DOES LOCAL TAX FINANCING OF PUBLIC SCHOOLS PERPETUATE INEQUALITY? *

Raj Chetty and John Friedman

Harvard University and NBER

National Tax Association Session: Tax Policy Lessons from Administrative Tax Records

* The data used in this paper were accessed through contract TIRNO-09-R-00007 with the Statistics of Income (SOI) Division at the US Treasury. Jessica Laird and Colin Sullivan provided outstanding research assistance. Financial support from the Lab for Economic Applications and Policy at Harvard and the National Science Foundation is gratefully acknowledged.

I Introduction

Public schools in the United States have traditionally received been funded by local taxes. Although the share of revenues from local sources has fallen from 82% to 44% over the past century, as states have taken up some of the burden from individual districts, local financing remains a crucial source of funding. In particular, property taxes that are set and collected at the school-district level accounted for 34% of funding for public schools in 2007. Richer communities not only have larger property tax bases, but many have higher tax rates. These financing arrangements have led public school expenditures, as well as public school quality, to vary dramatically across cities and states For instance, per-pupil expenditures in Connecticut (the richest state) were nearly than double that in Mississippi (the poorest state) in 2007.

These disparities in school quality are particularly important for public policy because of their potential impact on inequality. If richer families sort into school districts that spend more to produce higher quality education, and quality of education has a causal impact on children's incomes, then local financing of public schools may propagate income inequality. The observation that there may be a link between inequality and local financing of public schools is not new; for instance Benabou (1996) explores the implications of local school finance for inequality and growth, while Loury (1981) directly models the dynamics of income across generations as families invest in the human capital of their children. But the literature has lacked an estimate of the causal impact of school quality on children's incomes and thus has been unable to quantify the potential importance of this mechanism.

In this paper, we use the estimates from the STAR experiment analyzed by Chetty et al. (2010) to overcome this hurdle. We first show that the correlation between parents' and kids' incomes in our sample matches other estimates in the literature at around 0.29. Each $1,000 of parental income increases kids' earnings at age 27 by $110. We then estimate the fraction of this correlation that runs through school quality in two steps. First, we estimate the relationship between parental income and school quality in Kindergarten: for each $1,000 of parental income, school quality increases by 0.74% of a standard deviation. Second, we take results from Chetty et al. (2010) to estimate the causal effect of school quality on adult earnings. Combining these estimates, we conclude that school quality differences could explain as much as 40% of the correlation between parents' and children's incomes in the United States. While this particular number relies on several strong assumptions that we explain below, our analysis suggests that financing schools through property taxes rather than other forms of taxation could have substantial impacts on inequality in the long run. We conclude by discussing several specific tax policy reforms that have been under consideration in the U.S. that could potentially reduce the propagation of inequality through the school system.

II Institutional Background and Data

Institutional Background. We analyze data from Project STAR, the largest and most widely studied education intervention conducted in the United States. STAR was a randomized experiment conducted in Tennessee from 1985 to 1989. STAR randomly assigned 11,571 students across 79 schools to kindergarten classrooms that varied in size, teachers, and other characteristics. For a more detailed summary of the STAR experiment, see Word et al. (1990), Krueger (1999), and Finn et al. (2007).

Variable Definitions and Summary Statistics. We measure adult outcomes of Project STAR participants and characteristics using administrative data from United States tax records. Selected variables from tax records were accessed through a research contract with the Statistics of Income (SOI) division of the Internal Revenue Service in the U.S. Department of the

Treasury. 95.0% of STAR records were linked to the tax data using an algorithm based on standard identifiers. We take two key pieces of data on income from the tax records for the purposes of this analysis: parental household Income and children's Income. We always express all monetary variables in 2009 dollars, adjusting for inflation using the Consumer Price Index.

We define parental household income as average Adjusted Gross Income (winsorized at $252,000, the 99th percentile) from 1996-1998, when the children are 16-18 years old. For parents who do not file, we define household income as zero. For divorced parents, this income measure captures the total resources available to the household claiming the child as a dependent (including any alimony payments), rather than the sum of the individual incomes of the two parents. By this measure, mean parent income is $48,010 (standard deviation of $41,622) for STAR students whom we are able to link to parents. This is less than the average for all parents of similarly aged children in the US. Using a 0.25% random sample from the entire US population of households with children born in 1979-1980, we estimate a national average of $65,660 (standard deviation of $53.844).

The data for children's income come from W-2 forms, yielding information on earnings for both filers and non-filers. We define earnings in each year as the sum of earnings on all W-2 forms filed on an individual's behalf. We windsorize earnings in each year at $100,000 to eliminate outliers; fewer than 1% of individuals in the STAR sample report earnings above $100,000 in a given year. To increase precision, we use average (inflation indexed) earnings from year 2005 to 2007 as an outcome measure. The mean individual earnings for the STAR sample in 2005-2007 (when the STAR students are 25-27 years old) is $15,912 (standard deviation of $15,558). This earnings measure includes zeros for the 13.9% of STAR students who report no income between 2005-07. The mean level of earnings in the STAR sample is lower than in the same cohort in the US population, as expected given that Project STAR targeted more disadvantaged schools. Using the same 0.25% random sample as for parents, we estimate a national average of $20,500 (standard deviation of $19,541) for students in this cohort.

III Do School Quality Differences Propagate Income Inequality?

We begin by exploring the intergenerational correlation of income in the STAR data itself in Figure 1. To create this figure, we grouped the data into 20 equal-sized bins based on parental income, and then plotted average child income within each bin. We also show the best fit line (the solid line), which we fit on the underlying data. Each $1,000 of parental household income predicts an additional $110 , or 0.69%, of children's income.

The correlation between three-year average household income and three-year average child income is 0.29. This number is slightly lower than existing estimates in the literature (Solon (1992), Bjorklund and Jantii (1997), and Solon (1999)), most likely because we measure the income over different time periods for households and children.

There are many channels through which parental income might correlate with children's income. Higher income parents may transfer earnings ability to children through genetics, or through environmental conditions. Richer parents may live in healthier, safer neighborhoods, or may place more emphasis on academic achievement in their children's upbringing. They may also have greater network connections or social capital that generates higher earnings for their children. Finally, their children may attend higher quality schools.

We now proceed to estimate the share of intergenerational transmission that runs through the final channel. We do so in two steps. First, we estimate the causal effect of school quality on students' adult earnings. Second, we estimate the relationship between parental income and the appropriate measure of school quality.

The Effect of School Quality on Student Earnings. There are two ways to measure school quality. The first is to explicitly account for specific covariates, such as teacher experience or peer quality, that vary across schools. This approach depends heavily on accurate and complete measures of all aspects of school quality. The recent education literature (e.g. Hanushek 2003) suggests that this is very difficult, since aspects of school quality, such as the ability of teachers to maintain control of a class, are difficult to observe.

Instead, we follow the modern literature on teacher quality (e.g., Rockoff (2004), Kane and Staiger (2010)) and measure the ex post impact of school quality through test scores. This approach combines the effect of many different aspects of school quality --- both observed and unobserved --- into a single summary statistic. In order to avoid a mechanical correlation between a given student's test scores and adult earnings, we define our measure of quality for each student as the average test scores of her classroom peers. This measure of class quality is measured in the same percentile point units as the test scores. Since the test score is measured at the end of the year, better scoring peers might reflect higher quality peers, but it also could reflect smaller classrooms, better teachers, or any other score-improving aspect of school quality. In order to take greatest advantage of the randomization across classrooms, we use only the first grade in which a student appears in the STAR experiment. 55% of the students first enter in Kindergarten, with 20%, 14% and 11% in grades 1-3. For ease of exposition, we refer to this as "Kindergarten Classroom Quality".

In Chetty et al. (2010), we construct this measure of Kindergarten Class Quality and explore its effect on adult outcomes. The standard deviation of this measure across the STAR sample is

9 percentile points. Importantly for this analysis, we estimate that a 1 percentile point increase in Kindergarten Class Quality increases a students's earnings at age 27 by $50, or 0.3% of the sample average. This effect is constant across different subgroups of the population, including students from richer or poorer households. Therefore, a uniform increase in school quality for all students has no effect on income inequality.

The Effect of Parent Income on School Quality. Project STAR included students from 79 schools across Tennessee. These schools varied considerably in demographics; for instance, the fraction free-lunch students varies across schools from 2% to 99%. 14% of schools are 100% black, while 19% of schools are all white. Although it is difficult to benchmark this amount of variation to a nationally representative sample, the schools in STAR are not concentrated in one particular socioeconomic group. We now investigate the extent to which richer parents live in better quality school districts. Note that this relationship need not be the causal effect of parental income; all that matters is the extent to which higher income parents, for whatever reason, place their children in better schools.

Figure 2 displays the relationship between parental income and school quality, as measured by peer test scores. Each $1,000 of parental income predicts an additional 0.067 percentile points (0.74% of a SD) of Kindergarten Class Quality. The relationship is fairly constant throughout the parental income distribution.

In order to combine these two estimates and calculate the importance of school quality for the intergenerational correlation of income, we must make an important assumption about the precise form in which Kindergarten Class Quality affects student earnings. The estimate from Chetty et al. (2010) is a weighted average combination of the effect of many different aspects of class quality: teachers, class size, peers, and other class-level shocks such as classroom noise.

The variation in quality across schools that correlates with household income, however, may reflect a different aspect of class quality. For instance, suppose that effects in Chetty et al. (2010) are driven primarily by differences in teacher quality. If quality differences across schools are instead driven by the quality of peers, then school quality may not, in fact, be an important channel for intergenerational income correlations. For this exercise, we must assume that there is no such mismatch. Therefore we must make one of two strong assumptions. First, we could assume that all aspects of class quality contribute equally to the effect on earnings. Under this assumption, the precise mix of class-quality drivers in each case would not matter. Alternatively, we may allow for heterogeneous affects from different aspect of class quality, but then we must assume that the variation between classrooms within schools is similar to that between schools. Under either assumption, the estimate from Chetty et al. (2010) is the correct causal estimate of the impact of school quality differences that correlate with household income.

Under either of these assumptions, we can combine our two estimates to calculate the effect of parental income on kids' income through the channel of Kindergarten Class Quality. We estimate that an additional $1,000 of parental income implies an extra $3.34=0.067*$50 of student earnings at age 27. Kindergarten Class Quality alone explains approximately 3% of the correlation between parents' and kids' incomes.

IV School-Quality and Intergenerational Income Inequality

To illustrate the potential magnitude of the school quality channel, we make the following strong assumptions in addition to the previously mentioned assumption about how class quality impacts earnings.

First, we assume that the relationship between household income and school quality is constant across grades. In practice, the relationship between household income and school quality may differ across grades. The smaller size of elementary schools, relative to high schools, might permit greater assortative matching by household income in younger grades. On the other hand, there may be more observable differences across schools in older grades on which parents can select.

Second, we assume that the effect of school quality on adult earnings is constant across grades. Again, the effect of school quality could either increase or decrease as children age in practice.

Finally, we assume that class quality combines linearly across years. Class quality improvements across different grades might combine in many different ways in the education production function. For instance, class quality in different grades could be substitutes, if all a child needs is one good teacher to make a difference. Alternatively, quality in different grades might be complementary, if a particularly high quality experience makes a child more able to take advantage of future high quality classes.

Under these assumptions, the total effect of $1,000 of parental income on kids' income through better schools is 13×$3.34=$43.36, 40% of the $110 impact observed in the raw data. This estimate implies that better quality schooling may account for 40% of the inter-generational correlation between parents' and kids' incomes. That is, the correlation of income across generations would fall from 0.29 to 0.18 if all children had access to schools of the same quality. To help visualize the magnitude of this effect, we add to Figure 1 the counterfactual relationship between parents' and children's incomes removing all differences in school quality (shown by the dashed line).

V Conclusion: Implications for Tax Policy

The assumptions invoked in the preceding calculation are quite strong and thus our quantitative results must be interpreted as largely illustrative. Nevertheless, the calculations show tax policies which generate more uniform school quality across the U.S. could substantially reduce inequality. We conclude by summarizing some specific tax policies that may be well suited to addressing this issue in view of our results.

One approach is to offer tax credits to families investing in primary and secondary education. These policies could take the form of a tax-deferred savings account. One such existing policy is the Coverdell Education Savings Accounts Program. These accounts allow families to save up to $2,000 annually, which may be spent tax-free on qualified educational expenses from primary and secondary schools (in addition to higher education institutions). Take-up of these accounts has been quite low, and high income families are more likely to take advantage of these accounts than poorer families (Dynarski 2004). If properly targeted towards low income families, however, these accounts could be a powerful force towards leveling the school quality gradient. Congress has also considered, though not yet enacted, proposals to provide refundable tax credits for primary and secondary school expenses. For instance, the Hope Plus Scholarship Act of 2003 proposed expanding the Hope Scholarship to apply to primary and secondary school expenses. Both of these policies would not only encourage low-income families to invest in their children's education but also foster competition among schools that might itself raise quality.

Another approach is to use tax credits to invest directly in low income areas. There are a number of place-based federal tax policies in current law. For instance, the New Markets Tax Credit (NMTC) Program was established as part of the Community Renewal Tax Relief Act of 2000 and authorized the Treasury to allocate tax credits to Community Development Entities.

These credits were allocated in four rounds of competitive bidding, with the most recent round in 2010; many of the transactions were used to better the economic standing of distressed communities (Rubin and Stankiewicz 2005). To date, the program has made 495 awards totaling $26 billion. Many banks and financial companies are using these credits to fund charter schools. To the extent that these schools offer children from lower SES backgrounds access to higher quality education, the NMTC program may have substantial impacts on poverty and inequality in the long run. Another example is the Empowerment Zone Program, which designates businesses in high-poverty areas for employment credits and low-cost loan eligibility. This program started in 1994 with 6 zones; the program has since expanded to more than 40 cities nationwide. Similar programs, targeted at charter schools in high-poverty areas, could increase school quality in the lowest performing districts.

The most ambitious approach to tackling the problems raised in this study is to change the tax base used to finance primary education in the United States completely. For instance, one might consider financing schools via state-level income taxes rather than property taxes and having states allocate money uniformly across school districts. While such reforms may raise many other challenges in implementation, our results suggest that such systemic changes are likely to have the largest impacts on inequality in the long run.

There are a number of important questions that remain for future research in order to evaluate these tax policies. First, researchers should seek to relax the assumptions we make above and obtain a more precise understanding of how property tax finance of schools affects inequality. Second, these estimates rely on a small number of grades and students; future work should explore important dimensions of potential heterogeneity in the long-run causal effects of education. For instance, at what ages are educational interventions most fruitful and which types of tax credits would be most effective? Among the many potential aspects of education at which tax policies could be targeted --- better teachers, better peers, smaller classes, technology in the classrooms --- which have the highest long-run return? Finally, which tax policies (e.g. tax credits for charter schools or changes in property tax finance) are most effective at increasing the long-run outcomes of students from low socioeconomic status backgrounds?

References

Benabou, Roland. Heterogeneity, Stratification, and Growth: Macroeconomic Implications of Community Structure and School Finance. American Economic Review 86 (June 1996):584-609.

Bjorklund, Anders, and Markus Jantii. Intergenerational Income Mobility in Sweden Compared with the United States. The American Economic Review 87 (December 1997): 1009-1018.

Chetty, Raj, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, Danny Yagan. How Does Your Kindergarten Classroom Affect Your Earnings? Evidence From Project Star. Forthcoming, Quarterly Journal of Economics.

Dynarski, Susan. Who Benefits from the College Saving Incentives? Income, Educational Expectations and the Value of the 529 and the Coverdell. National Tax Journal 57 (June 2004): 359-383.

Finn, Jeremy D., Jayne Boyd-Zaharias, Reva M. Fish, and Susan B. Gerber, Project STAR and Beyond: Database User's Guide. Lebanon: Heros, inc., 2007.

Hanushek, Eric A. The Failure of Input-Based Schooling Policies. Economic Journal 113 (February 2003): F64-F98.

Kane, Thomas, and Douglas O. Staiger. Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation. Cambridge, MA: National Bureau of Economic Research, 2008. NBER Working Paper W14607.

Krueger, Alan B. Experimental Estimates of Education Production Functions. Quarterly Journal of Economics 114 (May 1999): 497-532.

Loury, Glen C. Intergenerational Transfers and the Distribution of Earnings. Econometrica 49 (July 1981): 843-867.

Rockoff, Jonah E. The Impact of Individual Teachers on Student Achievement: Evidence from Panel Data. American Economics Review 94 (May 2004): 247-252.

Rubin, Julia Sass, and Gregory M. Stankiewicz. The New Markets Tax Credit Program: A Midcourse Assessment. Community Development Investment Review, San Francisco Federal Reserve (2005, Issue 1): 1-11.

Solon, Gary. Intergenerational Income Mobility in the United States. The American Economic Review 82 (June 1992): 393-408

Word, Elizabeth, John Johnston, Helen P. Bain, D. Dewayne Fulton, Charles M. Achilles, Martha N. Lintz, John Folger, and Carolyn Breda. The State of Tennessee's Student/Teacher

Achievement Ratios (STAR) Project: Technical Report 1985-1990. Tennessee State Department of Education, 1990.