NEW EVIDENCE ON THE LONG-TERM IMPACTS OF TAX CREDITS[1]

Raj Chetty and *John N. Friedman,* Harvard University and NBER
Jonah Rockoff, Columbia University and NBER

INTRODUCTION

Many tax policy provisions direct subsidies to low income families with children under the belief that such redistribution not only helps these households' immediate circumstances but also provides for better opportunities for children. How large are the long-term impacts of tax policies? In this paper, we provide some suggestive evidence on this issue by focusing on one particular channel through which taxes might have long-term impacts: education. While taxes could have long-term impacts in many ways, education is thought to be a particularly important pathway for such effects. One practical advantage of studying the education channel is that test scores provide excellent short-term metrics of progress for nearly all children. Furthermore, previous research suggests that test scores provide a good proxy for the long-term outcomes of children as young as five (e.g. Chetty et al., 2011; Heckman et al., 2010). If cash transfers increase test scores, and those increases have a causal impact on adult outcomes, then examining the impact of tax credits on children's test scores provides an easy way to examine the long-run impact of income on children without needing to observe children for many years into adulthood.

Two recent papers have examined the short-term impacts of tax credits on test scores, but these papers have reached conflicting conclusions. Dahl and Lochner find large but imprecisely estimated effects on children's test scores from the increase in the Earned Income Tax Credit between 1994 and 1996. In contrast, Jacob and Ludwig (2007) find a precisely estimated zero effect from the outcome of a randomized housing subsidy lottery in Chicago. Neither of these studies has attempted to quantify the impact of test score improvements on long-term outcomes such as earnings, which is our focus here.

We analyze the impacts of taxes on educational achievement and earnings by combining two datasets to form a large sample linking student educational records with family background and income. The first data set contains administrative data from a large urban school district. These data include information on the test scores of children in grades three through eight. The data also include a rich set of individual characteristics, including age, gender, race, ethnic background, and English proficiency. The second dataset includes selected data from U.S. tax records for all families in the school district sample. These data provide precise information on the eligibility of families for various federal and state credits such as the Earned Income Tax Credit (EITC) and Child Tax Credit (CTC).

The ideal empirical setup for this question would link quasi-experimental variation in the receipt of tax credits across families directly with long-term data following those children into adulthood. Unfortunately, the electronic U.S. tax records do not cover enough years to perform this analysis. Instead, we conduct the analysis of the long-term impacts of tax credits in two parts. We identify the impact of cash transfers on children's test scores. We first exploit the nonlinearities in the transfer schedule present in the EITC and CTC. By controlling for a flexible function between family income and educational achievement that is smooth across all income ranges, we isolate the nonlinear variation in income from tax credits.

Using the variation from nonlinearities, we estimate that a $1000 tax credit increases student test scores by 6-9 percent of a standard deviation. These effects are larger in math (9.3 percent) than in reading (6.2 percent) and are larger for students in middle school (8.5 percent) than in elementary school (7.3 percent). Because this approach rests on relatively strong identification assumptions, the estimates in this paper should be viewed as only illustrative. Further research using more robust non-parametric identification strategies is needed before one can draw strong policy lessons from this analysis.

Combining our estimates of the impacts of tax credits on scores and estimate of the impacts of test scores on earnings from other work (Chetty et al. 2011b), we find that each dollar of income through tax credits increases NPV earnings by more than one dollar. These results suggest that a substantial fraction of the cost of tax credits may be offset by earnings gains in the long run. Hence, when analyzing the costs and benefits of policies such as the Earned Income or Child Tax Credit, policy makers should carefully consider the potential impacts of these programs on future generations.

The remainder of this paper is organized as follows. In Section 2, we describe the specifics of the tax policies we examine and the data used in our analysis. Section 3 provides our main findings regarding the impact of income transfers on student test scores. Section 4 combines these estimates with those from Chetty et al. (2011b) to calculate the impact of cash transfers on students' long-run outcomes.

DATA AND INSTITUTIONAL BACKGROUND

Data

We draw information from two administrative databases: students' school district records and information on their parents from U.S. tax records.

We obtain information on students, including enrollment history, test scores, and teacher assignments from the administrative records of a large urban school district. These data span the school years 1988-1989 through 2008-2009 and cover roughly 2.5 million children in grades 3-8.[2] Variables from this data set include test scores for approximately 18 million student-subject-year observations for grades three through eight. The data set also contains information on ethnicity, gender, age, receipt of special education services, and limited English proficiency for the school years 1989 through 2009. The mean age at which students are observed is 11.6. Within our analysis sample, 3 percent of students receive special education services, while 10 percent have limited English proficiency. Roughly 80 percent of students are eligible for free or reduced price lunches.

We obtain data on students' household characteristics from income tax returns (e.g., form 1040). We linked the students to tax records using an algorithm based on standard identifiers (date of birth, state of birth, gender, and names) that is described in Chetty et al. (2011a). We then find households that claimed the students as dependents in the years for which we have school data. Variables for this analysis from the tax data include household characteristics for the students, which allow us to determine household eligibility for the EITC and CTC. A detailed description of the data set and variables is given in Chetty et al. (2011), who use these data to study the long-term impacts of Project STAR

The linked analysis dataset has one row per student per subject (math or English) per school year. It contains 5.98 million student-year-subject observations and roughly 5.31 million test scores; 4.80 million observations have teacher links; 89.2 percent of the observations in the school district data are matched to the tax data. The match rate is uncorrelated with teacher assignment, suggesting that the small degree of attrition is unlikely to produce significant bias.

Earned Income Tax Credit

The Earned Income Tax Credit (EITC) is a refundable credit paid to households with positive income. The defining feature of the EITC is the pyramid-shaped schedule of the credit, displayed in figure 1a.[3] Households receive an income subsidy for earnings up to a certain threshold. For all years in our sample, households with one dependent receive a 34 percent credit up to $8,970 of income for a maximum credit of $3,050.[4] Households with two children or more receive a 40 percent credit up to $12,590 for a maximum credit of $5,036.[5] The credit is then phased out once households earn more than $16,690; the phase-out rate is 16 percent for households with one child, and 21 percent for households with two or more children.[6] Families in our sample earned an average of $1,520 from the EITC. Note that this figure excludes households that do not qualify for the EITC. Approximately 65 percent of families in our sample qualify for the credit.

Child Tax Credit

For families with earnings below an income threshold, the Child Tax Credit (CTC) provides a partially refundable credit for each eligible dependent. Figure 1b depicts the credit schedule for a single filer for 2001-2008. The size of the basic credit is constant below the income threshold; after passing the threshold, the phase-out rate is 5 percent. The income threshold is $75,000 for singles and $110,000 for married households filing jointly. The CTC offered $400 per child (up to two) in 1998 (the first year of the credit), $500 in 1999-2001, and $1,000 from 2002 on.

Before 2001, the CTC was nonrefundable. Since many low-income families owe no income tax, they could not benefit from the CTC. Beginning in 2001, the CTC became partially refundable,

Figure 1: **EITC and CTC Schedules**

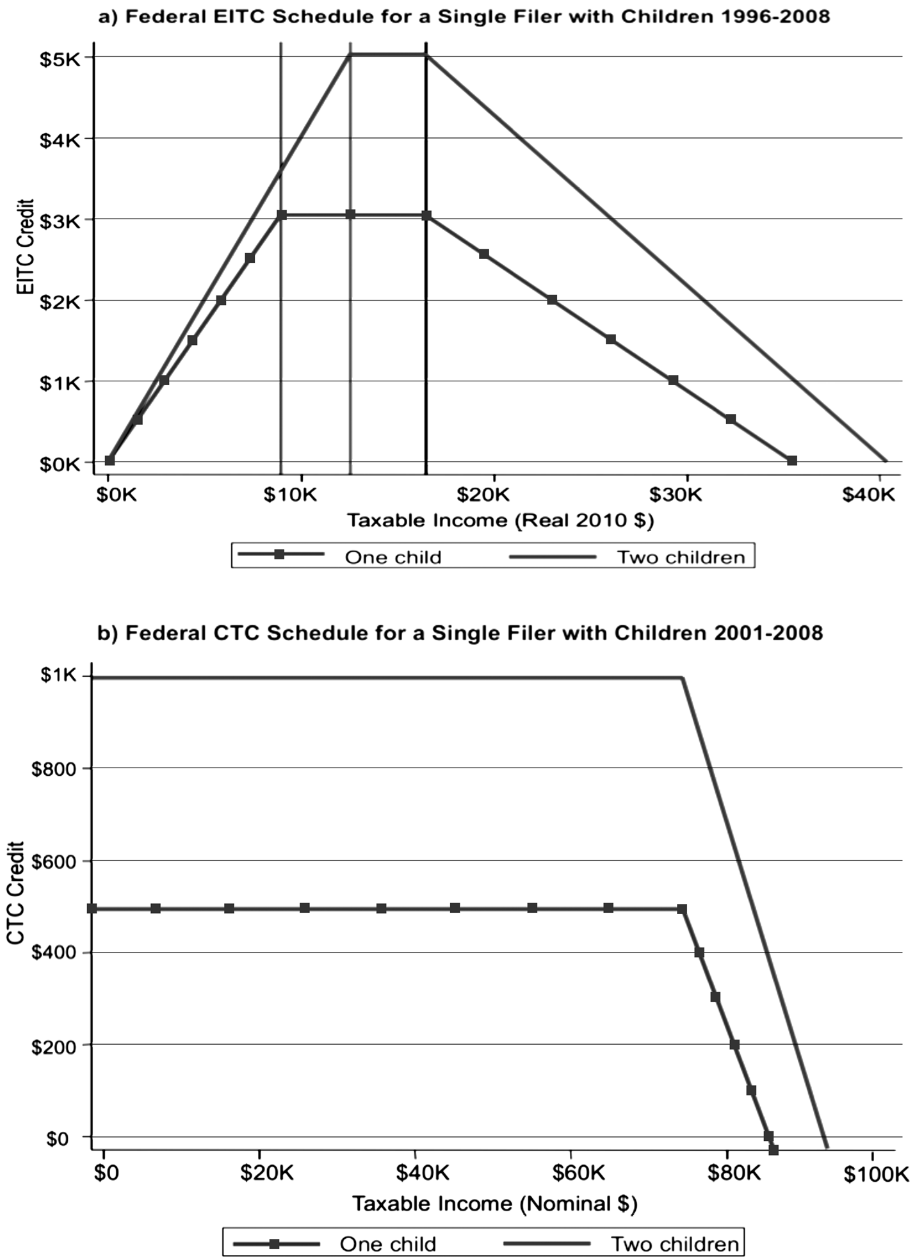

Notes: Panel A displays the Earned Income Tax Credit schedule for 2008 for households filing as head of household with children. Panel B displays the maximum Child Tax Credit for which a household is eligible. In both panels, the green line shows the schedule for those with one dependent; the red line shows the schedule for those with two dependents. All monetary values are in 2010 dollars.

where the newly refundable portion was called the Additional Child Tax Credit. Households were able to claim a refundable credit up to 15 percent of their income above an income threshold of $12,050.[7]

On average in our sample, families qualify for $606 from the Child Tax Credit, which is the nonrefundable portion of the credit, plus an additional $537 from the Additional Child Tax Credit. Approximately 82 percent of families qualified for either the Child Tax Credit or the Additional Child Tax Credit. Combining both credits, we find that families received an average of $1,652 from the EITC and CTC combined. This represents approximately 60 percent of total tax credits.

ESTIMATES OF THE IMPACT OF TAX CREDITS ON STUDENT ACHIEVEMENT

Estimating Equation and Identification Assumptions

Both the EITC and CTC have highly nonlinear schedules. In contrast, other determinants of a child's achievement change more smoothly throughout the income distribution. Our basic estimating equation is therefore

$$\textbf{(1)} \quad A_{ift}=\alpha+\phi(AGI_{ft}) + \beta * CREDIT_{ft} + \gamma X_{ift} + \varepsilon_{ift}$$

for student i in family f in year t, where A_{ift} is achievement on the standardized test at the end of the year, $\phi(\cdot)$ is a smooth function of family *AGI*, *CREDIT* is the combined EITC and simulated CTC payments to family f in year t, and X_{ift} is a vector of individual and family characteristics.[8]

Our key identification assumption is that the smooth function $\phi(\cdot)$ captures the entire relationship between simultaneous parental income and achievement other than that driven through the receipt of federal credits. In practice, the EITC provides most of the identification in our study. The key identification question may therefore be restated as: Do children of families earning between roughly $10,000 and $30,000 in AGI over perform in school, relative to the trend determined by their higher and lower scoring peers? Although we believe this assumption is plausible, we cannot be fully confident that it holds in practice. Hence, our estimates should not be viewed as definitive measures of the impacts of tax policies on test scores but rather as a suggestive indication of the magnitude of effects one might expect.

Graphical Evidence

We begin by plotting the cross-sectional patterns of the two key variables: household income and student achievement. Figure 1 plots two series: average scores, as a function of contemporaneous household income, and simulated credits as a function of AGI. Overall, scores are sharply increasing with household income, with an average slope of approximately 0.01. This implies that each $10,000 of income increases scores by roughly 0.1 SDs. The relationship between scores and income is generally quite smooth and slightly concave, except between $10,000 and $30,000.

In the second series, the shape of the EITC is clear at the lower end of the income distribution, where the simulated credit first rises and then sharply falls as a function of AGI. Above about $40,000, the simulated credit flattens, reflecting the constant credit available in this income range through the CTC. The credit then falls gradually again once incomes begin to rise above the threshold for the CTC. This figure also makes clear that the main identification in this paper comes from the EITC, rather than the CTC. The EITC appears as a dramatic increase and decrease of available credit, while the CTC appears as a slight decline in the credit. The reason for this difference is apparent in the phase-in and phase-out rates present in each program. The CTC is a constant credit with a 5 percent phase-out rate at the end. In contrast, the EITC provides phase-in and phase-out rates that are several times higher. As a result, the marginal effects of the CTC are simply too small to be noticeable on the necessary scale.

Above $40,000, both series are smooth and roughly linear. Below there, however, each series deviates from the otherwise smooth pattern. It is in this range that the EITC more than doubles the credit available to households with children. And it is also in this range that children appear to over-perform significantly relative to the income-achievement gradient established in the rest of the figure. Furthermore, the break in linearity in each figure occurs at the same place. Just as the simulated credit available through the EITC begins to increase, student achievement turns up from projected path, suggesting that the change in achievement may be due to the EITC. We now proceed to explore the relationship more formally.

Figure 2: **Average Score and Total Credits vs. AGI**

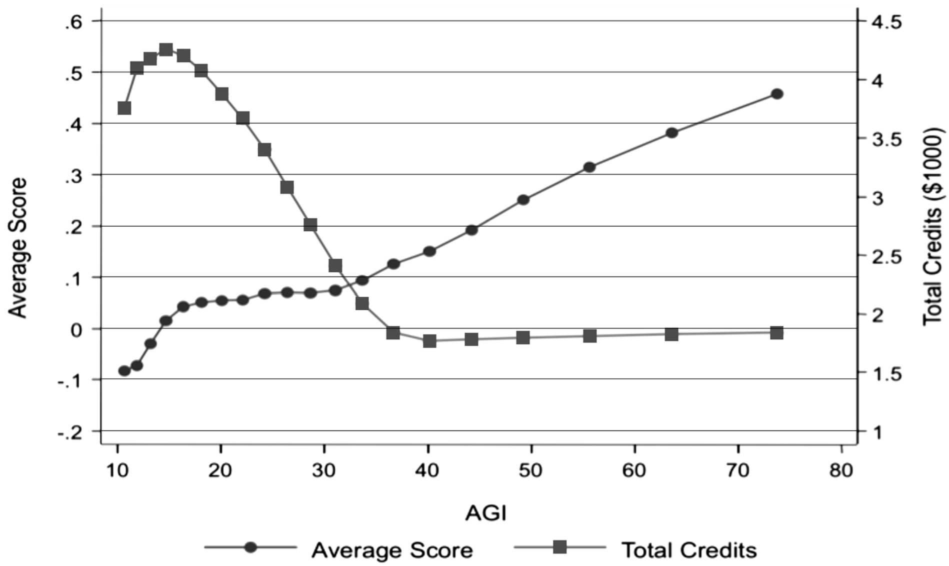

Notes: This figure plots the average test score and credit amount for each bin of household AGI. To construct this figure, we group student-year observations into 20 equal-sized (5 percentile point) bins and plot the mean test score for each bin. All monetary values are expressed in real 2010 dollars.

Regression Estimates

We estimate equation (1) in table 1 panel A. Column 1 estimates the most parsimonious specification, including only a linear control for household AGI. The coefficient is 0.075 SD, and is highly significant with a standard error of 0.002, implying a t-stat of approximately 37. Column 2 increases the flexibility of the AGI control function. Now using a cubic function of AGI as the function $\phi(\cdot)$ in equation (1) we estimate exactly the same coefficient and standard error.[9] Intuitively, the best-fit curve for the relationship between income and score is highly linear, even when the regression allows for more flexibility. Therefore the linear and cubic specifications yield nearly identical results.

Column 3 presents our primary specification. In it we control for a quintic polynomial of AGI, as well as the vector of individual characteristics described above. These additional controls increase the coefficient of interest slightly to 0.08 SD. The coefficient increases because individuals who receive lower amounts of credits not only have more household income, but also tend to be from households with married parents and mothers who gave birth at a later age. Each of these additional characteristics also predicts higher test scores; when controlling for them, tax credits appear to have an even larger impact on achievement.

Figure 3 represents the regression in column 3 of table 1 in scatterplot form.[10] Intuitively, figure 3 presents a non-parametric version of the key regression coefficient in column 3. The linear fit appears approximately correct, and the relationship is not driven by outliers in either direction. Our estimated coefficient of 0.08 is large when compared with the cross-sectional impact of income, though consistent with the estimates in Dahl and Lochner (forthcoming). It is also worth reiterating the strong assumption on the cross-sectional pattern of test scores and household income on which our identification strategy depends. This relationship must hold constant across high and low-income households in order for our identification method to be valid, and thus these results should be interpreted with caution.

Columns 4 and 5 repeat the specification in column 3, separating out math and reading tests. The results suggest that income from tax credits has a larger impact on math scores than reading

Table 1
Estimates of the Effect of Tax Credits on Test Scores

Dependent Variable: Test Score

Panel A: Main Results

	Linear AGI	*Cubic AGI*	*Full Controls*	*Math*	*Reading*
	(1)	*(2)*	*(3)*	*(4)*	*(5)*
Simulated Credits	0.075	0.075	0.080	0.093	0.062
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Observations	3,006,098	3,006,098	3,006,098	1,533,339	1,472,759

Panel B: Heterogeneity by Grade

	Grade					
	3	4	5	6	7	8
Simulated Credits	0.073	0.069	0.078	0.082	0.085	0.090
	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.006)
Observations	638,903	623,058	555,894	476,987	396,939	314,280

Panel C: Heterogeneity over Time

	Year					
	2003	2004	2005	2006	2007	2008
Simulated Credits	0.037	0.051	0.066	0.097	0.100	0.104
	(0.005)	(0.005)	(0.005)	(0.004)	(0.004)	(0.005)
Observations	249,290	330,098	397,792	461,961	480,449	467,816

Notes: *Panel A:* Each column reports coefficients from an OLS regression, with standard errors in parentheses. In all columns, the dependent variable is student test score, normalized within grade-year cell. Columns 1-3 include the entire sample and differ only in the control variables included. Column 1 includes only a linear control for household AGI. Column 2 includes only a quadratic polynomial control for household AGI. Column 3 includes a five-degree polynomial in household AGI, as well as all school and household control characteristics. Column 4 replicates the regression in Column 3 using only math scores. Column 5 replicates the regression in column 3 using only reading scores. *Panel B and Panel C:* Each column reports coefficients from an OLS regression, with standard errors in parentheses. The columns replicate column 3 from panel A for each grade separately. See the notes to panel A for details.

scores. Each $1,000 in income generates a 0.093 SD increase in math scores, but only a 0.062 SD increase in reading scores. This could be because math scores are the best proxy for underlying academic achievement, or because reading scores are a particularly bad proxy for achievement in a low-income population where English is often a second language.

Table 1 panel B investigates heterogeneity of these effects across grades. Each column replicates the specification in column 3 of table 1 panel A, restricting to a single grade. We find that the impact of federal credits increases in later grades, though the effect is non-monotonic. This finding suggests that, as students age, they are better able to take advantage of the benefits afforded through income transfers.

Table 1 panel C presents an analysis of heterogeneity across years. We find that the impact of federal tax credits on achievement increases sharply from 2003, when the effect estimate is only 0.037, through 2006 when we estimate the effect at 0.097. The coefficients then level off through 2007 and 2008.

CONCLUSION: LONG-TERM EFFECTS OF TAX CREDITS

We now combine the evidence from our preceding analysis to answer our original question: what are the long-term impacts of tax credits on earnings? Our estimates from the first part of our analysis imply that a $1000 tax credit increases a child's test score by 6 percent of a standard

Figure 3: **Average Test Score vs. Credit**

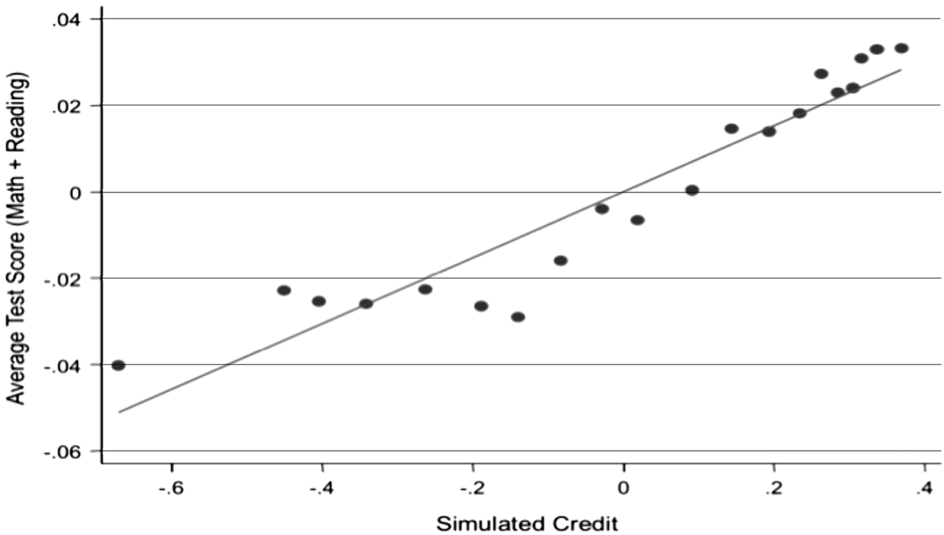

Notes: This figure represents the data underlying the regression in column 3 of table 1. We regress both average test score and household credit amount on the full vector of control variables, including a five-degree polynomial in household income, student English proficiency, receipt of special education status, age, and gender, as well as the household background characteristics including a dummy variable for married filing status, a dummy variable for the difference between the age of the claiming parent and dependent less than 20 years, a dummy variable for home ownership, and average savings in tax-deferred account. We add back the sample means of each variable, and then group observations into 20 equal-sized (5 percentile point) bins and plot the mean residual test score in each bin. The best-fit line is that fitted on the underlying individual data.

deviation (taking the more conservative estimate for reading scores from table 1 panel A). This is a relatively large effect; for comparison, the standard deviation of teacher impacts on achievement is approximately 10 percent of a standard deviation, which is similar to the estimates of Dahl and Lochner (forthcoming).

These score gains themselves have no direct economic interpretation, as we do not know how test score gains translate into earnings gains. Ideally, we would directly analyze the long-term impacts of the EITC or CTC on children's future earnings, but our data do not cover a long enough time period to permit such an analysis. As a feasible alternative, Chetty et al. (2011b) evaluated the effects of a different intervention – better teachers — to understand how test score gains translate into earnings gains. Under the assumption that score increases generated from these different policies have the same long-run effects, the research in that paper allows us to achieve the original objective of understanding the long-run impacts of cash grants through tax policy on children's long-run outcomes.

Combining the results from the two papers suggests that a tax credit to families with young children generates a significant dollar earnings gain over a student's lifetime. We estimate that a 1 SD increase in test scores raises earnings by approximately 9 percentage points. Hence, a \$1,000 tax credit would raise a child's lifetime earnings by 0.09 x 0.05 = 0.54 percentage points. The dollar gains in lifetime earnings are of the same order of magnitude as the cost of the tax credit because a small percentage increase in earnings over a lifetime adds up to a large sum in present value.

Taken at face value, these findings imply that there are substantial returns to public policies that help poor families with children. Consider, for instance, the expansion of the EITC in 2009 to pay an additional credit to families with 3 children. This policy was passed in 2009 as part of the American Recovery and Reinvestment Act for two years, and the 2011 budget made this change permanent. These results suggest that this reform may have increased the NPV earnings of children of these families by more than 5 percent.

Although this analysis has used the federal EITC for identification, the findings apply equally to state EITC programs. Many states offer an EITC that is defined as a percentage of the federal credit; in 2009, this percentage varies from 0 percent in 28 states to as much as 43 percent (for families with 3 children in Wisconsin). But these programs have come under pressure as states grapple with declining revenue streams. For instance, Michigan planned to increase the state EITC from 10 percent to 20 percent in 2009, but the legislature froze the credit at 10 percent of the federal EITC. This freeze saved the state about $100 million but also deprived children from poor families of more than $100 million in NPV earnings. These gains, which come far in the future, are often difficult to represent vividly in public debates.

Many states also wrestle with the choice to make state EITC credits refundable. Of the earned income credits in the 22 states plus the District of Columbia, all but four are fully refundable. Credits that are not refundable hit those families with the lowest income, and therefore the least tax to offset. The impacts on long-term outcomes are likely highest among these poorest families because of credit constraints.

There are many caveats that one must keep in mind when interpreting these results. First, our estimates of the impacts of tax credits on test scores rely on very strong assumptions on the cross-sectional pattern of test scores and household income. This relationship must hold constant across high and low-income households in order for our identification strategy to be valid. Second, the long-run impact of test score increases from different sources may vary considerably. In the extreme, increase student test scores through cheating should not have any long-run impact (and might even have a negative one). In this case, we must assume that the long-run effect of a higher score that comes from a better teacher is the same as that from an increase in tax credits for a child's household. Finally, our results do not shed light on the mechanism through which an increase in tax credits aids student achievement. Because of these limitations, policy makers should exercise caution when extrapolating evidence from the tax credits we have studied to predict the likely impact of future credits. The most important lesson of our analysis is that tax policy could have substantial long-term impacts and that future research should focus on analyzing these issues further.

Notes

1 The tax data were accessed through contract TIRNO-09-R-00007 with the Statistics of Income (SOI) Division at the U.S. Internal Revenue Service. Sarah Griffis, Jessica Laird, and Heather Sarsons provided outstanding research assistance. Financial support from the Lab for Economic Applications and Policy at Harvard and the National Science Foundation is gratefully acknowledged.

2 For simplicity, we refer below to school years by the year in which the spring term occurs, e.g., the school year 1988-89 is 1989.

3 The EITC schedule differs by filing status. Figure 1a shows the schedule for single filers.

4 All figures are quoted in 2010 dollars.

5 In tax years after 2008, the EITC included additional payments to households claiming three or more children; since our sample period ends in 2008, though, this recent reform is not relevant for our analysis.

6 Eligible dependents must live in the household for more than six months during a given tax year, and must remain either under 19 or full-time students under 24 for the entire tax year. Beginning in 2002, Congress lengthened the "plateau" range, while leaving the phase-out rate unchanged. In the first three years after the reform, the phase-out begins for married households at $17,690; in 2005-2007, the phase-out period begins at $18,690; and in 2008, the phase-out begins at $19,690.

7 For example, consider a family with two children and $22,050 of taxable income that owed $300 in tax payments (after the EITC). Under the original CTC, this family would claim $300 in offsetting credit but could not claim more. Under the Additional CTC, the family could claim an additional amount equal to 0.15 × ($22,050 - $12,050) = $1,500. The family would thus receive a CTC equal to $1800 in total.

8 These controls include English proficiency, receipt of special education status, age, and gender, as well as the household background characteristics including a dummy variable for married filing status, a dummy variable for the difference between the age of the claiming parent and dependent less than 20 years, a dummy variable for home ownership, and average savings in tax-deferred account. In practice, we use a five-degree polynomial to estimate the smooth function $\phi(\cdot)$. We have also run similar specifications using higher-order polynomials, as well as smoothed splines (i.e. splines that have continuous derivatives at knot-points), and the results are unaffected.

9 The standard error is slightly lower than before, though it is the same when rounded for significant digits.

[10] We regress both achievement and simulated credit on the polynomial in AGI and other controls, and then take residuals. We then group observations into 20 bins based on the size of the tax credit residual and plot the mean achievement for students in each bin.

References

Chetty, Raj, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan. How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project STAR. *Quarterly Journal of Economics* CXXVI 4 (2011): 1593-1660.

Chetty, Raj, John N. Friedman and Jonah E. Rockoff. New Evidence on the Long-Term Impacts of Tax Credits. IRS White Paper, 2011a.

The Long-Term Impacts of Teachers: Teacher Value-Added and Student Outcomes in Adulthood. NBER WP 17699, 2011b.

Dahl, Gordon B. and Lance Lochner (Forthcoming). The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit. *American Economic Review*.

Heckman, James, Lena Malofeeva Rodrigo Pinto, and Peter Savelyev. Understanding the Mechanisms Through Which an Influential Early Childhood Program Boosted Adult Outcomes. Unpublished manuscript, University of Chicago, 2010.

Jacob, Brian A. and Jens Ludwig. The Effects of Housing Vouchers on Children's Outcomes. University of Chicago Working Paper, 2001.