The Black-White Education-Scaled Test-Score Gap in Grades K-7*

Timothy N. Bond† Kevin Lang‡

September 2, 2013

Abstract

The ordinality of test scores presents several difficulties when measuring the black-white achievement gap. We address these by estimating in each grade the expected black-white education gap conditional on observed test scores, thus creating a scale with interval properties. We find no racial component in the evolution of the achievement gap through the first eight years of schooling and that most, if not all, of the gap can be explained by socioeconomic differences. Our results suggest that the rising racial test gap found by previous studies is likely due to excessive measurement error in the early grades.

*We are grateful to participants in the Harvard Inequality Seminar, conferences at Linnaeus University, Northwestern and Oberlin, seminars at UC San Diego and UC Irvine and an informal brown bag lunch at Purdue, Ivan Fernandez-Val, Jon Guryan, Jesse Rothstein, and Justin Tobias for helpful comments and suggestions. The usual caveat applies.

†Department of Economics, Krannert School of Management, Purdue University, 403 W. State Street, West Lafayette, IN 47907; email: tnbond@purdue.edu

‡Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215; email:lang@bu.edu

1 Introduction

Test scores are ordinal measures of achievement. They provide a rank ordering of students but cannot measure by "how much" one student outperforms another; the published scales lack interval properties. This fact is widely accepted among experts in education measurement.[1] However, economists who use test scores in research typically ignore this, assigning arbitrary interval properties to test scales which lack them.

We address the problem of ordinality with respect to measuring the black-white test gap during the early years of schooling. To convert the test scores to an interval scale, we use longitudinal data to estimate, in each grade, the expected future black-white education gap conditional on test scores.[2] In contrast to studies using arbitrary scales, we find no evidence of a racial component in the evolution of achievement through the first eight years of schooling. Black students perform no worse in seventh grade than would have been expected based on their kindergarten scores. Further, we find little evidence of a racial component to achievement at all. Once we control for a few socioeconomic variables, most of the gaps disappear.

In their influential studies of the black-white achievement gap, Fryer and Levitt (2004, 2006) find that blacks score similarly to whites on tests at the beginning of kindergarten, but fall substantially behind by third grade. Fryer and Levitt rely on a conventional normalization that places equal value on standard deviations from the mean test score. Our previous work (Bond and Lang, forthcoming) shows that their result is very sensitive to this scale choice. Using the same data set we showed that alternative scales suggest the true growth of the achievement gap ranges from zero to double that shown by Fryer and Levitt (2006).

Scale choice is important in other contexts such as teacher or program evaluation. Lang (2010) points out that renormalizing each year's scores to have a standard deviation of one can cause artificial "fade-out" because the true variance of achievement is likely to increase over time. Under the assumption that the transformed test scores are a linear function of the underlying true scores, Cascio and Staiger (2012) demonstrate this phenomenon empirically. Together Bond and Lang and Cascio and Staiger suggest that scale choice can have important effects on the policy conclusions we draw from changes in measured achievement.

It is possible to mitigate but not eliminate the arbitrariness of such scales by tying them to an external metric. Temperature is measured on an interval scale when related to energy

[1] See, for example, Stevens (1946) and Thorndike (1966) for early references. For examples of approaches to comparing achievement using ordinality alone see Braun (1988), Holland (2002), Ho and Haertel (2006) and Reardon (2008). Modern scoring methods like Item-Response Theory (IRT) only provide a more precise ordinal ranking of test-takers; they do not attempt to impose an interval scale (Lord, 1975).

[2] Cunha and Heckman (2008) and Cunha, Heckman, and Schennach (2010) tie test scores to adult outcomes in a different context.

but not to pleasantness. If we are interested in temperature because of its relation to energy, we can treat it as measured on an interval scale. In this paper, we rescale test scores in each grade so that a one unit change in the scale corresponds to a one-year difference in predicted education. This produces an interval scale with respect to this one external measure but not necessarily with respect to others such as predicted income.

When we measure achievement in terms of predicted educational attainment, black children's kindergarten reading test scores predict that they will obtain .7 years less education than whites. When we instead make predictions based on kindergarten math scores, blacks are predicted to obtain a full year less education than whites. When we measure education not in years but in the associated average log earnings, blacks lag behind whites by around 10 percent. In all cases, the gap is unchanged if we make our predictions based on later test scores. If anything, the evidence points to blacks doing better than expected rather than worse as they progress through school.

We show that rescaling tests based on adult outcomes creates a shrinkage estimator of each student's future achievement. Since our question of interest concerns group, not individual, differences in achievement, a simple average of these scores excessively shrinks the estimate of the black-white achievement gap. We use an instrumental variables procedure to correct for this excess shrinkage. Without the adjustment, the patterns look similar to that in Fryer and Levitt (2006) because excess shrinkage is greatest in the early years of schooling. This suggests that measurement error is a greater problem on early childhood tests and provides a potential explanation for the Fryer and Levitt (2006) results.[3]

Economists have often found different black-white achievement gaps when looking at different tests. The Fryer/Levitt results differed starkly from earlier work that suggested the achievement gap emerges before schooling (e.g., Jencks and Phillips, 1998). In our data, the pre-kindergarten Peabody Picture Vocabulary Test (PPVT) shows a much larger gap than the kindergarten Peabody Individual Achievement Test (PIAT) reading and math tests when scaled in standard deviations as is customary. Previous work suggests differences in the gap across tests reflect differences in test content (Murnane et al, 2006) and test scale (Bond and Lang, forthcoming). Remarkably, our rescaling reveals similar achievement gaps on the PIAT and PPVT tests. While we cannot correct the PPVT for excess shrinkage, this result suggests that differences in the gap across tests may reflect differences in the degree of measurement error, particularly in the early school years.

We find that our "education-scaled" test gap in the early years, particularly in math, is at

[3] As discussed by Junker, Schofield, and Taylor (2012), economists have also frequently ignored issues of measurement error in analysis of test scores. They show that using the reliability estimates from an underlying IRT model to correct regression estimates can have a large impact on their magnitude. Boyd et al (2012) show that measurement error is noticeably larger than suggested by reliability estimates.

least as large if not larger than the actual gap in educational attainment. At the same time, barring some surprising narrowing of the black-white earnings gap, the "earnings/education-scaled" test-gap in kindergarten through seventh grade is less than the future earnings gap, suggesting either that test scores contain information beyond their effect on education, as argued in Neal and Johnson (1996), or continued labor market discrimination.

Strikingly, much, and in some cases all, of the education-scaled gaps can be explained by a small number of controls representing the child's early environment. Results that condition on sociodemographics should be treated with great caution due to the sociological fallacy (Jensen, 1969). However this suggests that our previous inability to explain the test gap by environmental factors may have reflected scaling decisions. The achievement gap may be due to racial differences in socioeconomics rather than a specific racial component in human capital acquisition or the environment more generally.

It is important to understand what our results do and do not mean. It would be easy to interpret our results as saying that the entire black-white education gap is due to pre-school factors. This is true only in the sense that the entire gap can be predicted on the basis of kindergarten test scores. But it is not true if it is interpreted to mean that subsequent events do not affect the gap. To the extent that low kindergarten scores predict future attendance at lower quality schools, less future parental support, etc., the gap is "explained" by factors known in kindergarten. But this does not mean that differences in these factors no longer matter. Instead, our results tell us that blacks do no better or worse, on average, than would be predicted by their early test scores.

This finding is important for two reasons. First, even if it were possible to measure the change in the magnitude of the gap based on an absolute scale, say the Fundamental Test of Intellect and Learning (or futiles), we would want to know whether the evolution of the gap was predictable from its initial condition. If individuals whose achieved skill differs by 2 futiles at the end of kindergarten normally differ by 6 futiles by the end of third grade, then a growth in the black-white gap from 2 to 6 over the same period does not point us to a race-based explanation for the growing gap. Our approach addresses this directly. Second, although most previous work has implicitly recognized this concern by normalizing the standard deviation of test scores in each grade to 1, it has ignored the effect of measurement error. Our approach allows us to correct for measurement error, which we show declines noticeably over the first few years of school. This, in turn, implies that the standard approach over-estimates the growth of the achievement gap. Finally, our previous work (Bond and Lang, forthcoming) shows that arbitrary scaling decisions lead to very different conclusions about the evolution of the achievement gap.

2 Theoretical Framework

2.1 Intuition: Adult Outcomes as a Shrinkage Estimator

When we observe that two children have two different test scores, the fact that these scores lack interval properties can prevent us from drawing many conclusions. If Sally scores a 12 and Billy scores a 10, we know that Sally performed at a higher level than Billy, but without other data its very difficult to quantify what the size of that difference means. Likewise if one year later, Sally scores a 13 and Billy scores a 12, we do not know if Billy gained ground relative to Sally. One could easily imagine examples where Billy had fallen further behind Sally in economically important ways. For instance, suppose the exam was a college entrance exam, and 13 was the minimum score for admission.

To allow these cross-person and overtime comparisons, we propose to create an interval scale using longitudinal data that includes adult outcomes. Suppose we could observe from a large number of observations that, on average, students who scored a 10 received 11 years of education, students who scored an 11 received 11.1 years of education, students who received a 12 received 11.3 years of education, and students who received a 13 obtained 12 years of education. We could then say that based on the first test we predict that Sally will obtain .3 years of education more than Billy, and that based on the second test Sally will obtain .7 more than Billy. Therefore Billy lost ground relative to Sally.

In creating a scale such as this where we use the average of future performance, we are inherently using a shrinkage estimator. Shrinkage estimators are commonly used in Bayesian statistics and can be thought of in two complementary ways. They are biased estimators that have lower mean-squared error than unbiased estimators. They are estimators that use additional information to improve unbiased estimators.

Suppose that we have a test that is relatively poor indicator of future performance. The test scores are virtually uncorrelated with adult outcomes. This could be for two reasons. First, the test could be testing skills that are not economically relevant, in which case it is unlikely we would want to use this test as a guide for policy. Second, the test could involve a good deal of measurement error.[4] In each of these cases, there will be little dispersion once we transform the scores into predicted future outcomes. The tests are poor indicators of future achievement, and we would not want to put great weight on them in making such a future prediction. Our estimates are "shrunk" to be closer to the mean.

Thus, in the course of making our transformation, we are correcting for the oft ignored problem of measurement error in test scores. Our new scale will put less weight on the results

[4] We omit a third situation in which we could have chosen an adult outcome that is not economically relevant.

of noisy tests when measuring the difference in achievement between two individuals.

2.2 Intuition: Excessive Shrinking in Group Comparisons

Our focus in the previous subsection was on comparisons across individuals. The focus of our paper is different. We wish to use individual data to estimate differences in group means. Because we now have many observations of our group, we would ideally like to shrink our estimates less. We have more confidence in our estimate of average group performance even if we have little confidence in our estimates of individual performance. Thus, taking a simple average of our individually rescaled scores would be incorrect. An example from outside the realm of test measurement may help illustrate this point.

Consider the game of baseball and suppose we are asked to predict the difference in batting average (hits/at bats) for two players based on ten official at bats. We observe that the first player has two hits in ten at bats for an average of .200, while the second has three hits in ten at bats for an average of .300. Based on our data the maximum likelihood estimator would project that these two players would finish the season with a .100 difference in batting average. But given that in each season players typically have 500 or more at bats, even those with no statistical training are unlikely to make this prediction. Instead they are likely to state that baseball is a game with a lot of variance, and the season is long. Therefore it is most likely that both players will end the season with a batting average around .255, the average for the league.[5]

Now suppose instead that we are comparing two teams of 25 players each of whom has had 10 at bats, and that the proper shrinkage estimator for an individual player with 10 at bats is .252 + .001*(number of hits). We take the average of the shrunk scores for each team. Team A's batting average using the shrunk scores is .254. Its players have an average of two hits each. The team has a total of 50 hits in 250 at bats. If the true mean batting average for the team is .254, the odds of having only 50 hits in 250 at bats is only about 2.5 percent.[6] Suppose that team B's batting average using the shrunk scores is .255. The team has a total of 75 hits in 250 at bats. If its true team batting average is .255, the odds of it getting 75 or more hits in 250 at bats are only about 5 percent. The likelihood that the true team batting averages differ by only .001 is very small. Using the average of the shrunk

[5] Of course, given additional information such as that the first player was a notoriously poor batter, Mario Mendoza, while the second player was Ted Williams, one of the greatest batters in the history of baseball, most people would put even more weight on their prior and say with a great deal of confidence that the second player will perform better. In this case we would be shrinking our estimate to the individual's career mean rather than the league's.

[6] The probability that a randomly chosen player would have five hits out of ten is also about .02, but the probability that at least one of fifty players would have five hits is about .66.

estimates understates the gap.

The same logic applies when we have many students taking tests. When basing our projection for an individual off one single poor test score, we will not be very confident that the student will have a worse than average outcome. Test scores contain some randomness; they measure ability with error. But if we consistently observe individuals from one group obtaining a poor score on the test, we will be more confident that that group as a whole will have a worse outcome. We may not be confident about which individual from that group will have a poor outcome, but we can be fairly confident that the group will have poor outcomes on average. In essence, averaging individual shrinkage estimators is throwing away the information we obtain by repeated observations of members of the group.

If there is a genuine gap in the average performance of two exogenously selected groups, then the average of the individually shrunk gaps will be smaller than the true average gap. In section 2.4, we discuss how we can estimate and undo this excess shrinkage.

There is one last complication we must discuss. So far, we have assumed that the number of observations with each test score/grade combination is large. In practice, this will not be the case. Therefore, in addition to the other sources of measurement error discussed in the literature, our transformed scales will be subject to sampling error and thus will be the shrunk estimates plus sampling error. As the number of observations gets large, this sampling error will go to 0. The importance of sampling error is an issue we will address empirically.

2.3 Intuition: Interpreting Outcome-Scaled Scores

It is important to emphasize that scores that are rescaled to adult outcomes do not directly measure an individual's skill set at the time of testing. For example, if we predict that a kindergartner with a test score of 256 will earn $85,000 at age 45, we do not mean that we would expect her to earn $85,000 if she reached the age of 45 with her current skill set. Instead we mean that the average person with her skill set in kindergarten will, at age 45, have a skill set associated with earning $85,000.

Suppose that we find that children who enter kindergarten earn $25,000 at age 45, on average, if they are only able to recognize and identify six letters but earn $85,000 if they can read and understand relatively simple passages. Does the skill gap between individuals with these scores grow between kindergarten entrance and age 45? The answer, by our definition, is "no." Of course, we could define the skill gap in other ways, such as the difference between what they could earn at age 45 with their current skill set, but since almost no kindergartners have the cognitive and non-cognitive skills that would allow them to work even in a sheltered work environment, the kindergarten gap by this definition is 0, and it is not informative to say that it grows.

The skill gap that we observe early in life may directly affect future skill acquisition. Students with high test scores at age 5 may be tracked into an accelerated program that allows them to acquire new skills more quickly. Or students with high test scores may continue to reap the benefits of the factors that enabled them to earn high initial test scores. Both of these factors will be captured by our measure. This is a fundamental limit of tying test scores to adult outcomes.

Thus, if the gap between two groups based on predicted education does not change over time, the skill gap based on some ideal scale might nevertheless change. What we know is that the difference in performance does not change in a way that was not predicted by their earlier scores. Neither group has deviated from the path predicted by kindergarten performance.

Observing a large but constant gap between blacks and whites is consistent with a world in which by some absolute (and, in our view, undeterminable) metric blacks have only slightly lower average skills than whites at school entry but are subsequently assigned to worse schools which exacerbate these differences. However, it must then also be the case that, *conditional on skill at entry,* school quality is similar for blacks and whites. In other words, low-skill whites and blacks follow the same skill trajectory.

To illustrate this point, consider two examples. Suppose that, on average, blacks have poorer reading skills than whites at entry, and the government institutes an effective intervention targeting low-skill readers in the second grade. As a consequence, children with poor reading skills in kindergarten generally benefit from remediation in second grade and eventually complete more education. Of course this effective intervention will lower the black-white test score gap in grades two and beyond. However, it will also raise the average education associated with low reading test scores in kindergarten and first grade and therefore close the black-white achievement gap in these grades since it improves the trajectory of early low-skill readers. Suppose instead that the intervention targets only low-skill *black* readers. Because blacks are a small part of the population, this intervention has only a small effect on the average eventual completed education of all students with low early test scores and therefore little effect on the black-white test score gap in kindergarten and first grade. It will have a more substantial effect on the black-white achievement gap from second grade onwards. After participating in the program, black students have better outcomes than would be projected on the basis of their kindergarten and first-grade test scores alone.

Assuming no other race-related factors, we would appropriately conclude in the first case that race did not predict a change in the gap and in the second case that the gap narrowed.

2.4 A More Formal Presentation

Suppose we are interested in the difference in average "achievement" between blacks and whites in a given grade. This poses two immediate problems. The first is that we cannot observe achievement directly. The second is that achievement has no natural scale.

To solve the latter problem, suppose we are interested in achievement because it predicts future levels of education. We can then normalize achievement at a given time to be in units of expected completed schooling, S, so that for each individual i in grade g,

$$S_i = A_{ig} + \varepsilon_{ig} \tag{1}$$

where, A_{ig} is units of normalized achievement and ε_{ig} is a mean zero error term that reflects determinants of educational attainment that arise after the measurement of achievement. We assume that $E\left(A_{ig}\varepsilon_{ig}\right) = 0$ and $E\left(\varepsilon_{ig}\varepsilon_{jg}\right) = 0$ for $i \neq j$. The interpretation of ε is important. We will therefore discuss it in greater detail later in this section.

We assume that we have access to a series of test scores. In each grade the test score, τ_{ig}, is a function of A_{ig} and some noise component ν_{ig},

$$\tau_{ig} = \tau_g(A_{ig}) + \nu_{ig}. \tag{2}$$

The function τ_g reflects the fact that each test is scored in some arbitrary fashion so that it is not necessarily linear in (normalized) units of achievement while ν_{ig} reflects the measurement error associated with any test.

We further assume that measurement error is uncorrelated over time, so that $E[\nu_{ir}\nu_{is}] = 0, \ \ \forall r \neq s$. As Boyd, Lankford, Loeb and Wyckoff (2012) discuss in detail there are a number of factors that contribute to measurement error other than those captured by test publishers' estimates of reliability. Some of these such as luck or how the student was feeling on a particular day are very likely to be uncorrelated over time. This assumption is less obvious in the case of the items or domains included in the exam. We will present evidence that suggests that such serial correlation is unlikely to be a significant problem in our data.

We derive an "education-normalized scale" in the following fashion. Suppose that τ is discrete, as it is in our data. We define the scale by the population mean of S_i at τ_{ig} :

$$s_g(q) = \frac{\sum_{\tau_{ig}=q} S_i}{N_q} \tag{3}$$

where N_q is the number of individuals in that grade with a score of $\tau = q$. Thus $s_g\left(q\right)$ is the average education ultimately attained by individuals with a score of q on the test in grade

g.

We note that in practice this approach will add sampling error to the other sources of measurement error because we have only a finite number of observations at each score in each grade. In part to address this issue, we also provide one set of estimates based on a kernel estimator that combines test scores from three tests.

Henceforth we drop the subscript g when doing so will not cause confusion.

It is tempting to define the education-normalized test score gap by the difference in the mean of s for the two groups. However, this will be incorrect. Suppose

$$\tau_i = A_i + \nu_i \tag{4}$$

and that A and ν are independent and normally distributed with variances σ_A^2 and σ_v^2, respectively. Note that we are fortunate in this example because the test scores have already been scaled to equal the achievement scale.

A standard result from statistical theory gives

$$E\left(A|\tau = a\right) = \beta_1 a + \left(1 - \beta_1\right) \overline{A} \tag{5}$$

where

$$\beta_1 = \frac{\sigma_A^2}{\sigma_A^2 + \sigma_\nu^2}. \tag{6}$$

Now because $s\left(a\right)$ is just the average of A given that τ equals a, by the law of large numbers

$$plim_{N(a)\to\infty} s(a) = E\left(A|\tau = a\right) = \beta_1 a + \left(1 - \beta_1\right) \overline{A}. \tag{7}$$

While the conclusion that s_i is the shrinkage estimate of $\left(A|\tau_i\right)$ does not rely on the assumption of normality, we will maintain this assumption for the rest of the example. Suppose we have a large number of test scores from group c. Then

$$s_c \equiv \frac{\Sigma_{i\in c} s_i}{N_c} = \beta_1 \frac{\Sigma_{i\in c}\tau_i}{N_c} + \left(1 - \beta_1\right) \overline{A} \tag{8}$$

$$= \beta_1 \frac{\Sigma_{i\in c}\left(A_i + \nu_i\right)}{N_c} + \left(1 - \beta_1\right) \overline{A} \tag{9}$$

where N_c is the number of members of group c. But

$$plim_{N_c\to\infty} \left(\beta_1 \frac{\Sigma_{i\in c}\left(A_i + \nu_i\right)}{N_c} + \left(1 - \beta_1\right) \overline{A}\right) = \beta_1 A_c + \left(1 - \beta_1\right) \overline{A} \tag{10}$$

where A_c is the mean achievement of group c.

By the same logic

$$plim_{N_c \to \infty} s_c - plim_{N_b \to \infty} s_b = \beta_1 (A_c - A_b). \tag{11}$$

To find a consistent estimate of the differences in achievement between the two groups, we need to augment the difference between their mean education-scaled test scores by a factor of β^{-1}.

To address this overcorrection, we approximate the relation between s and A by a linear function

$$s_i = \beta_0 + \beta_1 A_i + \mu_i. \tag{12}$$

The linear relation is exact under the normality assumption but need not be otherwise. In principle, we could allow for a more general relation.

If we observed A_i, we could estimate β_1 by regressing s_i on A_i. We do not observe A, but we do observe S which, from (1), is a noisy measure of A. We estimate

$$s_i = \beta_0 + \beta_1 S_i + \varepsilon_i \tag{13}$$

Because $S_i = A_i + \varepsilon_i$ with $E(A\varepsilon) = 0$, the measurement error is classical. Therefore, we can estimate β_1 consistently if we can find a suitable instrument for S. A natural instrument for S is s_{g-1}, the (renormed) test score from a prior test. However, the renorming includes S_i and therefore is correlated with ε_i. Therefore, we construct a "leave-one-out" instrument which is the average eventual educational attainment of all *other* individuals with the same test score on the prior test

$$s^*_{ig-1} = \frac{\sum_{g_{jg-1}=q, j \neq i} S_{ig-1}}{N_q - 1}. \tag{14}$$

s^*_{ig-1} is correlated with A_{ig} since achievement is persistent.

Therefore we estimate the black-white achievement gap by

$$\widehat{\beta_1^{-1}} (s_w - s_b) = \frac{\Sigma S_i s^*_{ig-1}}{\Sigma s_{ig} s^*_{ig-1}} (s_w - s_b). \tag{15}$$

Note that if measurement error is positively correlated over time, we will underestimate β_1^{-1} and therefore the magnitude of the test-score gap. In finite samples, this will be a problem because some individuals will earn the same test score as each other in both grade g and grade $g-1$. Their completed schooling enters the calculation of both s_{ig} and s^*_{ig-1} creating correlation in the error measurement. Asymptotically this correlation goes to zero as both the overall measurement error and correlated measurement error go to zero. We return to a discussion of the importance of small sample bias later.

2.5 Interpretation and the Martingale Property

It is important to remember exactly what our estimates mean and how they should be interpreted. A_{ig} is predicted educational attainment based on achievement in grade g. This interpretation has important implications for what we should find in the data. By definition

$$A_{ig+1} = A_{ig} + \omega_{ig+1}$$

where ω_{ig+1} is the innovation in achievement between grades and

$$E\left(A_{ig}\Sigma_{t=1}^{T}\omega_{ig+t}\right) = 0.$$

Thus, A is a martingale and s is a martingale augmented with measurement error. As discussed in Farber and Gibbons (1996), this means that the covariance of the test scores

$$\begin{aligned}\sigma_{g,g+t} &= E\left(A_o - \overline{A}_0 + \Sigma_{j=1}^{g+t}\omega_j + \mu_{g+t}\right)\left(A_o - \overline{A}_0 + \Sigma_{j=1}^{g}\omega_j + \mu_g\right)\\ &= \sigma_{A_0}^2 + \Sigma_{j=1}^{g}\sigma_{\omega_j}^2 + \sigma_{\mu_{g,g+t}}.\end{aligned}$$

The first two terms are independent of t. Under the assumption that measurement error is uncorrelated over time, the last term is 0 except when t equals 0. Note that in contrast, the covariance is increasing in g. Therefore, the model implies that the lower triangle of the covariance matrix is constant for all terms in a column below the diagonal and increasing from left to right. We can therefore cast light on the importance of serial correlation of the measurement error by examining the covariance matrix of the test scores.

3 Data

The Children of the National Longitudinal Survey of Youth (CNLSY) is a biennial survey of children born to women surveyed in the National Longitudinal Survey of Youth 1979 cohort (NLSY79). The NLSY79 is a longitudinal survey that has followed a sample of 12,686 youths who were age 14 through 21 in December 1978. The survey includes a nationally representative sample, as well as oversamples of blacks, Hispanics, military personnel, and poor whites. The military and poor white oversamples were dropped from later surveys.

Since 1986, the children of women from the NLSY79 have been surveyed and assessed every other year. Separate questions are asked for children and young adults. Children are eligible to enter the childhood sample at birth and advance to the young adult sample at age 15. As of 2010, a total of 11,506 children born to 4,931 unique mothers had been surveyed.

Our focus is on the Peabody Individual Achievement Tests (PIAT). Children were given three PIAT assessments in each survey in which they were age five through fourteen. The PIAT Mathematics (PIAT-M) measures mathematics skill as typically taught in school. It is comprised of 84 multiple choice questions on a wide range of topics from number recognition to trigonometry. The PIAT Reading Recognition (PIAT-RR) is an oral reading test which assesses children's ability to recognize letters and read single words. The PIAT Reading Comprehension (PIAT-RC) is a test of a child's ability to understand sentences. The PIAT-RC is administered only if the child's score on the PIAT-RR is sufficiently high.[7]

We also examine the Peabody Picture Vocabulary Test (PPVT). The PPVT is a test of receptive vocabulary designed to assess general aptitude. The CNLSY currently administers this test to children at age four or five and age eleven, but due to variation in this policy over time, we observe PPVT scores for children as young as three. We are interested in the PPVT primarily as a measure of achievement before entering grade-school. Therefore we restrict our analysis of the PPVT to those who took the test before age five.

While the survey is a panel by year, we are interested in the racial achievement gap by grade. To convert our data to such a panel, we drop any child we observe in the same grade over multiple surveys. Because the survey was conducted biennially, this restriction binds if the child spent three years in the same grade and thus affects only a handful of individuals. We focus only on the black-white test gap, and drop members of other races. These modifications leave us with an unbalanced panel of 7,343 children born to 3,318 mothers.

The sample is not nationally representative because children born before 1982, when the mothers were age seventeen through twenty-five, are observed only during their later childhood, while those born in later years are observed only during their early childhood. To correct for this non-representativeness, we create custom weights for each grade-test designed to make that subsample nationally representative.[8] Individuals with a valid PIAT-RR raw score below the threshold for taking the RC are included in the construction of the weights for the PIAT-RC but are excluded from the analysis. This avoids putting undue weight in the early grades on a small number of low achieving students who advance to the RC due to randomly high scores. The RC results should, therefore, be interpreted as representative of the population that would have scored sufficiently well on the RR to take the RC exam in that grade. Note that we should view gaps based on the RC with caution especially

[7] From 1986 to 1992, the threshold was a raw score of 15 on the PIAT-RR. This threshold was subsequently raised to 18.

[8] We are grateful to Jay Zagorsky of the Center for Human Resources Research for providing us with the program.

in kindergarten but also in first grade because the students taking the exam are not fully representative of the overall student population.

Table 1 shows the gap on the age-adjusted percentile scale for each test in each grade, using the custom weights discussed above.[9] To ease comparison with other studies, in table 1 we follow convention and normalize the scores in each grade to have mean 0 and standard deviation 1. Because some children leave the child sample and enter the young adult sample during 8th grade, we restrict our attention to kindergarten through grade seven.

Each test tells a different story about the black-white test gap. In kindergarten blacks are .65 standard deviations behind whites on the math test. This gap rises only very slightly through seventh grade. The two reading gaps are initially only very modest but grow to roughly the magnitude of the math gap by third grade. The PPVT, administered earlier than the PIAT tests, shows a gap of over one standard deviation, larger than the gap on any PIAT test in any grade.

Taken together these tests reflect the myriad of findings in the black-white test score gap literature. The reading tests show the pattern demonstrated by Fryer and Levitt (2004, 2006) for the test administered as part of the Early Childhood Longitudinal Survey (Kindergarten sample). The PPVT gap appears similar to that in Jencks and Phillips (1998), while the PIAT-M shows a nearly constant gap that is smaller than the one observed on the PPVT.

As discussed in Bond and Lang (forthcoming), these test gaps are based on arbitrary scaling decisions. Plausible order-preserving transformations of the scales can produce startling different results. The bounds we established in that paper are very large. There we show that without placing more structure on the scales, the gap could be small to modest and decreasing from kindergarten through third grade or small to nonexistent in kindergarten but growing to substantial by the end of third grade, or somewhere in between.

In this paper, we resolve this indeterminacy by relating the scores to economic outcomes, in particular educational attainment. To do so, we construct a sample of 3,853 children who are observed in the panel after age 22 and for whom we know highest grade completed. We lose roughly one-half of our observations on each test, but still have over 1000 observations for all but the earliest PIAT-RC. We again construct custom weights so that each of these test-grade samples is nationally representative.

Table 2 repeats table 1 for this subsample. The magnitudes of the test gaps are generally similar to the full sample, though at times somewhat smaller. This probably reflects the fact that children who are 22 by 2010 were born no later than 1988 when the mothers were 23 to

[9] This scale represents the percentile of the distribution each child's raw score is in for their three month age group. Note that since we are grouping children of different ages within the same grade together, this means that younger children may have higher percentile scores than older children within the same grade despite having answered fewer questions correctly.

31 and thus were born to relatively young mothers. By restricting the age of the mothers, we reduce the socioeconomic differences between black and white mothers.[10] Nevertheless, the patterns mimic those in table 1: a math test gap that grows only very slowly, a growing reading gap, and a pre-schooling PPVT gap that is larger than that on any subsequent test.

Since we use this sample only to translate test scores into an education scale, the test score gap for the older sample has no direct significance. The real risk is that, because our sample with completed education was born to young mothers, the relation between test scores and educational attainment for this group may not be representative of the entire population in a way that biases our estimate of the "education test-score" gap. It will be apparent that the test scores for the older group are lower than for the sample as a whole, but it is difficult to determine whether this causes any bias in our estimates of the test score gap.

Not surprisingly given past research, we observe a racial gap in educational attainment. Table 2 also displays the difference in average educational attainment between blacks and whites for each test-grade sample. We observe gaps that are generally between .70 and .85 years of education, depending on the sample. This is somewhat higher than we observe for their parents' generation (.70) in the NLSY79 adult sample. It is unclear whether this reflects a change in the gap or the nonrepresentativeness of our older sample.

Our empirical approach depends on the assumption that measurement error is uncorrelated over time. Our model implies that the covariance between a test score in period t and all subsequent test scores $cov(s_g s_{g+2j})$ should be a constant for all $j = 1, 2, 3...$[11] and that $cov\left(s_{g+k} s_{g+k+2j}\right)$ should be nondecreasing in k for k positive. Appendix table A shows the unweighted covariance matrix of the test scores. We have relatively few years for which we can test this hypothesis. In all cases the covariance terms are much smaller than the variances. While we have not formally tested the hypothesis that $cov(s_g s_{g+2j})$ is constant for all test and grade combinations, it does not appear to be severely violated. This suggests that the correlation in measurement error induced by some individuals sharing the same scores in tests in years g and $g - 2$ is unlikely to be a serious concern. We address this concern directly later in the paper.

[10] On one measure of background, mother's AFQT percentile score, the gap between blacks and whites grows by about .03 standard deviations per year increase in mother's age at child's birth. Moreover, white mothers tend to be older than black mothers.

[11] We use $2j$ instead of j because tests are generally administered two years apart.

4 Empirical Implementation

In order to obtain estimates of s_g for each grade-test combination, we use all individuals in our sample with a valid score for that grade/test and for whom we observe educational attainment after the age of 22. We then calculate average educational attainment by score for that sample. We apply the results of this rescaling to the entire sample. We interpolate s_g for any test scores not present in the over 22 sample. This produces a score on the new scale for each individual with a valid test score on that grade-test.

Figures 1-3 show the relation between the transformed score and the base percentile on each test. While the underlying relation should be strictly increasing, not surprisingly, given the small number of observations with a particular score on a given test, there is notable imprecision in our point estimates. There is, however, a clear overall positive relation between test performance and educational attainment, s, on each test in each grade.

We first estimate the gap between blacks and whites using the s_g scales. However as discussed above, these scales over-correct for measurement error when applied to group averages and thus understate the gap. We correct this by estimating the relation between schooling and the s_g scores. If schooling were a perfect measure of achievement/ability, this would provide an estimate of how much our s_g measure understates achievement. However schooling is achievement measured with error, and so this will attenuate our correction towards zero. We correct this by using the lagged s_g values as instruments. Because the survey is given biennially, we use two year (grade) lagged test scores. For the first grade and kindergarten scores, we use the childhood PPVT s. We also use the PPVT as an instrument for the second grade PIAT-RC due to the small size and selected nature of the sample of children who advance to that test in kindergarten. Each instrument is calculated using the leave-one-out method to avoid correlation arising from the use of the individual's eventual schooling attainment in creating the s scale.[12]

We bootstrap the standard errors. Particularly in the early grades, the distribution of the bootstrap estimates tends to be skewed. Therefore, we present the 95% confidence intervals for all of our estimates, which will be valid under weaker assumptions than required for the use of the normal approximation.

The scale discussed thus far assumes that we value all years of education equally. There are many other possible choices. We do not attempt to consider the full range of alternatives, which would lead to a bounding exercise similar to that in Bond and Lang (forthcoming).

[12]It would be possible to estimate β by using the older sample to calculate $\sigma_{Ss_{g-2}}$ and the full sample to calculate $\sigma_{s_g s_{g-2}}$. This would probably increase the precision of our estimates somewhat, but we are concerned that because the older sample is more homogeneous, calculating covariances from two different samples would be problematic.

Instead, we consider one quasi-monetary scale in which we scale education by the associated mean log annual earnings. While it would be more natural to relate test scores directly to wages or earnings, our sample is too young for this exercise to be informative, and we therefore rely on the indirect approach.

Using 2007 data from the American Community Survey (ACS), we calculate the average log annual earnings by years of education for white males born in 1967.[13] The ACS and CNLSY education categories do not line up exactly, particularly among those with more than a high school diploma. We assign all CNLSY observations whom we observe with 13-15 years of education with the average log earnings of those in the ACS who are either college dropouts or associate's degree holders. We likewise assign those with 17 or 18 years of education the average of those in the ACS with a master's degree, and those with more than 18 years of education are assigned the average of doctoral and professional degree holders. We exclude from our calculations those who earn less than $6,000 in salary income.[14] We repeat our estimates replacing years of education with the average log earnings values to compute s_g. Early in our research, we also experimented with a scale based on mean earnings rather than mean log earnings. The results with the two measures were broadly similar, and we did not pursue this approach further.

5 Results

5.1 Estimated Achievement Gaps

Table 3 shows the test score gaps as measured by s_g for each PIAT grade-test. The bootstrapped 95 percent confidence intervals are in brackets. These scores have a clear interpretation with respect to adult outcomes: the average expected educational attainment of children with the black distribution of PPVT scores is .88 years lower than that of children with the white distribution. When measured this way, each of the PIAT tests shows a similar pattern. There is some growth in the gap over the first few years of education, but the gap stabilizes by third grade and remains roughly constant through seventh grade. Blacks, however begin much further behind in math than in reading. Based on their math tests in kindergarten, blacks are expected to obtain .55 fewer years of education than do whites, compared to a gap of only .20 years on the reading recognition test. This difference in the gap closes rapidly, so that by the third grade blacks are .67 years of expected education behind in math and .60 and .61 on the reading recognition and comprehension tests. We

[13] We use 2007 to avoid using earnings data from the recent recession years.

[14] Many of these are small business owners whose income is calculated separately in the ACS.

remind the reader that the reading comprehension results in the earliest grades should be treated with caution because many students in these grades do not perform sufficiently well on the reading recognition test to advance to the comprehension test. Therefore the results for the reading comprehension test are based on a selected sample. Nevertheless the patterns for the two reading tests are similar.

These results do not account for the over-correction that takes place when a shrinkage estimate based on a single test score is applied to a group average. Table 4 reports corrected gaps using the IV strategy discussed above. Strikingly, after correcting for excess shrinkage, the three tests show a consistent story. There is no evidence in any test that the black-white test gap grows over time. On the math test, the kindergarten black-white test gap projects that, on average, blacks will obtain 1.24 fewer years of education than whites do. This gap is substantially larger than the black-white education gap observed in the data, and is consistent with Lang and Manove (2011), who show that blacks obtain more education than whites do conditional on test scores. By seventh grade the gap has, in fact, decreased to .89 years of education, though we cannot reject that it is unchanged. The reading recognition test shows a gap of .64 years of education in kindergarten and remains flat at .68 years in seventh grade.

We are unable to estimate the gap on the reading comprehension test at kindergarten with any precision. While our estimates suggest that this test, as scaled by educational attainment, is mostly noise, we cannot precisely pin down the size of the bias this creates, and thus our confidence interval spans 10 years of education. Using the first grade as our reference point then, we again see no evidence of growth in the test gap through seventh grade. As noted above, however, this is still a somewhat selected sample. Roughly 15% of first graders do not score well enough on the PIAT-RR to take the PIAT-RC. This number is less than 1% in second grade, when we observe a .91 year education gap in performance. From this reference point, the gap falls to .72 by seventh grade, a decline similar to the one on the math test, although this change is again not statistically significant.

It is striking to compare tables 3 and 4. Measurement error and thus the implicit shrinkage on the test declines dramatically as students progress through school. On the math test, the adjustment factor is about 125 percent in kindergarten but only about 25 percent in seventh grade. Similarly, on the reading recognition test the adjustment factor goes from about 3 in kindergarten to .2 in seventh grade.

We note that as Murnane et al (2006) argued and our earlier paper (Bond and Lang, forthcoming) confirmed with other scales, the gap on the early PPVT test is much higher than on the PIAT. Our estimate of the unadjusted gap on the PPVT is .88 years of education. While this is higher than all of our unadjusted gaps, it is somewhat lower than the adjusted gap on the PIAT-M at entry and about the size of some of our early estimates of the reading gap. While we cannot adjust the PPVT gap for measurement error, one plausible explanation for the difference between the early PIAT and PPVT estimates is that the latter test suffers from much less measurement error.

Consistent with this interpretation, the covariance between the PPVT and the two reading tests (see appendix table A) increases sharply between kindergarten and third grade from .21 to .39 for reading recognition and from .14 to .37 for reading comprehension. Note that this is only possible if the PIAT reading tests are doing a better job of capturing skills already acquired by the time the children took the PPVT.[15] In contrast the correlation between the PPVT and math PIAT is roughly constant, going from .34 to .35.

Similarly, we might expect the correlation between child's test score and mother's performance on the Armed Forces Qualifying Test, often used as a measure of general intelligence, would decline as children progress through school. In fact, this correlation increases from kindergarten to second grade for each of the PIAT tests (not shown). While greater measurement error on the kindergarten test than on the second grade test is not the only possible explanation for this regularity, it is surely one of the simplest.

These results show the achievement gap when test scores are calibrated using education and treating all years of education as equally valuable. It is natural to ask whether the results would be similar using other important metrics such as wages or earnings. Unfortunately, the sample of respondents in the CNLSY for whom we have wage data is small and not representative. Therefore, as discussed above, we instead scale education by the earnings associated with each level of education, a non-linear transformation of the education scale.

Table 5 shows the measurement-error corrected results from this exercise. The results confirm the patterns obtained when using completed education to scale the test scores. There is little evidence of a growing achievement gap between blacks and whites. The math test suggests that, given their performance in kindergarten, blacks will earn roughly 17% less than whites do and shows no significant change through seventh grade. While the size of the gap fluctuates across grades, any evidence for a change in the gap is in the direction of blacks catching up rather than falling behind.

The gaps implied by the reading tests are similar and, if anything, lower than those derived from the math test. Still in neither case does table 5 suggest that the gap grows as children progress through school.

Nevertheless, there is also a striking difference between the results in tables 4 and 5.

[15] A possible explanation for this increased ability to capture these skills is that they are more correlated with the more advanced skills of third graders than with the sorts of skills generally developed by the end of kindergarten.

Assuming even a 10 percent return to education, the gaps in math in table 4 suggest a (log) wage gap on the order of .08 to .1. The math gaps in table 5 are all above this range as are the slightly smaller estimated reading gaps. Measured by this dollar metric, the test score gap appears substantially larger.

To address the concern that our results are driven by differences between blacks and whites in both test scores and educational attainment, tables 6 and 7 repeat tables 4 and 5, but use only whites in the calculation of the rescaled test scores. Our estimated gaps are similar whether we include or exclude blacks in the re-scaling of the scores although the latter are less precise.

5.2 Unified Measure of Achievement

Thus far we have considered using adult outcomes as a way to scale the individual subject tests. In this subsection we combine information from all three tests to estimate a single measure of achievement in each grade by forming a conditional expectation of future achievement

$$E[A_{ig}|T_{ig}] = h(T_{ig})$$

where T is the set of tests available for student i and h is the conditional expectation function. Analogous to our earlier discussion, we do not observe achievement directly but observe eventual educational attainment, which reflects achievement in grade g. Following the theory laid out previously, if we can estimate h, we can use instrumental variables to create corrected achievement gaps for each grade.

We estimate h using a multivariate kernel Nadarya-Watson regression estimator. For a set of test scores T, the estimator creates weights for each observation based on the closeness of its test scores to T. The estimator then uses these weights to form a weighted average of the outcome variable (in our case, education). Thus we can generate an expected outcome conditional on the full set of tests.

The weights depend on the choice of kernel function and bandwidth. We select a multivariate Gaussian kernel. For each point, the kernel weights observations around the point so that the density is multivariate normal. The choice of kernel is inconsequential; however the bandwidth is not (Blundell and Duncan, 1998). In a multivariate Gaussian kernel, the bandwidth essentially determines the variance of the density. For bandwidth selection, we follow Silverman's (1986) rule of thumb, so that the bandwidth is proportional to the variance of the distribution in the data.

As previously noted, many children do not advance to the reading comprehension test during the first two years of school. To account for this, we estimate the conditional expectations separately for those who did and did not take the RC exam. In the remaining grades, the very small sample of children who do not advance to the reading comprehension exam is dropped from the analysis.

Table 8 displays the results of this exercise. The first column shows our achievement gap estimates using only the differences in conditional expectations, not adjusting for excess shrinkage. We see the familiar pattern of a rising initial achievement gap. Based on their performance on all tests in kindergarten, blacks are projected to obtain .41 fewer years of education than whites do. This gap rises quickly, however, to .73 years in second grade and remains roughly constant thereafter. In this respect, the results are more similar to those in table 3 for math than for either of the reading tests. This may reflect the poor ability of the early reading tests to predict educational attainment.

However, once we correct for excess shrinkage, the growth in the gap again disappears. The estimates in column two project a future racial difference in educational attainment of .82 years in kindergarten, with little change through seventh grade. Once again, the projected education gap is at least as high if not higher than the actual education gap, consistent with Lang and Manove (2011).

The results in table 8 are broadly consistent with those in table 4. In every grade except 7th, the estimated gap when we use all three tests lies within the range of the gaps produced by using each of the three tests individually. However, the confidence intervals are consistently tighter when we use all three tests, and our estimates appear meaningful even for kindergarten and first grade.

Moreover, the gap averages about .9 years of education, almost exactly what we obtain using the early PPVT test. This suggests that the difference between the results using the PPVT and PIAT tests may not be their content but simply greater measurement error in the latter although we cannot test this directly

In table 9, we repeat the exercise but instead scale the tests to represent the education-predicted mean log earnings of each score. We find similar results to those of table 8. The achievement gap remains steady at about a 12% earnings difference, which is on par with that shown for the math achievement tests in table 5. Our estimates are also generally more precise than in table 5, though the improvement in precision is not nearly as substantial as with the education-scaled scores.

5.3 Correlated Measurement Error

As discussed above, in some cases more than one individual gets the same pair of scores on, for example, the first and third grade math tests. Suppose that Linda and Mike both got 28th percentile in first grade and 36th percentile in third grade. Then both Linda and Mike's eventual education enter the calculation of the mean education associated with a 36 in third grade. Moreover, when we instrument for Mike's third grade education-scaled score with the mean education of everyone else with a 28 in first grade, Linda's education will also enter that calculation. This creates correlated measurement error in finite samples.[16]

To cast light on the importance of this small sample bias for our sample, we ran four simulations in which we took our actual data and added additional error to the education levels. We added a mean zero normal error with standard deviations of 1, 2, 3 and 4. Since the standard deviation of education conditional on test scores is a little less than 2 in most grades, we in effect experimented with increasing the sampling variance by 50-500 percent.

We conducted the simulation 100 times and compared the mean estimates with our actual estimates. The differences caused by this increase in the sampling error were sufficiently modest that in no case were we able to reject that the simulations produced estimates that, on average, were equal to those obtained with the actual data. And the differences between the mean simulated and actual coefficients were also visually modest, suggesting that small sample bias due to correlated measurement error is not a major concern.[17]

5.4 Achievement Gaps and Sociodemographics

One of the key findings in Fryer and Levitt (2004, 2006) was that the early test gap could be "explained" by a small set of sociodemographic controls. Our earlier work (Bond and Lang, forthcoming) showed that while the gaps after controlling for sociodemographic factors were still sensitive to scale choice, they were much more robust than the raw gaps. In tables 9 and 10, we explore the impact of sociodemographics on our "education-scaled" test gaps.

We select a set of sociodemographic controls from the CNLSY to account for differences in the early childhood environment. We include mother's education and age at first birth, and the child's birth weight. We also include a set of controls for the child's home environment from age 0-2: log family income, log hours worked per week by the mother, whether the child ever lived in a household below the poverty line and categorical variables for number of books in the household, amount of cuddly and plush toys, frequency with which the mother reads to the child, whether the child sees a father-figure daily, and frequency of eating dinner with both parents. When we had multiple observations of these variables between age 0 and 2 we used the mean for income and hours worked, and the median category for the categorical

[16] Asymptotically there will be lots of such pairs but their mean deviation from expected education will go to 0, so the IV estimator is consistent.

[17] In one case in the experiment which added N(0,16) error, there was a noticable difference between the mean estimate of the experiment and table 4, but the variance around this estimate was much too large to be meaningful.

variables.[18] From the year in which the test is administered we control for whether the child sees a father-figure daily and whether there are ten or more children's books in the household, as well as family income and mother's hours worked and poverty status. This set of controls is based on the ones used in the CNLSY by Lang and Sepulveda (2008) to closely match those used by Fryer and Levitt (2004, 2006) in the ECLS-K although it is probably somewhat more extensive that the latter.

We compute the education-scaled test scores and their measurement error corrections as before and then add these controls to our regression to estimate the controlled test gap. Tables 10 and 11 show the results for the education- and mean log earnings-scaled test scores, respectively. While we lack precision in our estimates of the education-scaled gaps, there is no evidence that the controlled gap increases with schooling. Our estimates using the mean log earnings-scaled test scores are more precise and tell the same story. Relative to table 5, our controls reduce the gap on every test and in every grade, sometimes quite substantially. In fact, at no point using this scale is the test gap in reading recognition statistically significant once we control for early childhood environment.

One must always be careful in the interpretation of achievement gaps conditional on sociodemographics. As pointed out by Jensen (1969), environment may reflect heritable factors. However, our results in table 11, in particular, suggest that the frequently observed racial test gaps may reflect a common effect of environment on test scores, and not a specific race-based environmental disadvantage.

6 Summary and Conclusions

Tying test scores to educational outcomes has a remarkable effect. Whether we use years of education or the associated mean-log earnings, we find no evidence of a racial component in the evolution of the black-white test gap. As they progress from kindergarten through seventh grades, blacks, on average, perform about as well on achievement tests as predicted by their initial performance. Further, most if not all of the racial difference in achievement can be explained by a modest number of socioeconomic controls. Our findings suggest that the skill gap after kindergarten does not evolve in a race-specific way and that the test score gap may be primarily a sociodemographic gap.

As we have noted, without a common external metric there is no way to distinguish whether the skill gap is larger or small in kindergarten than in seventh grade. Without an external metric, we would have to depend on judgments such as whether the difference

[18] If children had a median category in between two discrete categories, a new category was created for them.

between being able to count and being able to add is greater or less than the difference between the ability to solve a single equation in one unknown and the ability to solve two simultaneous equations in two unknowns. Our approach says that the first gap is larger than the second if and only if it is associated with a greater gap in subsequent skill development. If instead, based on some personal judgment, we decide that the second gap is really larger, our results still tell us that blacks do not fall further behind than can be predicted based on their early scores. In our example, students who can count but not add are likely to becomes students who can solve one but not two equations regardless of their race.

However, the weight of our evidence suggests that the difference between our results and Fryer and Levitt (2006) is not the choice of scale but because early childhood tests are measured with substantial error. Our shrinkage estimators put much less confidence in early childhood tests than in later tests. Either the early tests do not measure something that predicts future educational attainment, or they do so with little precision. Even using other scales, it is likely that the pattern observed by Fryer and Levitt (2006) simply reflects that the ECLS tests are better indicators of the achievement of older children.

We cannot rule out that tying test scores to a different outcome could lead to a different result. Perhaps a scale based on adult wages would show a rising or falling black-white achievement gap through schooling. At this time such data do not exist to make that exercise feasible. The fact that two different interval scales yield the same results gives us some confidence that our results would be robust to other such metrics.

Our results give a new and important perspective on the black-white achievement gap. When measured based on predicted future outcomes, there is no racial component in the evolution of the achievement gap through at least the first eight years of schooling. Previous results that show otherwise are likely to have been due to arbitrary scaling, measurement error, or inconsistent skill measurements.

References

Blundell, Richard and Alan Duncan, "Kernel Regression in Empirical Microeconomics," *Journal of Human Resources,* 33 (1998), 62-87.

Bond, Timothy and Kevin Lang, "The Evolution of the Black-White Test Score Gap in Grades K-3: The Fragility of Results," *Review of Economics and Statistics*, forthcoming.

Boyd, Donald, Hamilton Lankford, Susanna Loeb and James Wyckoff, "Measuring Test Measurement Error: A General Approach," National Bureau of Economic Research Working Paper 18010, 2012.

Braun, Henry I., "A New Approach to Avoiding Problems of Scale in Interpreting Trends in Mental Measurement Data," *Journal of Educational Measurement* 25 (1988), 171-191.

Cascio, Elizabeth U. and Douglas O. Staiger, "Knowledge, Tests, and Fadeout in Educational Intervention," NBER working paper no. 18038 (2012).

Cunha, Flavio and James J. Heckman, "Formulating, Identifying, and Estimating the Technology of Cognitive and Noncognitive Skill Formation," *Journal of Human Resources* 43 (2008), 738-782.

Cunha, Flavio, James J. Heckman, and Susanne M. Schennach, "Estimating the Technology of Cognitive and Noncognitive Skill Formation," *Econometrica* 78 (2010), 883-931.

Farber, Henry S. and Robert Gibbons, "Learning and Wage Dynamics," *Quarterly Journal of Economics* 111 (November 1996):1007-47.

Fryer, Roland G., Jr. and Steven D. Levitt, "Understanding the Black-White Test Score Gap in the First Two Years of School,"*Review of Economics and Statistics* 86 (2004), 447-64.

Fryer, Roland G., Jr. and Steven D. Levitt, "The Black-White Test Score Gap Through Third Grade," *American Law and Economics Review* 8 (2006), 249-81.

Ho, Andrew D., "A Nonparametric Framework for Comparing Trends and Gaps Across Tests," *Journal of Educational and Behavioral Statistics* 34 (2009), 201-228.

Ho, Andrew D. and Edward H. Haertel., "Metric-Free Measures of Test Score Trends and Gaps with Policy-Relevant Examples," CSE Report no. 665 (2006).

Holland, Paul W., "Two Measures of Change in the Gaps Between the CDFs of Test-Score Distributions," *Journal of Education and Behavioral Statistics* 27 (2002), 3-17.

Jencks, Christopher and Meredith Phillips, "The Black-White Test Score Gap: An Introduction," in Christopher Jencks and Meredith Phillips (Eds.), *The Black-White Test Score Gap* (Washington, DC: Brookings Institution Press, 1998).

Jensen, Arthur R., "How Much Can We Boost IQ and Scholastic Achievement?" *Harvard Educational Review* 39 (1969) 1-123.

Junker, Brian, Lynn Steurle Schofield and Lowell J. Taylor, "The Use of Cognitive Ability Measures as Explanatory Variables in Regression Analysis," *IZA Journal of Labor Eco-*

nomics, 1:4 (October 2012).

Lang, Kevin, "Measurement Matters: Perspectives on Education Policy from an Economist and School Board Member," *Journal of Economic Perspectives* 24 (2010), 167-181.

Lang, Kevin and Michael Manove, "Education and Labor Market Discrimination," *American Economic Review*, 101(4) (June 2011): 1467-1496.

Lang, Kevin and Carlos E. Sepulveda, "The Black-White Test Score Differential," unpublished (2008).

Lord, Frederic M., "The 'Ability' Scale in Item Characteristic Curve Theory," *Psychometrika* 40 (1975), 205-217.

Murnane, Richard J., John B. Willett, Kristen L. Bub and Kathleen McCartney, "Understanding Trends in the Black-White Achievement Gaps during the First Years of School," *Brookings-Wharton Papers on Urban Affairs* (2006), 97-135.

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

Reardon, Sean F., "Thirteen Ways of Looking at the Black-White Test Score Gap," Stanford University Institute for Research on Education Policy and Practice working paper no. 2008-08 (2008).

Silverman, B.W., *Density estimation for statistics and data analysis.* Vol. 26. Chapman & Hall/CRC, 1986.

Stevens, S.S. "On the Theory of Scales of Measurement," *Science* 10 (1946), 677-680.

Thorndike, Robert L., "Intellectual Status and Intellectual Growth," *Journal of Educational Psychology*, 57(3) (June 1966): 121-127.

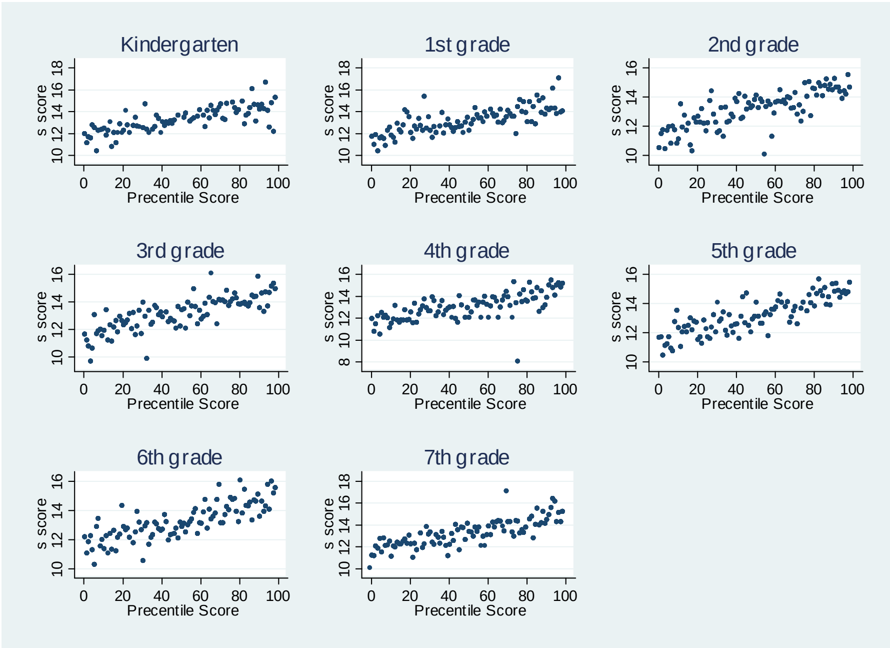

Figure 1: PIAT:M Transformation

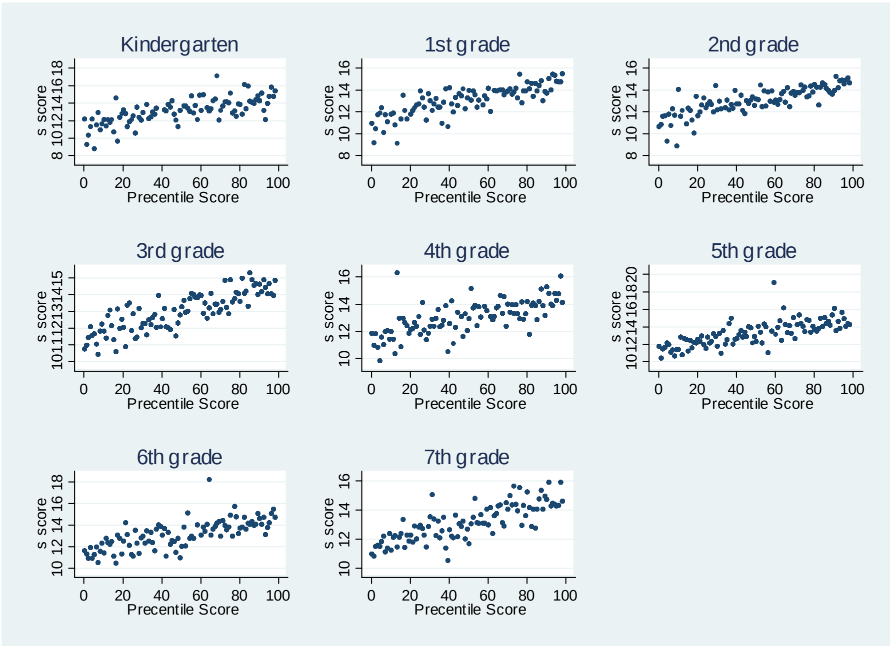

Figure 2: PIAT: RR Transformation

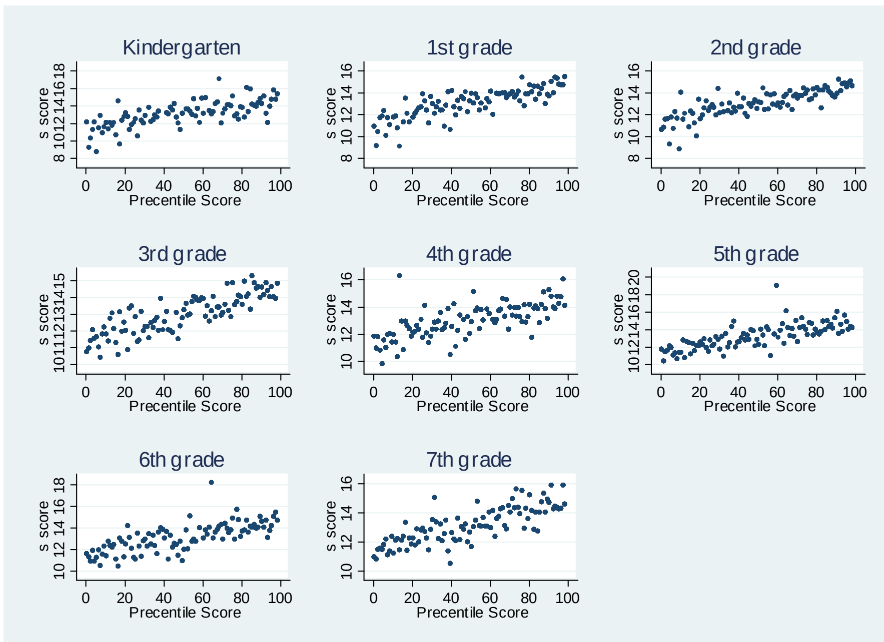

Figure 3: PIAT:RC Transformation

Table 1: Descriptive Statistics - Full Sample

Pre Age-5 PPVT			
Test Gap		1.15	
		(0.05)	
Observations		3657	
	(1)	(2)	(3)
	Math	Read-RR	Read-RC
Kindergarten			
Test Gap	0.65	0.19	0.19
	(0.06)	(0.06)	(0.13)
Observations	2877	2835	1221
First Grade			
Test Gap	0.66	0.42	0.39
	(0.05)	(0.05)	(0.06)
Observations	2893	2888	2478
Second Grade			
Test Gap	0.74	0.60	0.61
	(0.05)	(0.05)	(0.05)
Observations	2858	2885	2781
Third Grade			
Test Gap	0.73	0.62	0.67
	(0.04)	(0.04)	(0.04)
Observations	2889	2861	2811
Fourth Grade			
Test Gap	0.79	0.65	0.66
	(0.04)	(0.04)	(0.05)
Observations	2864	2729	2702
Fifth Grade			
Test Gap	0.71	0.57	0.61
	(0.04)	(0.04)	(0.04)
Observations	2734	2785	2755
Sixth Grade			
Test Gap	0.81	0.64	0.72
	(0.04)	(0.04)	(0.04)
Observations	2590	2594	2572
Seventh Grade			
Test Gap	0.74	0.59	0.69
	(0.04)	(0.05)	(0.04)
Observations	2475	2477	2466

SOURCE: Children of the National Longitudinal Survey of Youth. Test gaps are difference between average white and average black perecentile score measured in standard deviations. Custom weights are used so that each test-grade sample is nationally representative.

Table 2: Descriptive Statistics - Over 22 Sample

Pre Age-5 PPVT			
Test Gap		1.23	
		(0.08)	
Observations		1866	

	(1) Math	(2) Read-RR	(3) Read-RC
Kindergarten			
Test Gap	0.63	0.10	0.06
	(0.07)	(0.08)	(0.16)
Education Gap	0.86	0.85	1.13
	(0.13)	(0.13)	(0.19)
Observations	1480	1446	661
First Grade			
Test Gap	0.67	0.40	0.40
	(0.06)	(0.06)	(0.08)
Education Gap	0.76	0.78	0.83
	(0.11)	(0.11)	(0.13)
Observations	1544	1536	1281
Second Grade			
Test Gap	0.69	0.52	0.52
	(0.06)	(0.06)	(0.06)
Education Gap	0.76	0.72	0.70
	(0.12)	(0.12)	(0.12)
Observations	1638	1633	1572
Third Grade			
Test Gap	0.67	0.57	0.61
	(0.06)	(0.06)	(0.06)
Education Gap	0.85	0.86	0.86
	(0.12)	(0.12)	(0.12)
Observations	1587	1580	1552
Fourth Grade			
Test Gap	0.72	0.58	0.60
	(0.05)	(0.06)	(0.06)
Education Gap	0.77	0.78	0.76
	(0.11)	(0.11)	(0.11)
Observations	1612	1580	1587
Fifth Grade			
Test Gap	0.68	0.52	0.60
	(0.06)	(0.06)	(0.06)
Education Gap	0.84	0.84	0.82
	(0.11)	(0.12)	(0.12)
Observations	1562	1558	1539
Sixth Grade			
Test Gap	0.74	0.56	0.63
	(0.06)	(0.06)	(0.06)
Education Gap	0.85	0.85	0.85
	(0.12)	(0.12)	(0.12)
Observations	1447	1446	1434
Seventh Grade			
Test Gap	0.71	0.57	0.67
	(0.06)	(0.06)	(0.05)
Education Gap	0.67	0.69	0.69
	(0.11)	(0.11)	(0.11)
Observations	1453	1450	1444

Table 3: Raw Difference in Expected Grade Completion conditional on Test Score

	(1) Math	(2) Read-RR	(3) Read-RC
Pre-Age 5 PPVT		0.88	
		[1.12, 0.65]	
Kindergarten	0.55	0.20	0.26
	[0.33, 0.72]	[0.08, 0.36]	[-0.02, 0.52]
First Grade	0.50	0.35	0.32
	[0.37, 0.66]	[0.19, 0.46]	[0.17, 0.48]
Second Grade	0.72	0.58	0.48
	[0.56, 0.96]	[0.36, 0.77]	[0.26, 0.60]
Third Grade	0.67	0.60	0.61
	[0.52, 0.84]	[0.49, 0.78]	[0.46, 0.75]
Fourth Grade	0.70	0.56	0.58
	[0.57, 0.88]	[0.40, 0.71]	[0.44, 0.77]
Fifth Grade	0.69	0.47	0.51
	[0.54, 0.85]	[0.30, 0.61]	[0.33, 0.63]
Sixth Grade	0.70	0.58	0.60
	[0.55, 0.90]	[0.41, 0.79]	[0.40, 0.76]
Seventh Grade	0.71	0.54	0.57
	[0.54, 0.85]	[0.38, 0.70]	[0.44, 0.75]

Difference between average white and average black predicted education conditional on test score for each grade-test combination. Bootstrapped 95 percent confidence intervals in brackets. Conditional predicted education computed for those who are observed over age 22 and applied to the full sample. All results are weighted to be nationally representative.

Table 4: Measurement Error Adjusted Difference in Ability in Units of Predicted Education

	(1) Math	(2) Read-RR	(3) Read-RC
Kindergarten	1.24	0.64	1.32
	[0.65, 2.07]	[0.17, 1.68]	[-3.04, 7.41]
First Grade	1.01	0.88	0.64
	[0.57, 1.54]	[0.38, 1.40]	[0.25, 1.15]
Second Grade	1.05	0.81	0.91
	[0.50, 1.52]	[0.35, 1.37]	[0.43, 1.48]
Third Grade	1.02	0.65	0.69
	[0.47, 1.55]	[0.41, 0.96]	[0.23, 1.09]
Fourth Grade	1.05	0.57	0.71
	[0.68, 1.56]	[0.29, 0.78]	[0.12, 1.06]
Fifth Grade	0.81	0.60	0.67
	[0.52, 1.08]	[0.32, 0.75]	[0.36, 0.87]
Sixth Grade	0.91	0.74	0.81
	[0.62, 1.18]	[0.50, 1.07]	[0.48, 1.06]
Seventh Grade	0.89	0.68	0.72
	[0.54, 1.12]	[0.43, 0.90]	[0.37, 1.19]

Difference between average white and average black predicted education conditional on test score for each grade-test combination corrected for measurement error by instrumental variables. Bootstrapped 95 percent confidence intervals in brackets. Conditional predicted education computed for those who are observed over age 22 and applied to the full sample. All kindergarten and first grade tests, and the second grade Read-RC use predicted education conditional on test score for the PPVT as an instrument, while the remaining tests use that measure lagged two grades. All results are weighted to be nationally representative.

Table 5: Difference in Ability in Education-Predicted Log Earnings

	(1) Math	(2) Read-RR	(3) Read-RC
Kindergarten	0.17	0.09	0.15
	[0.07, 0.40]	[0.00, 0.41]	[-0.36, 0.94]
First Grade	0.12	0.11	0.08
	[0.06, 0.17]	[0.04, 0.19]	[0.03, 0.14]
Second Grade	0.13	0.12	0.12
	[0.06, 0.20]	[0.06, 0.23]	[0.05, 0.23]
Third Grade	0.12	0.09	0.09
	[0.05, 0.18]	[0.05, 0.13]	[0.03, 0.16]
Fourth Grade	0.13	0.07	0.08
	[0.07, 0.20]	[0.02, 0.10]	[-0.01, 0.13]
Fifth Grade	0.11	0.08	0.09
	[0.06, 0.15]	[0.04, 0.10]	[0.05, 0.12]
Sixth Grade	0.11	0.09	0.10
	[0.07, 0.16]	[0.06, 0.14]	[0.05, 0.13]
Seventh Grade	0.12	0.10	0.10
	[0.07 0.15]	[0.05, 0.13]	[0.05, 0.17]

Difference between average white and average black mean log-earnings of predicted education conditional on test score for each grade-test combination corrected for measurement error by instrumental variables. Bootstrapped 95 percent confidence intervals in brackets. Conditional predicted education computed for whites who are observed over age 22 and applied to the full sample. All kindergarten and first grade tests, and the second grade Read-RC use log-earnings predicted education conditional on test score for the PPVT as an instrument, while the remaining tests use that measure lagged two grades. All results are weighted to be nationally representative.

Table 6: Measurement Error Adjusted Difference in Ability in Units of Predicted White Education

	(1) Math	(2) Read-RR	(3) Read-RC
Kindergarten	1.14	0.62	1.26
	[0.58, 2.17]	[0.12, 2.27]	[-4.23, 8.81]
First Grade	1.01	0.82	0.74
	[0.56, 1.55]	[0.31, 1.32]	[0.28, 1.26]
Second Grade	1.07	0.87	0.79
	[0.57, 1.57]	[0.29, 1.53]	[0.35, 1.63]
Third Grade	0.97	0.64	0.66
	[0.52, 1.59]	[0.40, 0.99]	[0.24, 1.15]
Fourth Grade	1.12	0.53	0.72
	[0.69, 1.58]	[0.28, 0.76]	[0.02, 1.04]
Fifth Grade	0.79	0.55	0.63
	[0.51, 1.10]	[0.29, 0.75]	[0.33, 0.82]
Sixth Grade	0.84	0.74	0.75
	[0.54, 1.12]	[0.48, 1.01]	[0.43, 1.02]
Seventh Grade	0.87	0.66	0.68
	[0.51, 1.13]	[0.36, 0.93]	[0.38, 1.20]

Difference between average white and average black predicted education for whites conditional on test score for each grade-test combination corrected for measurement error by instrumental variables. Bootstrapped 95 percent confidence intervals in brackets. Conditional predicted education computed for whites who are observed over age 22 and applied to the full sample. All kindergarten and first grade tests, and the second grade Read-RC use predicted education conditional on test score for the PPVT as an instrument, while the remaining tests use that measure lagged two grades. All results are weighted to be nationally representative.

Table 7: Difference in Ability in Education-Predicted White Log Income

	(1) Math	(2) Read-RR	(3) Read-RC
Kindergarten	0.16	0.09	0.14
	[0.06, 0.41]	[-0.11, 0.47]	[-0.35, 0.81]
First Grade	0.12	0.11	0.09
	[0.05, 0.19]	[0.03, 0.20]	[0.03, 0.17]
Second Grade	0.13	0.14	0.10
	[0.06, 0.20]	[0.06, 0.27]	[0.03, 0.22]
Third Grade	0.12	0.09	0.09
	[0.04, 0.21]	[0.06, 0.13]	[0.02, 0.17]
Fourth Grade	0.13	0.06	0.09
	[0.07, 0.20]	[0.02, 0.10]	[-0.02, 0.14]
Fifth Grade	0.11	0.08	0.09
	[0.07, 0.15]	[0.04, 0.10]	[0.04, 0.12]
Sixth Grade	0.11	0.09	0.09
	[0.06, 0.15]	[0.05, 0.13]	[0.05, 0.13]
Seventh Grade	0.11	0.09	0.10
	[0.06, 0.15]	[0.05, 0.13]	[0.04, 0.17]

Difference between average white and average black mean log-earnings of predicted education for whites conditional on test score for each grade-test combination corrected for measurement error by instrumental variables. Bootstrapped 95 percent confidence intervals in brackets. Conditional predicted education computed for whites who are observed over age 22 and applied to the full sample. All kindergarten and first grade tests, and the second grade Read-RC use log-earnings predicted education conditional on test score for the PPVT as an instrument, while the remaining tests use that measure lagged two grades. All results are weighted to be nationally representative.

Table 8: Difference in Predicted Education Using All Tests

	(1) Unadjusted	(2) IV Adjusted
Kindergarten	0.41	0.82
	[0.23, 0.55]	[0.43, 1.12]
First Grade	0.49	0.93
	[0.33, 0.65]	[0.64, 1.25]
Second Grade	0.73	0.98
	[0.56, 0.88]	[0.64, 1.25]
Third Grade	0.75	0.88
	[0.64, 0.91]	[0.70, 1.15]
Fourth Grade	0.79	0.93
	[0.66, 0.98]	[0.73, 1.19]
Fifth Grade	0.71	0.79
	[0.52, 0.82]	[0.54, 0.90]
Sixth Grade	0.80	0.97
	[0.69, 1.01]	[0.77, 1.22]
Seventh Grade	0.74	0.85
	[0.56, 0.88]	[0.62, 1.08]

Difference between average white and average black predicted education conditional on all test scores for each grade-test combination. Bootstrapped 95 percent confidence intervals in brackets. Column 2 estimates are corrected for measurement error by instrumental variables. Conditional predicted education computed for those who are observed over age 22 using a multivariate kernal regression and applied to the full sample. Kindergarten and first grade use the predicted education conditional on test score for the PPVT as an instrument, while the remaining grades use that measure lagged two grades. All results are weighted to be nationally representative.

Table 9: Difference in Education-Predicted Log Income Using All Tests

	(1) Unadjusted	(2) IV Adjusted
Kindergarten	0.05	0.11
	[0.00, 0.08]	[0.01, 0.20]
First Grade	0.07	0.12
	[0.03, 0.11]	[0.05, 0.21]
Second Grade	0.09	0.12
	[0.06, 0.12]	[0.05, 0.21]
Third Grade	0.10	0.12
	[0.07, 0.13]	[0.08, 0.18]
Fourth Grade	0.10	0.11
	[0.07, 0.14]	[0.04, 0.17]
Fifth Grade	0.09	0.11
	[0.04, 0.12]	[0.05, 0.14]
Sixth Grade	0.10	0.12
	[0.07, 0.14]	[0.07, 0.17]
Seventh Grade	0.10	0.12
	[0.06, 0.13]	[0.07, 0.18]

Difference between average white and average black mean log-earnings of predicted education conditional on all test scores for each grade-test combination. Bootstrapped 95 percent confidence intervals in brackets. Column 2 estimates are corrected for measurement error by instrumental variables. Conditional predicted education computed for those who are observed over age 22 using a multivariate kernal regression and applied to the full sample. Kindergarten and first grade use the predicted education conditional on test score for the PPVT as an instrument, while the remaining grades use that measure lagged two grades. All results are weighted to be nationally representative.

Table 10: Conditional Difference in Ability in Units of Predicted Education

	(1) Math	(2) Read-RR	(3) Read-RC
Kindergarten	0.61	-0.04	0.55
	[-0.12, 1.52]	[-0.86, 1.21]	[-5.41, 5.49]
First Grade	0.38	-0.19	0.01
	[-0.34, 0.83]	[-1.07, 0.38]	[-0.66, 0.41]
Second Grade	0.53	0.11	0.06
	[0.02, 1.11]	[-0.53, 0.53]	[-0.79, 0.55]
Third Grade	0.63	0.22	0.34
	[0.13, 1.23]	[-0.06, 0.60]	[0.01, 0.73]
Fourth Grade	0.39	0.22	-0.02
	[-0.00, 0.98]	[-0.13, 0.49]	[-0.46, 0.43]
Fifth Grade	0.50	0.31	0.69
	[0.17, 1.03]	[-0.12, 0.74]	[0.30, 1.25]
Sixth Grade	0.26	0.18	0.26
	[-0.20, 0.73]	[-0.26, 0.61]	[-0.21, 0.76]
Seventh Grade	0.48	0.08	0.26
	[0.07, 0.83]	[-0.50, 0.36]	[-0.14, 0.71]

Opposite of coefficient on black indicator in regression on predicted education conditional on test score for each grade-test combination corrected for measurement error by instrumental variables. Boostrapped 95 percent confidence intervals in brackets. Each regression includes controls for mother's education and age at first birth, child's birthweight, and household conditions at age 2 including log family income, log mother's hours worked, books, frequency of mother reading to child, mother's philosophy on children's learning, amount of toys in the household, whether the child sees the dad daily, and frequency of eating dinner with both parents Conditional predicted education computed for those who are observed over age 22 and applied to the full sample. All kindergarten and first grade tests, and the second grade Read-RC use predicted education conditional on test score for the PPVT as an instrument, while the remaining tests use that measure lagged two grades. All results are weighted to be nationally representative..

Table 11: Conditional Difference in Ability in Education-Predicted Log Income

	(1) Math	(2) Read-RR	(3) Read-RC
Kindergarten	0.09	-0.02	0.00
	[-0.01, 0.27]	[-0.17, 0.16]	[-0.88, 0.89]
First Grade	0.04	-0.02	-0.00
	[-0.06, 0.10]	[-0.16, 0.06]	[-0.10, 0.06]
Second Grade	0.06	0.02	0.01
	[-0.01, 0.14]	[-0.08, 0.09]	[-0.10, 0.09]
Third Grade	0.07	0.03	0.04
	[0.00, 0.15]	[-0.01, 0.08]	[0.00, 0.10]
Fourth Grade	0.04	0.03	0.00
	[-0.01, 0.12]	[-0.01, 0.07]	[-0.05, 0.06]
Fifth Grade	0.07	0.04	0.09
	[0.02, 0.14]	[-0.02, 0.07]	[0.04, 0.17]
Sixth Grade	0.04	0.03	0.04
	[-0.02, 0.10]	[-0.03, 0.09]	[-0.02, 0.11]
Seventh Grade	0.06	0.01	0.04
	[-0.00, 0.11]	[-0.08, 0.06]	[-0.02, 0.10]

Opposite of coefficient on black indicator in regression on log-earnings of predicted education conditional on test score for each grade-test combination corrected for measurement error by instrumental variables. Bootstrapped 95 percent confidence intervals in brackets. Each regression includes controls for mother's education and age at first birth, child's birthweight, and household conditions at age 2 including log family income, log mother's hours worked, books, frequency of mother reading to child, mother's philosophy on children's learning, amount of toys in the household, whether the child sees the dad daily, and frequency of eating dinner with both parents Conditional predicted education computed for those who are observed over age 22 and applied to the full sample. All kindergarten and first grade tests, and the second grade Read-RC use log-earnings of predicted education conditional on test score for the PPVT as an instrument, while the remaining tests use that measure lagged two grades. All results are weighted to be nationally representative..

Table A: Test Score Covariance Matrix

	PPVT	Grade K	Grade 1	Grade 2	Grade 3	Grade 4	Grade 5	Grade 6	Grade 7
PPVT	0.81								
MATH									
Grade K	0.34	1.04							
Grade 1	0.32		1.02						
Grade 2	0.39	0.52		1.31					
Grade 3	0.35		0.52		1.23				
Grade 4	0.34	0.45		0.66		1.20			
Grade 5	0.37		0.51		0.71		1.34		
Grade 6	0.40	0.46		0.64		0.66		1.29	
Grade 7	0.36		0.50		0.65		0.74		1.26
READING RECOGNITION									
Grade K	0.21	1.07							
Grade 1	0.28		1.09						
Grade 2	0.33	0.40		1.11					
Grade 3	0.39		0.67		1.23				
Grade 4	0.30	0.37		0.66		1.23			
Grade 5	0.35		0.61		0.77		1.50		
Grade 6	0.34	0.37		0.70		0.69		1.39	
Grade 7	0.35		0.61		0.77		0.81		1.43
READING COMPREHENSION									
Grade K	0.13	0.78							
Grade 1	0.29		1.05						
Grade 2	0.35	0.19		1.07					
Grade 3	0.37		0.49		1.24				
Grade 4	0.29	0.12		0.52		1.16			
Grade 5	0.32		0.43		0.56		1.14		
Grade 6	0.41	0.20		0.54		0.55		1.28	
Grade 7	0.39		0.37		0.56		0.54		1.16

Covariances are calculated using all available observations for each individual cell and are unweighted. Covariances of tests taken 1, 3, 5, or 7 years apart are not shown because the sample is surveyed every two years.