

The Impacts of Neighborhoods on Intergenerational Mobility Childhood Exposure Effects and County-Level Estimates

Raj Chetty and Nathaniel Hendren, Harvard University

To what extent are children's opportunities for upward economic mobility shaped by the neighborhoods in which they grow up? We study this question using data from de-identified tax records on more than five million children whose families moved across counties between 1996 and 2012. The study consists of two parts. In part one, we show that the area in which a child grows up has significant causal effects on her prospects for upward mobility. In part two, we present estimates of the causal effect of each county in the United States on a child's chances of success. Using these results, we identify the properties of high- vs. low-opportunity areas to obtain insights into policies that can increase economic opportunity.

Part 1: Do Neighborhoods Matter for Economic Mobility?


In previous work (Chetty, Hendren, Kline, and Saez 2014), we documented substantial variation in rates of upward income mobility across commuting zones (aggregations of counties analogous to metropolitan areas) in the United States. This geographic variation could be driven by two very different sources. One possibility is that neighborhoods have *causal effects* on upward mobility: that is, moving a given child to a different neighborhood would change her life outcomes. Another possibility is that the observed geographic variation is due to systematic differences in the types of people living in each area, such as differences in race or wealth. Distinguishing between these two explanations is essential to determine whether changing neighborhood environments is a good way to improve economic mobility or whether policy makers should focus on other types of interventions.

The ideal experiment to test between these two explanations and identify the causal effects of neighborhoods would be to randomly assign children to different neighborhoods and compare their incomes in adulthood. We use a quasi-experimental approximation to this experiment that relies on differences in the timing of when families move across areas.

Figure 1 illustrates our approach and results. As an example, consider a set of families who move from Cincinnati to Pittsburgh. Children who grow up in low-income families (at the 25th percentile of the national distribution) in Cincinnati from birth have an income of \$23,000 on average at age 26, while those in Pittsburgh have an income of \$28,000. Now consider the incomes of children whose families moved from Cincinnati to Pittsburgh at some point in their childhood. Figure 1 plots the fraction of the difference in income between Pittsburgh and Cincinnati that a child will on average obtain by moving at different ages during childhood. Children who were nine years old at the time of the move (the earliest age we can analyze given available data) capture 50% of this difference, leading to an income of approximately \$25,500 as adults. Children who move from Cincinnati to Pittsburgh at later ages have steadily declining incomes, relative to those who moved at younger ages. Those whose families moved after they were 23 experience no gain relative to those who stayed in Cincinnati permanently.

Figure 1 shows that every extra year a child spends in a better environment – as measured by the outcomes of children already living in that area – improves her outcomes, a pattern we term a *childhood exposure effect*. We find equal and opposite exposure effects for children whose families moved to worse areas. Further, we find analogous exposure effects for a broad range of other outcomes, including college attendance and the probability of having a teenage birth.

FIGURE 1
Effects of Moving to a Different Neighborhood on a Child’s Income in Adulthood



Notes: This figure plots the percentage gain from moving to a better area by the age at which the child moves. For example, children who move at age 9 have outcomes that are about 50% between the outcomes of children who grow up permanently in the origin and destination areas.

The key assumption underlying the analysis shown in Figure 1 – the assumption that is necessary to make it as good as the ideal randomized experiment – is that families who move from Cincinnati to Pittsburgh when their children are young are comparable to those who move when their children are older. This assumption would not hold if, for instance, families who move to better areas when their children are young are more educated or have higher wealth than families who move later.

We implement a series of tests to assess the validity of this assumption and evaluate the robustness of our quasi-experimental methodology. First, we compare siblings within the same family, and show that the difference in siblings’ outcomes is proportional to the difference in their exposure to better environments. When a family with two children moves from Cincinnati to Pittsburgh, the younger child does better than the older child on average. Second, we show that one obtains similar estimates of exposure effects when analyzing families displaced by events outside their control, such as natural disasters or local plant closures.

Finally, we exploit differences in cities’ effects across subgroups to develop sharper tests for exposure effects. For example, some areas – such as those with high crime rates – generate significantly worse outcomes for boys than girls. We find that when a family with a boy and a girl moves to such an area, their son’s outcomes worsen in proportion to the number of years he grows up there, but their daughter’s outcomes change much less. Similarly, some areas are particularly good at producing “superstars” – children who reach the top 10% of the income distribution – even though they don’t produce better outcomes on average. We find that children who move to such areas when young are themselves more likely to become superstars, but do not have higher incomes on average.

Since it is unlikely that other factors would reproduce all of these patterns, we conclude that the pattern in Figure 1 reflects the causal effect of neighborhoods on children’s long-term outcomes. This result has several important policy implications. First, it shows that the neighborhood environment during *childhood* is a key determinant of a child’s long-term success. This suggests that policy makers seeking to improve mobility should focus on improving childhood environments (e.g., by improving local schools) and not just on the strength of the local labor market or availability of jobs. Second, Figure 1 shows that the incremental benefits of exposure to a better area do not vary with a child’s age. Moving to a better area at age 9 instead of 10 produces the same incremental improvement in earnings as moving to that area at age 15 instead of 16. This finding is particularly important in light of recent discussions about early childhood interventions, as it shows that there are significant returns to improving children’s environments even at older ages.

Part 2: County-Level Estimates of Causal Exposure Effects

The first part of our study establishes that neighborhoods matter for intergenerational mobility, but does not directly identify the causal effect of any given area. In the second part of our analysis, we estimate the causal childhood exposure effect of every county in the U.S. by studying the outcomes of children who moved between counties at different ages.

To understand how we estimate these effects, consider families in the New York metro area. If we were to find that children who moved from Manhattan to Queens at a young age do better as adults, we can infer that Queens has positive causal exposure effects relative to Manhattan. Building on this logic, we use data on movers across the full set of counties in the U.S. to estimate the effect of spending an additional year of childhood in each county. We construct these estimates separately by parent income level, permitting the effects of each area to vary with the family’s income.

Table 1 shows the causal effects of the top 10 and bottom 10 counties among the 100 largest counties in the U.S for children growing up in families at the 25th percentile of the national income distribution. The estimates represent the percentage change in earnings from spending an additional year of one’s childhood in the relevant county relative to the national average.

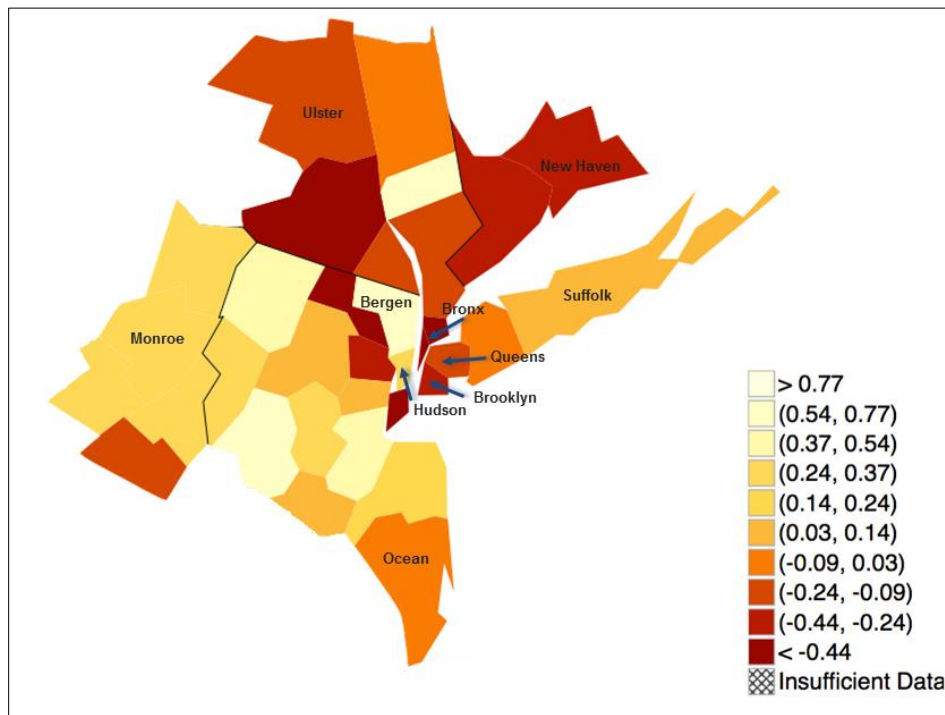
TABLE 1
Causal Exposure Effects: Top 10 and Bottom 10 Among the 100 Largest Counties
For Children with Parents at 25th Percentile of the Income Distribution

Rank		Δ Earnings (%) per year of exposure	Rank		Δ Earnings (%) per year of exposure
1	DuPage, IL	0.76%	91	Pima, AZ	-0.61%
2	Snohomish, WA	0.72%	92	Bronx, NY	-0.62%
3	Bergen, NJ	0.71%	93	Milwaukee, WI	-0.62%
4	Bucks, PA	0.66%	94	Wayne, MI	-0.63%
5	Contra Costa, CA	0.61%	95	Fresno, CA	-0.65%
6	Fairfax, VA	0.60%	96	Cook, IL	-0.67%
7	King, WA	0.57%	97	Orange, FL	-0.67%
8	Norfolk, MA	0.54%	98	Hillsborough, FL	-0.67%
9	Montgomery, MD	0.52%	99	Mecklenburg, NC	-0.69%
10	Middlesex, NJ	0.43%	100	Baltimore City, MD	-0.86%

For example, each additional year that a child spends growing up in DuPage County, IL raises her household income in adulthood by 0.76%. This implies that growing up in DuPage County from birth – i.e., having about 20 years of exposure to that environment – would raise a child’s earnings by 15% relative to the national average. In contrast, every extra year spent in the city of Baltimore *reduces* a child’s earnings by 0.86% per year of exposure, generating a total earnings penalty of approximately 17% for children who grow up there from birth.¹

There is considerable variation across counties even within metro areas. Figure 2 presents a map of the causal exposure effects for counties in the New York City area for children growing up in families at the 25th percentile. The estimates range from an earnings loss of -0.54% per year of childhood spent in Manhattan (New York County) to an earnings gain of 0.25% per year in Hudson County, NJ and 0.71% per year in Bergen County, NJ. Concretely, this implies that children in low-income families who move from Manhattan to Hudson County, NJ when they are born earn 16% more as adults on average.²

Figure 2: Causal Exposure Effects by County in the New York Combined Statistical Area
For Children with Parents at 25th Percentile of the Income Distribution



Notes: This figure shows the percentage change in household earnings caused by spending an additional year growing up in each county for children with parents at the 25th percentile of the national income distribution. Lighter colored areas are areas that generate larger earnings gains. To download statistics for your county, visit www.equality-of-opportunity.org

¹ These estimates are based on data for children born between 1980-86 and who grew up in the 1980’s and 1990’s. We find that neighborhoods’ effects generally remain stable over time, but some cities have presumably gotten better in the 2000’s, while others may have gotten worse.

² Most families at the 25th percentile of the national distribution (roughly a household income of \$30,000 for a family with teenage children) who live in Manhattan are in Harlem. Hence, the comparison is effectively between the effects of growing up in Harlem vs. an area with relatively low house prices in New Jersey.

The causal effects of counties are typically smaller in percentage terms for children who grow up in high-income families, but remain substantial. For instance, for children growing up in families in the top 1% of the income distribution, we estimate that every extra year of childhood spent in Manhattan reduces their earnings by 1.08% relative to Westchester. Areas that produce better outcomes for children in low-income families are, on average, no worse for those from high-income families. This finding suggests that the success of the poor need not come at the expense of the rich, implying that social mobility is not a “zero-sum game.”

Neighborhoods matter more for boys than girls. For example, every extra year of childhood exposure to Baltimore reduces earnings by 1.39% for low-income boys, but only 0.27% for girls. Areas with high crime rates and a large fraction of single parents generate particularly negative outcomes for boys relative to girls. There are also significant gender differences related to marriage rates. For example, Northern California generates high levels of *individual* earnings for girls, but produces lower levels of *household* income because fewer children get married in their 20s.

Our estimates of causal effects at the county and commuting zone (CZ) level are strongly correlated with the raw estimates of intergenerational mobility reported in Chetty, Hendren, Kline, and Saez (2014), but there are several significant differences. For example, children who grow up in New York City have above-average rates of upward mobility. However, the *causal* effect of growing up in New York City on upward mobility – as revealed by analyzing individuals who move into and out of New York – is negative relative to the national average. This negative effect of growing up in New York is masked when one simply studies the average outcomes of children who grow up there because families who live in New York tend to have unusually high rates of upward mobility. In particular, New York has a very large share of immigrants, and we find that immigrants have higher rates of upward mobility independent of where they live. This example shows that part of the variation in mobility across areas is driven simply by the characteristics of the people who live in those areas, which is why it is important to identify each area’s causal effect as we do in this study.

What are the properties of areas that improve upward mobility? Within a given commuting zone, we find that counties that have higher rates of upward mobility tend to have five characteristics: they have less segregation by income and race, lower levels of income inequality, better schools, lower rates of violent crime, and a larger share of two-parent households.

We also find that areas with a larger African-American population tend to have lower rates of upward mobility. These spatial differences amplify racial inequality across generations: we estimate that one-fourth of the gap in intergenerational mobility between blacks and whites can be attributed to the counties in which they live.

Lastly, we examine whether one has to pay a higher rent to live in an area with greater upward mobility. In the nation as a whole, we find weak correlations between rents and upward mobility. However, in large metro areas – especially those with high levels of segregation and sprawl – counties that offer better prospects of upward mobility are much more expensive. For example, Chicago has one area with a high level of upward mobility – DuPage County – which is also one of the most expensive counties in the area. There are, however, some “bargains” even in the largest cities: for example, Hudson County in the New York metro area and Snohomish County in the Seattle area both offer high levels of upward mobility with relatively low house prices.

The high housing prices that families often must pay to achieve better outcomes for their children may partially explain the persistence of poverty in large American cities. One approach to addressing this problem is to provide subsidized housing vouchers that enable families to move to better (e.g., lower-poverty) neighborhoods. In a companion paper (Chetty, Hendren, and Katz 2015), we show that the Moving to Opportunity experiment – which randomly assigned families subsidized

housing vouchers to move to low poverty areas – significantly improved long-term outcomes for children who moved at young ages, providing direct support for such policies.

Of course, given limits to the scalability of policies that seek to move families, one must also find methods of improving neighborhood environments in areas that currently generate low levels of mobility. Our study does not directly identify which policies are most successful in achieving this goal, but our findings provide support for policies that reduce segregation and concentrated poverty in cities (e.g., affordable housing subsidies or changes in zoning laws) as well as efforts to improve public schools.

The broader lesson of our analysis is that social mobility should be tackled at a local level by improving childhood environments. Much remains to be learned about the best ways to make such improvements. We hope the county-level data constructed here will ultimately offer new solutions to increase opportunities for disadvantaged youth throughout the United States.

Works Cited

Chetty, R., N. Hendren, P. Kline, and E. Saez. “Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States.” *Quarterly Journal of Economics* 129(4): 1553-1623, 2014.

Chetty, R. N. Hendren, and L. Katz. “The Long-Term Effects of Exposure to Better Neighborhoods: New Evidence from the Moving to Opportunity Experiment” Harvard University Working Paper, 2015.