NBER WORKING PAPER SERIES

THE IMPACTS OF NEIGHBORHOODS ON INTERGENERATIONAL MOBILITY I:
CHILDHOOD EXPOSURE EFFECTS

Raj Chetty
Nathaniel Hendren

Working Paper 23001
http://www.nber.org/papers/w23001

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2016, Revised May 2017

The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects
Raj Chetty and Nathaniel Hendren
NBER Working Paper No. 23001
December 2016, Revised May 2017
JEL No. H0,J0,R0

ABSTRACT

We show that the neighborhoods in which children grow up shape their earnings, college attendance rates, and fertility and marriage patterns by studying more than seven million families who move across commuting zones and counties in the U.S. Exploiting variation in the age of children when families move, we find that neighborhoods have significant childhood exposure effects: the outcomes of children whose families move to a better neighborhood – as measured by the outcomes of children already living there – improve linearly in proportion to the amount of time they spend growing up in that area, at a rate of approximately 4% per year of exposure. We distinguish the causal effects of neighborhoods from confounding factors by comparing the outcomes of siblings within families, studying moves triggered by displacement shocks, and exploiting sharp variation in predicted place effects across birth cohorts, genders, and quantiles to implement overidentification tests. The findings show that neighborhoods affect intergenerational mobility primarily through childhood exposure, helping reconcile conflicting results in the prior literature.

Raj Chetty
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305
and NBER
chetty@stanford.edu

Nathaniel Hendren
Harvard University
Department of Economics
Littauer Center Room 235
Cambridge, MA 02138
and NBER
nhendren@gmail.com

I Introduction

To what extent are children's economic opportunities shaped by the neighborhoods in which they grow up? Despite extensive research, the answer to this question remains debated. Observational studies have documented significant variation across neighborhoods in economic outcomes (e.g., Wilson 1987, Jencks and Mayer 1990, Massey 1993, Sampson et al. 2002, Sharkey and Faber 2014). However, experimental studies of families that move have traditionally found little evidence that neighborhoods affect economic outcomes (e.g., Katz et al. 2001, Oreopoulos 2003, Ludwig et al. 2013).

Using de-identified tax records covering the U.S. population, we present new quasi-experimental evidence on the effects of neighborhoods on intergenerational mobility that reconcile the conflicting findings of prior work and shed light on the mechanisms through which neighborhoods affect children's outcomes. Our analysis consists of two papers. In this paper, we measure the degree to which the differences in intergenerational mobility across areas in observational data are driven by causal effects of place. In the second paper (Chetty and Hendren 2017), we build on the research design developed here to construct estimates of the causal effect of growing up in each county in the United States on children's long-term outcomes and characterize the features of areas that produce good outcomes.

Our analysis is motivated by our previous work showing that children's expected incomes conditional on their parents' incomes vary substantially with the area (commuting zone or county) in which they grow up (Chetty, Hendren, Kline, and Saez 2014).[1] This geographic variation in intergenerational mobility could be driven by two very different sources. One possibility is that neighborhoods have causal effects on economic mobility: that is, moving a given child to a different neighborhood would change his or 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 demographics or wealth.

We assess the relative importance of these two explanations by asking whether children who move to areas with higher rates of upward income mobility among "permanent residents" have better outcomes themselves.[2] Since moving is an endogenous choice, simple comparisons of the

---

[1]We characterize neighborhood (or "place") effects at two geographies: counties and commuting zones (CZs), which are aggregations of counties that are similar to metro areas but cover the entire U.S., including rural areas. Naturally, the variance of place effects across these broad geographies is a lower bound for the total variance of neighborhood effects, which would include additional local variation.

[2]We define "permanent residents" as the parents who stay in the same commuting zone (or, in the county-level analysis, the same county) throughout the period we observe (1996-2012).

outcomes of children whose families move to different areas confound causal effects of place with selection effects (differences in unobservables). We address this identification problem by exploiting variation in the *timing* of children's moves across areas.[3] We compare the outcomes of children who moved to a better (or worse) area at different ages to identify the rate at which the outcomes of children who move converge to those of the permanent residents.[4] The identification assumption underlying our research design is that the selection effects (children's unobservables) associated with moving to a better versus worse area do not vary with the age of the child when the family moves. This is a strong assumption, one that could plausibly be violated for several reasons. For instance, families who move to better areas when their children are young may be more educated or invest more in their children in other ways. We present evidence supporting the validity of this identification assumption after presenting a set of baseline results.

In our baseline analysis, we focus on families with children born between 1980 and 1988 who moved once across commuting zones (CZs) between 1997 and 2010. We find that on average, spending an additional year in a CZ or county where the mean income rank of children of permanent residents is 1 percentile higher (at a given level of parental income) increases a child's income rank in adulthood by approximately 0.04 percentiles. That is, the outcomes of children who move converge to the outcomes of permanent residents in the destination at a rate of 4% per year of childhood exposure. Symmetrically, moving to an area where permanent residents have worse outcomes reduces a child's expected income by 4% per year. When analyzing children who move more than once during childhood, we find that children's outcomes vary in proportion to the amount of time they spend in each area rather than the specific ages during which they live in those areas.

Together, these results imply that neighborhoods have substantial *childhood exposure effects*: every additional year of childhood spent in a better environment improves a child's long-term outcomes. Convergence is linear with respect to age: moving to a better area at age 8 instead of 9 is associated with the same increase in income as moving to that area at age 15 instead of 16. The exposure effects persist until children are in their early twenties. Extrapolating over the duration of childhood, from age 0 to 20, the 4% annual convergence rate implies that children who move at birth to an area with one percentile higher incomes among permanent residents would pick up

---

[3]Several recent studies have used movers to identify causal effects of places on other outcomes using event-study designs, comparing individuals' outcomes before vs. after they move (e.g., Chetty et al. 2013, Finkelstein et al. 2016). We use a different research design because we naturally do not have pre-move data on income in adulthood when studying the impact of moving during childhood.

[4]Throughout the paper, we refer to areas where children have better outcomes in adulthood as "better" neighborhoods. We use this terminology without any normative connotation, as there are of course many other amenities of neighborhoods that may be relevant from a normative perspective.

about 80% of that effect themselves. We find similar childhood exposure effects for several other outcomes, including rates of college attendance, marriage, and teenage birth. We also find similar exposure effects when families moves across counties.

As noted above, the identification assumption underlying the interpretation of the 4% convergence rate as a causal exposure effect is that the potential outcomes of children who move to better vs. worse areas do not vary with the age at which they move. We use three approaches to evaluate this assumption: controlling for observable factors, isolating plausibly exogenous moves triggered by aggregate displacement shocks, and implementing a set of outcome-based placebo tests. The first two approaches are familiar techniques in the treatment effects literature, while the third exploits the multi-dimensional nature of the treatments we study to implement overidentification tests.

To implement the first approach, we begin by controlling for factors that are fixed within the family (e.g., parent education) by including family fixed effects.[5] This approach identifies exposure effects from comparisons between siblings, by asking whether the *difference* in outcomes between two siblings who move to a new area is proportional to their age difference interacted with permanent residents' outcomes in the destination. We estimate an annual exposure effect of approximately 4% per year with family fixed effects, very similar to our baseline estimate. These sibling comparisons address confounds due to factors that are fixed within families, but they do not account for *time-varying* factors, such as a change in family environment at the time of the move that directly affects children in proportion to exposure time independent of neighborhoods. We cannot observe all such time-varying factors, but we do observe two particularly important characteristics of the family environment in each year: income and marital status. Controlling flexibly for changes in income and marital status interacted with the age of the child at the time of the move has no impact on the exposure effect estimates.

The preceding results rule out confounds due to observable factors such as income, but they do not address potential confounds due to unobservable factors. In particular, whatever event endogenously induced a family to move (e.g., a wealth shock) could also have had direct effects

---

[5]The idea of using sibling comparisons to better isolate neighborhood effects dates to the seminal review by Jencks and Mayer (1990). Plotnick and Hoffman (1996) and Aaronson (1998) implement this idea using data on 742 sibling pairs from the Panel Study of Income Dynamics, but reach conflicting conclusions due to differences in sample and econometric specifications. Several studies also use sibling comparisons to identify critical periods that shape immigrants' long-term outcomes (Basu 2010, van den Berg et al. 2014). Our approach differs from these studies in that we focus on how the difference in siblings' outcomes *covaries* with the outcomes of permanent residents in the destination neighborhood, whereas the studies of immigrants estimate the mean difference in siblings' outcomes as a function of their age gap. This allows us to separate the role of neighborhood exposure from changes within the family that also generate exposure-dependent differences across siblings, such as changes in income or wealth when a family moves to a new country.

on their children's outcomes. Our second approach addresses the problem of bias associated with endogenous choice by focusing on a subset of moves that are more likely to be driven by exogenous aggregate shocks. In particular, we identify moves that occur as part of large outflows from ZIP codes, often caused by natural disasters or local plant closures. We replicate our baseline design within this subsample of displaced movers, comparing the outcomes of children who move to different destinations at different ages. We obtain similar exposure effect estimates for displaced households, mitigating concerns that our baseline estimates are biased by omitted variables correlated with a household's choice of when to move.[6]

Although the evidence from the first two approaches strongly supports the validity of the identification assumption, each of these approaches itself rests on assumptions – selection on observables and exogeneity of the displacement shocks – that could themselves potentially be violated. We therefore turn to a third approach – a set of placebo (overidentification) tests that exploit heterogeneity in permanent residents' outcomes across subgroups – that in our view provides the most compelling method of assessing the validity of the research design.

We begin by analyzing heterogeneity across birth cohorts. Although outcomes within CZs are highly persistent over time, some places improve and others decline. Exploiting this variation, we find using multivariable regressions that the outcomes of children who move to a new area converge to the outcomes of permanent residents of the destination in their *own* birth cohort, but are unrelated to those of the preceding and subsequent birth cohorts (conditional on their own birth cohort's predictions). Such cohort-specific convergence is precisely what one would expect in the causal exposure effect model, but it would be unlikely to emerge from sorting or other omitted variables because the cohort-specific effects are only realized with a long time lag, after children grow up.

We implement analogous placebo tests by exploiting variation in the *distribution* of outcomes across areas. For instance, low-income children who spend their entire childhood in Boston and San Francisco have similar outcomes on average, but children in San Francisco are more likely to end up in the upper or lower tail of the income distribution (i.e., either in the top 10% or not employed). The causal exposure effects model predicts convergence not just at the mean but across the entire distribution; in contrast, it would be unlikely that omitted variables (such as changes in parent wealth) would happen to perfectly replicate the entire distribution of outcomes in each area

---

[6]We eliminate variation due to individuals' endogenous choices of *where* to move in these specifications by instrumenting for each household's change in neighborhood quality using the average change in neighborhood quality of those who move out of the ZIP code during the years in our sample.

in proportion to exposure time. In practice, we find quantile-specific convergence: controlling for mean outcomes, children's outcomes converge to predicted outcomes in the destination across the distribution in proportion to exposure time, at a rate of about 4% per year.

Finally, we implement placebo tests exploiting heterogeneity in permanent residents' outcomes across genders. While outcomes are highly correlated across genders, there are some places where boys do worse than girls (e.g., areas with concentrated poverty) and vice versa. When a family with a daughter and a son moves to an area that is especially good for boys, their son does better than their daughter in proportion to the number of years they spend in the new area. Once again, if our findings of neighborhood exposure effects were driven by sorting or omitted variables, one would not expect to find such gender-specific convergence in outcomes unless families are fully aware of the exact gender differences in outcomes across areas and sort to neighborhoods based on these gender differences.

Our findings yield three lessons. First, place matters for intergenerational mobility: the differences we see in outcomes across neighborhoods are largely due to the causal effect of places rather than differences in the characteristics of their residents. Second, place matters largely because of differences in childhood environment, rather than the differences in labor market conditions that have received attention in previous studies of place. Moving to a better area just before entering the labor market has little impact on individual's outcomes, suggesting that place-conscious policies to promote upward mobility should focus primarily on improving the local childhood environment rather than conditions in adulthood. Third, each year of childhood exposure matters roughly equally; there is no "critical age" after which the returns to living in a better neighborhood fall sharply. This result is germane to recent policy discussions regarding early childhood interventions, as it suggests that improvements in neighborhood environments can be beneficial even in adolescence.

Our results help explain why previous experimental studies – most notably, the Moving to Opportunity (MTO) Experiment – failed to detect significant effects of moving to a better neighborhood on economic outcomes. Prior analyses of the MTO experiment focused primarily on the effects of neighborhoods on adults and older youth (e.g., Kling et al. 2007), because data on the long-term outcomes of younger children were unavailable. In a companion paper (Chetty, Hendren, and Katz 2016), we link the MTO data to tax records and show that the MTO data exhibit childhood exposure effects consistent with those identified here. In particular, Chetty, Hendren, and Katz (2016) find substantial improvements in earnings and other outcomes for children whose

families received experimental vouchers to move to low-poverty neighborhoods at young ages. In contrast, children who moved at older ages experienced no gains or slight losses.[7]

More generally, our findings imply that much of the neighborhood-level variation in economic outcomes documented in previous observational studies does in fact reflect causal effects of place, but that such effects arise through accumulated childhood exposure rather than immediate impacts on adults. The idea that exposure time to better neighborhoods may matter has been noted since at least Wilson (1987) and Jencks and Mayer (1990), and has received growing attention in observational studies in sociology (Crowder and South 2011, Wodtke et al. 2011, Wodtke 2013, Sharkey and Faber 2014). We contribute to this literature by presenting quasi-experimental estimates of exposure effects, addressing the concerns about selection and omitted variable bias that arise in observational studies (Ludwig et al. 2008). Although we find evidence of childhood exposure effects that are qualitatively consistent with the observational studies, we find no evidence of exposure effects in adulthood either in this study or our MTO study, contrary to the patterns observed in observational data (Clampet-Lundquist and Massey 1993).

Our findings are also consistent with recent studies that use other research designs – random assignment of refugees (Damm and Dustmann 2014), housing demolitions (Chyn 2016), and selection corrections using group characteristics (Altonji and Mansfield 2016) – to show that neighborhoods have causal effects on children's long-term outcomes. The present analysis complements these studies and Chetty et al.'s (2016) re-analysis of the MTO experiment in two ways. First, it sheds light on the mechanisms underlying neighborhood effects by delivering precise estimates of the magnitude and linear age pattern of childhood exposure effects. Second, it develops a scalable method to estimate neighborhood effects in all areas, even those where randomized or natural experiments are unavailable.[8]

---

[7]One important distinction between the two studies is that the analysis sample in the present quasi-experimental study consists entirely of families who moved across commuting zones, whereas the MTO experiment compares families who moved with families who did not move at all or stayed in an area similar to where they lived before. As a result, the analysis here identifies the effects of moving to better vs. worse areas *conditional* on moving to a different area, whereas the MTO analysis compares the effects of moving vs. staying in a given area. The exposure effect estimates here thus net out any fixed disruption costs of moving to a different area, whereas such costs are not netted out in the MTO experiment. This distinction may explain why Chetty, Hendren, and Katz (2016) find slightly negative effects for children who move at older ages in the MTO data, whereas we estimate positive exposure effects of moving to a better area (conditional on moving) at all ages here.

[8]Our estimates of neighborhood exposure effects are based on households who choose to move to certain areas. The effects of moving a randomly selected household to a new area may differ, since households that choose to move to a given area may be more likely to benefit from that move. The fact that exposure effects are similar within the subset of displaced households and are symmetric for moves to better and worse areas suggests such heterogeneity in exposure effects is limited, but further work is needed to understand how exposure effects vary with households' willingness to move.

This paper is organized as follows. Section II describes the data. Section III presents our empirical framework, starting with a description of differences in intergenerational mobility across areas and then specifying our estimating equations. Section IV presents baseline estimates of neighborhood exposure effects and discusses the mechanisms through which neighborhoods affect children's incomes. Section V presents tests evaluating our identification assumption. Section VI presents estimates of exposure effects for other outcomes. Section VII concludes.

II Data

We use data from federal income tax records spanning 1996-2012. The data include both income tax returns (1040 forms) and third-party information returns (e.g., W-2 forms), which contain information on the earnings of those who do not file tax returns. Because our empirical analysis is designed to determine how much of the geographic variation in intergenerational mobility documented by Chetty et al. (2014) is due to causal effects of place, we use an analysis sample that is essentially identical to the "extended sample" used in Chetty et al. (2014). Online Appendix A of Chetty et al. (2014) gives a detailed description of how we construct the analysis sample starting from the raw population data. Here, we briefly summarize the key variable and sample definitions, following Section III of Chetty et al. (2014).

II.A Sample Definitions

Our base dataset of children consists of all individuals who (1) have a valid Social Security Number or Individual Taxpayer Identification Number, (2) were born between 1980-1988, and (3) are U.S. citizens as of 2013.[9] We impose the citizenship requirement to exclude individuals who are likely to have immigrated to the U.S. as adults, for whom we cannot measure parent income. We cannot directly restrict the sample to individuals born in the U.S. because the database only records current citizenship status.

We identify the parents of a child as the first tax filers (between 1996-2012) who claim the child as a child dependent and were between the ages of 15 and 40 when the child was born.[10] If the child is first claimed by a single filer, the child is defined as having a single parent. For simplicity, we assign each child a parent (or parents) permanently using this algorithm, regardless

---

[9]For selected outcomes that can be measured at earlier ages, such as teenage labor force participation rates, we extend the sample to include more recent birth cohorts, up to 1996.

[10]We impose the 15-40 age restriction to limit links to grandparents or other guardians who might claim a child as a dependent.

of any subsequent changes in parents' marital status or dependent claiming.

If parents never file a tax return, we do not link them to their child. Although some low-income individuals do not file tax returns in a given year, almost all parents file a tax return at some point between 1996 and 2012 to obtain a tax refund on their withheld taxes and the Earned Income Tax Credit (Cilke 1998). As a result, approximately 94% of the children in the 1980-1988 birth cohorts are claimed as a dependent at some point between 1996 and 2012. The fraction of children linked to parents drops sharply prior to the 1980 birth cohort because our data begins in 1996 and many children begin to the leave the household starting at age 17 (Chetty et al. 2014, Online Appendix Table I). This is why we limit our analysis to children born during or after 1980.

Our full analysis sample includes all children in the base dataset who are born in the 1980-88 birth cohorts for whom we are able to identify parents and whose mean parent income between 1996-2000 is strictly positive.[11] We divide the full sample into two parts: *permanent residents* (or stayers) and *movers*. We define the permanent residents of each commuting zone (CZ) c as the subset of parents who reside in a single CZ c in all years of our sample, 1996-2012. The movers sample consists of individuals in the full sample who are not permanent residents.

In our baseline analysis, we focus on the subset of individuals who live in CZs with populations in the 2000 Census above 250,000 (excluding 19.6% of the observations) to ensure that we have adequately large samples to estimate permanent residents' outcomes (the key independent variables in our analysis) precisely. There are approximately 24.6 million children in the baseline analysis sample for whom we observe outcomes at age 24 or later, of whom 19.5 million are children of permanent residents.

II.B Variable Definitions and Summary Statistics

In this section, we define the key variables we use in our analysis. We measure all monetary variables in 2012 dollars, adjusting for inflation using the headline consumer price index (CPI-U). We begin by defining the two key variables we measure for parents: income and location.

Parent Income. Our primary measure of parent income is total pre-tax income at the household level, which we label parent family (or household) income. In years where a parent files a tax return, we define family income as Adjusted Gross Income (as reported on the 1040 tax return) plus tax-exempt interest income and the non-taxable portion of Social Security and Disability benefits.

---

[11]We limit the sample to parents with positive income (excluding 1.5% of children) because parents who file a tax return – as is required to link them to a child – yet have zero income are unlikely to be representative of individuals with zero income and those with negative income typically have large capital losses, which are a proxy for having significant wealth.

In years where a parent does not file a tax return, we define family income as the sum of wage earnings (reported on form W-2), unemployment benefits (reported on form 1099-G), and gross social security and disability benefits (reported on form SSA-1099) for both parents.[12] In years where parents have no tax return and no information returns, family income is coded as zero.[13] Income is measured prior to the deduction of income taxes and employee-level payroll taxes, and excludes non-taxable cash transfers and in-kind benefits.

In our baseline analysis, we average parents' family income over the five years from 1996 to 2000 to obtain a proxy for parent lifetime income that is less affected by transitory fluctuations (Solon 1992). We use the earliest years in our sample to best reflect the economic resources of parents while the children in our sample are growing up.[14] Because we measure parent income in a fixed set of years, the age of the child when parent income is measured varies across birth cohorts. We account for this variation by conditioning on the child's birth cohort throughout our analysis.

Parent Location. In each year, parents are assigned ZIP codes of residence based on the ZIP code from which they filed their tax return. If the parent does not file in a given year, we search W-2 forms for a payee ZIP code in that year. Non-filers with no information returns are assigned missing ZIP codes. For children whose parents were married when they were first claimed as dependents, we always track the mother's location if marital status changes. We map parents' ZIP codes to counties and CZs using the crosswalks and methods described in Chetty et al. (2014, Online Appendix A).

Next, we define the outcomes that we analyze for children.

Income. We define child family income in exactly the same way as parent family income. We measure children's annual incomes at ages ranging from 24-30 and define the child's household based on his or her marital status at the point at which income is measured. For some robustness

---

[12]The database does not record W-2's and other information returns prior to 1999, so non-filer's income is coded as 0 prior to 1999. Assigning non-filing parents 0 income has little impact on our estimates because only 3.1% of parents in the full analysis sample do not file in each year prior to 1999 and most non-filers have very low W-2 income (Chetty et al. 2014). For instance, in 2000, the median W-2 income among non-filers in our baseline analysis sample was $0.

[13]Importantly, these observations are true zeros rather than missing data. Because the database covers all tax records, we know that these individuals have 0 taxable income.

[14]Formally, we define mean family income as the mother's family income plus the father's family income in each year from 1996 to 2000 divided by 10 (or divided by 5 if we only identify a single parent). For parents who do not change marital status, this is simply mean family income over the 5 year period. For parents who are married initially and then divorce, this measure tracks the mean family incomes of the two divorced parents over time. For parents who are single initially and then get married, this measure tracks individual income prior to marriage and total family income (including the new spouse's income) after marriage. These household measures of income increase with marriage and naturally do not account for cohabitation; to ensure that these features do not generate bias, we assess the robustness of our results to using individual measures of income.

checks, we analyze individual income, defined as the sum of individual W-2 wage earnings, UI benefits, SSDI payments, and half of household self-employment income.

Employment. We define an indicator for whether the child is employed at a given age based on whether he has a W-2 form filed on his behalf at that age. We measure employment rates starting at age 16 to analyze teenage labor force participation.

College Attendance. We define college attendance as an indicator for having one or more 1098-T forms filed on one's behalf when the individual is aged 18-23. Title IV institutions – all colleges and universities as well as vocational schools and other post-secondary institutions eligible for federal student aid – are required to file 1098-T forms that report tuition payments or scholarships received for every student. The 1098-T forms are available from 1999-2012 and are filed directly by colleges independent of whether an individual files a tax return. Comparisons to other data sources indicate that 1098-T forms capture more than 95% of college enrollment in the U.S. (Chetty et al. 2017).

Teenage Birth. For women, we define an indicator for teenage birth if they are listed as a parent on a birth certificate when they are between the ages of 13 and 19, using data from the Social Security Administration's DM-2 database.[15]

Marriage. We define an indicator for whether the child is married at at a given age based on the marital status listed on 1040 forms for tax filers. We code non-filers as single because linked CPS-IRS data show that the vast majority of non-filers below the age of 62 are single (Cilke 1998).

Summary Statistics. Table I reports summary statistics for our analysis sample and various subgroups used in our CZ-level analysis; Online Appendix Table I presents analogous statistics for the sample used in our county-level analysis. The first panel reports statistics for permanent residents in our full analysis sample who live in CZs with more than 250,000 people. The second panel considers the 4.4 million children who moved between CZs with more than 250,000 people (excluding children whose parents moved more than three times between 1996-2012, who account for 3% of the observations). The third panel focuses on our primary analysis sample of *one-time movers*: children whose parents moved exactly once across CZs between 1996-2012, are observed in the destination CZ for at least two years, and moved at least 100 miles.[16] There are 1.6 million

---

[15]The total count of births in the SSA DM-2 database closely matches vital statistics counts from the Center for Disease Control prior to 2008; however, the DM-2 database contains approximately 10% fewer births between 2008-2012. Using an alternative measure of teenage birth that does not suffer from this missing data problem – in which we define a woman as having a teen birth if she ever claims a dependent who was born while she was between the ages of 13 and 19 – yields very similar results (not reported). We do not use the dependent-claiming definition as our primary measure of teenage birth because it only covers children who are claimed as dependents by their mothers.

[16]We impose these restrictions to eliminate moves across CZ borders that do not reflect a true change of location. We measure the distance of moves as the distance between the centroids of the origin and destination ZIPs (obtained from www.boutell.com/zipcodes). We show the robustness of our results to using alternative cutoffs for minimum

children in this one-time movers sample.

While our analysis does not require movers to be comparable to permanent residents, we find that movers and permanent residents have similar characteristics. Median parent family income is $61,300 for permanent residents, compared to $58,700 for one-time movers. Children of permanent residents have a median family income of $35,600 when they are 30 years old, compared with $33,600 for one-time movers. Roughly 70% of children of permanent residents and one-time movers are enrolled in a college at some point between the ages of 18 and 23. 11% of daughters of permanent residents and one-time movers have a teenage birth.

III Empirical Framework

In this section, we first present a descriptive characterization of intergenerational mobility for children who grow up in different areas in the U.S. We then formally define our estimands of interest – childhood exposure effects – and describe the research design we use to identify these exposure effects in observational data.

III.A Geographical Variation in Outcomes of Permanent Residents

We conceptualize "neighborhood" effects as the sum of place effects at different geographies, ranging from broad to narrow: commuting zones, counties, ZIP codes, and census tracts. In the main text of this paper, we focus on variation across commuting zones (CZs). CZs are aggregations of counties based on commuting patterns in the 1990 Census constructed by Tolbert and Sizer (1996). There are 741 CZs in the U.S.; on average, each CZ contains 4 counties and has a population of 380,000. We replicate the results reported in the main text at the county level in Online Appendix C. We focus on variation across relatively broad geographic units to maximize statistical precision because some of our research designs require large sample sizes to discern fine variation in permanent residents' outcomes across subsamples.

We characterize the outcomes of children who spent their entire childhoods in a single CZ by focusing on children of "permanent residents" – parents who stay in the same CZ between 1996 and 2012.[17] Importantly, our definition of permanent residents conditions on *parents'* locations,

---

population size and move distances in Online Appendix A.

[17]Because our data start in 1996, we cannot measure parents' location over their children's entire childhood. For the 1980 birth cohort, we measure parents' location between the ages of 16 and 32; for the 1991 birth cohort, we measure parents' location between 5 and 21. This creates measurement error in children's childhood environment that is larger in earlier birth cohorts. Fortunately, we find that our results do not vary significantly across birth cohorts, and in particular remain similar for the most recent birth cohorts. The reason such measurement error turns out to be modest empirically is that most families who stay in a given area for several years tend not to have moved

not children's locations in adulthood. The CZ where a child grew up may differ from the CZ where he lives when we measure his earnings in adulthood.

Since places can have different effects across parent income levels and over time, we characterize children's mean outcomes conditional on their parents' income separately for each CZ c and birth cohort s. Chetty et al. (2014) show that measuring incomes using percentile *ranks* (rather than dollar levels) has significant statistical advantages. Following their approach, we define child i's percentile rank y_i based on his position in the *national* distribution of incomes relative to all others in his birth cohort. Similarly, we measure the percentile rank of the parents of child i, $p(i)$, based on their positions in the national distribution of parental income for child i's birth cohort.

Let $\bar{y}_{pcs}$ denote the mean rank of children with parents at percentile p of the income distribution in CZ c in birth cohort s. Figure I illustrates how we estimate $\bar{y}_{pcs}$ for children born in 1980 to parents who are permanent residents of the Chicago CZ. This figure plots the mean child rank at age 30 within each percentile bin of the parent income distribution, $E[y_i | p(i) = p]$. The conditional expectation of a child's rank given his parents' rank is almost perfectly linear, a property that is robust across CZs (Chetty et al. (2014), Online Appendix Figure IV). Exploiting linearity, we parsimoniously summarize the relationship between children's mean income ranks and their parents' ranks by regressing children's ranks on their parents' ranks in each CZ c and birth cohort s:

$$y_i = \alpha_{cs} + \psi_{cs} p_i + \varepsilon_i. \tag{1}$$

We then estimate $\bar{y}_{pcs}$ using the fitted values from this regression:

$$\bar{y}_{pcs} = \hat{\alpha}_{cs} + \hat{\psi}_{cs} p. \tag{2}$$

For example, in Chicago, $\bar{y}_{25,c,1980} = 40.1$ for children growing up at the 25th percentile of the national income distribution and $\bar{y}_{75,c,1980} = 59.3$ for children growing up at the 75th percentile.

Figure II maps children's mean income ranks at age 30 by CZ for children with parents at the 25th percentile (Panel A) and 75th percentile (Panel B); analogous maps at the county level are presented in Online Appendix Figure I. We construct these maps by dividing CZs into deciles based on their estimated value of $\bar{y}_{25,c,s}$ and $\bar{y}_{75,c,s}$, with lighter colors representing deciles with higher mean outcomes. As documented by Chetty et al. (2014), children's outcomes vary substantially across CZs, especially for children from low-income families. Chetty et al. (2014, Section V.C) discuss the spatial patterns in these maps in detail. Here, we focus on investigating whether the

---

in the past either. For example, among families who stayed in the same CZ c when their children were between ages 16-24, 81.5% of them lived in the same CZ when their children were age 8.

variation in these maps is driven by causal effects of place or heterogeneity in the types of people living in different places.

III.B Definition of Exposure Effects

Our objective is to determine how much a child's potential outcomes would improve on average if he were to grow up in an area where the permanent residents' outcomes are 1 percentile point higher. We answer this question by studying children who move across areas to estimate childhood *exposure effects*. We define the exposure effect at age m as the impact of spending year m of one's childhood in an area where permanent residents' outcomes are 1 percentile point higher.

Formally, consider a hypothetical experiment in which we randomly assign children to new neighborhoods d starting at age m for the rest of their childhood. The best linear predictor of children's outcomes y_i in the experimental sample, based on the permanent residents' outcomes in CZ d $(\bar{y}_{pds})$, can be written as

$$y_i = \alpha + \beta_m \bar{y}_{pds} + \theta_i, \tag{3}$$

where the error term θ_i captures family inputs and other determinants of children's outcomes. Since random assignment guarantees that θ_i is orthogonal to $\bar{y}_{pds}$, estimating (3) using OLS yields a coefficient β_m that represents the mean impact of spending year m of one's childhood onward in an area where permanents residents have 1 percentile better outcomes. We define the exposure effect at age m as $\gamma_m = \beta_m - \beta_{m+1}$.[18] Note that if income y_i is measured at age T, $\beta_m = 0$ for $m > T$, as moving after the outcome is measured cannot have a causal effect on the outcome.

Estimating the exposure effects $\{\gamma_m\}$ is of interest for several reasons. First, a positive effect (at any age) allows us to reject the null hypothesis that neighborhoods do not matter, a null of interest given prior experimental evidence. Second, $\{\gamma_m\}$ is informative about the ages at which neighborhood environments matter most for children's outcomes. Third, the magnitude of $\beta_0 = \sum_{t=0}^{T} \gamma_m$ – the impact of assigning children to better neighborhood from birth – provides an estimate of the degree to which the differences in children's outcomes across areas are due to place effects vs. selection. If place effects are homogeneous across children within birth cohorts and parent income groups, $\beta_0 = 0$ would imply that all of the variation across areas is due to selection, while $\beta_0 = 1$ would imply that all of the variation reflects causal effects of place. More generally, the magnitude of β_0 tells us how much of the differences across areas in Figure II rub off on children

---

[18]For simplicity, we do not allow β_m to vary across parent income percentiles p in our baseline analysis, thereby estimating the average exposure effect across families with different incomes. We estimate (3) separately by parental income level in Online Appendix Table III.

who are randomly assigned to live there from birth.

Although β_0 sheds light on the causal effect of places *on average*, we caution that it does not identify the causal effect of any given area on a child's potential outcomes. The causal effect of growing up in a given CZ c will generally differ from the mean predicted impact based on permanent residents' outcomes $(\alpha + \beta_0 \bar{y}_{pds})$ because selection and causal effects will vary across areas. We build on the methodology developed in this paper to estimate the causal effect of each CZ and county in the second paper in this series (Chetty and Hendren 2017).

III.C Estimating Exposure Effects in Observational Data

We estimate exposure effects by studying families who move across CZs with children of different ages in observational data. In observational data, the error term θ_i in (3) will generally be correlated with $\bar{y}_{pds}$. For instance, parents who move to a good area may have latent ability or wealth that produces better child outcomes. Estimating (3) in an observational sample of families who move exactly once yields a regression coefficient

$$b_m = \beta_m + \delta_m,$$

where $\delta_m = \frac{cov(\theta_i, \bar{y}_{pds})}{var(\bar{y}_{pds})}$ is a standard selection effect that measures the extent to which parental inputs and other determinants of children's outcomes for movers covary with permanent residents' outcomes. Fortunately, the identification of exposure effects does not require that *where* people move is orthogonal to child's potential outcomes. Instead, it requires that *when* people move to better versus worse areas is orthogonal to children's potential outcomes, as formalized in the following assumption.

Assumption 1. Selection effects do not vary with the child's age at move: $\delta_m = \delta$ for all m.

Assumption 1 allows for the possibility that the families who move to better areas may differ from those who move to worse areas, but requires that the extent of such selection does not vary with the age of the child when the parent moves. Under this assumption, we obtain consistent estimates of exposure effects $\gamma_m = \beta_m - \beta_{m+1} = b_m - b_{m+1}$ from (3) even in observational data because the selection effect δ cancels out when estimating the exposure effect. We can estimate the selection effect δ itself by examining the outcomes of children whose families move *after* their income is measured, e.g. at age $a \geq 30$ if income is measured at age $T = 30$. Because moves at age $a > T$ cannot have a causal effect on children's outcomes at age 30, $b_m = \delta$ for $m > T$ under Assumption 1. Using the estimated selection effect, we can identify the causal effect of moving to

a better area at age m as $\beta_m = b_m - b_{T+1}$ and thereby identify β_0, the causal effect of growing up from birth in an area with 1 percentile better outcomes.

Of course, Assumption 1 is a strong restriction that may not hold in practice. We therefore evaluate its validity in detail after presenting a set of baseline estimates in the next section.

IV   Baseline Estimates of Childhood Exposure Effects

This section presents our baseline estimates of exposure effects $\{\gamma_m\}$. We begin by presenting a set of semi-parametric estimates of $\{\gamma_m\}$ using specifications that condition on origin fixed effects and correspond most closely to the hypothetical experiment described in Section III.B. We then present estimates from parametric models that show how movers' outcomes can be parsimoniously modeled as a linear combination of the outcomes of permanent residents in origins and destination. Finally, we present a set of supplementary results that shed light on the mechanisms through which neighborhoods affect children's outcomes.

In our baseline analysis, we focus on children whose parents moved across CZs exactly once between 1996 and 2012 and are observed in the destination CZ for at least two years. We also restrict attention to families who moved at least 100 miles to exclude moves across CZ borders that do not reflect a true change of neighborhood and limit the sample to CZs with populations above 250,000 to minimize sampling error in the estimates of permanent residents' outcomes $\bar{y}_{pds}$. We show that the findings are robust to alternative cutoffs for population size and move distance in Online Appendix A, and present estimates that include families who move more than once in Online Appendix B.

In prior work (Chetty et al. 2014), we found that the intergenerational correlation between parents' and children's incomes stabilizes when children turn 30, as college graduates experience steeper wage growth in their 20s (Haider and Solon 2006). Measuring income at age 30 limits us to estimating exposure effects only after age 15 given the time span of our dataset.[19] Fortunately, measuring income at earlier ages (from 24-30) turns out not to affect the exposure effect estimates. The reason is that our estimates of b_m are identified by comparing the incomes of children who move with the incomes of permanent residents in the destination at the *same age*. The incomes of permanent residents serve as goalposts that allow us to measure the degree of convergence in incomes at any age, even before we observe children's permanent income. For example, if a given

---

[19]The most recent birth cohort for which we observe income at age 30 (in 2012) is the 1982 cohort; since our data begin in 1996, we cannot observe moves before age 15.

area c sends many children to college and therefore generates relatively low incomes at age 24, we will obtain a *higher* estimate of b_m if a child who moves to area c has a *low* level of income at age 24. We therefore measure income at age 24 in our baseline specifications to estimate exposure effects for the broadest age range.[20]

IV.A Semi-Parametric Estimates

To begin, consider the set of children whose families moved when they were exactly m years old. We analyze how these children's incomes in adulthood are related to those of permanent residents in their destination CZ using the following linear regression:

$$y_i = \alpha_{qos} + b_m \Delta_{odps} + \varepsilon_{1i}, \tag{4}$$

where y_i denotes the child's income rank at age 24, α_{qos} is a fixed effect for the origin CZ o by parent income decile q by birth cohort s and $\Delta_{odps} = \bar{y}_{pds} - \bar{y}_{pos}$ is the difference in predicted income rank (at age 24) of permanent residents in the destination versus origin for the relevant parent income rank p and birth cohort s. Equation (4) can be interpreted as an observational analog of the specification in (3) that we would ideally estimate in experimental data.[21]

Figure III presents a non-parametric binned scatter plot corresponding to the regression in (4) for children who move at age $m = 13$. To construct the figure, we first demean both y_i and Δ_{odps} within the parent decile (q) by origin (o) by birth cohort (s) cells in the sample of movers at age $m = 13$ to construct residuals: $y_i^r = y_i - E[y_i|q, o, s]$ and $\Delta_{odps}^r = \Delta_{odps} - E[\Delta_{odps}|q, o, s]$. We then divide the Δ_{odps}^r residuals into twenty equal-size groups (ventiles) and plot the mean value of y_i^r vs. the mean value of Δ_{odps}^r in each bin.

Figure III shows that children who move to areas where children of permanent residents earn more at age 24 themselves earn more when they are 24. The relationship between y_i and Δ_{odps} is linear. The regression coefficient of $b_{13} = 0.615$, estimated in the microdata using (4), implies that a 1 percentile increase in $\bar{y}_{pds}$ is associated with a 0.615 percentile increase in y_i for the children who move at age 13.

---

[20]We show below that we obtain similar estimates when measuring income at later ages (from 26-30) over the overlapping range of ages at which children move. We do not study income before age 24 because many children are enrolled in college at earlier ages and because we find that exposure effects persist until age 23 when income is measured at any point between 24 and 30. We study college attendance as a separate outcome in Section VI.

[21]We use parent income deciles rather than percentiles to define the fixed effects α_{qos} to simplify computation; using finer bins to measure parent income groups has little effect on the estimates. Conditional on parent percentile, origin, and birth cohort, the variation in Δ_{odps} is entirely driven by variation in the destination outcomes ($\bar{y}_{pds}$). Hence, b_m is identified from variation in $\bar{y}_{pds}$, as in (3), up to the approximation error from using parent deciles instead of exact percentiles.

Building on this approach, we estimate analogous regression coefficients b_m for children whose parents move at each age m from 9 to 30. We estimate $\{b_m\}$ using the following specification:

$$y_i = \alpha_{qosm} + \sum_{m=9}^{30} b_m I(m_i = m)\Delta_{odps} + \sum_{s=1980}^{1987} \kappa_s I(s_i = s)\Delta_{odps} + \varepsilon_{2i}, \tag{5}$$

where α_{qosm} is an origin CZ by parent income decile by birth cohort by age at move fixed effect and $I(x_i = x)$ is an indicator function that is 1 when $x_i = x$ and 0 otherwise. This specification generalizes (4) by fully interacting the age at move m with the independent variables in (4). In addition, we permit the effects of Δ_{odps} to vary across birth cohorts (captured by the κ_s coefficients) because our ability to measure parent's locations during childhood varies across birth cohorts. We observe children's locations starting only at age 16 for the 1980 cohort, but starting at age 8 for the 1988 cohort. This leads to greater measurement error in Δ_{odps} for earlier birth cohorts, which can confound our estimates of b_m since the distribution of ages at move is unbalanced across cohorts (see Online Appendix A for further details). By including cohort interactions, we identify $\{b_m\}$ from within-cohort variation in ages at move.[22]

Figure IVa plots estimates of b_m from (5). The estimates exhibit two key patterns: *selection effects* after age 24 and *exposure effects* before age 24. First, the fact that $b_m > 0$ for $m > 24$ is direct evidence of selection effects ($\delta_m > 0$), as moves after age 24 cannot have a causal effect on income at 24. Families who move to better areas have children with better unobservable attributes. The degree of selection δ_m does not vary significantly with m above age 24: regressing b_m on m for $m \geq 24$ yields a statistically insignificant slope of 0.001 (s.e. $= 0.011$). This result is consistent with Assumption 1, which requires that selection does not vary with the child's age at move. The mean value of δ_m for $m \geq 24$ is $\delta = 0.126$, i.e. families who move to an area where permanent residents have 1 percentile better outcomes have 0.126 percentile better outcomes themselves purely due to selection effects. Assumption 1 allows us to extrapolate the selection effect of $\delta = 0.126$ back to earlier ages $m < 24$, as shown by the dashed horizontal line in Figure IVa, and thereby identify causal exposure effects at earlier ages.

This leads to the second key pattern in Figure IVa, which is that the estimates of b_m decline steadily with the age at move m for $m < 24$. Under Assumption 1, this declining pattern constitutes evidence of an exposure effect, i.e. that moving to a better area earlier in childhood generates larger long-term gains.[23] The linearity of the relationship between b_m and the age at move m in Figure

---

[22]To avoid collinearity, we omit the most recent cohort interaction with Δ_{odps} (the 1988 cohort when income is measured at age 24). We show below that these cohort interactions have little impact on the estimates obtained from (5), but play a larger role in specifications that include family fixed effects.

[23]This declining pattern could also potentially be generated by critical age effects rather than effects that operate

IVa below age 24 implies that the exposure effect $\gamma_m = b_{m+1} - b_m$ is approximately constant with respect to age at move m. Regressing $\hat{b}_m$ on m for $m < 24$, we estimate an average annual exposure effect of $\gamma = 0.044$ (s.e. $= 0.003$). That is, the outcomes of children who move converge to the outcomes of permanent residents of the destination area at a rate of 4.4% per year of exposure until age 23.[24]

Because some children do not move with their parents, the estimates of b_m in (5) should be interpreted as intent-to-treat (ITT) estimates, in the sense that they capture the causal effect of moving (plus the selection effect) for children whose *parents* moved at age m. We can obtain treatment-on-the-treated (TOT) estimates for children who move with their parents by inflating the ITT estimates by the fraction of children who moved with their parents at each age m.[25] In Online Appendix Figure III, we show that the TOT estimate of the exposure effect is $\gamma^{TOT} = 0.040$. This estimate is very similar to our baseline estimate because virtually all children move with their parents below age 18 and roughly 60% of children move with their parents between ages 18-23. Because the treatment effects converge toward zero as the age at move approaches 23, inflating the coefficients by 1/0.6 at later ages has little impact on exposure effect estimates.

IV.B   Parametric Estimates

Equation (5) includes more than 200,000 fixed effects (α_{qosm}), making it difficult to estimate in smaller samples and introduce additional controls such as family fixed effects. As a more tractable alternative, we estimate a model in which we control parametrically for the two key factors captured by the α_{qosm} fixed effects: (1) the quality of the origin location, which we model by interacting the predicted outcomes for permanent residents in the origin at parent income percentile p_i with birth cohort fixed effects and (2) disruption costs of moving that may vary with the age at move and parent income, which we model using age at move fixed effects linearly interacted with parent

---

in proportion to exposure time. We present evidence in Section IV.C below supporting the interpretation of these results as exposure effects.

[24]Figure IVa is identified from variation in movers' destinations holding their origin fixed. An alternative approach is to exploit variation in origins, holding destinations fixed. Online Appendix Figure II presents estimates of b_m identified from variation in origins by replacing the origin (α_{qosm}) fixed effects in (5) with destination (α_{qdsm}) fixed effects. The resulting estimates yield a qualitative pattern that is the mirror image of those in Figure IVa: the later the family moves to the destination, the more the child's outcomes match the permanent residents in the origin, up to age 23. The estimated exposure effect of 0.030 is smaller than the estimates above because we measure children's origins with greater error than destinations, as our location data is left-censored. This is why we focus on variation in destinations in most of our specifications.

[25]We identify children who move with their parents based on whether they ever file a tax return, receive a W-2 form, or attend a college in the destination CZ.

income percentile p_i. This leads to the following regression specification:

$$
\begin{aligned}
y_i \;=\; & \sum_{s=1980}^{1988} I(s_i = s)(\alpha_s^1 + \alpha_s^2 \bar{y}_{pos}) + \sum_{m=9}^{30} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p_i) \\
& + \sum_{m=9}^{30} b_m I(m_i = m)\Delta_{odps} + \sum_{s=1980}^{1987} \kappa_s^d I(s_i = s)\Delta_{odps} + \varepsilon_{3i},
\end{aligned}
\tag{6}
$$

The first two terms of this specification control for origin quality and disruption effects. The third term represents the exposure effects of interest and the fourth consists of cohort interactions with Δ_{odps} to control for differential measurement error across cohorts, as in (5).[26]

Figure IVb plots the coefficients $\{b_m\}$ obtained from estimating (6). The coefficients are very similar to those obtained from the more flexible specification used to construct Figure IVa. Regressing the b_m coefficients on m for $m \leq 23$, we obtain an average annual exposure effect estimate of $\gamma = 0.038$ (s.e. $= 0.002$). This estimate is similar to that obtained from the fixed effects specification because controlling for the quality of the origin using the permanent residents' outcomes is adequate to account for differences in origin quality. Put differently, movers' outcomes can be modeled as a weighted average of the outcomes of permanents residents in the origin and destination, with weights reflecting the amount of childhood spent in the two places.

When measuring income at age 24, we cannot determine whether b_m stabilizes after age 24 because moving after age 24 has no causal effect on income or because we measure income at that point. In Online Appendix Figure IV, we replicate the analysis measuring income at ages 26, 28, and 30 in addition to age 24. All of these series display very similar patterns of exposure effects in the overlapping age ranges, showing that our estimates of b_m are insensitive to the age at which we measure children's incomes in adulthood. In particular, all four series decline linearly at a rate of approximately $\gamma = 0.04$ until age 23 and are flat thereafter. These results imply that neighborhood exposure before age 23 is what matters for income in subsequent years.

The kink at age 23 motivates the baseline regression specification that we use for the rest of our analysis. We parameterize both the exposure and selection effects shown in Figure IV linearly, replacing the non-parametric $\sum_{m=9}^{30} b_m I(m_i = m)\Delta_{odps}$ term in (6) with two separate lines above

---

[26]In addition to having much fewer fixed effects, this specification uses variation in both the quality of the origin ($\bar{y}_{pos}$) and the destination ($\bar{y}_{pds}$) to identify $\{b_m\}$. In contrast, the semi-parametric model in (5) is identified purely from variation in destinations because it includes origin fixed effects. Estimating a parametric model that identifies $\{b_m\}$ from variation in destinations by controlling for outcomes of permanent residents in the origin interacted with the age of the child at the time of the move ($\sum_{m=9}^{30} b_m I(m_i = m)y_{pos}$) yields very similar estimates.

and below age 23:

$$y_i = \sum_{s=1980}^{1988} I(s_i = s)(\alpha_s^1 + \alpha_s^2 \bar{y}_{pos}) + \sum_{m=9}^{30} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p_i) + \sum_{s=1980}^{1987} \kappa_s^d I(s_i = s)\Delta_{odps} \quad (7)$$
$$+ I(m_i \leq 23)\left(b_0 + (23 - m_i)\gamma\right)\Delta_{odps} + I(m_i > 23)(\delta + (23 - m_i)\delta')\Delta_{odps} + \varepsilon_{3i},$$

Estimating this specification directly in the microdata yields an average annual exposure effect $\gamma = 0.040$ (s.e. = 0.002), as shown in Column 1 of Table II.[27]

The estimates of γ are robust to alternative specifications and sample definitions. Columns 2 and 3 of Table II show that estimating γ using data only up to age 18 or 23 – i.e., excluding the data at older ages that identifies the selection effect in (7) – yields similar estimates of γ. Column 4 shows that excluding the cohort interactions, $\sum_{s=1980}^{1988} I(s_i = s)\alpha_s^2 \bar{y}_{pos}$ and $\sum_{s=1980}^{1987} \kappa_s^d I(s_i = s)\Delta_{odps}$, in (7) does not affect the estimate of γ significantly. Column 5 shows that we obtain an estimate of $\gamma = 0.041$ (s.e. = 0.002) when we measure movers' income ranks y_i and permanent residents' income ranks $\bar{y}_{pcs}$ at the individual rather than household level.

We replicate the analysis in Table II at the county level in Online Appendix Table V. We obtain slightly smaller exposure effect estimates of $\gamma \simeq 0.035$ at the county level, indicating that selection effects account for a larger fraction of the variance in permanent residents' outcomes at smaller geographies. This is intuitive, as families are more likely to sort geographically (e.g., to better school districts) within rather than across labor markets.

IV.C Mechanisms

In this subsection, we present a set of additional specifications that shed light on the mechanisms through which neighborhoods affect children's outcomes.

We begin by distinguishing the role of childhood environment from differences caused by variation in labor market conditions or local costs of living across areas. In Column 6 of Table II, we add fixed effects for the CZ in which the child lives at age 24 (when income is measured) to the baseline model. This specification compares the outcomes of children who live in the same labor market in adulthood but grew up in different neighborhoods. We obtain an annual exposure effect

---

[27]This coefficient differs slightly from the coefficient of $\gamma = 0.038$ that we obtain when regressing the coefficients b_m on m in Figure IVb because estimating the regression in the microdata puts different weights on each age (as we have more data at older ages), while estimating the regression using the b_m coefficients puts equal weight on all ages. The standard error in this and all subsequent specifications is also obtained from the regression in the microdata. To simplify computation, we report conventional (unclustered) standard errors. Clustering standard errors by family to account for correlated outcomes across siblings does not affect the standard errors appreciably. In addition, regressing the estimates of b_m on m in Figure IVb – which is analogous to clustering the standard errors by the age at move – also yields a s.e. of 0.002, showing that our inferences are not sensitive to the way in which standard errors are computed.

of $\gamma = 0.031$ in this specification, indicating that the majority of the exposure effect in our baseline specification is driven by differences in exposure to a better *childhood* environment, holding fixed labor market conditions.[28] This conclusion is consistent with the fact that moving to an area where permanent residents have higher income just before entering the labor market (e.g., in one's early 20s) has little effect on income, as shown in Figure IV.

Next, we examine heterogeneity in exposure effects across subsamples (Online Appendix Table III). Standard models of learning predict that moving to a better area will improve outcomes but moving to a worse area will not. In practice, the exposure effect for negative moves is larger than for positive moves: $\gamma = 0.030$ for moves to better CZs ($\Delta_{odps} > 0$), while $\gamma = 0.040$ for moves to worse CZs ($\Delta_{odps} < 0$).[29] Spending part of one's childhood in a good neighborhood does not make a child immune to subsequent deterioration in his or her neighborhood environment. We also find slightly larger exposure effects for children from above-median income families relative to below-median income families ($\gamma = 0.047$ vs. $\gamma = 0.031$).

Finally, we distinguish between two different mechanisms that could explain why moving to a better area at a younger age is more beneficial: exposure effects – the mechanism we have focused on above – and critical age effects. Critical age (or critical period) models predict that the impacts of moving to a different neighborhood vary with children's ages (e.g., Lynch and Smith 2005). For example, suppose that moving to a better neighborhood improves a child's network of friends with a probability that falls with the age at move and that once one makes new contacts, they last forever. In this model, neighborhood effects would decline with a child's age at move (as in Figure IV), but the duration of exposure to a better area would not matter for long-term outcomes. Alternatively, if better neighborhoods offer a positive treatment (such as better schooling) in each year of childhood, the key determinant of outcomes would be the total duration of exposure rather than the specific age at which a child moves. Distinguishing between these mechanisms can be important for policy: the critical age view calls for improving children's environments at certain key ages while the exposure view calls for a sustained improvement in environment throughout childhood.

A critical age model cannot be distinguished from an exposure effect model in a sample of one-

---

[28]This specification likely over-adjusts for differences in labor market conditions and underestimates γ because the CZ in which the child resides as an adult is itself an endogenous outcome that is likely related to the quality of a child's environment. For example, one of the effects of growing up in a good area may be an increased probability of getting a high-paying job in another city.

[29]Moreover, roughly an equal fraction of families with children move to CZs with better vs. worse outcomes; 48.7% move to CZs with $\Delta_{odps} > 0$. This contrasts with sorting models suggesting families with children would tend to sort to CZs that produce better outcomes.

time movers because a child's age at move is perfectly collinear with his duration of exposure to the new area. However, this collinearity is broken when families move multiple times. Intuitively, one can distinguish between the critical age and exposure mechanisms by considering children who move to an area with better permanent residents' outcomes $\bar{y}_{pds}$ but then move back to the place where they started. In this case, the exposure model predicts that children will experience gains that are proportional to the number of years they spent in the destination CZ, whereas the critical age model predicts that the gain will depend only upon the age at which the child first moves to the new area.

To implement this analysis, we first generalize the specification in (7) to include families who move more than once by replacing the Δ_{odps} terms with a duration-weighted measure of exposure to different areas over childhood (see Online Appendix B for details). This multiple movers specification yields an annual exposure effect estimate of $\gamma = 0.042$ (s.e. $= 0.001$) (Online Appendix Table IV, Column 2). We then test between the critical age and exposure mechanisms by controlling for the age of the child at the time of each move j interacted with the change in permanent residents' outcomes ($\Delta_{od(j)ps}$). This specification, which isolates variation in exposure that is orthogonal to the ages at which children move, yields an exposure effect estimate of $\gamma = 0.036$ (s.e. $= 0.005$) (Online Appendix Table IV, Column 4). The similarity between this estimate and our baseline estimate of $\gamma \simeq 0.04$ implies that what matters for children's incomes in adulthood is the total time spent in a given area (exposure) rather than the age at which one arrives in that area.[30]

IV.D Summary

Under our key identification assumption (Assumption 1), the empirical results in this section yield three lessons. First, place matters: children who move at earlier ages to areas where prior residents have higher incomes earn more themselves as adults. Second, place matters in proportion to the duration of childhood exposure. Every year of exposure to the better area during childhood contributes to higher income in adulthood. Third, each year of childhood exposure matters roughly equally. The returns to growing up in a better neighborhood persist well beyond early childhood.

All of these conclusions rest on the assumption that selection effects do not vary with the child's age at move. We evaluate the validity of this assumption in the next section.

---

[30]Critical age effects have been most widely documented in linguistic patterns and anthropometric measures (e.g., Singleton and Ryan 2004, Bleakley and Chin 2004, van den Berg et al. 2014). One potential explanation for why we do not find evidence of critical age effects here is that we focus on U.S. natives, for whom language acquisition is less of an issue.

V   Validation of Baseline Design

We assess the validity of our key identifying assumption – that the potential outcomes of children who move to better vs. worse areas do not vary with the age at which they move – using a series of tests that focus on different forms of selection and omitted variable bias. To organize the analysis, we partition the unobserved determinant of children's outcomes, represented by θ_i in equation (3), into two components: a component $\bar{\theta}_i$ that reflects inputs that are *fixed* within families, such as parent genetics and education, and a residual component $\tilde{\theta}_i = \theta_i - \bar{\theta}_i$ that may vary over time within families, such as parents' jobs.

We implement four tests for bias in this section. First, we address bias due to selection on fixed family factors $\bar{\theta}_i$ by comparing siblings' outcomes. Second, we control for changes in parents' income and marital status, two key time-varying factors $\tilde{\theta}_i$ that we observe in our data. Our remaining tests focus on *unobservable* time-varying factors, such as changes in wealth, that may have triggered a move to a better area. In our third set of tests, we isolate moves that occur due to displacement shocks that induce many families to move. Finally, we conduct a set of outcome-based placebo (overidentification) tests of the exposure effect model, exploiting heterogeneity in permanent residents' outcomes across subgroups to generate sharp testable predictions about how children's outcomes should change when they move to different areas. In our view, this last approach, although least conventional, provides the most compelling evidence that the identifying assumption holds and that neighborhoods have causal exposure effects on children's long-term outcomes.

V.A   Sibling Comparisons

If families with better unobservables (higher $\bar{\theta}_i$) move to better neighborhoods at earlier ages, Assumption 1 would be violated and our estimated exposure effect $\hat{\gamma}$ would be biased upward. We control for differences in such family-level factors $\bar{\theta}_i$ by including family fixed effects when estimating (6). For example, consider a family that moves to a better area with two children, who are ages m_1 and m_2 at the time of the move. When including family fixed effects, the exposure effect γ is identified by the extent to which the difference in siblings' outcomes, $y_1 - y_2$, covaries with the difference in their ages interacted with the change in permanent residents' outcomes, $(m_1 - m_2)\Delta_{odps}$.

Figure Va replicates Figure IVb, adding family fixed effects to equation (6). The linear decline in the estimated values of b_m until age 23 is very similar to that in the baseline specification.

Children who move to a better area at younger ages have better outcomes than their older siblings. Regressing the b_m coefficients on m for $m \leq 23$ yields an average annual exposure effect estimate of $\gamma = 0.043$ (s.e. $= 0.003$), very similar to our estimates above.

The selection effect (i.e., the level of b_m after age 24) falls from $\delta = 0.23$ in the baseline specification to $\delta = 0.01$ (not significantly different from zero) with family fixed effects.[31] Family fixed effects thus reduce the *level* of the b_m coefficients by accounting for differential selection in which types of families move to better vs. worse areas, but do not affect the *slope* of the b_m coefficients. This is precisely what we should expect if selection effects in where families choose to move do not vary with children's ages when they move, as required by Assumption 1.

Column 7 of Table II shows that adding family fixed effects to the linear specification in equation (7) and estimating the model directly on the micro data yields an estimate of $\gamma = 0.044$. Other variants of this regression specification, analogous to those in Columns 2-6 of Table II, all yield very similar estimates of γ, with one exception: excluding the cohort interactions with $\bar{y}_{pos}$ and Δ_{odps}, as in Column 4, yields $\gamma = 0.031$ (Column 8 of Table II). The reason that the estimate of γ falls in this specification is that we observe children's origin locations for fewer years in earlier birth cohorts, as discussed in Section IV.A. The missing data on origins increases the level of the selection effect δ in earlier cohorts (see Online Appendix A). Because we only observe moves at older ages for children in earlier cohorts, these differences across cohorts induce a positive correlation between δ_m and m, biasing our estimate of γ downward. This bias is magnified in the specifications with family fixed effects because they are identified purely by comparing the outcomes of children in different birth cohorts, whereas our baseline specifications also compare children in the same birth cohort whose parents move at different times. Including cohort interactions with Δ_{odps} eliminates this bias by permitting a separate selection term δ for each cohort.[32]

In sum, we continue to find childhood exposure effects of $\gamma \simeq 0.04$ when comparing siblings' outcomes, implying that our design is not confounded by differences in the types of families who move to better areas when their children are younger.

V.B Controls for Time-Varying Observables

The research design in Figure Va accounts for bias due to fixed differences in family inputs $\bar{\theta}_i$, but it does not account for time-varying inputs $\tilde{\theta}_i$. For example, moves to better areas may be triggered

---

[31]δ is identified even with family fixed effects because Δ_{odps} varies across birth cohorts.

[32]The attenuation bias in γ is further amplified in CZs with smaller populations, where Δ_{odps} is measured with greater error (see Online Appendix A and Appendix Table VI).

by events such as job promotions that directly affect children's outcomes in proportion to their time of exposure to the destination. Such shocks could bias our estimate of β upward even with family fixed effects.

Prior research has focused on changes in parents' income and marital status as two key factors that may induce moves and also directly affect children's outcomes in adulthood (e.g., Jencks and Mayer 1990). We can directly control for these two time-varying factors in our data, as we observe parents' incomes and marital status in each year from 1996-2012. We control for the effects of changes in income around the move when estimating (6) by including controls for the change in the parent's income rank from the year before to the year after the move interacted with indicators for the child's age at move. The interactions with age at move permit the effects of income changes to vary with the duration of childhood exposure to higher vs. lower levels of parent income. Similarly, we control for the impact of changes in marital status by interacting indicators for each of the four possible changes in the mother's marital status in the year before vs. after the move (married to unmarried, unmarried to married, unmarried to unmarried, and married to married) with indicators for the child's age at move.

Figure Vb replicates Figure Va, controlling for all of these variables in addition to family fixed effects. Controlling for changes in parent income and marital status has little effect on the estimates of $\{b_m\}$. The estimates of $\gamma = 0.042$ and $\delta = 0.015$ are virtually identical to those when we do not control for these time-varying factors. Column 9 of Table II confirms that including these controls in a linear regression estimated on the micro data yields similar estimates.

These results show that changes in income and family structure are not a significant source of bias in our design. However, other unobserved factors could still be correlated with moving to a better or worse area in a manner that generates omitted variable bias. The fundamental identification problem is that any unobserved shock that induces child i's family to move to a different area could be correlated with parental inputs θ_i. These changes in parental inputs could potentially increase the child's income y_i in proportion to the time spent in the new area even in the absence of neighborhood effects. For example, a wealth shock might lead a family to both move to a better neighborhood and increase investments in the child in the years after the shock, which could improve y_i in proportion to exposure time independent of neighborhood effects. In the next two subsections, we address concerns about bias due to such unobserved factors.

V.C Displacement Shocks

One approach to accounting for unobservable shocks is to identify moves where we have some information about the shock that precipitated the move. Suppose we identify families who were forced to move from an origin o to a nearby destination d because of an exogenous shock such as a natural disaster. Such displacement shocks can induce differential changes in neighborhood quality as measured by permanent residents' outcomes (Δ_{odps}). For instance, Hurricane Katrina displaced families from New Orleans (an area with relatively poor outcomes compared to surrounding areas), leading to an increase in average neighborhood quality for displaced families ($\Delta_{odps} > 0$). In contrast, Hurricane Rita hit Houston, an area with relatively good outcomes, and may have reduced neighborhood quality ($\Delta_{odps} < 0$). If these displacement shocks do not have direct exposure effects on children that are correlated with Δ_{odps} – e.g., the direct effects of the disruption induced by hurricanes does not covary with neighborhood quality changes – then Assumption 1 is satisfied and we obtain unbiased estimates of the exposure effect γ by focusing on displaced families. Conceptually, by isolating a subset of moves caused by known exogenous shocks, we can more credibly ensure that changes in children's outcomes are not driven by unobservable factors.[33]

To operationalize this approach, we first identify displacement shocks based on population outflows at the ZIP code level. Let K_{zt} denote the number of families who leave ZIP code z in year t in our sample of one-time movers and $\bar{K}_z$ mean outflows between 1996 and 2012. We define the shock to outflows in year t in ZIP z as $k_{zt} = K_{zt}/\bar{K}_z$.[34]

Though many of the families who move in subsamples with large values of k_{zt} do so for exogenous reasons, their destination d is still an endogenous choice that could lead to bias. For example, families who choose to move to better areas (higher $\bar{y}_{pds}$) when induced to move by an exogenous shock might also invest more in their children. To reduce potential biases arising from the endogenous choice of destinations, we isolate variation arising from the *average* change in neighborhood quality for individuals who are displaced. Let $E[\Delta_{odps}|q,z]$ denote the difference in the mean predicted outcome in the destination CZs relative to the origin CZ for individuals in origin ZIP code z and parent income decile q (averaging over all years in the sample, not just the year

---

[33]This research design is related to Sacerdote's (2012) analysis of the effects of Hurricanes Katrina and Rita on student test score achievement. Although we use similar variation, we do not focus on the direct effects of the displacement itself, but rather on how children's long-term outcomes vary in relation to the outcomes of permanent residents in the destination to which they were displaced.

[34]Searches of historical newspaper records for cases with the highest outflow rates k_{zt} reveal that they are frequently associated with events such as natural disasters or local plant closures. Unfortunately, there is insufficient power to estimate exposure effects purely from the events identified in newspapers.

of the shock). We instrument for the difference in predicted outcomes in each family's destination relative to origin (Δ_{odps}) with $E[\Delta_{odps}|q,z]$ and estimate the linear specification in (7) using 2SLS to identify the exposure effect, γ_{IV}.[35]

Figure VI presents the results of this analysis. To construct this figure, we take ZIP-year cells with above-median outflows ($k_{zt} > 1.17$) and divide them into population-weighted bins based on the size of the shock k_{zt}.[36] The first point in Figure VI shows the 2SLS estimate of the annual exposure effect γ_{IV} using all observations with k_{zt} greater than its median value (1.17). The second point shows the estimate of γ_{IV} using all observations with k_{zt} at or above the 52nd percentile. The remaining points are constructed in the same way, increasing the threshold by 2 percentiles at each point, with the last point representing an estimate of γ_{IV} using data only from ZIP codes in the highest two percentiles of outflow rates. The dotted lines show a 95% confidence interval for the regression coefficients.

If the baseline estimates were driven entirely by selection, γ_{IV} would fall to 0 as we limit the sample to individuals who are more likely to have been induced to move because of an exogenous displacement shock. But the coefficients remain quite stable at $\gamma_{IV} \simeq 0.04$ even when we restrict to moves that occurred as part of large displacements. That is, when we focus on families who move to a better area for what are likely to be exogenous reasons, we continue to find that children who are younger at the time of the move earn more as adults.

These findings support the view that our baseline estimates of exposure effects capture the causal effects of neighborhoods rather than other unobserved factors that change when families move. Moreover, they indicate that the treatment effects of moving to a different area are similar for families who choose to move for idiosyncratic reasons and families who are exogenously displaced. This result suggests that the exposure effects identified by our baseline design can be generalized to a broader set of families beyond those who choose to make a particular move.

---

[35] This approach does not fully eliminate the scope for selection bias, as biases from the endogenous choice of destinations could persist if there is unobserved heterogeneity across areas experiencing displacement shocks. However, it reduces the scope for selection bias by focusing on moves induced by aggregate displacement shocks and eliminating variation in Δ_{odps} due to individual choice, which is more likely to be correlated with unobservables θ_i than the area-level variation in $E[\Delta_{odps}|q,z]$. By testing if the estimate of γ remains stable when we use an estimator that reduces the scope for selection, we can gauge whether our baseline estimate of γ is biased.

[36] To ensure that large outflows are not driven by areas with small populations, we exclude ZIP-year cells with less than 10 children leaving in that year.

V.D Outcome-Based Placebo Tests

As a final approach to test for bias due to unobservable factors, we implement placebo tests that exploit the heterogeneity in permanent residents' outcomes across subgroups. We exploit variation along three dimensions: birth cohorts, quantiles of the income distribution, and child gender. The causal exposure effect model predicts precise convergence of a child's outcome to permanent residents' outcomes for his or her *own* subgroup. In contrast, we argue below that omitted variable and selection models would not generate such subgroup-specific convergence under plausible assumptions about parents' information sets and preferences. The heterogeneity in permanent residents' outcomes thus gives us a rich set of overidentifying restrictions to test whether neighborhoods have causal effects.[37] We consider each of the three dimensions of heterogeneity in turn.

Birth Cohorts. Although permanent residents' outcomes are generally very stable over time, outcomes in some areas (such as Oklahoma City, OK) have improved over time, while others (such as Sacramento, CA) have gotten worse.[38] Such changes could occur, for instance, because of changes in the quality of local schools or other area-level characteristics that affect children's outcomes. We exploit this heterogeneity across birth cohorts to test for confounds in our baseline research design.

Under the causal exposure effect model, when a child's family moves to destination d, the difference in permanent residents' outcomes $\Delta_{odp,s(i)}$ for that child's own birth cohort $s(i)$ should predict his or her outcomes more strongly than the difference in outcomes Δ_{odps} for other cohorts $s \neq s(i)$. Intuitively, what matters for a child's outcome is a neighborhood's quality for his own cohort, not the neighborhood's quality for younger or older cohorts. In contrast, it is unlikely that other unobservables θ_i will vary sharply across birth cohorts s in association with Δ_{odps} because the fluctuations across birth cohorts are realized only in adulthood and thus cannot be directly observed at the time of the move.[39] Therefore, by testing whether exposure effects are predicted by a child's own vs. surrounding cohorts, we can assess the importance of bias due to unobservables.

We implement this analysis by estimating the baseline specification in (7), replacing the change in permanent residents' outcomes for the child's own cohort, $\Delta_{odp,s(i)}$, with analogous predictions for adjacent birth cohorts $s(i) + t$, $\Delta_{odp,s(i)+t}$ (see Online Appendix D for details). The series in red triangles in Figure VII plots the exposure effect estimates ($\tilde{\gamma}_t$) obtained from these regressions,

---

[37] In addition to being useful for identification, these results are also of direct interest in understanding the heterogeneity of place effects across subgroups.

[38] The autocorrelation of $\bar{y}_{pcs}$ with $\bar{y}_{pc,s-1}$ across children's birth cohorts is 0.95 at the 25th percentile of the parent income distribution.

[39] For instance, a family that moves with a 10 year old child will not observe $\bar{y}_{pds}$ for another 14 years (if income is measured at age 24).

with t ranging from -4 to 4. The estimates of $\tilde{\gamma}_t$ are similar to our baseline estimate of $\gamma = 0.040$ for the leads and lags, consistent with the high degree of serial correlation in permanent residents' outcomes. The series in blue circles plots analogous coefficients $\tilde{\gamma}_t$ when all the cohort-specific predictions from the four years before to the four years after the child's own cohort are included *simultaneously*. In this specification, the coefficients on the placebo exposure effects ($\tilde{\gamma}_t$ for $t \neq 0$) are all very close to zero and not statistically significant.[40] However, the exposure effect estimate for the child's own cohort remains at approximately $\gamma = 0.04$ even when we control for the surrounding cohorts' predictions and is significantly different from the estimates of $\tilde{\gamma}_t$ for $t \neq 0$ ($p < 0.001$).

The evidence in Figure VII strongly supports the view that the change in children's outcomes is driven by causal effects of exposure to a different place. Intuitively, it is unlikely that a correlated shock – such as a change in wealth when the family moves – would covary precisely with cohort-level differences in place effects, as manifested in the outcomes of children of permanent residents. Formally, this test relies on the assumption that if unobservables θ_i are correlated with exposure to a given cohort $s(i)$'s place effect (proxied for by permanent residents' outcomes), they must also be correlated with exposure to the place effects of adjacent cohorts t:

$$Cov(\theta_i, m\Delta_{odp,s(i)}|X) > 0 \Rightarrow Cov(\theta_i, m\Delta_{odpt}|X, m\Delta_{odp,s(i)}) > 0, \tag{8}$$

where X represents the vector of fixed effects and other controls in (7). Under this assumption, the findings in Figure VII imply that our estimates of γ reflect causal neighborhood effects (which are cohort-specific) rather than omitted variables, which are *not* cohort-specific under (8).

Quantiles: Distributional Convergence. Places differ not only in children's mean outcomes, but also in the *distribution* of children's outcomes. For example, children who grow up in low-income families in Boston and San Francisco have comparable mean ranks, but children in San Francisco are more likely to end up in the tails of the income distribution than those in Boston. If neighborhoods have causal exposure effects, we would expect convergence in movers' outcomes not just at the mean but across the entire distribution in proportion to exposure time. In contrast, it is less plausible that omitted variables such as wealth shocks would perfectly replicate the distribution of outcomes of permanent residents in each CZ.[41] Therefore, testing for quantile-specific convergence can distinguish the causal exposure effect model from omitted variable explanations.

---

[40]A test of the joint hypothesis that all $\tilde{\gamma}_t = 0$ for all $t \neq 0$ yields a p-value of 0.251.

[41]Families are unlikely to be able to forecast their child's eventual quantile in the income distribution, making it difficult to sort precisely on quantile-specific neighborhood effects. Even with such knowledge, there is no ex-ante reason to expect unobserved shocks such as changes in wealth to have differential and potentially non-monotonic effects across quantiles, in precise proportion to the outcomes in the destination.

To implement these tests, we begin by constructing predictions of the probability of having an income in the upper or lower tail of the national income distribution at age 24 for children of permanent residents in each CZ c. In each CZ, we regress an indicator for a child being in the top 10% of the distribution or an indicator for not being employed on parent income rank p using an equation analogous to (1), including a quadratic term in parental income rank p to account for the nonlinearities in tail outcomes identified in Chetty et al. (2014). We then calculate the predicted probability of being non-employed π^U_{pcs} and being above the 90th percentile π^{90}_{pcs} using the fitted values from these regressions, as in (2).

In Table III, we estimate exposure effect models analogous to (7) using these distributional predictions instead of mean predictions. In Columns 1-3, the dependent variable is an indicator for having income in the top 10% of the income distribution. Column 1 replicates the baseline specification in (7), using $\Delta^{90}_{odps} = \pi^{90}_{pds} - \pi^{90}_{pos}$ instead of the mean prediction $\Delta_{odps} = \bar{y}_{pds} - \bar{y}_{pos}$ as the key independent variable (see Online Appendix D for the exact regression specifications). We obtain an exposure effect estimate of $\gamma = 0.043$ per year in this specification. Column 2 uses the change in the predicted mean rank, Δ_{odps}, instead. Here, we obtain a statistically significant estimate of 0.024, as expected given the high degree of correlation in permanent residents' outcomes across quantiles: places where more children reach the top 10% also tend to have better mean outcomes. In Column 3, we include both the quantile prediction Δ^{90}_{odps} and the mean prediction Δ_{odps}, identifying the coefficients purely from differential variation across quantiles within CZs. The coefficient on the quantile prediction remains unchanged at approximately $\gamma = 0.04$, while the coefficient on the mean prediction is not significantly different from 0.

Columns 4-6 of Table III replicate Columns 1-3, using an indicator for non-employment as the dependent variable and the prediction for non-employment Δ^U_{odps} instead of Δ^{90}_{odps} as the key independent variable. As in the upper tail, children's probabilities of being in the lower tail of the income distribution are fully determined by the quantile-specific prediction rather than the mean prediction. In Column 6, the coefficient on the non-employment prediction Δ^U_{odps} is $\gamma = 0.043$, while the placebo coefficient on the mean rank prediction is -0.002.

In short, we find evidence of distributional convergence: controlling for mean outcomes, children's outcomes converge to predicted outcomes in the destination across the distribution in proportion to exposure time, at a rate of approximately 4% per year.[42] Since omitted variables such

---

[42] The rate of convergence need not be identical across all quantiles of the income distribution because the prediction for permanent residents at each quantile π^{90}_{pcs} could reflect a different combination of causal effects and sorting. The key test is whether the prediction for the relevant quantile has more predictive power than predictions at the mean

as wealth shocks would be unlikely to generate such distributional convergence, this finding again supports the view that the convergence in mover's outcomes is driven by causal effects of place. Formally, assume that if unobservables θ_i are correlated positively with exposure to place effects on upper (or lower) tail outcomes π_{pcs}^q, they must also be correlated with exposure to the place effects on mean incomes (proxied for by permanent residents' outcomes):

$$Cov(\theta_i, m\Delta_{odps}^q | X^q) > 0 \Rightarrow Cov(\theta_i, m\Delta_{odps} | X^q, m\Delta_{odps}^q) > 0. \tag{9}$$

Under this assumption, the findings in Table III imply that our estimates of γ reflect causal place effects (which are quantile-specific) rather than omitted variables, which are not quantile-specific under (9).

Gender. Finally, we conduct an analogous set of placebo tests exploiting heterogeneity in permanent residents' outcomes by child gender. We begin by constructing gender-specific predictions of the mean household income ranks of children of permanent residents by estimating (1) separately for male and female children, which we denote by $\bar{y}_{pcs}^m$ and $\bar{y}_{pcs}^f$. Places that are better for boys are generally better for girls as well: the (population-weighted) correlation of $\bar{y}_{pcs}^m$ and $\bar{y}_{pcs}^f$ across CZs is 0.93 at the median ($p = 50$).[43] We exploit the residual variation across genders to conduct placebo tests analogous to those above, based on the premise that unobservable shocks are unlikely to have gender-specific effects.

In Table IV, we estimate exposure effect models analogous to (7) with separate predictions by gender. Column 1 replicates (7) using the gender-specific prediction Δ_{odps}^g instead of the prediction that pools both genders. We obtain an exposure effect estimate of $\gamma = 0.038$ per year in this specification. In Column 2, we use the prediction for the other gender Δ_{odps}^{-g} instead. Here, we obtain an estimate of 0.031, as expected given the high degree of correlation in outcomes across genders. In Column 3, we include predictions for both genders, identifying the coefficients purely from differential variation across genders within CZs. In this specification, the coefficient on the own gender prediction is $\gamma = 0.03$, three times larger than the other-gender prediction, which is close to zero.[44]

---

or other quantiles.

[43]Online Appendix Figure V presents choropleth maps of $\bar{y}_{pcs}^m - \bar{y}_{pcs}^f$ at $p = 25$ and $p = 75$. For low-income families ($p = 25$), outcomes for boys are relatively worse than those for girls in areas with higher crime rates, a larger fraction of single parents, and greater inequality (Chetty et al. 2016).

[44]It is not surprising that the other gender prediction remains positive, as the prediction for the other gender may be informative about a place's effect for children of a given gender due to measurement error. In general, finding a 0 effect on the "placebo" prediction is sufficient but not necessary to conclude that there is no sorting under an assumption analogous to (8).

One may be concerned that families sort to different areas based on their child's gender, which – unlike the quantile and cohort-specific variation used above – is known at the time of the move. To address this concern, Columns 4-6 of Table IV replicate Columns 1-3 including family fixed effects. The own-gender prediction remains a stronger predictor of children's outcomes than the other-gender prediction even when we compare siblings' outcomes within families. Column 7 shows that this remains the case when we restrict the sample to families that have at least one boy and one girl, for whom differential sorting by gender is infeasible.

The gender-specific convergence documented in Table IV supports the causal exposure effects model under an assumption analogous to (8), namely that the unobservable θ_i does not vary differentially across children of different genders within a family. This assumption requires that families who move to areas that are particularly good for boys do not systematically invest more in their sons relative to their daughters, a restriction that would hold if, for instance, families do not have different preferences over their sons' and daughters' outcomes. Under this assumption, the gender-specific convergence in proportion to exposure time must reflect causal place effects.

V.E   Summary

The results in this section show that various refinements of our baseline design – such as including family fixed effects or exploiting cohort- or gender-specific variation – all yield annual exposure effect estimates of $\gamma \simeq 0.04$. These findings imply that any omitted variable θ_i that generates bias in our estimate of the exposure effect γ must: (1) operate within the family in proportion to exposure time (family fixed effects); (2) be orthogonal to changes in parental income and marital status (controls for observables); (3) persist in the presence of moves induced by displacement shocks (displacement shock analysis); and (4) precisely replicate permanent residents' outcomes by birth cohort, quantile, and gender in proportion to exposure time (outcome-based placebo tests). We believe that plausible omitted variables are unlikely to have all of these properties. We therefore conclude that our estimate of $\gamma \simeq 0.04$ is an unbiased estimate of the annual childhood exposure effect. Extrapolating over the duration of childhood, from age 0 to 20, this estimate implies that growing up from birth in an area with one unit higher income ranks in observational data increases a given child's income rank by about 0.8 units (80% of the observational difference).

An auxiliary implication of the results in this section is that the simple baseline design of comparing families who move with children of different ages is not confounded by selection and omitted variable biases. Although there is clear evidence of selection in terms of *where* families

move – as shown by the estimate of $\delta > 0$ in Figure IV – we find no evidence of differential selection based on *when* families move to a better vs. worse area (at least after their children are nine years old).[45] This finding implies that research designs exploiting variation in the timing of moves can be used to identify the causal effects of neighborhoods in observational data, providing a scalable tool for identifying neighborhood effects even in the absence of randomized experiments.

VI Other Outcomes

In this section, we estimate neighborhood effects for several other outcomes beyond income: college attendance, marriage, teenage birth, and teenage employment. This analysis provides further evidence on the types of outcomes that are shaped by neighborhoods and illustrate how neighborhoods affect behavior before children enter the labor market.

Figure VIII replicates Figure IVb using college attendance and marriage as the outcomes. In Panel A, we replicate the specification in equation (6), replacing Δ_{odps} with $\Delta^C_{odps} = C_{pds} - C_{pos}$, where C_{pcs} is the fraction of children who attend college at any point between ages 18 and 23 (among children of permanent residents in CZ c in birth cohort s with parental income rank p). In Panel B, we replace Δ_{odps} with $\Delta^M_{odps} = M_{pds} - M_{pos}$, where M_{pcs} is the fraction of children who are married at age 26.

We find evidence of childhood exposure effects until age 23 for both of these outcomes. Moving to an area with higher college attendance rates at a younger age increases a child's probability of attending college. Likewise, moving at a younger age to an area where permanent residents are more likely to be married increases a child's probability of being married. Using parametric models analogous to (7), the estimated annual exposure effect for college attendance is comparable to our estimates for income ($\gamma = 0.037$) and is slightly smaller for marriage ($\gamma = 0.025$).

In Figure IX, we analyze outcomes measured while children are teenagers. Panel A considers teen birth, defined as being listed as a parent on a birth certificate prior to age 20. We construct gender-specific predictions of teenage birth rates and plot estimates from the baseline specification in (6), replacing Δ_{odps} with $\Delta^z_{odpsg} = z_{pdsg} - z_{posg}$, where z_{pcsg} is the fraction of children of permanent residents with parental income p in CZ c, cohort s, and gender g who have a teenage birth. For both boys and girls, there are clear childhood exposure effects: moving at an earlier age to an area

---

[45]Such differential selection might be small because the outcomes of children of permanent residents $\bar{y}_{pcs}$ are not highly correlated with mean parent incomes across areas (Chetty et al. 2014). As a result, moving to a better area for children (higher $\bar{y}_{pcs}$) is not systematically associated with parents finding higher-paying jobs, mitigating what might be the most important confounding factor for our design.

with a higher teen birth rate increases a child's probability of having a teenage birth. The gradient is especially steep between ages 13 and 18, suggesting that a child's neighborhood environment during adolescence may play a particularly important role in determining teen birth outcomes.

In Panels B-D of Figure IX, we analyze neighborhood effects on teenage employment rates. In these figures, the outcomes are indicators for employment (based on having a W-2 form filed on one's behalf) at ages 16, 17, or 18. The key independent variable (corresponding to Δ_{odps}) in each of these figures is the difference in employment rates of children of permanent residents in the destination vs. the origin CZ at the relevant age (16, 17, or 18). For teen employment, we find discontinuous effects of moving just before employment is measured rather than continuous exposure effects. Children who move at age 15 to a CZ where more 16-year-olds work are much more likely to work at age 16 than children who make the same move at age 17. Making the same move at earlier ages (before age 16) further increases the probability of working at age 16, but the exposure effect is small relative to the jump at age 16 itself. Analogous jumps are observed at ages 17 and 18 when we measure employment at ages 17 and 18 (Panels C and D).[46] These jumps suggest that neighborhood effects may be partly driven by distinct experiences at different points of childhood, such as summer jobs that are available in a given area at certain ages. Such age-specific impacts may aggregate to produce the linear childhood exposure effects that shape outcomes in adulthood.

Although the mean income of individuals in an area is correlated with other outcomes such as college attendance and teenage birth rates, there is substantial independent variation in each of these outcomes. For example, permanent residents' mean income ranks at age 30 have a (population-weighted) correlation of 0.46 with college attendance rates for children with parents at $p = 25$ (Online Appendix Table VII). Hence, the finding that movers' outcomes converge to those of permanent residents on all of these dimensions constitutes further evidence that neighborhoods have causal effects, as it would be unlikely that unobserved confounds would generate such convergence on a spectrum of different outcomes.[47] Moreover, the fact that neighborhoods have causal effects on a wide variety of outcomes beyond earnings further suggests that the mechanism through which neighborhoods shape children's outcomes is not driven by labor market conditions but rather a set

---

[46]The magnitude of the $\{b_m\}$ coefficients in Panels B-D is approximately 0.8 at young ages and 0 after the age at which employment is measured. Under our identifying assumption of constant selection effects by age, this implies that children who move at birth pick up 80% of the differences in teenage employment rates across CZs observed for permanent residents.

[47]This logic is analogous to the tests for distributional convergence in Section V.D; here, we effectively test for convergence in the joint distribution of income and various other outcomes.

of environmental factors that shape behaviors throughout childhood.

VII   Conclusion

This paper has shown that children's opportunities for economic mobility are shaped by the neighborhoods in which they grow up. Neighborhoods affect children's long-term outcomes through childhood exposure effects: every extra year a child spends growing up in an area where permanent residents' incomes are higher increases his or her income. Movers' outcomes converge to those of permanent residents in the destination to which they move at a rate of approximately 4% per year of childhood exposure. Extrapolating this annual exposure effect over 20 years of childhood, children who move to a new area at birth will pick up roughly 80% of the difference in permanent residents' outcomes between their origin and destination. Much of the variation in intergenerational mobility observed across areas thus appears to be driven by causal effects of place rather than differences in the types of people living in those places.

These results motivate place-focused approaches to improving economic mobility, such as making investments to improve outcomes in areas that currently have low levels of mobility or helping families move to higher opportunity areas. Identifying specific policy solutions – i.e., the investments needed to improve mobility and the areas to which families should be encouraged to move – requires identifying the causal effect of each neighborhood and understanding what makes some areas produce better outcomes than others. The analysis in the present paper shows that differences in permanent residents' outcomes are predictive of neighborhoods' causal effects *on average*. However, it does not provide estimates of the causal effect of each area on children's outcomes, as the outcomes of permanent residents in any given area will reflect a different mix of selection and causal effects. We construct estimates of the causal effect of growing up in each CZ and county in the U.S. and characterize the properties of areas that produce good outcomes in the next paper in this series.

References

Aaronson, D. (1998). Using sibling data to estimate the impact of neighborhoods on childrens educational outcomes. *Journal of Human Resources 33*(4), 915–46.

Altonji, J. G. and R. K. Mansfield (2016). Estimating group effects using averages of observables to control for sorting on unobservables: School and neighborhood effects. *Working Paper*.

Basu, S. (2010). Age of entry effects on the education of immigrant children: A sibling study. *Available at SSRN 1720573*.

Bleakley, H. and A. Chin (2004). Language skills and earnings: Evidence from childhood immigrants. *Review of Economics and Statistics 86(2): 481-96*.

Chetty, R., J. N. Friedman, and E. Saez (2013). Using differences in knowledge across neighborhoods to uncover the impacts of the eitc on earnings. *American Economic Review 103*(7), 2683–2721.

Chetty, R., J. N. Friedman, E. Saez, N. Turner, and D. Yagan (2017). Mobility report cards: The role of colleges in intergenerational mobility. *Working Paper*.

Chetty, R. and N. Hendren (2017). The impacts of neighborhoods on intergenerational mobility ii: County-level estimates. *Working Paper*.

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

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *Quarterly Journal of Economics 129*(4), 1553–1623.

Chetty, R., N. Hendren, F. Lin, J. Majerovitz, and B. Scuderi (2016). Childhood environment and gender gaps in adulthood. *The American Economic Review Papers and Proceedings 106*(5), 282–288.

Chyn, E. (2016). Moved to opportunity: The long-run effect of public housing demolition on labor market outcomes of children. *Working Paper*.

Cilke, J. (1998). A profile of non-filers. *U.S. Department of the Treasury, Office of Tax Analysis Working Paper No. 78*.

Clampet-Lundquist, S. and D. S. Massey (1993). Neighborhood effects on economic self sufficiency: A reconsideration of the moving to opportunity experiment. *American Journal of Sociology 114.1 (2008): 107-143*.

Coleman, J. S. (1988). Social capital in the creation of human capital. *American Journal of Sociology 94*, pp. S95–S120.

Crowder, K. and S. J. South (2011). Spatial and temporal dimensions of neighborhood effects on high school graduation. *Social Science Research 40*(1), 87–106.

Damm, A. P. and C. Dustmann (2014). Does growing up in a high crime neighborhood affect youth criminal behavior? *American Economic Review 104*(6), 1806–32.

Finkelstein, A., M. Gentzkow, and H. Williams (2016). Sources of geographic variation in health care: Evidence from patient migration*. *The Quarterly Journal of Economics 131*(4), 1681.

Haider, S. and G. Solon (2006). Life-cycle variation in the association between current and lifetime earnings. *American Economic Review 96*(4), 1308–1320.

Jencks, C. and S. E. Mayer (1990). The social consequences of growing up in a poor neighborhood. Technical report, National Research Council.

Katz, L. F., J. R. Kling, and J. B. Liebman (2001). Moving to opportunity in boston: Early results of a randomized mobility experiment. *The Quarterly Journal of Economics 116*(2), 607–654.

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

Ludwig, J., G. J. Duncan, L. A. Gennetian, L. F. Katz, R. C. Kessler, J. R. Kling, and L. Sanbonmatsu (2013). Long-term neighborhood effects on low-income families: Evidence from moving to opportunity. *American Economic Review Papers and Proceedings 103(3): 226-31*.

Ludwig, J., J. B. Liebman, J. R. Kling, G. J. Duncan, L. F. Katz, R. C. Kessler, and L. Sanbonmatsu (2008). What can we learn about neighborhood effects from the moving to opportunity experiment? *American Journal of Sociology 114*, 144–188.

Lynch, J. and G. D. Smith (2005). A life course approach to chronic disease epidemiology. *Annual Review of Public Health 26*(1), 1–35. PMID: 15760279.

Massey, D. S. (1993). *American apartheid: Segregation and the making of the underclass*. Harvard University Press.

Oreopoulos, P. (2003). The long-run consequences of living in a poor neighborhood. *Quarterly Journal of Economics 118*(4), 1533–1175.

Plotnick, R. and S. Hoffman (1996). The effect of neighborhood characteristics on young adult outcomes: The effect of neighborhood characteristics on young adult outcomes: Alternative estimates. *Institute for Research on Poverty Discussion Paper no. 1106-96*.

Sacerdote, B. (2012). When the saints go marching out: Long-term outcomes for student evacuees from hurricanes katrina and rita. *American Economic Journal: Applied Economics 4*(1), 109–135.

Sampson, R. J., J. D. Morenoff, and T. Gannon-Rowley (2002). Assessing neighborhood effects: Social processes and new directions in research. *Annual Review of Sociology 28 (1): 443-478*.

Sharkey, P. and J. W. Faber (2014). Where, when, why, and for whom do residential contexts matter? moving away from the dichotomous understanding of neighborhood effects. *Annual Review of Sociology 40*, 559–79.

Singleton, D. and L. Ryan (2004). *Language Acquisition: The Age Factor*. Second language acquisition. Multilingual Matters.

Solon, G. (1992). Intergenerational income mobility in the united states. *American Economic Review 82*(3), 393–408.

Tolbert, C. M. and M. Sizer (1996). U.S. commuting zones and labor market areas: A 1990 update. *Economic Research Service Staff Paper 9614*.

van den Berg, G. J., P. Lundborg, P. Nystedt, and D.-O. Rooth (2014). Crtical periods during childhood and adolescence. *Journal of the European Economic Association 12*(6), 1521–1557.

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

Wodtke, G. T. (2013). Duration and timing of exposure to neighborhood poverty and the risk of adolescent parenthood. *Demography*.

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.

Online Appendix A. Sensitivity to Population and Distance Restrictions

This appendix assesses the sensitivity of our exposure effect estimates to alternative population and distance restrictions. We first discuss the impacts of these restrictions on our baseline specification. We then discuss their impacts on specifications with family fixed effects, focusing on the role of cohort interactions with Δ_{odps}.

Our baseline analysis restricts the sample to origin and destination CZs with more than 250,000 people based on the 2000 Census and requires the distance between the origin and destination ZIP codes to be more than 100 miles. Appendix Table II shows how our estimates of the exposure effect γ using the specification in (7) change when we vary these restrictions.

Column 1 repeats the baseline specification in column 1 of Table II as a reference. Columns 2-4 include moves of all distances, dropping the 100 mile restriction. Column 2 restricts to moves between origins and destinations with at least 50K people; column 3 restricts to moves to and from places with at least 250K people (our baseline population restriction); and column 4 restricts to move to and from CZs with at least 500K people. Columns 5-7 replicate columns 2-4 imposing our baseline restriction that moves are farther than 100 miles. Hence, column 6 is identical to column 1, as it imposes our baseline distance restriction of 100 miles and a 250K population restriction. Columns 8-10 repeat the same specifications as in 5-7, requiring that moves be farther than 200 miles.

We find that using more stringent population and distance restrictions yield estimates of γ similar to the baseline estimate of 0.04, but including moves across shorter distances or between places with smaller populations tends to lead to slightly attenuated coefficients, γ. For example, relaxing the population restriction from 250K to 50K or removing our distance restriction reduces the estimate of γ from 0.04 to 0.036 (s.e. = 0.002), as shown in columns 3 and 5. The reason for this attenuation is that permanent residents' outcomes $\bar{y}_{pcs}$ are estimated with sampling error. Including CZs with smaller populations naturally leads to less precise estimates of the change in permanent residents' outcomes Δ_{odps}, leading to an attenuated estimate of γ. Similarly, moves across short distances are more likely to reflect moves across two areas that have similar characteristics but are on different sides of a CZ border, again leading to measurement error in Δ_{odps} and attenuation in γ. Since the estimates stabilize once we impose a distance restriction of at least 100 miles and a population restriction of at least 250,000 people, we use these cutoffs in our baseline analysis.

Family Fixed Effect Estimates and Cohort Interactions. Population restrictions have a larger

effect on the estimates of γ in specifications that include family fixed effects, especially when we exclude cohort interactions with Δ_{odps}. Appendix Table VI illustrates this point by reporting estimates of γ with and without family fixed effects. The first two columns of this table replicate our baseline analysis, limiting the sample to CZs with more than 250,000 people, while the last two columns include CZs with more than 50,000 people. In each pair of columns, we show estimates both with and without family fixed effects. Panel A includes interactions of $\bar{y}_{pos}$ and Δ_{odps} with the child's birth cohort as shown in equation (7), while Panel B does not.

Relaxing the population restrictions reduces the estimate of γ with family fixed effects from 0.044 to 0.036, twice the change in the estimate without family fixed effects. The change is particularly large when we exclude the cohort interactions, where the estimate falls from 0.031 to 0.023. The family fixed effect estimates are particularly sensitive to cohort interactions and population restrictions for two reasons.

First, children's origin locations are censored in earlier cohorts because we do not observe children's locations prior to 1996. This censoring increases the estimated selection effect δ in earlier cohorts because families often return to where they used to live (or a similar area) after moving elsewhere. This increases the correlation between $\bar{y}_{pds}$ and children's outcomes because $\bar{y}_{pds}$ is correlated with the quality of the place where children lived prior to 1996. Recall that we cannot observe moves at early ages for children in earlier cohorts (e.g., we do not observe moves before age 16 for the 1980 cohort). As a result of these missing data issues, δ_m increases with m on average, violating Assumption 1 and biasing our estimate of γ downward. Including cohort interactions with Δ_{odps} resolves this problem by permitting a separate selection term δ for each cohort. Cohort interactions are especially important in the specifications with family fixed effects because those specifications are identified purely from comparisons across birth cohorts, whereas our baseline specifications also compare children in the same birth cohort whose parents move at different times.

Second, there are two distinct sources of variation that identify exposure effects: differences in children's ages at move m for a given value of Δ_{odps} and differences in Δ_{odps} across birth cohorts due to changes in place effects over time. The latter source of variation becomes more important in models with family fixed effects because the variance in ages between siblings is smaller than the variance in ages across children in different families. Since Δ_{odps} is measured with greater error across cohorts in small CZs, there is more attenuation in our estimate of γ when we include small CZs in models with family fixed effects. Again, this attenuation is mitigated by the inclusion

of cohort controls because we identify γ primarily from the age gaps in the siblings rather than cross-cohort variation in Δ_{odps} when including the cohort controls.

In sum, these results show that our qualitative conclusions are robust to alternative specifications, but provide justification for including cohort controls and imposing the population and distance restrictions used in our baseline specifications to minimize attenuation bias.

Online Appendix B. Estimates Using Families that Move Multiple Times

This appendix discusses how we estimate exposure effects when including families who move across CZs more than once.

We estimate specifications using families who move more than once using the sample in Panel B of Table 1. As in our baseline analysis, we restrict attention to CZs with populations above 250,000 to reduce attenuation bias from sampling error in $\bar{y}_{pcs}$. However, unlike in the baseline specification, we do not impose any distance requirements on the moves and do not require movers to stay in the destination CZ for at least two years, because our objective is to evaluate the stability of the exposure effect estimates when we expand the estimation sample to include a more representative set of moves.

Let $d(j)$ denote the jth destination location and let m_i^j denote the age of the child i when his parents move to destination j. We consider families that move across CZs up to $j = 3$ times (excluding the 3% of children whose families move across CZs more than 3 times between 1996-2012). Let E_i denote the total number of times the family moved when the child was under age 23, $E_i \in \{0, 1, 2, 3\}$. To identify separate effects when children are below vs. above age 23, define e_{ij} as the number of years child i's family spends in place j before he turns 23 and $e_{ij}^{>23}$ as the number of years that his family spends in place j after he turns 23.

For each place j in which the child lives, we define $\Delta_{od(j)ps}^j = \Delta_{od(j)ps} = \bar{y}_{pd(j)s} - \bar{y}_{pos}$ as the difference in the child's predicted outcome based on permanent residents in destination j and the mover's first observed (origin) CZ, o. We then estimate the following specification, which generalizes our baseline linear specification in equation ((7)) to permit multiple moves:

$$
\begin{aligned}
y_i &= \left[\sum_{j=1}^{3} \gamma^j e_{ij} \Delta_{od(j)ps}^j \right] + \sum_{s=1980}^{1988} I(s_i = s) \left(\alpha_s^1 + \alpha_s^2 \bar{y}_{pos} + \sum_{j=1}^{3} \kappa_s^j \Delta_{od(j)ps}^j \right) + \sum_{j=1}^{3} \delta_0^j I\left\{ e_{ij}^{>23} > 0 \right\} \Delta_{od(j)ps} + \sum_{j=1}^{3} \delta^j e_{ij}^{>23} \Delta_{od(j)ps}^j \quad (10) \\
&\quad + \sum_{j=0}^{3} I(E_i = j) \left(\zeta^{2,j} + \zeta^{3,j} p_i + \sum_{j' \leq j} [\zeta^{4,j,j'} e_{ij'} + \zeta^{5,j,j'} e_{ij'} p_i] \right) + \epsilon_{3,i}
\end{aligned}
$$

The parameters of interest in (10) are the coefficients $\{\gamma^j\}$ on the interaction between e_{ij} and $\Delta_{od(j)ps}^j$. We include three sets of additional controls to identify these parameters. First, we control

41

for the outcomes of permanent residents in the origin interacted with the child's birth cohort $(\alpha_s \bar{y}_{pos})$ and allow for a level impact of $\Delta^j_{od(j)ps}$ that may vary by cohort $(\kappa^j_s \Delta_{odps})$. Second, we control for the time spent in each area after age 23 by including interactions of $e^{>23}_{ij}$ with $\Delta_{od(j)ps}$ and a level effect, $I\left\{e^{>23}_{ij} > 0\right\}\Delta_{od(j)ps}$. Finally, we include a set of controls to capture potential disruption effects of multiple moves (Coleman 1988), shown in the second line of (10). These controls include indicators for the number of total moves when the child is under age 23 ($I\left\{E_i = j\right\}\zeta^{2,j}$) and interactions of these indicators with (a) parental income rank ($\zeta^{3,j}p_i$), (b) the number of years spent in each place before age 23 ($\sum_{j' \leq j}\zeta^{4,j,j'}e_{ij'}$), and (c) the interaction of exposure time with parental income ($\sum_{j' \leq j}\zeta^{5,j,j'}e_{ij'}p_i$).

Column 1 of Appendix Table IV presents the results of estimating (10). We estimate a coefficient of $\gamma^1 = 0.042$ (s.e. $= 0.001$) for the first destination, $\gamma^2 = 0.036$ (s.e. $= 0.004$) for the second destination, and $\gamma^3 = 0.032$ (s.e. $= 0.006$) for the third destination. Column 2 presents an estimate of γ from a specification that constrains the estimates $\gamma^j = \gamma$ for all j by replacing $\left[\sum_{j=1}^3 \gamma_j e_{ij}\Delta^j_{od(j)ps}\right]$ in equation (10) with $\left[\gamma \sum_{j=1}^3 e_{ij}\Delta^j_{od(j)ps}\right]$. This specification yields an estimate of $\gamma = 0.042$ (s.e. $= 0.001$).

Columns 3 and 4 add controls for the age of the child at the time of move j interacted with $\Delta_{od(j)ps}$, $\sum_{j=1}^3 m^j_i \alpha^3_j + m^j_i \alpha^4 \Delta_{od(j)ps}$, to the specification in equation (10). We obtain estimates of γ^j ranging between 0.035 and 0.038, and a constrained coefficient of $\gamma = 0.036$ (s.e. $= 0.005$). The standard errors of the estimates rise in these specifications because of the correlation between children's ages at move and exposure times.

Online Appendix C. County-Level Estimates

In this appendix, we present estimates of childhood exposure effects γ using counties as the geographic unit of analysis instead of CZs. To construct these estimates, we first estimate the mean ranks of children of permanent residents in each county ($\bar{y}_{pcs}$) using the sample of children whose parents stay in a single county from 1996-2012 (see Appendix Table I, Panel A for summary statistics for this sample). We estimate $\bar{y}_{pcs}$ using the same method as that in Figure I, exploiting the linearity of the rank-rank relationship within counties. We then regress the outcomes of children who move on $\bar{y}_{pcs}$ interacted with their age at move, using specifications analogous to (7) estimated at the county level. As in our baseline analysis, we focus on children whose families move exactly once across counties and are observed in the destination county for at least two years. We also limit the sample to origin and destination counties with populations (in the 2000 Census) of at

least 250,000.

Appendix Table V presents estimates of γ using variants of the specification in (7). In Column 1, we replicate the baseline specification in Column 1 of Table II at the county level. To match that specification and isolate the effect of changing the geographic unit, we restrict the sample to moves across counties that are in different CZs and that exceed 100 miles based on the distance between the origin and destination ZIP code (see Appendix Table I, Panel B for summary statistics). We obtain an estimate of $\gamma = 0.037$ (s.e. $= 0.003$), slightly lower than our baseline estimate of $\gamma = 0.040$ at the CZ level. Column 2 shows that adding family fixed effects to the specification in Column 1 yields an estimate of $\gamma = 0.033$ (s.e. $= 0.011$), not significantly different from the baseline estimate. Column 3 shows that adding time-varying controls for changes in income and marital status, as in Column 9 of Table II, yields a coefficient of $\gamma = 0.032$ (s.e. $= 0.011$).

In Columns 4-7 of Appendix Table V, we turn to moves across counties *within* rather than across CZs (see Appendix Table I, Panel C for summary statistics). Here, we drop the restriction that moves must exceed 100 miles and include all within-CZ moves between counties that have more than 250,000 residents. Column 4 replicates the specification in Column 1. Here, we obtain an estimate of $\gamma = 0.022$, significantly lower than the estimate of cross-CZ moves of 0.037. This reduction implies that a larger fraction of the variation in $\bar{y}_{pcs}$ across counties within a CZ reflects sorting rather than causal effects. This is intuitive insofar as families are more likely to sort in ways that are correlated with their children's outcomes within a given labor market rather than between labor markets.

Estimates of exposure effects at the county-within-CZ level are less precise, particularly in specifications that include family fixed effects. To maximize precision, in Column 5, we present an estimate of γ from a specification that uses multiple observations per child, including income ranks at all observed ages for ages 24 and above. In this specification, we replace $\bar{y}_{pcs}$ with $\bar{y}_{pcsa}$, where a indexes the age at which the child's income is measured, and use y_{ia} as the dependent variable, where y_{ia} denotes child i's rank in his cohort's income distribution at age a. This specification yields an estimate of $\gamma = 0.027$ (s.e. $= 0.003$). Column 6 adds family by age of income measurement (a) fixed effects to the specification in Column 5. This specification yields a point estimate of $\gamma = 0.029$, which is indistinguishable from the baseline estimate but is also indistinguishable from 0 given the standard error of 0.025. Unfortunately, there is insufficient power to obtain precise estimates from specifications with family fixed effects at the county-within-CZ level, even when attempting to maximize power by using income data from all available years. Column 7 adds time-varying

controls for changes in parental income and marital status, which yields similar estimates to those in Column 6.

In sum, we find exposure effects at the county level similar to those at the CZ level, but with slightly smaller values of the annual exposure effect γ, particularly across counties within a given CZ. This finding implies that a larger fraction of the variation in observed rates of intergenerational mobility across counties within a labor market (CZ) is due to sorting than between labor markets.

Online Appendix D. Specifications for Outcome-Based Placebo Tests

This section presents details of the specifications used for the outcome-based placebo tests described in Section V.D.

Cohort-Specific Convergence. The estimates plotted in triangles in Figure VII are each obtained from a separate regression of a child's mean rank on predictions based on permanent residents' outcomes for a different birth cohort. Each of these regressions replicates the baseline specification in equation (7), replacing Δ_{odps} and $\bar{y}_{pos}$ with Δ_{odpt} and $\bar{y}_{pot}$ for each birth cohort $t = s + a$ for $a \in \{-4, -3, -2, -1, 1, 2, 3, 4\}$ surrounding a given child's birth cohort (s).

The estimates plotted in circles are obtained from a single regression. Starting from the specification in equation (7), we add six variables to the regression for each cohort $t = s + a$ for $a \in \{-4, -3, -2, -1, 1, 2, 3, 4\}$. The first five variables replicate the linear parameterization of the exposure effect for cohort t: $I\{m_i \leq 23\}\Delta_{odpt}$, $I\{m_i > 23\}\Delta_{odpt}$, $I\{m_i \leq 23\}m_i\Delta_{odpt}$, $I\{m_i > 23\}m_i\Delta_{odpt}$, and $\bar{y}_{pot}$. These variables are populated for any child, i, in cohort $s(i)$ for which we observe permanent resident predictions, $\bar{y}_{pct}$.

If we cannot observe Δ_{odpt} or $\bar{y}_{pot}$ (e.g. if $t = 1979$ or $t = 1987$), we set each of these first five variables to zero and we include an indicator, $I_a = I\{\text{cohort } s(i) + a \text{ is missing}\}$ for each $a \in \{-4, -3, -2, -1, 1, 2, 3, 4\}$. For example, since our data covers the 1980-86 cohorts, a child in the 1983 cohort would have non-zero entries for the permanent resident predictions for the three years surrounding 1983, but would have zero entries for the permanent resident predictions along with indicators $I_a = 1$ for $a = -4$ and $a = 4$.

The resulting specification is

$$y_i = \sum_{s=1980}^{1988} I(s_i = s)(\alpha_s^1 + \alpha_s^2 \bar{y}_{pos}) + \sum_{m=9}^{30} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p_i) + \sum_{s=1980}^{1987} \kappa_s^d I(s_i = s)\Delta_{odps}$$
$$+ I(m_i \leq 23)\left(b_0 + (23 - m_i)\gamma\right)\Delta_{odps} + I(m_i > 23)(\delta + (23 - m_i)\delta')\Delta_{odps}$$
$$+ \sum_{a \in \{-4,-3,-2,-1,1,2,3,4\}} I\{m_i \leq 23\}\left(b_0^a + (23 - m_i)\gamma_a\right)\Delta_{odp,s+a} + I(m_i > 23)(\delta_a + (23 - m_i)\delta_a')\Delta_{odp,s+a}$$
$$+ \sum_{a \in \{-4,-3,-2,-1,1,2,3,4\}} \alpha_a \bar{y}_{po,s+a} + \omega_a I_a + \varepsilon_i.$$

Quantiles: Distributional Convergence. For the distributional convergence analysis shown in Columns 1-3 of Table III, we again begin with the baseline linear specification in equation (7). However, we now define the outcome, y_i^{90}, to be an indicator for having income above the 90th percentile. Letting π_{pcs}^{90} denote the predicted probability of having income above the 90th percentile in CZ c for parental income level p and cohort s, we define $\Delta_{odps}^{90} = \pi_{pds}^{90} - \pi_{pos}^{90}$ as the difference in permanent residents' outcomes in the destination relative to the origin.

Column 1 of Table III reports the coefficient γ from a regression analogous to the baseline specification in equation (7), replacing the mean rank outcomes with indicators for income above the 90th percentile:

$$y_i^{90} = \sum_{s=1980}^{1988} I(s_i = s)(\alpha_s^1 + \alpha_s^2 \pi_{pos}^{90}) + \sum_{m=9}^{30} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p_i) + \sum_{s=1980}^{1987} \kappa_s^d I(s_i = s)\Delta_{odps}^{90}$$
$$+ I(m_i \leq 23)\left(b_0 + (23 - m_i)\gamma\right)\Delta_{odps}^{90} + I(m_i > 23)(\delta + (23 - m_i)\delta')\Delta_{odps}^{90} + \varepsilon_i.$$

In column 2, we use the same dependent variable as above (y_i^{90}), but use the mean ranks of permanent residents as the key independent variables in the second line of the equation above:[48]

$$y_i^{90} = \sum_{s=1980}^{1988} I(s_i = s)(\alpha_s^1 + \alpha_s^2 \pi_{pos}^{90}) + \sum_{m=9}^{30} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p_i) + \sum_{s=1980}^{1988} \kappa_s^d I(s_i = s)\Delta_{odps}^{90}$$
$$+ I(m_i \leq 23)\left(b_{mean,0} + (23 - m_i)\gamma_{mean}\right)\Delta_{odps} + I(m_i > 23)(\delta_{mean} + (23 - m_i)\delta'_{mean})\Delta_{odps} + \alpha^3 \bar{y}_{pos} + \varepsilon_i.$$

In column 3, we combine the independent variables used in columns 1 and 2 and estimate the following specification:

$$y_i^{90} = \sum_{s=1980}^{1988} I(s_i = s)(\alpha_s^1 + \alpha_s^{2,p90} \pi_{pos}^{90}) + \sum_{m=9}^{30} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p_i) + \sum_{s=1980}^{1987} \kappa_s^d I(s_i = s)\Delta_{odps}^{90})$$
$$+ I(m_i \leq 23)\left(b_0 + (23 - m_i)\gamma\right)\Delta_{odps}^{90} + I(m_i > 23)(\delta + (23 - m_i)\delta')\Delta_{odps}^{90}$$
$$+ I(m_i \leq 23)\left(b_{mean,0} + (23 - m_i)\gamma_{mean}\right)\Delta_{odps} + I(m_i > 23)(\delta_{mean} + (23 - m_i)\delta'_{mean})\Delta_{odps} + \alpha^3 \bar{y}_{pos} + \varepsilon_i.$$

---

[48]In column 1, we omit the 1988 cohort from the set of cohort interactions ($\kappa_s^d I(s_i = s)\Delta_{odps}^{90}$) since it would be collinear with the Δ_{odps}^{90} terms in the second line. However, in column 2, where we replace Δ_{odps}^{90} with Δ_{odps}, we include the 1988 cohort interaction since this collinearity no longer arises.

Columns 4-6 repeat the specifications in columns 1-3 using an indicator for being unemployed rather than the indicator for having income above the 90th percentile.

Gender-Specific Convergence. For the gender-specific convergence analysis in Table IV, we first construct gender-specific predictions of the mean outcomes of children of permanent residents in each CZ and birth cohort by estimating (1) separately for male and female children, which we denote by $\bar{y}_{pcs}^m$ and $\bar{y}_{pcs}^f$. For a child of gender $g \in \{m, f\}$, we define the difference in predicted outcomes for children of the same gender between the destination and origin as $\Delta_{odps}^g = \bar{y}_{pds}^g - \bar{y}_{pos}^g$.

Column 1 of Table IV replicates the baseline specification in equation (7), replacing Δ_{odps} and $\bar{y}_{pos}$ with Δ_{odps}^g and $\bar{y}_{pos}^g$. Column 2 replicates equation (7), using Δ_{odps}^{-g} and $\bar{y}_{pos}^{-g}$, the predictions based on the opposite gender $(-g)$, as the key independent variables:

$$
\begin{aligned}
y_i &= \sum_{s=1980}^{1988} I(s_i = s)(\alpha_s^1 + \alpha_s^{2g}\bar{y}_{pos}^g) + \sum_{m=9}^{30} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p_i) + \sum_{s=1980}^{1988} I(s_i = s)(\kappa_s^{dg}\Delta_{odps}^g) \\
&\quad + I(m_i \leq 23)\left(b_{other,0} + (23 - m_i)\gamma_{other}\right)\Delta_{odps}^{-g} + I(m_i > 23)(\delta_{other} + (23 - m_i)\delta_{other}')\Delta_{odps}^{-g} + \alpha^3 \bar{y}_{pos}^{-g} + \varepsilon_i.
\end{aligned}
$$

Column 3 includes both the own-gender and other-gender predictions, yielding the following specification:

$$
\begin{aligned}
y_i &= \sum_{s=1980}^{1988} I(s_i = s)(\alpha_s^1 + \alpha_s^{2g}\bar{y}_{pos}^g) + \sum_{m=9}^{30} I(m_i = m)(\zeta_m^1 + \zeta_m^2 p_i) + \sum_{s=1980}^{1987} I(s_i = s)(\kappa_s^{dg}\Delta_{odps}^g) \\
&\quad + I(m_i \leq 23)\left(b_0 + (23 - m_i)\gamma\right)\Delta_{odps}^g + I(m_i > 23)(\delta + (23 - m_i)\delta')\Delta_{odps}^g \\
&\quad + I(m_i \leq 23)\left(b_{other,0} + (23 - m_i)\gamma_{other}\right)\Delta_{odps}^{-g} + I(m_i > 23)(\delta_{other} + (23 - m_i)\delta_{other}')\Delta_{odps}^{-g} + \alpha^3 \bar{y}_{pos}^{-g} + \varepsilon_i,
\end{aligned}
$$

Columns 4-6 replicate Columns 1-3 including family fixed effects. Column 7 replicates the specification in Column 6, restricting the sample to households with at least one child of each gender.

TABLE I

Summary Statistics for CZ Permanent Residents and Movers

Variable	Mean (1)	Std. Dev. (2)	Median (3)	Num. of Obs. (4)
A. Permanent Residents: Families who do not Move Across CZs				
Parent family income	89,909	357,194	61,300	19,499,662
Child family income at 24	24,731	140,200	19,600	19,499,662
Child family income at 26	33,723	161,423	26,100	14,894,662
Child family income at 30	48,912	138,512	35,600	6,081,738
Child individual income at 24	20,331	139,697	17,200	19,499,662
Child married at 26	0.25	0.43	0.00	12,997,702
Child married at 30	0.39	0.49	0.00	6,081,738
Child attends college between 18-23	0.70	0.46	1.00	17,602,702
Child has teen birth (females only)	0.11	0.32	0.00	9,670,225
Child working at age 16	0.41	0.49	0.00	13,417,924
B. Families who Move Between 1-3 Times Across CZs				
Parent family income	90,468	376,413	53,500	4,374,418
Child family income at 24	23,489	57,852	18,100	4,374,418
Child family income at 26	31,658	99,394	23,800	3,276,406
Child family income at 30	46,368	107,380	32,500	1,305,997
Child individual income at 24	19,091	51,689	15,600	4,374,418
Child married at 26	0.25	0.43	0.00	2,867,598
Child married at 30	0.38	0.49	0.00	1,305,997
Child attends college between 18-23	0.66	0.473	1.00	3,965,610
Child has teen birth (females only)	0.13	0.33	0.00	2,169,207
Child working at age 16	0.40	0.49	0.00	3,068,421
C. Primary Analysis Sample: Families who Move Exactly Once Across CZs				
Parent family income	97,064	369,971	58,700	1,553,021
Child family income at 24	23,867	56,564	18,600	1,553,021
Child family income at 26	32,419	108,431	24,500	1,160,278
Child family income at 30	47,882	117,450	33,600	460,457
Child individual income at 24	19,462	48,452	16,000	1,553,021
Child married at 26	0.25	0.43	0.00	1,016,264
Child married at 30	0.38	0.49	0.00	460,457
Child attends college between 18-23	0.69	0.46	1.00	1,409,007
Child has teen birth (females only)	0.11	0.32	0.00	769,717
Child working at age 16	0.39	0.49	0.00	1,092,564

Notes: The table presents summary statistics for the samples used in our CZ-level analyses. The full analysis sample of children consists of all individuals in the tax data who (1) have a valid Social Security Number or Individual Taxpayer Identification Number, (2) were born between 1980-1988, and (3) are U.S. citizens as of 2013. We report summary statistics for three subsets of this sample. Panel A shows statistics for permanent residents -- children whose parents do not move across CZs throughout our sample window (1996-2012) -- who live in CZs with more than 250,000 people based on the 2000 Census. Panel B shows statistics for families who moved once, twice, or three times across CZs with more than 250,000 people from 1996-2012. Panel C shows statistics for our primary analysis sample: children whose families moved exactly once across CZs with more than 250,000 people, are observed in the destination CZ for at least 2 years, and moved at least 100 miles (based on their ZIP codes). Parent family income is the average pre-tax household income from 1996-2000, measured as AGI for tax filers and using information returns for non-filers. Child family income is measured analogously at various ages, while child individual income is defined as the sum of individual W-2 wage earnings, UI benefits, SSDI payments, and half of household self-employment income. Marital status is defined based on the marital status listed on 1040 forms for tax filers; non-filers are coded as single. College attendance is defined as having a 1098-T form filed on one's behalf at any point between the ages of 18 and 23. Teenage birth is defined (for women only) as having a child between the ages of 13 and 19, using data from the Social Security Administration's DM-2 database. We define an indicator for working at age 16 based on having a W-2 form filed on one's behalf at that age. All dollar values are reported in 2012 dollars, deflated using the CPI-U. See Section II for further details on variable and sample definitions.

TABLE II
Childhood Exposure Effect Estimates

Dependent Variable: Child's Income Rank at Age 24

Specification:	Pooled	Age ≤ 23	Age < 18	No Cohort Controls	Individual Income	Child CZ FE	With Family Fixed Effects		
							Baseline	No Cohort Controls	Time-Varying Controls
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Exposure Effect (γ)	0.040	0.040	0.037	0.036	0.041	0.031	0.044	0.031	0.043
	(0.002)	(0.002)	(0.005)	(0.002)	(0.002)	(0.002)	(0.008)	(0.005)	(0.008)
Num of Obs.	1,553,021	1,287,773	687,323	1,553,021	1,553,021	1,473,218	1,553,021	1,553,021	1,553,021

Notes: This table reports estimates of annual childhood exposure effects on children's income ranks at age 24 (γ). The estimates can be interpreted as the impact of spending an additional year of childhood in a CZ where children of permanent residents have 1 percentile point higher income ranks at age 24. Standard errors are shown in parentheses. Each column reports estimates from a regression of a child's income rank at age 24 on the difference between permanent residents' predicted ranks in the destination vs. the origin, interacted with the age of the child at the time of the move (m). We permit separate linear interactions for m≤23 and m>23, and report the coefficient on the interaction for m≤23. Each regression also includes additional controls specified in equation (7). Permanent residents' predicted ranks are constructed using linear regressions of children's ranks on parents' ranks in each CZ and birth cohort, as shown in Figure I. Column 1 reports the estimate of γ from equation (7) using all children in the primary analysis sample of one-time movers, defined in the notes to Table I (Panel C). Columns 2 and 3 restrict the sample to those who move at or before age 23 or 18. In Column 4, we exclude the cohort interactions with the predicted outcomes of permanent residents in the origin and destination location and instead include a single control for the predicted outcomes of permanent residents in the origin. Column 5 replicates Column 1, using individual income ranks (rather than household income ranks) to measure both the child's outcome and the predicted outcomes of permanent residents in the origin and destination. Column 6 adds fixed effects for the child's CZ in 2012 to the specification in Column 1; note that this specification has a slightly smaller number of observations because ZIP code information is missing for some children in 2012 (e.g. a child with no earnings or taxable income). Column 7 adds family fixed effects to the baseline specification in Column 1. Column 8 adds family fixed effects to the specification in Column 4 that does not include cohort-varying intercepts. Column 9 adds controls for changes in parental marital status and income rank in the year before versus after the move, along with their interactions with the age of the child at the time of the move and indicators for moving above and below age 23, to the specification in Column 7.

TABLE III
Childhood Exposure Effect Estimates: Distributional Convergence

Dep. Var:	Upper-Tail: Child in Top Decile			Lower-Tail: Child Not Employed		
	(1)	(2)	(3)	(4)	(5)	(6)
Distributional Prediction	0.043		0.040	0.041		0.043
	(0.002)		(0.003)	(0.003)		(0.003)
Mean Rank Prediction (Placebo)		0.024	0.003		0.018	-0.002
		(0.002)	(0.003)		(0.002)	(0.003)
Num. of Obs.	1,553,021	1,553,021	1,553,021	1,553,021	1,553,021	1,553,021

Notes: This table reports estimates of annual childhood exposure effects (γ) for upper-tail and lower-tail outcomes: being in the top 10% of the cohort-specific income distribution at age 24 or not being employed. Standard errors are shown in parentheses. Column 1 reports estimates from a regression of an indicator for being in the top 10% on the difference between permanent residents' predicted probabilities of being in the upper-tail in the destination vs. the origin, interacted with the age of the child at the time of the move (m). We permit separate linear interactions for $m \leq 23$ and $m > 23$, and report the coefficient on the interaction for $m \leq 23$. The regression also includes additional controls analogous to those in equation (7); see Online Appendix D for details. Column 2 continues to use the indicator for being in the top 10% as the dependent variable, but uses the difference between permanent residents' predicted *mean* ranks in the destination vs. the origin instead of their predicted upper-tail probabilities on the right hand side of the regression. Column 3 includes both the upper-tail (distributional) prediction as well as the mean rank prediction in the same regression. Columns 4-6 replicate Columns 1-3 using an indicator for being unemployed at age 24 as the outcome and the difference between the permanent residents' predicted probabilities of being unemployed in the destination vs. the origin as the key independent variable. Employment is defined as an indicator for having a W-2 filed on one's behalf at age 24. In all columns, the sample consists of all children in the primary analysis sample of one-time movers, defined in the notes to Table I (Panel C).

TABLE IV
Childhood Exposure Effect Estimates: Gender-Specific Convergence

Dependent Variable: Child's Income Rank at Age 24

	No Family Fixed Effects			With Family Fixed Effects			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Own Gender Prediction	0.038 (0.002)		0.030 (0.003)	0.031 (0.006)		0.027 (0.006)	0.030 (0.007)
Other Gender Prediction (Placebo)		0.031 (0.002)	0.010 (0.003)		0.016 (0.005)	0.017 (0.005)	0.009 (0.007)
Sample	Full Sample			Full Sample			2-Gender HH
Num. of Obs.	1,552,898	1,552,898	1,552,898	1,552,898	1,552,898	1,552,898	490,964

Notes: This table reports estimates of annual childhood exposure effects (γ) using gender-specific permanent resident predictions. Standard errors are shown in parentheses. In all columns, the dependent variable is the child's family income rank at age 24. In Columns 1-6, the sample consists of all children in the primary analysis sample of one-time movers, defined in the notes to Table I (Panel C). Column 1 replicates Column 1 of Table II, replacing the predicted outcomes based on all permanent residents in the origin and destination with predictions based on the outcomes of children who have the same gender as the child who moves. Column 2 replicates Column 1, replacing the own-gender predicted outcomes with the predicted outcomes of the opposite gender. Column 3 combines the variables in Columns 1 and 2, including both the own-gender and other-gender (placebo) predictions. Columns 4-6 replicate 1-3 including family fixed effects. Column 7 replicates column 6, restricting the sample of movers to families with at least one child of each gender. Each regression also includes additional controls analogous to those in equation (7); see Online Appendix D for details.

Summary Statistics for County Permanent Residents and Movers

Variable	Mean (1)	Std. Dev. (2)	Median (3)	Num. of Obs. (4)
A. Permanent Residents: Families who do not Move Across Counties				
Parent family income	99,124	426,497	64,500	11,355,962
Child family income at 24	24,449	174,380	19,300	11,355,962
Child family income at 26	33,868	201,955	26,000	8,639,565
Child family income at 30	50,241	152,582	35,800	3,495,902
Child individual income at 24	20,741	175,277	17,300	11,355,962
Child married at 26	0.22	0.41	0.00	7,558,611
Child married at 30	0.36	0.48	0.00	3,495,902
Child attends college between 18-23	0.73	0.45	1.00	10,275,008
Child has teen birth (females only)	0.10	0.31	0.00	5,636,590
Child working at age 16	0.40	0.49	0.00	7,860,060
B. Families who Move Exactly Once Across Counties in Different CZs				
Parent family income	112,141	456,094	60,600	595,244
Child family income at 24	23,931	69,737	18,400	595,244
Child family income at 26	32,774	139,408	24,500	440,653
Child family income at 30	49,424	103,587	33,800	172,614
Child individual income at 24	19,987	56,835	16,200	595,244
Child married at 26	0.22	0.42	0.00	387,036
Child married at 30	0.36	0.48	0.00	172,614
Child attends college between 18-23	0.72	0.45	1.00	541,627
Child has teen birth (females only)	0.11	0.31	0.00	296,024
Child working at age 16	0.37	0.48	0.00	422,630
C. Families who Move Exactly Once Across Counties Within the Same CZ				
Parent family income	93,059	450,928	57,000	617,502
Child family income at 24	24,006	68,559	18,700	617,502
Child family income at 26	32,993	75,520	25,200	457,140
Child family income at 30	49,974	108,248	35,000	179,856
Child individual income at 24	20,679	61,101	17,000	617,502
Child married at 26	0.19	0.40	0.00	401,484
Child married at 30	0.34	0.47	0.00	179,856
Child attends college between 18-23	0.72	0.45	1.00	561,846
Child has teen birth (females only)	0.10	0.30	0.00	309,108
Child working at age 16	0.38	0.49	0.00	437,646

Notes: The table presents summary statistics for the samples used in our county-level analysis. We present statistics for three subsamples of the full analysis sample described in the notes to Table I. Panel A shows statistics for permanent residents -- children whose parents do not move across counties throughout our sample window (1996-2012) -- who live in counties with more than 250,000 people based on the 2000 Census. Panel B shows statistics for families who moved exactly once across counties in two different CZs between 1996-2012, each of which had more than 250,000 people, and who moved at least 100 miles (based on their ZIP codes). Panel C shows statistics for children whose families moved exactly once across counties of more than 250,000 people within the same CZ. The sample of families who move is restricted to those who are observed in the destination county for at least 2 years in both Panels B and C. See notes to Table I for variable definitions.

APPENDIX TABLE II
Exposure Effect Estimates: Sensitivity to Population and Distance Restrictions

	Baseline Specification (1)	No Distance Restriction			Distance > 100 Miles (Baseline)			Distance > 200 Miles		
		Pop > 50K (2)	Pop > 250K (3)	Pop > 500K (4)	Pop > 50K (5)	Pop > 250K (6)	Pop > 500K (7)	Pop > 50K (8)	Pop > 250K (9)	Pop > 500K (10)
Exposure Effect (γ)	0.040 (0.002)	0.033 (0.001)	0.036 (0.002)	0.038 (0.002)	0.036 (0.001)	0.040 (0.002)	0.041 (0.002)	0.038 (0.002)	0.040 (0.002)	0.042 (0.002)
Num of Obs.	1,553,021	3,002,272	2,143,525	1,562,615	2,062,277	1,553,021	1,165,153	1,655,105	1,288,816	991,657

Notes: This table reports estimates of annual childhood exposure effects (γ) with alternative sample restrictions. Standard errors are shown in parentheses. Column 1 replicates the baseline specification in Column 1 of Table II, which restricts the sample to origin and destination CZs with more than 250,000 people based on the 2000 Census and requires the distance between the origin and destination ZIP codes to be more than 100 miles. Columns 2-10 replicate the specification in Column 1, varying these distance and population restrictions. See notes to Table II for further details.

APPENDIX TABLE III
Heterogeneity in Exposure Effects Across Subgroups

	Baseline	By Parental Income		Positive vs. Negative Moves	
		Above Median Income	Below Median Income	Moves to Better CZs	Moves to Worse CZs
	(1)	(2)	(3)	(4)	(5)
Exposure Effect (γ)	0.040	0.047	0.031	0.030	0.040
	(0.002)	(0.003)	(0.003)	(0.004)	(0.004)
Num of Obs.	1,553,021	803,189	749,832	783,936	769,085

Notes: This table reports estimates of annual childhood exposure effects (γ) for various subsamples. Standard errors are shown in parentheses. Column 1 replicates the baseline specification in Column 1 of Table II. Columns 2 and 3 replicate Column 1, restricting the sample of movers to those with above-median or below-median income (parent rank above or below 0.5). Note there are more observations in the above-median income specification because the higher-income families move at higher rates. Columns 4 and 5 replicate Column 1, restricting the sample to moves in which the outcomes of permanent residents in the destination CZ are higher (positive moves) or lower (negative moves) than those of permanent residents in the origin CZ at the relevant parent income level. See notes to Table II for further details.

APPENDIX TABLE IV
Exposure Effect Estimates Using Families who Move Multiple Times

	Generalized Baseline Specification		Critical Age versus Exposure Effects	
	Separate Coefficients	Constrained Coefficient	Separate Coefficients	Constrained Coefficient
	(1)	(2)		
1st Destination Exposure Effect	0.042 (0.001)		0.036 (0.000)	
2nd Destination Exposure Effect	0.036 (0.004)		0.038 (0.000)	
3rd Destination Exposure Effect	0.032 (0.006)		0.035 (0.000)	
Average Exposure Effect		0.042 (0.001)		0.036 (0.005)
Controls for Age at Move x $\Delta_{od(j)ps}$			X	X
Num of Obs.	4,374,418	4,374,418	4,374,418	4,374,418

Notes: This table reports estimates of annual childhood exposure effects (γ) analogous to those reported in Table II using the expanded sample of movers who move up to 3 times across CZs, defined in the notes to Table I (Panel B). Standard errors are shown in parentheses. Column 1 presents estimates of the exposure effect of the 1st, 2nd, and 3rd destination by permitting separate exposure effect coefficients for each move using the specification in Appendix B. Column 2 presents an estimate of the exposure effect restricting the coefficient to be the same across all moves. Columns 3 and 4 replicate 1 and 2, adding controls for the child's age at the time of move to destination j interacted with the change in permanent resident outcomes ($\Delta_{od(j)ps}$) as described in equation (10). See Appendix C for further details on these specifications.

Annual Exposure Effect Estimates at the County Level

Specification:	Across CZ Moves			Within CZ Moves			
	Baseline	Family FE	Family FE & Time-Varying Controls	Age 24	Age ≥ 24	Family FE	Family FE & Time-Varying Controls
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Exposure Effect (γ)	0.037	0.033	0.032	0.022	0.027	0.029	0.029
	(0.003)	(0.011)	(0.011)	(0.003)	(0.003)	(0.025)	(0.025)
Num of Obs.	595,244	595,244	595,244	617,502	2,900,311	2,900,311	2,900,311

Notes: This table reports exposure effect estimates at the county level analogous to those presented in Table II, using the samples of one-time county movers defined in the notes to Appendix Table I (Panels B and C). Standard errors are shown in parentheses. Column 1 presents estimates at the county level using a specification analogous to that in Column 1 of Table II, replacing the CZ-level predictions with county-level predictions based on permanent residents. In this specification, we restrict the sample to moves across counties that are in different CZs and require that both the origin and destination counties have populations exceeding 250,000 (based on the 2000 Census). We also require that the move exceed 100 miles based on the distance between the origin and destination ZIP code. Column 2 adds family fixed effects to the specification in Column 1, as in Column 7 in Table II. Column 3 adds time-varying controls to the specification in Column 2, including changes in parent income in the year before versus after the move and indicators for transitions in marital status, all interacted with the child's age at the time of the move, as in Column 9 of Table II. Columns 4-7 consider the set of individuals who move across counties within CZs, again requiring both origin and destination counties to have populations above 250,000 but dropping the distance restriction. Column 4 replicates the specification in Column 1 on this sample. Column 5 presents an estimate from a specification that uses multiple observations per child, including income ranks at all observed ages for ages 24 and above. Column 6 adds family-by-age of outcome fixed effects to the specification in Column 5. Column 7 adds time-varying controls for changes in parental income and marital status to the specification in Column 6, as in Column 3. Standard errors are clustered by child in Columns 4-7 to account for the multiple observations per child.

APPENDIX TABLE VI
Family Fixed Effect Estimates: Sensitivity to Cohort Controls and Population Restrictions

	Baseline (Pop > 250K)		Pop > 50K	
	No Fam FE	Fam FE	No Fam FE	Fam FE
	(1)	(2)	(3)	(4)
A. With Cohort-Varying Intercept				
Exposure Effect (γ)	0.040	0.044	0.036	0.036
	(0.002)	(0.008)	(0.001)	(0.006)
Ratio of Fam FE to No Fam FE Ests.	1.10		0.98	
B. Without Cohort-Varying Intercept				
Exposure Effect (γ)	0.036	0.031	0.032	0.023
	(0.002)	(0.005)	(0.001)	(0.003)
Ratio of Fam FE to No Fam FE Ests.	0.86		0.71	
Num of Obs.	1,553,021	1,553,021	2,062,277	2,062,277

Notes: This table reports estimates of annual childhood exposure effects (γ) using alternative specifications. Standard errors are shown in parentheses. Column 1 of Panel A replicates the baseline specification shown in Column 1 of Table 2. Column 2 of Panel A replicates the baseline specification with family fixed effects shown in Column 7 of Table 2. Columns 1 and 2 of Panel B replicate these specifications excluding the cohort interactions in equation (7), replicating the specifications in Columns 4 and 8 of Table 2. Columns 3 and 4 replicate the specifications in Columns 1 and 2, extending the sample to include moves in which the origin and destination CZ has at least 50,000 residents (instead of 250,000 residents as required in our baseline specifications). See the notes to Table II for further details.

APPENDIX TABLE VII

Correlations of Permanent Resident Outcomes Across CZs

	Fam. Inc. Rank at Age 24	Fam. Inc. Rank at Age 26	Fam. Inc. Rank at Age 30	Indiv. Inc. Rank at Age 26	Rank at Age 24 in Top 10%	Emp. at Age 24	Females Ind Inc. Rank at Age 24	Males Ind Inc. Rank at Age 24	College Attend Age 18-23	Married at Age 26	Teen Emp. at Age 16
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
A. Below-Median Income Parents (p = 25)											
Child Family Income Rank at Age 24	1.00										
Child Family Income Rank at Age 26	0.97	1.00									
Child Family Income Rank at Age 30	0.87	0.96	1.00								
Child Indiv. Inc. Rank at Age 26	0.67	0.77	0.86	1.00							
Child Rank at Age 24 in Top 10%	0.85	0.83	0.71	0.39	1.00						
Child Employed at Age 24	0.71	0.67	0.59	0.44	0.59	1.00					
Child Rank at Age 24 for Females	0.93	0.97	0.95	0.83	0.74	0.59	1.00				
Child Rank at Age 24 for Males	0.96	0.98	0.92	0.69	0.87	0.71	0.91	1.00			
College Attendance Age 18-23	0.15	0.29	0.46	0.63	-0.08	0.05	0.40	0.18	1.00		
Married at Age 26	0.64	0.55	0.38	-0.08	0.76	0.49	0.44	0.64	-0.35	1.00	
Teen Employment at Age 16	0.36	0.28	0.19	0.17	0.30	0.57	0.26	0.28	-0.09	0.22	1.00
B. Above-Median Income Parents (p = 75)											
Child Family Income Rank at Age 24	1.00										
Child Family Income Rank at Age 26	0.94	1.00									
Child Family Income Rank at Age 30	0.73	0.87	1.00								
Child Indiv. Inc. Rank at Age 26	0.48	0.56	0.79	1.00							
Child Rank at Age 24 in Top 10%	0.79	0.81	0.53	0.09	1.00						
Child Employed at Age 24	0.67	0.64	0.59	0.50	0.30	1.00					
Child Rank at Age 24 for Females	0.91	0.98	0.83	0.50	0.79	0.60	1.00				
Child Rank at Age 24 for Males	0.93	0.98	0.87	0.58	0.79	0.66	0.92	1.00			
College Attendance Age 18-23	-0.21	-0.06	0.24	0.30	-0.30	0.00	-0.04	-0.08	1.00		
Married at Age 26	0.53	0.51	0.11	-0.42	0.75	0.21	0.55	0.46	-0.37	1.00	
Teen Employment at Age 16	0.57	0.51	0.46	0.44	0.22	0.80	0.49	0.52	-0.03	0.11	1.00

Notes: This table presents correlations of outcomes for children of permanent residents across CZs. Teen employment is defined using only cohorts 1983-1986. All correlations are weighted by population in the 2000 Census. See notes to Table I for the definition of the sample of permanent residents and definitions of the variables used in this table.

FIGURE I: Mean Child Income Rank vs. Parent Income Rank for Children Raised in Chicago
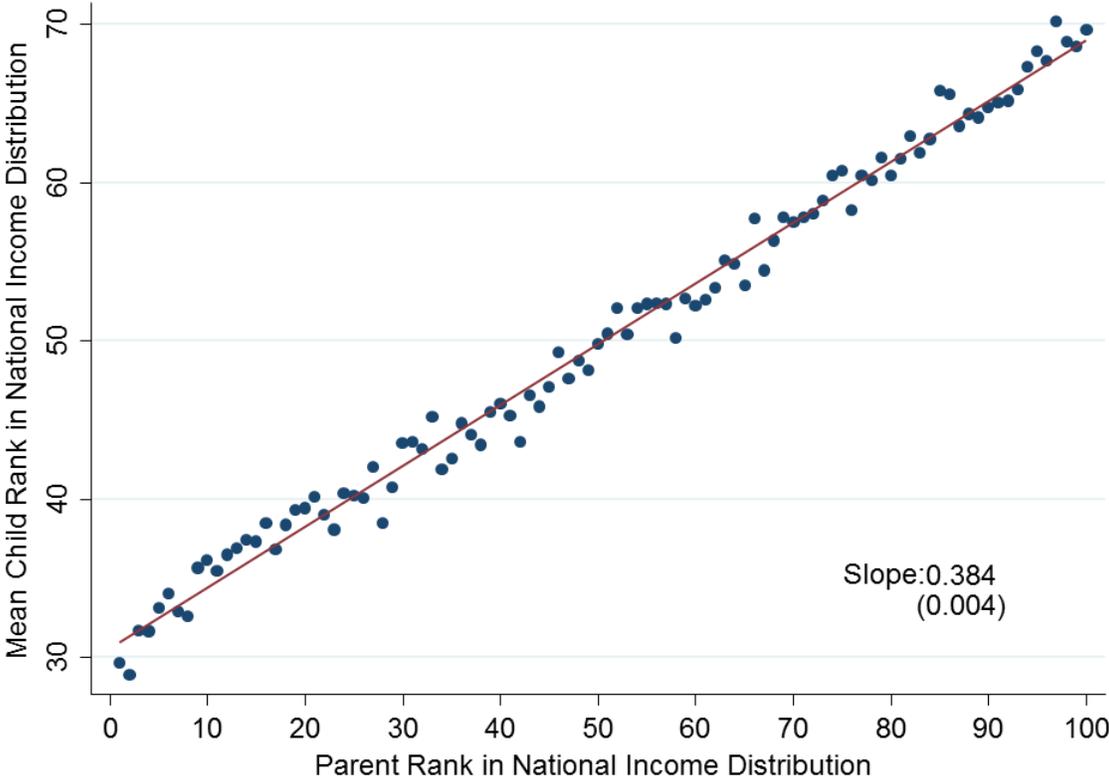
Notes: This figure presents a binned scatter plot of the relationship between children's income ranks and parent income ranks for children raised in Chicago. The points on the figure plot the mean rank of children within each parental income percentile bin. The best-fit line is estimated using an OLS regression on the underlying micro data. The figure also reports the slope of the best-fit line (the rank-rank slope), along with the standard error of the estimate (in parentheses). The sample includes all children in the 1980 birth cohort in our analysis sample whose parents were permanent residents of the Chicago commuting zone during the sample period (1996-2012). Children's incomes are measured at the household (i.e., family) level at age 30; parents' incomes are defined as mean family income from 1996-2000. Children are assigned ranks based on their incomes relative to all other children in their birth cohort. Parents' are assigned ranks based on their incomes relative to other parents of children in the same birth cohort.

FIGURE II: Mean Income Ranks for Children of Permanent Residents

A. For Children with Parents at the 25th Percentile
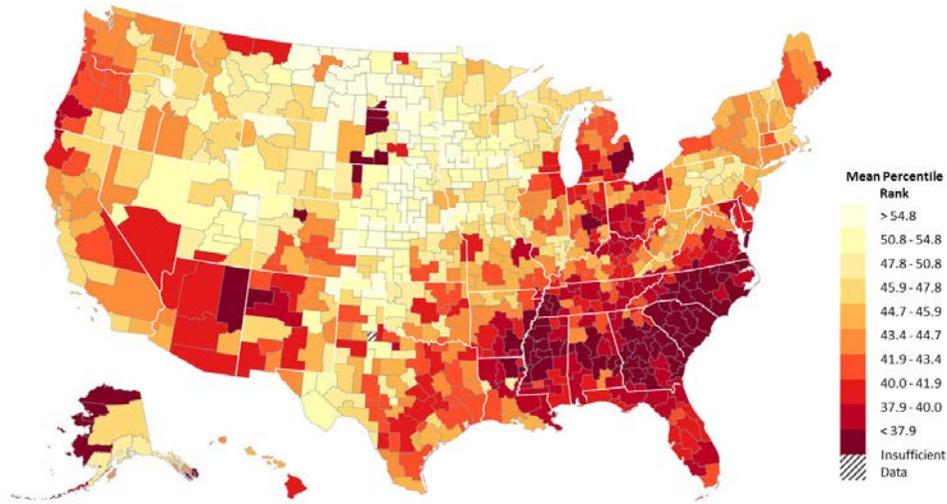
B. For Children with Parents at the 75th Percentile
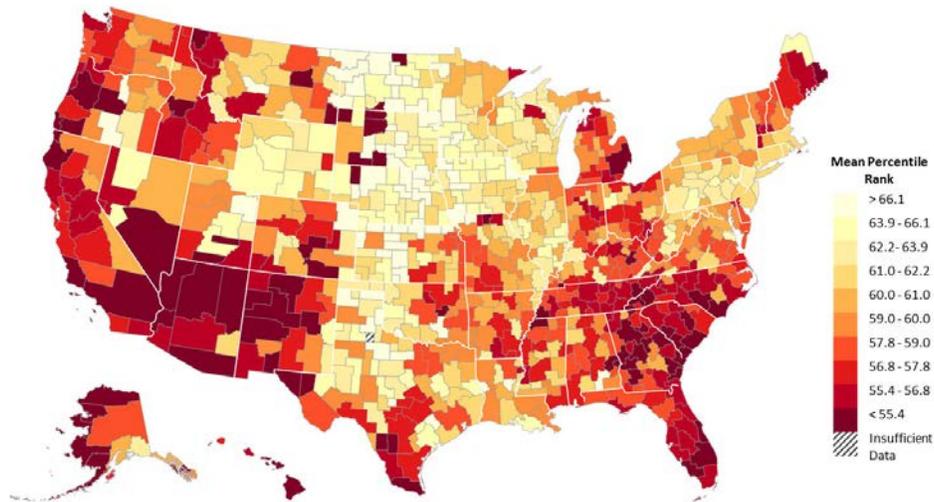
Notes: These maps plot children's mean percentile ranks at age 30 conditional on having parents at the 25th percentile (Panel A) and 75th percentile (Panel B). The maps are constructed by grouping CZs into ten deciles and shading the areas so that lighter colors correspond to higher outcomes for children. Areas with fewer than 10 children, for which we have insufficient data to estimate outcomes, are shaded with the striped pattern. The sample includes all children in the 1980 birth cohort in our analysis sample whose parents are permanent residents (i.e., whose parents do not move across CZs between 1996-2012). To construct these estimates, we first regress children's family income ranks on a constant and their parents' family income ranks separately for each CZ and birth cohort. We then define the predicted income rank for children with parents at percentile p in CZ c in birth cohort s ($\bar{y}_{pcs}$) as the intercept + p times the slope of this regression. Panel A reports the predicted child rank for parents at $p = 25$, which corresponds to an annual household income of \$30,000. Similarly, Panel B reports the predicted child rank for parents at $p = 75$, which corresponds to an annual household income of \$97,000. See notes to Figure I for details on definitions of parent and child income ranks.
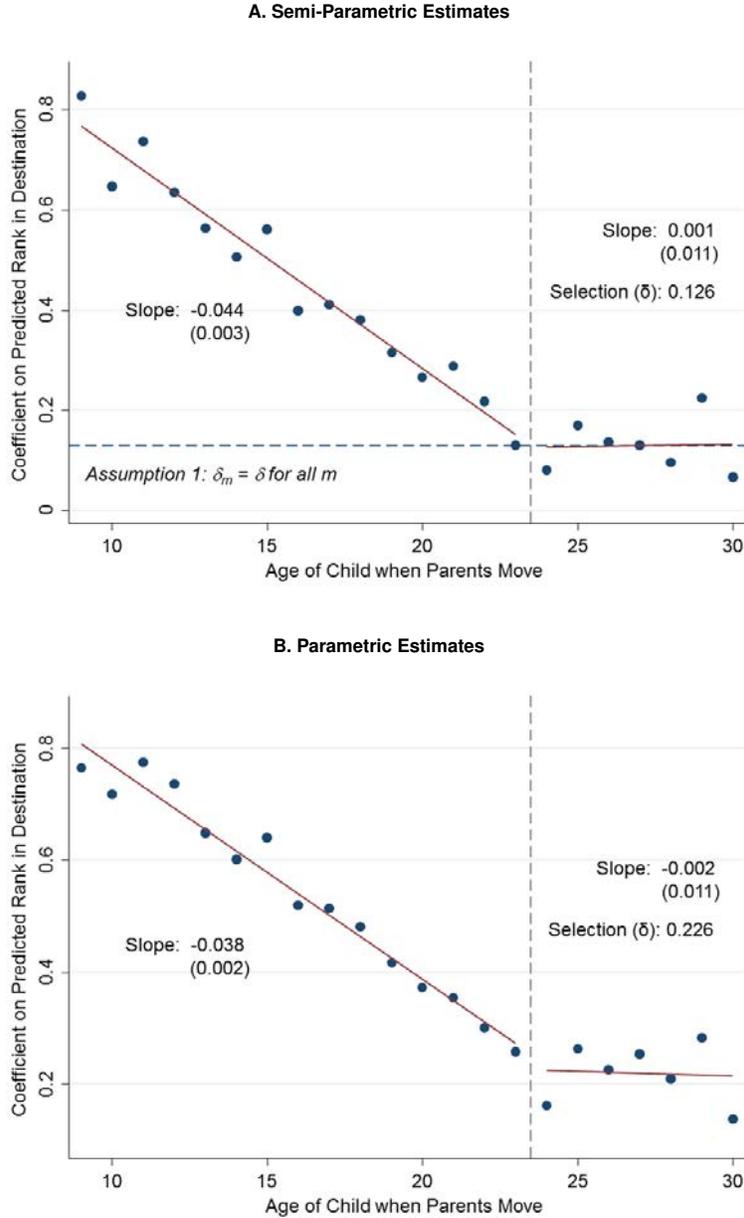
FIGURE III

Movers' Outcomes vs. Predicted Outcomes Based on Permanent Residents in Destination
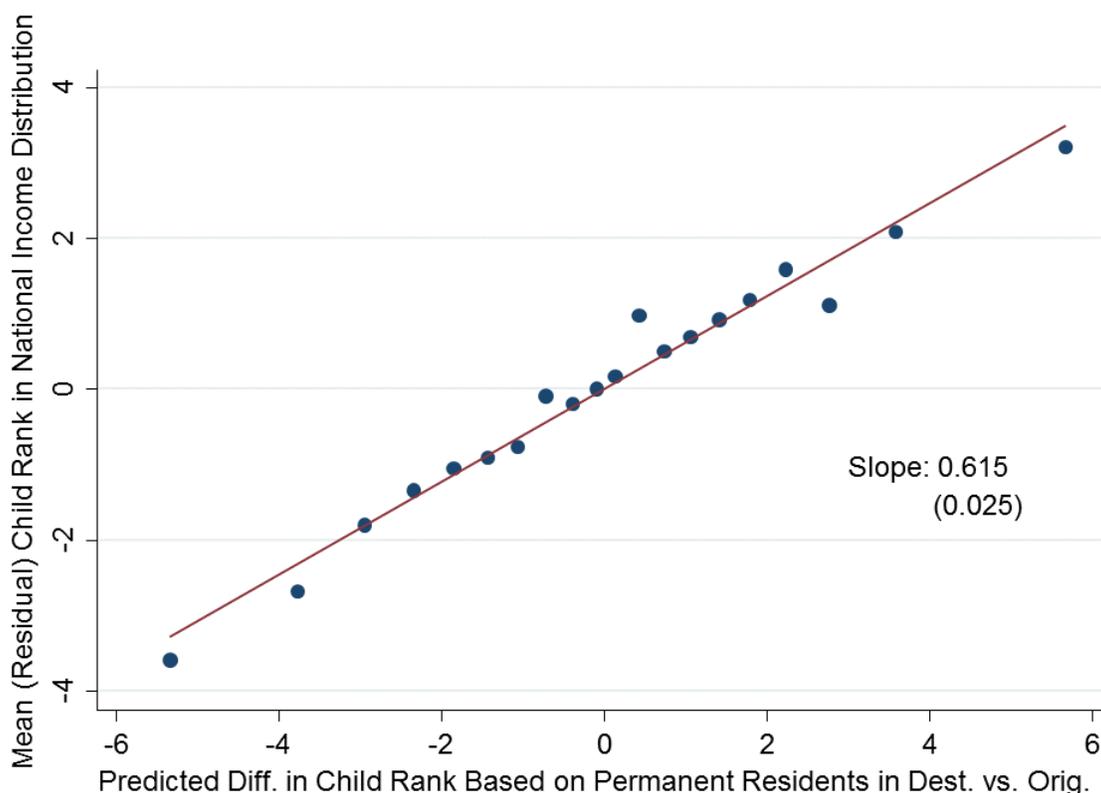
Notes: This figure presents a binned scatter plot depicting the relationship between the income ranks of children who moved to a different CZ at age 13 and the differences in the outcomes of permanent residents in the destination vs. origin CZ. The sample includes all children in the 1980-88 birth cohorts whose parents moved when the child was 13 years old and moved only once between 1996-2012. Children's family income ranks y_i are measured at age 24. Permanent residents' predicted ranks for each parent income percentile p, CZ c, and birth cohort s ($\bar{y}_{pcs}$) are constructed using the methodology described in the notes to Figure I. To construct the figure, we first demean both y_i and $\Delta_{odps} = \bar{y}_{pds} - \bar{y}_{pos}$ within the parent decile (q) by origin (o) by birth cohort (s) cells in the sample of movers at age $m = 13$ to construct residuals: $y_i^r = y_i - E[y_i|q, o, s]$ and $\Delta_{odps}^r = \Delta_{odps} - E[\Delta_{odps}|q, o, s]$. We then divide the Δ_{odps}^r residuals into twenty equal-size groups (ventiles) and plot the mean value of y_i^r vs. the mean value of Δ_{odps}^r in each bin. The slope of the best-fit line, which corresponds to b_{13} in equation (4), is estimated using an OLS regression on the underlying microdata, with standard error in parentheses.

FIGURE IV: Childhood Exposure Effects on Income Ranks in Adulthood

A. Semi-Parametric Estimates

Slope: 0.001
(0.011)

Selection (δ): 0.126

Slope: -0.044
(0.003)

Assumption 1: $\delta_m = \delta$ for all m

(y-axis) Coefficient on Predicted Rank in Destination

(x-axis) Age of Child when Parents Move

B. Parametric Estimates

Slope: -0.002
(0.011)

Selection (δ): 0.226

Slope: -0.038
(0.002)

(y-axis) Coefficient on Predicted Rank in Destination

(x-axis) Age of Child when Parents Move

Notes: Panel A plots estimates of the coefficients $\{b_m\}$ vs. the child's age when the parents move (m) using the semi-parametric specification in equation (5), measuring children's incomes at age 24. The sample includes all children in the primary analysis sample whose parents moved exactly once between 1996-2012. The $\{b_m\}$ coefficients can be interpreted as the effect of moving to an area where permanent resident outcomes are 1 percentile higher at age m. They are estimated by regressing the child's income rank in adulthood y_i on $\Delta_{odps} = \bar{y}_{pds} - \bar{y}_{pos}$, the difference between permanent residents' predicted ranks in the destination vs. the origin, interacted with each age of the child at the time of the move m. We include origin CZ by parent income decile by birth cohort by age at move fixed effects when estimating this specification. Panel B plots estimates from the parametric specification in equation (6), measuring children's incomes at age 24. This specification replicates the specification used in Panel A, replacing the fixed effects with indicators for the child's age at the time of the move interacted with parent income rank and predicted outcomes for permanent residents in the origin interacted with birth cohort fixed effects. The dashed vertical lines separate the data into two groups: age at move $m \leq 23$ and $m > 23$. Best-fit lines are estimated using unweighted OLS regressions of the $\{b_m\}$ coefficients on m separately for $m \leq 23$ and $m > 23$. The slopes of these regression lines are reported along with standard errors (in parentheses) on the left side of each panel for $m \leq 23$ and on the right side for $m > 23$. The magnitudes of the slopes for $m \leq 23$ represent estimates of annual childhood exposure effects. The parameter δ is defined as the mean value of the b_m estimates for $m > 23$; this parameter represents a selection effect because moves after age 24 cannot affect income measured at age 24. In Panel A, the dashed horizontal line shows the value of the selection effect δ; the identification assumption underlying the analysis is that the selection effect δ does not vary with the child's age at move m.

FIGURE V: Childhood Exposure Effects on Income Ranks with Additional Controls

A. With Family Fixed Effects

B. With Family Fixed Effects and Time-Varying Controls
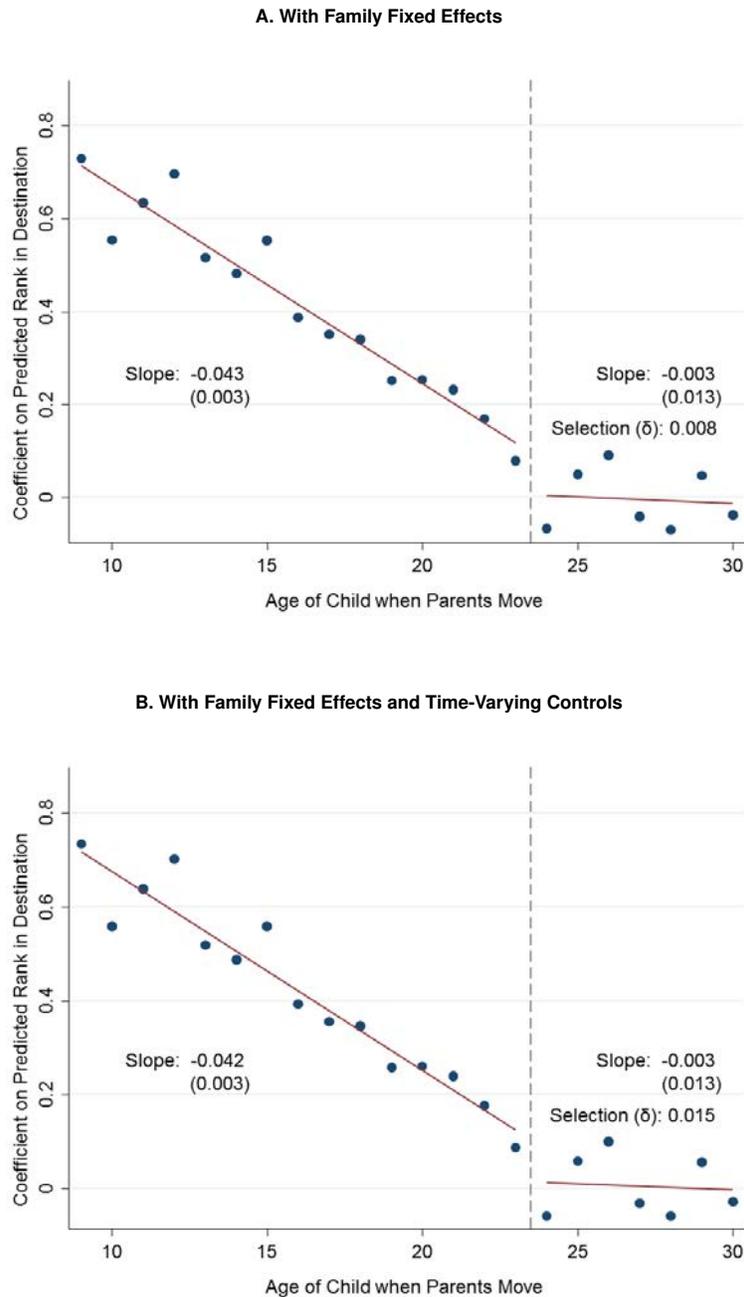

Notes: This figure replicates Figure IVb using specifications analogous to equation (6) that include family fixed effects (Panel A) and both family fixed effects and controls for changes in marital status and parental income around the time of the move (Panel B). To control for changes in parental income, we first construct parental income ranks by child's birth cohort and calendar year. We then interact the differences in parental ranks in the year before versus after the move with the child's age at the time of the move along with interactions with indicators for moving above versus below age 23. To control for changes in marital status, we construct indicators for being always married, getting divorced, or being never married in the year before the move and the year after the move (getting married is the omitted category). We then interact these marital status indicators with the child's age at the time of the move along with interactions with indicators for moving above versus below age 23. See notes to Figure IV for additional details on the construction of the figure.

FIGURE VI: Exposure Effect Estimates Using Displacement Shocks
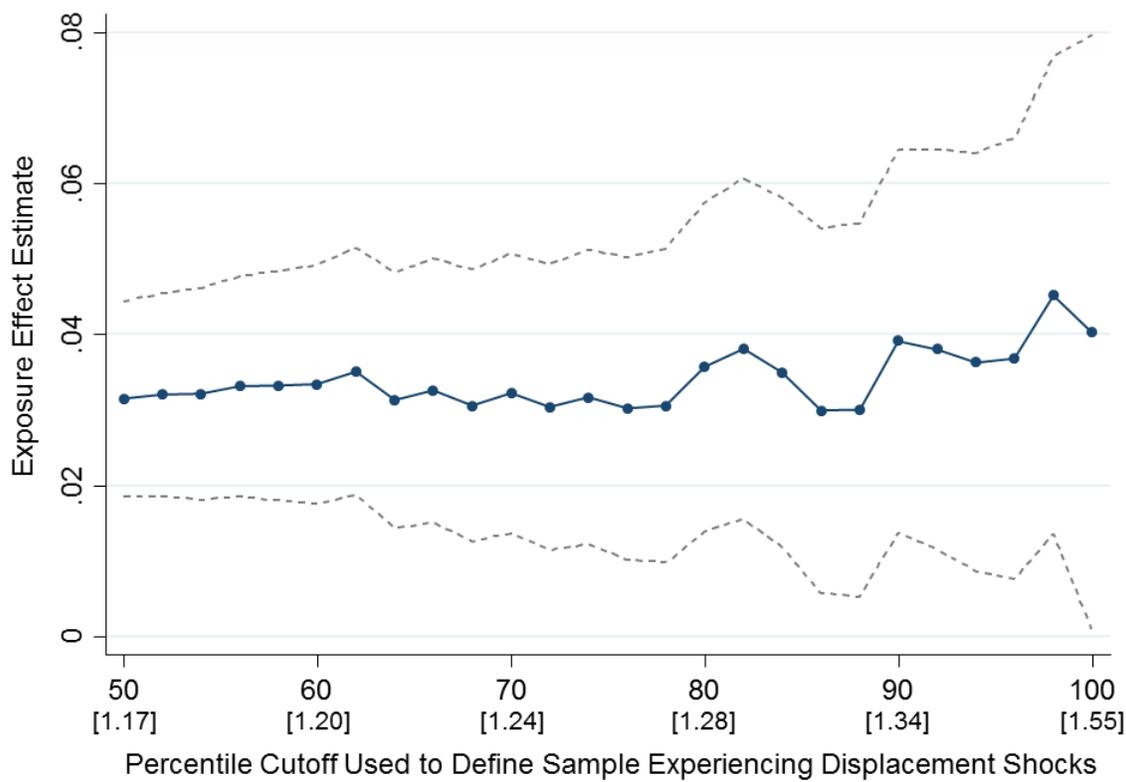
Notes: This figure presents estimates of annual childhood exposure effects (γ) for the subset of areas that experience displacement shocks, defined as ZIP code by year cells that have large outflows in the number of residents. We measure outflows by defining K_{zt} as the number of families who leave ZIP code z in year t in our one-time movers sample and $\bar{K}_z$ as mean outflows between 1996 and 2012. We define the shock to outflows in year t in ZIP z as $k_{zt} = K_{zt}/\bar{K}_z$. We then take ZIP-year cells with above-median outflows ($k_{zt} > 1.17$) and divide them into 25 population-weighted bins based on the size of the shock k_{zt}. For each subset of observations with values of k_{zt} above the percentile threshold listed on the x axis, we estimate γ using equation (7), instrumenting for the change in predicted outcomes based on permanent residents Δ_{odps} with the average change in predicted outcomes for movers from the origin ZIP, $E[\Delta_{odps}|q,z]$. We define $E[\Delta_{odps}|q,z]$ as the mean value of Δ_{odps} for each parental income decile q, pooling across all years and all movers out of ZIP code z. The figure plots the resulting estimates of γ vs. the percentile threshold cutoff for the sample. The dashed lines show 95% confidence intervals for the estimates. The mean value of the outflow shock k_{zt} used in each subsample is shown in brackets below the percentile thresholds.

## FIGURE VII: Exposure Effect Estimates Based on Cross-Cohort Variation
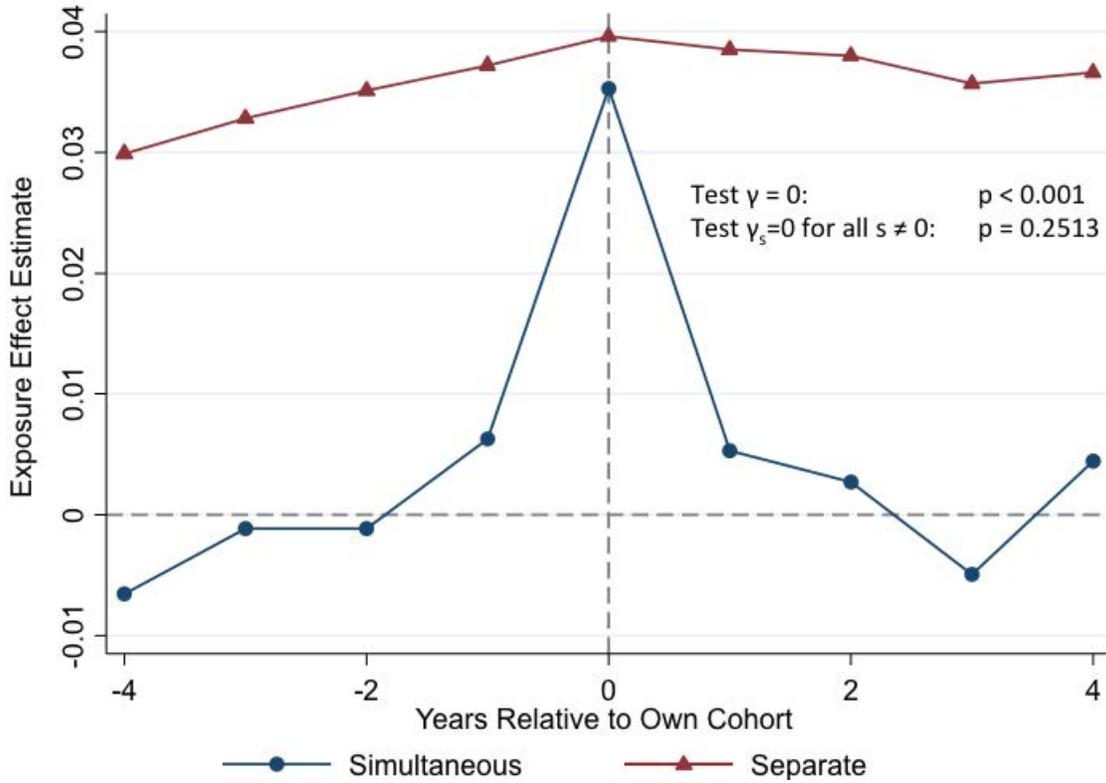
Notes: This figure presents estimates of the annual childhood exposure effect on children's income ranks in adulthood using permanent resident predictions for the child's own birth cohort and surrounding "placebo" birth cohorts. The series in triangles plots estimates of the exposure effect γ_t from nine separate regressions analogous to that in equation (7), using permanent resident predictions from cohort $s + t$ (where t ranges between -4 and 4) as the key independent variables and the outcomes of children in birth cohort s as the dependent variable. By construction, the exposure effect estimate for $t = 0$ (highlighted by the dashed vertical line) corresponds to the baseline estimate of $\gamma = 0.040$ in Column 1 of Table II. The series in circles plots estimates from a single multivariable regression that simultaneously includes all nine permanent resident predictions $t = -4, ..., 4$ and plots the coefficient on the interaction of the child's age at the time of the move m with $\Delta_{odp,s+t}$, the difference between permanent residents' predicted ranks in the destination vs. the origin in cohort $s + t$. The figure also reports p-values from two hypothesis tests: the hypothesis that γ (the estimate using the actual cohort, $t = 0$) equals zero in the simultaneous specification and the hypothesis that all other coefficients γ_{s+t} excluding the own-cohort coefficient are equal to zero. See Online Appendix D for further details on the regression specifications.

FIGURE VIII: Exposure Effects on College Attendance and Marriage

A. College Attendance (Age 18-23)

B. Marriage (Age 26)
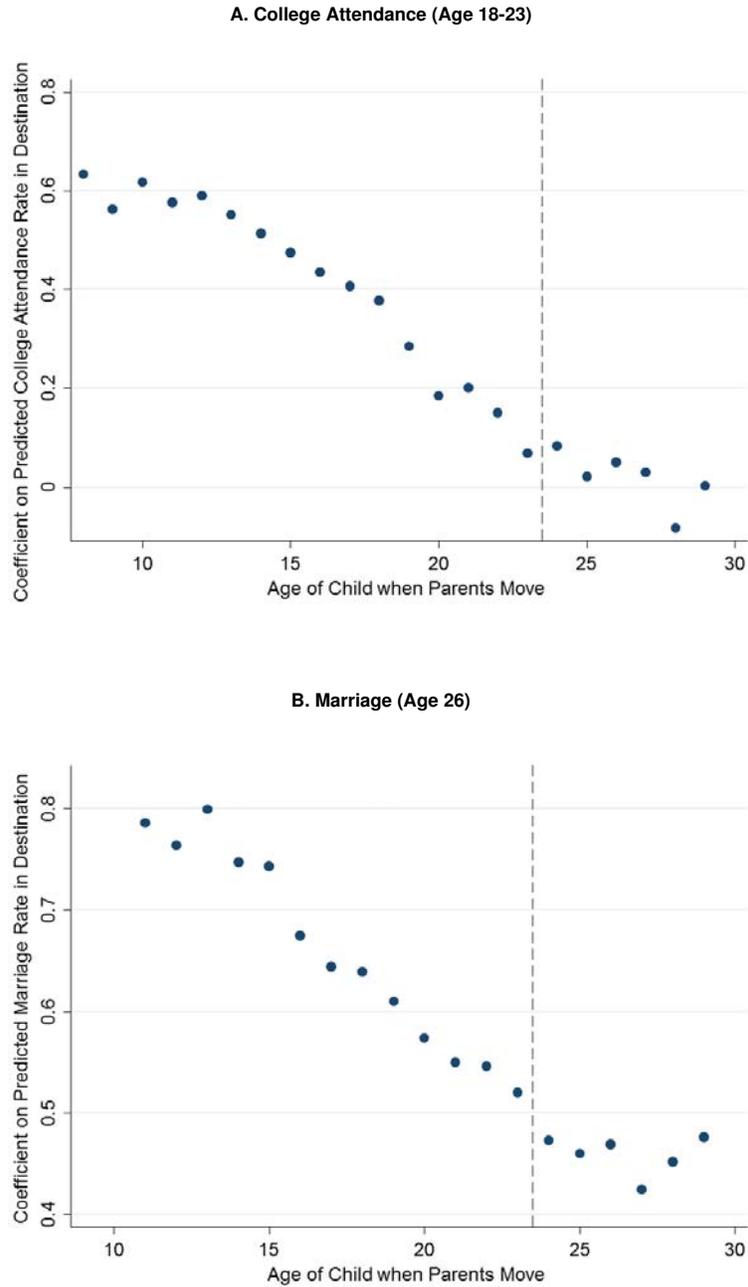
Notes: This figure plots exposure effects for college and marriage outcomes using an approach analogous to that in Figure IVb. In Panel A, we replicate the specification in equation (6), using an indicator for college attendance at any age between 18-23 as the dependent variable instead of the child's income rank and replacing the key independent variable $\Delta_{odps} = \bar{y}_{pds} - \bar{y}_{pos}$ with the difference between permanent residents' college attendance rates in the destination vs. the origin. The coefficients that are plotted can therefore be interpreted as the effect of moving to an area where permanent residents' college attendance rates are 1 percentage point higher at age m. We require that the child be observed between ages 18 and 23 to define college attendance; because we observe college attendance in years 1999-2012, we obtain estimates for children who move between the ages of 8 and 29. In Panel B, we replicate the baseline specification in equation (6), replacing the child's outcomes with an indicator for being married at age 26 and replacing $\Delta_{odps} = \bar{y}_{pds} - \bar{y}_{pos}$ with the difference between permanent residents' marriage rates in the destination vs. the origin. See notes to Figure IV for further details on the construction of this figure.

FIGURE IX: Exposure Effects on Teenage Birth and Employment

A. Teenage Birth

B. Employment at Age 16

**C. Employment at Age 17**

D. Employment at Age 18
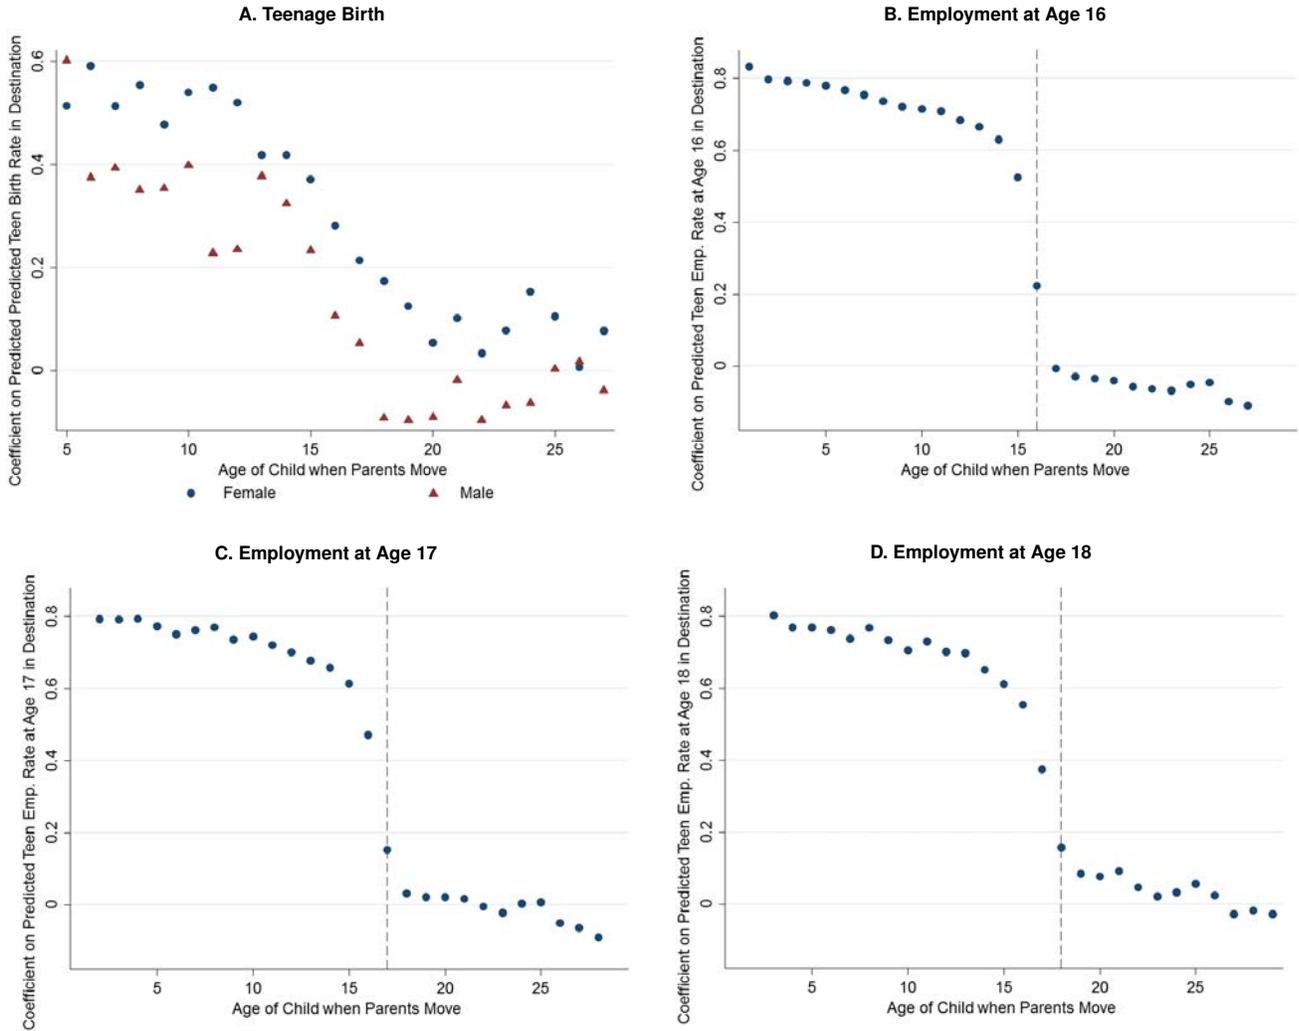
Notes: This figure plots exposure effects for teenagers' outcomes using an approach analogous to that in Figure IVb. Panel A replicates the parametric specification in equation (6), using teenage birth as the dependent variable and replacing the key independent variable $\Delta_{odps} = \bar{y}_{pds} - \bar{y}_{pos}$ with the difference between permanent residents' teen birth rates in the destination vs. the origin. We define teenage birth as having a child between the ages of 13 and 19 using data from the Social Security Administration's DM-2 database, and estimate separate specifications for males and females who have a child. Panels B-D replicate the parametric specification in equation (6), using an indicator for working at ages 16, 17, or 18 (based on having a W-2) as the dependent variable and replacing the key independent variable $\Delta_{odps} = \bar{y}_{pds} - \bar{y}_{pos}$ with the difference between permanent residents' teen employment rates in the destination vs. the origin at the corresponding age. The coefficients that are plotted can therefore be interpreted as the effect of moving at age m to an area where permanent residents' teen employment rates are 1 percentage point higher at ages 16, 17, or 18. The age at which teen employment is measured is shown by the vertical dashed line in each panel; since moves after the age at which employment is measured cannot have a causal effect, the coefficients to the right of the dashed lines reflect selection.

ONLINE APPENDIX FIGURE I
Mean Income Ranks for Children of Permanent Residents by County

**A. For Children with Parents at the 25th Percentile**
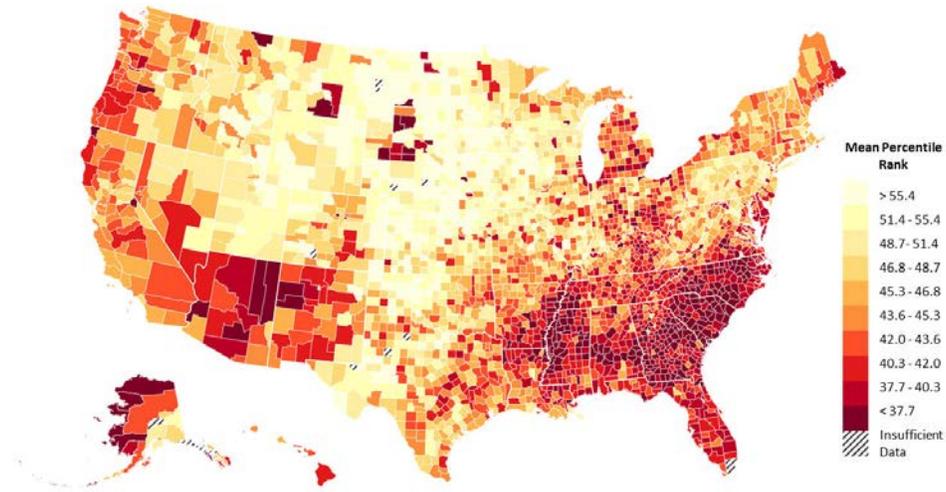
B. For Children with Parents at the 75th Percentile
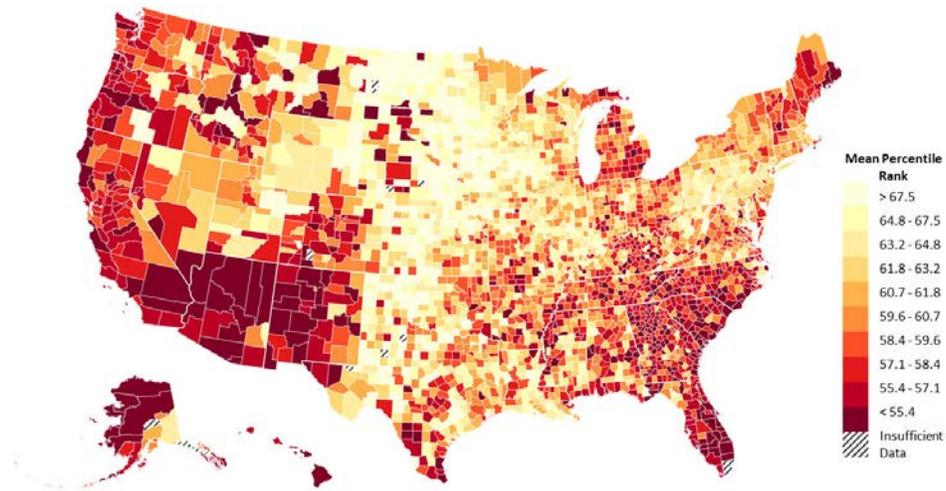
Notes: These maps plot children's predicted income percentile ranks at age 30 conditional on having parents at the 25th percentile (Panel A) and 75th percentile (Panel B) at the county level, analogous to the CZ-level maps shown in Figure II. The maps are constructed by grouping counties into ten deciles and shading the areas so that lighter colors correspond to higher outcomes for children. Areas with fewer than 10 children, for which we have insufficient data to estimate outcomes, are shaded with the striped pattern. The sample includes all children in the 1980-1982 birth cohorts in our analysis sample whose parents are permanent residents of a given county (i.e., whose parents do not move across counties between 1996-2012). See notes to Figure II for further details on the construction of this figure.

ONLINE APPENDIX FIGURE II

Childhood Exposure Effects using Variation in Origin CZs
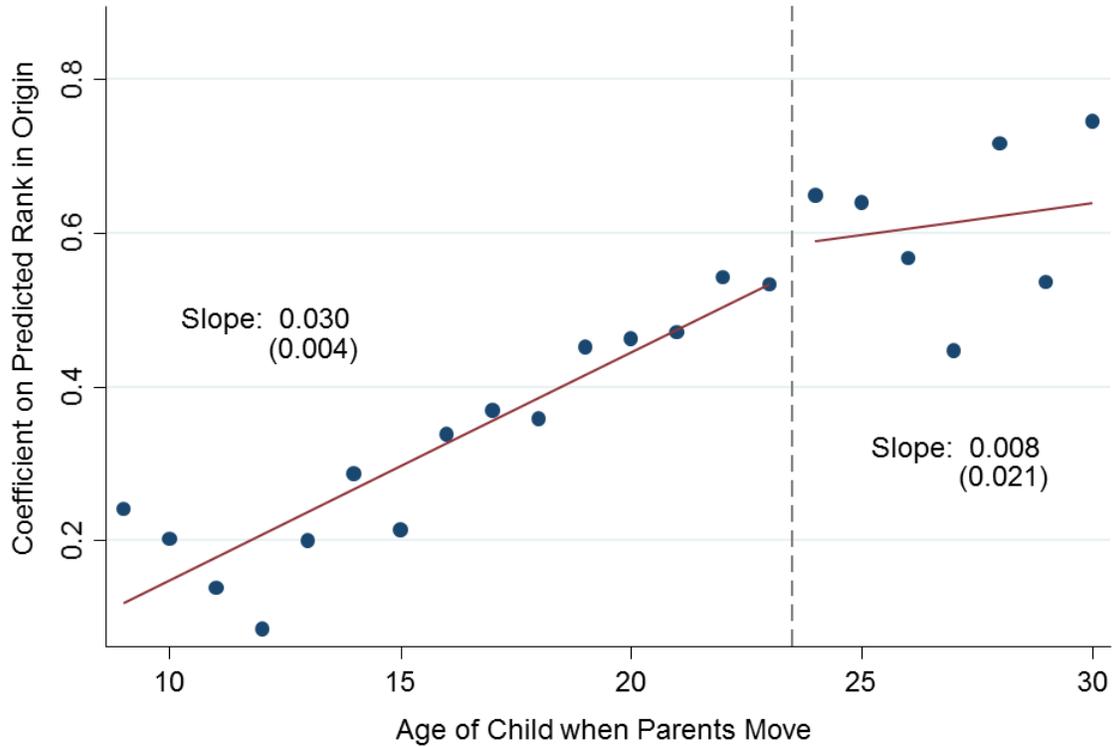
Notes: This figure presents estimates of b_m using the semi-parametric specification in equation (5), replacing the α_{qosm} fixed effects (origin by parent income decile by birth cohort by age at move) with α_{qdsm} fixed effects (destination by parent income decile by birth cohort by age at move). The estimates of b_m are thus identified from variation across permanent resident's outcomes in the origin from which children move rather than the destinations to which they move. We also report estimates from unweighted OLS regressions of the estimated coefficients b_m on the age at move m for $m \leq 23$ and $m > 23$, as in Figure IV. The slope for $m \leq 23$ represents an estimate of the annual childhood exposure effect that is identified from differences in exposure to an origin CZ with 1 percentile better permanent resident outcomes. See notes to Figure IV for further details on the construction of this figure.

ONLINE APPENDIX FIGURE III
Exposure Effects Adjusting for Children's Rates of Moving with Parents
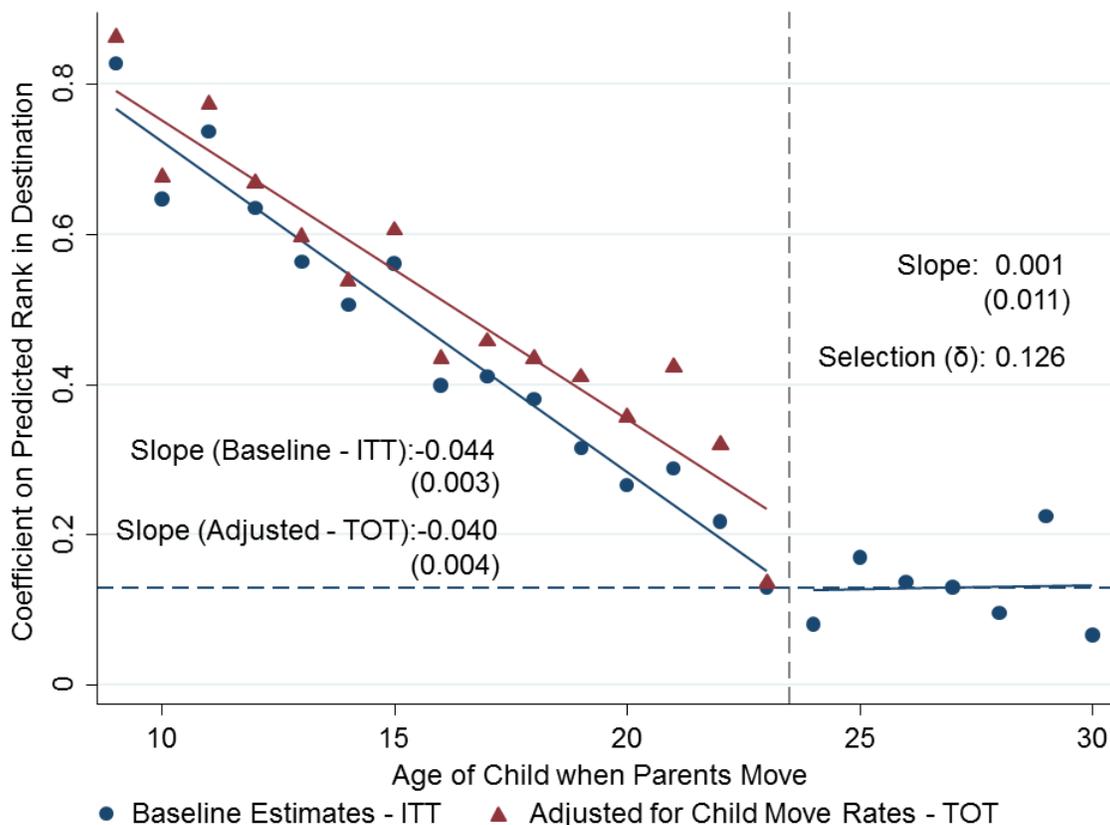
Notes: The series in circles plots estimates of the coefficients $\{b_m\}$ vs. the child's age when the parents move (m) using the semi-parametric specification in equation (5), replicating Figure IVa exactly. The series in triangles presents analogous estimates of b_m which adjust for the fact that not all children move with their parents. This adjusted series is constructed as follows. First, at each child age m, we estimate the fraction of children who follow their parents to the new destination when their parents move, ϕ_m, as the fraction of children who (a) file a tax return in the destination CZ, (b) have a form W-2 mailed to an address in the destination CZ, or (c) attend a college (based on 1098-T filings by colleges) in the destination CZ. For parents who move when their child is age $m \leq 23$, we then plot $b_m^{TOT} = \frac{b_m - \delta}{\phi_m} + \delta$, where $\delta = 0.126$ is the estimated selection effect (shown by the horizontal dashed line) and b_m are the baseline "intent-to-treat" estimates shown in circles. The b_m^{TOT} estimates, plotted in triangles, can be interpreted as "treatment-on-the-treated" exposure effects for children who actually move with their parents. We also report estimates from unweighted OLS regressions of the estimated coefficients b_m and b_m^{TOT} on the age at move m for $m \leq 23$, as in Figure IV. The slope using b_m^{TOT} as the dependent variable represents a "treatment-on-the-treated" estimate of the magnitude of the annual childhood exposure effect, γ^{TOT}. See notes to Figure IV for further details.

ONLINE APPENDIX FIGURE IV

Sensitivity of Childhood Exposure Effects to Measuring Income at Older Ages
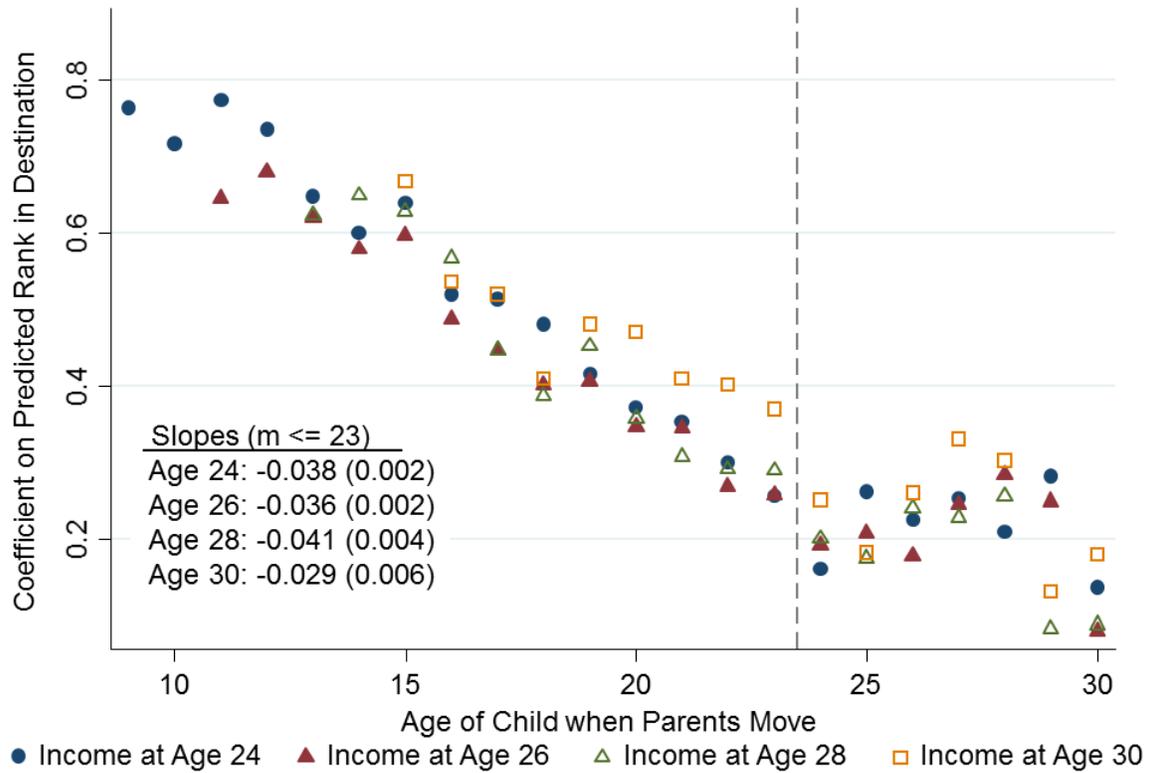

Notes: This figure replicates the parametric in equation (6) presented in Figure IVb, measuring children's incomes at ages 24, 26, 38, and 30. Coefficients from an unweighted regression of the estimated coefficients b_m on the age at move m for $m \leq 23$, which represent estimates of the magnitudes of annual childhood exposure effects, are reported. Because childrens' locations are not observed before 1996, the earliest age at which movers can be identified varies across specifications. See notes to Figure IV for further details on the construction of this figure.
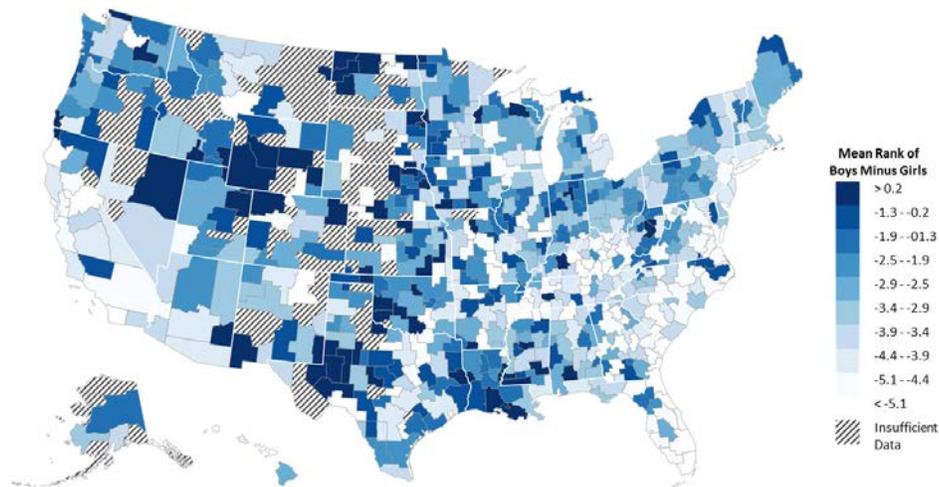
ONLINE APPENDIX FIGURE V
Gender Differences in Mean Income Ranks for Children of Permanent Residents

A. For Children with Parents at the 25th Percentile
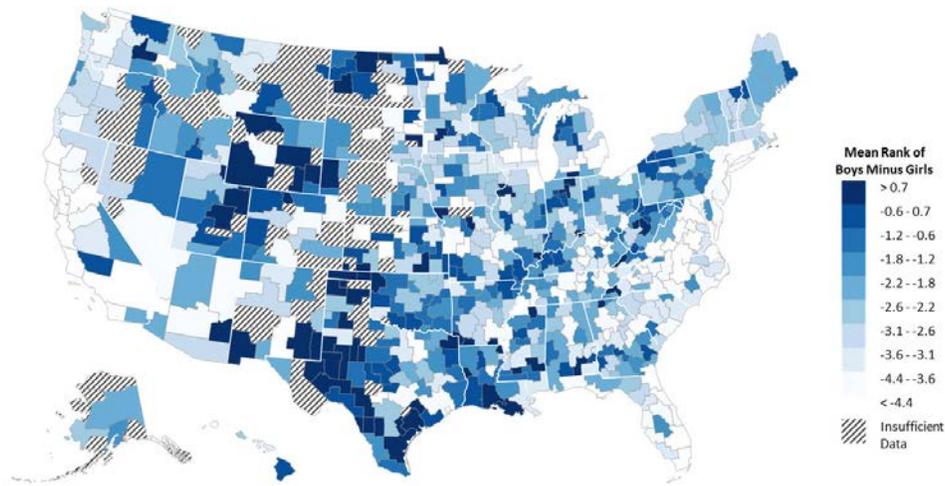
B. For Children with Parents at the 75th Percentile

Notes: These maps plot the difference between sons' and daughters' mean household income ranks at age 30 ($\bar{y}_{pcs}^{m} - \bar{y}_{pcs}^{f}$) conditional on having parents at the 25th percentile (Panel A) and 75th percentile (Panel B). The maps are constructed by grouping CZs into ten deciles and shading the areas so that darker colors correspond to better outcomes for boys relative to girls. Areas with fewer than 250 male or female children, for which we have insufficient data to estimate outcomes, are shaded with the striped pattern. The sample includes all children in the 1980-82 birth cohorts in our analysis sample whose parents are permanent residents of a given CZ (i.e., whose parents do not move across CZs between 1996-2012). Mean ranks are constructed using the same method as in Figure I, separately by gender; see notes to Figure I for further details.