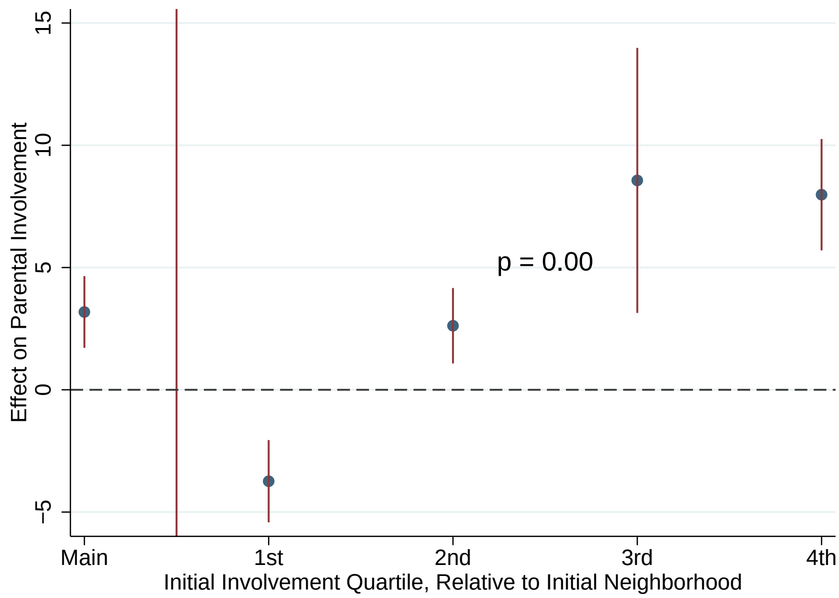

Notes:
Instrument for neighborhood crime rate: interaction between 1990 neighborhood black share, and closures in all other neighborhoods last year.
Neighborhood quality measured using family's current neighborhood. Instrument measured using family's initial (Wave 1) neighborhood.
Controls: parental employment/marital status, occupancy type, subject fixed effect.
Spikes indicate 90% confidence intervals.
P-value displayed for test of joint equality across all categories.

Figure A3: Heterogeneous Effects of Relocations Received on Child Externalizing Behavior, by Initial Neighborhood / Initial Involvement

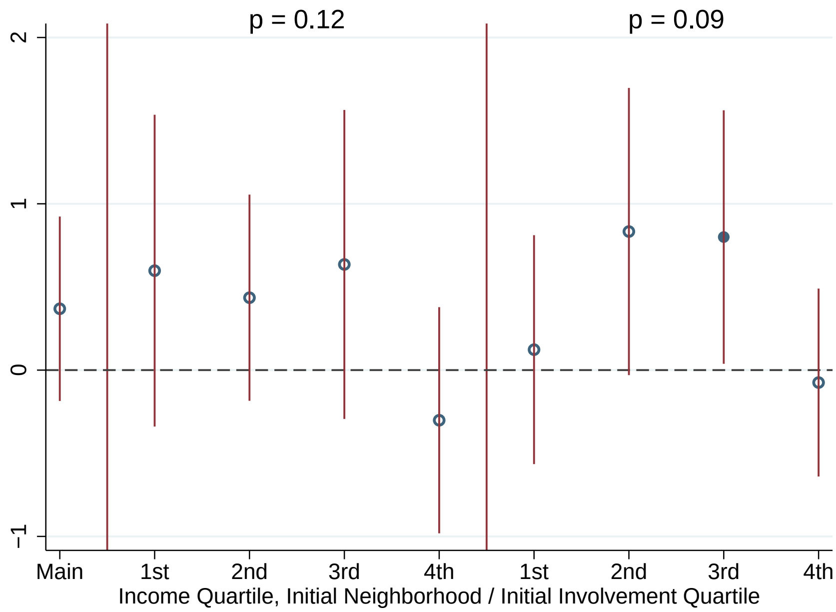

Notes:
Neighborhood income of Wave 1 neighborhood measured in 1990.
Points denote coefficient on interaction between listed categorical variable and relocations received last year.
Controls: cumulative prior relocations (up to $t-2$) parental employment, subject fixed effect, neighborhood effect, year dummies.
Spikes indicate 90% confidence intervals.
P-value displayed for test of joint equality across all categories.

Figure A4: Heterogeneous Effects of Relocations Received on Child Internalizing Behavior, by Initial Neighborhood / Initial Involvement

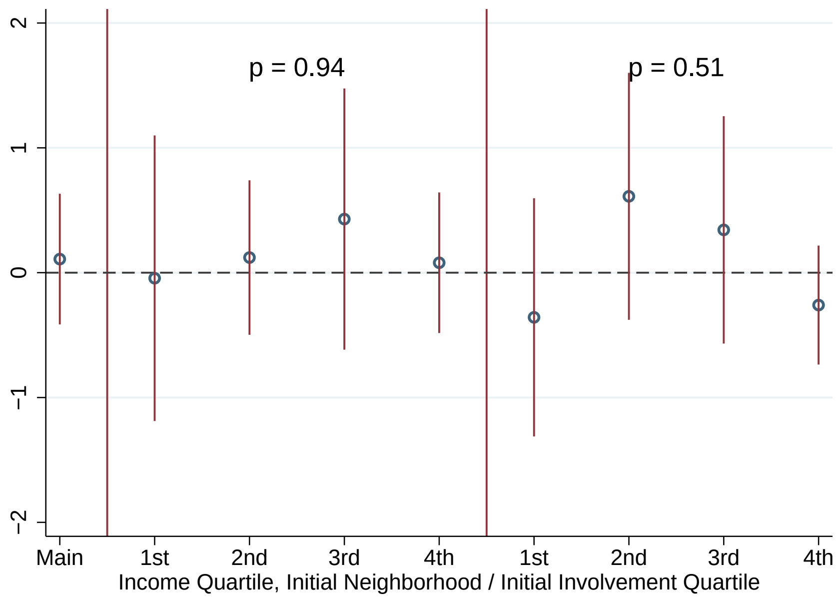

Notes:
Neighborhood income of Wave 1 neighborhood measured in 1990.
Points denote coefficient on interaction between listed categorical variable and relocations received last year.
Controls: cumulative prior relocations (up to $t-2$) parental employment, subject fixed effect, neighborhood effect, year dummies.
Spikes indicate 90% confidence intervals.
P-value displayed for test of joint equality across all categories.

Figure A5: Heterogeneous Effects of Relocations Received on Child WRAT, by Initial Neighborhood / Initial Involvement

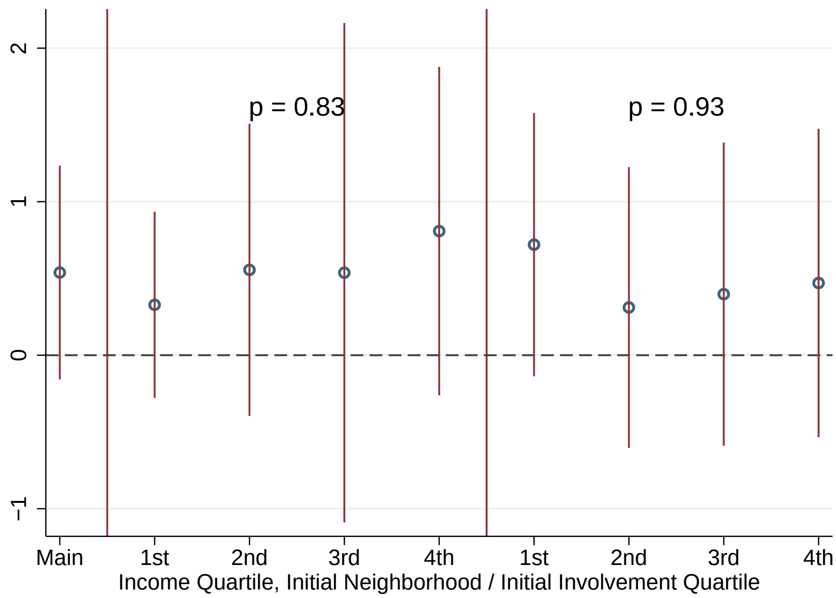

Notes:
Neighborhood income of Wave 1 neighborhood measured in 1990.
Points denote coefficient on interaction between listed categorical variable and relocations received last year.
Controls: cumulative prior relocations (up to $t-2$) parental employment, subject fixed effect, neighborhood effect, year dummies.
Spikes indicate 90% confidence intervals.
P-value displayed for test of joint equality across all categories.

References

Aliprantis, D. and D. Hartley (2015). Blowing it up and knocking it down: The local and city-wide effects of demolishing high concentration public housing on crime. *Journal of Urban Economics 88*, 67–81.

Amato, P. R. and F. Rivera (1999). Paternal involvement and children's behavior problems. *Journal of Marriage and the Family*, 375–384.

Andersson, F., J. Haltiwanger, M. J. Kutzbach, G. Palloni, H. Pollakowski, and D. H. Weinberg (2016). Childhood housing and adult earnings: A between-siblings analysis of housing vouchers and public housing. *US Census Bureau Center for Economic Studies Paper No. CES-WP-13-48R*.

Austen, B. (2012). The last tower. *Harper's Magazine*.

Avvisati, F., B. Besbas, and N. Guyon (2011). Parental involvement in school: A literature review. *Revue d'économie politique 120*(5), 759–778.

Avvisati, F., M. Gurgand, N. Guyon, and E. Maurin (2014). Getting parents involved: A field experiment in deprived schools. *The Review of Economic Studies 81*(1), 57–83.

Bennett, L., J. L. Smith, and P. A. Wright (2015). *Where are poor people to live?: Transforming public housing communities*. Routledge.

Bergman, P. (2015). Parent-child information frictions and human capital investment: Evidence from a field experiment. *SSRN Working Paper*.

Case, A. C. and L. F. Katz (1991). The company you keep: The effects of family and neighborhood on disadvantaged youths. Technical report, National Bureau of Economic Research.

CHA (2011). The plan for transformation: An update on relocation. Technical report, Chicago Housing Authority.

Chetty, R., N. Hendren, and L. F. Katz (2016). The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *The American Economic Review 106*(4), 855–902.

Chicago Tribune (Oct 13, 2002). On moving day, past gets packed away at Cabrini.

Chyn, E. (2016). Moved to opportunity: The long-run effect of public housing demolition on labor market outcomes of children. *Unpublished paper. University of Michigan, Ann Arbor*.

Cunha, F., J. J. Heckman, and S. M. Schennach (2010). Estimating the technology of cognitive and noncognitive skill formation. *Econometrica 78*(3), 883–931.

Desmond, M. and T. Shollenberger (2015). Forced displacement from rental housing: Prevalence and neighborhood consequences. *Demography 52*(5), 1751–1772.

Fan, X. and M. Chen (2001). Parental involvement and students' academic achievement: A meta-analysis. *Educational psychology review 13*(1), 1–22.

Fryer Jr, R. G., S. D. Levitt, and J. A. List (2015). Parental incentives and early childhood achievement: A field experiment in Chicago Heights. Technical report, National Bureau of Economic Research.

Garza, M. M. (1999). 9 high-rises at Taylor Homes slated to close. *Chicago Tribune*.

Gaviria, A. and S. Raphael (2001). School-based peer effects and juvenile behavior. *Review of Economics and Statistics 83*(2), 257–268.

Gelber, A. and A. Isen (2013). Children's schooling and parents' behavior: Evidence from the Head Start Impact Study. *Journal of Public Economics 101*, 25–38.

Gershoff, E. T. (2002). Corporal punishment by parents and associated child behaviors and experiences: a meta-analytic and theoretical review. *Psychological bulletin 128*(4), 539.

Green, C. L., J. M. Walker, K. V. Hoover-Dempsey, and H. M. Sandler (2007). Parents' motivations for involvement in children's education: An empirical test of a theoretical model of parental involvement. *Journal of educational psychology 99*(3), 532.

Guryan, J., E. Hurst, and M. Kearney (2008). Parental education and parental time with children. *The Journal of Economic Perspectives*, 23–46.

Hoffmann, J. P. and M. J. Dufur (2008). Family and school capital effects on delinquency: Substitutes or complements? *Sociological Perspectives 51*(1), 29–62.

Hsin, A. (2012). Is biology destiny? birth weight and differential parental treatment. *Demography 49*(4), 1385–1405.

Hunt, D. B. (2009). *Blueprint for disaster: The unraveling of Chicago public housing.* University of Chicago Press.

Imberman, S. A., A. D. Kugler, and B. I. Sacerdote (2012). Katrina's children: Evidence on the structure of peer effects from hurricane evacuees. *The American Economic Review 102*(5), 2048–2082.

Jacob, B. A. (2004). Public housing, housing vouchers, and student achievement: Evidence from public housing demolitions in Chicago. *American Economic Review*, 233–258.

Jacob, B. A., M. Kapustin, and J. Ludwig (2015). The impact of housing assistance on child outcomes: Evidence from a randomized housing lottery. *The Quarterly Journal of Economics 130*(1), 465–506.

Jencks, C. and S. E. Mayer (1990). The social consequences of growing up in a poor neighborhood. *Inner-city poverty in the United States 111*, 186.

Jeynes, W. H. (2007). The relationship between parental involvement and urban secondary school student academic achievement a meta-analysis. *Urban education 42*(1), 82–110.

Katz, L. F., J. R. Kling, and J. Liebman (2005). Bullets don't got no name: Consequences of fear in the ghetto.

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

Lens, M. C., I. G. Ellen, and K. O'Regan (2011). Do vouchers help low-income households live in safer neighborhoods? evidence on the housing choice voucher program. *Cityscape*, 135–159.

Leventhal, T. and J. Brooks-Gunn (2003). Moving To Opportunity: an experimental study of neighborhood effects on mental health. *American Journal of Public Health 93*(9), 1576–1582.

Liu, J. (2004). Childhood externalizing behavior: theory and implications. *Journal of child and adolescent psychiatric nursing 17*(3), 93–103.

Morenoff, J. D., R. J. Sampson, and S. W. Raudenbush (2001). Neighborhood inequality, collective efficacy, and the spatial dynamics of urban violence. *Criminology 39*(3), 517–558.

Moser, W. (2012). Chicagos Lathrop Homes: A reverse Potemkin village on the river. *Chicago Magazine*.

Nakajima, R. (2007). Measuring peer effects on youth smoking behaviour. *The Review of Economic Studies 74*(3), 897–935.

Oakley, D. and K. Burchfield (2009). Out of the projects, still in the hood: The spatial constraints on public-housing residents' relocation in Chicago. *Journal of Urban Affairs 31*(5), 589–614.

Patacchini, E. and Y. Zenou (2011). Neighborhood effects and parental involvement in the intergenerational transmission of education. *Journal of Regional Science 51*(5), 987–1013.

Pavan, R. (2013). On the production of skills and the birth order effect. Technical report, working paper. University of London, Royal Holloway.

Pop-Eleches, C. and M. Urquiola (2013). Going to a better school: Effects and behavioral responses. *American Economic Review 103*(4), 1289–1324.

Popkin, S. J., V. E. Gwiasda, L. M. Olson, D. P. Rosenbaum, and L. Buron (2000). *The Hidden War: Crime and the Tragedy of Public Housing in Chicago*. Rutgers University Press.

Popkin, S. J., M. J. Rich, L. Hendey, C. Hayes, J. Parilla, and G. Galster (2012). Public housing transformation and crime: Making the case for responsible relocation. *Cityscape*, 137–160.

Sacerdote, B. et al. (2011). Peer effects in education: How might they work, how big are they and how much do we know thus far? *Handbook of the Economics of Education 3*, 249–277.

Sampson, R. J., S. W. Raudenbush, and F. Earls (1997). Neighborhoods and violent crime: A multilevel study of collective efficacy. *Science 277*(5328), 918–924.

Sandler, D. H. (2012). Is public housing a public bad? externalities of Chicago's public housing demolitions. *Job Market Paper*.

Terry, D. (1995). Chicago housing agency to be taken over by U.S. *The New York Times*.

The Economist (11 July, 1998). Chicago's public housing projects.

Thompson, M. (2006). Relocating from the distress of Chicago public housing to the difficulties of the private market: How the move threatens to push families away from opportunity. *Northwestern Journal of Law & Social Policy 1*, 267.

Todd, P. E. and K. I. Wolpin (2007). The production of cognitive achievement in children: Home, school, and racial test score gaps. *Journal of Human capital 1*(1), 91–136.

Weisleder, A. and A. Fernald (2013). Talking to children matters: Early language experience strengthens processing and builds vocabulary. *Psychological science 24*(11), 2143–2152.

Wodtke, G. T., D. J. Harding, and F. Elwert (2011). Neighborhood effects in temporal perspective the impact of long-term exposure to concentrated disadvantage on high school graduation. *American Sociological Review 76*(5), 713–736.