Borrowing Constraints on Families with Young Children[1]

Elizabeth M. Caucutt
Lance Lochner

Raw differences in educational attainment by family income are dramatic. For example, Ellwood and Kane (2000) find that among the high school graduating class of 1992, those with parents in the top income quartile have a 30 percentage point higher probability of attending a post-secondary institution than those from the bottom income quartile. However, most of this gap disappears after controlling for adolescent achievement or cognitive test scores and family background (Cameron and Heckman 1998, 2001; Carneiro and Heckman 2002; Ellwood and Kane 2000). This has led many economists to conclude that short-term borrowing constraints at college ages are not an important determinant of college attendance and completion decisions. That is, increasing credit access for college students is unlikely to affect college enrollment rates by more than a few percentage points and will do little to reduce the attendance gaps by family income or race. By the end of high school, it may be too late to help many youths from disadvantaged backgrounds. It appears that, because they are ill-prepared to do so, poor and minority youths choose to attend college at lower rates than wealthier white youths.[2]

Indeed, there is a wealth of indirect evidence that skill investments are complementary over the life cycle and that a failure to make adequate investments during early childhood years reduces incentives in later years to invest through college attendance or high school completion (Cunha et al., forthcoming). The importance of cognitive-achievement scores in determining high-school-dropout or college-attendance rates is one indicator of the complementarity of early and late investments. Using data on children from the National Longitudinal Survey of Youth (NLSY), figures 1 and 2 show the importance of adolescent achievement scores, as measured by a combined math and reading Peabody Individual Achievement Test (PIAT) score at ages 13–14, in determining high-school-dropout and college-attendance rates. After controlling for family background and family income (measured when children were ages 15–18), a one-standard-deviation increase in PIAT scores at ages 13–14 is associated with a 12 percentage point decline in high school dropout rates and a 15 percentage point increase in college attendance.

This study investigates the role of family income and borrowing constraints in determining early investments in children and youth achievement scores. As figure 3 shows, youths raised in families in the bottom third of the income distribution are much less likely to be among the highest PIAT test scorers (at ages 13–14) than are those in middle- and high-income groups.[3] While more than 50 percent of all 13- to 14-year-olds in the top tercile of the income distribution are in the top third of the test-score distribution, fewer than 20 percent of those in the bottom income tercile managed such scores. These findings raise the natural question: To what extent do family borrowing constraints during early childhood and adolescence influence early investments in children, cognitive achievement levels, and ultimately college attendance and completion?

While a number of studies have recently examined the effects of credit constraints on college-going behavior (see Carniero and Heckman 2002 for a summary of the empirical literature), very little attention has been paid to the role of borrowing constraints when children are younger.[4] Yet it seems possible that constraints at early ages play a more important role in determining investment decisions for a number of reasons. First, most empirical studies indicate that early investments in children produce high long-term payoffs (see Karoly et al., 1998 or Blau and Currie, forthcoming, and references

Figure 1: Fraction of Youth Who Drop Out of High School by Adolescent Achievement Scores

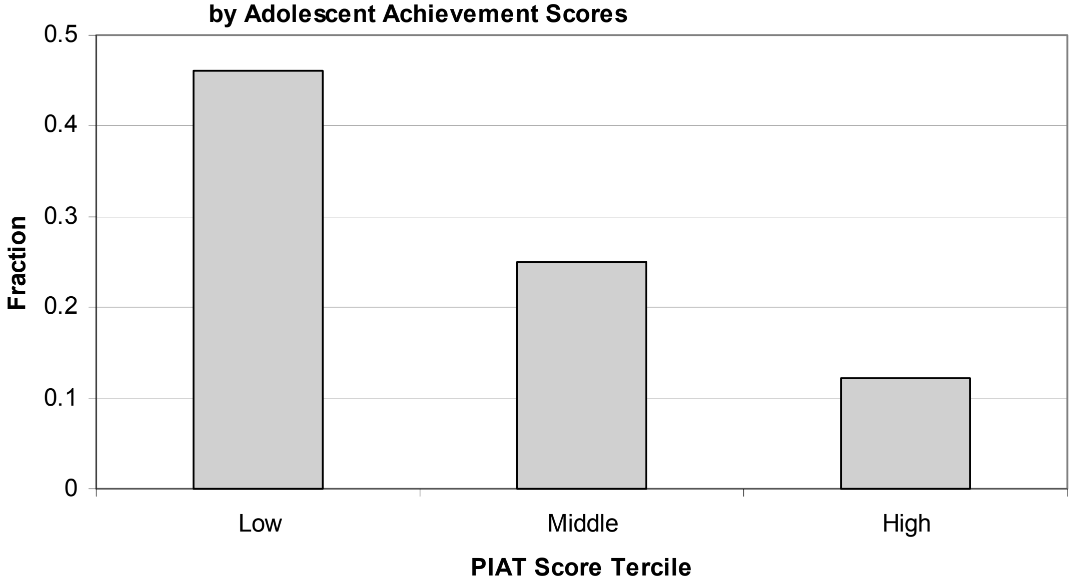

Figure 2: Fraction of Youth Who Attend College by Adolescent Achievement Scores

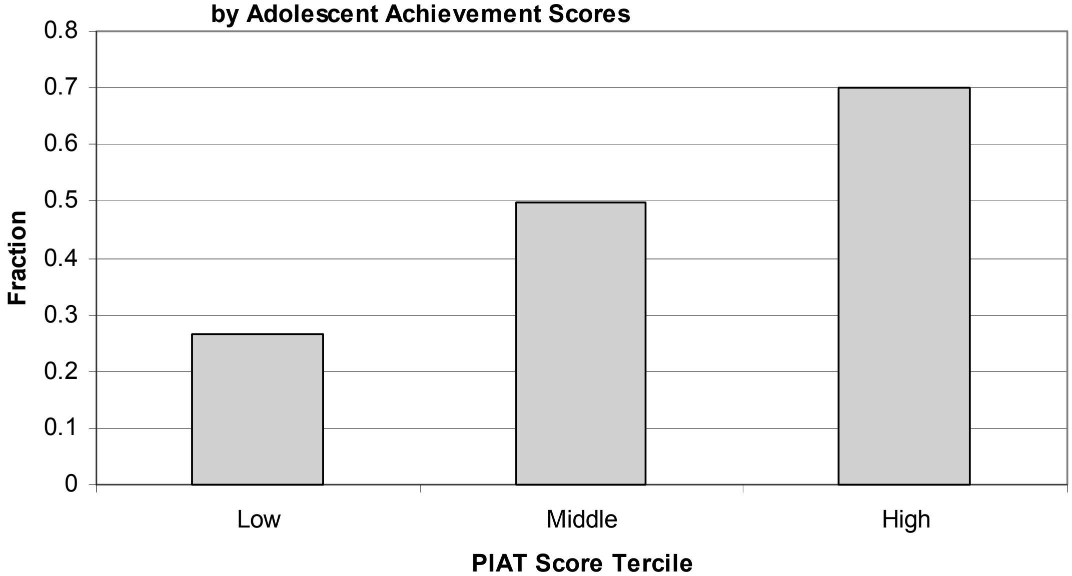

therein). Randomized studies of early intervention and preschool programs for disadvantaged children have estimated large benefit–cost ratios, suggesting that many poor families are not making investments in their children even though those investments would more than pay for themselves in the long run. Not only are children able to learn quickly when they are young, but early learning begets later learning as emphasized by Cunha et al., forthcoming. Second, recent studies have shown that increases in family income lead to increases in the test scores of adolescents and young children (e.g., Blau 1999; Duncan and Brooks-Gunn 1997; Levy and Duncan 1999; and Dahl and Lochner 2005). A few of these studies suggest that family income is more important at earlier ages (e.g., Duncan and Brooks-Gunn 1997; Levy and Duncan 1999). Third, parents age with their children and, as both children and parents age, parental resources tend to increase with the accumulation of human capital and the associated rise

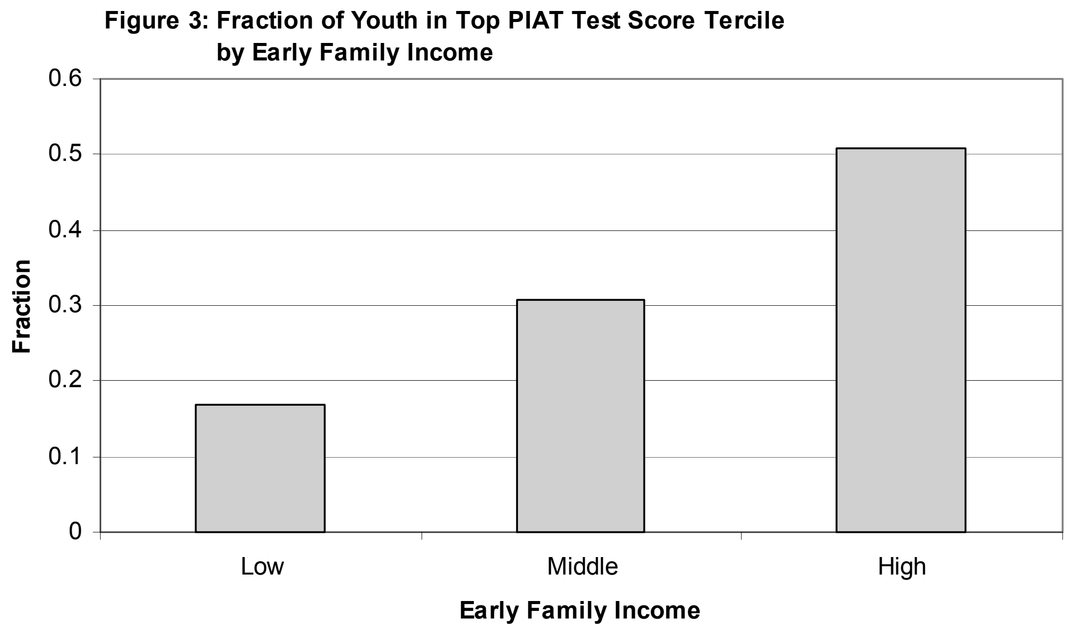

in earnings. Fourth, despite generous government student-loan programs for college-age students and their families in the United States and other developed countries, governments have not traditionally offered loans to parents of young children to help finance earlier human-capital investments.

While the direct costs of public elementary and secondary education are fully subsidized, a good education through high school is not free. In many U.S. communities, parents must choose between sending their child to poor public schools and paying for their child to attend better private schools. Alternatively, parents may choose between high-cost communities with good public schools and low-cost neighborhoods with poor ones. Other investments in young children can also be costly. Preschool programs in the United States can cost as much as attendance at a top university. While the government does not fully neglect poor preschool-age children, the quality of publicly provided preschool programs (e.g., Head Start) is far below what it could be (Zigler 1994; Blau and Currie, forthcoming). Expenditures on computers and books also add up. Finally, parental time is an important, yet costly, input that poor parents may be unable to afford.

In order to better understand the role of borrowing constraints, we distinguish between *intergenerational* borrowing constraints, which would prevent parents from borrowing against their children's future earnings, and *intragenerational* borrowing constraints, which would prevent individuals from borrowing against their own future income. Intuitively, the former implies that only the present value of lifetime parental income affects child success, while the latter implies that the timing of parental income (over the lifecycle) also matters. We use data on children from the NLSY to test for the latter form of constraint, estimating the effects of family income at different ages of the child on the math and reading achievement of those children at ages 5–14. We interpret evidence that the timing of income matters as evidence that *intragenerational* borrowing constraints distort investment decisions.[5]

INTERGENERATIONAL AND INTRAGENERATIONAL BORROWING CONSTRAINTS

Parents may be *intergenerationally* constrained from borrowing against their children's future income; that is, parents may not be able to pass on debts to their children. This is more like a generational budget constraint than a borrowing constraint, but it does imply that family income may affect child achievement. Importantly, this type of constraint, if binding, suggests that the present discounted value of lifetime family income will be an important determinant of child achievement. However, if parents can save and borrow against their own future earnings, the timing of their income should not be important.

The second, stronger constraint—an *intragenerational* one—more closely matches the standard idea of a borrowing constraint. Parents may be unable to borrow against their own future earnings in order to smooth consumption or to make investments in their children at young ages. If parents are borrowing constrained in this way, the timing of their income will be important. That is, constrained parents who earn a smaller share of their lifetime income at younger ages will tend to invest less in their children when they are young than will parents who earn a larger share of their income early on.

The first type of constraint is, in theory, easy to test. Do children from wealthier families perform better than children from poorer families? In practice, however, this type of test is difficult to implement since innate abilities of children may be correlated with the abilities and lifetime earnings of their parents. Additionally, rich parents and poor parents may differ in many unobservable ways that have little to do with income. This problem has plagued past work on credit constraints and college-going behavior as discussed in Carniero and Heckman (2002). So, while we think that intergenerational constraints are almost certainly important, we focus attention on the effects of intragenerational constraints.

Testing for intragenerational borrowing constraints relies on examining how the timing of income matters conditional on the discounted value of lifetime income. This amounts to comparing the test scores of children in families with the same total lifetime earnings but with different income profiles. If families are unaffected by borrowing constraints, children from families who earn a larger share of their income early on should perform as well as children from families who earn more of their income late in their careers. If borrowing constraints are binding, the first set of children should perform better than the second. For example, in the absence of intragenerational constraints, a child raised in a family that earns $20,000 for the first 10 years of a child's life and then $40,000 for the next 10 years should perform as well as a child raised in a family earning $40,000 for the first ten years and $20,000 the next ten (ignoring inflation and discounting). If intragenerational constraints are important, the latter child should perform better. By holding lifetime income constant, we can reduce concerns that a positive correlation between parents' and children's abilities will bias our results.[6]

It is important to recognize that income earned earlier should be worth more in a present-value sense. That is, income earned today is worth $1 + r$ times income earned next year (in the absence of credit constraints), where r is the annual interest rate. It is necessary to account for this when examining the effects of income timing. Therefore, we discount income earned by the family so that it is measured in present-value terms as of the child's birth year.[7] After this adjustment, the absence of intragenerational borrowing constraints suggests that the timing of income should be irrelevant conditional on the discounted present value of lifetime family income.

We examine three types of evidence on the importance of income timing. First, we test whether the slope of a family's income profile significantly affects test scores conditional on the discounted present value of lifetime family income over a long time span. Second, we estimate the effects of past income on children's achievement-test scores to see whether income earned at a child's earlier ages has a different effect on test scores than does income earned at later ages. Third, we estimate whether future income has the same effect on test scores as do past and current income. The first test examines the role of income timing over the past, present, and future; the second compares the effects of income earned at different points in the past; and the third compares the effects of past income with future income. As we next show, these tests point to the existence of intragenerational borrowing constraints (as well as intergenerational constraints): family income earned at a child's earlier ages has significantly larger effects on a child's test scores than does income earned at later ages.

EVIDENCE ON BORROWING CONSTRAINTS AND CHILD ACHIEVEMENT

The National Longitudinal Survey of Youth

We use data on children from the NLSY and the main NLSY sample of mothers. These data are ideal for studying the effects of family income on children because they enable us to link children to their mothers and follow families over many years. For children, biannual measures of family background and cognitive and behavioral assessments are available from 1986 to 2000. Detailed longitudinal demographic, educational, and labor market information for the mothers is available annually from 1979 through 1994 and biannually thereafter. Additionally, a widely used measure of cognitive ability—the Armed Forces Qualifying Test (AFQT)—is available for mothers in the NLSY. The NLSY oversamples minority and poor white families, which provides a larger sample of the families more likely to face borrowing constraints. We use data drawn from more than 7,000 interviewed children born to over 3,500 interviewed mothers.

Since family-income measures are available annually from 1979 to 1994 and biannually thereafter, for children born after 1979, we were able to compile an income history for almost every year since birth. Our empirical analysis first uses the Consumer Price Index for Urban Consumers (CPI-U) to deflate all income measures to year-2000 dollars. We then created two measures of family income that we use for our empirical analysis. The first simply discounts income at each age of the child back to the first year of that child's life, using a 5 percent interest rate. This puts all income measures for a child on an equal basis for comparison—a dollar at age 10 should be just as valuable as a dollar at age one after discounting in this way. The second income measure we created is a measure of average "lifetime income," which is simply the average of all discounted income measures for all observed periods from birth to the final survey date. This measure is fixed for each child and does not vary over time.

We analyze PIAT math and reading scores, collected biannually from 1986 to 2000, for children ages 5–14. The assessments measure mathematics ability, oral-reading ability, and the ability to derive meaning from printed words. To simplify our exposition, we focus on a combined math and reading test score, which places a weight of 50 percent on the math score and 25 percent each on of the reading-comprehension and reading-recognition scores. To make these scores more easily interpretable, we created standardized test scores with a mean of zero and a standard deviation of one.[8]

Empirical Tests of Intragenerational Borrowing Constraints

In the absence of credit constraints, the slope of a family's life-cycle-income profile should have no affect on child achievement after controlling for the discounted present value of lifetime income. The presence of intragenerational borrowing constraints suggests that children raised in families that obtain more of their income earlier will perform better; thus, the slope of income profiles should be negatively related to achievement. We test this proposition by regressing PIAT scores at all observed ages on the discounted present value of lifetime family income and the slope of family-income profiles (as well as a

Table 1: Effects of Lifetime Income and the Slope of Family-Income Profiles on PIAT Scores

	1	2	3
Slope of Family-Income Profiles	**-0.1689**	**-0.1187**	**-0.1629**
	(0.0838)	**(0.0796)**	**(0.0793)**
Average "Lifetime Income"	0.2199	0.0964	0.0624
	(0.0062)	(0.0065)	(0.0070)
Controls:			
Basic Background Measures		X	X
Additional Controls			X

Notes:
(1) Income is measured in $10,000 and is discounted at an annual rate of $r = 0.05$.
(2) All specifications control for the child's age, race, and gender and for the mother's age.
(3) Average "lifetime income" refers to the discounted present value of all available income measures after birth of the child.
(4) Basic background measures include the mother's education, AFQT score, and family background (foreign born, rural residence at age 14, living with both parents at age 14); the highest grade completed by the child's grandparents; and year dummies.
(5) Additional controls include the number of adults and children/dependents in the household, current marital status of the mother, age and education of the spouse, and whether or not the mother is currently living with her parents.

variety of family background characteristics). These results are presented in table 1 for a few different specifications. Consistent with intragenerational credit constraints, we generally find negative effects of the slope on children; although, the estimates are not always statistically significant.[9] To interpret these estimates, it is useful to note that incomes grow about $2,300 more per year for those with high earnings growth (90th percentile) relative to those with low earnings growth (10th percentile). The most conservative estimates in table 1 suggest that this difference is associated with a 0.027 standard deviation difference in PIAT test scores.

We next estimate the effects of family income earned at all earlier ages (allowing those effects to vary by age) on PIAT scores, controlling for a variety of individual and family background characteristics as well as average discounted lifetime income. Here, we are interested in determining whether the effect of income on subsequent PIAT scores depends on the age at which that income was received. In the absence of intragenerational borrowing constraints, income earned at all ages (once appropriately discounted) should have the same effect on child achievement. By contrast, when borrowing constraints are binding for some families, income received at earlier ages should have a greater effect than income received at later ages.

We explore the effects of income at different ages assuming that the effects of income on a child's test score depend linearly on the age at which income is earned and the age at which the achievement test is taken. Specifically, we assume that the effects of family income for child i earned at age j, $I_{i,j}$, on an achievement test score at age a, $T_{i,a}$ is given by

$$\frac{\partial T_{i,a}}{\partial I_{i,j}} = \alpha_1 + \alpha_2 j + \alpha_3 a$$

In this case, α_2 tells us how income earned at different ages affects subsequent child test scores, while α_3 tells us at what ages test scores respond most to changes in past income. Table 2 reports

Table 2: Effects of Past Income on PIAT Scores

	1	2	3	4
Long-Term Effects:				
α_1	-0.0140	-0.0102	-0.0101	0.0260
	(0.0031)	(0.0038)	(0.0064)	(0.0098)
α_2	**-0.0013**	**-0.0011**	**-0.0011**	**-0.0020**
	(0.0006)	**(0.0008)**	**(0.0008)**	**(0.0008)**
α_3	0.0027	0.0025	0.0025	-0.0019
	(0.0004)	(0.0005)	(0.0006)	(0.0011)
Temporary Effect of Current Income		-0.0075	-0.0075	-0.0044
		(0.0136)	(0.0139)	(0.0139)
Avg. "Lifetime Income"			-0.0006	-0.3610
			(0.0272)	(0.0785)
Avg. "Lifetime Income" * Current Age				0.0476
				(0.0097)

Notes:
(1) Long-term effect of income at age J on child PIAT score at age A is $\alpha_1 + \alpha_2 J + \alpha_3 A$.
(2) Income is measured in $10,000 and is discounted at an annual rate of $r = 0.05$.
(3) Average "lifetime income" refers to the discounted present value of all available income measures after birth of the child.
(4) All specifications control for the child's age, race, and gender; mother's age, AFQT score, education, family background (foreign born, rural residence at age 14, living with both parents at age 14), marital status, and the age and education of her spouse; the highest grade completed by the child's grandparents; number of adults and children/dependents in the household; whether or not the mother is currently living with her parents; and year dummies.

Table 3: Effects of Past, Current, and Future Income on PIAT Scores

	1	2	3
Past and Current Income	0.1698	0.0850	0.0619
	(0.0118)	(0.0114)	(0.0116)
Future Income	0.0634	0.0221	0.0111
	(0.0075)	(0.0072)	(0.0072)
Controls:			
Basic Background Measures		X	X
Additional Controls			X

Notes:
(1) Income is measured in $10,000 and is discounted at an annual rate of r = 0.05.
(2) All specifications control for age, race, and gender of the child and age of the mother.
(3) Basic background measures include the mother's education, AFQT score, and family background (foreign born, rural residence at age 14, living with both parents at age 14); the highest grade completed by the child's grandparents; and year dummies.
(4) Additional controls include the number of adults and children/dependents in the household, current marital status of the mother, age and education of her spouse, and whether or not the mother is currently living with her parents.

estimates of α_1, α_2, and α_3 using the NLSY data. Note that specifications 2–4 also account for an additional affect of current income on contemporaneous test scores, while specifications 3 and 4 incorporate an independent effect of discounted lifetime income.[10] Focusing on the estimates of α_2, which are all negative, we see that income earned at later ages has a smaller effect on subsequent child achievement scores than does income earned at earlier ages. These estimates are statistically significant at the 0.05 level in specifications 1 and 4. Consistent with intragenerational borrowing constraints, the estimates suggest that shifting $10,000 in family income from age ten to the first year of a child's life would increase subsequent test scores by .01 to .02 standard deviations. More generally, the estimates suggest that past income has a positive effect on test scores at ages when the math and reading tests were administered (recall that tests were not administered to children before age five). The final specification suggests that lifetime income has growing effects as a child ages. That is, children from wealthier families (as measured by lifetime family income) perform increasingly well over time relative to children from less fortunate families.

The permanent-income hypothesis does not distinguish between income earned in the past and that earned in the future. If individuals are reasonably certain about their future income prospects and unaffected by intragenerational borrowing constraints, income earned in the past, present, and future should all affect child test scores equally. Because the NLSY offers panel data over a long time period, it is possible to observe family income measured both before and after some tests are taken by children. Table 3 reports estimates from regressions of children's test scores on the present value of past and current income as well as the present value of all future income. Columns 2 and 3 control for additional family-background variables as described earlier. All of these estimates suggest that past and current income have a significantly greater effect on test scores than does future income. In the final column, the coefficient estimate for future income is not statistically different from zero, while the effect of past income is significantly positive at 0.06. This result is consistent with both intergenerational and intragenerational borrowing constraints.

An obvious concern with this approach is uncertainty about future earnings. If individuals are completely uncertain about future income, actual realizations of that income process should have no effect on current decisions or outcomes, even if individuals are not borrowing constrained. However, since future income primarily represents income earned over the next one to five years, it seems unlikely that it is all that uncertain for most families. To examine the role of uncertainty more formally, we ask whether the results of table 3 hold for families with fairly predictable income profiles. To measure the predictability of income, we estimate family-specific log income regressions on age and age-squared, using all income observations after the child's birth. We can then compute two potential measures of uncertainty (or variability, at least): (i) R^2 statistics that measure the fraction of the variance in log income that can be explained by age and age-squared alone, and (ii) the square root of the mean squared error (RMSE) from the regression. We separate individuals according to these measures of ability and separately run regressions of children's test scores on past/current and future family income. If the insignificant coefficient on future income in table 3 is due primarily to uncertainty in future earnings, we expect that future income should have effects similar to past/current income among those with predictable earnings profiles (i.e., a high R^2 statistic or a low RMSE). This is not the case. Using either measure of predictability, we find that past/current income has a significantly greater effect than future income among those with highly predictable income profiles. In a combined measure that takes only those individuals with an R^2 above 0.75 and those within the lowest quartile of RMSE, we find a more dramatic difference in coefficients on past/current and future income than we observe in table 3.[11] While predictability in an *ex post* sense (as implied by our measures) does not necessarily imply a high degree of predictability in an *ex ante* sense (i.e., low uncertainty), these results are at least consistent with a more important role for intragenerational borrowing constraints than for uncertainty.

CONCLUSIONS

While none of these "tests" for intragenerational borrowing constraints are perfect, we view the combination of all three sets of results as convincing evidence that some families with young children are constrained. At the very least, this evidence suggests that a better understanding of the role credit constraints play in determining families' investments in children is warranted.

The existence of intragenerational borrowing constraints suggests a positive role for government policy. An inexpensive way to address such concerns may be to expand borrowing opportunities for families with younger children. Simply allowing young parents to borrow against their future earnings in order to pay for early-childhood-development programs or to finance private-school tuition should help alleviate intragenerational constraints. For example, a program modeled on the federal PLUS loan program for parents of college students could be extended to qualifying parents of younger children. It is important to note, however, that expanding borrowing opportunities in this way is not likely to address intergenerational borrowing constraints, which are also likely to be important. Dealing with poor parents' inability to borrow against their children's future earnings prospects is a more complicated problem, which is likely to require a redistribution of resources from wealthier families to poorer families. Subsidies for early childhood programs and private schooling can alleviate problems with borrowing constraints, but they come at a sizeable cost. To the extent that subsidies are granted to all children, wealthier families will tend to respond to such incentives by overspending on childhood investments at high costs to taxpayers. Thus, a more efficient approach may require targeting subsidies to lower-income families that are most likely to be affected by borrowing constraints of one form or the other. But this requires redistribution from the middle and upper classes to the less fortunate. Such redistribution is not necessary to address intragenerational borrowing constraints since they can be alleviated by expanded borrowing opportunities for younger parents.

ENDNOTES

[1] Lance Lochner based his remarks at the conference on this paper.

[2] See Carneiro and Heckman 2002 or Cunha et al., forthcoming, for a detailed discussion of these issues.

[3] Income measures are based on the discounted present value of family income from the child's birth through age 12. Income is deflated using the Consumer Price Index for Urban Consumers and discounted at an annual rate of 5 percent. See the data discussion in the "Evidence on Borrowing Constraints and Child Achievement" section for further details.

[4] Only recently have economists (e.g., Restuccia and Urrutia 2004; Caucutt and Lochner 2005; and Cunha and Heckman 2005) begun to consider multiple investment periods at young ages. However, only Caucutt and Lochner (2005) examine the role of early versus late borrowing constraints. Restuccia and Urrutia (2004) abstract from financial asset accumulation, while Cunha and Heckman (2005) shut down late borrowing altogether and do not focus on early borrowing constraints.

[5] If investments are perfectly substitutable over time, then the timing of income may not matter, even if a family is borrowing constrained. Strictly speaking, our tests offer a joint test against perfect substitutability and borrowing constraints.

[6] One might expect more able parents to earn more of their income later because they invest more in their human capital early on. Then, if abler parents have abler children (a problem in the test of the intergenerational constraint), children in families earning more of their income at later ages should be innately abler. This suggests that, in the absence of credit constraints, children from families earning their income earlier should perform worse, on average—the opposite of the prediction based on credit constraints. Thus, a positive intergenerational correlation in ability may make it difficult to find evidence of intragenerational borrowing constraints.

[7] We use a discount rate of $r = 0.05$; however, other reasonable rates yield similar conclusions.

[8] See Caucutt and Lochner (2005) for a detailed discussion of the data.

[9] Estimates of the coefficient on the slope of income profiles are likely to be biased upward for two reasons. First, since we only use income over a limited number of years in computing average lifetime income, this number will tend to be too low for those with a steeper slope relative to those with a flatter income profile. Because average lifetime income has a positive effect on children, this mismeasurement will tend to bias estimates of the coefficient on the slope of income profiles upward. However, this is unlikely to have much effect on our estimates since income far into the future is heavily discounted. A second potential source of bias would arise if unobserved differences across families are related to the slope of family income profiles. If family income profiles are rising because parents are accumulating human capital, then parents with the steepest earnings profiles will tend to accumulate the most human capital. If those parents also tend to invest more in their children or if their investments are more productive, this will produce upwardly biased estimates of the coefficient on the slope of the income profile.

[10] These effects are estimated based on the following regression:

$$T_{i,a} = X_{i,a}\gamma_x + \bar{I}_i\gamma_I + \alpha_0 I_{i,a} + \alpha_1\left(\sum_{j=0}^{a} I_{i,j}\right) + \alpha_2\left(\sum_{j=0}^{a} jI_{i,j}\right) + \alpha_3\left(\sum_{j=0}^{a} aI_{i,j}\right) + \varepsilon_{i,a},$$

where $X_{i,a}$ represents background characteristics for child i at age a, and $\bar{I}_i$ represents average family income from the child's birth through the final period of observation. The inclusion of $I_{i,a}$ allows for an additional temporary effect of contemporaneous income on test scores, so

$$\frac{\partial T_{i,a}}{\partial I_{i,a}} = \alpha_0 + \alpha_1 + (\alpha_2 + \alpha_3)a.$$

[11] For these individuals, the coefficient on past and current income is 0.149 (standard error of 0.053), while the coefficient on future income is –0.035 (standard error of 0.030).

REFERENCES

Blau, D., and J. Currie. Forthcoming. “Who’s Minding the Kids? Preschool, Day Care, and After-School Care,” in E. Hanushek and F. Welch, eds., *Handbook of Education Economics*. North-Holland, Amsterdam.

Cameron, S., and J. Heckman. 1998. “Life Cycle Schooling and Dynamic Selection Bias: Models and Evidence for Five Cohorts of American Males,” *Journal of Political Economy* 106(2): 262–333.

Cameron, S., and J. Heckman. 2001. “The Dynamics of Educational Attainment for Black, Hispanic, and White Males,” *Journal of Political Economy* 109(3): 455–99

Carneiro, P., and J. Heckman. 2002. “The Evidence on Credit Constraints in Post-Secondary Schooling,” *Economic Journal* 112: 705–34.

Caucutt, E., and L. Lochner. 2005. “Early and Late Human Capital Investments, Borrowing Constraints, and the Family,” Working Paper.

Cunha, F., and J. Heckman. 2005. “The Technology of Skill Formation,” Working Paper.

Cunha, F., J. Heckman, L. Lochner, and D. Masterov. Forthcoming. “Interpreting the Evidence on Lifecycle Skill Formation,” in E. Hanushek and F. Welch, eds., *Handbook of Education Economics.* North-Holland, Amsterdam.

Dahl, G., and L. Lochner. 2005. “The Impact of Family Income on Child Achievement,” Working Paper.

Duncan, G., and J. Brooks-Gunn. 1997. *Consequences of Growing Up Poor*. New York: Russell Sage Foundation.

Ellwood, D., and T. Kane. 2000. *Who is Getting a College Education? Family Background and the Growing Gaps in Enrollment.* New York: Russell Sage Foundation.

Karoly, L., et al. 1998. *Investing in Our Children: What We Know and Don’t Know about the Costs and Benefits of Early Childhood Interventions.* Santa Monica: The RAND Corporation.

Levy, D., and G. Duncan. 1999. “Using Sibling Samples to Assess the Effect of Childhood Family Income on Completed Schooling,” Working Paper.

Restuccia, D., and C. Urrutia. 2004. “Intergenerational Persistence of Earnings: The Role of Early and College Education,” *American Economic Review* 94(5): 1354–78.

Zigler, E. 1994. “Reshaping Early Childhood Intervention to be a More Effective Weapon against Poverty,” *American Journal of Community Psychology* 22(1): 37–47.