The Economic Impact of Tax Expenditures: Evidence from Spatial Variation Across the U.S.*

Raj Chetty, Harvard University and NBER
Nathaniel Hendren, Harvard University and NBER
Patrick Kline, UC Berkeley and NBER
Emmanuel Saez, UC Berkeley and NBER

April 2015

Abstract

This paper develops a framework to study the effects of tax expenditures on intergenerational mobility using spatial variation in tax expenditures across the United States. We measure intergenerational mobility at the local (commuting zone) level based on the correlation between parents' and children's earnings. We show that the level of local tax expenditures (as a percentage of AGI) is positively correlated with intergenerational mobility and that this correlation is robust to introducing controls for local area characteristics. To understand the mechanisms driving this correlation, we analyze the largest tax expenditures in greater detail. Here, we find that the level and the progressivity of state income taxes are positively correlated with intergenerational mobility. Finally, we find significant positive correlations between state EITC policy and intergenerational mobility. We conclude by discussing other applications of this methodology to evaluate the net benefits of tax expenditures.

I Introduction

Tax expenditures – the exemption of certain economic activities from taxation – account for over a trillion dollars of annual federal, state, and local government spending.[1] Reducing tax expenditures is thus a potentially powerful way to reduce budget deficits. However, tax expenditures may also provide important benefits to the economy – e.g., by stimulating entrepreneurship and growth, increasing equality of opportunity, or providing better access to health care. This tradeoff makes it important to identify the costs and benefits of major tax expenditures to determine which expenditures are most valuable.

Given the importance of this question, a large literature studying the impacts of tax expenditures has developed over the past several decades. Previous work investigating the impacts of tax expenditures has largely relied on analysis at the national level. For example, a recent volume edited by Poterba (2011) includes several studies evaluating the economic effects of tax expenditures and reviews the voluminous previous literature on this issue. Virtually all of these studies exploit time series variation in federal tax expenditures – e.g., in the allowance for mortgage interest deductions or other tax credits. The limitation of such studies is that time series variation in tax expenditures is naturally correlated with many other factors that may affect the economy, such as changes in other government policies or the strength of the economy. Thus evidence on the benefits of tax expenditures remains limited.

In this paper, we address these empirical challenges using differences in tax expenditures across cities in the U.S. to identify the benefits of tax expenditures. There is considerable local variation in tax policy and expenditures that arises from variation in local policies that interact with the federal tax code. For instance, because state income taxes are deductible for federal tax purposes, states with higher income tax rates effectively receive larger tax expenditures than those that have higher sales taxes instead. Such local variation provides useful counterfactuals for outcomes in the absence of tax expenditures and thus can yield much sharper estimates of the impacts of tax expenditures.

To harness the power of spatial variation, we use selected information from de-identified federal income tax records spanning 1996-2012. These data provide information on a variety of economic outcomes of interest at a high level of spatial granularity. Such data are essential for the approach we propose here because one cannot obtain precise estimates of outcomes of interest within each

---

[1]This total refers to the sum of individual tax expenditure estimates and does not take into account interactions among different tax expenditures.

city from publicly available survey datasets.

There are a variety of outcomes that one could study to evaluate the efficacy of tax expenditures. For instance, one can investigate whether tax expenditures raise local income levels, increasing homeownership rates, change educational outcomes, affect mortality rates, stimulate new business starts, etc. As a first step in this research agenda, we focus on intergenerational mobility as the outcome of interest in this paper. We focus on intergenerational mobility because many tax expenditures are loosely motivated by the goal of expanding opportunities for upward income mobility for low-income families. For example, deductions for education and health costs, progressive federal tax deductions for state income taxes, and tax credits aimed at low-income families such as the Earned Income Tax Credit (EITC) all are targeted toward providing increased resources to low income families with children. Are these tax expenditures effective in promoting income mobility?[2]

We conduct our analysis using local statistics on tax expenditures and intergenerational mobility. For tax expenditures, we primarily draw upon existing data sources on tax expenditures by area. We focus on three sets of measures. First, we explore overall tax expenditures and the progressivity of tax expenditures. Second, we explore specific tax expenditures including mortgage interest deductions and state tax rates. Third, we explore local state EITC policies.

In contrast to measures of tax expenditures, statistics on intergenerational mobility do not exist at a fine geographic level in the U.S. Therefore, we construct new measures of intergenerational mobility in this study. We explore in detail various potential measures of mobility, such as the relationship between log parent income and log child income. However, we show log specifications suffer significant bias and non-robustness because of the omission of zero incomes and nonlinearities in the relationship between parent and child income. Therefore, we adopt a rank-rank specification similar to that used by Dahl and DeLeire (2008). We rank children based on their incomes relative to other children in the same birth cohort, including those with zero income. We rank parents of these children based on their incomes relative to other parents with children in these birth cohorts. The relationship between mean child ranks and parent ranks is almost perfectly linear. The slope of the rank-rank relationship is 0.341, i.e. a 10 percentile point increase in parent rank is associated with a 3.41 percentile increase in a child's income rank. We also show that children's college attendance and teenage birth rates are linearly related to parent income ranks. A 10 percentile point increase in parent rank is associated with a 6.8 percentage point increase in children's college

---

[2]Prior work on tax expenditures (e.g. Rosen 1985, Clotfelter 1985) has investigated other types of responses to specific expenditures such as the mortgage interest deduction or charitable giving. However, this paper is the first to analyze the impacts of such expenditures on intergenerational mobility.

enrollment rates. In short, the use of ranks generates a highly robust measure of intergenerational mobility.

We use this rank-rank specification to characterize variation in intergenerational mobility across commuting zones (CZs) in the U.S. Commuting zones are geographical aggregations of counties based on commuting patterns that are similar to metro areas but cover the entire U.S., including rural areas (Tolbert and Sizer, 1996). We assign children to commuting zones based on where they lived at age 16 – i.e., where they grew up – irrespective of whether they left that CZ afterward. When analyzing CZ's, we continue to rank both children and parents based on their positions in the *national* income distribution, which allows us to measure absolute outcomes as we discuss below.

Using these statistics, we analyze the relationship between tax expenditures and intergenerational mobility. We find that both the level and progressivity of CZ tax expenditures are positively correlated with higher levels of intergenerational mobility. These relationships are robust to the inclusion of a broad range of local demographic controls. An increase of overall tax expenditures by 1% of AGI in a CZ decreases the parent-child income correlation by .5 percentage points, relative to the mean correlation of 0.33. In addition, we find evidence that the tax expenditure components of mortgage interest deductions, state income taxes, and state EITCs each have individually positive effects on intergenerational mobility. The progressivity of overall tax expenditures and state income taxes also have a robust, significant relationship with higher intergenerational mobility.

We also explore the potential bias from aggregating our measures of mobility up to the CZ level. Although we prefer aggregation to the CZ level to reduce the impact of geographic sorting within neighborhoods, CZs are typically very large geographic regions; for example, the Boston CZ includes all of Cape Cod. Therefore, we assess the robustness of our results by constructing mobility statistics at the county level. We find broadly similar results at the county level.

Overall, our results suggest that tax expenditures aimed at low-income taxpayers can have significant impacts on economic opportunity. Hence, the short-term fiscal gains from reducing such expenditures must be weighed against the potentially large long-term costs of reduced income growth for low income individuals.

Our results build on and contribute to an extensive empirical literature on intergenerational mobility, reviewed by Solon (1999), Grawe and Mulligan (2002), and Black and Devereux (2011). Several studies have compared mobility across counties using a log-log specification and have found that relative mobility is lower in the U.S. than in other developed countries (e.g., Bjorklund and Jäntti 1997, Corak and Heisz 1995, Jäntti et al. 2006, Smeeding et al., 2011, Corak 2013). Our

estimates of the IGE in the U.S. as a whole are similar to those found in prior work, with the exception of Mazumder's (2005) widely cited estimates, which imply much lower levels of intergenerational mobility than we find here for reasons we explain in Section III.B below. Our analysis is most closely related to contemporaneous work by Graham and Sharkey, who use data from the National Longitudinal Survey of Youth and the Panel Study of Income Dynamics to estimate relative mobility using log-log specifications in a subset of cities in the U.S. Their estimates are correlated with ours, but are naturally less precise due to limitations in sample size and do not permit an assessment of absolute mobility.

Our approach of within-country comparisons offers two advantages over the cross-country comparisons that have been the focus of prior comparative work. First, differences in measurement and econometric methods make it difficult to reach definitive conclusions from cross-country comparisons (Solon 2002, Black and Devereux 2011). The income measures and covariates we analyze here are all measured using the same data sources across all CZ's, making it easier to draw comparisons and learn about the determinants of mobility. Second, and more importantly, we can characterize both relative and absolute mobility across CZ's by using national ranks to measure children's outcomes. The cross-country literature has focused on differences in relative mobility across countries, partly because it is difficult to compare the absolute standard of living across very different economies. Although the literature on cross-country differences in economic growth had characterized differences in mean absolute living standards across nations, much less is known about how the prospects of children from low-income families vary across countries when measured on a common absolute scale (Ray 2010). In particular, the greater degree of relative mobility in European countries than the U.S. could be partly driven by worse outcomes for high-income children rather than better outcomes for low-income children.

In addition to providing new evidence on the role of tax expenditures in mobility, our analysis contributes to the literature by using new data to compute intergenerational mobility at the local level. These CZ measures of tax expenditures and intergenerational mobility are provided in the Online Appendix for future research into the impact of tax expenditures on intergenerational mobility. All CZ and county level statistics are disclosed only for CZ and counties with more than 250 observations in our linked parents-child dataset.

The remainder of the paper proceeds as follows. Section 2 describes the tax expenditure and intergenerational mobility data used for our analysis. Section 3 presents the main analysis of the relationship between tax expenditures and intergenerational mobility, including an extensive

discussion of potential confounding factors. Section 4 concludes and outlines directions for future research.

II   Measuring Tax Expenditures and Intergenerational Mobility

We draw on two primary measures of data. First, we utilize measures of local tax expenditures. Such data is readily available from existing sources we describe in the next subsection. Second, we construct new measures of intergenerational mobility. This data is not available at a local level; therefore, we discuss in detail our construction of such measures from tax records.

II.A   Geographical Definitions

Our empirical approach relies on variation in tax expenditures and intergenerational mobility across areas. To implement such an approach, we must first define our notion of geography. In our baseline analysis, we study spatial variation across commuting zones (CZs). CZs are aggregations of counties based on commuting patterns in the 1990 census constructed by Tolbert and Sizer 1996. CZs are roughly similar in size to metropolitan statistical areas (MSA), but cover the entire U.S. (including rural areas). There are 741 CZ's in the United States; on average, each CZ contains 4 counties and has a population of 379,786.

The choice of geographical unit used is to some extent arbitrary: one could also construct statistics by state, county, or ZIP code. We use CZs because they provide an economically meaningful definition of a local labor market area whose characteristics might affect economic outcomes (Dorn 2009, Autor and Dorn 2013). To assess the robustness of our findings, we also construct analogous statistics on intergenerational mobility by MSA and county in the Appendix tables. Appendix Table I provides a correlation of our baseline CZ mobility estimates with the measures derived using these alternative geographic definitions. In Table IX and Table X, we show that our primary results on the correlations of tax expenditures are robust to these alternative definitions of geography, as discussed below.

II.B   Local Tax Expenditure Data

Local CZ tax expenditure data are constructed from the Internal Revenue Service SOI Individual Income Tax Statistics ZIP Code Data from Internal Revenue Service (2008). This publicly available dataset includes 5-digit ZIP Code totals for 2008 number of returns, Adjusted Gross Income (AGI), total itemized deductions, mortgage interest deductions, and federal Earned Income Tax Credit

(EITC) by seven AGI classes.[3]

To measure local tax expenditures, total AGI and total itemized deductions are aggregated across all AGI classes and combined to the CZ level. CZ total itemized deductions are then measured as a percentage of total AGI, effectively resulting in CZ mean total itemized deductions as a percentage of AGI. Figure I maps this measure of overall CZ tax expenditures by dividing CZs into ten equally sized deciles. Darker areas represent areas with higher total itemized deductions relative to AGI. To measure the progressivity of tax expenditures, we use the difference in total itemized deductions as a percentage of AGI for the lowest AGI class, under $10,000, and the highest AGI class, $200,000 and over. The progressivity of tax expenditures varies greatly across CZs. Figure II maps this measure by CZ. Darker areas represent regions with more progressive local tax expenditures.

Mortgage interest deductions are aggregated from ZIP-5 AGI class totals to CZ overall mortgage interest deduction totals and measured as a percentage of AGI. Mean mortgage interest deductions for each AGI class are calculated by dividing CZ total mortgage interest deductions by the number of returns. Inequality of mortgage interest deductions by CZ is measured as the difference between the mean mortgage interest deduction for the top AGI class ($200,000 and over) and the lowest AGI class (under $10,000). To avoid mechanical scaling effects, we control for local housing prices in our analysis of mortgage interest deductions. A categorical variable indicating median house price bracket is obtained from the 2000 Census for each ZIP-5, and combined with housing counts to obtain CZ measures of local housing prices.

State income marginal tax rates for the 2008 tax year are obtained from The Tax Foundation (2012).[4] We use the marginal tax rate for individuals with taxable incomes of $40,000 to measure the overall level of state income taxes. To measure the progressivity of state income tax rates, we compute the difference in the marginal tax rate for the top bracket specified for the given state and the marginal tax rate for incomes of $20,000. New Jersey has the most progressive state income taxes by this measure, with a 7.22% difference in marginal tax rate for taxable incomes over $500,000 and taxable incomes of $20,000.

In addition, we construct a measure of the exposure to the EITC. Because state EITC policies changed significantly over the period when children in our sample were growing up, we define a

---

[3]The AGI classes are "Under $10,000", "$10,000 to $25,000", "$25,000 to $50,000", "$50,000 to $75,000", "$75,000 to $100,000", "$100,000 to $200,000", and "$200,000 and more."

[4]See http://taxfoundation.org/article/state-individual-income-tax-rates-2000-2012 for state individual income tax rates in 2008.

measure of mean exposure to the state EITC as the mean state EITC rate between 1981 and 2001, when the children in our sample were between the ages of 0 and 20.[5]

Finally, we construct two additional tax measures. First, we construct the mean local tax rate in each CZ. This is given by the total tax revenue collected at the county or lower level in the CZ (based on the 1992 Census of Governments) divided by total household income in the CZ based on the 1990 Census. Second, we construct a measure of tax progressivity, which is given by the difference in the top state income tax rate and the state income tax rate for individuals with taxable income of $20,000 in 2008 based on data from Tax Foundation.

II.C Data on Intergenerational Mobility

Although regional measures of tax expenditures already exist, we must construct new local measures of intergenerational mobility. To do so, 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 give us information on the earnings of those who do not file tax returns. We provide a detailed description of how we construct our analysis sample starting from the raw population data in Appendix A. Here, we briefly summarize the key variable and sample definitions. Note that in what follows, the year always refers to the tax year (i.e., the calendar year in which the income is earned).

II.C.1 Sample Definitions

Our base dataset of children consists of all individuals in Social Security Administration records who (1) have a valid Social Security Number or Individual Taxpayer Identification Number, (2) were born between 1980-1991, and (3) are U.S. citizens as of 2013. 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.[6]

We identify the parents of the children in the base dataset using information on dependent claiming on filed tax returns (form 1040). The parents of a child are defined as the first tax filers (over the period 1996-2012) who claim the child as a child dependent. To exclude cases where a grandparent of a sibling is the child's custodian, we require that the tax filer who claims the child

---

[5]We assign state-years without a state EITC a state EITC rate of 0 when computing this mean.

[6]The database only records current citizenship status; hence, we cannot directly restrict the sample to individuals born in the United States. Because of immigration and naturalization, the number of individuals in our sample in the 1981 birth cohort exceeds the size of the 1980 US birth cohort (as recorded in vital statistics data) by 4.6%, as shown in Appendix Table II. See Appendix A for further details.

(the female in the case of joint filers) is between the ages of 15-40 when the child was born. If the child is first claimed by a single filer, the child is defined as having a single parent. We assign each child a parent (or parents) permanently using this algorithm, ignoring any subsequent changes in parents' marital status. This approach simplifies computation and is most likely to identify the structure of the household when the child was growing up, since we link most children to parents relatively late in their childhood.

If parents never file a tax return, we cannot link them to their child. Although some low-income individuals do not file tax returns in a given year, almost all non-institutionalized working-age adults with children file a tax return at some point between 1996 and 2012.[7] As a result, we are able to identify parents for approximately 95% of the children in the birth cohorts we analyze (Appendix Table II).[8] Appendix Table II reports the fraction of children linked to parents by birth cohort. This fraction rises from approximately 90% for the early birth cohorts in our base sample to nearly 99% for the most recent birth cohorts. This is because we have more opportunities to link younger children to their parents. For children in the 1980 birth cohort, the earliest age at which we observe parents' tax returns is age 16 (in 1996). Because many children begin to leave the household around age 18, match rates fall sharply prior to the 1980 birth cohort, as shown in Appendix Table II. We therefore limit our analysis to children born during or after 1980. We also replicate our main results restricting the sample to more recent birth cohorts to ensure that our results are not biased by selective attrition.

Our primary analysis sample, which we refer to below as the *core sample*, includes all children in the base dataset who are (1) born in birth cohorts 1980-82, (2) for whom we are able to identify parents, and (3) whose mean parent income over the years 1996-2000 is strictly positive (which excludes 1.2% of children).[9] Since parent income is the independent variable throughout our analysis, our results can be interpreted as applying to the population of parents with positive mean income over a five year interval. For some analyses – such as the measurement of college and teenage birth – we use the *extended core sample*, which imposes the same restrictions as the core sample, but includes all birth cohorts 1980-1991. The core sample contains approximately 10 million children,

---

[7]Most non-filers are individuals with zero income or retired individuals receiving solely social security benefits (Cilke, 1998). Individuals with positive income have incentives to file to obtain a tax refund on their withheld taxes and the EITC.

[8]Chetty et al. (2014 forthcoming) present further evidence that one can identify parents for virtually all children who grew up in the U.S. by showing that 98% of children enrolled in a large school district in grades 3-8 can be linked to parents in the tax data.

[9]We limit the sample to parents with positive income because parents who file a tax return (as required to link 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.

while the extended core sample contains approximately 44 million children.

Statistics of Income Sample. Because we can only reliably link children to parents starting with the 1980 birth cohort in the population tax data, we can only measure earnings of children up to age 32 (in 2012) in the full sample. To evaluate whether estimates of intergenerational mobility would change significantly if earnings were measured at later ages, we supplement our analysis using a random cross-section of filed tax returns maintained by the Statistics of Income (SOI) division of the Internal Revenue Service prior to 1996. The SOI cross-sections provide identifiers for dependents claimed on tax forms starting in 1987, allowing us to link parents to children back to the 1971 birth cohort in this smaller sample using an algorithm analogous to that described above (see Appendix A for further details). The SOI cross-sections are stratified random samples of tax returns with a sampling probability that rises with income; using sampling weights, we can calculate statistics representative of the national distribution. After linking parents to children in the SOI sample, we use population tax data to obtain data on income for children and parents and impose the same restrictions used to define the core sample above to obtain our SOI analysis sample. The SOI sample contains 130,000 children for cohorts 1971-1981, making it sufficiently large to estimate intergenerational mobility at the national but not local level.

II.C.2 Variable Definitions and Summary Statistics

In this section, we define the key variables we use to measure intergenerational mobility.

Parent Income. Following Lee and Solon (2009), our primary measure of parent income is total pre-tax income at the household level, which we label *parent family income.* More precisely, 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. This income measure includes labor earnings and capital income as well as unemployment insurance, social security, and disability benefits. It excludes non-taxable cash transfers such as TANF and SSI, in-kind benefits such as food stamps, all refundable tax credits (such as the EITC), and any non-taxable pension contributions (e.g., to 401(k)'s). 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.[10] We cannot observe

---

[10]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. We verify that this is not an important source of bias by showing that we obtain very similar results when defining parent income only using post-1999 data.

self-employment income for non-filers and therefore code it as zero; given the strong incentives for individuals with children to file created by the EITC, most non-filers likely have very low levels of self-employment income. In years where parents have no tax return and no information returns, family income is coded as zero. Income is always measured prior to the deduction of individual income taxes and employee-level payroll taxes.[11]

In our baseline analysis, we average parents' family income using 5 years from 1996 to 2000 to reduce noise due to 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. When computing mean income, we measure the mean family income of the fixed set of parents identified using the algorithm described above.[12]

We evaluate the robustness of our findings using a measure of individual parent income instead of family income. For single parents, parent family income coincides with individual income. For married parents, we define each parent's individual earnings as the sum of wage earnings from form W-2, unemployment benefits from form 1099-G, and Social Security and Disability benefits from form SSA-1099 for that individual. Individual earnings excludes capital and other non-labor income. To incorporate these sources of income, we turn to the 1040 tax return and add half of family non-labor income – defined as total family income minus total family earnings reported on form 1040 – to each parent's individual earnings. We divide non-labor earnings equally between spouses because we cannot identify which spouse earns non-labor income from the 1040 tax return.

Child Income. We define child family income in exactly the same way as parent family income. In our baseline analysis, we average child family income over the last two years in our data (2011 and 2012). We use the most recent two years because the children in our sample are all born after 1980, and income in the early 30's provides a better measure of lifetime income than income at earlier ages (Haider and Solon 2006). We report results using alternative years to assess the sensitivity of our findings. For children, we define household income based on current marital status rather than marital status at a fixed point in time. Because family income varies with marital status, we also

---

[11]To remove the effects of outliers and measurement error in our results, we use the SOI data to verify top earners when available and otherwise replace adjusted gross income with wage amounts. This affects than 0.117% of the sample.

[12]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 decrease by a half if previously married parents divorce, which provides part of the motivation for assessing the robustness of our results using individual measures of income.

10

report results using individual earnings and income measures for children, constructed in exactly the same way as for parents.

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 18-21. Title IV institutions – all colleges and universities as well as vocational schools and other postsecondary institutions eligible for federal student aid – are required to file 1098-T forms that report tuition payments or scholarships received for every student. Because the 1098-T forms are filed directly by colleges independent of whether an individual files a tax return, we have complete records on college attendance for all children. However, we have no information about college completion or degree attainment because the data are based on tuition payments. The 1098-T data are available from 1999-2011.[13] Comparisons to other data sources indicate that 1098-T forms capture college enrollment quite accurately (Chetty, Friedman, and Rockoff 2013, Section 3.1).

Teenage Birth. We use teenage birth as an outcome measure exclusively for women. To measure teenage birth, we first identify all women who claim a dependent when filing their taxes at any point before 2012. We observe dates of birth and death for all dependents and tax filers until the end of 2012 as recorded by the Social Security Administration. We use this information to identify women who ever claim a dependent who was born while the mother was a teenager (between the ages of 13 and 19 as of 12/31 the year the child was born). This measure is an imperfect proxy for having a teenage birth because it is based purely on dependent claiming.[14] Nevertheless, the aggregate level and spatial pattern of teenage births in our data is closely aligned with estimates based on the American Community Survey.[15]

Table I reports summary statistics for the core sample. We express all monetary variables in 2012 dollars, adjusting for inflation using the consumer price index (CPI-U). Median parent family income is \$60,129. Among the 30.6% of children matched to single parents, 72.0% are matched to a female parent. Children in our core sample have a median family income of \$34,975 when they are approximately 30 years old. 6.1% of children have zero income in both 2011 and 2012. 58.9%

---

[13]Colleges are not required to file 1098-T forms for students whose qualified tuition and related expenses are waived or paid entirely with scholarships or grants; similarly, colleges have varying policies about filing 1098-T forms for international students. However, the forms are frequently available even for such cases, presumably because of automated reporting to the IRS by universities.

[14]In particular, our measure does not capture births to mothers who never file a tax return before 2012 or children who are claimed by other relatives (such as their grandparents). Our definition would also miscategorize other dependents who are not biological children, but were born between 13 and 19 years after the female who claims them as a dependent, as teenage births.

[15]15.9% of women in our core sample are teen mothers in our data; the corresponding number is 14.6% in the 2003 ACS data. The unweighted correlation between the state-level teenage birth rates in the tax data and the ACS is 0.80.

are enrolled in a college at some point between the ages of 18 and 21 and 15.9% of women have a teenage birth.

In Appendix C, we show that the distribution of child income in our data lines up closely with corresponding estimates in the Current Population Survey (CPS) and American Community Survey (ACS). This confirms that our dataset covers roughly the same nationally representative population as previous research on intergenerational mobility.

II.D   National Statistics

In order to construct measures of intergenerational mobility at the local level, we first must identify the appropriate metric for studying mobility using data aggregated to the national level. In principle, there are many ways to characterize the relationship parent and child income. A key contribution of this section is the establishment of a linear-in-ranks specification that allows for a parsimonious characterization of mobility in terms of an intercept and slope in parent income. We also show this measure is a highly robust measure of mobility by comparing it to two alternative mobility outcomes, namely college attendence and teenage birth.

II.D.1   Baseline Estimates

In our baseline analysis, we use the core sample (1980-82 birth cohorts) and measure parent income as mean family income from 1996-2000 and child income as mean family income in 2011-12 (as in Lee and Solon 2009), when children are approximately 30 years old. Figure IIIa presents a binned scatter plot of the mean family income of children versus the mean family income of their parents. To construct this figure, we divide the x axis into 100 equal-sized (percentile) bins and plot mean child income vs. mean parent income in each bin. This binned scatter plot provides a non-parametric representation of the conditional expectation of child income given parent income. The regression coefficients and standard errors reported in this and all subsequent binned scatter plots are estimated on the underlying microdata using OLS regressions. Figure IIIa shows that the conditional expectation of children's income given parents' income is strongly concave. Below the 90th percentile of parent income, a \$1 increase in parent family income is associated with a 33.5 cent increase in average child family income. In contrast, above the 90th percentile, a \$1 increase in parent income is associated with only a 1.6 cent increase in child income.

Partly motivated by this non-linearity, much of the empirical literature has estimated regressions of log child income on log parent income. The slope of this regression measures the elasticity of child

income with respect to parent income, commonly termed the intergenerational income elasticity (IGE). We implement this specification in the first column of row 1 of Table II, excluding children with zero income as in prior work. We obtain an IGE estimate of 0.344, similar to the estimates of Solon (1992) and Zimmerman (1992) in survey data when using multi-year income averages.

The remaining columns in the first row of Table II replicate the log-log specification for alternative samples analyzed in the prior literature. Columns 2-5 split the sample by the child's gender and the parents' marital status in the year they first claim the child. Column 6 replicates Column 1 for the extended sample of 1980-85 birth cohorts. Column 7 restricts to the sample to children whose parents were within 5 years of median parent age at child birth.[16] The IGE estimates are similar for males and females, but are lower when we condition on marital status.

While the log-log specification is a familiar and intuitive benchmark, it suffers from two shortcomings that are illustrated in Figure IIIb. First, the relationship between log child income and log parent income is highly non-linear, consistent with the findings of Corak and Heisz (1999) in Canadian tax data. This is illustrated in the series in blue circles in Figure IIIb, which plots mean log child income vs. mean log family income by percentile bin, constructed using the same method as Figure IIIa. Because of this non-linearity, the IGE is sensitive to the point of measurement in the income distribution. For example, restricting the sample to bins between the 10th and 90th percentile of parent income (denoted by the vertical dashed lines in the graph) yields a considerably higher IGE estimate of 0.452.

Second, the log-log specification discards observations with zero income. The red series in triangles in Figure IIIb plots the fraction of children with zero income by parental income bin. This fraction varies from 18% among the poorest families to less than 5% among the richest families. Dropping children with zero income overstates the degree of intergenerational mobility by ignoring the lower labor force participation rates of children from low-income families. The way in which these zeros are treated can change the IGE dramatically. For instance, including the zeros by assigning those with zero income an income of \$1 (so that the log of their income is zero) raises the estimated IGE to 0.618, as shown in row 2 of Table II. If instead we treat those with 0 income as having an income of \$1,000, the estimated IGE becomes 0.413. Hence, small differences in the way in which children's income is measured at the bottom of the distribution can substantially change IGE estimates.

---

[16]Precisely, mother's age at child birth must be between 24-28 and father's age at child birth must be between 26-30. For single parents, only one of these restrictions is used.

To address these shortcomings of the log-log specification, we use a *rank-rank* specification similar to that proposed by Dahl and DeLeire (2008). We define the percentile rank of the parent to be their rank in the national distribution of parent incomes in our core sample divided by the number of parents. For example, a parent rank of 0.4 means that parent is at the 40th percentile of the distribution of parent incomes in our core sample.[17] Similarly, we define the percentile rank of each child as their rank in the national distribution of child incomes within their birth cohort divided by the size of their birth cohort. Importantly, this definition allows us to include zeros in child income, who are coded as having the lowest rank.[18] Unless otherwise noted, we hold the definition of these ranks fixed using the aggregate distribution, even when analyzing subgroups.

Figure IV presents a binned scatter plot of the mean percentile rank of children vs. their parents' percentile rank. The conditional expectation of a child's rank given his parents' rank is almost perfectly linear, except for a small amount of curvature at the very bottom of the parent income distribution. An OLS regression implies that a one percentage point increase in parent rank is associated with an increase of roughly 0.341 percentage points in the expected rank of the child, as reported in row 4 of Table II. While such a relationship may seem quite surprising, Appendix Table III verifies that this relationship holds at most quantiles of the child income distribution. For example, the median child rank as a function of the parental income rank is roughly linear, as is the 25th and 75th percentile of a child's rank. As a result, it is clear that this linearity is not spurious, but rather a structural property of the relationship between parent and child income.

Previous literature has highlighted two major potential sources of biases in estimates of inter-generational mobility: lifecycle bias and attenuation bias. In Appendix E, we carefully document that our rank-rank specification does not suffer from these sources of bias. We also show that the linearity property in parental rank also holds for two highly correlated outcomes with earnings: the chances of attending college and for the probability of having a teenage birth. As a result, this rank-rank specification will be our primary measure of mobility that we seek to estimate at a more local level.

II.E Spatial Variation in Mobility

We now turn to characterizing the variation in intergenerational mobility across areas within the U.S. We begin by defining measures of geographic location. Next, we define two concepts of

---

[17]Because parent ranks are defined on a child-level dataset, some parents may appear multiple times in the dataset.
[18]In the case of ties, we define the rank as the mean rank for the individuals in that group. For example, if 10% of the cohort has zero income, all children with zero income would receive a rank of 0.05.

intergenerational mobility – relative mobility and absolute mobility – that we measure at the CZ level using a rank-rank specification. We then present baseline estimates of relative and absolute mobility by CZ. Finally, we assess the robustness of these estimates to alternative measures.

II.E.1 Geographical Assignment

We *permanently* assign each child to a single CZ based on the ZIP code from which his or her parent filed their tax return in the first year the child was claimed as a dependent.[19] We interpret this CZ as the area where a child grew up. Because our data begin in 1996, location is measured in 1996 for 95.9% of children in our sample. For children in our core sample of 1980-82 birth cohorts, we therefore typically measure location when children were approximately 15 years old. For the children in the more recent birth cohorts in our extended sample, location is measured at earlier ages. Using these more recent cohorts, we find that 83.5% of children live in the same CZ at age 16 as they did at age 5. Furthermore, we verify that the spatial patterns for the outcomes we can measure at earlier ages (college attendance and teenage birth) are quite similar if we define CZ's based on location at age 5 instead of age 16. We therefore believe that our baseline definition of location provides a reasonably informative measure of the area where a child grew up.

Importantly, the CZ where a child grew up does not necessarily correspond to the CZ he lives in as an adult when we measure his income (at age 30) in 2011-12. In our core sample, 61.9% of children live in a different CZ as adults. A key strength of our data is that they allow us to systematically track children who move across the country, which is particularly important when studying income mobility because the probability of moving rises with parent and child income.

II.E.2 Measures of Relative and Absolute Mobility

In our baseline analysis, we measure mobility at the CZ level using the core sample (1980-82 birth cohorts) and the definitions of parent and child family income described in Section III.A. Importantly, we continue to rank both children and parents based on their positions in the *national* income distribution (rather than the distribution within their CZ), using exactly the same ranks as in Figure IV. By measuring outcomes on a fixed national scale, we can compare both relative and absolute outcomes across CZs.

Figure VIa illustrates our measures of relative and absolute mobility. This figure plots mean child rank vs. parent rank as in Figure IV, restricting the core sample to children who grew up in

---

[19]We use a crosswalk from 1999 ZIP codes to CZs that is constructed as described in Appendix IV.

the Salt Lake City, UT (circles) and Charlotte, NC commuting zones (triangles). The rank-rank relationship is virtually linear in both of these CZ's, as at the national level. This linearity of the rank-rank relationship is a remarkably robust property across CZs, as illustrated in Appendix Figure I.[20]

Exploiting this approximate linearity, we can summarize the conditional expectation of a child's rank given his parents' rank in each CZ using just two parameters: a slope and the intercept. Formally, let y_{ic} denote the national income rank (among children in his birth cohort) of child i who grew up in CZ c. Similarly, let x_{ic} denote his parent's rank in the income distribution of parents in the core sample. For each CZ c, we estimate a slope and intercept using an OLS regression of child rank on parent rank in the microdata:

$$y_{ic} = \alpha_c + \beta_c x_{ic} + \varepsilon_{ic} \tag{1}$$

Let $\bar{y}_{pc}$ denote the expected rank of a child whose parent's national income rank is p in CZ c.

One way to measure intergenerational mobility is to ask, "What are the outcomes of children from low-income families relative to those of children from high-income families?" This question has been the focus of most prior research on intergenerational mobility (Solon 1999, Black et al. 2011). We define the degree of *relative mobility* in CZ c as the difference between the expected outcomes of children born to parents at the top and bottom of the income distribution:

$$\bar{y}_{100,c} - \bar{y}_{0,c} = 100 \times \beta_c.$$

In Salt Lake City, $\bar{y}_{100} - \bar{y}_0 = 26.4$: children at the top of the distribution have an expected outcome that is 26.4 percentiles above the children of the poorest parents. Charlotte exhibits much less relative mobility (i.e., much greater persistence of income across generations). In Charlotte, $\bar{y}_{100} - \bar{y}_0 = 39.7$.

A different way to measure intergenerational mobility is to ask, "What are the outcomes of children from families of a given income level in absolute terms?" We define the degree of *absolute mobility* at percentile p in CZ c as the expected rank of a child whose parent's rank is p:

$$\bar{y}_{pc} \equiv E\left[y_{ic}|x_{ic} = p, c\right] \simeq \alpha_c + \beta_c p \tag{2}$$

---

[20]Appendix Figure I plots the rank-rank relationship for the 20 largest CZs along with the ratio between the local fraction of parents in each percentile bin and the national fraction. The relationship between expected child rank and parent rank remains linear even across CZs with very different parent income distributions.

Absolute mobility can be calculated at any percentile p using the estimates of β_c and intercept α_c that we report below for each CZ. Given academic and public interest in the outcomes of disadvantaged youth, we focus much of our discussion on average absolute mobility for children from families with below-median parent income ($E[y_{ic}|x_{ic} < p = 50, c]$), which we term "absolute upward mobility." Because the rank-rank relationship is linear, the average rank of children with below-median parent income equals the average rank of children with parents at the 25th percentile ($\bar{y}_{25}$), illustrated by the dashed vertical line in Figure VIa.[21]

Absolute upward mobility is $\bar{y}_{25} = 46.2$ in Salt Lake City, compared with $\bar{y}_{25} = 35.8$ in Charlotte. That is, among families earning \$28,800 – the 25th percentile of the national parent family income distribution – children who grew up in Salt Lake City are on average 10 percentile points higher in their birth cohort's income distribution at age 30 than children who grew up in Charlotte.

Absolute mobility is higher in Salt Lake City not just for below-median families, but at all percentiles p of the parent income distribution. The gap in absolute outcomes is largest at the bottom of the income distribution and nearly zero at the top. This result has two important implications that we investigate more systematically below. First, the greater relative mobility in Salt Lake City comes from better absolute outcomes at the bottom of the distribution rather than worse outcomes at the top. Second, the difference in outcomes across the CZ's is greater for children from low income families than for those from high income families.

Figure VIb shows an analogous comparison of intergenerational mobility in San Francisco and Chicago. San Francisco has substantially higher relative mobility than Chicago: $\bar{y}_{100} - \bar{y}_0 = 25.0$ in San Francisco vs. $\bar{y}_{100} - \bar{y}_0 = 39.3$ in Chicago. However, unlike in the Salt Lake City vs. Charlotte comparison, San Francisco does not dominate Chicago in terms of absolute mobility across the income distribution. Below the 60th percentile, children in San Francisco have better outcomes than those in Chicago; above the 60th percentile, the reverse is true. Part of the greater relative mobility in San Francisco comes from worse outcomes for children from high-income families rather than better outcomes for children from low income families.

The comparisons in Figure VI illustrate the importance of measuring both relative and absolute mobility from a normative perspective. Any social welfare function based on mean rank outcomes that respects the Pareto principle would rank the outcomes in Salt Lake City above Charlotte. But normative comparisons of San Francisco and Chicago depend on the weight one puts on relative

---

[21]Absolute mobility is related to but conceptually distinct from the local growth rate in a CZ because it tracks the outcomes of individuals who move out of the CZ and measures mean child ranks at a specific percentile of the national parent income distribution rather than integrating over the entire local income distribution.

vs. absolute mobility (or, equivalently, on the weights one places on absolute mobility at each percentile p). Changes in relative mobility by themselves have ambiguous normative implications. For instance, a policy that reduces the earnings of children from high-income families while keeping other children's earnings fixed increases relative mobility. But such a change would *reduce* social welfare under a normative criterion that respects Pareto efficiency. In contrast, an increase in absolute mobility at percentile p holding fixed absolute mobility at other income levels would increase welfare under any such criterion. Hence, evaluating policies that raise absolute mobility at certain income levels while reducing absolute mobility at others requires measures of both relative and absolute mobility.

There are many potential ways of measuring mobility. Indeed, any measure of the extent to which parental income predicts children's outcomes in a CZ can be interpreted as a measure of relative mobility. In analyzing the relationship between tax expenditures and mobility, we will be careful to assess robustness of our results to alternative definitions.

To implement this robustness analysis, we also compute non-parametric quintile transition matrices, as in Corak and Heisz (1999), Hertz (2006), and Jäntti et al. (2006). Let π_c^{mn} denote the probability that a child is in national income quintile m conditional on their parent being in national income quintile n for children who grew up in CZ c. The quintile transition matrix for CZ c is the set of 25 probabilities $\{\pi_c^{mn}\}$ for $m = 1, ..., 5$ and $n = 1, ..., 5$. Note that by definition, $\sum_{m=1}^{5} \pi_c^{mn} = 1$. To improve precision, we estimate π_c^{mn} pooling the 1980-1985 birth cohorts. We use the baseline family income definitions for both parents and children.

Appendix Table IV reports transition matrix estimates for each CZ. These transition matrices characterize the joint distribution of parent and child income within each CZ non-parametrically. Combined with the marginal distributions of parent and child income at the national level (reported in Appendix Table V and Appendix Table VI), these tables enable our robustness analysis discussed in Section F.

II.E.3   Baseline Estimates by CZ

We estimate (1) using OLS separately for each CZ to calculate absolute upward mobility ($\bar{y}_{25,c} = \alpha_c + 0.25\beta_c$) and relative mobility ($100 \times \beta_c$) by CZ. Columns 6 and 7 of Appendix Table IV report estimates of absolute and relative mobility by CZ using the core sample and the baseline family income definitions.[22]

---

[22] The mean population-weighted standard error of our estimates of β_c across CZs is 0.006; the mean unweighted standard error is 0.023. Because these standard errors are quite small relative to the variance in β_c across areas, we

Absolute Upward Mobility. Panel A of Figure VII illustrates the spatial variation in absolute upward mobility. We construct this heat map by dividing CZs into deciles based on their estimated value of $\bar{y}_{25,c}$. The deciles are color coded so that lighter colors represent areas with higher levels of absolute upward mobility $\bar{y}_{25,c}$.[23] There is substantial variation in upward mobility across CZs. CZs in the top decile have $\bar{y}_{25,c} > 52$. In these areas, children who grow up in families at the 25th percentile of the national income distribution can expect to have incomes above the national median on average – a substantial degree of upward mobility. At the other extreme, CZs in the bottom decile have $\bar{y}_{25,c} < 37.4$. Note that the 37th percentile of the family income distribution for children at age 30 is \$22,900, while the 52nd percentile is \$35,500; hence, the difference in upward mobility across areas translates to large differences in children's incomes. Pooling all CZs, the unweighted standard deviation (SD) of $\bar{y}_{25,c}$ is 5.97; the population-weighted SD is 3.41.

The map exhibits three spatial patterns. First, upward income mobility varies substantially at the regional level. The Southeast exhibits the lowest rates of upward mobility, while the Great Plains has the highest rates of intergenerational mobility. The West Coast and Northeast also have high rates of upward mobility, though not as high as the Great Plains.

Second, there is substantial within-region variation as well. Using unweighted CZ-level regressions of the upward mobility estimates on Census division and state fixed effects, we estimate that 53% of the cross-CZ variance in absolute upward mobility is within the nine Census divisions and 36% is within states. For example, many parts of Texas exhibit relatively high rates of upward mobility, unlike much of the rest of the South. Ohio exhibits much lower rates of upward mobility than nearby Pennsylvania. The statistics also pick up much more granular variation in upward mobility, illustrating their accuracy. For example, South Dakota generally exhibits very high levels of upward mobility, with the exception of a few areas in the Southwest corner of the state. These areas are the largest Native American reservations in the United States and are well known to suffer from very high rates of persistent poverty.

The third generic pattern is that urban areas tend to exhibit lower levels of intergenerational mobility than rural areas on average. For instance, children from low-income families who grow up in the Chicago area have significantly lower incomes at age 30 than those who grow up in rural areas

---

ignore the estimation error in β_c and $\bar{y}_{25,c}$ in what follows.

[23]We cannot estimate mobility for 32 CZs in which we have fewer than 250 children in the core sample, shown in grey in the maps in Figure VII. These CZs account for less than 0.05% of the U.S. population in the 2000 Census. In Appendix Figure II, we present a version of this map in which we pool data from the 1980-85 cohorts to estimate mobility for the CZ's that have fewer than 250 observations in the core (1980-82) sample. The estimates of mobility in the CZ's with missing data are quite similar to those in neighboring CZ's, consistent with the spatial autocorrelation evident in the rest of the map.

in Illinois. On average, urban areas – which we define as CZs that intersect MSAs – have upward mobility of $\bar{y}_{25,c} = 41.7$, while rural areas have $\bar{y}_{25,c} = 45.8$. In interpreting this comparison, it is important to recall that our definition of geography is based on where children grew up, not where they live as adults. Many of the children who grow up in rural areas and move up in the income distribution live in urban areas as adults. On average, 44.6% of children who grow up in rural areas live in urban areas at age 30. Among those who move from the bottom quintile of the national income distribution to the top quintile, the corresponding statistic is 55.2%.

Table III shows statistics on intergenerational mobility for the 50 largest CZs by population. Among these cities, absolute upward mobility ranges from 46.2 in the Salt Lake City area to 35.8 in Charlotte. There is considerable variation even between nearby cities: Pittsburgh is ranked second in terms of upward mobility among large metro areas, while Cleveland – approximately 100 miles away – is ranked in the bottom 10. The cities with the highest levels of upward mobility are quite geographically dispersed: they include Pittsburgh, San Francisco, and Boston. Upward mobility is especially low in certain cities in the "Rust Belt" such as Indianapolis and Columbus and cities in the Southeast such as Atlanta and Raleigh. The fact that children who grow up in low-income families in Atlanta and Raleigh fare poorly is perhaps especially striking because these cities are generally considered to be booming cities in the South with relatively high rates of job growth.

In Column 5, of Table III, we consider an alternative measure of upward mobility based on the non-parametric quintile transition matrix described above: the probability that a child born to a family in the bottom quintile of the national income distribution reaches the top quintile of the national income distribution ($\pi_c^{5,1}$). The ranking of areas based on this statistic is very similar to that based on the mean rank measure of upward mobility. The probability that a child reaches the top fifth of the income distribution conditional on having parents in the bottom fifth is 4.4% in Charlotte, compared with 10.8% in Salt Lake City and 12.9% in San Jose. Similarly, the odds of moving from the bottom fifth to the top fifth are twice as large in Pittsburgh as in nearby Cleveland. Note that if parent income played no role in determining children's outcomes, the quintile transition probability would be $\pi_c^{5,1} = 20\%$. Hence, the variation in rates of upward mobility across areas is large relative to the maximum range of 0 to 20%.

In Column 6 of Table III, we consider another measure of absolute upward mobility: the probability that a child has family income above the poverty line conditional on having parents at the 25th percentile. To construct this statistic, we first regress an indicator for having family income above the federal poverty line in 2012 on parent rank in the national income distribution in each

CZ. We define household size as the maximum household size in 2010-11, where household size is defined as 1 plus an indicator for being married plus the number of dependents claimed. The poverty line threshold is defined as $11; 170 + (household size - 1)*3,960. We then calculate the predicted fraction of children above the poverty line for parents at the 25th percentile based on the slope and intercept in each CZ. This statistic also generates very similar rankings across CZs, confirming that our results are not sensitive to the way in which we measure upward mobility.

Relative Mobility. Panel B of Figure VII presents a heat map of relative mobility. This map is constructed in the same way as Panel A, dividing CZs into deciles based on the rank-rank slope β_c. In this map, lighter areas denote areas with greater relative mobility (lower β_c). Relative mobility also varies substantially across areas. In the CZs in the bottom decile of relative mobility, children born to the highest-income parents are more than 40.2 percentile points higher in the income distribution on average than those born to the lowest-income parents. In the CZs in the top decile of relative mobility, the corresponding gap is almost half as large, at less than 23.5 percentiles.

The geographical patterns in relative mobility in Panel B are very similar to those for absolute upward mobility in Panel A. The unweighted correlation across CZs between the two measures is -0.68; the population-weighted correlation is -0.61. Areas with greater relative mobility tend have better absolute outcomes for children from poor families rather than worse outcomes for children from high income families.

To investigate the connection between absolute mobility and relative mobility more systematically, let $\bar{y}_{pc} \equiv E\left[y_{ic}|x_{ic} = p, c\right]$ denote the expectation of the child's rank given a parent rank of p in CZ c as above. We estimate $\bar{y}_{pc}$ in each CZ non-parametrically as the mean value of y_{ic} for children in each percentile bin of parent income $p = 0, ..., 99$. For each value of the 100 values of p, we then regress $\bar{y}_{pc}$ on relative mobility β_c using an unweighted OLS regression with one observation per CZ:

$$\bar{y}_{pc} = a + \gamma_p \beta_c.$$

In this equation, γ_p measures the association across CZs between a 1 unit increase in β_c (i.e., greater intergenerational persistence) and mean absolute outcomes of children whose parents were at the p^{th} percentile of the national income distribution. A negative coefficient ($\hat{\gamma}_p < 0$) implies that CZs with greater relative mobility generate better outcomes for children who parents are at percentile p of the national distribution on average.

Figure IXa plots the coefficients $\hat{\gamma}_p$ at each parent income percentile p along with a linear fit to the coefficients γ_p. The coefficients $\hat{\gamma}_p$ are increasing with p: CZs with greater relative mobility

21

(reductions in β_c) produce better outcomes for children from lower income families. The linear function crosses 0 at $p = 85.1$, implying that for $p < 85.1$, the coefficient $\gamma_p < 0$. That is, increases in relative mobility (lower β_c) are associated with better outcomes for children who grow up in families below the 85th percentile on average. For families at the 85th percentile, differences in relative mobility across CZs are uncorrelated with the mean percentile rank of the child. For families in the top 15%, living in a CZ with greater relative mobility is associated with worse outcomes on average for children, as $\gamma_p > 0$ in this range. Observe that γ_p reaches only 0.2 for the richest families but is nearly -0.8 for the poorest families. Hence, differences in relative mobility across CZ's are associated with much larger differences in absolute mobility for children from low-income families than high-income families.[24]

Figure IXb presents a schematic that illustrates the intuition underlying the preceding results. This figure plots hypothetical rank-rank relationships in two representative CZs, one of which has more relative mobility than the other. Figure IXa implies that in such a pairwise comparison, the rank-rank relationship "pivots" at the 85th percentile on average. That is, the absolute outcomes of children born to families below the 85th percentile are better on average in the CZ with greater relative mobility. This explains why the maps of absolute mobility at $p = 25$ and relative mobility in Figure VII look similar.

Because the pivot point is very high in the income distribution, differences in relative mobility have a smaller effect on children's percentile ranks in high-income families than low-income families.[25] One potential explanation for this pattern is that the rich are able to insulate themselves from differences in local conditions to a greater extent than the poor. If the differences in relative mobility across areas are caused by differences in local policies, this result suggests that one may be able to improve the outcomes of children from poor families without hurting children from high income families significantly.

---

[24]If the rank-rank relationship were strictly linear, the relationship plotted in Figure VIIIa would be perfectly linear and $\gamma_{100} - \gamma_0 = 1$ mechanically. The slight deviation from linearity at the bottom of the distribution evident in Figure VI generates the curvature in Figure VIIIa and the slight deviation of $\gamma_{100} - \gamma_0$ from 1.

[25]It bears emphasis that this result applies to percentile ranks rather than mean income levels. Because the income distribution has a thick upper tail, a given difference in percentile ranks translates to a much larger difference in mean incomes in the upper tail of the income distribution. The probability that children of affluent parents become very high income "superstars" may therefore differ significantly across areas, an interesting question that we defer to future research.

II.E.4 Robustness and Sample Selection: The Role of Mortality

One concern with our construction of income rank measures of mobility in each CZ is that differences in mortality are known to vary across regions. Indeed, one potential channel through which tax expenditures could impact mobility is through a mitigation of health inequality. However, we only observe incomes on parents and children who are alive. This potentially induces a significant selection problem that both biases our estimates of mobility by CZ and also masks the impact of tax expenditures on mobility.

Therefore, we also construct measures of mortality rates by income quantiles by CZ. In Appendix Table IV, we compute the 3 year mortality rate for adults age 25-60 by income quintile within each CZ . Fortunately, we find that health inequality across areas has little explanatory power towards our measures of mobility. This suggests sample selection problems are not driving our results and therefore we feel comfortable studying our baseline measures of income mobility for our main analysis.

In addition to robustness to mortality, Appendix F assesses the robustness of the spatial pattern in mobility to alternative measures and samples. Overall, we find our measures of mobility are highly robust to alternative specifications.

III The Impact of Tax Expenditures on Mobility

Having established robust measures of intergenerational mobility at the CZ level, we now turn to primary analysis of the impact of tax expenditures on intergenerational mobility. We start with the link between overall tax expenditures and intergenerational mobility and then turn to specific components of tax expenditures.

III.A Overall Tax Expenditures

We first analyze overall tax expenditures. We are interested in the effects of both the level of tax expenditures and the progressivity of tax expenditures on intergenerational mobility. We measure the level tax expenditures at the local level (CZ) as the ratio of aggregate itemized deductions to aggregate AGI in the CZ and measure progressivity of tax expenditures as the difference in the percentage of aggregate itemized deductions relative to aggregate AGI in the CZ for top bracket taxpayers (AGI above \$200,000) to low bracket taxpayers (AGI below \$10,000).

Figure X displays a binned scatterplot of the relationship between CZ aggregate tax expenditures as a percentage of AGI in 2008 and the CZ IGE as measured by the correlation between

parent rank income and child rank income. To generate the binned scatterplot, we group CZs into centiles (one-hundred equal-sized bins) on tax expenditures as a percentage of AGI, weighting by CZ population. The dots represent the weighted means of the IGE and tax expenditure measure. The best-fit line is calculated from a regression on the CZ level data and shows a negative relationship between the local level of tax expenditures and the rank-rank correlation.

Next, using a similar structure, Figure XI displays a binned scatterplot of the relationship between progressivity of CZ tax expenditures and IGE as measured by the correlation between parent rank income and child rank income (CZs with over 300% difference in tax expenditures are excluded from the figure and best-fit line). Again, the best-fit line shows a negative relationship between the progressivity of tax expenditures and the rank-rank correlation.

The negative relationships depicted in Figures X and XI suggest that places with higher or more progressive tax expenditures have more inter-generational mobility, i.e., a lower correlation between parents' income rank and children' income rank. To formally measure the effects of different tax expenditures on intergenerational mobility, we use OLS regressions of the form:

$$\beta_c = a + b * EXPEND_c + \gamma X_c + \epsilon_c$$

for CZ c, where β_c is relative mobility in the CZ (i.e. the parent rank-child rank correlation using within-CZ income centile ranks described above), $EXPEND_c$ is the measure of tax expenditures of interest described above, and X_c is a vector of CZ characteristic controls including CZ median income and percentage of the population that is a 4-year-college graduate, white, black, Hispanic and other.

Table IV reports estimates of b for the level and progressivity of tax expenditures as a percentage of AGI, weighting by the population in each CZ. Column 1 reports the results of a regression of IGE on CZ tax expenditures. The coefficient is negative and significant; CZs with higher tax expenditures have significantly lower parent-child income correlation, i.e. higher intergenerational mobility. A one standard deviation increase in CZ percentage tax expenditures, 4.09% of AGI, decreases CZ parent-child income correlation by .5 percentage points, relative to the CZ national mean of 0.34. This result is robust to inclusion of demographic controls in column 2. The coefficient however is significantly smaller when state fixed effects are included in column 3.

To study the progressivity of tax expenditures, in Columns 4-6 we replicate the analysis in Columns 1-3 using the difference in mean percentage tax expenditures for filers under $10,000 AGI

and over \$200,000 AGI.[26] Progressivity of tax expenditures has a similar effect on intergenerational mobility. A one standard deviation increase in CZ difference between lowest and highest bracket tax expenditures, 45.3% of AGI, decreases the correlation between parent and child incomes by 0.07 percentage points relative to the CZ national mean of 0.34. Including demographic controls and state fixed effects decreases the magnitude of the coefficient, but the effect remains significant.

Overall CZ levels of tax expenditures and progressivity of tax expenditures are positively related to intergenerational mobility. Our analysis demonstrates that places with high and more progressive tax expenditures have lower correlation of parent-child mobility and higher intergenerational mobility. Tax expenditures include a large number of different tax components, which may individually have different impacts on intergenerational mobility. For this reason, we turn to analysis of three specific tax expenditure components: mortgage interest deductions, state income taxes, and state EITCs.

The relationship between tax expenditures and intergenerational mobility may not be causal if the OLS identification assumptions fail to hold. Omitted factors may explain both higher local tax expenditures and greater intergenerational mobility. The potential problems with a causal interpretation of our results should be kept in mind throughout our analysis of specific tax expenditures.

III.B   Specific Tax Expenditures

Tax expenditures include a number of components. The two most important ones quantitatively are (1) mortgage interest deductions, (2) state and income local tax deductions. Hence, we focus on these two tax expenditures.

Mortgage interest deductions. Mortgage interest deductions are the largest federal tax expenditure. These deductions may impact economic opportunity by providing opportunities for credit-constrained middle and low-income families to become homeowners.

In Columns 1-3 of Table V, we report estimates for the effect of CZ mortgage interest deductions on intergenerational mobility. We find a negative and statistically significant effect of CZ mortgage interest deductions on parent rank-child rank correlation that is robust to the inclusion of demographic controls and state fixed effects. The effect is comparable in size to the effect of overall tax expenditures: a one standard deviation increase in mortgage interest deductions as a percentage of AGI decreases the parent-child IGE by 1.31 percentage points, relative to the national mean of 0.34. Columns 4-6 repeat the analysis using inequality of mortgage interest deductions as

---

[26]Three outlier CZs with very low AGI totals (904, 101, and 830) are excluded from the regressions in Table IV, Columns 4-6.

measured by the level difference in mean mortgage interest deductions for the highest and lowest AGI classes, including controls for local housing prices from 2000 Census estimates. The basic regression reported in Column 4 yields a statistically insignificant coefficient – implying that areas with relatively larger mortgage interest deductions by high relative to low income taxpayers, i.e. more regressive mortgage interest deductions, are more economically mobile. This result is not statistically significant; however, the inclusion of demographic controls and state fixed effects (Columns 5 and 6) does render it statistically significant, but with a smaller magnitude than the direct relationship to the level of average mortgage interest deductions in Columns 1-3.

In sum, there is some evidence that CZs with larger mortgage interest deductions as a fraction of AGI are more economically mobile. It is possible, however, that this relationship is not causal if mortgage deductions are correlated with other omitted factors related to intergenerational mobility. Further research isolating quasi-experimental variation in mortgage deductions is needed to understand the causal impacts of such deductions more precisely.

State income tax rates. Itemized deductions include state and local income taxes. State and local income taxes depend on both the level of income of individuals and the local or state income tax rate. Hence, this component of itemized deductions is naturally endogenous to income. To eliminate this endogeneity issue, we focus instead on the state tax rate policy. We measure the level of state income taxes by the marginal tax rate in the state for a taxable income level of $40,000 and measure progressivity of the state income tax with the difference between the state top marginal tax rate and state marginal tax rate for a taxable income of $20,000.

Table VI presents results of an analysis of state income tax rates and IGE. In Column 1, we find that a 1% increase in state income tax rate decreases the intergenerational income correlation, i.e. increases intergenerational mobility, by 0.2pp (compared to the mean of 0.34). This coefficient falls slightly (but insignificantly) when including demographic controls in Column 2. In Columns 3-4, we find that states with more progressive individual income tax rates have statistically significant higher intergenerational mobility (negative coefficients), which is again robust to the inclusion of demographic controls.

Both the level of state income taxes and its progressivity positively affect mobility (i.e. lower the IGE correlation). A natural potential explanation for this relationship could be alleviating credit constraints by taxing higher incomes and redistributing toward credit constrained lower incomes with higher educational expenditures. However, the relationship may not be entirely causal if these aspects of state taxes are correlated with other characteristics that could partly drive the results.

For example, states with higher and more progressive state taxes may also have other state policies promoting economic opportunity and mobility.

III.C   Local Policy: State EITC

To further analyze the role of local income tax policy, we next focus on the largest state tax level transfer program, the state EITC. The federal EITC is a refundable tax credit aimed at low-income families. Eligibility for the federal EITC is determined by total earnings and the number of qualifying children.[27] Twenty-three states and the District of Columbia offered state EITCs in 2008, motivated by evidence of the impacts of the federal EITC on outcomes for low-income taxpayers (see Meyer 2010 for a survey of the literature). State EITCs "piggyback" on the federal EITC and offer a fixed percentage of the federal credit.[28]

As discussed in Section 2.2, our primary measure of the State EITC is the mean exposure to the EITC for a child when growing up. We also consider a measure of the State EITC rate in year 2008. Table VII reports the results. We find a robust positive relationship between the size of state EITC top-ups and intergenerational mobility. A 1 standard deviation increase in exposure to the EITC corresponds to a 0.14pp reduction in the rank-rank slope (and hence an increase in mobility), relative to a base of 34pp. Column 2 of Table VII shows a smaller relationship if one uses the EITC rate in year 2008 instead of the exposure-weighted EITC rate. Finally, column 3 shows that most of this negative correlation is driven precisely by the EITC exposure, not so much by the EITC rate in 2008. This is consistent with the idea that exposure to the EITC when growing up leads to greater intergenerational mobility. However, as discussed below, one needs to also take into account the potential that other confounding factors may be both correlated with state EITC policies and intergenerational mobility.

III.D   Specification Tests and Robustness Checks

As discussed in Section 2, there are many potential ways to define geography when studying the relationship between tax expenditures and intergenerational mobility. In this section, we show that our general results are robust to alternative definitions of mobility and alternative geographical boundaries. For simplicity, we conduct the analysis by looking at three tax variables defined in

---

[27]See IRS Publication 596 (Internal Revenue Service 2011) for details on federal program eligibility and rules.

[28]Minnesota offers a varying rate of the federal EITC credit depending on income and Wisconsin offers a varying state EITC based on the number of children. For our analysis, Minnesota is assigned its average rate of 33% and Wisconsin is assigned the 4% rate for single child families. For more information on state EITCs, see Levitis and Koulish (2008).

Section 2.2: the local tax rate, the of EITC exposure, and tax progressivity (difference between top state tax rate and tax rate for $20,000 of taxable income).

To begin, we explore the relationship between quintile transition matrices and intergenerational mobility. In contrast to relative or absolute mobility, we measure mobility as the chance someone from the bottom fifth of the income distribution grows up to reach the top fifth of the income distribution.

The results are presented in Table VIII. Here, we find significant positive relationships with both the local tax rate and EITC exposure (0.346 and 0.152 respectively). Both variables in the regression are normalized by their standard deviations, so the coefficient of 0.346 says that a 1 standard deviation increase in the local tax rate correlates with a 0.346 standard deviation increase in the chance of reaching the top of the income distribution for someone whose parents were in the bottom quintile. Fortunately, these results are quite similar to the results we obtain using our measures of absolute and relative mobility, and suggest our findings of a positive relationship between tax expenditures and mobility are robust to alternative measures of mobility.

An additional, and potentially very significant, concern with our estimates is that our measures of intergenerational mobility are constructed at a very aggregated level (e.g. the Boston CZ includes Cape Cod). As a result, we may be missing the majority of the variation in intergenerational mobility across the US, which would lead us to provide underpowered and potentially downwardly biased measures of the impact of tax expenditures on intergenerational mobility.

As a result, we conduct a robustness analysis using measures of intergenerational mobility at the county level, as opposed to the CZ level. These measures are constructed exactly analogous to the CZ measures, except using the county as the measure of geography. Fortunately, the tax expenditure and tax statistics are readily available at the county level. For reference, the corresponding mobility statistics are presented in Appendix Table VII.

Table IX focuses on the correlation of the same three tax expenditure measures with absolute upward mobility. Here, we again find significant effects of both the local tax rate and the EITC exposure. Column 4 shows that these two factors are both significant in a multivariate regression; and the tax progressivity coefficient remains insignificant. In short, the evidence suggests the spatial patterns we have identified between tax expenditures and intergenerational mobility are robust to a finer definition of geography.

Finally, an alternative concern is that much of our analysis includes mobility measures from fairly small, lesser populated areas. Therefore, Appendix Table VIII reproduces our analysis using

MSA border definitions instead of CZs. While MSAs are a common and natural measure of local geography and are in general of similar size to CZs, they do not cover the entire U.S. Indeed, we have only 381 observations using the MSA definition of geography.

Table X reports the results. Here again we find a significant relationship between both the local tax rate and EITC exposure and absolute upward mobility, with coefficients loosely similar to the other specifications (0.149 and 0.322). In the multivariate regression, the EITC exposure measure continues to be statistically significant. However, the local tax rate is still positive but now statistically insignificant, perhaps due to the smaller sample size using MSAs.

In short, our results suggest a strong and robust spatial relationship between tax expenditures and intergenerational mobility.

III.E   Controlling for Confounding Factors

While the analysis heretofore is promising, they are all subject to the caveat that alternative correlated factors could be driving both variation in tax expenditures and variation in intergenerational mobility. Such an endogeneity problem would lead to inconsistent estimates of the impact of tax expenditures on mobility.

In order to examine in more depth whether omitted variable bias can affect our estimates, we analyze how adding observable local area variables can affect the correlation between the tax policy variables and intergenerational mobility that we documented above. Any factor that varies across areas and has a direct impact on intergenerational mobility can potentially confound our OLS analysis of the relationship between tax expenditures and intergenerational mobility. It is therefore imperative that we attempt to control for such factors in our analysis.

III.E.1   Race and Income

Perhaps the most obvious spatial pattern is that intergenerational mobility is lower in areas with a larger minority population, such as the Southeast. We therefore begin our analysis of potential confounding factors by exploring the role of race in upward mobility. Of course, we do not observe race in the data. However, we use census data on racial shares by zipcode to identify people who live in areas where a large fraction of the population identifies as non-hispanic white. For a given thresholds, x, we re-estimate our measures of relative mobility by CZ on the subsample of people in zipcodes (zip5s) for which at least $x\%$ declare their ethnicity to be non-hispanic white, where x ranges between 0 and 100.

The results are presented in Figure XII. If the variation in upward mobility across areas were entirely driven by heterogeneity in outcomes across race at the individual level, the coefficient in Figure XII would fall to zero as x converged to 1, as illustrated by the dashed line. Intuitively, if all of the spatial variation were driven by individual-level differences in race, there would be no spatial variation left in a purely white sample. The data reject this hypothesis: even in the subsample of people who live in zipcodes where more than 95% of individuals declare themselves to be non-hispanic white, the regression coefficient remains at 0.89. Put differently, all races living in areas with larger minority populations have lower rates of upward income mobility. This suggests that properties of the place (such as variations in tax expenditures) are driving the spatial differences in mobility, not differences in race.

In addition to race, an additional confounding factor is the role of differences in income levels across areas. For example, Table III reveals that many cities with high mobility tend to be thought of as richer cities (e.g. Boston, San Francisco, etc.). To explore this further, we construct mean income characteristics by CZ (reported in Appendix Table IV) and control for these characteristics in our main regressions. For robustness, we include separate controls for both the mean child incomes in our sample and the mean parent incomes in our sample. We have also explored additional controls reported in Appendix Table IV and obtained similar results.

Table XI reports the results. The first four columns of Table XI illustrate that our main results are robust to controlling for mean parent income in the CZ. Both the local tax rate and the EITC exposure variable remain highly statistically significant. Moreover, the measure of tax progressivity remains positive and of similar magnitude (0.204), but not statistically different from zero.

Columns 5-8 report controls for mean child family income. Here, the results are quite different than controlling for parent income, but broadly consistent with the story that tax expenditures has an impact on intergenerational mobility. None of the three variables are statistically different from zero (0.026, 0.005, and -0.026 are all insignificant from zero). However, there is good reason to believe that controlling for child income is "over-controlling". Indeed, one expects tax expenditures to positively affect mobility *through an increase in child income*. Hence, controlling for average child income in a CZ, one does not find a robust positive relationship between mobility and these measures.

III.E.2   Other Factors

Having established that our results are not confounded by differences in income levels across areas, we now analyze a broader set of potential confounding factors. We guide our search around a set of mechanisms that have been discussed in the sociology and economics literature: (1) race and its geographic concentration within the CZ, (2) income inequality, (3) primary school quality, (4) access to higher education, (5) migration and networks, (6) social capital, and (7) family structure. Because most of these factors are slow-moving and we have estimates of intergenerational income mobility for essentially one birth cohort, we focus on cross-sectional correlations rather than changes over time. For most covariates, we use data in year 2000, when the children in our sample were roughly 18 years old, because many variables (e.g., school quality and access to local higher education) cannot be consistently measured in earlier years. We verify that we obtain similar results using data from 1990 for selected variables. Most of these statistics are publicly available. However, we also include several measures of confounding variables from tax data that are in Appendix Table IV (see the notes at the bottom of this table for details).

We begin by exploring the extent to which these factors could be confounding our results. Table XII reports the raw correlation of our measures of intergenerational mobility with variables that are potential confounding factors. To guide this comparison, the first sets of rows presents the results for several of the tax variables discussed above: the local tax rate, tax progressivity, and EITC exposure.

The first 5 columns report correlations with absolute upward mobility (as opposed to relative mobility) using several specifications: a baseline specification with no controls (Column 1), a state FE specification that includes state fixed effects (Column 2), a specification that weights each CZ proportional to its population (Column 3), a specification using only urban areas that overlap with an MSA (Column 4), and a specification that includes controls for the share of minorities in the CZ and income growth between 1990 and 2006-2010 drawn from the ACS. Column 6 reports the results with relative mobility but to provide robustness relative to Sections 3.1-3.4, we do not weight CZs by the fraction of the population nor do we include any additional controls.

The results here reveal five sets of factors that have the strongest correlations with mobility, and hence the strongest potential to confound our analysis of the impact of tax expenditures on mobility: the geographic concentration of race and income within a CZ, income inequality, primary school quality, social capital, and family structure.

31

To explore the extent to which these factors could be confounding our results, Table XIII presents a multivariate regression that includes the three main tax expenditure variables and variables from each of these five categories. For simplicity we focus on absolute upward mobility and three measures of tax expenditures: State EITC rates, the progressivity of the local tax schedule, and the local tax rate. Column 1 reports a positive impact of EITC exposure on upward mobility – consistent with the EITC impacts on relative mobility reported above. Column 2 includes all three of these tax expenditure measures. Here, we see they all have positive effects, albeit the tax progressivity measure is not statistically significant. Column 3 reports coefficients on the potential confounding factors: income inequality (as measured by the Gini coefficient computed on those with income below the 99th percentile), the geographic concentration of races within the CZ, school quality as measured by the high school dropout rate, a social capital index, and the fraction of single parents in the CZ.

Controlling for these five potential confounding factors reveals some surprising results. Although the EITC exposure is significant with no controls, it is no longer significant after controlling for these potentially confounding factors (Column 4). However, in Column 5 we show that measures of tax progressivity and the local tax rate are significantly related to upward mobility, even after controlling for all five of these potential confounding factors. Taken together, this suggests tax expenditures can have significant impacts on intergenerational mobility.

III.F   Further Evidence on Causal Effects: Cross-Area Movers

In addition to potentially confounding factors explored in the previous section, another potential issue is that the variation in mobility rates across areas does not reflect a difference in causal outcomes. Put differently, if it's just the case that different types of people live in different places (e.g. "low mobility people" live in low mobility areas, etc.) then there would be little or no scope for areas with low mobility to benefit from greater tax expenditures that aim to promote upward mobility.

In this subsection, we present an analysis that explores the scope of the potential impact of place-based tax credits on children's outcomes. We do so by examining the extent to which the spatial variation in mobility (which we know is correlated with differences in tax expenditures) reflects the causal impact of the location in which a child grew up on his/her economic outcomes in adulthood.

To identify the causal impact of locations on children's' outcomes, our approach utilizes a new

quasi-experimental design that analyzes the experience of children that move to different places at different stages of their lives when growing up. For example, suppose two children from the same location and same parental income level move to two different places. Figure XIIIa shows that the outcomes of these children who move before age 13 are correlated with the expected outcomes of children with the same parental income in the destination location.[29] This consistent with the idea that variation in tax expenditure policies across areas can have a causal effect on children's outcomes. Figure XIIIb plots the slope of this relationship by the age of the child at the time of the move. Here, we see a pattern precisely consistent with places having causal effects on children's outcomes. A child that moves at age 11 has outcomes in adulthood that are correlated at 0.77 with the destination location expected outcomes. Every additional year of exposure corresponds to capturing roughly 4 percentage points of the difference in outcomes between permanent residents in the origin and destination CZ. Extrapolating these results to birth, this suggests that differences in kid outcomes across places are roughly 75% causal. Table XVIII presents coefficients for the regression estimates shown in Figure XIIIb across a wide range of alternative samples and specifications (documented at the bottom of the table) and illustrates illustrates the robustness of this patter.

Table XVII parameterizes the relationship in Figure XIIIb using a piecewise linear specification that pivots at age 23. We report the slope coefficient normalized by 23, so that the 0.0396 slope corresponds to the fraction of the mobility characteristics of the stayers the mover could inherit by moving one year earlier to the destination. The table again illustrates the robustness of this exposure time pattern. Across a wide set of specifications of alternative income definitions, samples, and controls, we estimate coefficients around 0.04.

A priori, one primary concern with our empirical approach is selection bias. This would be

---

[29]Appendix Table XVIII provides the summary statistics on the samples used for this graph and the subsequent analysis in this section. Our baseline analysis restricts to moves greater than 100 miles between CZs with populations of at least 250,000 people based on the 2000 Census. We focus on the sample of movers who have only 1 origin and 1 destination CZ and stay in the destination for at least 2 years (i.e. move prior to 2011 in our sample). This results in a sample of 6.9M movers, roughly 3.2M of which we observe at ages 24 and above. Appendix Table XVII illustrates that our baseline results are robust to alternative sample selections.

Also, note we are measuring place outcomes now using the average outcomes of the stayers (families who do not move from a CZ over 1996-2012) in a CZ, not the entire population. On average, stayers are a fairly well-defined population. 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. Appendix Figure VIIIa illustrates that the relationship between child and parent ranks remains linear when restricting the calculation to the stayers. Appendix Figure VIIIb illustrates that the spatial pattern of mobility is similar for the stayers as opposed to the entire population. Moreover, we estimate the population weighted standard deviation of these outcomes are 3.6 percentiles at the CZ level for below-median income parents and 2.8 for above-median income parents. The correlation in outcomes between below and above-median income families is 0.56. Appendix Figure VIIIb also reports estimates at the county level for the New York Combined Statistical Area.

a problem if good parents move to good places disproportionately when their children are young (sorting), or if good things happen to families that move to good places (e.g. parents get a good job, etc.). These concerns motivate a set of additional robustness tests, all of which suggest that the basic exposure-time design provides a valid measure of the causal effects of places.

First, we control for sources of bias using family fixed effects and controls for changes in parental income and changes in marital status. Figure XIV reports the baseline design after including family fixed effects and controls for changes in marital status. The slopes are also reported in Table XVII (Row 25), where they can be compared to the baseline (row 1 or Table XVII). The results are virtually unchanged when identifying these effects within the family across siblings, which is inconsistent with the story that the results are driven by sorting patterns of families.

Second, we exploit natural experiments that generate moves. To do so, we identify large outflows from counties and illustrate that our estimates are quite similar when restricting to moves generated in these instances. We define the percentage outflow as the difference between the yearly outflow and the average outflow across years 1997-2010, normalized by the variance of the outflows across years in the zipcode. We then repeat our slope estimates (Row 1 of Table XVII) separately for samples that restrict to larger county outflows, ranging from above 50% to top 2% of the flow distribution within a parent-decile by zipcode cell.[30] The results are presented in Table XX. We instrument for the change in the child's income rank with the average change in place effects for movers from that zipcode, generating plausibly exogenous sources of variation in place impacts across CZs. Table XX shows that across a wide range of sample selections (i.e. larger displacement shocks), we find consistent estimates between 0.03 and 0.04, similar to our baseline results.

Third, we can conduct several placebo tests. A key prediction of the hypothesis that the slope represents the causal effect of place is that the place effect for various outcomes should best predict the outcomes. For example, places differ both in the mean ranks of child outcomes and in the fraction of childrens that reach the top 10% of the income distribution. . If parents are sorting, then they should be sorting on various measures of the place so that the place effect on both employment and mean ranks should be correlated with employment or child ranks.[31] In contrast,

---

[30]We have verified that many of these large outflows occur in areas that experienced plant closures and the result of Hurrican Katrina in 2005. The mean sample size within a parent decile-by-zipcode-by-year cell is 42 (median is 25). To ensure the results are not driven by a bias towards OLS due to the many instruments problem, we have replicated the analysis restricting to cells with at least 50 children and obtained similar results that are statistically indistinguishable from results presented here.

[31]Places differ not just in children's mean outcomes, but also in the distribution of outcomes. For example, consider children who grow up in Boston and San Francisco in families at the 25th percentile of the national income distribution. In both of these CZs, children's mean percentile rank at age 24 is 0.46 . However, children in San Francisco are more likely to end up in the upper or lower tail of the income distribution. The probability of reaching

if the pattern is causal, one would expect the corresponding outcome-based place effects to be sufficient predictors of child outcomes. To illustrate this, Figure XV plots the relationship between unemployment and both mean ranks of stayers and mean unemployment of stayers. As illustrated, the unemployment place effect predicts unemployment; and conditional on the employment effect the mean rank effect has minimal predictive content. Panels A and B of Table XXI reports placebo tests for other outcomes, including the event of earning above 90th percentile of the earnings distribution.[32]

Similarly, one can also exploit the fact that average outcomes for children vary across genders in different areas. Figure XVI reports the difference between absolute upward mobility for boys versus girls, with darker shades indicating worse places for boys relative to girls. Although mobility for girls and boys is highly correlated (0.9 for median-income families) across CZs, there is still considerable variation in the relative outcomes of boys versus girls across CZs.[33] If the exposure effects reflect the causal effects of places, we would expect the exposure impacts to be driven by variation in the child's own gender place effect, not the place effect for the other gender. Table XXI reports the relationship between the expected outcomes based on a child's own gender versus the opposite gender. The table clearly shows that the child's own gender place effects are driving the results. Hence, for the results to be driven by time-varying sorting patterns, it must be the case that parents are sorting to places in a gender-specific manner. Put differently, it must be the case that if good things happen to parents who have girls then they go to places that are good for girls; if good things happen to parents who have boys, then they go to places that are good for boys. In short, this type of sorting pattern is highly unlikely, and is more likely to reflect the causal impact places have on child outcomes. Columns 4-7 of Table XXI illustrates the robustness of this result to the inclusion of family fixed effects.

Fourth, we can exploit the fact that place effects vary by cohort. Some places are improving; others declining over time.[34] If places are having causal effects, one would expect the best predictor of these outcomes to be the expected outcomes in the child's own cohort, as opposed to neighboring

---

the top 10% is 7.3% in San Francisco vs. 5.9% in Boston; the corresponding probabilities for the bottom 10% are 15.5% and 11.7%.

[32]Table XIX reports additional statistics, analogous to the datapoints in Figure XV, for the series corresponding to Table XXI.

[33]For example, the difference in outcomes between males versus females is high in Syracuse and Albany, NY (i.e. comparatively good for males versus females), and low in Milwaukee, WI (i.e. comparatively good for females relative to males).

[34]Place effects are generally quite stable across cohorts: the autocorrelation of 0.95 at p=25 and 0.92 at p=75 . Good places in one year are, on average, good places in the next year. However, outcomes in some areas (such as Oklahoma City) have improved over time, while others (such as Sacramento) have gotten worse.

cohorts. In contrast, if families are selecting locations based on their general mobility characteristics, it's unlikely that at the time of the move they are sorting differentially with respect to the later outcomes of children in their own child's cohort relative to children, say, who are one year older in the destination neighborhood. Figure XVII regresses the child's rank outcomes on the expected outcomes for their own cohort and other neighboring cohorts. As seen in the figure, the impact of the place is driven entirely by the child's own cohort, consistent with the hypothesis that the impact is driven by the causal effect of places.

Fifth, while most of our outcomes of interest are earnings later in life (and thus are not varying at high frequency around the time of the move), we can also look for impacts of moving on outcomes readily measurable around the time of the move. To do so, we focus on teen labor. As shown in Table XII, areas with a higher share of teen labor rates tend to have higher mobility (the correlation is 0.631). Figure XVIII asks whether children that go to areas where more children work at age A tend to work at that age. As shown, children that move to areas where more children work (e.g. in summer jobs) tend to have discontinuously higher rates of working around the time of the move. Children that move at age 14 or 15 to a place where most 16 year olds work are much more likely to be working at age 16 than children that move to the same location at age 16 or 17[35]. This suggests either places are having causal effects on teen labor rates, or the unobserved variables are varying discontinuously around the time of the move, which is arguably quite unlikely.

Other Geographic Levels   In addition to analyzing moves at the CZ level, we also examine the role of moving at smaller geographies such as county. Appendix Table XVII illustrates the robustness of the basic design: children that move to areas that have higher outcomes for children tend to do better the earlier they move to the good destination. We estimate a slopes generally in the 0.03-0.04 range, consistent with our findings at the CZ level. Moreover, we find minimal attenuation when controlling for family fixed effects at these levels. This suggests our focus on commuting zone, as opposed to more de-aggregated definitions of geography, is not biasing our conclusions that the spatial pattern of intergenerational mobility reflects the causal effect of places.

Direct Estimates of Exposure Effects   Our estimates suggest that roughly 75% of the variation in mobility across areas is causal, in the sense that moving to areas where the prior residents are

---

[35]We also explore impacts on teen birth outcomes the birth certificate database. Comparing the data to population birth records from the CDC suggests that the 2008-2012 records appear to miss roughly 10% of births in the U.S. But, reassuringly we find similar exposure effect patterns highly correlated with our baseline estimates.

doing better will do better. To the extent to which tax expenditures vary across these areas, as illustrated above, this illustrates the potential role of these expenditures in promoting mobility. However, if places have different effects on movers than prior residents, our focus on the 'stayers' as the measure of how well children are doing potentially understates the true variation in the extent to which places and hence variation in tax expenditures explain mobility.

As a complementary approach, we refine our estimation strategy as follows. Consider the set of people that move from place A to B. How much better do the children do that move at a young age, which generates significant exposure time, as opposed to those that move at an older age and hence have less exposure time to place B (and more exposure to place A). We estimate these exposure time effects for each pair of origins and destinations in the data. Appendix Table XVI presents the estimates for each pair of CZs. For example, the coefficient of 0.44 in column (3) means that a child at the bottom of the parental income distribution that spends 1% more of their childhood in CZ 200 (Morristown, TN) relative to CZ 100 (Johnson City, TN) will on average have incomes that are 0.44 percentiles higher in the national income distribution. The coefficient in column (4) illustrates how this relationship varies with parental income. The coefficient of -1.0 indicates that the impact of moving to CZ 200 relative to CZ 100 is 0.44 - 1 = -0.56. In short, Morristown, TN increases incomes of poor children that move from Johnson City, TN poor children but reduces incomes of rich children that move between these locations. Columns (6) and (7) repeat the exercise using college attendence as the outcome as opposed to the child's rank in the income distribution. The patterns are similar, but generally more precise because college attendence can be measured at a younger age (here, we use age 19), which generates larger sample sizes.

Building on this analysis, we estimate a full set of fixed effects of each CZ and county. To do so, we simply replace the permanent resident outcomes in our regressions above with separate fixed effects for each geographic location. Appendix Table XVIII (part 3) reports the summary statistics for the estimation samples at the CZ and County level. To construct a robust picture of the exposure effects on children's outcomes, we estimate these fixed effects using a wide range of specifications and alternative outcomes. Appendix Table IV contains the estimates by CZ, and Appendix Table VII contains the estimates by CZ.[36]

---

[36]Each fixed effect begins with the prefix "Bj", followed by the parent income percentile (e.g. p25 corresponds to parents at the 25th percentile of the income distribution), followed by the geography (cz or cty), and the specification. For example, Bj_p25_cz_kr26_cc2 corresponds to the fixed effect of the CZ on a child's income rank at age 26 (kr26). In addition to the fixed effects based on movers, we also include outcomes for permanent residents. For example, e_rank_b_p25_kr26_cc2 corresponds to the expected rank outcomes for a child's income at age 26 (kr26) with parents at the 25th percentile of the income distribution (p25). The Appendix tables contain measures of mobility on family and individual income, separately for males and females, provide estimates by cohort, single and two-parent households,

Broadly, our estimated fixed effects confirm the robustness of our initial findings of the spatial patterns of mobility. We estimate that the standard deviation of place effects on a child's income rank at age 26 is 0.15 percentiles per year at the commuting zone level, 0.11 across counties within a CZ for children in below-median income families. And, we find that across all of the specifications they are positively correlated with the baseline measures of mobility used in the regressions in our primary analysis in Tables IV-XI. As a result, we consider this robustness check verifying our initial findings of the spatial patterns of mobility and their relationship to tax expenditures.

III.G The Mortgage Interest Deduction, the Low Income Housing Credit, and the Moving to Opportunity Experiment

One of our key findings, illustrated in Table V and discussed above, is the positive correlation of intergenerational mobility with the mortgage interest tax deduction. This is consistent with a potential role of providing these tax credits to owner occupied housing. However, there were two key concerns about drawing conclusions about the mortgage interest deduction correlations. First, these are correlations and may not be causal. Second, it's not clear whether the tax expenditures need to be tied to home ownership, or whether a tax expenditure for housing costs (regardless of ownership) would provide similar benefits for mobility. Indeed, similar questions of efficacy in promoting mobility can be asked of other government programs like the Low Income Housing Tax Credit (LIHTC) and Section 8 Housing Choice Vouchers, which are designed to increase demand and supply for affordable rental housing for low-income families. For example, under LIHTC, developers receive tax credits for specific projects that fulfill one of two conditions: at least 20% of all units must be occupied by tenants with incomes below 50% of the Area Median Gross Income (AMGI), or at least 40% of units must be occupied by tenants with incomes below 60% of AMGI.

To provide guidance on these issues, this section explores the impact of the Moving to Opportunity experiment.[37] This experiment randomly allocated Section 8 housing vouchers in 1994-1997 to families in 5 cities: Baltimore, Boston, Chicago, Los Angeles and New York. There was one control group and two treatment groups. In one treatment group, individuals were offered a Section 8 housing voucher. In the second treatment group, individuals were offered the voucher but required to move to a census tract with less than 10% poverty rates.

By randomly allocating individuals to receive the vouchers, this project does not suffer from

---

impacts on marriage, and impacts on college attendance. All estimates based on less than 250 observations are removed.

[37]See http://www.nber.org/mtopublic/ for an overview of the MTO experiment.

the selection issues of our correlational analysis above. One can interpret the difference between the treatment and control groups as the causal impact of providing the voucher to obtain housing. Indeed, much previous work has explored the impact of this program. In general, they find no significant impact of these policies on the earnings of adults and also found no significant impacts on outcomes for the children, such as test scores.

Here, we revisit this experiment with the ability to look at its impacts on earnings later in life for the children who obtained the most exposure to the new neighborhoods as a result of their housing vouchers. To do so, we utilize data provided to us by the Department of Housing and Urban Development (HUD) to link the families in this program to the tax data. The movers analysis above suggests that places have causal effects in proportion to exposure time to the location. As a result, one would expect larger impacts of the MTO experiment on those for whom it generated larger differences in exposure to good areas – namely those that were younger at the time of random assignment. Therefore, we focus on two subsets of children: those that were pre-teenagers at the time of random assignment (age ≤ 13) and teenages (age 13-18) at random assignment. HUD collected social security numbers (SSNs) prior to RA for 90% (11,892) of the individuals who participated in the MTO experiment and were born in or before 1991. The MTO records were linked to the tax data by SSN. Of the MTO records with a valid SSN, 99% (11,780) were successfully linked to the tax data. To protect confidentiality, individual identifiers were removed from the linked dataset prior to the statistical analysis. We match 86.4% of younger children and 83.8% of the older MTO children to the tax data.

Table XXII discusses the match rates and presents the summary statistics from a combination of tax data and HUD surveys. As illustrated in column (2) and (3), the match rates are not significantly different in the treatment groups (section 8 and experimental groups) relative to the control group. Moreover, as shown in the remaining rows of Table XXII, the children's characteristics do not vary significantly across treatment and control groups, as one would expect given the random assignment of treatment status.

We focus our analysis on several child outcomes, including household and individual income at ages 24-28, college attendence at ages 18-20, college quality at ages 18-20, various neighborhood characteristics of the zipcode in which the children live, fertility outcomes of female children, and tax payments. Part B of Table XXII presents the summary statistics of these outcomes for the two samples of children (age ≤ 13, and age 13-18 at random assignment).

Table XXIII presents the main set of results and comprehensive robustness analysis. The first

column illustrates the voucher take-up rate of 48% in the experimental group, and impacts on the neighborhood poverty rate. Panel A presents the impact on household income, individual income, employment, and college attendence for the sample of children less than age 13 at random assignment. Here, we find significant impacts of the experimental treatment, which W2- earnings impacts of $1,399 per year. We find a 2.5pp increase in college attendence, relative to a baseline of 16.5% for the control group. The mean college quality conditional on attending college for younger children in the control group is $31,409, while the quality for all those who do not attend college is $18,867. Moreover, we find large impacts on measures of neighborhood quality for the children at later ages – for example, we see a 1.6pp decrease in the poverty rate of the zipcodes in which these children live at age 25. Short term impacts on the quality of the children's neighborhood are also noticeable: the mean poverty rate in ZIP post random assignment to age 18 was 5.84% lower for the experimental group than for the control group.Thus the experimental ITT estimate is -1.6 percentage points, about one-third as large as the treatment effect on the average poverty rate in the ZIP code where the individual lived in childhood.

We find no significant impacts on the likelihood that the children have teen births; however, in the event the children do have a teen birth, we find that it is less likely that the father is absent on the birth certificate, consistent with positive impacts on family structure. Finally, we find significant impacts on the taxes paid by these children. We find a 5.7% increase in filing rates and a $183.90 increase in federal taxes collected per year.[38] This suggests that the total cost of providing these housing vouchers is significantly lower than the "sticker price" of the voucher because of its positive effects on children's tax revenue collected 15 years after the program.

Panel B replicates the analysis for the sample of children who were age 13-18 at random assignment. Consistent with the hypothesis that places have causal effects on children in proportion to exposure time, we find generally smaller effects on this sample. For example, we find an earnings impact of -$761, in contrast to the $1,399 in the age 13 group. The remaining panels repeat the analysis for pooled age groups (Panel C), those above age 18 at random assignment (Panel D). In Panel E, we replicate prior results of the MTO Final Impacts Evaluation. Similar to the MTO Final Impacts Evaluation we find no impact when analyzing available data from 2008. This is because one would have had only 552 observations on earnings for children who were less than 13 at RA in 2008. Finally we also explore heterogeneity in the effects by tax year as well as by gender (Panel

---

[38]The mean non-W-2 earnings for tax filers in the control group – which is a plausible upper bound for non-filers – is $1,721.

G) and race (reported in the HUD survey) and heterogeneity across the five site locations (Baseline with controls for the baseline pre-determined covariates, Panel E and H, Columns IM-JE). We also test for subgroup heterogeneity and confirm our results are not simply driven by multiple hypothesis testing (Panel F). In general, the patterns on these subgroups are consistent with the general patterns, albeit estimated with less precision. Panel G reports estimates separately by gender and illustrates the results are fairly similar for both male and female children, contrasting with earlier work suggesting greater impacts on females than males.

Columns EE-IK repeat Panels A-D using a robust set of control variables in the regressions; this illustrates that the results in Panels A-D are robust to the inclusion of additional control variables. Columns IL-JE reports the heterogeneity in the results by gender and race, which suggest the impacts are fairly similar across sites and genders, with some evidence of larger impacts in sites with higher initial baseline poverty rates. Panel G (Columns JG-JQ) decompose the baseline results into various income components, including w-2 vs. other sources of income such as unemployment benefits. Panel H (Columns JR-KC) illustrate the robustness of the choice of age ¡= 12 cutoff for splitting our groups. Columns KE-ONU replicate panels A-D using a "treatment on the treated" estimation approach as opposed to an "intention to treat" approach. These effects are virtually identical to the main specifications, except that they're scaled by ~180% to reflect the fact that not everyone who was given the opportunity to receive a housing voucher took up the voucher. Columns NW-OP present results separately that vary the age of the child's income at the time of measurement. Here, we observe that the impacts primarily arise on incomes measured at or above age 24, consistent with MTO improving life-cycle earnings patterns. Columns OR-PD presents estimates of the baseline impacts on subgroups of varying age categories, illustrating generally larger effects in younger age groups for those below age 13 at random assignment. Finally, Columns PF-PP illustrates the cumulative exposure impact of MTO for adults (Columns PF-PJ) and the associated impact on individual earnings (Columns PL-PP). Consistent with earlier MTO work, we find no evidence of an exposure impact on earnings for adults.

Overall, the results suggest targeted housing vouchers can improve upward mobility for children from poor families, and also suggest that programs like LIHTC could potentially be targeted more effectively if their goal is to improve upward mobility. While Baum-Snow and Marion (2009) show that LIHTC is effective in increasing low-income housing, in order to qualify, at least 50% of households in the tract must have incomes below 60% of AMGI. By encouraging the development of low-income housing in high-poverty areas, LIHTC housing may not benefit social mobility as

much as if they were build in less impoverished areas. However, further research is required for a full understanding of the impact of the LIHTC on upward mobility.

In terms of the mortgage interest deduction, we showed above it has a significant correlation with mobility. But, the tagging of the tax expenditure to the mortgage interest, as opposed to the overall cost of housing regardless of ownership, may not be essential for promoting intergenerational mobility. More generally, our results highlight the potential role of place-based tax credits for promoting intergenerational mobility.

III.H   Quantifying the Benefits of Tax Expenditures

Overall, these results suggest that tax expenditures aimed at low-income taxpayers can have significant impacts on intergenerational mobility. As a result, the short-term fiscal gains from reducing such expenditures must be weighed against the potentially large long-term costs of reduced income growth for low income individuals. Indeed, if tax expenditures today provide long-run benefits for children in the future, they potentially significantly alter the appropriate long-run cost analysis of tax expenditures.

As a final step in our analysis, we provide some guidance on how to quantify this trade-off. We compare the benefits of improved outcomes for children from low-income families through improved mobility to the cost of providing direct transfers to these children through modifications to the income tax schedule.[39] To do so, we construct a measure of the marginal cost of providing an additional dollar of tax expenditures to an individual at each point along the income distribution. We follow the standard optimal tax approach (Saez 2001, Bourguignon and Spadaro 2012) and write this marginal cost as a function of the elasticity of the income distribution, the elasticity of taxable income, and the marginal tax rate (the precise formulas are provided for each variable in Appendix Table IX, and the summary statistics for the samples used conditional on each marginal tax rate are provided in Table XIV).[40]

Assuming the impact of tax expenditures on mobility represents a causal relationship, our quantification analysis suggests significant benefits of tax expenditures to parents today, as quantified by future positive potential impacts of tax expenditures in the future. As shown in Table XIV, the marginal cost of providing an additional dollar of tax expenditures to below-median earners is generally around $1.10 (broader elasticity estimates yield a range of $1.05-$1.2) per $1 mechanical

---

[39]Here, "tax schedule" is inclusive of tax expenditures, such as EITC benefits.

[40]Because hypothetical modifications to the tax schedule would affect all filers, we estimate these statistics on the sample of all filers between age 25 and 60 to capture the relevant marginal cost of providing transfers through the tax schedule.

change in tax expenditures. In contrast, the cost of providing an additional dollar of tax expenditure to an above median earner is closer to \$0.70 for every \$1 in mechanical increase in tax expenditures.[41] This suggests the long-run fiscal cost of tax expenditures that generate an above median earning child is roughly $\frac{1.1-0.7}{1.1} = 36\%$ lower than an equivalent tax expenditures that generates a below-median earning child. Put differently, the results suggests incorporating the impact of tax expenditures on intergenerational mobility is crucial to accurately calculate the total cost of changes to tax expenditure policies.

IV   Discussion and Conclusion

In this paper, we combine local CZ data on tax expenditures and local tax return income data to investigate the relationship between tax expenditures and economic opportunity. Our results demonstrate consistent and fairly robust relationships between higher local tax expenditures and lower intergenerational elasticity (IGE), i.e. higher economic mobility. This pattern emerges both in considering overall tax expenditures and individual analyzes of mortgage interest deductions, state income taxes, and state EITCs. The progressivity of tax expenditures and state income taxes have the strongest correlations with intergenerational mobility. Overall, our results suggest that local variation in tax expenditures plays a significant role in explaining variation in intergenerational mobility across the US.

Our analysis also makes two contributions that may be useful for further research on tax expenditures and issues related to income mobility. First, we have constructed new geographic data on intergenerational mobility, which provides measures of local economic opportunity by CZ. Future researchers can use this mobility data to analyze its determinants and improve our understanding of the role of tax policy in affecting economic opportunity. To assess the causal effects of tax expenditures, future research could focus on isolating exogenous changes in tax policy, and especially local tax policy, and analyzing local outcomes using quasi-experimental research designs.

Second, the broader contribution of this paper lies in illustrating the potential of a spatial research design to gain insight into the impacts of tax expenditures. This design exploits local variation in tax policies and previously unavailable local level data on outcomes to identify policy impacts. Future research can extend this research design to study a broad range of important outcomes including innovation, housing markets, labor markets, and other indicators of well-being

---

[41]As shown in Table XIV, the precise reason for the lower cost is primarily driven by the difference in the curvature in the income tax schedule, combined with evidence that the elasticity of taxable earnings is roughly constant across the income distribution.

to provide a more comprehensive perspective on the benefits of tax expenditures.

Technical Appendices

A   Data Construction

To construct our core and extended sample, we begin with the universe of individuals in the Data Master-1 file, which includes year of birth and gender for all persons with a Social Security Number or Individual Taxpayer Identification Number. We restrict this sample to all individuals who are current US citizens as of March 2013.[42] We further restrict to individuals who are alive through the end of 2012. The resulting dataset contains 47.8 million children across all cohorts 1980-1991 (see Appendix Table II).

To each child, we assign a parent as the first person(s)[43] who claim the child as a dependent[44] on the 1040 tax form with one modification: we require matched parents to be a reasonable age. In the case of a potential match to married parents or single mothers, we require the mother to be age 15-40 at birth; in the case of a match to a single father, we require the father to be age 15-40 at birth. If no such eligible match occurs in 1996, we search subsequent years until a valid match is found. Given the match to parent(s), we hold this definition of parents fixed regardless of subsequent dependent claims or changes in marital status.[45]

To each parent and child, we assign yearly income measures using the IRS Databank, a balanced panel of individuals with data from multiple tax forms for 1996-2012. As discussed in the main text, we define family income as adjusted gross income plus tax-exempt interest income and the non-taxable portion of Social Security and Disability benefits using data directly from the 1040 filings. 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 each parent. This provides our yearly measure of parent family income. We compute child family income in exactly the same manner. [46]

Mean parent family income is equal to the sum of the mother's family income plus the father's family income in each year from 1996 to 2000, divided by 10 for those with two parents and

---

[42]The Data Master-1 file does not contain historical citizenship status and thus we can only restrict to a sample who are currently US citizens as of the time at which we access the data.

[43]If a parent is married but filing separately, we assign both parents.

[44]Children can be claimed as a dependent only if they are aged less than 19 at the end of the year (or aged less than 24 if students and any age if fully disabled). Legally, a dependent child must be the child, step child, adopted child, foster child, brother or sister, or a descendant of one of these (for example, a grandchild or nephew). Children are claimed by their custodian parent (i.e., the parent with whom they live for over half the year). Furthermore, the custodian parent must provide more than 50% of the support to the child. Hence, working children who support themselves for more than 50% cannot be claimed as dependents (see IRS Publication 501 for complete details).

[45]For example, a child matched to a married parents in 1996 who divorce in 1997 will always be considered matched to the two original parents. Conversely, a child matched to a single parent in 1996 that marries in 1997 will be considered matched to a single parent, though spouse income will be included through the family income measures.

[46]We augment income data from the IRS Databank with the SOI sample of returns from 1987-2011 to improve precision for high income outliers. The probabiliy of being in the SOI sample increases with income, and approaches 1 for the highest income individuals or those whose adjusted gross income exceeds $5 million. We therefore use these returns to determine the validity of high income individuals identified by the databank. If an individual's adjusted gross income exceeds $10,000,000 we look for the individual in the SOI sample; if present, we use the SOI measure of adjusted gross income and wage income as reported on a F1040 return. If not, we replace the adjusted gross income with the total wages reported on the filed F1040 contained in the databank. This patch affects 0.017% of individuals (or, equivalently, 1.7% of individuals in the top 1% of the income distribution) and avoids skewed results due to data aberrations in upper tail incomes. Note that the IRS Databank includes tax year 2012 where the SOI sample does not. We therefore topcode 2012 income at $100,000,000 for all individuals in 2012.

dividied by 5 for those with only one matched 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 is equivalent to individual income prior to marriage and total family income (including the new spouse's income) after marriage. Similarly, mean child family income is defined as the sum of the child's family income between 2011-2012, divided by 2.

By construction, these household measures of income naturally are affected by marriage. Moreover, we also cannot capture co-habitation in the tax data. These reasons provide additional motivation for our robustness assessment using individual measures of income.[47]

B   Statistics of Income Sample

To analyze earlier birth cohorts, we use a separate sample of repeated cross sections for tax years 1987-2011constructed by the Statistics of Income division at the IRS.[48] The SOI files are stratified annual files with a weight equal to the inverse of the sampling probability. IRS (2013) provides a detailed description of the sampling procedure. Because sampling of a given tax filer is not independent across years, the simplest way to obtain a representative sample of children by cohort of birth is to observe each cohort in only one year. Hence, we first restrict our analysis to children matched to parents at age 16. Formally, for children born in year $t = 1971, .., 1991$, we consider solely the SOI file for year $t + 16$ (i.e., 1987 for the 1971 cohort) and select all children born in year t claimed in a tax return in the year $t + 16$ SOI file.

The parent(s) are defined as the people claiming the child as a dependent on the tax return where the child appears.[49] This generates a representative sample of all children from cohorts 1971 to 1991 who are claimed as dependents when they are aged 16. The sample includes about 4,000 children for cohorts 1971 to 1981, 6000-8000 children for cohorts 1982 to 1988, and 12,000 for cohorts 1989-1991.[50] Comparing the sample we obtain to official annual birth statistics, our weight SOI sample of children is equal to about 70% of birth cohorts 1971-2, about 80% of the birth cohorts 1973 to 1979, and about 90-95% of the cohorts 1980 -1991.[51] Given this sample of matched children at age 16, we then repeat this matching procedure using match ages of 12-15 (instead of age 16). Because of the SOI sampling design, each of these age-match samples is representative. To increase precision, we pool those samples together which allows us to obtain further matches in cases where parents did not file when their children were 16 years old.[52] Given these parent-child

---

[47]As discussed in the main text, parent family income coincides with individual income for single parents. For married parents, we define each parent's individual earnings as the sum of wage earnings from form W-2, unemployment benefits from form 1099-G, and Social Security and Disability benefits from form SSA-1099 for that individual. Individual earnings excludes capital and other non-labor income. To incorporate these sources of income, we turn to the 1040 tax return and add half of family non-labor income – defined as total family income minus total family earnings as reported on a filed form 1040– to each parent's individual earnings. We divide non-labor earnings equally between spouses because we cannot identify which spouse earns non-labor income from the 1040 tax return.

[48]We begin with the 1987 sample because the dependent reporting requirements in the Tax Reform Act of 1986 allow the IRS to link children to parents.

[49]In the case of married filing separately, we can track both parents because a married filing separately return requires identifying the spouse. The child receives the same weight as the tax return in the SOI file since the probability of the child to appear in the sample is exactly the same as the probability of the tax return to appear in the sample (dependents can be claimed on one tax return only).

[50]This increase is due to the expansion of the sampling size of the SOI files over time.

[51]This increase is perhaps due to the fact that claiming children became more advantageous over time due to the expansion of refundable credits (such as the EITC or the child tax credit).

[52]Note that these samples partly overlap and hence the same child-parent match can appears several times in the

matches, we attach the same parent and child income measures used in the core sample, drawn from the IRS Databank. Parent income continues to be the average of income over 1996-2000; child income is the average of income over 2011-2012.

C   Comparison to Survey Datasets

To verify that our core sample is nationally representative, we computed select moments of the distributions of child family income and earnings and compared them to corresponding measures in two representative surveys often used for the study of income distribution: the 2011-2012 CPS and the 2011 ACS. Appendix Table X presents the results of this exercise. To assess the possible role of selection biases in our matching procedure, we compute tax based statistics both before and after merging children to their parents. Because our match rate is high, the differences between these two samples is small, though it is clear that children who lack valid parent matches have lower earnings.

The sum of the sampling weights in our survey based samples provide estimates of the size of the target population being sampled. This population is slightly larger than the number of observations in our tax samples, which is to be expected given that some households do not hold jobs or earn any income. Accordingly, the fraction with zero earnings is somewhat lower in the tax samples than in surveys.

Mean earnings are lower in the tax data than survey measures but median earnings are very close. This follows because, in contrast to the tax data, the survey based income measures are topcoded. Perhaps surprisingly, the interquartile range (P75-P25) of earnings is also similar across the three data sources. If survey data were reported with classical measurement error, we would expect the interquartile range to be larger in survey sources. However, it is well known that survey reports of income exhibit "mean reverting" measurement error which has the effect of reducing variability (Bound and Krueger, 1991; Bound, Brown, and Mattheowitz, 1999). Moreover, as demonstrated by Kline and Santos (2013), survey non-response tends to follow a U-shaped pattern with very high and low earning individuals being least likely to provide earnings responses, which can also reduce variability. The quantiles of family income also line up well across the three data sources, with the tax based moments strongly resembling those from the ACS, in particular. This may be a result of the higher response rate for earnings questions in the ACS.

In sum, the close correspondance between the distribution of income in our core tax sample and that found in representative surveys suggests to us that our data closely approximate the nationally representative target population of interest. Although our income measure differs slightly from measures that have been used in the prior literature, we show in the next section that our national results on intergenerational mobility line up closely with those found using survey datasets.

D   Assignment to CZs

Children are assigned ZIP codes of residence based on their parents' ZIP code on the form 1040 used to match the parent to the child.[53] For most children in the core sample, the match year is 1996, so the ZIP code corresponds to where their parents lived in 1996. Our primary source of mappings from ZIP codes to Commuting Zones is the 1999 Census crosswalk between ZIP codes and counties, further aggregated to Commuting Zones using David Dorn's county-to-CZ crosswalk.

---

pooled sample. This does not introduce bias, but we correct standard errors by always clustering at the child level.

[53]In a case where a parent filed a F1040 claiming the child but did not report a valid ZIP code, we search information returns (such as W2 and 1099G) for a valid ZIP code in that year.

The counties in the U.S. Census Bureau crosswalk and in David Dorn's crosswalk are not identical because they correspond to county definitions at different points in time; in particular the U.S. Census Bureau crosswalk includes changes between 1990 and 1999. We account for changes that affected 200 or more people.[54] Using the 1999 Census file, we identify the CZ of 38,839 ZIP codes. To better track individuals residing in ZIP codes that have been created since 1999, we add an additional 477 ZIP codes not valid in 1999 but appearing in the more up-to-date 2011 HUD-USPS crosswalk. [55]

Of 9,864,965 children with non-missing ZIP codes in our core sample, 9,778,071 could be assigned a childhood CZ using just ZIP codes valid in 1999; an additional 2,718 are assigned a CZ based on a ZIP code valid in 2011 but not in 1999. Some of our alternative specifications require tracking children's locations into adulthood using their most recently reported ZIP codes, eg, for the purpose of deflating income by local cost of living or obtaining mobility estimates for children that continue to reside in their childhood CZ's. Of 9,834,975 non-missing child ZIP codes in adulthood, 9,537,283 could be matched to a CZ from a ZIP code in use in 1999 and 198,317 were matched using a ZIP code created after 1999.

E   Robustness of Rank-Rank Relationship

In this appendix, we evaluate the robustness of our estimates of the degree of intergenerational persistence in permanent income to alternative specifications. We begin by evaluating two potential sources of bias emphasized in prior work: lifecycle bias and attenuation bias.

Lifecycle Bias. Prior research has shown that measuring children's income at early ages can understate intergenerational persistence in lifetime income because children with high lifetime incomes have steeper earnings profiles when they are young (Haider and Solon, 2006, Grawe, 2006, Solon 1999). Intuitively, children from high-income families are more likely to be in college or graduate school in their 20's, and their earnings at early ages therefore understate their lifetime income. To evaluate whether our baseline estimates suffer from such lifecycle bias, Figure Va plots estimates of the rank-rank slope by the age at which the child's income is measured. To construct this figure, we measure children's income as mean family income in 2011-2012 and parent income as mean family income between 1996-2000, as in our baseline analysis. We then replicate the OLS regression of child income rank on parent income rank for each birth cohort between 1980-1990. For children in the 1980 birth cohort, we measure earnings at age 32 in 2012; for those in the 1990 cohort, we measure earnings at age 22.[56]

As expected, the rank-rank slope rises very steeply in the early 20's as children enter the labor force. Children's income ranks at age 22 are essentially uncorrelated with their parent's income rank. The rank-rank slope stabilizes around age 30: it increases by 2.1% from age 30 to 31 and just 0.2% from age 31 to 32. This pattern explains why the estimates of intergenerational persistence in the extended sample of 1980-85 birth cohorts reported in Column 6 of Table II are lower that the estimates for the core sample of older birth cohorts.

Because we can only reliably link children born during or after 1980 to parents, we cannot observe children's income beyond age 32 in the population data. However, as described in Section

---

[54]This information is available at http://www.census.gov/geo/reference/county-changes.html.

[55]A salient example of recently created ZIP codes is the 2007 three-way split of Manhattan's ZIP code 10021, resulting in the creation of codes 10065 and 10075. This change, due to population growth, placed 50,000 Manhattan residents into new ZIP codes.

[56]We vary birth cohort and hold the year of income measurement fixed to eliminate calendar year effects. We obtain very similar results if we instead track a single cohort and vary age by measuring earnings in different calendar years.

II.A, a 0.1 percent stratified random sample of tax returns constructed by the IRS Statistics of Income division allows us to link children to parents going back to the 1971 birth cohort. The series in triangles in Figure Va replicates the analysis above within the SOI sample (using sampling weights to recover estimates representative of the population).[57] The estimates in the SOI sample are very similar to those in the full population prior to age 32. After age 32, the estimates remain roughly constant. These findings confirm that rank-rank correlations exhibit little or no lifecycle bias provided that income is measured after age 30, as in our baseline measure.[58]

An analogous lifecycle bias can arise if parent income is measured at very old or young ages. This is typically less of a concern than the age at which children's income is measured because most panel datasets linking parents to children, including ours, contain information on parents' income in the middle of their lifecycle. Nevertheless, to ensure that our estimates are not sensitive to the age at which parent income is measured, in Appendix Figure IIIb we plot the rank-rank slope using the core sample, varying the 5-year window used to measure parent income from a starting year of 1996 (when parents are 41 years old on average) to 2007 (when parents are 55 years old). As expected, the rank-rank estimates exhibit virtually no variation with the age of parent income measurement within this range.

Attenuation Bias. Income in a single year is a noisy measure of permanent income because of transitory shocks and measurement error. Solon, 1992, Zimmerman, 1992, and Mazumder (2005) show that using multi-year means of parent income – which provide more accurate measures of lifetime income – generate significantly higher estimates of intergenerational persistence in survey data. To evaluate whether our baseline estimates suffer from such attenuation bias, Figure Vb plots estimates of the rank-rank slope, varying the number of years used to calculate mean parent family income. To construct this figure, we measure children's income as mean family income in 2011-2012 and use the core sample of 1980-82 birth cohorts. We then replicate the OLS regression of child income rank on parent income rank, varying the number of years used to calculate mean parent income from 1 year (1996 income only) to 17 years (1996-2012).

Consistent with the findings of Solon, 1992, we find that the rank-rank slope rises when we increase the number of years used to measure parent income from one to five. However, the rank-rank slope based on five years (0.341) is only 3.2% larger than the slope based on one year of parent income (0.330). This 3.2% attenuation bias is considerably smaller than the 33% change in the IGE (from 0.3 to 0.4) reported by Solon, 1992 when using a five-year average instead of one year of data. We find less attenuation bias for three reasons: (1) income is measured with less error in the tax data than in the PSID, (2) we use family income measures rather than individual income, which fluctuates more across years, and (3) we use a rank-rank specification rather than a log-log specification, which is more sensitive to income fluctuations at the bottom of the distribution.[59]

The rank-rank slope is virtually unchanged by adding more years of data beyond five years: the estimated slope using 15 years of data to measure parent income (0.350) is only 2.8% larger than the baseline slope of 0.341 using 5 years of data.[60] This finding differs sharply from that of Mazumder

---

[57]As above, we measure child income in 2011-12 and vary the age at which income is measured by varying birth cohorts. We measure parent income using the same 5-year average over 1996-2000.

[58]Estimates of the IGE using the traditional log-log specification also stabilize around age 30; see Appendix Table XIII for estimates in the baseline sample and Appendix Table XIV for estimates on the SOI sample.

[59]In addition, Appendix Figure IIIb repeats this exercise after censoring kid income on p10-p90, and illustrates the same pattern holds. Similarly, Appendix Figure IIId repeats this exercise averaging child income; and Appendix Figure IIIC repeats this exercise using the college attendence gradient. In short, all of our results are highly robust to using varying number of years of income. Although this robustness contrasts with literature using survey datasets, it likely reflects the high quality of the income data in the tax data relative to survey data.

[60]We also report the sensitivity of log-log specifications to the number of years of data used to measure parents' income in Appendix Table XV

(2005), who argues that even 5-year averages of parent earnings exhibit substantial attenuation bias because of long-lasting transitory shocks to income. Mazumder obtains IGE estimates as high as 0.6 when using 15-year averages of parent income matched SIPP-SSA administrative data, 54% larger than his 4-year pooled estimate of 0.388.[61] We believe our results differ from Mazumder's findings because we directly observe income for all individuals in our data, whereas Mazumder imputes parent income based on race and education for up to 60% of the observations in his sample to account for top-coding in Social Security records.[62] These imputations are analogous to instrumenting for parent income using race and education, an approach known to yield higher estimates of the IGE, perhaps because parents' education directly affects children's earnings (e.g., Solon, 1992, Zimmerman, 1992, Mulligan, 1997, Solon, 1999). Because the SSA earnings limit is lower in the early years of his sample, Mazumder imputes income for a larger fraction of observations when he averages parent income over more years (Mazumder 2005, Figure 3). As a result, Mazumder's estimates effectively converge toward IV estimates as he uses more years to calculate mean parent income, explaining why his estimates rise so sharply with the number of years used to measure parent income.[63] Consistent with this explanation, when he drops imputed observations, his IGE estimates increase much less with the number of years used to measure parent income (Mazumder 2005, Table 6).[64] Irrespective of the source of Mazumder's findings, we can be confident that in tax records without imputed earnings, our baseline approach of averaging parents' income over five years yields a reliable estimate of intergenerational mobility.

Prior studies have focused on measurement error in parent income rather than children's income because only the former generates attenuation bias in the standard OLS log-log regression specification, insofar as the transitory measurement error in child income is unrelated to parent income. This is not true in a rank-rank specification because the measurement error in children's income generates misclassification error in their ranks, which can attenuate the rank-rank slope.[65] To evaluate the magnitude of this bias, we analyze the impact of varying the number of years used to measure the child's income in Appendix Figure IIId. The rank-rank slope increases only modestly when increasing the number of years used to compute child family income, with no detectable change once one averages over at least two years, as in our baseline measure.

Alternative Income Definitions. In rows 5-7 of Table II, we explore the robustness of the baseline rank-rank estimate to alternative definitions of child and parent income. One concern with our baseline family income measure is that children from two-parent households may have better outcomes than those raised by single parents simply because they have more parental input,

---

[61]See Table 4, row 1 on page 246 of Mazumder (2005).

[62]We continue to find little difference between IGE's based on five-year vs. fifteen-year averages of parent income when using a log-log IGE specification similar to that estimated by Mazumder. We obtain a log-log IGE of 0.366 using a 15-year average of parent family income, closer to Solon (1992) original estimates than Mazumder's more recent estimates.

[63]Mazumder also reports simulations of earnings processes showing that attenuation bias in the IGE should be substantial even when using five-year averages. However, he calibrates the parameters of the earnings process in his simulation based on estimates from survey data, which have much more noise than administrative records. If one replicates Mazumder's simulations using a smaller variance share for transitory shocks, one obtains results similar to those in Figure Vb, with little attenuation bias in estimates based on five-year averages.

[64]To be clear, Mazumder himself acknowledges the potential bias due to imputation, as he recommends in his conclusion that "future research should attempt to verify the results here using long-term measures of permanent earnings from other sources that do not require the kind of imputations that were necessary in this study." We simply follow his recommendation here.

[65]Formally, measurement error in the child's rank is correlated with the parent's rank because the support of the child's rank is bounded. Intuitively, children from the highest-ranked families are more likely to be under-ranked than over-ranked due to measurement error in their income. As a result, measurement error in children's or parent's income generally leads to attenuation bias in a rank-rank specification.

independent of income. Since two-parent households will have higher family income on average, this could lead us to overstate the intergenerational persistence of income holding fixed parents' marital status. To evaluate the magnitude of this effect, in row 5 we define the parent's rank based on the individual income of the parent with higher income, eliminating the mechanical variation in family income driven by the number of parents in the household.[66] The rank-rank correlation falls by approximately 10%, from 0.341 to 0.312. The same logic explains why we find a similar 10% reduction when we condition on marital status in the baseline specification using family income (columns 4 and 5 of row 4). The impact of using individual parent income instead of family income is modest because (1) most of the variation in parent income across households is not due to differences in marital status and (2) the mean ranks of children with married parents are only 4.6 percentile points higher than those with single parents.

In row 6 of Table II, we repeat this exercise for children. Here, the concern is that children of higher income parents may be more likely to marry, again raising observed persistence in family income relative to individual income. Using individual income to measure the child's rank has differential impacts by the child's gender, consistent with findings of Chadwick and Solon (2002). For male children, using individual income instead of family income reduces the rank-rank correlation from 0.336 in the baseline specification to 0.317, a 6% reduction. For female children, using individual income reduces the rank-rank correlation from 0.346 to 0.257, a 26% reduction. Note, we also confirm that all of our findings below hold if we use individual income for male children, the specification used in much of the earlier literature (Solon 1999).

Finally, in row 7 of Table II, we define the child's income based on labor earnings (excluding capital income and other non-labor income). There is virtually no difference between the rank-rank estimates in row 7 and row 6, indicating that the vast majority of the persistence in income across generates comes from labor earnings.

College attendence, college quality, and teen birth. We supplement our analysis of intergenerational income mobility by studying the relationship between parent income and two intermediate outcomes for children: college attendance and teenage birth.

The series in circles in Figure VIIIa presents a binned scatter plot of the college attendance rate of children vs. the percentile rank of parent family income using the core sample. College attendance is defined as attending college in one or more years between the ages 18 and 21. The relationship between college attendance rates and parental income rank is again virtually linear, with a slope of 0.675. That is, moving from the lowest-income to highest-income parents increases the college attendance rate by 67.5 percentage points, similar to the estimates reported by Bailey and Dynarski (2011) using survey data.

The series in triangles in Figure VIIIa plots college quality ranks vs. parent ranks. We define a child's college quality rank based on the mean earnings at age 30 of students who attended each college at age 20. The 54% of children who do not attend college at age 20 are included in this analysis and are assigned the mean rank for the non-college group, which is approximately 54/2 = 27. The relationship between college quality rank and parent income rank is convex because most children from low-income families do not attend college and hence increases in parent income have little impact on college quality rank at the bottom. To account for this non-linearity, we regress college quality ranks on a quadratic function of parent income rank and define the gradient in college quality as the difference in the predicted college quality rank for children with parents at the 75th percentile and children with parents at the 25th percentile. The P25-75 gap in college quality ranks is 19.1 percentiles in our core sample.

Figure VIIIb plots teenage birth rates for female children vs. parent income ranks. Teenage

---

[66] We define the top earner in a family based on mean income from 1996 to 2000.

birth is defined (for females only) as having a child when the mother is aged 13-19. There is a 29.8 percentage point gap in teenage birth rates between children from the highest- and lowest-income families.

These correlations between intermediate outcomes and parent income ranks do not vary significantly across subsamples or birth cohorts, as shown in rows 9-11 of Table I. The strength of these correlations indicates that much of the divergence between children from low vs. high income families emerges well before they enter the labor market, consistent with the findings of prior work (e.g., Neal and Johnson 1996, Cameron and Heckman 2001, Bhattacharya and Mazumder 2011).

Robustness to additional cohorts

One additional concern with our mobility statistics is that they are unique to the 1980-82 cohort and do not represent a stable measure of mobility. To explore this, Appendix Table XII compares the quintile transition matrix of parents and children using the 1980-1982 cohorts and the 1980-1985 cohorts. Here, we can see remarkable stability. For example, roughly 7.5% of the 1980-82 cohort starts out in the bottom quintile of parental income and ends up in the top quintile; the analogous figure is 7.8% of the 1980-1985 cohorts. Overall, this table clearly shows that our focus on the 1980-82 cohort, as opposed to more extended cohort definitions, does not significantly bias our estimates.

We conclude that our baseline measure of intergenerational persistence accurately captures the degree of persistence in lifetime income across generations and focus primarily on this measure in the remainder of the paper.

F   Robustness of Spatial Patterns

In Table XV, we assess the robustness of the spatial patterns in mobility documented in Section V along four dimensions: (1) changes in sample definitions, (2) changes in income measures, (3) adjustments for differences in the cost-of-living and growth rates across areas, and (4) the use of alternative statistics to measure relative and upward mobility. The first number in each cell of this table reports the correlation across CZs of a baseline mobility measure (using child family income rank and parent family income rank in the core sample) with an alternative mobility measure described in each row. The second number in each cell reports the ratio of the standard deviation of the alternative measure to the baseline measure. Note that we do not report the ratio of standard deviations for statistics that are measured in different units relative to the corresponding baseline measure.

Column 1 reports the unweighted correlation (and SD ratio) across CZs between our baseline measure of absolute upward mobility ($\bar{r}_{25,c}$) and the corresponding alternative measure of $\bar{r}_{25,c}$. Column 2 replicates Column 1 for relative mobility (β_c). Columns 3 and 4 replicate Columns 1 and 2 weighting CZs by their population in the 2000 Census.

Sample Definitions. In the first section of Table XV, we assess the robustness of the spatial patterns to changes in the sample definition, as we did at the national level in Table II. Rows 1 and 2 restrict the sample to male and female children, respectively. Rows 3 and 4 consider the subsamples of married parents and single parents. The correlations of both absolute and relative mobility in these subsamples with the corresponding baseline measures is typically above 0.9.

In row 5, we replicate the baseline specifications using the 1983-85 birth cohorts (whose incomes are measured at age 27 on average in 2011-12). In row 6, we consider the 1986-88 birth cohorts instead. The intergenerational mobility estimates across CZs for these later birth cohorts are very highly correlated with the baseline estimates. This result has three implications. First, it demonstrates that the reliability of CZ-level estimates is quite high across cohorts; in particular,

sampling error or cohort-specific shocks do not lead to much fluctuation in the CZ-level estimates. Second, because the later cohorts are linked to parents at earlier ages (as early as age 8), we conclude that the spatial patterns in intergenerational mobility are not sensitive to the precise age at which we link children to parents or measure their geographical location. Finally, because the earnings of later cohorts are measured at earlier ages, we conclude that one can detect the spatial differences in mobility even when measuring earnings quite early in children's careers.

In row 7, we restrict the sample based on the age of parents at the birth of the child. We limit the sample to children whose mothers are between the ages of 24-28 and fathers are between 26-30 (a five year window around the median age of birth). The intergenerational mobility estimates in this subsample are very highly correlated with the baseline estimates, indicating that the cross-area differences in income mobility are not biased by differences in the age of child birth for low income individuals.

In row 8, we assess the extent to which the variation in intergenerational mobility comes from children who succeed and move out of the CZ as adults vs. children who stay within the CZ. To do so, we restrict the sample to the 62% of children who live in the same CZ in 2012 as where they grew up. Despite the fact that this sample is endogenously selected on an ex-post outcome, the mobility estimates remain very highly correlated with those in the full sample. Apparently, areas such as Salt Lake City that generate high levels of upward income mobility do so not just by sending successful children to other CZs as adults but also by helping children move up in the income distribution within the area.

In row 9, we restrict the sample to the 88% of children in the core sample who are not claimed as dependents by other individuals in subsequent years after they are linked to the parents we identify. We obtain very similar estimates for this "unique parent" subsample, indicating that the spatial pattern of our mobility estimates is not distorted by measurement error in linking children to their parents.

Income Definitions. In the second section of Table XV, we evaluate the sensitivity of the spatial patterns to alternative definitions of income. The definitions we consider match those in the robustness analysis in Table II. In row 10 of Table XV, we define parent income as the income of the higher earner rather than total family income to evaluate potential biases from differences in parent marital status across areas. In row 11, we measure the child's income using individual income instead of family income to assess the effects of differences in the child's marital status. In row 12, we use the child's individual earnings (excluding capital and other non-labor income). In row 13, we replicate the specification in row 11 for male children, using individual income for the child and family income for the parent. Row 14 replicates row 13, but defines parent income as the income of the higher earner instead. In row 15, we define parent income using data from 1999-2003 (when we have data from W-2's) instead of 1996-2000. All of these definitions produce very similar spatial patterns in intergenerational mobility: correlations with the baseline measures exceed 0.9 in most cases.

Adjustments for Cost-of-Living and Growth Rates. The third section of Table XV considers a set of other factors that could bias comparisons of intergenerational mobility across areas. One natural concern is that our estimates of upward mobility may be affected by differences in prices across areas. To evaluate the importance of differences in cost of living, we construct a CZ-level price index using the American Chamber of Commerce Research Association (ACCRA) price index for urban areas combined with information on housing values, population density, and CZ location (see Appendix A for details). We then divide parents' income by the price index for the CZ where their child grew up and the child's income by the price index for the CZ where he lives as an adult (in 2012) to obtain real income measures.

Row 16 of Table XV shows that the measures of intergenerational mobility based on real incomes

are very highly correlated with our baseline measures (see also Appendix Figure V). The reason that cost-of-living adjustments have little effect is that prices affect *both* the parent and the child. Intuitively, in high-priced areas such as New York City, adjusting for prices reduces the child's absolute rank in the national real income distribution. But adjusting for prices also lowers the real income rank of parents living in New York City. As a result, the degree of upward mobility – i.e., the difference between the child's rank and the parent's rank – is essentially unaffected by adjusting for local prices.

The preceding logic assumes that children always live in the same cities as their parents. In practice, some children move to areas with higher prices (e.g. from rural Iowa to New York City). Our measures of upward mobility are affected by the cost of living adjustment in such cases, but they are not sufficiently frequent to have a large impact on our estimates. The correlation between the cost of living in the child's CZ at age 30 and the parent's CZ is 0.77, and the correlation between a child's nominal percentile rank and the local price index is only 0.10. As a result, cost of living adjustments end up having a minor impact on the difference between child and parent income and thus have little effect on our mobility statistics.

Next, we assess the extent to which economic growth is responsible for the spatial variation in upward mobility. In row 17, we define parent income as mean family income in 2011-12, the same years in which we measure child income. Insofar as local economic growth raises the incomes of both parents and children, this measure nets out the effects of growth on mobility. Both the upward and relative mobility measures remain very highly correlated with the baseline measures, suggesting that differences in local economic growth drive relatively little of the spatial variation in mobility.

As an alternative approach to accounting for growth shocks, we regress our measures of mobility on the CZ-level growth rate from 2000-2010 and calculate residuals.[67] Row 18 of Table XV shows that the correlation of the growth-adjusted relative mobility measures with the baseline measures exceeds 0.9; the correlations for absolute mobility exceed 0.8. Note that these growth-adjusted measures over-control for exogenous growth shocks insofar as growth is partly a consequence of factors that generate upward income mobility in an area. Hence, the finding that even controlling for growth rates directly does not significantly change the spatial pattern of intergenerational mobility supports the view that most of the variation in mobility across areas is not due to exogenous growth shocks in the 2000's.[68]

Alternative Statistics for Mobility. One potential concern with our approach is that using national ranks may misrepresent the degree of relative mobility within the local income distribution, which may better reflect a child's opportunities. To address this concern, in row 19 of Table XV, we measure relative mobility using local ranks. We rank parents relative to other parents living in the same CZ and children relative to other children who grew up in the same CZ (no matter where they live as adults). We define relative mobility as the slope of the local rank-rank relationship.[69] Relative mobility based on local ranks is very highly correlated with relative mobility based on national ranks. This is because local ranks are approximately a linear transformation of national

---

[67]We measure income in 2000 using the Census and in 2008 using the 5-year American Community Survey, averaged over 2006-2010. We calculate household income per working age adult as aggregate income in a CZ divided by the number of individuals aged 16-64 in that CZ. Annualized income growth is calculated as the annual growth rate implied by the change in income over the 8 year period; we use 8 years because 2008 is the midpoint of 2006-2010.

[68]The fact that college and teenage birth gradients are similar to income mobility gradients provides further evidence that growth shocks in the 2000s do not generate the differences in mobility across areas, as college and teenage birth are measured around 2000. These results also show that the spatial patterns are unlikely to be driven by differences in reporting of taxable income.

[69]We cannot study absolute mobility with local ranks because both child and parent ranks have a mean of 50 by definition: if one child moves up in the local distribution, another must move down.

ranks.

Finally, we consider two alternative measures of absolute upward mobility. In row 20, we measure absolute upward mobility based on the probability that the child rises from the bottom quintile of parent income to the top quintile of child income, as in Column 5 of Table III. In row 21, we measure absolute upward mobility as the probability that a child has family income above the poverty line conditional on having parents at the 25th percentile. To construct this statistic, we first regress an indicator for having family income above the federal poverty line in 2012 on parent rank in the national income distribution in each CZ.[70] We then calculate the predicted fraction of children above the poverty line for parents at the 25th percentile based on the slope and intercept in each CZ. The spatial patterns in both of these measures – which are also shown in the maps in Appendix Figure VI and VII – are very similar to those in our mean-rank based measure of upward mobility, with correlations across CZs above 0.9 in both cases.

Finally, a more technical confounding factor in accurately measuring mobility is differences across regions in mortality rates. If the relationship between mortality and income differs across space it could induce a selection problem, as we only observe incomes on parents and children who are alive. As a robustness check we compute the 3 year mortality rate for adults age 25-60 by income quintile within each CZ [Appendix Table IV]. We find that health inequality across areas has little explanatory power towards our measures of mobility.

Overall, we conclude that the geographical patterns in Figure VIIa are highly robust to alternative methods of measuring intergenerational mobility.

G   Comparison to Mitnik et al. (2014)

Mitnik et al. (2014) propose a new measure of the intergenerational elasticity that is more robust to the treatment of small incomes. In this appendix, we compare the traditional definition of the IGE (Solon 1999, Black and Devereux 2011) to the new measure proposed by Mitnik et al. We first show that the traditional IGE can be interpreted as the average elasticity of child income with respect to parent income in a model with heterogeneous elasticities, while Mitnik et al.'s new measure is a dollar-weighted (i.e., child-income-weighted) average of the same elasticity. We then compare estimates of the dollar-weighted IGE to estimates of the traditional IGE in our data and to the estimates of Mitnik et al.

Setup. Let Y_i denote the level of child income and X_i denote the level of parent income. Let $F_{Y|X=x}(y)$ denote the conditional distribution of Y given X, which we assume is differentiable with respect to x at all $(y, x) > 0$. Define the conditional quantile function (CQF) as the inverse of the CDF:

$$q(x, \tau) = F_{Y|X=x}^{-1}(\tau)$$

for $\tau \in [0, 1]$.[71] The CQF gives the quantiles of the conditional distribution of Y_i given X_i; for example, $q(x, .5)$ is the median of Y_i when $X_i = x$.

We can use the CQF to represent Y_i as:

$$Y_i = q(X_i, U_i),$$

---

[70]We define household size as the maximum household size in 2010-11, where household size is defined as 1 plus an indicator for being married plus the number of dependents claimed. The poverty line threshold is defined as $11,170+$(household size - 1)$\times\$3,960$.

[71]At mass points, we define $q(x, \tau) \equiv \min\{y : F_{Y|X=x}(y) \geq \tau\}$.

where $U_i|X_i \sim \text{Uniform}(0,1)$. Hence, the conditional mean of child income given parent income can be written as a function of the CQF:

$$E[Y_i|X_i = x] = E[q(X_i, U_i)|X_i = x] = \int_0^1 q(x, \tau) d\tau.$$

Define the elasticity of a given quantile of the child's income distribution with respect to parent income around a parent income level x as

$$\sigma(x, \tau) \equiv \frac{dq(x, \tau)}{dx} \frac{x}{q(x, \tau)} = \frac{q_x(x, \tau)x}{q(x, \tau)}.$$

In general, the elasticity will vary across quantiles τ.[72] We now show that traditional estimates of the intergenerational elasticity (e.g., Solon 1992) and the new estimator proposed by Mitnik et al. (2014) can be interpreted as different averages of the elasticities $\sigma(x, \tau)$.

Traditional IGE. The intergenerational elasticity at a given parent income level x, which we denote by $IGE(x)$, is defined as the impact of an increase in log parent income (starting from x) on expected log child income:

$$
\begin{aligned}
IGE(x) = \frac{dE[\log Y_i|X_i = x]}{d\log x} &= \frac{d}{d\log x} \int_0^1 \log q(x, \tau) d\tau \\
&= \int_0^1 \frac{d}{d\log x} \log q(x, \tau) d\tau \\
&= \int_0^1 \sigma(x, \tau) d\tau \\
&= \bar{\sigma}(x)
\end{aligned}
$$

where $\bar{\sigma}(x) \equiv E[\sigma(X_i, U_i)|X_i = x]$. If we interpret the IGE as the average of $IGE(x)$ across levels of parent income x, we obtain

$$IGE = E[\bar{\sigma}(x)] = \int_{-\infty}^{\infty} \int_0^1 \sigma(x, \tau) d\tau dF_X(x),$$

where $F_X(.)$ is the marginal distribution of X_i. Hence, the traditional IGE can be interpreted as the average elasticity of child income with respect to parent income across quantiles and parent income levels.

Mitnik et al. IGE. Mitnik et al. (2014) propose an alternative approach to estimating the IGE that switches the order of the log and the expectation relative to the traditional approach. They define the IGE as the impact of an increase in log parent income (starting from x) on the log of expected child income:

$$IGE_W(x) \equiv \frac{d\log E[Y_i|X_i = x]}{d\log x}.$$

---

[72]Naturally, this elasticity is only defined for quantiles where $q(x, \tau) > 0$; the standard empirical practice in the prior literature (e.g., Solon 1992) has been to exclude children with zero income for this reason.

To see how their estimand relates to the traditional IGE, observe that

$$
\begin{aligned}
IGE_W(x) &= \frac{d}{d\log x}\log\int_0^1 q\left(x,\tau\right)d\tau \\
&= \frac{\int_0^1 q_x\left(x,\tau\right)xd\tau}{\int_0^1 q\left(x,\tau\right)d\tau} \\
&= \frac{\int_0^1 q\left(x,\tau\right)\sigma\left(x,\tau\right)d\tau}{\int_0^1 q\left(x,\tau\right)d\tau} \\
&= E\left[\omega\left(X_i,U_i\right)\sigma\left(X_i,U_i\right)|X_i=x\right]
\end{aligned}
$$

where $\omega\left(X_i,U_i\right) \equiv \frac{q(X_i,U_i)}{E[q(X_i,U_i)|X_i]}$ is a set of quantile specific weights which sum to one for each value of X_i. Averaging $IGE_W(x)$ across levels of parent income x, Mitnik et al.'s statistic can be written as $IGE_W = E\left[IGE_W\left(x\right)\right]$. The parameter $IGE_W(x)$ is a weighted average of the elasticity $\sigma(x,\tau)$ across quantiles τ, with weights that are an increasing function of the child's income. Higher quantiles get larger weights, in proportion to their dollar value; the weights approach zero as the child's income approaches 0. In this sense, the IGE_W statistic defined by Mitnik et al. is a dollar-weighted mean of the IGE across quantiles.

The traditional IGE and the dollar-weighted IGE_W proposed by Mitnik et al. are two different parameters. The "correct" parameter depends on the question one seeks to answer. If one's goal is to estimate IGE_W, then the traditional IGE estimate will in general yield a biased estimate of IGE_W. Conversely, if one's target is to estimate the traditional IGE (e.g., for comparison to prior estimates in the literature), then IGE_W will in general be biased.

As Mitnik et al. note, one statistical benefit of the dollar-weighted IGE is that it is not sensitive small changes in incomes at the bottom of the distribution, such as recoding zero income as \$1. Intuitively, dollar-weighted elasticities are not sensitive to the impacts of parent income on children's income at the bottom of the distribution. In contrast, person-weighted estimates are very sensitive to changes in incomes at the bottom of the distribution, because such changes can affect elasticities at the lowest quantiles substantially. The traditional IGE is unstable because the elasticity of child income with respect to parent income is ill-defined at quantiles with zero income.

Empirical Estimates. In large samples, we can estimate $E\left[Y_i|X_i=x\right]$ non-parametrically as shown in Figure IIIa. It is therefore straightforward to estimate IGE_W simply by regressing the non-parametric estimates of $\log E\left[Y_i|X_i=x\right]$ on $\log x$.[73] The series in circles in Appendix Figure IVa plots $\log E\left[Y_i|X_i=x\right]$ vs. $\log x$ simply by taking logs of the data points shown in Figure IIIa.[74] As a reference, we also replicate the non-parametric plot of $E\left[\log Y_i|X_i=x\right]$ vs. $\log x$ from Figure IIIb (excluding children with zero income) in the series in triangles.

The two series are very similar, implying that nonparametric estimates of $\frac{d\log E[Y_i|X_i=x]}{d\log x}$ and $\frac{dE[\log Y_i|X_i=x]}{d\log x}$ are very similar for most values of x. The dollar-weighted IGE estimate (including children with zero income) is $IGE_W = 0.335$, virtually identical to the traditional IGE estimate of $IGE = 0.344$ obtained when we exclude children with zero income. Between the 10th and 90 percentiles, the dollar-weighted IGE is 0.414, while the traditional IGE is 0.452.

In Appendix Figure IVb, we report dollar-weighted IGE estimates by the age of the child to assess the extent of lifecycle bias in the dollar-weighted IGE estimates. We find that the dollar-weighted IGE also stabilizes around age 30: the estimated IGE_W is 2.1% higher at age 32 than

---

[73]Mitnik et al. use a Poisson pseudo-maximum-likelihood (PPML) estimator to estimate IGE_W in survey data, which approximates $\log E\left[Y_i|X_i=x\right]$ in large samples.

[74]Unlike in Figure IIIa, we include the top bin (the top 1% of parents) in this figure.

age 31 (0.343 vs. 0.336).

Mitnik et al. (2014) obtain larger estimates of the dollar-weighted IGE (around 0.5) in their SOI sample. Although both studies use similar data from tax records, there are several small methodological differences between Mitnik et al.'s approach and our approach. A useful direction for future work would be to investigate which of these differences in responsible for the differences in the IGE estimates.
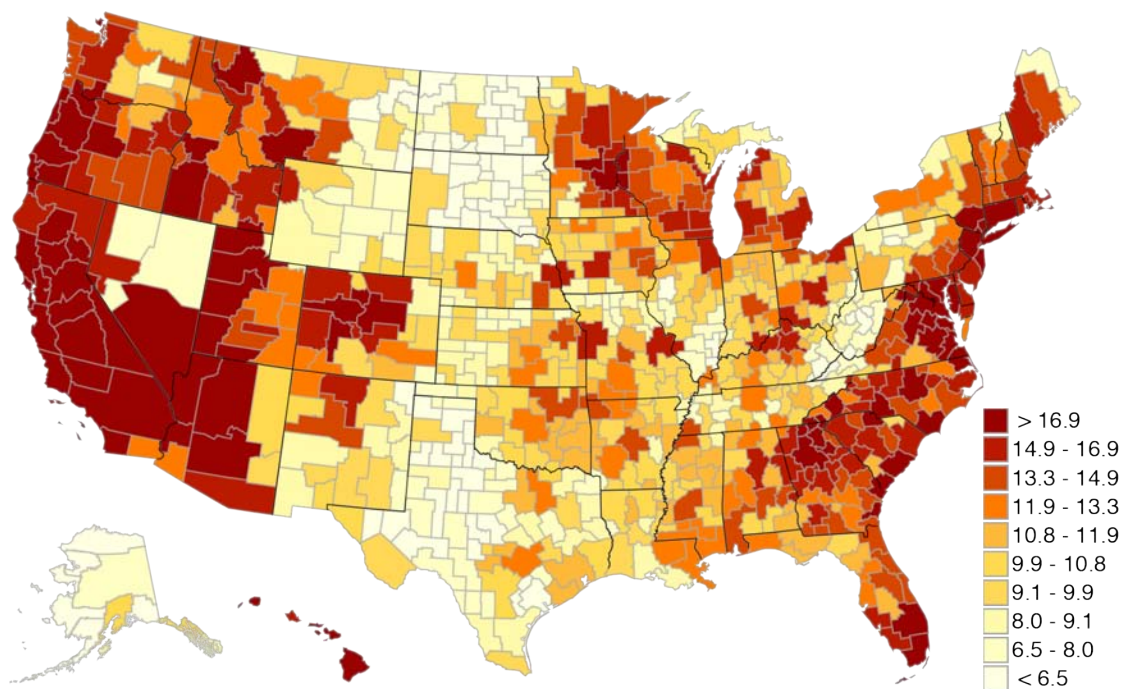
References

Autor, David H. and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review, 103* (5): 1553–97.

Bailey, Martha J. and Susan M. Dynarski. 2011. "Gains and Gaps: Changing Inequality in U.S. College Entry and Completion." NBER Working Paper 17633, National Bureau of Economic Research, Inc.

Baum-Snow, Nathaniel and Justin Marion. 2009. "The Effects of Low Income Housing Tax Credit Developments on Neighborhoods." *Journal of Public Economics, 93*, 654–666.

Bhattacharya, Debopam and Bhashkar Mazumder. 2011. "A nonparametric analysis of black–white differences in intergenerational income mobility in the United States." *Quantitative Economics, 2* (3): 335–379.

Bjorklund, Anders and Markus Jäntti. 1997. "Intergenerational Income Mobility in Sweden Compared to the United States." *American Economic Review, 87* (5): 1009–18.

Black, Dan, Seth Sanders, Evan Taylor, and Lowell Taylor. 2011. "The Impact of the Great Migration on Mortality of African Americans: Evidence from the Deep South." Unpublished Univ. of Chicago mimeo.

Black, Sandra E. and Paul J. Devereux. 2011. "Recent Developments in Intergenerational Mobility." in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 4, Elsevier, chapter 16, pp. 1487–1541.

Bourguignon, François and Amedeo Spadaro. 2012. "Tax–benefit revealed social preferences." *Journal of Economic Inequality, 10* (1): 75–108.

Cameron, Stephen V. and James J. Heckman. 2001. "The Dynamics of Educational Attainment for Black, Hispanic, and White Males." *Journal of Political Economy, 109* (3): 455–499.

Chadwick, Laura and Gary Solon. 2002. "Intergenerational Income Mobility Among Daughters." *American Economic Review, 92* (1): 335–344.

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff. 2014 forthcoming. "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood." *American Economic Review.*

Cilke, James. 1998. "A Profile of Non-Filers." U.S. Department of the Treasury, Office of Tax Analysis Working Paper No. 78.

Clotfelter, Charles T. "Tax Incentives and Charitable Giving: Evidence from a Panel of Tax-payers." in "Proceedings of the Federal Reserve Bank of Boston" 1985. pp. 187–221.

Corak, Miles. 2013. "Income Inequality, Equality of Opportunity, and Intergenerational Mobility." *Journal of Economic Perspectives, 27* (3): 79–102.

Corak, Miles and Andrew Heisz. 1995. "The Intergenerational Income Mobility of Canadian Men." *Canadian Business Economics*, p. 59–69.

Corak, Miles and Andrew Heisz. 1999. "The Intergenerational Earnings and Income Mobility of Canadian Men: Evidence from Longitudinal Income Tax Data." *Journal of Human Resources*, *34* (3): 504–533.

Dahl, Molly W. and Thomas DeLeire. 2008. *The association between children's earnings and fathers' lifetime earnings: estimates using administrative data*, University of Wisconsin-Madison, Institute for Research on Poverty.

Graham, Bryan and Patrick Sharkey. 2013. "Mobility and the Metropolis: How Communities Factor Into Economic Mobility." *Washington D.C.: Economic Mobility Project, The Pew Charitable Trusts.*

Grawe, Nathan D. 2006. "Lifecycle bias in estimates of intergenerational earnings persistence." *Labour Economics*, *13* (5): 551–570.

Grawe, Nathan D. and Casey B. Mulligan. 2002. "Economic Interpretations of Intergenerational Correlations." *Journal of Economic Perspectives*, *16* (3): 45–58.

Haider, Steven and Gary Solon. 2006. "Life-Cycle Variation in the Association between Current and Lifetime Earnings." *American Economic Review, 96* (4): 1308–1320.

Hertz, Tom. 2006. "Understanding mobility in America." *Center for American Progress Discussion Paper.*

Internal Revenue Service. 2008. "SOI Tax Stats - Individual Income Tax Statistics - ZIP Code Data (SOI): 2008." Technical Report, Washington, DC: government printing press.

Internal Revenue Service. 2013. "Statistics of Income: Individual Income Tax Returns, 2012." Technical Report, Washington, DC: government printing press.

Jäntti, Markus, Bernt Bratsberg, Knut Røed, Oddbjørn Raaum, Robin Naylor, Eva Österbacka, Anders Björklund, and Tor Eriksson. 2006. "American Exceptionalism in a New Light: A Comparison of Intergenerational Earnings Mobility in the Nordic Countries, the United Kingdom and the United States." IZA Discussion Paper 1938, Institute for the Study of Labor (IZA).

Lee, Chul-In and Gary Solon. 2009. "Trends in Intergenerational Income Mobility." *The Review of Economics and Statistics, 91* (4): 766–772.

Levitis, Jason and Jeremy Koulish. 2008. "State Earned Income Tax Credits: 2008 Legislative Update."

Mazumder, Bhashkar. 2005. "Fortunate Sons: New Estimates of Intergenerational Mobility in the United States Using Social Security Earnings Data." *The Review of Economics and Statistics, 87* (2): 235–255.

Meyer, Bruce D. 2010. "The Effects of the Earned Income Tax Credit and Recent Reforms." in "Tax Policy and the Economy, Volume 24" NBER Chapters, National Bureau of Economic Research, Inc, pp. 153–180.

Mulligan, Casey B. 1997. *Parental priorities and economic inequality*, Chicago: University of Chicago Press.

Neal, Derek A and William R Johnson. 1996. "The Role of Premarket Factors in Black-White Wage Differences." *Journal of Political Economy*, *104* (5): 869–95.

Ray, Debraj. 2010. "Uneven Growth: A Framework for Research in Development Economics." *Journal of Economic Perspectives*, *24* (3): 45–60.

Rosen, Harvey S. 1985. "Housing subsidies: Effects on housing decisions, efficiency, and equity." *Handbook of Public Economics*, *1* (1): 375–420.

Saez, Emmanuel. 2001. "Using Elasticities to Derive Optimal Income Tax Rates." *Review of Economic Studies*, *68* (1): 205–29.

Smeeding, Timothy M., Markus Jäntti, and Robert Erikson. 2011. *Persistence, privilege, and parenting: the comparative study of intergenerational mobility*, New York: Russell Sage Foundation.

Solon, Gary. 1992. "Intergenerational Income Mobility in the United States." *American Economic Review*, *82* (3): 393–408.

Solon, Gary. 1999. "Intergenerational Mobility in the Labor Market." in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 3, Elsevier, pp. 1761–1800.

Solon, Gary. 2002. "Cross-Country Differences in Intergenerational Earnings Mobility." *Journal of Economic Perspectives*, *16* (3): 59–66.

Tolbert, Charles M. and Molly Sizer. 1996. "U.S. Commuting Zones and Labor Market Areas: A 1990 update." *Economic Research Service Staff Paper*, *9614*.

Zimmerman, David J. 1992. "Regression toward Mediocrity in Economic Stature." *American Economic Review*, *82* (3): 409–29.

FIGURE I: Overall Tax Expenditures
(% Tax Expenditures/AGI)

Legend:
- > 16.9
- 14.9 - 16.9
- 13.3 - 14.9
- 11.9 - 13.3
- 10.8 - 11.9
- 9.9 - 10.8
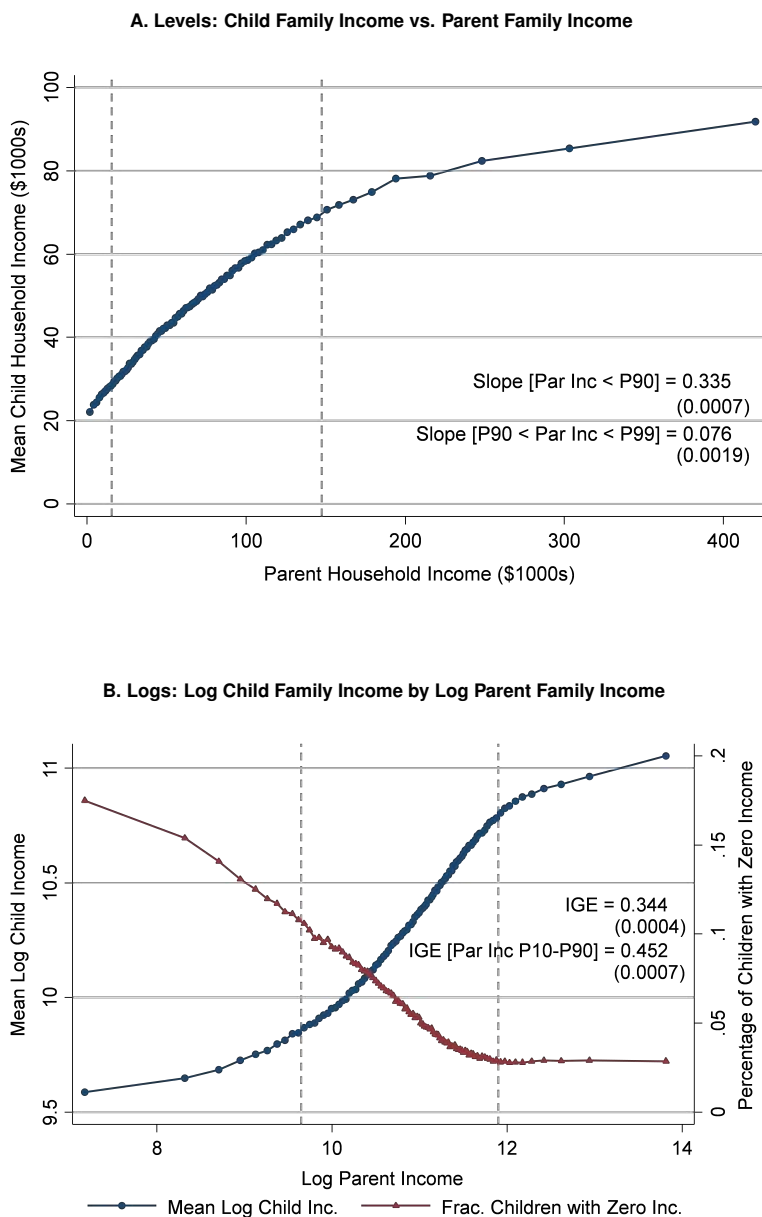- 9.1 - 9.9
- 8.0 - 9.1
- 6.5 - 8.0
- < 6.5

Notes: This figure maps CZ mean tax expenditures as a percentage of average Adjusted Gross Income (AGI). Darker areas represent higher tax expenditures. Data are from the IRS Statistics of Income ZIP Code Individual Income Statistics (2008).

FIGURE II: Progressivity of Tax Expenditures
(% Tax Expenditures/AGI, for Lowest - Highest AGI Class)

■	> 44.3
■	25.7 - 44.3
■	15.5 - 25.7
■	10.4 - 15.5
■	7.2 - 10.4
■	4.8 - 7.2
■	2.4 - 4.8
■	-0.1 - 2.4
■	-3.4 - -0.1
■	< -3.4
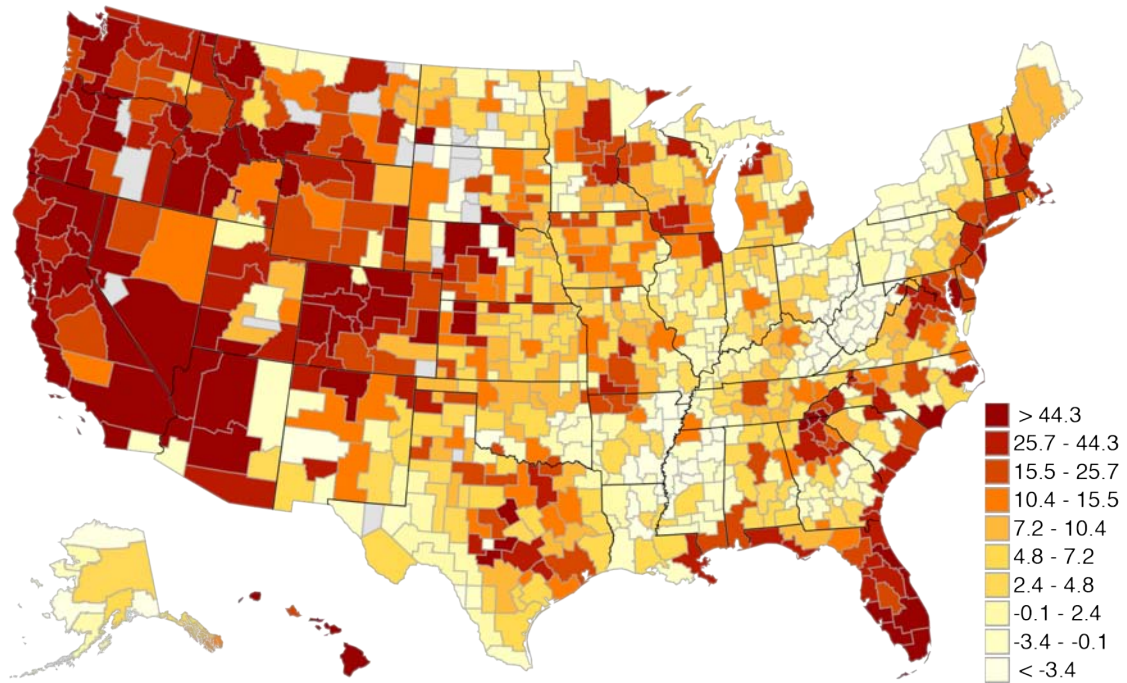
Notes: This figure maps CZ difference in mean tax expenditures as a percentage of average Adjusted Gross Income (AGI) for individuals with less than $10,000 AGI and individuals with over $200,000 AGI. Darker areas represent more progressive tax expenditures. Data are from the IRS Statistics of Income ZIP Code Individual Income Statistics (2008).

FIGURE III: The Association Between children's and Parents' Income in the United States

A. Levels: Child Family Income vs. Parent Family Income

Slope [Par Inc < P90] = 0.335
(0.0007)
Slope [P90 < Par Inc < P99] = 0.076
(0.0019)

B. Logs: Log Child Family Income by Log Parent Family Income

IGE = 0.344
(0.0004)
IGE [Par Inc P10-P90] = 0.452
(0.0007)

Mean Log Child Inc. — Frac. Children with Zero Inc.

Notes: Panel A depicts mean child family income (y-axis) against mean parent family income (x-axis). Panel B depicts mean child log family income (left y-axis) and the fraction of children with zero family income (right y-axis) against mean parents' log family income (x-axis). The graphs are based on the core sample of all children born in 1980-82 who are current U.S. citizens and can be matched to parents through tax return data (see Table 1). Child income is the average of 2011-2012 (when the child was around 30) while parent income is the average of 1996-2000 (when the child was a teenager). All dollar values are reported in 2010 dollars. To construct the figures, we group parents by percentiles of income (excluding parents with zero or negative income) and calculate the mean of the y-axis variable in each percentile. In panel B, children with zero family income are naturally excluded in the log-income means. The 10th and 90th percentile of parents' income are depicted in dashed vertical lines. Panel A excludes the top percentile for readability. Both panels show a nonlinear relationship between parent and child income with lower slopes in the upper tail in panel A and both the upper and lower tails in panel B. The slopes obtained using OLS regressions on the binned data are reported. In panel A, we report separate slopes below and above the 90th percentile. In panel B, we report slopes of the log-log regression in the full sample and for parents between the 10th and 90th percentiles. Panel B estimates are conventional intergenerational income elasticities. Such elasticities are sensitive to excluding the tails for parents' income and the treatment of children with zero income.

FIGURE IV: The Association Between children's Percentile Rank and Parents' Percentile Rank
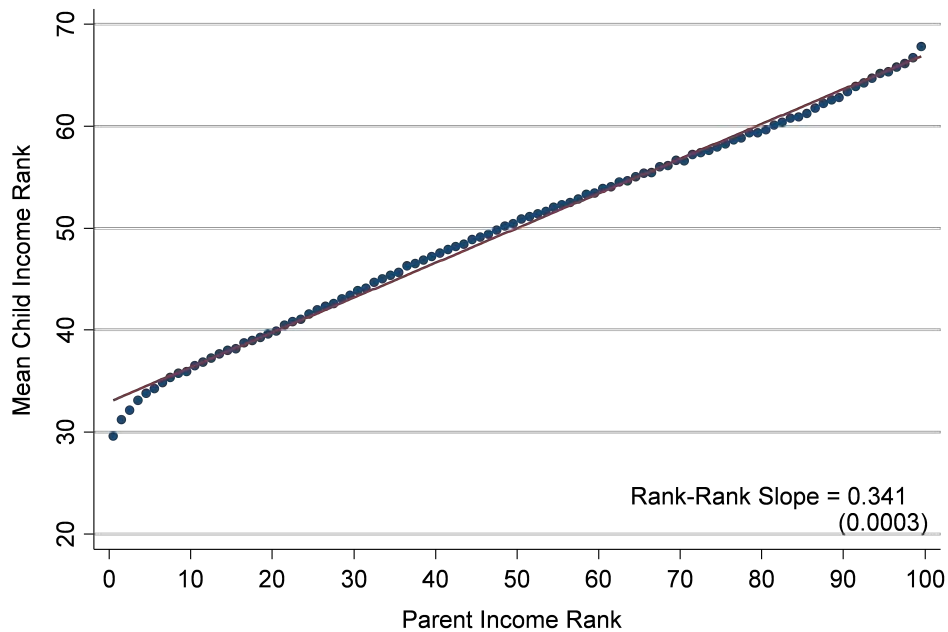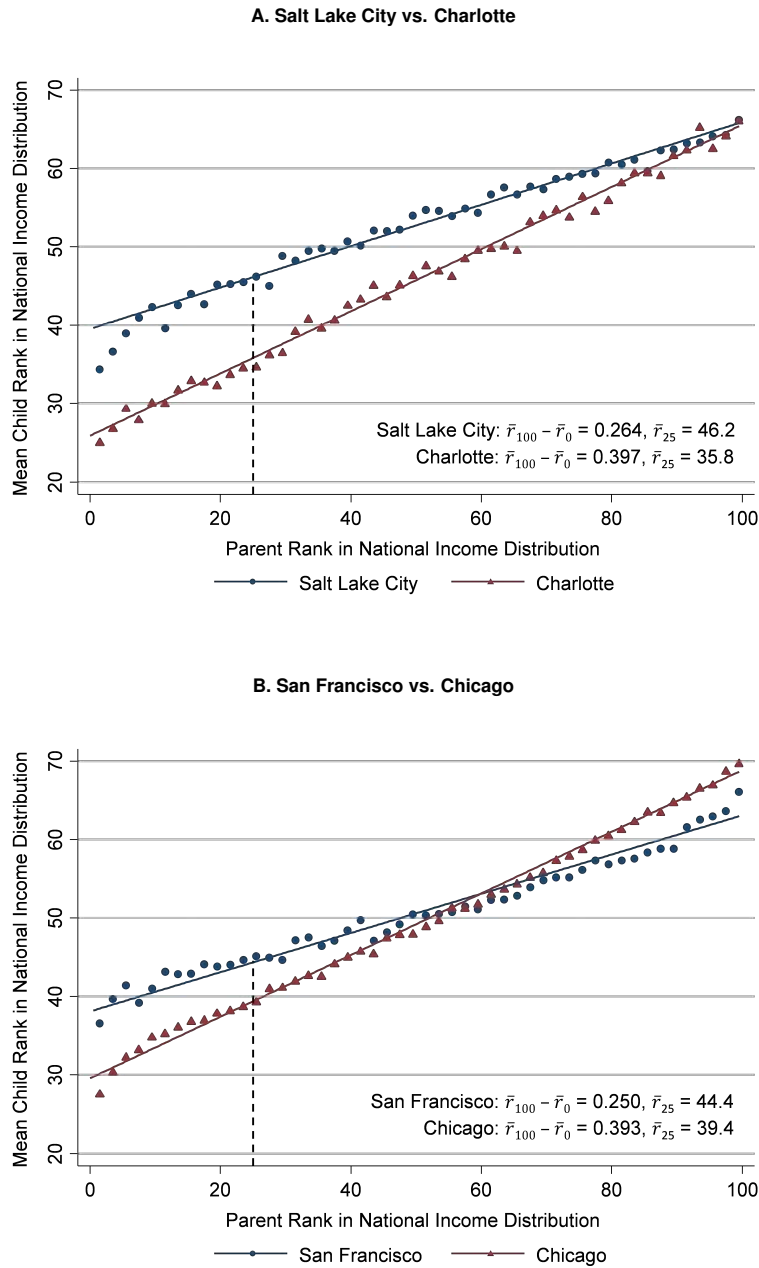

Notes: The figure depicts mean family income rank of children (y-axis) by percentile rank of family income of parents' (x-axis). The US series are based on the core sample of all children born in 1980-82 who are current U.S. citizens and can be matched to parents through tax return data (see Table 1). Child income is the average of 2011-2012 (when the child was around 30) while parent income is the average of 1996-2000 (when the child was a teenager). To construct the figure, we group parents by percentiles of income (excluding parents with zero or negative income) and calculate the mean rank of child family income for each percentile of parents' income. The rank-rank relationship is very close to linear along the full distribution of parental income with a slope of 0.341. Slopes are estimated with an OLS regression on the microdata.

FIGURE V: Robustness of Intergenerational Mobility Estimates

A. Lifecycle Bias: Rank-Rank Slopes by Age of Child

B. Attenuation Bias: Rank-Rank Slopes by Number of Years Used to Measure Parent Income
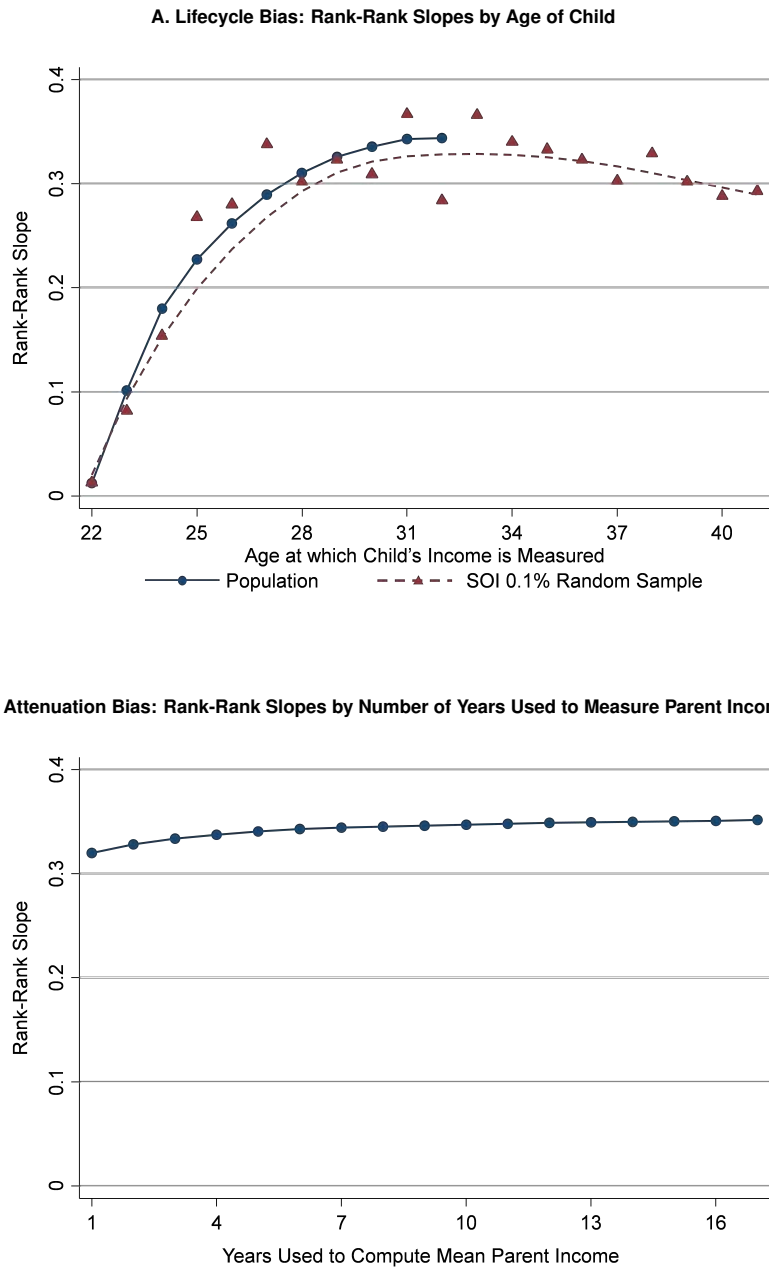
Notes: The figure depicts the robustness of the rank-rank slope coefficient from Figure 2A to changes in the age at which child income is measured (Panel A) and the number of years used to measure parent's income (Panel B). Child income is always the average of 2011-2012. Each point in Panel A shows the slope coefficient from an OLS regression of child income rank on parent income rank as we vary the child cohort and hence the age at which child income is measured. The first point corresponds to the 1990 cohort, who are around 22 when income is measured. The last point corresponds to the 1980 cohort, who are around 32 when income is measured. The slope increases sharply from zero at age 22 (when many children are still in college) to 0.34 at age 32. The slope increases very slowly after age 29 suggesting that measuring child income at age 30-32 is acceptable. In Panel B, we always consider the core sample of children born in 1980-2 who are current U.S. citizens and can be matched to parents through tax return data (see Table 1). Each point shows the slope coefficient from the same rank-rank regression, but here we vary the number of years used to compute mean parent income. The point for one year uses data for 1996 only. The point for two years adds data for 1997. We continue adding data until 2011, which gives us 16 years of data. The slope increases only modestly when increasing the number of years from 1 to 5 and is very stable for 5 or more years suggesting that using 5 years for parents' income is acceptable.
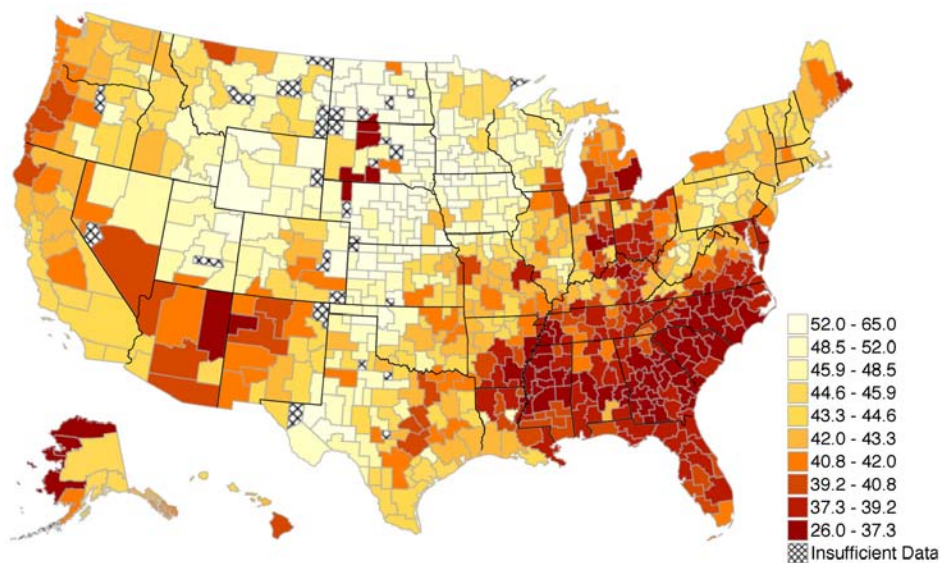
A. Salt Lake City vs. Charlotte

Salt Lake City: $\bar{r}_{100} - \bar{r}_0$ = 0.264, $\bar{r}_{25}$ = 46.2
Charlotte: $\bar{r}_{100} - \bar{r}_0$ = 0.397, $\bar{r}_{25}$ = 35.8

Mean Child Rank in National Income Distribution

Parent Rank in National Income Distribution

Salt Lake City   Charlotte

B. San Francisco vs. Chicago

San Francisco: $\bar{r}_{100} - \bar{r}_0$ = 0.250, $\bar{r}_{25}$ = 44.4
Chicago: $\bar{r}_{100} - \bar{r}_0$ = 0.393, $\bar{r}_{25}$ = 39.4

Mean Child Rank in National Income Distribution

Parent Rank in National Income Distribution

San Francisco   Chicago

Notes: The figure depicts mean child family income rank (y-axis) by percentile rank of parent family income (x-axis) for selected cities (Metropolitan Statistical Area). Panel A depicts Salt Lake City and Charlotte. Panel B depicts San Francisco and Chicago. The series are based on the core sample of all children born in 1980-82 who are current U.S. citizens and can be matched to parents through tax return data (see Table 1). Child income is the average of 2011-2012 (when the child was around 30) while parent income is the average of 1996-2000 (when the child was a teenager). Cities are based on the location of parents (when the child was a dependent). Parent and child percentile ranks is always defined at the national level, not the city level. To construct the figures, we group parents by national percentiles of income (excluding parents with zero or negative income) and calculate the mean national rank of child family income for each percentile of parrents' income. To reduce noise, percentiles are grouped in pairs. Both panels include the best linear fit. Dashed orange lines highlight our two intergenerational mobility concepts. The difference between the heights at the 100th percentile of parent income and the 0th percentile of parent income corresponds to Relative Mobility ($Y_{100} - Y_0$) for the same city. The height of the vertical line at the 25th percentile of parent income corresponds to Absolute Upward Mobility (Y_{25}) for the city in blue circles. Because of linearity, Absolute Upward Mobility (Y_{25}) is the same as ($Y_{[0,50]}$) which is the average rank of children with parents in the bottom half of the distribution reweighted using the national distribution.

FIGURE VII: The Geography of Intergenerational Mobility in the U.S.

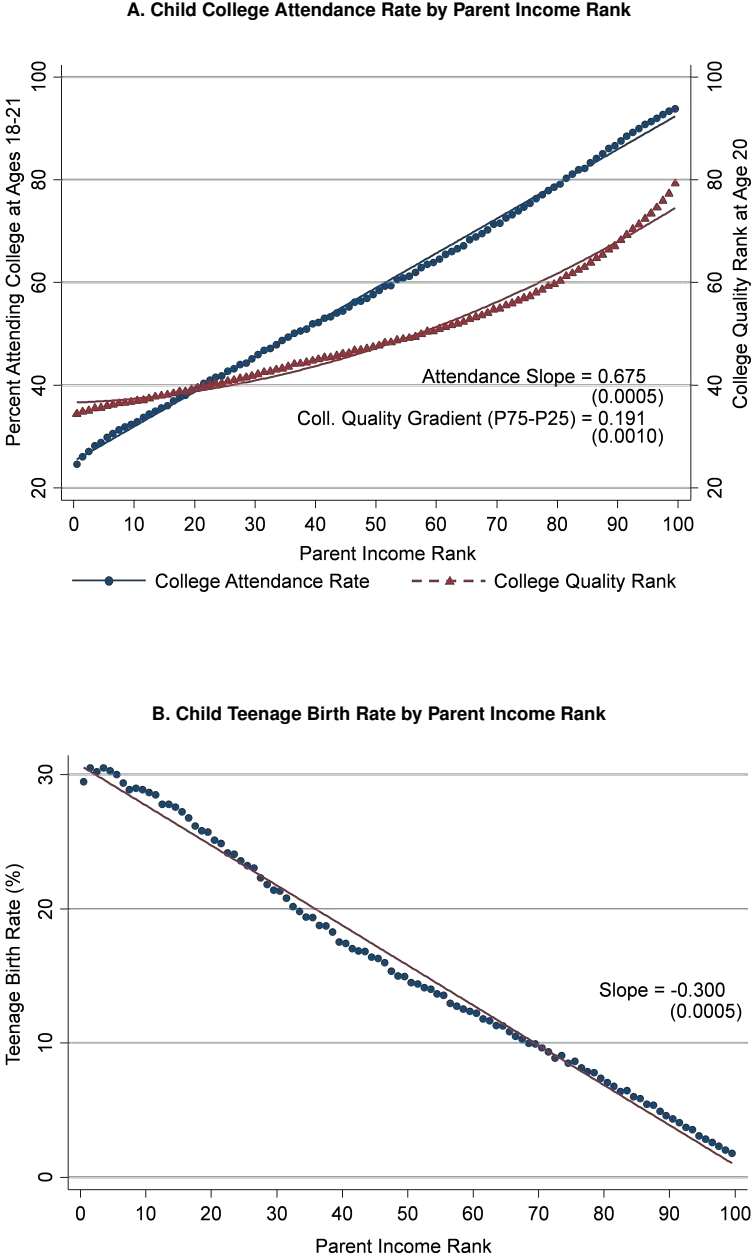**A. Absolute Upward Mobility: Mean Child Rank for Below-Median Parents** (Y_{25}) **by CZ**

	52.0 - 65.0
	48.5 - 52.0
	45.9 - 48.5
	44.6 - 45.9
	43.3 - 44.6
	42.0 - 43.3
	40.8 - 42.0
	39.2 - 40.8
	37.3 - 39.2
	26.0 - 37.3
	Insufficient Data

B. Relative Mobility: Rank-Rank Slopes ($Y_{100} - Y_0$) by CZ

	0.404 - 0.508
	0.381 - 0.404
	0.360 - 0.381
	0.346 - 0.360
	0.330 - 0.346
	0.312 - 0.330
	0.292 - 0.312
	0.270 - 0.292
	0.240 - 0.270
	0.068 - 0.240
	Insufficient Data
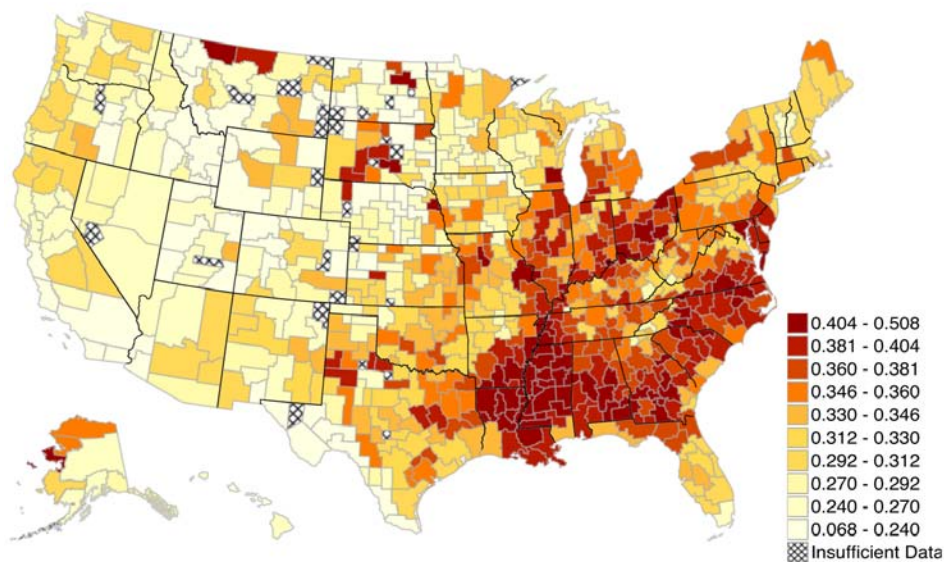
Corr. with baseline $\bar{r}_{25}$ = -0.68 (unweighted), -0.61 (pop-weighted)

Notes: The figure maps two of our baseline measures of intergenerational mobility by Commuting Zone (CZ). The series are based on the core sample of all children born in 1980-82 who are current U.S. citizens and can be matched to parents through tax return data (see Table 1). Child income is the average of 2011-2012 (when the child was around 30) while parent income is the average of 1996-2000 (when the child was a teenager). CZs are based on the location of parents (when the child was a dependent). In each CZ, we run an OLS regression of child income rank (at the national level) on a constant and parent income rank (at the national level). Panel A depicts Absolute Upward Mobility defined as the intercept + .25*slope (Y_{25}), which corresponds to the fitted value for parent income at the 25th percentile (see Figure 5). A light color implies more mobility. Panel B depicts Relative Mobility defined as the rank-rank slope coefficient ($Y_{100} - Y_0$) for each CZ (see Figure 5). A light color implies more mobility. Panels A and B show that both mobility measures are strongly correlated (the population-weighted correlation coefficient is -0.62).

FIGURE VIII: Gradients of College Attendance and Teenage Birth by Parent Family Income Rank

A. Child College Attendance Rate by Parent Income Rank

Attendance Slope = 0.675
(0.0005)
Coll. Quality Gradient (P75-P25) = 0.191
(0.0010)

College Attendance Rate —●— — —▲— — College Quality Rank

B. Child Teenage Birth Rate by Parent Income Rank

Slope = -0.300
(0.0005)

Notes: The figure depicts child college attendance rate (in Panel A) and child teenage birth rate for female children (in Panel B) by percentile rank of parent family income. The series are based on the core sample of all children born in 1980-82 who are current U.S. citizens and can be matched to parents through tax return data (see Table 1). Parent income is the average of 1996-2000 (when the child was a teenager). College attendance is defined as ever attending college from age 18 to 21, where attending college is defined as presence of a 1098-T form. Teenage birth is defined as having a child while being aged 19 or less. To construct the figures, we group parents by percentiles of income (excluding parents with zero or negative income) and calculate the mean of the y-axis variable in each percentile. The regression coefficients, standard errors, and best-fit lines are computed on the micro data. The rank-rank relationship is very close to linear in both panels. The slope for college attendance is extremely steep at 0.68. The slope of teenage birth rate is -0.22.

FIGURE IX: Relationship between Absolute and Relative Mobility

A. Empirical Pivot Point

A. Pivot Point Schematic
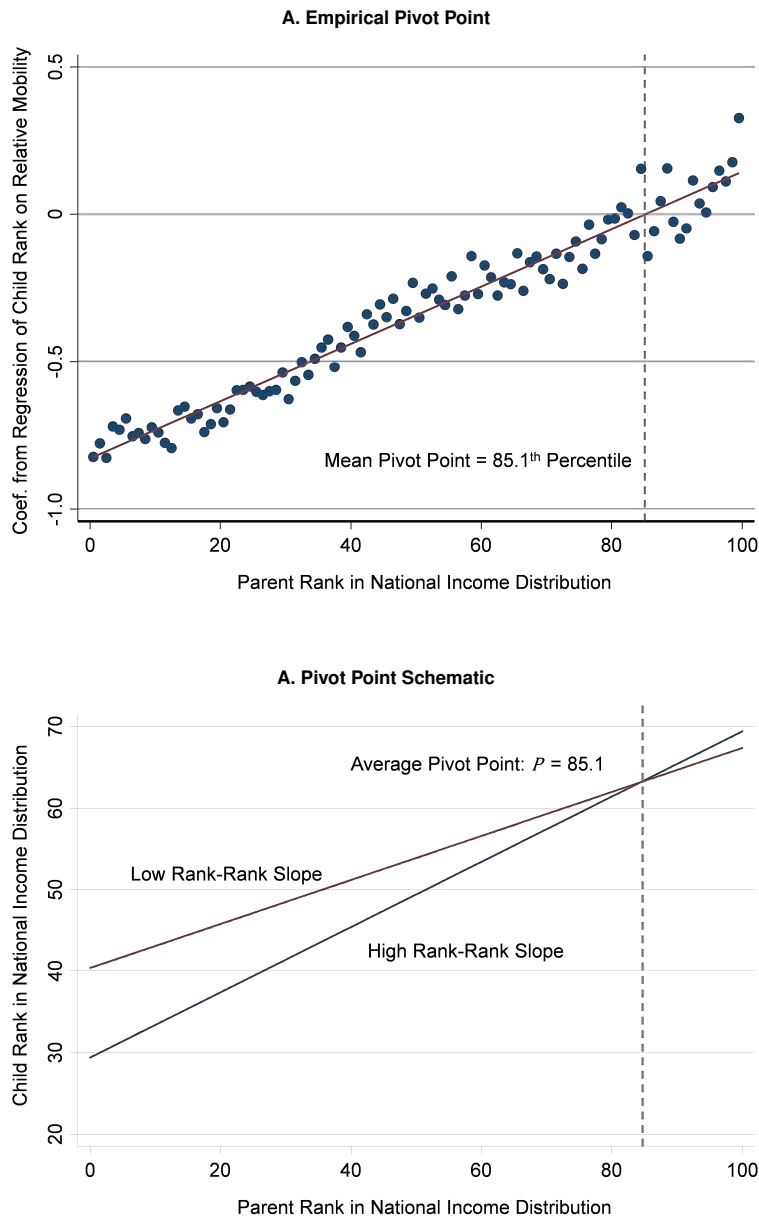
Notes: This schematic graph illustrates the relationship between absolute and relative mobility. On average, our empirical results show that there is more variation in outcomes across cities for children with parents at the bottom of the distribution (represented by the intercept with the y-axis at the 0th percentile in parent income in the rank-rank diagram) than for children with parents at the top of the distribution. On average, outcomes vary the least around percentile 85 for parents as if the rank-rank linear relationship across CZs were "pivoting" around parents' percentile 85 as depicted in the schematic.

FIGURE X: CZ Tax Expenditures and Intergenerational Income Correlation
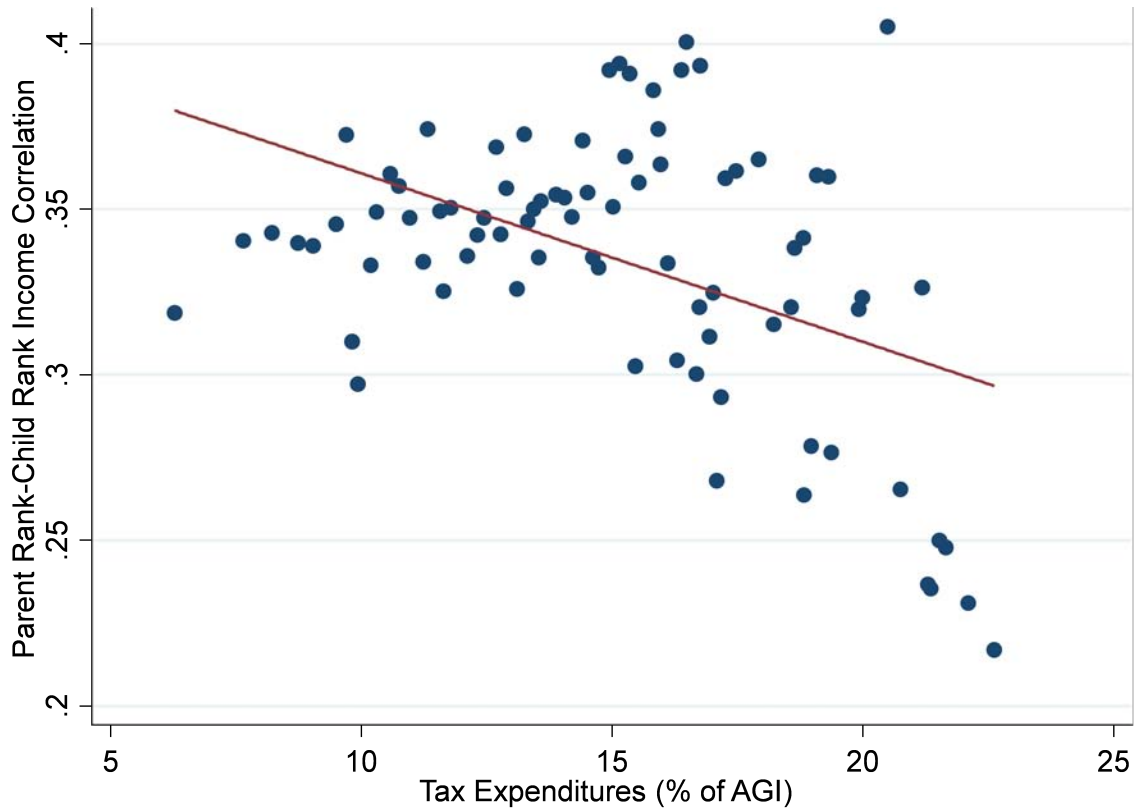
Notes: This figure displays a binned scatterplot of the relationship between CZ aggregate tax expenditures as a percentage of AGI in 2008 and the CZ IGE as measured by the correlation between parent rank income and child rank income. See Section 2.2 for more details on the construction of the tax expenditure and IGE measures. To generate the binned scatterplot, we group CZs into centiles (one-hundred equal-sized bins) on tax expenditures as a percentage of AGI, weighting by CZ population. The dots represent the weighted means of the IGE and tax expenditure measure. The best-fit line is calculated from a regression on the CZ level data.

FIGURE XI: CZ Progressivity of Tax Expenditures and Intergenerational Income Correlation
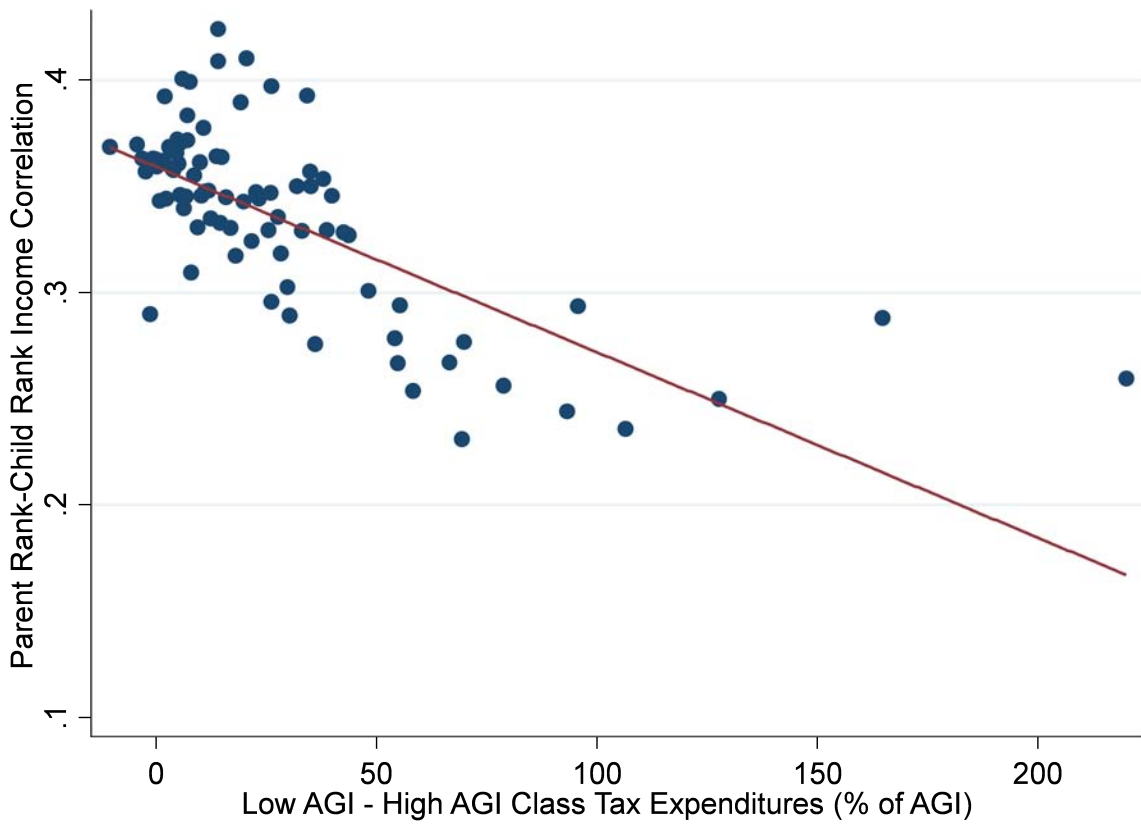
Notes: This figure displays a binned scatterplot of the relationship between progressivity of CZ tax expenditures and IGE as measured by the correlation between parent rank income and child rank income. Eight CZs with over 300% difference in tax expenditures are excluded from the figure and best-fit line. See Section 2.2 for more details on the construction of the tax expenditure and IGE measures. See notes to Figure 4 for further explanation of construction of the binned scatterplot.

FIGURE XII: Impact of Changing Racial Composition of Sample on CZ-Level Estimates of Upward Mobility
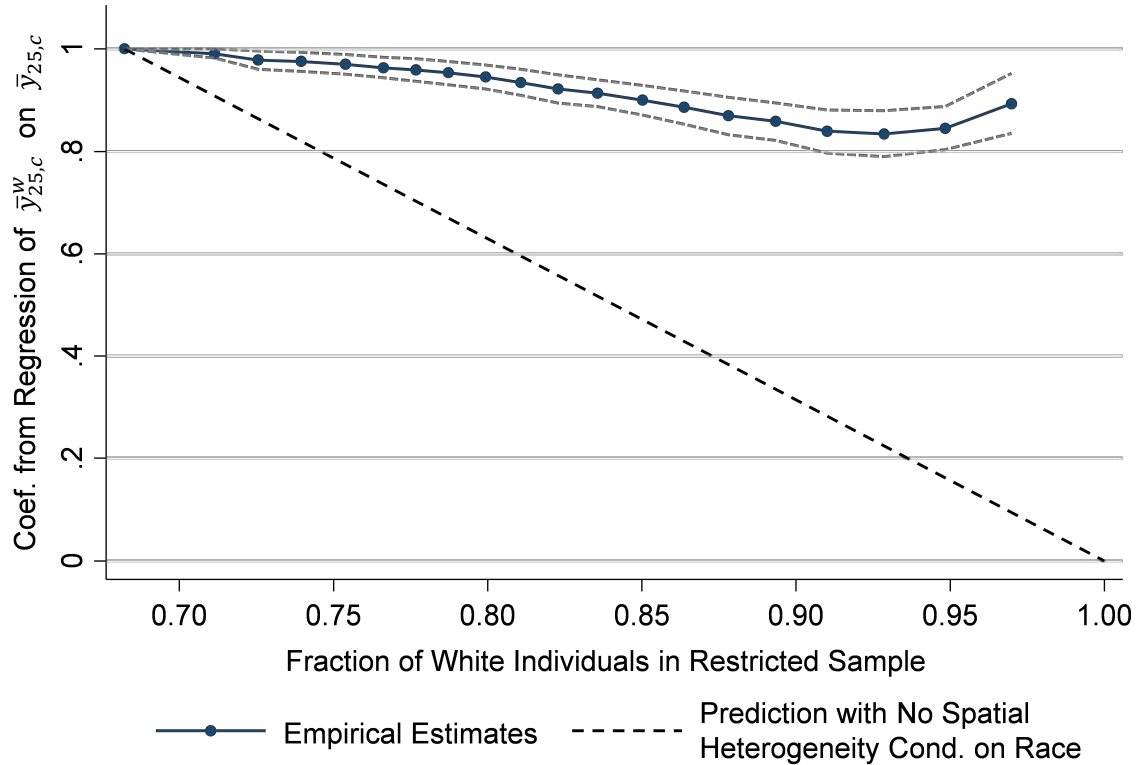
Notes: This figure plots the relationship between our main IGE rank-rank slope estimates and rank-rank slope estimates computed restricting to zip codes with a varying minimum threshold of the share of white residents. Fraction white and black is computed in the 2000 Census using 5-digit ZCTAs that are then matched to 5-digit address ZIP codes in 1999 (see text for complete details).We first compute Absolute Upward Mobility estimates by CZ limiting the sample to 5-digit zip codes with at least x% of white residents (based on the 2000 Census), which we denote as the x%+ white sample. The graph reports the regression coefficient from an unweighted regression of Upward Mobility estimates in the full core sample and in the x%+ white sample for x ranging between 0 and 95%. The confidence interval is based on standard errors that are clustered at the state level. The regression coefficient is above 0.8 even for high x showing that whites living in an CZ with many non-whites tend to experience the same mobility patterns as the full population in the CZ.

FIGURE XIII: Impact of Moving to a Different Area on Child's Outcomes

A. Impact of Neighborhood Quality on Childrens' Outcomes

Slope : 0.628
(0.048)
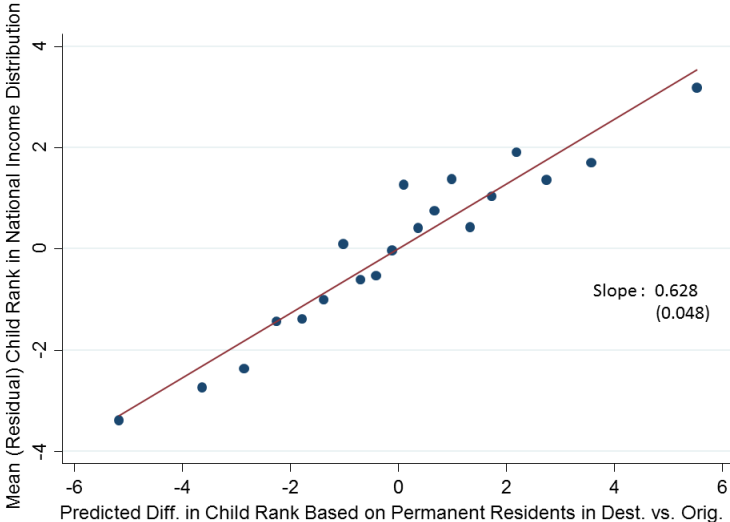
B. Impact of Neighborhood Quality on Childrens' Outcomes by Child's Age at Move

δ: 0.226

Slope: -0.038
(0.002)

Slope: -0.002
(0.011)
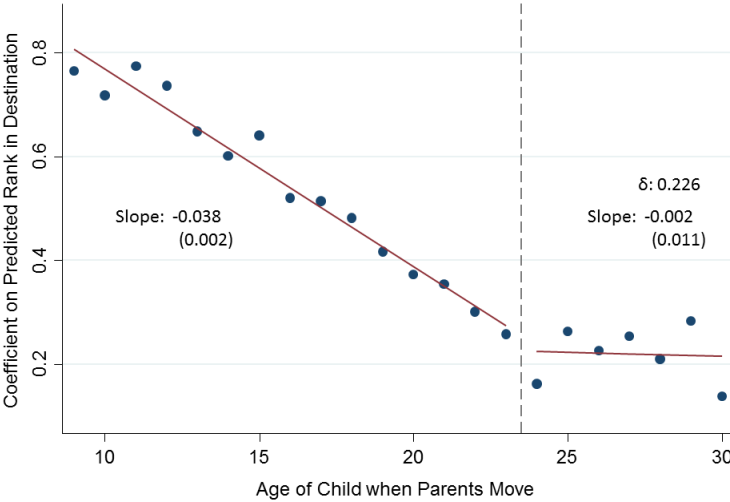
Notes: The top panel depicts the actual mean child income rank against the expected income rank of the child given the parent income and CZ of residence for children that moved before age 13. The bottom panel shows the effect of moving to a better city on child outcome by age at which the move takes place.

FIGURE XIV: Impact of Moving to a Different Area on Child's Outcomes (Family Fixed Effects)
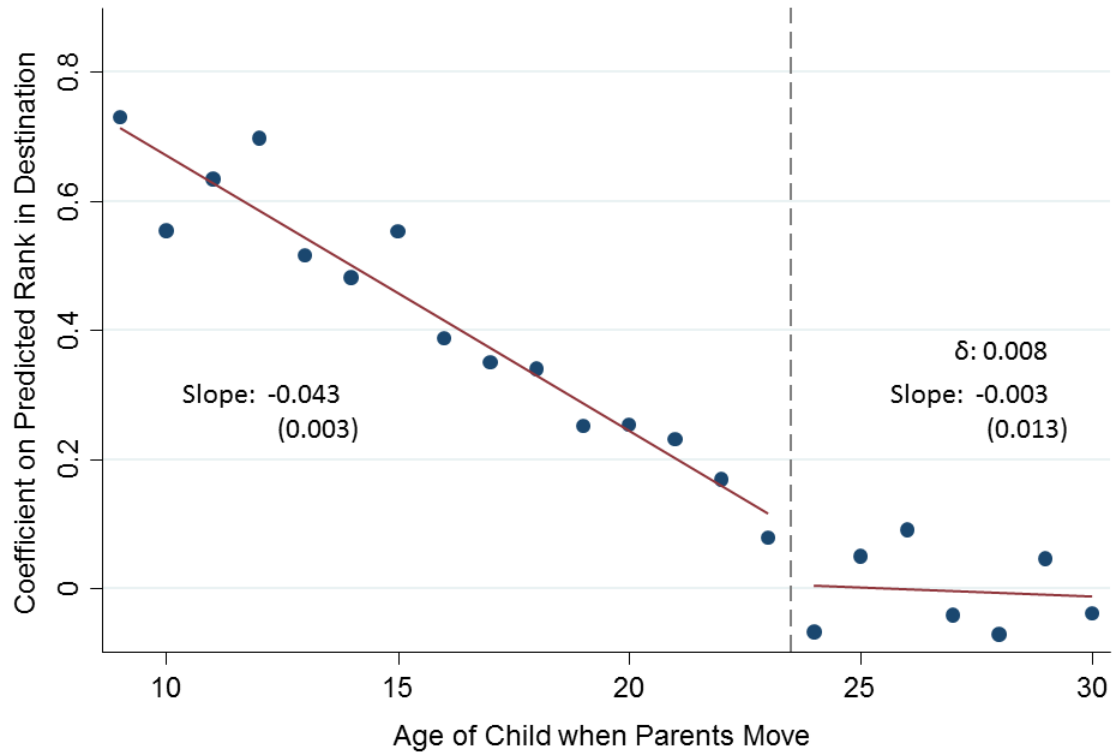
Notes: This figure replicates Figure XIIIb but controls for changes in marital status that are coincident with moves and includes family fixed effects.

FIGURE XV: Stayers Employment vs. Mean Rank as Predictors of Movers Employment
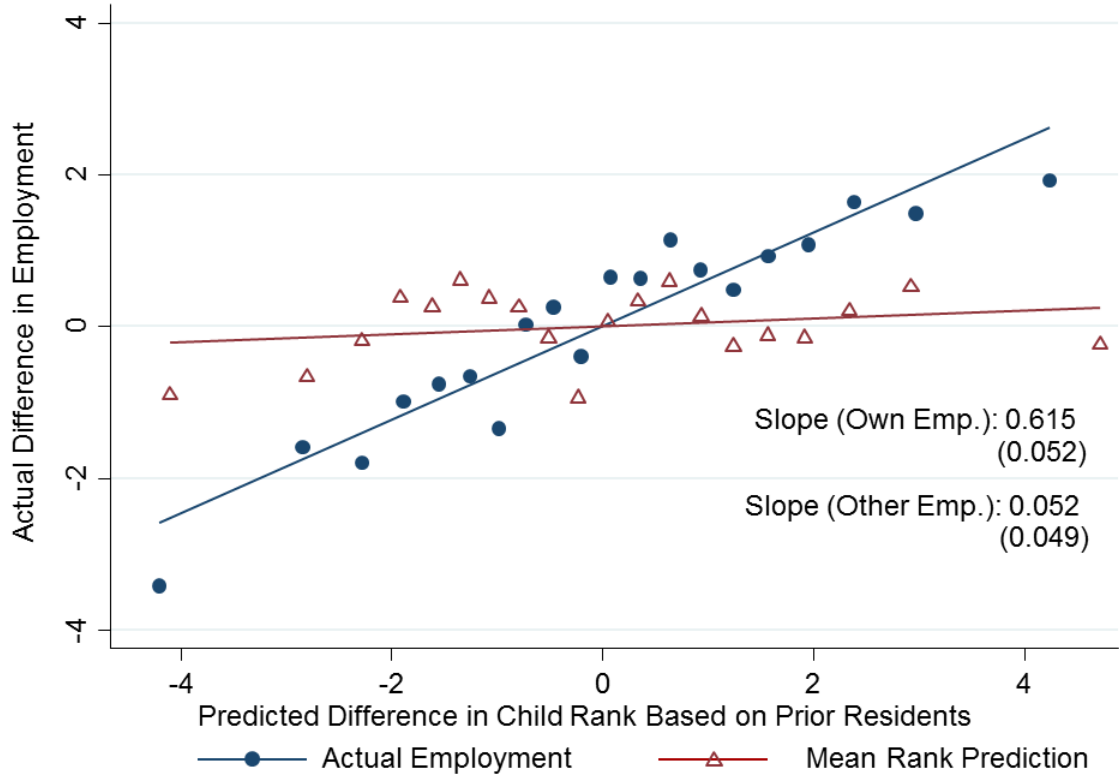
Notes: This figure shows that the employment rate of stayers is a sufficiently good predictor of the employment outcomes of movers and that income ranks add very little predictive power beyond employment rates. The figure uses the core sample of movers. The points correspond to binned means (in vingtile groups) from a regression of an indicator for being employed in 2011-12 on the expected difference in employment rates and child income ranks based on estimates from stayers. The regression also controls for the expected employment rate and child income rank of stayers in the origin.

FIGURE XVI: Difference in Absolute Upward Mobility for Boys vs. Girls
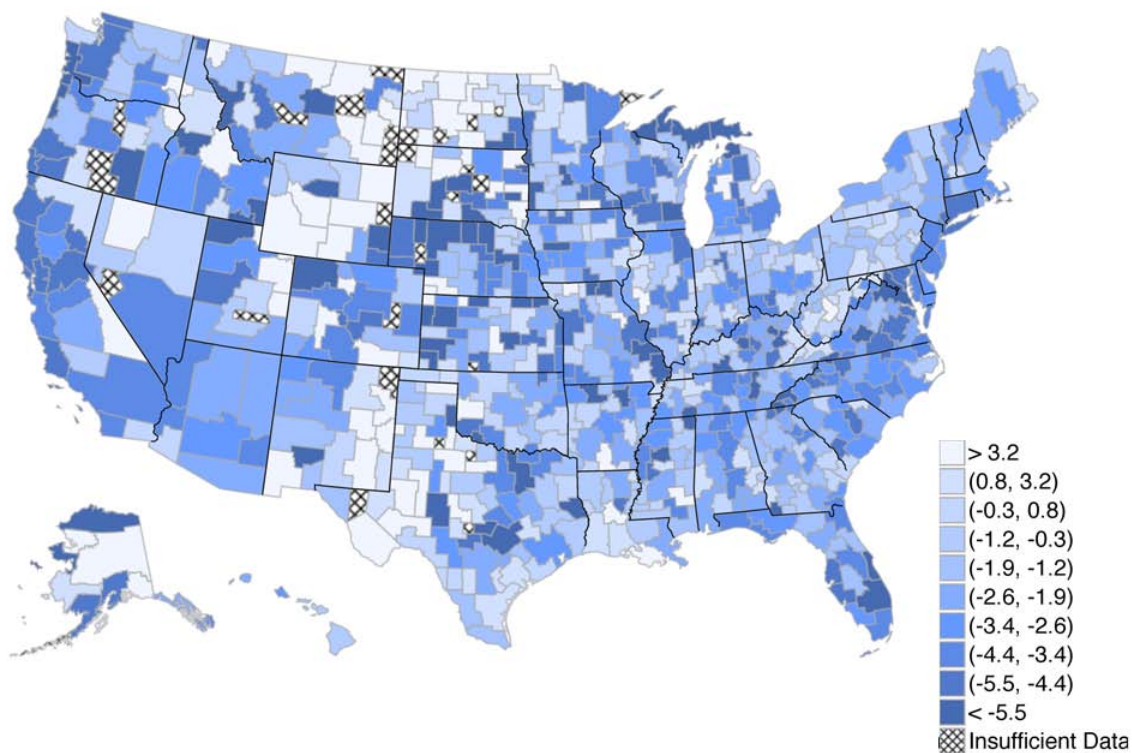
Notes: This map shows the difference in absolute upward mobility for boys and girls. The series are based on the core sample of all children born in 1980-82 who are current U.S. citizens and can be matched to parents through tax return data (see Table 1). Child income is the average of 2011-2012 (when the child was around 30) while parent income is the average of 1996-2000 (when the child was a teenager). CZs are based on the location of parents (when the child was a dependent). To construct absolute upward mobility for boys, we restrict the sample to boys and replicate the methodology used in Figure VIIa (similarly for girls).

FIGURE XVII: Cohort Specific Predictions and Actual Outcomes
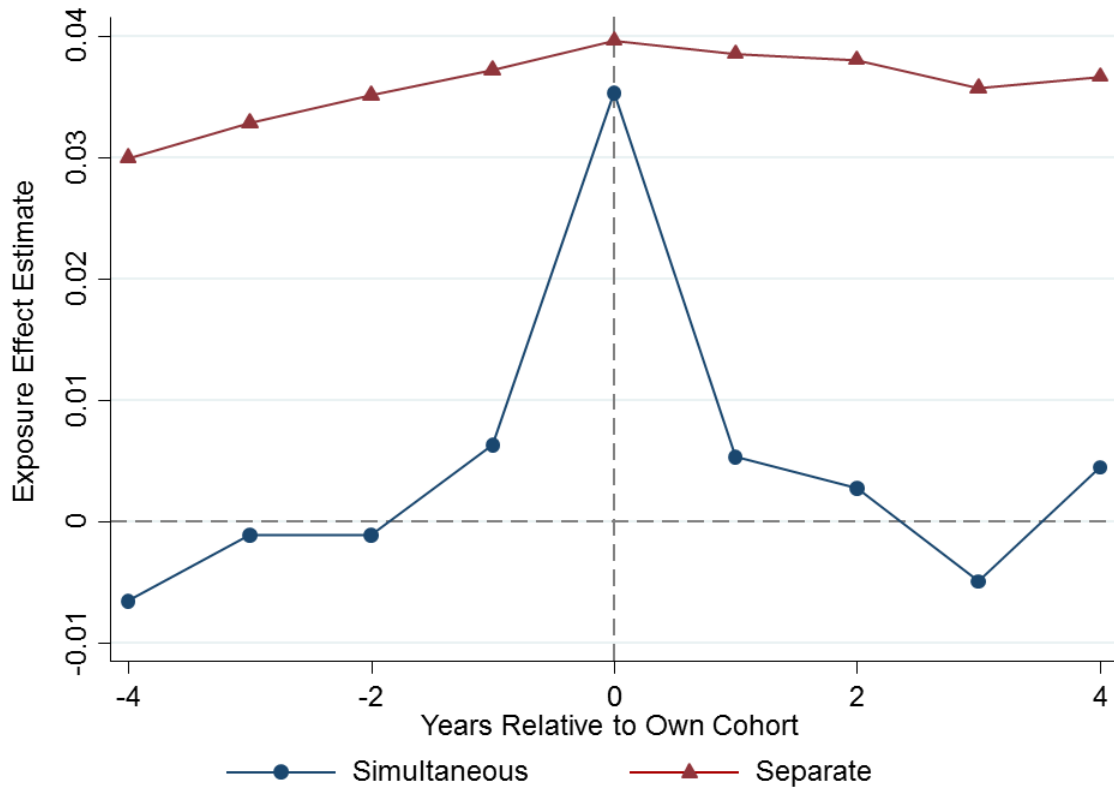
Notes: This figure shows that cohort specific predictions are more relevant for child outcomes than predictions using nearby cohorts. The figure uses the sample of movers. Each point is a coefficient from a regression of child income rank against cohort specific predictions. The series in circles plots the coefficient on place effects estimated, which is a cohort specific analog of the slope reported in Figure XIIIb (for age ≤ 23). The series in triangles plots the coefficient on the difference (between the destination and the origin) in expected child income rank based on stayers. The regressions also control for the slope in Figure XIIIb for age >23, the expected child income rank in their origin, time of move, and indicators for missing data.

FIGURE XVIII: Impact of Moving to a Different Area on Teen Labor
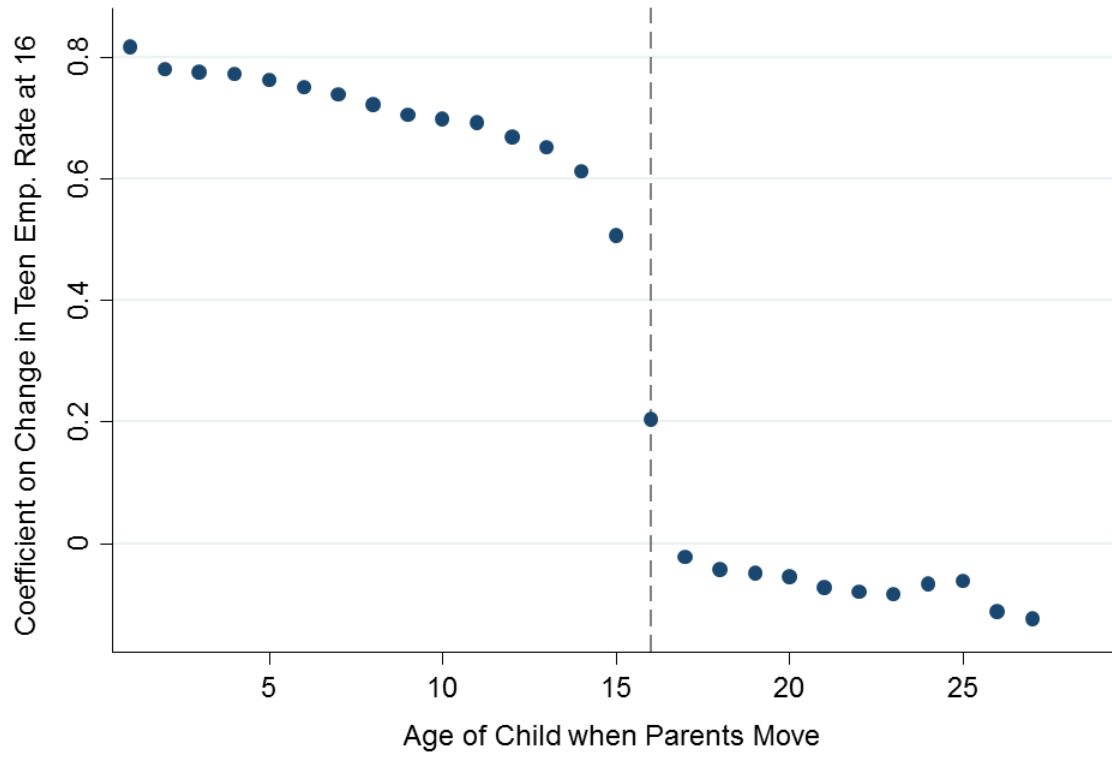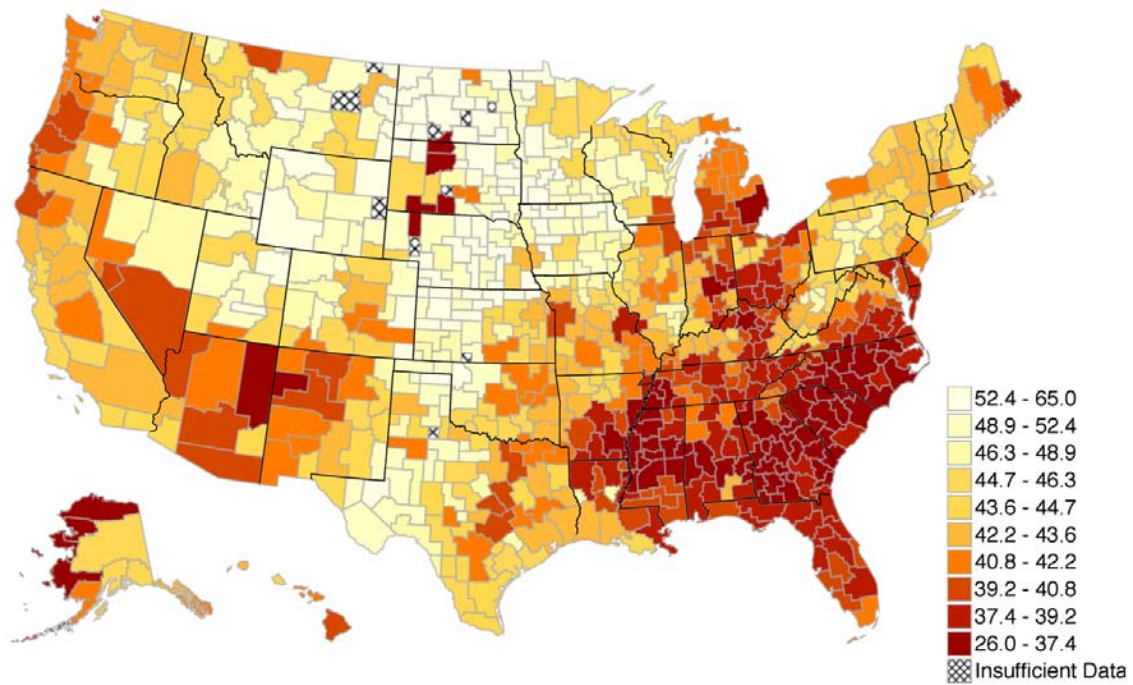
Notes: The bottom panel shows the effect on teen labor of moving to a better city on child outcome by age at which the move takes place.

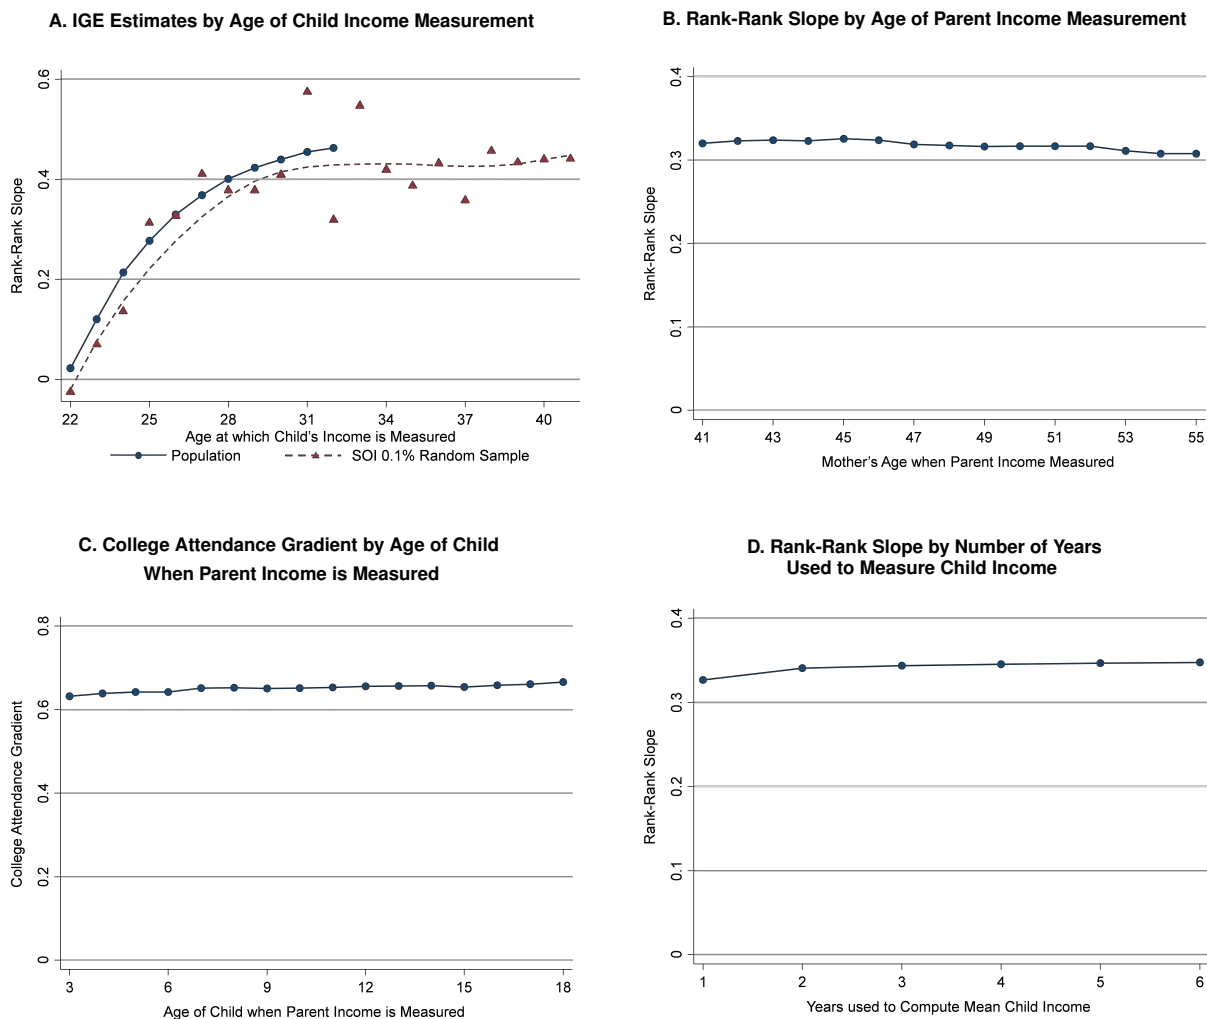APPENDIX FIGURE I: Rank-Rank Relationships and Relative Income Density in the 20 Largest Commuting Zones
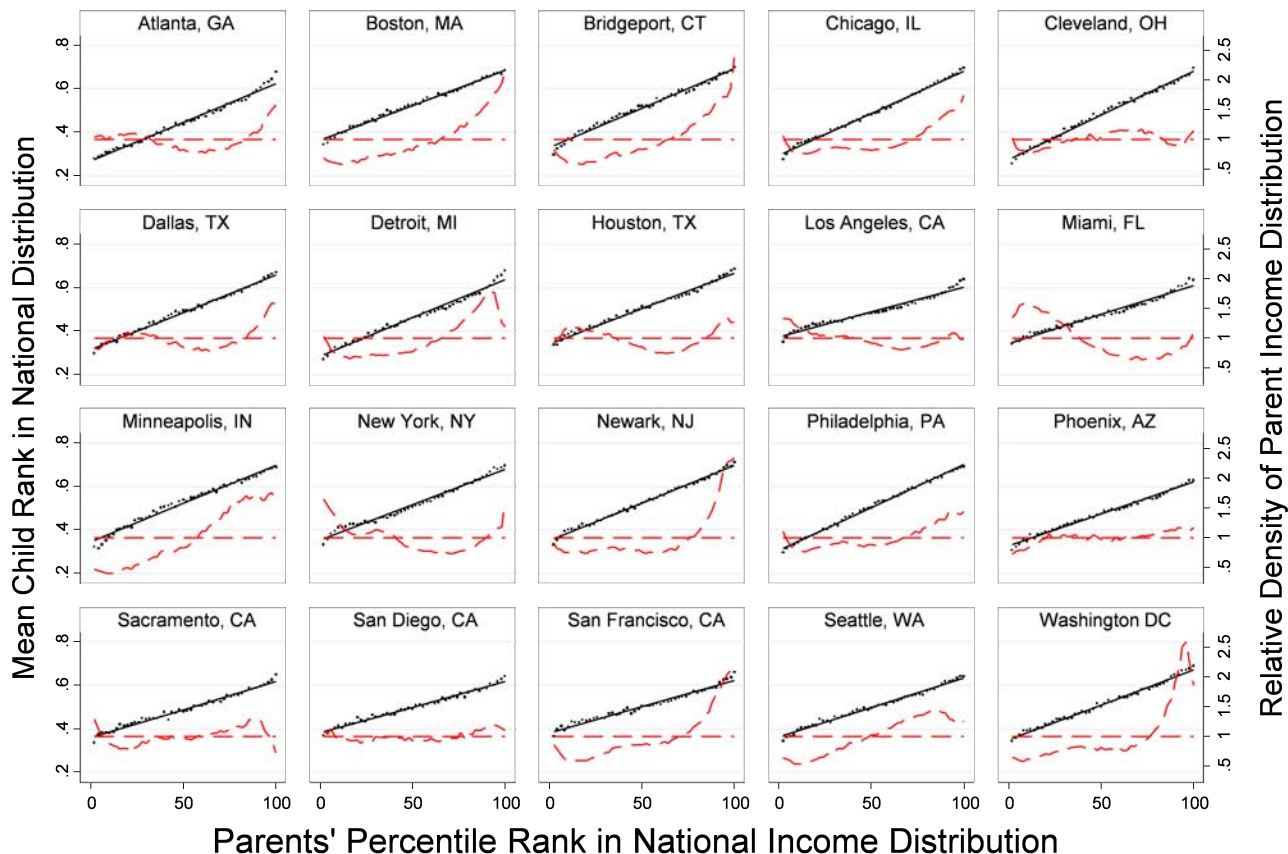
Notes: The figure depicts the rank-rank relationship and the relative parents' income distribution for the 20 largest CZs according to populations in the 2000 Census. The graphs are based on the core sample of all children born in 1980-1 who are current U.S. citizens and can be matched to parents through tax return data (see Table 1). In each CZ we group children into fifty bins according to their parents national income rank (the x-axis). Each of the fifty bins includes two percentiles. We calculate the mean child income rank in each bin in a black curve (left y-axis). We also plot the OLS linear fit to the binned data as a black line. The solid red curve shows the share of parents in each bin relative to the share in the national distribution (right y-axis). The red curve averages to one (denoted by the horizontal dashed line in each panel) in each city. The figure shows that the rank-rank linear fit is remarkably robust across cities, in spite of large differences in the local parental income distributions, which justifies using national ranks for estimating area based mobility measures.

APPENDIX FIGURE II: Absolute Upward Mobility 1980-82 and 1980-85 Cohorts

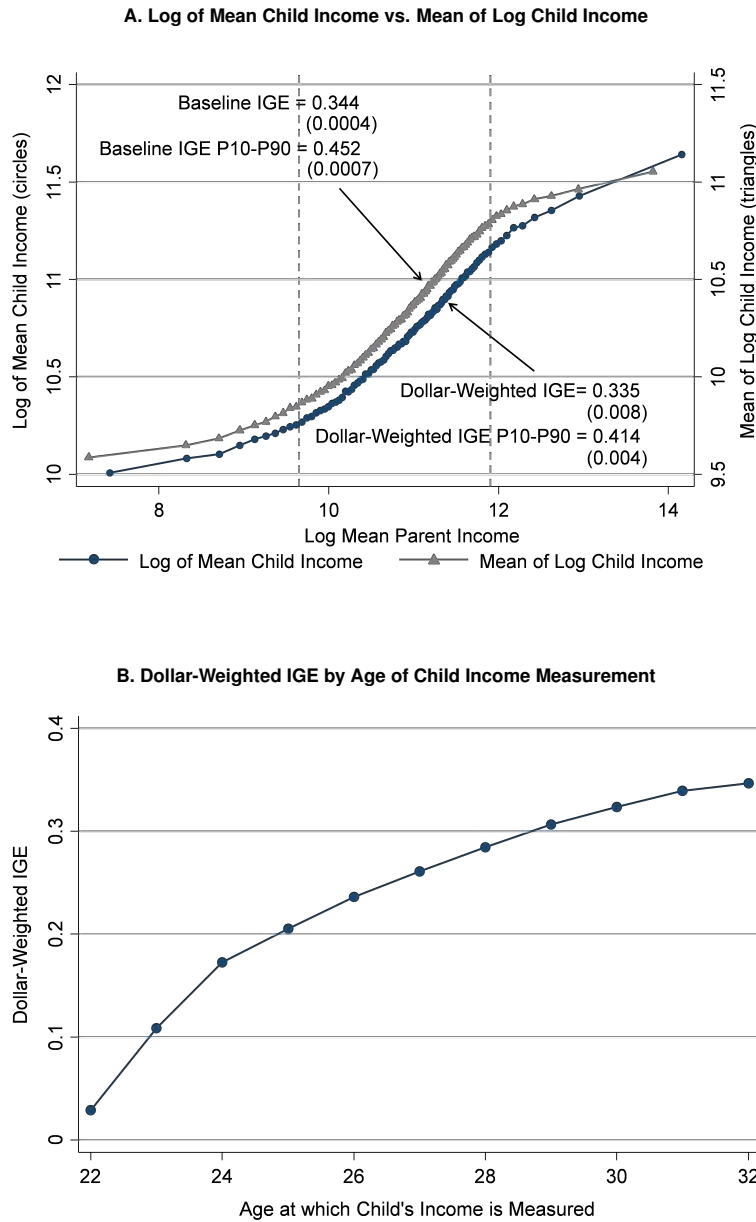|	
	52.4 - 65.0
	48.9 - 52.4
	46.3 - 48.9
	44.7 - 46.3
	43.6 - 44.7
	42.2 - 43.6
	40.8 - 42.2
	39.2 - 40.8
	37.4 - 39.2
	26.0 - 37.4
	Insufficient Data

Notes: The figure maps Absolute Upward Mobility by CZ. For the 709 CZs that have at least 250 children in the 1980-82 cohorts, the Absolute Mobility Measure is the same as in Figure VIa. For a further 22 CZs that have fewer than 250 children in the 1980-82 cohorts, but at least 250 children in the 1980-85 cohorts, we plot Absolute Upward Mobility by CZ computed on the 1980-85 cohorts sample - the details of the computation are equivalent to those describe for Figure VIa.

APPENDIX FIGURE III: Robustness of IGE Rank-Rank

A. IGE Estimates by Age of Child Income Measurement

B. Rank-Rank Slope by Age of Parent Income Measurement

C. College Attendance Gradient by Age of Child When Parent Income is Measured

**D. Rank-Rank Slope by Number of Years Used to Measure Child Income**

Notes: This figure evaluates the robustness of intergenerational mobility measures to lifecycle and attenuation bias. Panel A evaluates the robustness of the IGE to changes in the age at which child income is measured. Panel B evaluates the robustness of the rank-rank slope to changes in the age at which parent income is measured. Panel C evaluates the robustness of the college attendance gradient to the age of the child when parent income is measured. Panel D evaluates the robustness of the rank-rank slope to the number of years used to measure the child's income. In Panel A, we estimate the log-log IGE (excluding children with zero income), varying the age at which child income is measured. We restrict the sample to parents with income between the 10th and 90th percentile when estimating the IGE, as shown in Figure Ib. We measure child income in 2011-12 and analyze how the IGE varies across birth cohorts, as in Figure IIIa; see notes to that figure for further details. In Panel B, each point shows the slope coefficient from an OLS regression of child income rank on parent income rank (as in Figure IIa), using the core sample and varying the age at which parent income rank is measured. The first point measures parent income in 1996, when the mean age of parents is 41. The last point measures income in 2010, when parents are 55. Panel C reproduces Appendix Figure 2b from Chetty et al. (2014). In this figure, each point shows the slope coefficient from an OLS regression of an indicator for the child attending college at age 19 on parent income rank (similar to Figure IVa), varying the year in which parent income rank is measured from 1996 to 2011. In this series, we use data from the 1993 birth cohort, which allows us to analyze parent income starting when children are 3 years old in 1996. We list the age of the child on the x axis to evaluate whether the gradient differs when children are young (although parent age is of course also rising in lockstep). In Panel D, each point shows the slope coefficient from the same rank-rank regression as in Panel B using the core sample, but here we always use a five-year (1996-2000) mean to measure parent income and vary the number of years used to compute mean child income. The point for one year measures child income in 2012 only. The point for two years uses mean child income in 2011-12. We continue adding data for prior years; the 6th point uses mean income in years 2007-2012.

APPENDIX FIGURE IV: Dollar-Weighted vs. Traditional IGE Estimates

A. Log of Mean Child Income vs. Mean of Log Child Income

Baseline IGE = 0.344
(0.0004)
Baseline IGE P10-P90 = 0.452
(0.0007)

Dollar-Weighted IGE= 0.335
(0.008)
Dollar-Weighted IGE P10-P90 = 0.414
(0.004)

Log of Mean Child Income (circles)
Mean of Log Child Income (triangles)
Log Mean Parent Income

⬤ Log of Mean Child Income ▲ Mean of Log Child Income

B. Dollar-Weighted IGE by Age of Child Income Measurement

Dollar-Weighted IGE
Age at which Child's Income is Measured

Notes: This figure compares dollar-weighted (Mitnik et al. 2014) and traditional IGE estimates. Panel A is based on the core sample (1980-82 birth cohorts) and baseline family income definitions for parents and children. The series in circles (left axis) plots log of mean child income against log of mean parent income. The series is constructed by taking the logs of the points in Figure Ia; however, here we do not omit the top income bin. The slope coefficients, which correspond to the dollar-weighted IGE defined in Appendix C, and standard errors are estimated by OLS on the binned data. The series in triangles (right axis) reports the mean of log child income vs. the mean of log parent income (reproducing the series in Figure Ib). The slope coefficients and standard errors for the traditional IGE are estimated on the microdata. The dashed lines in Panel A show the 10th and 90th percentiles of the parent income distribution. Panel B shows how the dollar-weighted IGE varies with the age at which child income is measured. We estimate the dollar-weighted IGE by grouping parents into 100 bins based on their income rank and regressing the log of mean child income on the log of mean parent income across the 100 bins. The figure plots the slope from this regression vs. the age at which child income is measured. We measure child income in 2011-12 and analyze how the IGE varies across birth cohorts, as in Figure IIIa; see notes to that figure for further details. The first point corresponds to the children in the 1990 birth cohort, who are 21-22 when their incomes are measured in 2011-12 (denoted by age 22 on the figure). The last point corresponds to the 1980 cohort, who are 31-32 (denoted by 32) when their incomes are measured.

APPENDIX FIGURE V: Absolute Upward Mobility Controlling for Cost of Living
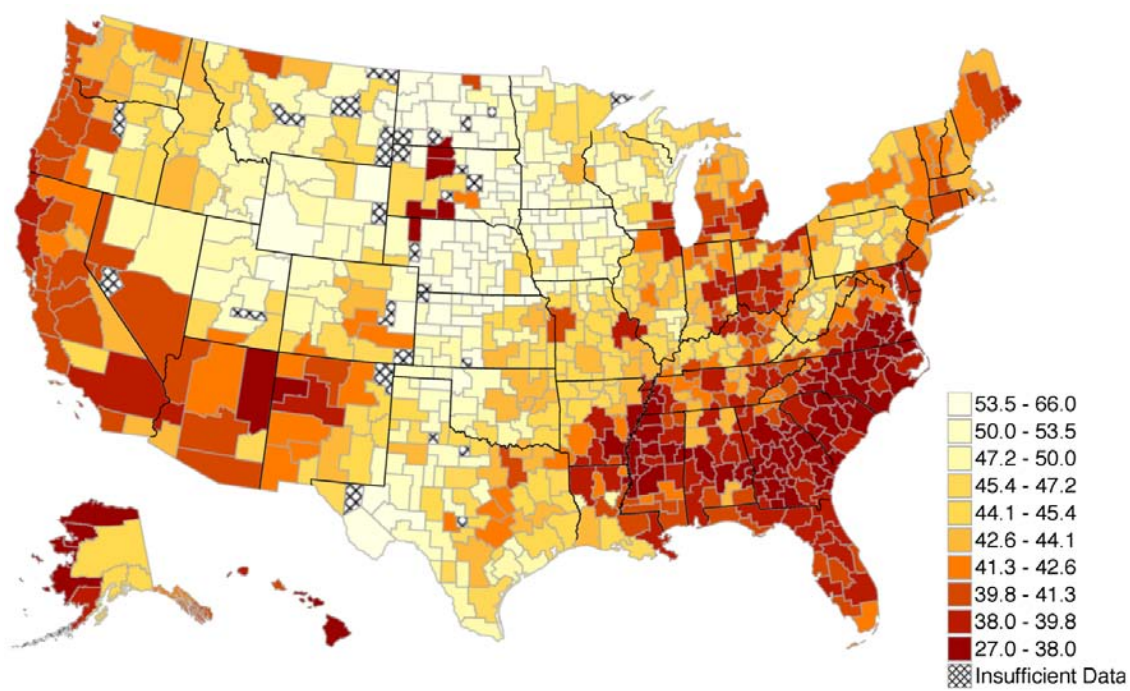
Corr. with baseline $\bar{r}_{25}$ = 0.98 (unweighted), 0.86 (pop-weighted)

Notes: This figure maps the spatial distribution of upward mobility across CZs adjusting income for cost of living. To compute cost of living adjusted upward mobility we deflate nominal family income (for both children and parents) by a cost of living index computed from the ACCRA index. The map is very close to the baseline Absolute Upward Mobility map from Figure VIIa, showing that our mobility measures are robust to using real income or nominal income.

APPENDIX FIGURE VI: Probability of Reaching Top Quintile Starting from Bottom Quintile

Corr. with baseline $\bar{r}_{25}$ = 0.91 (unweighted), 0.92 (pop-weighted)
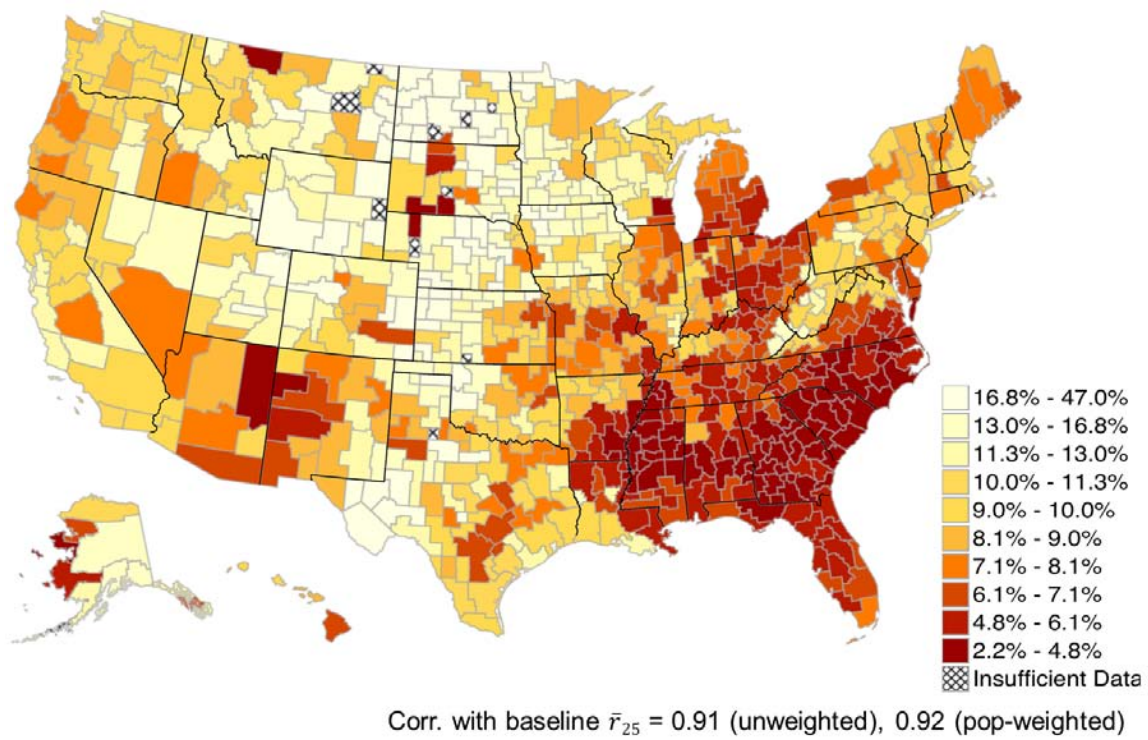
Notes: This figure maps the spatial distribution of the probability that a child reaches the top quintile of the income distribution conditional on having parents in the bottom quintile of the income distribution. The map is very close to the baseline Absolute Upward Mobility map from Figure VIIa, showing that our mobility measures are robust to using real income or nominal income.

APPENDIX FIGURE VII: Fraction of Children Above Poverty Line Given Parents at 25th Percentile

Corr. with baseline $\bar{r}_{25}$ = 0.94 (unweighted), 0.89 (pop-weighted)
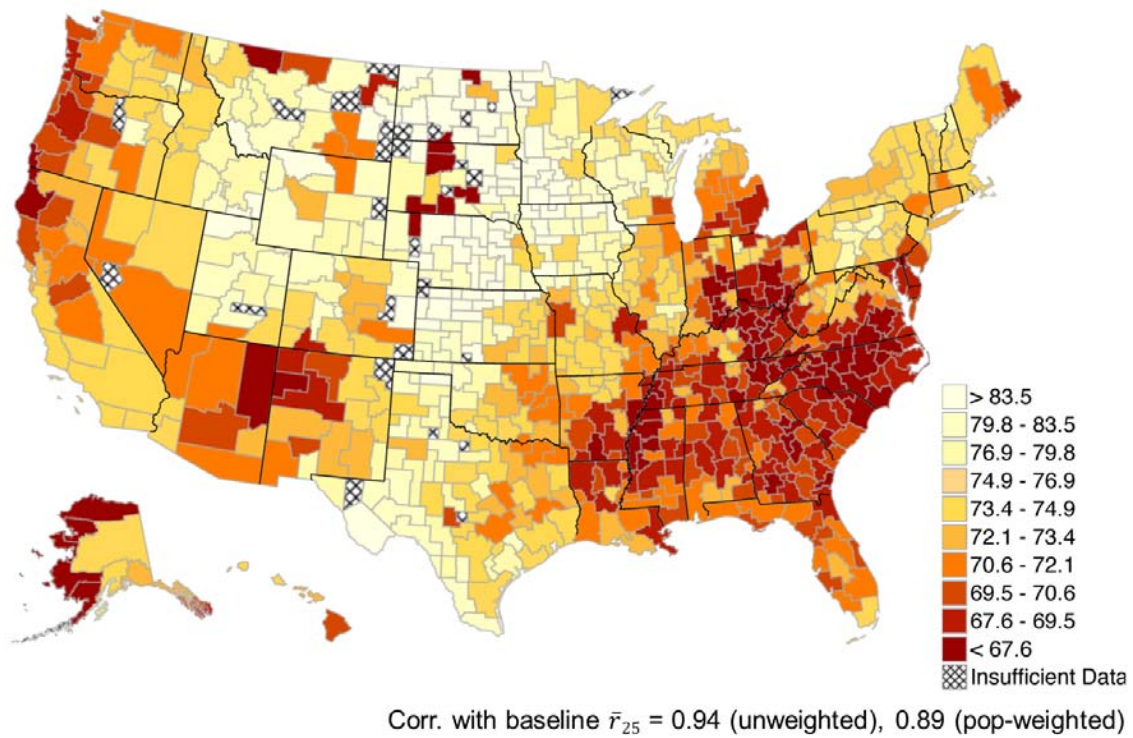
Notes: This figure maps the fitted values at parent rank 25 from a regression of an indicator for child family income being above the poverty line on parent income rank. The map is constructed by grouping CZs into ten deciles and shading the areas so that lighter colors correspond to higher mobility. Areas with fewer that 250 children in the core sample, for which we have inadequate data to estimate mobility, are shaded with the cross-hatch pattern. We report the unweighted and population-weighted correlation coefficient across CZs between these mobility measures and the baseline measure in Figure VIIa.
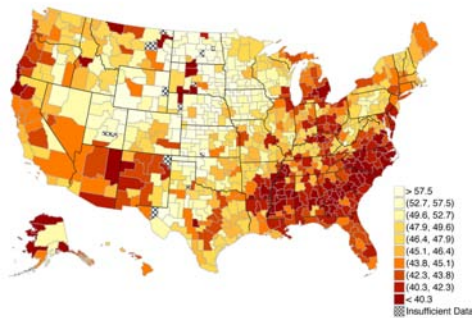
APPENDIX FIGURE VIII: Absolute Upward Mobility at Age 26 for Children of Stayers
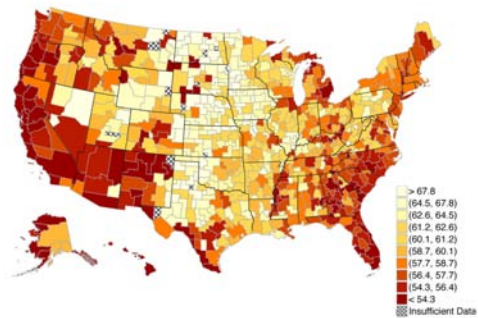
A. Mean Child Rank by Parent Rank in Chicago
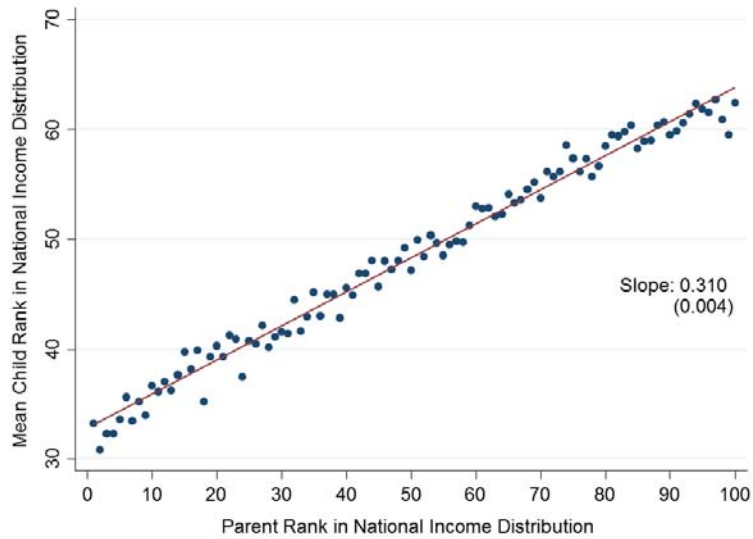
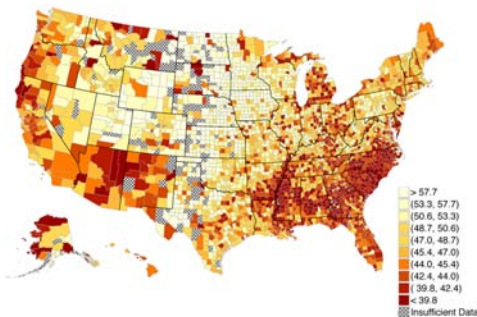
B. Absolute Upward Mobility Based on Stayers from 1996-2012

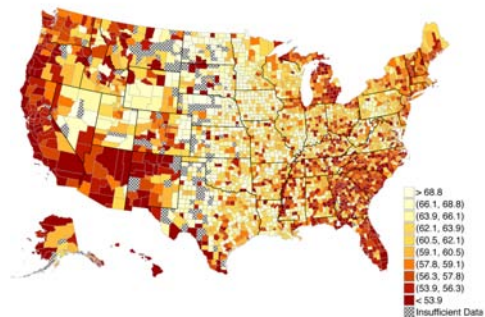**B.1. For Children with Parent at the 25th Percentile (CZ)**

B.2. For Children with Parent at the 75th Percentile (CZ)

B.3. For Children with Parent at the 25th Percentile (County)

B.4. For Children with Parent at the 75th Percentile (County)

Notes: Panel A replicates the series in triangles in Figure VI, but measures income ranks at age 26 and restricts the sample to children whose parents stayed in the same CZ from 1996-2012. Panel B replicates figure VIIa at the CZ and County level and at the 25th and 75th percentile, but estimates absolute upward mobility using income ranks at age 26 and restricts the sample to children whose parents stayed in the same CZ from 1996-2012.