Inequality of Educational Opportunity? Schools as Mediators of the Intergenerational Transmission of Income

Jesse Rothstein, *University of California, Berkeley*

Intergenerational income transmission varies across commuting zones (CZs). I investigate whether children's educational outcomes help to explain this variation. Differences among CZs in the relationship between parental income and children's human capital explain only one-ninth of the variation in income transmission. A similar share is explained by differences in the return to human capital. One-third reflects earnings differences not mediated by human capital, and 40% reflects differences in marriage patterns. Intergenerational mobility appears to reflect job networks and the structure of local labor and marriage markets more than it does the education system.

I. Introduction

Chetty et al. (2014a) use data on the universe of US tax filers to measure intergenerational income transmission—the strength of the association between parents' and children's incomes—at the fine geographic level and re-

I thank Audrey Tiew, Leah Shiferaw, Julien Lafortune, and Rachel Young for excellent research assistance as well as Charlie Brown, Avi Feller, Pat Kline, and conference and seminar participants at the National Bureau of Economic Research; New York University; the Society of Labor Economists; University of California, Irvine; University of California, Riverside; University of California, Santa Barbara; University of California, San Diego; Université du Québec à Montréal; and the Trans-Pacific Labor Seminar for helpful discussions and comments. I am grateful to the Russell Sage Foundation for financial support. Contact the author at rothstein@berkeley.edu. In-

[*Journal of Labor Economics*, 2019, vol. 37, no. S1]

0734-306X/2019/37S1-0010$10.00
Submitted April 19, 2017; Accepted September 14, 2018

veal massive heterogeneity across space: the gap in adult earnings between children from high- versus low-income families is nearly twice as large for children who grow up in Cincinnati as for those who grow up in Los Angeles. Little is known, however, about the mechanisms driving this variation.

There are many potential channels for intergenerational transmission, including differences in parenting practices between high- and low-income families, differences in explicit investments in children's education, differences in access to educational or other public institutions, and labor market institutions (such as insider hiring or spatial mismatch) that advantage children from high-income families regardless of their skills. These suggest quite different directions to look for potential policy interventions aimed at reducing transmission.

Chetty et al. (2014a) find that income transmission is negatively correlated with average test scores, high school completion, and school expenditures and is uncorrelated with average class size. This is suggestive, but these aggregate correlations are of limited value to understanding the mechanisms driving the variation across commuting zones (CZs). Income transmission is about differential outcomes of children from high- and low-income families, so to understand the mechanisms driving the between-area variation we need to understand how areas vary in the relative inputs received by and outcomes obtained by children from families with different incomes.

I investigate this directly, focusing on human capital accumulation as a potential mediator of intergenerational income transmission. If variation in school quality or in parenting practices related to learning is an important factor driving the variation in income transmission, we would expect that high income transmission areas would also be areas where the gap in educational outcomes between children from high- and low-income families is large. On the other hand, if parental income primarily helps children by, for example, buying them access to better labor market networks, then areas where poor children have good adult outcomes will not, in general, be areas where they do relatively well in school.

To measure children's educational outcomes, I rely primarily on the Education Longitudinal Study (ELS). This is a representative national sample that includes information about parental income, children's achievement (test scores) near the end of high school, and educational attainment and early adult earnings and income. The ELS data can be geocoded to CZs, the unit of geography considered by Chetty et al. (2014a).

The ELS contains only about 15,000 respondents, far too few to produce reliable income-achievement transmission measures for each of the 700 CZs in the United States. I show that this is not necessary in order to accomplish

formation concerning access to the data used in this paper is available as supplemental material online.

the more limited goal of measuring the across-CZ association between income-income transmission and income-achievement transmission. That association is identified even with small numbers of observations from each CZ—information can be pooled from many CZs with similar income-income transmission to identify the average income-achievement transmission among them even when the latter is not reliably estimated for any individual CZ. I develop an estimator for the slope of the CZ-level regression of intergenerational income transmission on income-achievement transmission—or for the correlation between the two—based on a mixed (random coefficients) model for the relationship between children's achievement and their parents' incomes.

I find that CZs vary substantially in the strength of transmission from parental income to children's twelfth-grade math scores (which I call "test score transmission") but that this is only weakly correlated across CZs with income transmission. Income transmission is more strongly correlated with the strength of transmission from parental income to children's educational attainment,[1] although the magnitude of the variation in the latter is not large enough to account for a large share of the former. These results suggest that differential inequities in access to good schools are not an important mechanism driving the across-CZ variation in income transmission; what role education does play seems to reflect access to higher education more than the quality of elementary and secondary schools.

This motivates me to consider other channels for intergenerational income transmission that may vary across CZs. One is the labor market return to skill. In every CZ, children from low-income families obtain less human capital than do children from higher-income families. As a result, differences in the return to human capital could drive differences in income transmission even if the skill acquisition technology were the same everywhere. Indeed, I find that the return to education varies substantially across CZs and explains as much of the variation in income transmission as do achievement and attainment gradients together. This points to labor market institutions as a potentially important factor.

I develop a decomposition that allows me to apportion the variation in CZ-level income transmission into four components: accumulation of human capital, the earnings returns to human capital, the residual component of earnings that is explained by parental income conditional on the child's measured human capital, and nonearnings components of the child's family

[1] This reproduces a Chetty et al. (2014a) result for college enrollment, discussed below. Another similar result comes from Kearney and Levine (2016), who find that high school dropout gaps by family status are stronger in more unequal states (which tend to have stronger income transmission). Kearney and Levine (2014) find that nonmarital childbearing is more common among women of low socioeconomic status in these states as well.

income (including spousal earnings and any nonlabor income). I find that the final component, spousal and unearned income, accounts for two-fifths of the relative advantage of children from high-income families in high-transmission CZs. This reflects differences in the likelihood of marriage or in the age of marriage rather than assortative matching or inheritances. Another one-third operates through children's residual earnings. Skill accumulation and the return to skill each represents only one-ninth of the total.

My analysis is purely observational; my estimates of the association between CZ-level income transmission and CZ-level transmission of parental income to children's test scores and other outcomes could be confounded by other CZ-level characteristics that are correlated with both. Keeping this caveat in mind, my results indicate that human capital plays a relatively small role in the geographic variation in the intergenerational transmission of income. Much of this variation appears to reflect differences in adult earnings of children with similar skills, perhaps due to labor market institutions (e.g., unions or other determinants of residual income inequality) or differences in access to good jobs (due, perhaps, to labor market networks or socially stratified labor markets). An even larger component is due to the use of family income–based (rather than individual earnings–based) measures of income transmission. This may be spurious, as differences in the likelihood of having spousal earnings, across income levels and across CZs, may simply reflect variation in age at marriage rather than true differences in opportunity across CZs and may not be indicative of children's economic success.

My results on the mechanisms driving the existing variation in income transmission do not translate directly into policy implications. It may be possible to increase opportunity through educational interventions even though education is not a primary channel explaining differences in current opportunity. Nevertheless, my results suggest that the space of policies worthy of consideration should be broader than this. Policies related to labor market opportunity and outcomes may be more important and merit at least as much attention.

II. Data

My analysis combines two sources of information: measured income transmission at the CZ level, from the analysis of tax data by Chetty et al. (2014a), and survey data with information about parental income and children's human capital attainment. I discuss these in turn.

A. Intergenerational Income Transmission

Chetty et al. (2014a) discuss several ways of defining intergenerational mobility. I focus on what they call "relative mobility," the advantage that a child from a high-income family has relative to a child from a low-income family in the same CZ in achieving a high income as an adult. Chetty and colleagues study children born between 1980 and 1982 and measure the income of child i in CZ c, y_{ic}, as the average family income, including any spousal earnings and nonlabor income, in 2011 and 2012, when the child is between 29 and 32. Children are linked to parents who claimed them as dependents in their late teens, and Chetty and colleagues define p_{ic} as the average parental family income (measured as the adjusted gross income plus tax-exempt interest and nontaxable Social Security benefits) for the parent(s) of child i in 1996–2000. Both children's and parents' incomes are scaled as national percentile ranks in the relevant distributions, without adjustment for family size or the number of earners. Children are assigned to the CZ in which their parents filed taxes in 1996, when the children were 14–16 years old.

Chetty and colleagues define relative mobility in CZ c as the coefficient θ_c from a bivariate regression:

$$y_{ic} = \alpha_c + p_{ic}\theta_c + e_{ic}. \tag{1}$$

Higher values of θ_c correspond to less mobility across generations, and I refer hereafter to θ_c as the strength of income transmission in the CZ. Universe tax data allow Chetty and colleagues to estimate θ_c extremely precisely at the CZ level. They find that $\theta_c = 0.43$ in Cincinnati, meaning that a 1 percentile difference in parental income is associated with a 0.43 percentile difference in children's eventual income on average in that city, and that in Los Angeles $\theta_c = 0.23$, implying a relationship between parent and child income that is only a bit more than half as strong as in Cincinnati.

Table 1 presents unweighted summary statistics for θ_c, extracted from the online data tables of Chetty et al. (2014a). The average of 0.33 indicates that in the average CZ, each 1 percentile increase in parental income is associated with one-third of a percentile increase in children's income. But the standard deviation of θ_c across CZs is 0.065. In 71 CZs, θ_c is less than 0.24, indicating parental income–child income relationships about one-quarter weaker than the average, while another 78 CZs have $\theta_c > 0.40$, about one-quarter larger than average. Among the 100 largest CZs, Santa Barbara has the weakest income transmission, and Cincinnati has the strongest.

B. Survey Data

To measure the transmission of parental income to children's human capital accumulation, I use the ELS (Ingels et al. 2014a). This is a nationally representative longitudinal sample of just over 19,000 tenth graders in 2002, corresponding roughly to the 1985–86 birth cohorts. Respondents were followed through 2012, when they were roughly 26. Math and reading scores are available in tenth grade, and math scores are available in twelfth grade. I also construct children's adult income, y_{ic}, as their self-reported 2011 family

Table 1
Summary Statistics for Commuting Zone (CZ)–Level Relative Mobility (Income Transmission)

	Full Sample		100 Largest CZs	
	(1)	(2)	(3)	(4)
N	709		100	
Mean	.325		.338	
Standard deviation	.065		.054	
Minimum	.068	Linton, ND	.215	Santa Barbara, CA
10th percentile	.240	Hutchinson, MN	.257	Bakersfield, CA
25th percentile	.280	Carroll, IA	.298	Manchester, NH
50th percentile	.330	Eagle Butte, SD	.348	Des Moines, IA
75th percentile	.373	Roanoke, VA	.382	Greenville, SC
90th percentile	.404	Vicksburg, MS	.398	Indianapolis, IN
Maximum	.508	Lake Providence, LA	.429	Cincinnati, OH
Correlations:				
1. Relative mobility for 1983–85 birth cohorts	.84		.98	
2. Causal mobility measure from Chetty and Hendren (2018)	.85		.91	
3. Relative mobility for college enrollment	.68		.70	

NOTE.—Statistics are computed at the CZ level, without weights, and pertain to the preferred "relative mobility" measure from Chetty et al. (2014a). Correlation 1 is with the relative mobility measure for the 1983–85 birth cohorts, from Chetty et al. (2014a). Correlation 2 is with the causal mobility measure from Chetty and Hendren (2018). Correlation 3 is with the CZ-level income–college enrollment transmission (the slope of college enrollment between 18 and 21 with respect to parental income percentile) for the 1980–82 birth cohorts, from Chetty et al. (2014a).

income (including spousal earnings and nonlabor income) when they were 25 or 26 years old. I assign students to CZs on the basis of their residential zip codes in the base-year survey, using information from subsequent surveys when this is missing.

I supplement the ELS with two similar panels. The Early Childhood Longitudinal Study, Kindergarten Cohort (ECLS-K; Tourangeau et al. 2009), sampled kindergartners in 1998–99 and followed them through eighth grade in 2007. Students are assigned to CZs on the basis of their eighth-grade residences.[2] The High School Longitudinal Study (HSLS; Ingels et al. 2014b) provides high school test scores for children born in roughly 1994–95. I assign them to the CZs they lived in during high school.

There are three major limitations of these samples for my purposes. Most importantly, sample sizes are well under 100 per CZ. Moreover, the surveys each use multistage sampling designs, with schools as one stage and then relatively large samples of students within each school.[3] This means that within-

[2] Where eighth-grade residences are unavailable, I use the location of the eighth-grade school, then the fifth-grade residence and school, then third grade, and so on.

[3] The regressions below account for CZ-level (or within-CZ) clustering but do not otherwise adjust for the survey designs. Most of my estimates are unweighted,

CZ heterogeneity is even more limited than the small sample sizes imply. A consequence is that it is necessary to pool information across CZs in order to obtain any precision at all about the relationship between parental income and later outcomes (Gelman and Hill 2006).

Second, none of the samples covers the 1980–82 birth cohorts analyzed by Chetty et al. (2014a). If income or test score transmission changed across cohorts, between-cohort comparisons may understate the relationship between the two. I explore sensitivity to this misalignment in two ways. First, I compare across the NCES samples, which differ in their distance from the cohorts of Chetty et al. (2014a). Results are quite similar, suggesting that between-cohort changes are not particularly important for my analysis.[4] Second, the appendix (available online; contains figs. A1–A3 and tables A1–A8) presents results that use two alternative measures of income transmission, one from Chetty et al. (2014a) for the 1983–85 birth cohorts—very close to the ELS cohorts—and one from Chetty and Hendren's (2018) mobility-based estimate based on children born between 1980 and 1991. None of the results presented here differ meaningfully when either alternative is used.

A final important limitation is that while all three studies included parental surveys, the parental income measures are extremely limited. The ELS collects only total family income, and only in the base year. The measure is categorical, with 13 bins (e.g., one corresponds to incomes between $25,000 and $35,000). The HSLS collected family income twice and reports it continuously. The ECLS data include a continuous measure in kindergarten and categorical measures in first and third grades.[5]

I also present an analysis of returns to education in American Community Survey (ACS) data. For maximum comparability with the measures of Chetty et al. (2014a), I use the 2010, 2011, and 2012 one-year public-use microdata samples and focus on the 253,852 individuals in these samples born between 1980 and 1982. I convert annual family incomes to percentiles within the ACS sample distribution. I do not have information about where

but results are generally robust to using student-level sampling weights when specifications permit it.

[4] Chetty et al. (2014b) find that national aggregate relative mobility has been quite stable across a range of birth cohorts (born 1971–93), but CZ-level measures might in principle vary across cohorts with little variation in the national aggregate. See also Aaronson and Mazumder (2008).

[5] I can assess the reliability of individual binned measures by comparing the same family's income across the three ECLS waves. As discussed below, I scale incomes as percentiles of the sample distribution. Percentiles constructed from the first- and third-grade wave bin midpoints are correlated 0.84; the kindergarten percentile is correlated 0.86 and 0.80 with the first- and third-grade measures, respectively. A percentile constructed from the average of the three is correlated 0.94–0.95 with the individual measures.

Table 2
Summary Statistics for Individual-Level Samples

	ELS (1)	ECLS (2)	HSLS (3)	ACS (4)
Birth year	1985–86	1992–93	1994–95	1980–82
Number of observations	15,240	19,940	21,440	330,366
Number of CZs	312	365	295	488
Female	.50	.48	.50	.50
Black	.14	.18	.17	.14
Hispanic	.16	.19	.22	.21
Asian	.04	.03	.03	.06
Other nonwhite	.05	.02	.08	.09
Parental income	61,417	51,789	70,464	
	(50,312)	(47,419)	(56,034)	
Test scores available for grades	10, 12	K, 1, 2, 3, 5, 8	9, 11	NA
	Age 26			Age 28–32
Post–high school outcomes:				
Any college	.84			.53
College completion (BA degree)	.33			.22
Years of education	14.0			13.3
	(1.8)			(2.8)
Marital status	.28			.47
Presence of working spouse	.24			.40
Earnings	25,451			29,508
	(24,672)			(32,477)
Family income	36,095			73,039
	(35,238)			(62,890)

NOTE.—Sample sizes and demographics are computed for the base-year sample for each survey and use sampling weights. Sample sizes in cols. 1–3 are rounded to the nearest 10. Standard deviations are in parentheses. ACS = American Community Survey; ECLS = Early Childhood Longitudinal Study; ELS = Education Longitudinal Study; HSLS = High School Longitudinal Study; NA = not available.

respondents lived as children, so I assign them to the CZ in which they lived at the time of the survey.

Summary statistics for the four microdata samples are reported in table 2.[6] Following Chetty et al. (2014a), my primary analysis converts incomes, earnings, and test scores to percentiles within the relevant samples; these have mean 50.0 and standard deviation 28.9 by construction. For parental income, I assign ELS categories to the midpoints of the national percentile range they span; in the HSLS and ECLS, I average incomes across the available waves, using bin midpoints as necessary, and construct percentiles of the distribution of averages. There is surely nonclassical measurement error

[6] Mean parental incomes vary across samples. This in part reflects inflation (I report nominal values) and life-cycle considerations (ECLS parents are on average younger when their incomes are collected than ELS or HSLS parents).

in each of the measures when scaled as percentiles.[7] There is no reason to expect the resulting bias to differ across CZs or across dependent variables, however. Figure A1 shows that percentile-percentile relationships between parental income and various children's outcomes are roughly linear in the ELS sample. Table A7 presents results that use alternative scales for parental income and child test scores.

Table 2 shows that 84% of ELS respondents report ever having attended a postsecondary educational institution by the age 26 survey. This is much higher than the 53% reporting some college or more in the ACS sample. In the tax data of Chetty et al. (2014a), 60% of children are recorded as attending college between 18 and 21. Only half of the college attenders in the ELS sample earned a college degree; in the ACS, the share without a degree is under 30%. It appears that the ELS is counting some respondents with weak postsecondary attachment as having attended college who would not be so counted in other data sets. I discuss this further below.

III. Conceptual Framework and Empirical Strategy

Section III.A lays out a simple model in which human capital mediates the relationship between parent and child income. Section III.B develops a methodology for estimating the key elements of the model—and in particular how they vary across CZs—with the limited available data. Section III.C describes a decomposition of the across-CZ variation in income transmission into various component mechanisms.

A. Children's Human Capital as a Mediator of Intergenerational Income Transmission

Let s_{ic} be a potential mediator of the relationship between parental income p_{ic} and children's income y_{ic}, such as the child's educational attainment or achievement. Its importance as a mediator depends on the strength of its relationship to p_{ic} and on the extent to which it accounts for the relationship between p_{ic} and y_{ic}. Assuming linear and additive relationships, we have a simple system:

$$s_{ic} = \alpha + p_{ic}\pi + u_{ic} \tag{2}$$

and

$$y_{ic} = \kappa + s_{ic}\lambda + p_{ic}\mu + v_{ic}. \tag{3}$$

[7] The ELS test scores—in math and reading in tenth grade and in math in twelfth grade—are point estimates of student proficiency from an item response theory (IRT) model. Measurement error does not bias student performance on the original IRT scale but will tend to compress gaps between groups on the percentile scale (Jacob and Rothstein 2016). This likely attenuates my estimates of income-to-achievement transmission but should not bias the between-CZ comparisons that are my primary interest.

These equations represent reduced-form transmission and are statistical projections, not causal models. The system is illustrated in figure 1. The coefficient π represents the importance of parental income as a determinant of s_{ic}; λ is the return to s_{ic} in children's incomes, while μ is the "direct" effect of parental income on children's income, not mediated by s_{ic}.

Plugging equation (2) into equation (3) and rearranging, we obtain

$$y_{ic} = (\kappa + \alpha\lambda) + p_{ic}(\pi\lambda + \mu) + (u_{ic}\lambda + v_{ic}). \tag{4}$$

Thus, the reduced-form transmission of parental income to child's income studied by Chetty et al. (2014a) is, in this framework, the sum of a component operating through the potential mediator and the direct effect:

$$\theta \equiv \pi\lambda + \mu. \tag{5}$$

Each of the coefficients in equations (2) and (3) may vary across CZs, and all may depend in important ways on local institutions and policies. For example, when s_{ic} is a human capital measure, a CZ with a bad system of public education, in which only those who can afford private school tuition (or, perhaps, a house in one of the few excellent public school districts) can obtain good schools for their children, would have a high π. A very unequal CZ labor market, with high wages for those with high human capital but few opportunities for those with little, would yield a high λ. Finally, a labor market in which strong family networks are needed to access good job opportunities would imply high μ. Thus, understanding which of these accounts—even descriptively—for the between-CZ variation in θ would in-

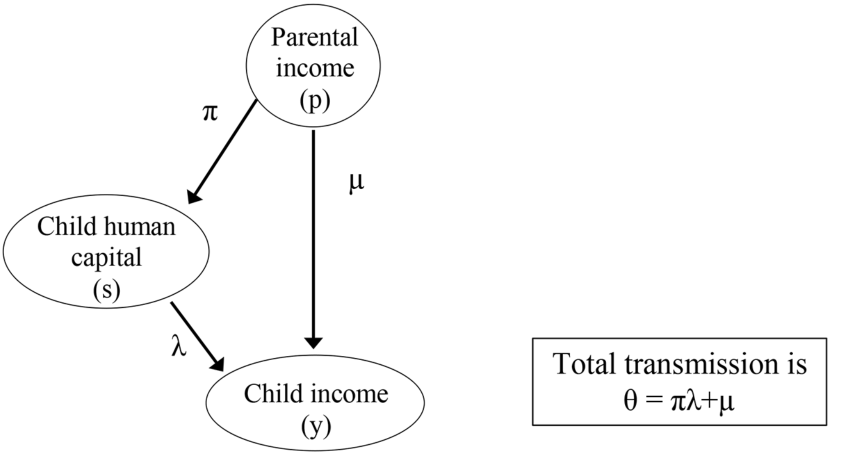

FIG. 1.—Path diagram of the transmission of parental income to child income, mediated by the child's human capital.

form further investigation of the determinants of intergenerational transmission.

B. Estimating the Importance of Test Score Transmission as a Mediator Using Across-CZ Variation

A barrier to estimating equations (2) and (3) is that few data sets contain all of the necessary measures. Tax data provide p_{ic} and y_{ic} for enormous samples but only one potential mediator, whether the child enrolls in college by age 21. Using this as s_{ic}, Chetty et al. (2014a) estimate π_c for each CZ and find that it is highly correlated with θ_c ($\rho = 0.68$). However, as I discuss below, for plausible λ the variation in π_c is too small to account for more than a small share of the variation in θ_c.

The ELS provides richer measures of human capital but for a sample that is far too small to permit reliable estimation of π_c for each CZ. In lieu of this, I develop a methodology for estimating the across-CZ bivariate regression of θ_c on π_c even when CZ-level samples are too small to permit estimation of π_c directly.[8] The slope of this regression can be interpreted as an (observational) estimate of λ. The R^2 provides an initial estimate of the share of the variance of θ_c that is attributable to the mediating role of human capital accumulation. In Section III.C, I develop a more careful decomposition.

My approach is built from the "reverse" projection of π_c onto θ_c:

$$\pi_c = \gamma + \theta_c\beta + \eta_c, \tag{6}$$

where $\beta = \mathrm{cov}(\theta_c, \pi_c)/\sigma_\theta^2$ is the across-CZ linear projection coefficient and η_c is orthogonal to θ_c by construction. If the terms of equation (6), including the residual variance σ_η^2, were known, it would be straightforward to obtain the forward regression of θ_c on π_c,

$$\frac{\mathrm{cov}(\theta_c, \pi_c)}{\sigma_\pi^2} = \frac{\mathrm{cov}(\theta_c, \pi_c)}{\sigma_\theta^2}\frac{\sigma_\theta^2}{\sigma_\pi^2} = \beta\frac{\sigma_\theta^2}{\sigma_\theta^2\beta^2 + \sigma_\eta^2}; \tag{7}$$

the R^2 (which is the same for the forward and reverse regressions),

$$R^2 = \beta^2\frac{\sigma_\theta^2}{\sigma_\pi^2} = \beta^2\frac{\sigma_\theta^2}{\sigma_\theta^2\beta^2 + \sigma_\eta^2}; \tag{8}$$

and the correlation between θ_c and π_c, $\mathrm{corr}(\theta_c, \pi_c) = (R^2)^{1/2}$. Note that θ_c and therefore σ_θ^2 are observed directly in Chetty et al.'s (2014a) computations from population data.

[8] I model s_{ic} as a mediator of y_{ic} and thus π_c as a mediator of θ_c. However, reverse causality from income transmission, θ_c, to human capital transmission, π_c, is possible. For example, it may be easier to attract high-ability college graduates into teaching in CZs with more equal labor markets, creating a causal path from economic mobility to gaps in children's outcomes.

To estimate equation (6), I return to equation (2), making explicit the variation in the coefficients across CZs:

$$s_{ic} = \alpha_c + p_{ic}\pi_c + u_{ic}. \tag{9}$$

Substituting equation (6) into equation (9), we obtain

$$s_{ic} = \alpha_c + p_{ic}(\gamma + \theta_c\beta + \eta_c) + u_{ic}. \tag{10}$$

Equation (10) is the basis of my analysis. It leads to four types of specifications. First, a simple national regression of s_{ic} on p_{ic} restricts $\beta = 0$ and $\sigma^2_\eta = 0$.

Second, I divide the nation into deciles of the distribution of θ_c and regress s_{ic} on p_{ic} separately in each decile d. By regressing the coefficients $\hat{\gamma}_d$ from these regressions on average income transmission in each decile, $\bar{\theta}_d$, I can test whether $\beta = 0$.[9]

Third, returning to pooled data, I can add to the national regression an interaction of p_{ic} with θ_c. This interaction coefficient estimates β in one step:

$$s_{ic} = \alpha_c + p_{ic}\gamma + (p_{ic}\theta_c)\beta + e_{ic}. \tag{11}$$

The error term here is $e_{ic} \equiv p_{ic}\eta_c + u_{ic}$; because both equation (9) and equation (6) are linear projections, it is orthogonal to p_{ic} and $p_{ic}\theta_c$ by construction. To ensure that γ and β are identified from within-CZ variation, I divide p_{ic} into its between- and within-CZ components and include main effects and θ_c interactions for each. I also remove the grand mean of θ_c to permit comparisons of γ to the simpler specification and include a main effect for $(\theta_c - \bar{\theta})$. Defining $\tilde{p}_{ic} \equiv p_{ic} - \bar{p}_c$ and $\tilde{\theta}_c \equiv \theta_c - \bar{\theta}$, the regression is

$$s_{ic} = \alpha_c + \tilde{p}_{ic}\gamma + \tilde{p}_{ic}\tilde{\theta}_c\beta + X_c\Omega + e_{ic}, \tag{12}$$

where X_c includes $\bar{p}_c$, $\bar{p}_c(\theta_c - \bar{\theta})$, and $(\theta_c - \bar{\theta})$. This is equivalent to simply controlling for X_c in equation (10). I explore ordinary least squares (OLS), (correlated) random effects, and fixed effects specifications for α_c, in each case reporting standard errors that are clustered at the CZ level to handle the dependence of e_{ic}.

Finally, my primary estimates are based on the full model (10), without restrictions. It can be seen as a random coefficients model, also known as a mixed model, with fixed coefficients γ and β and random coefficients α_c and η_c. If we assume that (α_c, η_c) and u_{ic} are each normally distributed and independent and identically distributed, equation (10) can be estimated by

[9] With more data one could use smaller cells. In the limit, with the first-stage regression estimated in each CZ separately, it estimates π_c, and the second-stage regression of $\hat{\pi}_c$ on θ_c is eq. (6) and estimates β. I present this analysis in the appendix, but it is poorly behaved in the small ELS sample.

maximum likelihood.[10] This yields an estimate not just of β but also of σ_η^2, and so it can be used to compute the forward regression (7). As in the fixed-coefficient specification, I separate $\bar{p}_c$ from $p_{ic} - \bar{p}_c$ to ensure that only within-CZ variation identifies the coefficients of interest.[11] The identifying assumption (beyond normality) is that η_c is orthogonal to θ_c. Recalling that η_c is the residual in equation (6), this simply means that the mixed model identifies only the observational regression of π_c on θ_c (and vice versa) and does not solve the causal inference problem.

One way to validate this strategy is to use the child's income as the skill measure—that is, let $s_{ic} \equiv y_{ic}$. This makes transmission to skill, π_c, identical to transmission to child's income, θ_c, thus ensuring that $\beta = 1$ and $\sigma_\eta^2 = 0$ in equation (6). I implement a version of this in table A2. Results are encouraging, although not perfect. The estimate of σ_η^2 is nearly identically zero—a result that does not occur for any of the other outcomes I examine below. The estimated β coefficient, however, is attenuated by about one-third from what was expected. I attribute this to the fact that this exercise mixes two different data sets—θ_c is measured in tax data, while I measure y_{ic} and p_{ic} in the ELS. As discussed above, the ELS measure of p_{ic} is of lower quality than the tax measure, while the ELS y_{ic} is measured at a younger age.

A different mixed model can be used to estimate the relationship between θ_c and the return to skill in the local labor market, defined as the coefficient of a regression of incomes on human capital:

$$y_{ic} = \tilde{\kappa}_c + s_{ic}\tilde{\lambda}_c + \tilde{v}_{ic}. \tag{13}$$

The standard omitted variables formula can be used to relate this reduced-form coefficient to the transmission coefficients from the path diagram in figure 1:

$$\tilde{\lambda}_c = \lambda_c + \frac{\mathrm{cov}(s_{ic}, p_{ic})}{V(s_{ic})}\mu_c = \lambda_c + \frac{\sigma_p^2}{\sigma_s^2}\pi_c\mu_c. \tag{14}$$

[10] Gelman and Hill (2006) discuss the estimation of models like eq. (10), which are referred to variously as mixed, hierarchical, random coefficient, or multilevel models. There is no fully satisfactory way to handle sampling weights in these models. Accordingly, I estimate them without weights. In simpler models, estimates are very similar with and without weights, so this is not likely to dramatically affect my results. In economics, it is common to estimate models like eq. (10) in two stages, as in n. 9. This does not require a normality assumption on (α_c, η_c), but estimation of σ_η^2 requires distinguishing what portion of the across-CZ variation in $\hat{\pi}_c$ is due to sampling error. As noted above, this is poorly behaved in the ELS sample. The mixed model approach can achieve better precision by pooling information across CZs.

[11] The random coefficient is specified to apply only to $p_{ic} - \bar{p}_c$, so the random intercept is $\tilde{\alpha}_c = \alpha_c + \bar{p}_c\eta_c$. I do not restrict the correlation between $\tilde{\alpha}_c$ and η_c. In most specifications, the estimated correlation is quite close to 1, suggesting that σ_α^2 is small.

As above, my starting point is the hypothetical reverse regression of $\tilde{\lambda}_c$ on θ_c: $\tilde{\lambda}_c = \gamma^{\tilde{\lambda}} + \theta_c\beta^{\tilde{\lambda}} + \eta_c^{\tilde{\lambda}}$. Substituting this into equation (13) yields a mixed model similar to equation (10):

$$y_{ic} = \tilde{\kappa}_c + s_{ic}\left(\gamma^{\tilde{\lambda}} + \theta_c\beta^{\tilde{\lambda}} + \eta_c^{\tilde{\lambda}}\right) + \tilde{v}_{ic}. \tag{15}$$

The parameters of this model can again be used to compute the regression of θ_c on $\tilde{\lambda}_c$.[12]

C. Decomposing the Across-CZ Variation in Income Transmission

The mixed models yield separate estimates of the relationships between θ_c and π_c and between θ_c and $\tilde{\lambda}_c$. Also of interest is the decomposition of the across-CZ variation in θ_c into its component parts. In this subsection I outline a strategy to decompose the separate contributions of variation in π_c, λ_c, and μ_c.

My starting point is the "structural" equation (3), allowing for CZ-level heterogeneity in the coefficients:

$$y_{ic} = \kappa_c + s_{ic}\lambda_c + p_{ic}\mu_c + v_{ic}. \tag{16}$$

The gradient of this with respect to parental income in CZ c is

$$\frac{dy_{ic}}{dp_{ic}}|_c = \frac{ds_{ic}}{dp_{ic}}|_c\lambda_c + \mu_c. \tag{17}$$

Given the definitions of θ_c and π_c, this is simply the decomposition defined earlier:

$$\theta_c = \pi_c\lambda_c + \mu_c. \tag{18}$$

The across-CZ gradient of equation (17) with respect to θ_c is

$$\frac{d^2y_{ic}}{dp_{ic}d\theta_c} = \frac{d^2s_{ic}}{dp_{ic}d\theta_c}\lambda_c + \frac{ds_{ic}}{dp_{ic}}|_c\frac{d\lambda_c}{d\theta_c} + \frac{d\mu_c}{d\theta_c} \tag{19}$$

or, using equation (18),

$$1 = \frac{d\pi_c}{d\theta_c}\lambda_c + \pi_c\frac{d\lambda_c}{d\theta_c} + \frac{d\mu_c}{d\theta_c}. \tag{20}$$

Each of the terms on the right side of equation (20) is interpretable as reflecting a distinct component of income transmission. The first term represents differences between high- and low-θ_c CZs in human capital accumulation

[12] In principle, one could estimate the relationships of θ_c with λ_c and μ_c via a version of this strategy by including a control for p_{ic} in eq. (13), with its own fixed and random coefficients. I have explored this model, but it is quite poorly behaved. I do use a fixed-coefficient version of this model in the decomposition discussed in Sec. III.C.

gaps between high- and low-income families, scaled by the return to human capital. But for scaling by λ_c, this term is identified by the β coefficient of the "reverse" regression (6) discussed above in Section III.B. It would be large if high-θ_c CZs offer less equal school quality to children from different family backgrounds.

The second term reflects covariance of the CZ-level return to skill with CZ-level income transmission, scaled by $\pi_c \equiv (ds_{ic}/dp_{ic})|_c$. This term would be large if high-θ_c CZs have higher returns to skill, producing better outcomes for children from high-income families who tend to obtain higher skill. The third term reflects differences in the transmission of parental income to children's incomes holding skills constant. This might be large if high-θ_c CZs have segmented labor markets or employment networks that allow high-income parents to ensure good outcomes for their children regardless of the children's skills.

To decompose the variation across cities, I use fixed values $\bar{\lambda}$ and $\bar{\pi}$ to scale the first and second terms:

$$\frac{d^2 y_{ic}}{dp_{ic}d\theta_c} = \frac{d^2 s_{ic}}{dp_{ic}d\theta_c}\lambda_c + \pi_c\frac{d\lambda_c}{d\theta_c} + \frac{d\mu_c}{d\theta_c} \tag{21}$$

$$\approx \frac{d^2 s_{ic}}{dp_{ic}d\theta_c}\bar{\lambda} + \bar{\pi}\frac{d\lambda_c}{d\theta_c} + \frac{d\mu_c}{d\theta_c}. \tag{22}$$

This leads to a three-step method for the decomposition. First, I estimate $\bar{\lambda}$ and $\bar{\pi}$ via pooled regressions of y_{ic} on s_{ic} and of s_{ic} on p_{ic}, respectively. Second, I estimate $d^2 s_{ic}/dp_{ic}d\theta_c = d\pi_c/d\theta_c = \beta$ via a regression of s_{ic} on p_{ic}, θ_c, and their interaction as in equation (11) above. Third, I regress y_{ic} on s_{ic}, p_{ic}, θ_c, and the two interactions $s_{ic} \times \theta_c$ and $p_{ic} \times \theta_c$. The interaction coefficients estimate $d\lambda_c/d\theta_c$ and $d\mu_c/d\theta_c$, respectively. As earlier, I include in each step the CZ means of the individual-level variables as well as CZ random effects.

Next, consider the left side of expression (21). In principle, $(dy_{ic}/dp_{ic})|_c$ is identically equal to θ_c, and its derivative with respect to θ_c is therefore 1. In practice, I rely on Chetty et al.'s (2014a) estimates of θ_c from tax data and measures of y_{ic} and p_{ic} from the ELS. In this blended data set, the regression of y_{ic} on p_{ic}, θ_c, and $p_{ic}\theta_c$ has an interaction coefficient well below 1. This reflects measurement differences between the ELS sample and the tax data. I take the empirical interaction coefficient as the target of my decomposition.

A final issue is that Chetty et al. (2014a) define y_{ic} as family income, including any spousal earnings and nonlabor income. There are thus two channels for each of the elements of the decomposition. The return to human capital in the second term, for example, includes both labor market and spousal market returns. These point to different structural factors of the CZ as explanations. Moreover, spousal market components of transmission may be artifacts of the fact that I measure children's income at a single point in time, when children are around 25: CZs in which children from high-income families typically marry young will have higher measured income transmission than do CZs in which these children typically marry later, but they may not meaningfully differ in the extent of available opportunity. Therefore, I separate children's incomes into their own earnings w_{ic} and the remaining component, reflecting spousal earnings and nonlabor income: $y_{ic} = w_{ic} + (y_{ic} - w_{ic})$. I apply the above decomposition only to the children's earnings and consider the reduced-form transmission of parental income to children's spousal and nonlabor income as a separate mechanism.

IV. Results

I present results in four parts. First, to lay the groundwork, I present national estimates of the path diagram in figure 1, using the ELS sample. Second, I use the variation in these estimates across CZs to identify the relationship between income transmission θ_c and human capital transmission π_c. Third, I use ELS and ACS data to examine variation in the reduced-form return to skill $\tilde{\lambda}_c$ across CZs, again relating it to CZ-level income transmission. Finally, I implement the decomposition described in Section III.C, attributing variation across CZs in income transmission to marriage market factors, skill accumulation, returns to skill, and direct transmission of parental income to children's earnings.

A. National Estimates

Table 3 presents estimates of the national relationships between parental income, children's human capital, and children's incomes, using the ELS sample. All specifications are weighted and include CZ fixed effects.

Columns 1 and 2 show the reduced-form relationships between parental income and children's twelfth-grade math scores and educational attainment, respectively. Each percentile of parental income is associated with 0.35 percentiles of children's achievement and with 0.019 years of education.[13] Column 3 presents an analysis of reduced-form income transmission. Each percentile of parental income is associated with 0.16 percentiles of children's adult incomes.

Columns 4–6 show regressions of the child's adult family income on the two human capital measures, first separately and then together. Each math score percentile is associated with 0.24 percentiles in additional income, while each year of education is associated with 3.6 percentiles. Each coefficient falls, as expected, when both measures are included together. Finally,

[13] Bradbury et al. (2015) find that the parental income coefficient is largely invariant to the age at which children's test scores are measured. I look across a larger range and find that it grows somewhat with age (see table A1). Bradbury et al. (2015) also compare results across four English-speaking countries. This exercise is similar in spirit to my comparison across CZs.

Table 3
Income Transmission Mediation Analysis at the National Level (Education Longitudinal Study [ELS])

	Dependent Variable						
	Twelfth-Grade Math Score (Percentile)	Years of Education	Family Income (Percentile)				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Parents' income (percentile)	.35	.019	.16				.07
	(.01)	(.001)	(.01)				(.01)
Twelfth-grade math score (percentile)				.24		.18	.17
				(.01)		(.02)	(.02)
Years of education					3.59	1.76	1.61
					(.19)	(.24)	(.24)
N	13,650	13,250	11,510	9,980	11,510	9,980	9,980
R^2	.19	.15	.08	.11	.10	.12	.12

Note.—Parental income, child family income, and twelfth-grade math scores are measured in percentiles of the national distribution and range from 0 to 100. All regressions use ELS sample weights (for the wave 1 survey in col. 1 and for the wave 3 survey in cols. 2–7) and include commuting zone (CZ) fixed effects. Standard errors are clustered at the CZ level. Sample sizes are rounded to the nearest 10.

column 7 includes both parental income and child human capital controls. Here, the parental income coefficient represents μ in the path diagram. This is 0.07, less than half of what it was without human capital controls in column 3 but nevertheless highly significant. The test score and education coefficients, representing λ, are only slightly reduced from column 6. The role of human capital as a mediating factor can be computed by multiplying these coefficients by the corresponding π coefficients in columns 1 and 2. This yields $0.35 \times 0.18 + 0.019 \times 1.76 = 0.098$, or roughly 60% of the total transmission in column 3.

The national analysis thus indicates that human capital is an important mediating factor in intergenerational income transmission. As we will see, human capital plays a much smaller role in explaining the across-CZ variation.

B. Transmission from Parental Income to Children's Human Capital across CZs

In this subsection, I examine variation across CZs in the transmission of parental income to children's human capital. I examine test scores first, then educational attainment.

1. *Twelfth-Grade Math Scores in the ELS*

As a first effort to explore across-CZ variation, I divide CZs into deciles based on θ_c. For each decile, I estimate a separate regression of children's test scores on parental income, with CZ fixed effects. Figure 2 plots the parental

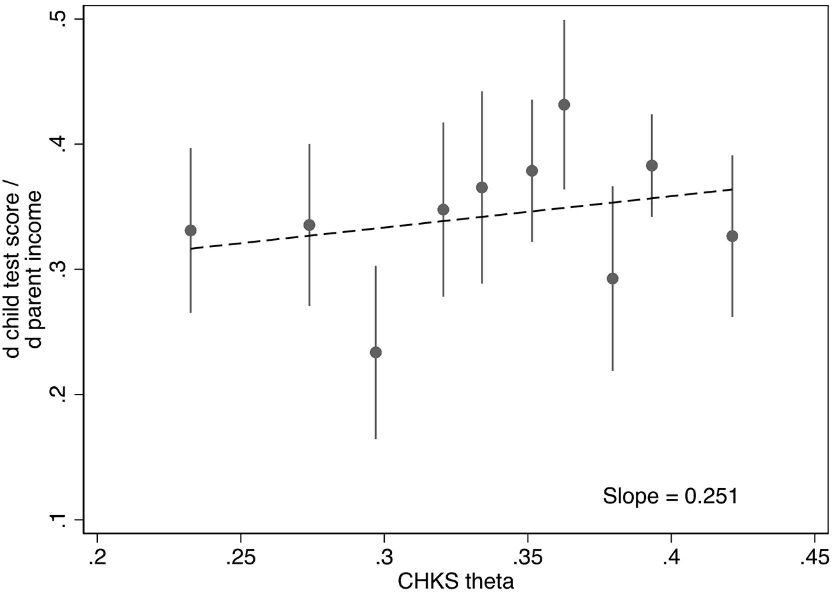

FIG. 2.—Parental income–child test score transmission, by commuting zone (CZ) income transmission (θ) decile. CZs are divided into deciles based on the income transmission (relative mobility) measure of Chetty et al. (2014a). The figure plots coefficients and 95% confidence intervals for regressions of child test score percentiles on parental income percentiles, estimated separately for each decile. Regressions include CZ fixed effects and use Education Longitudinal Study sampling weights. The dashed line shows an unweighted regression of the decile coefficient on the decile mean income transmission; its slope is shown in the lower right.

income coefficients and confidence intervals against the mean of θ_c in the decile. The slope of the best-fit line to this scatterplot, which corresponds to β in equation (10), is 0.25, indicating that parental income is more strongly associated with children's test scores in high income transmission than in low income transmission CZs. The implied difference between the third and eighth deciles—corresponding roughly to the interquartile range of θ_c—is a 0.021 increase in the slope shown in column 1 of table 3. However, the figure also indicates substantial variation around the best-fit line, likely more than could be explained by sampling error.

Table 4 presents parametric estimates of the interacted specifications (11) and (10). I begin with a specification that does not allow for a θ_c interaction, as in table 3, but show both within- and between-CZ coefficients.[14]

[14] The ELS is a multistage sample in which schools are sampled and students are sampled within schools. When I decompose $p_{ic} - \bar{p}_c$ into the deviation from the

Table 4
Transmission of Parental Income to Children's Twelfth-Grade Math Achievement (Education Longitudinal Study [ELS])

	(1)	(2)	(3)	(4)	(5)
Parental income − CZ mean	.35	.35	.34	.35	.33
	(.01)	(.01)	(.01)	(.01)	(.01)
CZ mean parental income	.69	.69	.70		.70
	(.04)	(.04)	(.04)		(.04)
CZ income transmission (θ)		−97.0	−74.4		−72.6
		(26.4)	(28.1)		(27.8)
(Parental income − CZ mean) × CZ income transmission (θ)		.32	.37	.32	.41
		(.21)	(.15)	(.21)	(.17)
CZ mean parental income × CZ income transmission (θ)		1.75	1.23		1.20
		(.53)	(.57)		(.56)
SD of parental income random coefficient (η)					.07
					(.02)
	CZ Effects				
	None	None	RE	FE	RE
Across-CZ distribution:					
SD of CZ income transmission (θ)		.057	.057	.057	.057
SD of parental income–test score transmission (π)		.018	.021	.018	.072
Coefficient of between-CZ regression of θ on π					.26
					(.12)
R^2					.11
Corr(θ, π)		1	1	1	.32
p-value, SD(η) = 0/corr(θ, π) = 1 (LR test)					<.01

NOTE.—The dependent variable in each column is the twelfth-grade math score in national percentile units (0–100). Parental income is also measured in percentiles (0–100). Commuting zone (CZ) income transmission is the relative mobility measure for the 1980–82 birth cohorts from Chetty et al. (2014a), de-meaned across CZs. CZ effects include random effects (RE) and fixed effects (FE). The RE specification in col. 3 is estimated via generalized least squares (GLS); the mixed model in col. 5 is estimated via maximum likelihood. Specifications in cols. 1, 2, and 4 are weighted using ELS sampling weights; cols. 3 and 5 are unweighted. Standard errors are clustered at the CZ level. The p-value in col. 5 is for a likelihood ratio (LR) test of the mixed model against a RE model with fixed coefficients (as in col. 3, although estimated by maximum likelihood rather than GLS). The number of observations (rounded to the nearest 10) is 13,650.

Columns 2–4 include interactions between parental income and CZ income transmission θ_c. Column 2 is an OLS specification without CZ-level variation in the intercept (but with a θ_c main effect), column 3 is a generalized least squares model with CZ random effects, and column 4 includes CZ fixed effects. The interaction coefficient from each specification estimates

school mean and the difference between school and CZ means, the across-CZ and within-CZ across-school coefficients are indistinguishable, and the within-school coefficient is much smaller. This is exactly what one would expect based on measurement error in p_{ic}, but it could also derive from sorting into schools on the basis of unobservables or school-based peer effects.

the coefficient β in equation (11); it measures the extent to which CZs with strong transmission from parental income to child income also exhibit strong transmission from parental income to children's test scores. This coefficient is 0.32 in the OLS and fixed effects specifications and 0.37 in the random effects specification; the difference reflects the use of sampling weights in the OLS and fixed effects specifications but not in the random effects specification. Each is a fair amount larger than the 0.25 estimate of β from figure 2 and is comparable in magnitude to the within-CZ parental income main effect. While income-achievement transmission is (in col. 3) 0.34 in the average CZ, it is 0.32 in a CZ at the 25th percentile of the θ_c distribution and 0.36 in a CZ at the 75th percentile. As in figure 2, these estimates are consistent with test scores being a meaningful, although not overwhelming, mediator of the between-CZ difference in the transmission of income across generations.

Column 5 presents the mixed model (10), allowing $\sigma_\eta^2 > 0$. This allows for variation across CZs in income-achievement transmission that is not predicted by the CZ's income-income transmission. I estimate $\sigma_\eta = 0.07$. The hypothesis that $\sigma_\eta^2 = 0$ is decisively rejected.[15] The lower portion of the table shows various summaries of the joint distribution of θ_c and π_c that is implied by the mixed model coefficients. The standard deviation of π_c is 0.072, much larger than in previous columns. A CZ at the 25th percentile of the π_c distribution has an income–test score transmission coefficient of 0.28, while one at the 75th percentile has a coefficient of 0.38. Most of this variation comes from the η_c component that is orthogonal to θ_c, however: the correlation between θ_c and π_c is only 0.32.

The slope of θ_c with respect to π_c is a statistically significant 0.26 (standard error, 0.12): on average, in CZs in which the test score advantage of children from rich families is 1 percentile larger than average, the adult income advantage is about 0.26 percentiles larger than in the average CZ. This is relatively small but consistent with the national evidence. Table 3, column 1, indicates that students whose test scores are 1 percentile above average tend to have adult incomes 0.24 percentiles above average. Here, we find that CZs in which the test score gap between high-income and low-income students is 1 percentile greater than in the average CZ have adult income gaps between those students that are on average 0.26 percentiles larger than in the average CZ. These are strikingly similar and appear to point to a meaningful role for student achievement as a mediator of income transmission.

[15] The null hypothesis that $\sigma_\eta = 0$ is on the boundary of the parameter space for the mixed model likelihood function. The test is a likelihood ratio test based on the comparison of col. 5 to the specification in col. 3, estimating the latter by maximum likelihood rather than by generalized least squares. Note that the null hypothesis that $\sigma_\eta^2 = 0$ corresponds to a perfect correlation between θ_c and π_c and to an R^2 of 1 in the regression of the former on the latter.

However, it is worth considering the magnitude of the across-CZ variation in test score transmission. I estimate that the standard deviation of π_c is 0.072. Take 0.24 as the return to a 1 percentile increase in test scores. Thus, each standard deviation of π_c drives an increase of θ_c of $0.072 \times 0.24 = 0.018$. But the standard deviation of θ_c is 0.057, more than three times as large. In other words, there is much more variability in CZ income-income transmission than can be operating through the test score channel: only 11% of the across-CZ variation in the former is explained by the latter. There are evidently other channels that account for the bulk of the geographic variation in income transmission; test scores—and the knowledge and skills that they represent—are a mechanism, but not the dominant one.

2. *Test Scores across Grades and Subjects*

By estimating the models in table 4 for test scores measured at different ages, I can explore whether the relative advantage of high-income children in high-transmission CZs appears to grow with age, as might be expected if schools play a role in income transmission. This analysis is likely sensitive to scaling decisions (Bond and Lang 2013, forthcoming). I scale scores at each age in national percentiles, but a 1 percentile advantage in kindergarten may not mean the same thing as a 1 percentile advantage in twelfth grade. Setting this issue aside, table 5 presents mixed model estimates for each of the available test scores from the ECLS, ELS, and HSLS. The β coefficients in column 2 are similar in magnitude across most of the specifications, although they are imprecisely estimated. The random component of the parental income coefficient (σ_η, in col. 3) is meaningful in each row, and column 6 indicates that the null hypothesis that $\sigma_\eta = 0$ is rejected in all but one case. The slope of θ_c with respect to π_c (col. 4) is modest and generally larger for reading than for math. It appears to grow somewhat with age, although this is not entirely consistent. Correlations between θ_c and π_c (col. 5) are quite low across grades and subjects, but again they are larger in later grades.

The pattern of results has several implications. First, there is some indication that the relative importance of parental income to student test scores in high income transmission CZs grows between kindergarten and high school, consistent with the hypothesis that differential access to school quality (rather than, say, parenting practices) is a mechanism contributing to differential income transmission. Second, there is substantial heterogeneity across CZs in the transmission of parental income to children's test scores that is not associated with CZ-level income transmission, indicating that the institutions or other CZ characteristics that contribute to test score transmission differ from those determining income transmission. Put somewhat differently, there is only a modest correlation across CZs between income-income and income–test score transmission, even in later grades, so different influences must be at work. Finally, results for the HSLS are

Table 5
Parental Income–Child Achievement Transmission across Grades, Cohorts, and Subjects

	Parental Income (1)	Parental Income × CZ Income Transmission (2)	SD of Parental Income Random Coefficient (η) (3)	Coefficient of Regression of Income Transmission (θ) on Test Score Transmission (π) (4)	Corr (θ, π) (5)	p-value, LR Test of SD(η) = 0 (6)
	A. Math Scores					
ECLS K (spring)	.35	.33	.08	.17	.24	<.01
	(.01)	(.25)	(.01)	(.13)		
ECLS G1 (spring)	.35	.08	.06	.09	.08	<.01
	(.01)	(.24)	(.01)	(.29)		
ECLS G3	.42	.13	.08	.08	.10	<.01
	(.01)	(.23)	(.01)	(.14)		
ECLS G5	.39	.31	.09	.13	.20	<.01
	(.01)	(.26)	(.01)	(.11)		
ECLS G8	.41	.22	.07	.16	.19	.01
	(.01)	(.22)	(.02)	(.17)		
HSLS G9	.30	.30	.05	.33	.32	.02
	(.01)	(.17)	(.01)	(.19)		
HSLS G11	.28	.60	.07	.30	.43	<.01
	(.01)	(.18)	(.01)	(.09)		
ELS G10	.31	.37	.06	.29	.33	<.01
	(.01)	(.16)	(.01)	(.13)		
ELS G12	.33	.41	.07	.26	.32	<.01
	(.01)	(.17)	(.02)	(.12)		
	B. Reading Scores					
ECLS K (spring)	.38	.16	.08	.09	.12	<.01
	(.01)	(.23)	(.01)	(.14)		
ECLS G1 (spring)	.38	.23	.06	.21	.22	<.01
	(.01)	(.22)	(.01)	(.20)		
ECLS G3	.40	.41	.06	.33	.37	<.01
	(.01)	(.21)	(.02)	(.19)		
ECLS G5	.39	.48	.06	.37	.42	<.01
	(.01)	(.21)	(.01)	(.15)		
ECLS G8	.39	.33	.05	.48	.40	.21
	(.01)	(.21)	(.02)	(.35)		
ELS G10	.30	.25	.07	.15	.19	<.01
	(.01)	(.18)	(.01)	(.11)		

NOTE.—Each row presents statistics from a single mixed model pertaining to a different test score (for a given sample, grade, and subject), each scaled as national percentile units (0–100). Parental incomes in cols. 1–3 are percentiles, deviated from the commuting zone (CZ) mean. Specifications are as in table 4, col. 5. See the note to table 4 for details. The number of observations (rounded to the nearest 10) ranges between 9,140 and 20,430. ECLS = Early Childhood Longitudinal Study; ELS = Education Longitudinal Study; G = grade; HSLS = High School Longitudinal Study; K = kindergarten; LR = likelihood ratio.

quite similar to those for the ELS, although the latter is much closer to the cohorts for which θ_c is computed, suggesting that cohort differences are unable to explain the weak relationship of income-income and income–test score transmission in the HSLS and ECLS.

3. *Educational Attainment*

I have thus far used test scores as a summary of children's human capital. An alternative is to focus on educational attainment. I consider two summaries of attainment as of the last ELS survey, around age 26: an indicator for a 4-year degree and the number of years of education. As discussed above, the ELS counts a surprisingly large share of students as having attended some college, and results for this outcome (presented in the appendix) are highly discrepant and appear to be driven by overmeasurement of college attendance in the ELS. The other two attainment summaries are closer to expectations (table 2).

Columns 1 and 3 of table 6 present estimates of the interacted specification (11) with CZ random effects for the two measures. (I report only the coefficients pertaining to within-CZ variation in parental income, although

Table 6
Parental Income–Child Educational Transmission (Education Longitudinal Study)

	College Graduation (0/100)		Years of Education at 26 (×100)	
	(1)	(2)	(3)	(4)
Parental income − CZ mean	.45	.45	1.85	1.86
	(.02)	(.02)	(.06)	(.06)
(Parental income − CZ mean) × CZ income transmission (θ)	.64	.74	2.30	2.35
	(.30)	(.29)	(1.12)	(1.09)
SD of parental income random coefficient (η)		.08		.22
		(.03)		(.13)
Across-CZ distribution:				
SD of CZ income transmission (θ)	.056	.056	.056	.056
SD of parental income–test score transmission (π)	.036	.086	.130	.254
Coefficient of between-CZ regression of θ on π		.32		.12
		(.19)		(.11)
R^2		.24		.27
Corr(θ, π)	1	.49	1	.52
p-value, SD(η) = 0/corr(θ, π) = 1 (LR test)		.15		.33

Note.—Specifications in cols. 1 and 3 are as in table 4, col. 3; those in cols. 2 and 4 are as in table 4, col. 5. All columns include controls for commuting zone (CZ) mean parental income, CZ income transmission, and their interaction. See the note to table 4 for details. The dependent variable in cols. 1–2 is scaled as 0 for non–college graduates and 100 for graduates; in cols. 3–4, the dependent variable is years of education multiplied by 100. Standard errors are clustered at the CZ level. The number of observations (rounded to the nearest 10) is 13,250. LR = likelihood ratio.

CZ means and an income transmission main effect are included as before.) Not surprisingly, parental income is strongly related to both measures of children's attainment. The interaction coefficient β is large and statistically significant for each outcome.

Columns 2 and 4 present the mixed model specifications. Likelihood ratio tests do not reject the restrictions that the parental income random coefficients are zero (i.e., $\sigma_\eta = 0$). Coefficients of regressions of income transmission on income-attainment transmission yield modest coefficients: CZs in which students from high-income families are 1 percentage point more likely to graduate from college (relative to students from low-income families) have adult income gaps between children from high- and low-income families that are 0.32 percentiles larger, and CZs in which the high-income children earn one more year of education have adult income gaps that are 12 percentiles larger. Neither of these is significantly different from zero.

The correlation between income transmission and attainment transmission is stronger than that for test scores, around 0.5. However, this is still quite far from 1; three-quarters of the variance in income transmission across CZs is unexplained by differences in transmission from parental income to children's higher education attainment. As in the earlier analysis of test scores, the evidence points to a role for educational attainment as a mechanism driving variation in intergenerational income transmission but does not indicate that it is an overwhelming factor.

The R^2 statistics in the lower portion of table 6 provide one way to measure the importance of the attainment channel. Variation across CZs in the transmission of parental income to educational attainment in years explains about one-quarter of the variability in CZ-level income transmission. As with test scores, another way to understand this is to use an estimate of the return to education to measure the importance of educational attainment as a mediator of income transmission. I begin with Chetty et al.'s (2014a) measure of transmission from parental income to college enrollment. They find that the standard deviation of π_c, across CZs, is 0.11 percentage points of college enrollment per percentile of family income, very similar to my estimate in column 2 of table 6. In a regression of family income percentiles on an indicator for some college in the ACS sample with CZ fixed effects, I find that those with some college or more have family incomes about 19.2 percentiles higher than those without college on average. This implies that a 1 standard deviation increase in π_c would drive only a $19.2 \times 0.11 = 0.02$ increase in θ_c, or less than one-third of a standard deviation of that variable. I obtain even smaller magnitudes when I use my estimates of transmission of parental income to other attainment measures. For example, column 4 of table 6 indicates that a 1 standard deviation of π_c is 0.0025 years of education per percentile of parental income. Column 5 of table 3 indicates that each year of education is associated with 3.6 additional percentiles of children's income.[16] Thus, a 1 standard deviation increase in π_c drives an increase of θ_c of $0.0025 \times 3.6 = 0.01$, or about one-sixth of a standard deviation. Although the transmission of parental income to children's income is correlated across CZs with transmission of parental income to children's educational attainment, the latter again appears not to be a primary mechanism for the former.

4. Robustness and Additional Results

The results given above indicate that CZs with stronger than average transmission of parental income to children's income tend to also have stronger than average transmission of parental income to children's test scores and educational attainment, but the relationships are not large enough to account for a large share of the variation in intergenerational income transmission. This basic conclusion is robust to a variety of different specification and measurement choices, explored in the appendix.

First, table A5 explores the sensitivity of these results to the choice of income transmission measure. Results are robust to using Chetty et al.'s (2014a) measure computed for the 1983–85 birth cohorts, which more closely corresponds to the ELS sample, or to using the more plausibly causal measure from Chetty and Hendren (2018).

Second, I show that the results are not driven by associations between parental income and children's race. Chetty et al. (2014a) document that θ_c is quite strongly correlated with the fraction black in the CZ, although they also find that an alternative measure computed solely from zip codes with very few black residents is quite similar. Table A6 augments the main mixed model specifications with controls for the child's own race and gender as well as interactions of race and gender with θ_c. These have little effect.

Third, I explore alternative scalings of parental income and children's test scores in table A7. The basic result of a weak relationship between CZ-level income transmission and CZ-level transmission from parental income to children's achievement is unchanged when I measure children's test scores as z-scores or as predicted adult earnings (Bond and Lang, forthcoming) or when I measure parental income in logs or as predicted children's test scores.

Overall, the basic results for achievement, attainment, and income transmission appear quite robust. They are suggestive that learning in school is not a key channel determining the across-CZ variation in income transmission but that access to higher education may be more important.

[16] I obtain slightly larger estimates when I use the ACS sample—4.1 using the full sample, or 5.9 when very high and very low levels of attainment are trimmed. These would not change the qualitative conclusion I draw here.

One possibility not yet considered is that math and reading test scores do not fully capture the impacts of better childhood environments. A growing literature in recent years has documented the importance of noncognitive skills as a component of human capital. Both the ECLS and the ELS contain batteries of questions aimed at identifying children's noncognitive skills, and I use these to assess whether high income transmission CZs tend to be CZs with large gaps in noncognitive skills between children from high- and low-income families (table A8). Results are mixed. The β coefficient on the parental income–CZ income transmission interaction is generally small and not statistically significant, and it frequently has the wrong sign. For about half of the available measures, there is statistically significant variation across CZs in the return to parental income (i.e., $\sigma_\eta \neq 0$). Overall, there is little indication that noncognitive skills are important mediators of income-to-income transmission. One set of results, however, tells a somewhat different story. In the ECLS, noncognitive skill measures are constructed both from children's survey responses and from teacher surveys. The measures based on teacher surveys do tend to yield strong associations with income transmission. It is not clear how to account for the discrepancy between teacher surveys and student self-reports—even when the concepts overlap (e.g., for externalizing problem behaviors), results are quite different. This may indicate that high-transmission CZs tend to be CZs in which teachers are more biased in their assessments of low-income children, but this is quite speculative.

C. Returns to Human Capital across CZs

The results described above have concerned the role of skills—achievement, attainment, and noncognitive skills—as mediators of the intergenerational transmission of income. In terms of figure 1, the results suggest that π_c is not a primary mechanism influencing variation in reduced-form transmission θ_c. This in turn implies that much of the variation in income transmission must be due to differences in the returns to human capital (i.e., in λ_c) or to direct effects of parental income on children's income not operating through human capital (i.e., to μ_c).

As an initial exploration of this, I examine variation in the return to skill across CZs. As before, I estimate mixed models, in this case allowing the return to human capital to vary both with the Chetty et al. (2014a) income transmission measure and independently across CZs. These models do not isolate the relationship between income transmission and λ_c from figure 1, as to do that I would need to examine the return to skill controlling for parental income. I simply examine the reduced-form return to skill, $\tilde{\lambda}_c = \lambda_c + (\sigma_p^2/\sigma_s^2)\pi_c\mu_c$. If I find that this is strongly associated with θ_c, which could indicate either that λ_c is a major component of the across-CZ variation in θ_c or that μ_c is.

Panel A of table 7 presents results for a sample of 28–32-year-olds surveyed in 2010–12 by the ACS and assigned to their current CZs.[17] Column 1 shows that each year of education, relative to the CZ mean, is associated with 5.3 percentiles of adult earnings. Columns 2–4 present models that include interactions between the individual's education and CZ-level income transmission. The interaction coefficient is positive and highly significant, indicating that the (reduced-form) return to education is larger in high income transmission CZs. Column 5 presents the mixed model, allowing for unexplained heterogeneity across CZs in the return to education. This heterogeneity term is substantial. The correlation between the CZ-level return to education and CZ-level income transmission is about 0.5, comparable to the attainment transmission results earlier and much larger than that for achievement transmission. The overall variability in returns to education across CZs (i.e., in $\tilde{\lambda}_c$) is substantial, with a standard deviation of 0.7 (compared with the mean of 5.3). Only about 30% of this is attributable to θ_c.

Panel B of table 7 presents a parallel analysis of returns to skill in the ELS data. Here, I combine my two human capital measures, constructing a skill index as the fitted value from a regression of children's earnings on their twelfth-grade math scores and indicators for each possible attainment, with CZ fixed effects. This skill index is strongly related to earnings, as expected.[18] It is much more strongly related in high-transmission CZs, with interaction coefficients that are notably larger than the main effects. Column 5 indicates, however, that there is a great deal of variation in the returns to skill that is orthogonal to income transmission, and the correlation between the two is only 0.3.

V. Decomposing the Across-CZ Variation in Income Transmission

The results thus far indicate that intergenerational income transmission is positively correlated across CZs with transmission from parental income to children's test scores and educational attainment and with the reduced-form

[17] I censor years of education at 9 and 17. Values outside this range are unusual. The earnings-education relationship is approximately linear within this range but not outside it. Table 7 shows results for the individual earnings percentile as the dependent variable, but results are similar when the family income percentile is used instead.

[18] In constructing the skill index, I measure children's earnings as a percentile of the adult income distribution for use in my decomposition below. Thus, a child with median earnings ($22,000 in the ELS sample) is assigned a percentile of 38, as $22,000 is the 38th percentile of the family income distribution. The dependent variable in table 7 is the percentile of the child earnings distribution, in which the same child would be assigned a percentile of 50. This explains why the coefficient is larger than 1 in col. 1.

Table 7
Returns to Education in the American Community Survey (ACS) and Education Longitudinal Study (ELS) Samples

	(1)	(2)	(3)	(4)	(5)
	A. Returns to Education in ACS Data				
Years of education − CZ mean	5.34	5.35	5.53	5.35	5.18
	(.08)	(.08)	(.07)	(.08)	(.05)
(Years of education − CZ mean) × CZ income transmission (θ)		3.97	4.58	3.97	6.95
		(1.05)	(1.04)	(1.05)	(.94)
SD of education random coefficient (η)					.62
					(.04)
	CZ Controls				
	None	None	RE	FE	RE
Across-CZ distribution:					
SD of CZ income transmission (θ)		.056	.056	.056	.056
SD of return to education ($\tilde{\lambda}$)		.222	.256	.222	.728
Coefficient of between-CZ regression of θ on $\tilde{\lambda}$					.04
					(.00)
R^2					.29
Corr(θ, $\tilde{\lambda}$)		1	1	1	.53
p-value, SD(η) = 0/corr(θ, $\tilde{\lambda}$) = 1 (LR test)					<.01
	B. Returns to Skills in ELS Data				
Skill index − CZ mean	1.09	1.09	1.07	1.06	1.08
	(.04)	(.04)	(.04)	(.04)	(.04)
(Skill index − CZ mean) × CZ income transmission (θ)		2.30	1.28	2.30	1.26
		(.67)	(.66)	(.68)	(.73)
SD of skill index random coefficient (η)					.21
					(.07)
	CZ Controls				
	None	None	RE	FE	RE
Across-CZ distribution:					
SD of CZ income transmission (θ)		.057	.057	.057	.057
SD of return to education ($\tilde{\lambda}$)		.130	.072	.130	.219
Coefficient of between-CZ regression of θ on $\tilde{\lambda}$					.08
					(.07)
R^2					.10
Corr(θ, $\tilde{\lambda}$)		1	1	1	.32
p-value, SD(η) = 0/corr(θ, $\tilde{\lambda}$) = 1 (LR test)					.14

NOTE.—In panel A, the sample consists of individuals born between 1980 and 1982 in the ACS 2010–12 one-year public-use microdata samples (N = 241,670). Respondents are assigned to their commuting zone (CZ) of current residence. The dependent variable is the child's earnings percentile (0–100). Years of education is naturally coded, with values below 9 or above 17 set to missing. In panel B, the sample is the ELS sample (N = 9,980). Skill index is the fitted value from a regression of children's earnings percentiles at age 25, scaled as a percentile of the family income distribution, on their twelfth-grade math score percentile and dummies for years of schooling completed, with CZ fixed effects (FE). Specifications match the corresponding columns of table 4. Columns 2, 3, and 5 include controls for CZ mean years of education (panel A) or skill index (panel B), the CZ income transmission, and their interaction; coefficients on CZ-level covariates are not reported. See the note to table 4 for details. LR = likelihood ratio; RE = random effects.

labor market returns to human capital. Some preliminary calculations indicate that neither the achievement nor the attainment relationship is large enough on its own to be a primary channel in overall income transmission, but I have not yet considered them together or quantified the contribution of the returns-to-skill effects. Moreover, the returns-to-skill estimates are reduced form and combine true returns to skill with any effect of parental income on children's income not operating through education (i.e., with μ_c). In this section, I explore decompositions of the across-CZ variation in income transmission that address these shortcomings.

As a preliminary, I explore income-income transmission, θ_c, in the ELS data. Measurement differences between the ELS and the tax data used by Chetty et al. (2014a) mean that the θ_c's implied by the ELS data differ somewhat from the tax data–based θ_c's reported by Chetty et al. (2014a)—although they are nearly perfectly correlated. I also show that marriage patterns and labor force participation are quantitatively important channels for intergenerational income transmission. This motivates me to extend the three-component path diagram from figure 1 by considering transmission into children's own earnings and into the other components of family income (spousal earnings and nonlabor income) separately. I decompose the transmission of parental income to children's earnings into the three components from the path diagram in figure 1 and equation (22): (*a*) children's skill accumulation by the end of school, (*b*) returns to skills, and (*c*) direct effects of parental income on children's earnings not operating through observed human capital. I then separately estimate the contribution from transmission of parental income into spousal earnings and nonlabor income. Figure 3 illustrates the expanded diagram.

A. Income Transmission in the ELS Sample

Table 8 presents mixed models, akin to those used earlier to examine transmission from parental income to children's achievement, where here the dependent variables are different components of ELS children's incomes.

In column 1, the dependent variable is the child's total family income as a percentile of the national distribution.[19] This is very nearly the same measure used by Chetty et al. (2014a) to construct their income transmission measures. Thus, we expect the π_c in this specification, the CZ-level transmission of parental income to children's income in the ELS sample, to be identical to the θ_c of Chetty et al. (2014a) but for differences in measurement between the ELS and the tax data. Indeed, I estimate a correlation of 0.99, and $\sigma^2_\eta \approx 0$. However, the scales are somewhat different: where one would expect an interaction coefficient $\beta = 1$, I instead estimate $\hat{\beta} = 0.64$ (standard error, 0.16). The implied regression of θ_c on π_c has a coefficient of

[19] Table A2 presents fixed-coefficient versions of this specification.

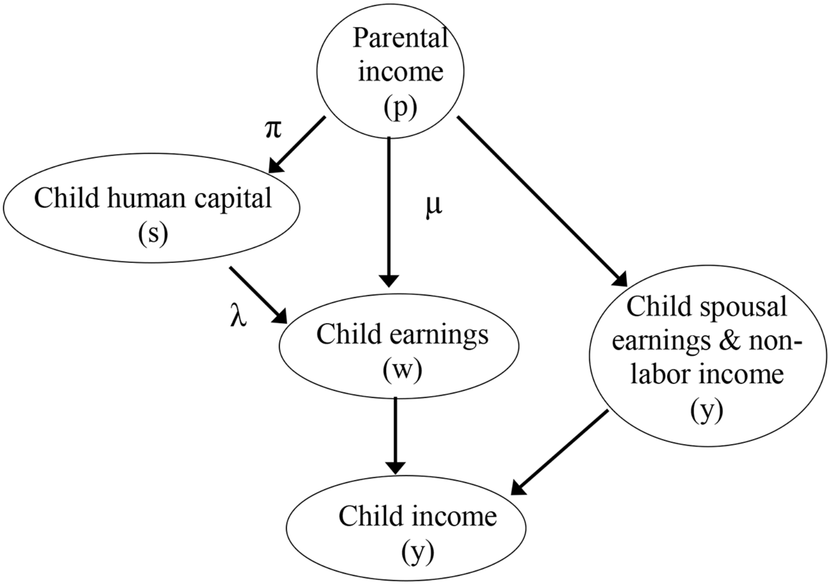

FIG. 3.—Path diagram with spousal earnings and nonlabor income.

1.52, although here the expected 1 is within the confidence interval. These results might reflect the lower quality of the ELS parental income measure relative to the tax data[20] or the fact that the ELS child income is measured at age 25, while in the tax data it is measured around age 30.[21]

Columns 2–4 replace the dependent variable with indicators (scored as 0 or 100) for positive own earnings, for being married, and for having positive spousal earnings. In each case, the interaction coefficient between parental income and CZ-level income transmission is positive and significant: in high-transmission CZs, children from high-income families are relatively more likely than children from low-income families to work, to be married, and to have a working spouse. In each case, the across-CZ correlation between income transmission and transmission of parental income to the outcome is around 0.5. Evidently, an important part of the variation in income transmission reflects labor force participation rather than solely differences

[20] I have explored specifications that instrument for parental income with parental education. Although the exclusion restriction is dubious, it does raise the β coefficient to around 1, consistent with bias from measurement error in parental income.

[21] Chetty et al. (2014a) find that average income transmission is lower when children's income is measured at younger ages, but they do not present evidence regarding cross-CZ variation.

in earnings conditional on participation; another important part relates to marital patterns.

Further light is shed by the gender breakdown in panel B. CZ-level income transmission is almost perfectly correlated with the CZ-level association between parental income and daughters' labor force participation, although there is little variation across CZs in the latter.

In column 5, I use the child's earnings as the dependent variable. Earnings are scaled here as a percentile of the family income distribution, to permit a direct comparison to column 1 (see n. 18). The p_{ic}-θ_c interaction coefficient is here only 0.38 (standard error, 0.14), reduced by nearly half from column 1. Evidently, a large part of the variation in measured income transmission, using the definitions of Chetty et al. (2014a), derives from components other than the child's own earnings—either spousal earnings or nonlabor income. This is particularly true for men. Column 6 adds nonlabor income (for both the child and the spouse, if present) into the income measure. Results are similar to those in column 5. The key interaction coefficient remains much lower than in column 1, especially for sons.

Spousal earnings, the only component of family income included in column 1 but not column 6, are clearly an important factor. This could reflect variation across CZs in the relative likelihood that children from high- and low-income families have working spouses, as seen in column 4, but it could also reflect differences in spousal earnings distributions conditional on work, as would occur if CZs vary in the degree of assortative mating. To assess the role of the latter, I shut off any assortative matching by assigning all working spouses the same earnings. I compute the average earnings across the entire sample for working spouses by gender—\$27,000 for women and \$41,000 for men—and use these for every working spouse in the sample, assigning 0 for those who are unmarried or have nonworking spouses. I then construct a family income as the sum of the child's actual earnings, any nonlabor income, and imputed spousal earnings. As before, this sum is converted to a percentile of the actual child family income distribution. Insofar as an important part of the variation in income transmission reflects differences in assortative mating, we would expect the β coefficient in column 7 to more closely resemble that in column 6 than that in column 1. This is not what I find—$\hat{\beta}$ here is even larger than in column 1. Evidently, differences in assortative mating are not contributing meaningfully to the across-CZ variation in family income transmission.

I interpret the results in table 8 as pointing to the importance of marriage as a mechanism driving between-CZ variation in measured income transmission. Nearly one-third of the across-CZ variation in income transmission is explained by differences in within-CZ gradients of marriage (at the time of the age 26 ELS follow-up survey) with respect to parental income. This may represent a spurious component of the variation in θ_c. It

Table 8
Transmission of Parental Income to Children's Income, Earnings, and Marital Status

	Child Family Income (1)	Own Earnings > 0 (0/100) (2)	Marital Status (0/100) (3)	Working Spouse (0/100) (4)	Child Earnings (5)	Child Earnings + Nonlabor Income (6)	Child Earnings + Nonlabor Income + Imputed Spousal Earnings (7)
	A. Full Sample						
Parental income − CZ mean	.17	.15	.00	.01	.15	.18	.15
	(.01)	(.01)	(.01)	(.01)	(.01)	(.01)	(.01)
(Parental income − CZ mean) × CZ income transmission (θ)	.64	.47	.61	.61	.38	.45	.76
	(.16)	(.25)	(.27)	(.27)	(.14)	(.13)	(.16)
SD of parental income random coefficient (η)	.006	.045	.052	.060	.020	.023	.015
	(.018)	(.023)	(.026)	(.021)	(.015)	(.015)	(.015)
Across-CZ distribution:							
SD of Chetty et al. (2014a) income-income transmission coefficient (θ)	.057	.057	.056	.056	.057	.057	.057
SD of ELS parental income-child outcome transmission	.037	.052	.062	.069	.029	.035	.045
Coefficient of regression of θ on ELS transmission	1.52	.55	.50	.40	1.41	1.21	1.18
	(.37)	(.42)	(.34)	(.23)	(.98)	(.68)	(.27)
R^2	.97	.26	.31	.24	.53	.55	.90
Correlation	.99	.51	.55	.49	.73	.74	.95
p-value, SD(η) = 0 (LR test)	.92	.46	.18	.13	.45	.31	.60

	B. By Gender						
Men:							
(Parental income − CZ mean) × CZ income transmission	.53	.49	.48	.66	.19	.27	.52
	(.23)	(.31)	(.39)	(.34)	(.22)	(.23)	(.23)
SD of ELS parental income–child outcome transmission	.033	.060	.073	.085	.032	.042	.030
Coefficient of regression of θ on ELS transmission	1.61	.43	.29	.29	.61	.48	1.85
	(.77)	(.51)	(.49)	(.30)	(.95)	(.55)	(.81)
R^2	.85	.21	.14	.19	.12	.13	.96
Correlation	.92	.46	.37	.44	.34	.36	.98
Women:							
(Parental income − CZ mean) × CZ income transmission	.65	.35	.89	.73	.49	.54	.87
	(.25)	(.31)	(.38)	(.39)	(.19)	(.17)	(.24)
SD of ELS parental income–child outcome transmission	.038	.020	.086	.096	.038	.037	.051
Coefficient of regression of θ on ELS transmission	1.47	2.85	.39	.25	1.06	1.24	1.08
	(.56)	(2.53)	(.22)	(.16)	(.50)	(.44)	(.26)
R^2	.96	.99	.34	.19	.51	.66	.94
Correlation	.98	.99	.59	.43	.72	.81	.97

Note.—All specifications are as in table 4, col. 5. All columns include main effects for commuting zone (CZ) mean parental income and CZ income transmission and their interaction. See the note to table 4 for details. In cols. 1 and 5–7, the dependent variable is a measure of child income, scaled as a percentile (0–100) of the child total family income distribution. In cols. 2–4, the dependent variable is an indicator, multiplied by 100. The number of observations (rounded to the nearest 10) ranges from 11,510 to 16,200 in panel A. ELS = Education Longitudinal Study; LR = likelihood ratio.

is not clear whether a two-earner couple should be seen as being as successful as a single person with the same family income. Moreover, the median age of marriage for the ELS cohorts is around 26 (US Bureau of the Census 2004), so it is quite possible that many people who are not married at age 26 or even at 30 will be later and will eventually be able to pool their earnings with their spouses to achieve much higher family incomes than I see in the age 26 survey.

Whether transmission operating through marriage is spurious or not, the interpretation of income-marriage transmission is quite different than that of income-earnings transmission, even though both may be statistically mediated by the child's human capital. Going forward, I separate children's family incomes into the child's own earnings and a second component combining spousal earnings and nonlabor income, and I focus on the mediating role of human capital for the former.

B. Decomposition of Income Transmission

Table 9 presents my analysis of the decomposition of across-CZ variation in income transmission into the four components indicated in figure 3: skill accumulation, as moderated by the average own-earnings return to skill; returns to skill, moderated by the average parental income gradient in skill accumulation; "direct" transmission of parental income to children's earnings conditional on human capital; and spousal and nonlabor income.

Column 1 presents the baseline income transmission analysis, using the family income percentile as the dependent variable. This specification is the same as in column 1 of table 8 but omits the random coefficient on parental income. Of interest is the interaction between $p_{ic} - \bar{p}_c$ and θ_c. This coefficient would be identically 1 if I used the same sample and income measures that were used by Chetty et al. (2014a) in their calculation of θ_c. My estimate is just over two-thirds of that.

Next, I decompose children's family incomes into the child's earnings and the remainder, reflecting spousal earnings and nonlabor income. I scale children's earnings as a percentile of the family income distribution, as in table 8, and then scale the remaining component as the increment to the family's income percentile that is obtained by adding spousal earnings and nonlabor income. Column 2 presents the analysis of children's earnings, using the same specification as in column 1. The interaction coefficient falls by nearly half, to 0.37—as in table 8, only a bit more than half of the across-CZ variation in parental income–child income transmission is attributable to variation in parental income–child earnings transmission.

Columns 3–5 decompose the transmission into child earnings into three components, reflecting skill accumulation, returns to skills, and direct transmission, using the methods introduced in Section III.C. In column 3, I show the component reflecting skill accumulation. I use the same skill index

Table 9
Decomposition of the Variation in Intergenerational Transmission

	Mechanism					
	Family Income	Own Earnings				Nonlabor and Spousal Income
	Total Transmission	Total Transmission	Skills	Return to Skills	Residual	Total Transmission
	DV: Child Income (1)	DV: Child Earnings (2)	DV: Child Skill Index (3)	DV: Child Earnings (4)	(5)	DV: Family Income Less Own Earnings (6)
A. Full Sample						
(Parental income − CZ mean) × CZ income transmission	.69	.37	.08	.23		.28
	(.17)	(.15)	(.06)	(.15)		(.11)
(Skill index − CZ mean) × CZ income transmission				.84		
				(.66)		
Scale factor:						
λ			.99			
π				.09		
Scaled component	.69	.37	.08	.07	.23	.28
Share of col. 1 (%)	100	54	11	11	33	41
B. By Gender						
Men:						
Scale factor			.80	.09		
Scaled component	.50	.31	.08	.16	.09	.21
Share of col. 1 (%)	100	63	16	32	18	42
Women:						
Scaled component	.78	.32	.05	−.03	.30	.43
Share of col. 1 (%)	100	41	7	−4	39	55

NOTE.—Each specification has controls for commuting zone (CZ) mean parental income, the individual deviation from that mean, income transmission, and an interaction between CZ mean income and CZ income transmission. Columns 4 and 5 report a single specification, which also includes the CZ mean of the skill index (see the note to table 7 for details) and its interaction with income transmission. See the main text for an explanation of scale factors and scaling of dependent variables (DVs).

used in table 7, combining twelfth-grade math scores and years of completed education, scaled as the predicted child earnings percentile. By construction, the return to this index in child earnings, $\bar{\lambda}$, is almost identically 1.[22] I repeat the random effects regression from column 2, replacing the child's actual earnings percentile with the skill index. Not surprisingly given the earlier results, the interaction term, which represents the first term of the decomposition (22), is small and not statistically significant. The point estimate of 0.08 implies that relative skill accumulation of children from high- and low-income families—and the earnings gap that it generates—accounts for only 11% ($=0.08/0.69$) of the differences in ELS income transmission between cities with low and high values of the Chetty et al. (2014a) transmission measure.

Columns 4 and 5 explore the role of returns to skill and direct transmission, respectively. These come from a single regression of the child's actual earnings on her skill index and parental income, each interacted with CZ income transmission. The skill-θ_c interaction coefficient estimates $\partial\lambda_c/\partial\theta_c$; the second term of the decomposition (22) can then be obtained by multiplying it by the coefficient of parental income in a pooled regression for children's skill, $\bar{\pi} = 0.09$. Thus, the second term, in column 4, is $0.84 \times 0.09 = 0.07$, indicating that differences in returns to skill account for another 11% of the variation in income transmission.[23] The third component of the decomposition (22), in column 5, is estimated by the parental income–θ_c interaction, 0.23. This indicates that differences in the relationship between parental income and child earnings, controlling for both the child's human capital and the CZ-level return to that human capital, account for one-third of the total variation in income transmission across CZs.

Finally, column 6 presents results for the portion of family income deriving from spousal earnings and nonlabor income. This is significantly more strongly related to parental income in CZs that Chetty et al. (2014a) measure as high transmission than in those measured as low transmission, and this accounts for 41% of the total variation in CZ-level income transmission. As table 8 indicates, this largely reflects differences in the likelihood of being married at age 26, not differences in assortative matching.

Panel B of table 9 reports the decomposition separately for boys and girls. The income transmission measure of Chetty et al. (2014a) better captures parent-daughter family income relationships than it does parent-son rela-

[22] The skill index is constructed on the basis of a weighted regression, but I estimate $\bar{\lambda}$ without weights for consistency with the unweighted random effects models in table 9. The resulting $\hat{\bar{\lambda}} = 0.99$.

[23] This component also captures differences in the accumulation of unobserved skills not measured by math scores or educational attainment: these constitute an omitted variable that is correlated with observed skill, so if some CZs have stronger gradients of unobserved skill with respect to parental income, they would appear to have higher returns to observed skill.

tionships. Transmission to the child's own earnings is similar for both, so it represents a larger share of the total for sons. For them, returns to skills are twice as important as skill accumulation. For daughters, the return to skill is actually negatively associated with θ_c, and skill accumulation is only trivially positively associated. All of the variation in transmission to earnings is operating through the direct component, controlling for human capital. This in part reflects variation in the relationship between parental income and daughters' labor force participation, as documented in table 8. The contribution of spousal earnings to between-θ_c differences in family income transmission is about twice as large for girls as for boys. Table 8 indicates that this is largely due to a stronger role of marital status for girls, not to greater assortative mating.

Overall, these results make clear that differences in skill accumulation—achievement and attainment—account for only a small share of the variation across CZs in income transmission. Marriage patterns are the largest single channel explaining family income transmission. For sons, returns to skills also play a meaningful role, while for daughters the transmission of parental income to children's earnings not mediated by human capital is more important.

VI. Conclusion

The pathbreaking work of Chetty et al. (2014a) showed that there is dramatic variation in intergenerational income mobility across geographic areas within the United States. This raises the intriguing possibility that we can identify policies that account for this variation and, by exporting these policies from high- to low-mobility areas, move closer to equality of opportunity.

Chetty et al. (2014a) presented suggestive correlations indicating that school quality might be an important contributing factor. This paper has investigated this suggestion by asking whether high- and low-income children's academic outcomes are more equal in areas where their adult economic outcomes are more equal. I find that there is statistically significant variation across CZs in the gradients of educational attainment, academic achievement, and noncognitive skills with respect to parental income. This variation is positively correlated with variation in income transmission across CZs, but the correlations are modest. Moreover, while substantial, the variation in human capital transmission is not large enough in magnitude to be a primary mechanism by which income is transmitted across generations.

I find that only about one-ninth of the across-CZ variation in intergenerational income mobility is attributable to differences in children's earnings deriving from differences in the accumulation of observed skills. A similar share is attributable to differences in the labor market returns to children's skills. About one-third is attributable to differences in the labor market return to parental income holding skills (and the returns to skills) constant. The remaining, largest portion derives from differences in spousal and nonlabor income, primarily reflecting differences in the likelihood of having a working spouse.

Taken together, these facts indicate that the education system makes only a modest contribution to variation in intergenerational income transmission. The evidence points to other factors as potentially more important, including cultural tendencies toward early marriage and local labor market factors that influence the labor force participation rate and the ability of children from high-income families to match into high-earnings jobs conditional on their education and skills.

This is not to say that school quality is not important for other reasons, of course, or even that it does not contribute to overall mobility in a way that is roughly constant across CZs. Nevertheless, further investigation into the determinants of local intergenerational mobility should expand from a near-exclusive focus on education to other potential mechanisms. One area for further attention is differences in the likelihood of marriage, either because CZs vary in the likelihood that romantic partners will be formally married or because of variation in partnership rates. In terms of earnings outcomes, other areas of interest include local income inequality and labor market institutions that influence it (e.g., unions) as well as factors influencing the strength of local labor market networks and the spatial and social stratification of the local market.

References

Aaronson, Daniel, and Bhashkar Mazumder. 2008. Intergenerational economic mobility in the United States, 1940 to 2000. *Journal of Human Resources* 43:139–72.

Bond, Timothy N., and Kevin Lang. 2013. The evolution of the black-white test score gap in grades K–3: The fragility of results. *Review of Economics and Statistics* 95:1468–79.

———. Forthcoming. The black-white education-scaled test-score gap in grades K–7. *Journal of Human Resources*.

Bradbury, Bruce, Miles Corak, Jane Waldfogel, and Elizabeth Washbrook. 2015. *Too many children left behind: The US achievement gap in comparative perspective*. New York: Russell Sage Foundation.

Chetty, Raj, and Nathaniel Hendren. 2018. The impacts of neighborhoods on intergenerational mobility II: County-level estimates. *Quarterly Journal of Economics* 133:1163–228.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014a. Where is the land of opportunity? The geography of intergenerational mobility in the United States. *Quarterly Journal of Economics* 129:1553–623.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, Emmanuel Saez, and Nicholas Turner. 2014b. Is the United States still a land of opportunity? Recent trends in intergenerational mobility. *American Economic Review* 104:141–47.

Gelman, Andrew, and Jennifer Hill. 2006. *Data analysis using regression and multi-level/hierarchical models*. New York: Cambridge University Press.

Ingels, Steven J., Daniel J. Pratt, Christopher P. Alexander, Donna M. Jewell, Erich Lauff, Tiffany Mattox, and David Wilson. 2014a. Education Longitudinal Study of 2002 third follow-up data file documentation. Washington, DC: National Center for Education Statistics, Institute of Education Sciences, US Department of Education.

Ingels, Steven J., Daniel J. Pratt, Deborah R. Herget, Jill A. Dever, Laura Burns Fritch, Randolph Ottem, James E. Rogers, Sami Kitmitto, and Steve Leinwand. 2014b. High School Longitudinal Study of 2009 (HSLS:09) base year to first follow-up data file documentation. Washington, DC: National Center for Education Statistics, Institute of Education Sciences, US Department of Education.

Jacob, Brian, and Jesse Rothstein. 2016. The measurement of student ability in modern assessment systems. *Journal of Economic Perspectives* 30:85–108.

Kearney, Melissa S., and Phillip B. Levine. 2014. Income inequality and early nonmarital childbearing. *Journal of Human Resources* 49:1–31.

———. 2016. Income inequality, social mobility, and the decision to drop out of high school. *Brookings Papers on Economic Activity*, Spring, 333–80.

Tourangeau, Karen, Christine Nord, Thanh Lê, Alberto G. Sorongon, and Michelle Najarian. 2009. Early Childhood Longitudinal Study, Kindergarten Class of 1998–99 (ECLS-K), combined user's manual for the ECLS-K eighth-grade and K–8 full sample data files and electronic codebooks. Washington, DC: National Center for Education Statistics, Institute of Education Sciences, US Department of Education.

US Bureau of the Census. 2004. Table MS-2. Estimated median age at first marriage, by sex: 1890 to present. https://www.census.gov/population/socdemo/hh-fam/ms2.pdf.