The Distribution of Teacher Quality and Implications for Policy

Eric A. Hanushek and Steven G. Rivkin

December 2011

Abstract

It has become commonplace to measure teacher quality in terms of teacher value-added. Operationally this means evaluating teachers according to the learning gains of students on various achievement tests. Existing research consistently shows large variations in teacher effectiveness, much of which is within schools as opposed to between schools. The policy implications of these variations are dramatic. But, the underlying statistical modeling has become the subject of intense research, in part because of the direct use of value-added measures in policy discussions.

Eric A. Hanushek
Stanford University, University of Texas at Dallas, NBER, and CESifo
hanushek@stanford.edu

Steven G. Rivkin
University of Illinois at Chicago, University of Texas at Dallas, NBER, and CESifo
sgrivkin@amherst.edu

1. Introduction

Children, parents, administrators, and policy-makers all tend to focus on the quality of teaching as the key determinant of school quality, but the failure of observed teacher characteristics, including education and experience, to account for much of the variation in student test scores or other outcomes presented researchers with an apparent conundrum. Although the absence of a strong relationship between outcomes and these characteristics is consistent with teacher quality not being an important determinant of learning, it is also consistent with the possibility that these quantifiable characteristics are simply poor measures of teacher effectiveness. In fact recent outcome-based estimates find substantial variation in teacher contributions to achievement, supporting the latter interpretation and the general emphasis on teacher quality.

Easily quantifiable characteristics explain little of the variation in teacher effectiveness, and this has important implications for the development of policies designed to raise the quality of instruction and to reduce unequal access to high quality teachers. First, neither a graduate degree nor additional years of experience past the initial year or two translate into significantly higher instructional effectiveness, bringing into question a salary structure based almost entirely on these two variables. Second, descriptions of unequal access to quality teachers as measured by experience, education, or other quantifiable characteristics fail to portray accurately any actual differences in the quality of instruction by student demographics, community characteristics, and specific schools. Third, the failure of quantifiable characteristics to explain much of the variation in teacher effectiveness suggests that efforts to raise the quality of instruction through more

stringent requirements for entering the teaching profession may be seriously misguided, particularly as they may discourage many from entering the profession by raising the cost of becoming a teacher. Fourth, the focus on student outcomes – highlighted by state and federal accountability systems including No Child Left Behind (NCLB) – has led to attempts to legislate better teachers, although these efforts have not circumvented the problems of defining an effective teacher.[1]

The analysis of teacher effectiveness has largely turned away from attempts to identify specific characteristics of teachers. Instead attention has focused directly on the relationship between teachers and student outcomes. This outcome-based perspective, now commonly called value-added analysis, takes the perspective that a good teacher is simply one who consistently gets higher achievement from students (after allowing for other sources of student achievement such as family influences or prior teachers). The underlying analysis has focused on statistical estimation that separates teacher influences from other factors, and most typically it has relied upon administrative data from schools.

The current discussion of value-added measures of teacher effectiveness is not so much about the conceptual approach as it is about the details of the analysis and the interpretation of the results. The debate has been heightened by the direct introduction of value-added estimates into teacher evaluations and personnel decisions (see, for example, Isenberg & Hock (2010)) and by public identification of teachers according to value-added in the newspapers (Song & Felch (2011)). This interaction with actual policy decisions plus the larger focus of recent education reforms on student outcomes has contributed to growing scholarly discussion of the validity of

---

[1] NCLB introduced a requirement for "highly qualified teachers" within schools serving disadvantaged students. This requirement was phrased in terms of qualifications as opposed to effectiveness in the classroom, and the definitions of highly qualified were left up to the separate states. As a result, most states simply inserted variants of the existing requirements for teacher certification. See Chubb (2009), Hanushek (2009).

outcome-based, or value-added, estimates of teacher quality. Rothstein (2010) provides evidence that sorting both among and within schools may bias estimates of teacher effectiveness, and Meghir & Palme (2005) emphasize the possibility that school and family responses endogenous to realized teacher quality contribute to observed achievement. Other analyses have emphasized the importance of measurement error in using test outcome data (e.g., Kane & Staiger (2002), McCaffrey et al (2009)).

We consider the existing evidence on teacher quality, the methods used to produce those estimates, and implications of the findings for the development of teacher policies. Section 2 summarizes baseline estimates of the variance in teacher quality based on estimates of teacher fixed effects and other empirical approaches. Section 3 issues related to the measurement of outcomes. Section 4 discusses potential impediments to the identification of the teacher quality variance caused by purposeful sorting of children and teachers into neighborhoods, schools, and classrooms and the endogenous responses by parents and school administrators to observed teacher effectiveness. Section 5 describes research that considers persistence of teacher effects over time. Section 6 provides a discussion of selected implications for education policy, while Section 7 has some final thoughts about teacher policy.

2. Outcome-based Estimates of the Variance in Teacher Effectiveness

Research into the determinants of student outcomes and the role of teachers has changed dramatically over the past two decades. The earliest work, generally traced back to the "Coleman Report" (Coleman et al (1966)), was typically cross-sectional analysis that relied on measuring the various inputs to education – from families, peers, and schools. Additionally, much of the early analysis relied upon data collected for other purposes and not specific to educational performance. More recently, the analysis has changed dramatically as data have

improved, as researchers have extended the simple analyses in a variety of ways, and as replication of various analyses has become more feasible.

From the start, the impact of teachers was simply viewed one of many possible inputs to achievement. The consistent finding over four decades of research – frequently called education production function research in economics – has been that the most commonly used indicators of quality differences are not closely related to achievement gain, leading some to question whether teacher quality really matters (see the review in Hanushek & Rivkin (2006)). [2] In reality, interpretations of research on teachers (and other school inputs) often confused the estimated effects of specific teacher characteristics with the overall contribution of teachers.

The focus on the measurement of teacher value-added to student achievement represents a shift from a research design focused on the link between student outcomes and specific teacher characteristics to a research framework using a less parametric approach to identify overall teacher contributions to learning. Using administrative data bases, some covering all of the teachers in a state, such research provides strong support for the existence of substantial differences in teacher effectiveness, even within schools.[3] Although this approach circumvents the need to identify specific teacher characteristics related to quality, the less parametric approach introduces additional complications and has sparked an active debate on the measurement and subsequent policy use of estimated teacher value-added.

The precise method of attributing differences in classroom achievement to teachers is the subject of considerable analysis as discussed below. We begin by briefly outlining the general

---

[2] The Coleman Report (Coleman et al (1966)) is generally acknowledged as the beginning of this entire line of research into the determinants of student achievement. It was commonly interpreted as finding that "schools don't matter;" instead, families and peers are the prime determinants of student achievement. While this study was deeply flawed (Hanushek & Kain (1972)), portions of this conclusion still are included in public policy discussions.
[3] The earliest academic research includes Hanushek (1971), Murnane (1975), and Armor et al (1976). Policy interest rose with the introduction of the ideas directly into state evaluations (Sanders & Horn (1994)).

analytical framework that forms the basis of much of the work in this area and then describe the range of results from recent efforts to measure the variance of teacher effectiveness.

As the available research discusses in great detail, the determinants of both student and teacher choices and the allocation of students among classrooms complicate efforts to isolate the contributions of teachers to learning. The discussion highlights the importance and difficulty of accounting for student and school differences that affect both own performance and the classroom environment, and the likely advantages of a value-added modeling approach that accounts explicitly for differences in the history of family, school, and community influences. It also recognizes, however, that the value-added framework does not address all potential impediments to consistent estimation of the quality of instruction and the variance in teacher quality. Specific methods mitigate some deficiencies and not others, although none may completely resolve the potential problems.

This section outlines the decision making processes of families, teachers, and principals and the potential implications of these choices for the estimation of teacher quality, presents an empirical framework that captures potential biases from these choices and other factors, and ends with a discussion of empirical evidence on the validity of specific estimation methods.[4] A central issue in the empirical analysis is that the relevant participants – families, teachers, and administrators – are not randomly assigned to schools and classrooms. Families choose a community and school, possibly trading off school quality with other housing amenities (see Hanushek & Yilmaz (2007), Hanushek et al (2011)). Teachers and administrators also make choices along a number of dimensions including geography, school, sector, and effort. These

---

[4] Ballou et al (2004) and McCaffrey et al (2004) provide general discussions of value-added estimation of teacher effects.

choices must be taken into account in order to identify the variation in teacher value-added to student achievement.

The basic educational model implicitly underlying educational production functions is that:

(1) $$A_{gi} = f(S_i, X_i, \mu_i)$$

where A_g is the achievement of student i in grade g (the subscript i is suppressed when not central to the discussion), **S** is a vector of school and peer factors, **X** is a vector of family and neighborhood inputs, and μ is individual ability.

One important facet of that enters is that education is a cumulative process. Achievement in grade g involves not only educational inputs in that grade but also the whole history of inputs that provided the basic knowledge that enters into the summative achievement in grade g. To fix some basic ideas and to match much of the empirical analysis, assume a linear model such that:

(2) $$A_{Gi} = \sum_{g=0}^{G} S_{gi}\varphi_g + \sum_{g=0}^{G} X_{gi}\gamma_g + \sum_{g=0}^{G} \mu_i + \varepsilon_i$$

where S_g and X_g are the vectors of school and family inputs during grade g, μ is assumed to be the constant input of ability, the φ and γ are unknown parameters, ε is a stochastic term representing unmeasured influences, and we observe how inputs accumulate through grade G.[5] Then, we can decompose this into grade G and prior grades. Further, if we put some structure on the parameters such that they follow a geometrically declining pattern for inputs in the more distant past – indicating that the impact of past inputs depreciates at a constant rate θ over time, we get an expression such as:

---

[5] Developments along these lines can be found in Boardman & Murnane (1979), Hanushek (1979), and, more recently, Todd & Wolpin (2003).

$$A_{Gi} = S_{Gi}\varphi + X_{Gi}\gamma + \mu_i + \varepsilon_{Gi}$$

$$(3) \quad + \sum_{g=0}^{G-1} S_{gi}\varphi(1-\theta)^{G-g} + \sum_{g=0}^{G-1} X_{gi}\gamma(1-\theta)^{G-g} + \sum_{g=0}^{G-1} \mu_i(1-\theta)^{G-g} + \sum_{g=0}^{G-1} \varepsilon_{gi}(1-\theta)^{G-g}$$

Note that, if this relationship holds across grades, the second line is simply $(1-\theta)A_{G-1,i}$ so that

we can write current achievement as a function of (depreciated) past achievement plus the inputs

during grade G:

$$(4) \quad A_{Gi} = S_{Gi}\varphi + X_{Gi}\gamma + (1-\theta)A_{G-1,i} + (\mu_i + \varepsilon_{Gi})$$

This expression for achievement obviously puts considerable structure on the relationship,

something that has been investigated in various forms as indicated below.

This expression is particularly useful because few if any data sets actually track the entire

past history of family and school inputs. Equation (4) suggests that it is possible to concentrate

on the inputs in grade G with A_{G-1} capturing the historical inputs.[6]

Because many data sets including administrative data have limited information on family

income and other family characteristics including parental education that are likely related to

both commitment to schooling and home resources available to support education, it is often

difficult to control directly for family heterogeneity. The availability of multiple observations for

students in panel data sets, however, makes possible alternative approaches to account for

student heterogeneity. Specifically, a variety of studies use measures of prior achievement,

---

[6] Again, this interpretation involves a number of assumptions about the additive nature of inputs. For example, it does not allow for the "one great teacher" who has a lasting impact on the pattern achievement (over and above the effect on achievement in the one year). It also assumes that knowledge from all sources depreciates at the same rate, something that comes up below in considering the persistence of teacher effects.

7

student fixed effects, or a combination of the two to account for stable differences among students not captured by the limited set of available variables (as suggested by equation (4)).

Empirical analyses of teacher value-added typically begin with a slightly modified version of the education production function in equation (4) that breaks out a term for teacher inputs, τ_j:

$$(5) \qquad A_G = (1-\theta)A_{G-1} + \tau_j + \mathbf{S}\boldsymbol{\varphi} + \mathbf{X}\boldsymbol{\gamma} + \varepsilon$$

where τ_j is a teacher fixed effect that provides a measure of teacher value-added for teacher j.[7] Most empirical analyses, while not explicit, also subsume ability effects (μ) from equation (4) into a composite error, ε, although some of the discussion of sorting into classrooms can be thought of as focusing on these ability effects.[8] Further, various analyses focus on the precise specification of the τ_j, such as whether it is the same for all students or there is some heterogeneous treatment effects.[9]

While we discuss further details of the estimation below, we begin simply with available estimates of the variation in teacher value-added (i.e., variation in τ_j).[10] The most conservative estimates try to guard against biases from unmeasured characteristics of families and teachers by concentrating on just the differences among teachers within a given school. This approach holds constant the aggregate motivation of families to be in a given school at the cost of ignoring any

---

[7] Alternative estimation forms, largely restricting θ, have pluses and minuses but are currently less frequently employed; see Meghir & Rivkin (2011) on conceptual issues and Hanushek et al (2009) on empirical implications.
[8] For the most part, all estimates employ OLS estimation in a panel context.
[9] One issue that is not explicitly considered here is that of peer interactions and the effect of these on achievement. While there is a substantial literature on this (see the review in Sacerdote (2011)), the consideration of peer interactions and of teacher quality have been almost completely separated in the literature – and we follow that history in this discussion.
[10] Again, this may be interpreted as the average treatment effect, and – as discussed below – some attempts have been made to look at specific dimensions of heterogeneous effects. These differences appear relatively small, however, compared to the overall estimated variation in teacher value-added.

variation in average teacher effectiveness across schools.[11]

Table 1 summarizes existing estimates of the standard deviation of teacher effectiveness (σ_τ) expressed in units of student achievement (normalized to a standard deviation of one).[12] Although covering a range of schooling environments across the U.S., these studies produce fairly similar estimates of the variance in teacher value-added: the average standard deviation for reading is 0.13 and for math is 0.17 and the distributions for both are fairly tight. Note also these estimates rely on just within-school variation in value-added, ignoring the surprisingly small between-school component. The between-school component is not typically considered because of potential sorting, testing, and other interpretative problems.[13]

The magnitudes of these estimates support the beliefs that teacher quality is an important determinant of school quality and achievement. For example, the math results imply that having a teacher at the 25[th] percentile as compared to the 75[th] percentile of the quality distribution would mean a difference in learning gains of roughly 0.2 standard deviations in a single year. This would move a student at the middle of the achievement distribution to the 58[th] percentile. The magnitude of such an effect is large both relative to typical measures of black-white or income achievement gaps of 0.7-1 standard deviation and compared to methodologically compelling estimates of the effects of a ten student reduction in class size of 0.1-0.3 standard deviations. We discuss the economic value of these differences in teacher effectiveness below.

---

[11] Families have been seen to choose schools based on school value-added (see Hanushek et al (2004a). The differences in demand shows up in housing price variations, although the measurement of school quality varies widely across studies (see Black & Machin (2011))

[12] Some questions remain about whether performance on different tests can be compared through this normalization. At the very least, such comparisons require assumptions about the psychometric property of the tests such as normality of the distribution. Some analysis, as described below, considers such issues directly.

[13] The study by Kane & Staiger (2008) is the one exception that does not exclude the between school component. Hanushek & Rivkin (2010b) provide estimates of the within and between school variation in value-added under different specifications.

3. Outcome Measurement Issues[14]

Test measurement error can complicate the outcome-based estimation of teacher quality. Not only does the test error introduce noise into the estimates of teacher effectiveness, but it also biases upward estimates of the variance in teacher quality. The multiple decisions that determine the allocation of students to teachers must be accounted for in order to produce consistent estimates of teacher quality. In addition, any endogenous responses of parents and administrators to realized teacher quality, responses that are generally not observed, likely move estimates of teacher effectiveness away from their true values.

Four test measurement issues receive considerable attention in both the research literature and broader policy discussion: 1) random measurement error; 2) focus of tests on particular portions of the achievement distribution; 3) cardinal v. ordinal comparisons of test scores; and 4) multidimensionality of education outcomes.

The fixed-effect estimator is the conditional mean performance for a teacher's students, given the prior achievement levels of the students and other factors explicitly considered, i.e., the S and X. This estimated value of teacher quality, $\hat{\tau}_j$, is the true value of teacher quality, τ_j, plus estimation error, v_j:

$$(6) \qquad \hat{\tau}_j = \tau_j + v_j$$

Considerable attention has been devoted to the role measurement error for student achievement.[15] If the only error in equation (5) comes from error in the individual measurement of achievement, the random test measurement error introduces sampling error into the estimates of the teacher fixed effects, v_j. When the test measurement errors are uncorrelated with true

---

[14] Ishii & Rivkin (2009) expand upon many of the issues considered in this and the next section.
[15] This attention to test measurement error can be traced back to early work on school accountability by Kane & Staiger (2002).

teacher quality and the expected value of v_j equals zero, the estimates of teacher effects will be unbiased, but there will be an upward bias in the estimated variance of teacher quality. In particular, the estimated variance based on raw test scores would equal $\mathrm{var}(\hat{\tau}) = \mathrm{var}(\tau) + \mathrm{var}(v)$. In this case, $\mathrm{var}(v)$ is a direct function of the variance of test measurement error.

All of the estimates of the variance in teacher effectiveness in Table 1 were adjusted in one way or another for test measurement error. One common approach is the empirical Bayes approach that adjusts the estimates by producing weighted averages of the actual estimates and the grand mean, where the weights are a function of the estimated variance.[16] (See, for example, Kane & Staiger (2002) and Jacob & Lefgren (2008)). The impact of this adjustment is to pull each of the individual estimates toward the grand mean and, as a result, reducing the variance of the estimates. As Kain and Staiger illustrate in a slightly different context, sample size is the primary determinant of differences in the error variance; consequently the grand mean tends to receive higher weight for teachers with fewer students.

Alternatively, assuming that the confounding measurement error is uncorrelated across years, the correlation of the fixed effect estimates for the same teachers across years would equal the ratio of the variance of true teacher quality divided by the variance of estimated teacher quality, $r_{12} = \mathrm{var}(\tau) / \mathrm{var}(\hat{\tau})$. Multiplication of the estimated variance of $\hat{\tau}$ by the estimated year-to-year correlation would provide a consistent estimate of the overall variance in teacher quality corrected for the random contributions of test error and the other factors (see Hanushek et al (2005), McCaffrey et al (2009)).

A potential problem with efforts to remove the non-teacher quality components of the variance estimate is variation in teacher quality over time or across students. For the most part,

---

[16] Morrison (1983) describes the empirical Bayes shrinkage estimator.

studies have defined teacher quality as the stable component of teacher effects on students. But, in a year when a teacher simply has a "bad year" for one reason or another, the teacher's impact on the students reflects this problem. The seriousness of this issue, perhaps including whether alternative definitions of teacher quality should be employed, is currently unknown – in part because it is difficult analytically to separate variations in teacher quality from random errors not associated with the teachers.[17] Similarly, heterogeneity in teacher quality by student demographic characteristics (perhaps due to the benefits of a same-race teacher) or skills (some teachers may be more effective with students in particular parts of the achievement distribution) will tend to increase the standard error in models that impose homogeneous effects.[18] The standard focus on the "persistent" component of teacher quality ignores heterogeneity along these dimensions, which may be appropriate depending on the use of the estimates.[19]

The second and third issues concerning assessment differences across the distribution and the appropriate interpretation of test results are related in both substance and proposed solutions. First, there are complications introduced by the coverage of a test. If the degree of overlap between the curriculum and test varies, then estimates of value-added will not produce consistent rankings of teacher effectiveness. Consider first a test that focuses on basic skills and a class filled with students working above grade level. In this case efforts and efficacy of the current teacher have no virtually no impact on test results, and thus the estimate of value-added for that teacher would not be based on true performance. More generally, the less overlap between test coverage and curriculum the less informative are the estimates.

---

[17] In some instances, researchers have attempted to consider whether the impact of a teacher changes with a move to another schools, thus allowing for variations in teacher quality; e.g., Jackson & Bruegmann (2009b), Goldhaber et al (2011, Jackson & Bruegmann (2009a). They provide mixed evidence on any changes in teacher performance.
[18] See Dee (2005), Hanushek et al (2005).
[19] When the approach to measurement error is an empirical Bayes method, ignoring these components tends to lead to greater shrinkage toward the grand mean for particular teachers.

Second, consider assumptions related to test scaling and meaning of test score changes in various parts of the test distribution. There is an ongoing debate among psychometricians over the construction of vertically scaled tests for which a one unit change in score has similar meaning across the distribution in terms of actual knowledge gained. (See, for example, Koretz (2008), Braun et al (2010)). Many argue that even those tests assumed to be vertically scaled actually do not satisfy this criterion, and many tests administered in schools do not even claim to have this property. Such concerns have led some to focus solely on the ordinal properties of test results, for example on percentile scores. Yet percentile score measures also do not avoid assumptions about scaling, as it takes a larger change in knowledge to move a given number of percentile points in the tails of the distribution. Therefore two similar effective teachers may produce very different average changes in student percentile rankings with different distributions of students.

One approach designed to address both scaling and test focus concerns is to compare each student with all students in the sample with similar prior test score performance.[20] There are several different manifestations of this type of model, but essentially teachers would be judged by how their students perform relative to similar students in other classrooms. For example, one could estimate a model of teacher fixed effects in which the percentile score would be regressed on dummy variables for each percentile of the prior year score (and potentially scores in previous years or other subjects), other controls and a full set of teacher dummy variables. Alternatively quantile treatment effect methods could be used to calculate the median (or other) score percentile given the prior year distribution of achievement. Percentile growth models, while discussed here for the evaluation of teacher effectiveness, have been introduced into the

---

[20] See, for example, Neal (2011). Moreover, such comparisons are embedded into some state accountability system such as Colorado (http://www.cde.state.co.us/index_home.htm).

accountability system of states (see, for example, the Colorado accountability system:

http://www.cde.state.co.us/index_home.htm).

The final issue concerns the interpretation of estimates of teacher productivity. Education

is a multi-dimensional enterprise in which schools seek to foster learning along a number of

academic as well as non-academic dimensions. Because teachers allocate time among various

endeavors in pursuit of a number of objectives, estimates of teacher quality should ideally

incorporate the time allocated to the type of instruction in question in order to generate

meaningful measures of effectiveness. If it takes one teacher twice as long as another to produce

a given amount of mathematics learning, that teacher should be judged as half as productive in

this area. Yet in reality a lack of information on teacher time allocation prevents adjustments for

time spent on a given task.[21]

4. Other Impediments to the Estimation of Teacher Quality

Although the outcome measurement issues have been rather thoroughly addressed,

another set of issues evokes considerable current discussion and disagreement. In particular, the

myriad decisions of families, administrators, teachers, and students that determine the

distributions of students and teachers in schools and classrooms raise concerns about the validity

of estimates of teacher value-added. This section begins with a discussion of the various levels of

choices and then turns to a discussion of the empirical evidence.

Purposeful Choices of Families, Teachers, and Schools

Residential location, school choice, classroom allocation, and academic support decisions

complicate the estimation of teacher quality, and researchers have used a variety of methods to

isolate the contributions of teachers. These include the use of student fixed effects and prior test

---

[21] Little analytical work has gone into identifying multiple outcomes, although it has entered directly into policy applications where alternative teacher assessments enter even where value-added measures are available and where multiple value-added measures are combined. See, for example, the IMPACT evaluation system in Washington, DC: http://www.dc.gov/DCPS/impact [accessed December 11, 2011].

scores to account for unobserved student heterogeneity and school fixed effects to account for unobserved differences in a range of school factors. Yet concerns remain that the endogenous sorting of children among classrooms biases estimates of both teacher quality and its variance. For example, Rothstein (2010) focuses on potential biases introduced by sorting on the basis of time-varying unobservables.[22]

The direction of any bias would depend upon the character of that sorting which would be determined largely by the principal's objective function. An egalitarian principal might place more disruptive children with a higher quality teacher, while a principal that desires to please the senior staff might give experienced teachers the more compliant, intellectually engaged children. These alternative allocation mechanisms produce very different distributions of achievement within schools and different patterns of bias in the absence of adequate controls for classroom heterogeneity.

The direction of any bias also depends upon the character of teacher sorting among communities and schools. Evidence on teacher preferences suggests that teachers tend to prefer schools with higher achieving students and appear to have heterogeneous preferences regarding school location and characteristics related to student race and ethnicity.[23] Survey evidence also suggests that principal behavior influences the probability a teacher remains in a school, which potentially introduces a link between teacher quality and principal behavior.[24] Together these findings suggest the existence of a positive correlation between teacher supply and factors that contribute to higher achievement.[25] To the extent that schools exploit superior position in the

---

[22] Note that this is not a story of peer interactions but one of unmeasured student attributes associated with teachers that lead to bias in equation (6).
[23] Boyd et al (2005) and Hanushek et al (2004b) describe differences in teacher sorting by race.
[24] Loeb et al (2005) discuss the relationship between teacher turnover and working conditions.
[25] Hanushek & Rivkin (2007).

teacher labor market by employing more effective teachers, unobserved student heterogeneity is likely to inflate the between school component of the estimated variance in teacher quality.

The possibility of endogenous school and family responses to realized teacher quality complicates estimation even further. In the absence of information on the prevalence of intervention teachers or family support, estimates of teacher effectiveness combine the actual teacher effect with these other influences. In the case of classroom support, it is likely that the allocation of resources is compensatory, meaning that it would bias estimates of teacher productivity toward the mean and estimates of the variance in teacher quality downward. With regards to family influences, Kim (2001), Todd & Wolpin (2003), and Meghir & Palme (2005) point out that estimates of school and teacher effects incorporate family responses that vary systematically with input quality. If families tend to contribute more time and money to academic support during periods in which they perceive the teacher as inadequate or of low quality, this will tend to bias downward estimates of the variation in teacher quality and bias estimates of the quality of specific teachers toward the mean. However, if one component of high quality teaching is raising the level of difficulty, higher quality teachers may actually increase the likelihood that students seek help outside of school, and estimates of teacher fixed effects could be biased away from zero and the estimated variance could be biased upwards. In either case, it is not adequate to represent families as fixed inputs into the education production process, meaning that even the inclusion of student fixed effects could fail to account for potentially important family influences.

Empirical Evidence on Validity of Estimates

This section considers evidence on the validity of non-experimental estimates of teacher value-added. We begin with a paper by Kane and Staiger that finds evidence in support of the validity of value-added estimates of teacher fixed effects from an analysis of data from a random

assignment experiment. Next we discuss research by Rothstein (2010) that finds evidence that widely used value-added specifications are subject to considerable bias and Koedel & Betts (2011) that extends Rothstein's analysis by considering samples with multiple years of data for each teacher rather than the single year used in Rothstein. This leads us briefly to discuss a paper by Guarino et al (2011) that examines the performance of various value-added models under different data generating processes. Finally, we consider the new test of bias based on teacher mobility introduced by Chetty et al. (2011).

Recognizing the threat of student sorting both within and between schools to the estimation of teacher value-added, Kane & Staiger (2008) use experimental data generated by a random assignment study of the National Board for Professional Teaching Standards Certification Program to investigate the validity of non-experimental estimates of teacher value-added. In the study, pairs of teachers are identified in each school, one with and the other without certification, and classrooms are randomly assigned to the pairs. The difference in average test scores of the classrooms is regressed on the difference in empirical Bayes estimates of value-added for the same pair of teachers based on multiple years of data from earlier years in order to examine the validity of the estimation based on non-experimental data. The authors develop a specification test of the hypothesis that the non-experimental estimates provide consistent estimates of the difference in teacher effectiveness. Results of the empirical analysis are consistent with the hypothesis that sorting on unobservables does not confound the estimates of teacher value-added based on observational data in their sample.

It should be noted that the test does have some limitations. First, given the small sample size, the test lacks power in that the acceptance region includes values that would reflect significant sorting on unobservables that would bias the estimates. Second, compensatory assignment of better teachers to more difficult students reduces the probability of not rejecting the null hypothesis of no sorting on unobservables when it is false Finally, the small group of

schools in which principals agreed to permit classes to be randomly assigned to teachers is unlikely to be representative, meaning that evidence of the validity of value-added estimates with this sample may not generalize beyond this sample.

The aforementioned analysis covers a small sample of teachers in schools that agreed to participate in an experiment, and this both weakens the power of the test and raises questions about the generalizability of the findings. By comparison, the two papers by Rothstein evaluate some widely used models of teacher value-added using a large sample of North Carolina schools drawn from state administrative data. In contrast to the findings presented in Kane & Staiger (2008), Rothstein finds evidence that time-varying unobserved heterogeneity introduces bias into the estimates of teacher value-added for a number of widely used specifications. Importantly, this finding emerges even for school fixed effect models that focus solely on within school variation in teacher effectiveness. It appears that the models in question do not account fully for non-random sorting into classrooms on the basis of time-varying factors, as Rothstein finds that the strict exogeneity assumption is violated in specifications that include student fixed effects.

Rothstein (2010) develops falsification tests for three widely used value-added models that investigate the presence of non-random classroom sorting on the basis of unobservables. The first model regresses test score gain on indicators for schools and teachers; the second model regresses test score on lagged test score and the same indicators; and the final model stacks the gain observations and regresses gain on teacher indicators and student fixed effects for a sample of students who do not switch schools.

Although the precise form of the test differs by specification, essentially Rothstein tests whether classroom teachers in grade g account for a significant portion of the variation in test score gains in earlier grades. He finds strong evidence that future teachers predict variation in

test score gains in models that do and do not restrict the coefficient on lagged test score to equal 1 despite the fact that future teachers cannot be causally related to achievement. Even the inclusion of student fixed effects does not eliminate the explanatory power of future teachers, meaning that the strict exogeneity condition is violated by dynamic sorting on unobservables. Importantly, lagged teacher indicators are also significantly related to achievement in models that control for lagged achievement, indicating that the dynamic sorting on unobservables is not entirely based on achievement score.

Koedel & Betts (2011) replicate Rothstein's analysis for the San Diego school district and then extend the empirical analysis in order to investigate the sensitivity of Rothstein's findings to the particular character of his sample. Similar to Rothstein they find evidence of substantial bias in estimates based on a single year of data for each teacher. However, they also find that restricting the sample to teachers with multiple years of data significantly attenuates the bias. This suggests that patterns of systematic sorting change over time, perhaps to compensate teachers who received a particularly difficult draw in the previous year.

One complication relevant for each of these analyses concerned the appropriate functional form for the value-added framework. As Rothstein (2010) points out, incorrect assumptions regarding the rate of knowledge depreciation, form of student heterogeneity, and other factors can introduce bias, and models differ along a number of important dimensions. Guarino et al (2011) attempt to clarify the nature of any biases and sensitivity to the underlying sorting patterns by examining estimates of teacher value-added from samples produced by data generating processes that differed according to various aspects of student sorting among classrooms.

Not surprisingly, the authors find that the magnitude and character of biases for different value-added models vary according to the underlying sorting mechanism. However, they also find that the simple lagged achievement value-added model without student fixed effects tended to outperform other more highly structural models across a range of sorting types. This occurred despite the fact that this model was in no cases the prescribed approach under the specified structural model of learning with cumulative effects and student heterogeneity. The authors conclude tentatively that this pattern of findings suggest that the flexibility of the simpler lagged achievement model tends to improve performance.

Chetty et al (2011), in a novel analysis of the long-run impacts of teachers that we discuss below, begin with value-added measures for a large number of teachers over time in a large metropolitan area. They then exploit school-grade-cohort changes in achievement that are associated with the movement of teachers with varying estimated value-added. The predictions of how cohort achievement will change based on their estimates of teacher quality are closely related to the observed changes, leading them to conclude that any biases in the estimated teacher effects are small. They find no evidence of either selection on observables or selection on unobservables.

5. Persistence of Teacher Effects

The degree to which teacher quality effects persist into the future constitutes an important determinant of the importance of differences in teacher productivity. If a year with a poor teacher can be easily reversed by a higher quality teacher in a future grade the cost of having an ineffective teacher would be low. However, if the effects of teacher quality cumulate into the future in a way that is difficult to undo, the cost would be much higher. A small but growing

number of papers either focus on or consider the issue of persistence. Prior to discussing the findings we consider some of the main conceptual issues related to the question of persistence.

Conceptual Framework

Standard education production function approaches impose specific assumptions on the rate of knowledge and therefore treatment effect depreciation. One of the most common is that the rate of depreciation does not depend on the source of the learning, be it family, teacher, or other factor. Yet this framework does not capture some of the main ideas put forth in the work specific to teacher including the potential tradeoff between teacher effects on short term knowledge that raises current test scores only and effects on longer term knowledge that foster learning in both the current and subsequent years. Therefore we introduce a framework that builds on the notion that teachers can affect test scores by influencing short term or long term knowledge.[26] Short term knowledge refers to information that affects the current test score but has no effect on future achievement such as information on the specific form in which questions are asked. Longer term knowledge refers to information that adds to the long-run stock of human capital. Empirically teacher effects on subsequent year test scores are assumed to identify their effects on long-term knowledge, while effects on current year scores reflect the effects on both long and short-term knowledge.

Equation (7), an expanded version of equation (5), represents achievement (A) for student i with teacher j in year Y as a function of achievement in year 1 (A_{iy_1}), a vector of student characteristics (X), a vector of school fixed effects for the current and previous years (φ), a vector of long-term teacher effects for the current and previous years (τ), a short term teacher effect for year y (λ) and a random error (ε).

---

[26] Jacob et al (2010) and Carrell & West (2010) use a similar framework.

$$(7) \qquad A_{ijY} = (1-\theta)A_{iy_1} + X_{iY}\gamma + \sum_{y=2}^{Y}\varphi_y + \sum_{y=2}^{Y}\tau_y + \lambda_j + \varepsilon_{ijY}$$

Because long and short term teacher effects cannot be observed, separate identification of these and the degree of persistence require multiple years of data and additional assumptions on top of those discussed in the previous section. Underlying Equation (7) is the assumption that short term effects disappear after the current year. Consequently estimates of teacher effects in the current year capture both long and short-term effects, while estimates of teacher effects on subsequent year achievement capture only the long-term component. If estimated effects of the year y-2 teacher are available for years y, y-1, and y-2 comparisons between y and y-1 provide information on the relative size of long and short-term effects and comparisons between y-1 and y-2 provide information on the depreciation of past teacher effects.

Empirical findings

A small but growing number of papers estimate the persistence of teacher effects, and we now discuss three papers from that literature: Jacob et al (2010); Konstantopoulos (2011); and Lockwood & McCaffrey (2009). These three papers take different approaches to identification of the rate of persistence, in part because of data differences. Not surprisingly, complications introduced by omitted variables or sample attrition raise concerns in each of these analyses.

Jacob et al (2010) develop a multi-stage empirical framework that isolates the long-term effect of teachers and find that only 20 percent of the initial effect persists beyond one year. First, estimates for the quality of teacher j for students in cohort c are derived by averaging the estimated teacher j fixed effects over all cohorts other than c; this removes unobserved cohort c specific classroom influences from the estimate of teacher quality. Next, these estimates of teacher quality are used to instrument for grade g achievement in a regression of achievement in grade g+1 on grade g achievement, classroom fixed effects, and other controls. The coefficient

22

on grade g achievement provides an estimate of the persistence of the teacher effects, as it captures the extent to which teacher induced changes in grade g achievement persist into the subsequent year. Then the authors substitute achievement in grade g-1 in order to estimate two year persistence and find little change over time in the rate of persistence.

This approach provides an appealing method for identifying the persistence in teacher effects, but a number of questions emerge regarding the interpretation of the estimates. One of these questions relate narrowly to the details of the method, while others relate to broader issues relevant to all of the studies discussed. We consider first the method-specific concern and then turn to the broader estimation issues.

The exclusion of a student's own cohort from the calculation of the prior year teacher's effectiveness mitigates biases introduced by unobserved peer or school influences, but it also imposes the strong assumption that the quality of a given teacher's work does not vary over time. If the quality of instruction fluctuates over time due to on-the-job training, changes in personal circumstances, changes in curriculum, or some other factor, performance in other years will provide a noisy measure of instructional quality, and the estimates of persistence will be attenuated.

The remaining estimation issues generalize beyond the specific method used in this paper and relate to both bias and interpretation. First, estimates of persistence that include current year teacher fixed effects in addition to prior year teacher fixed effects or quality estimates rely upon changes in classroom composition for identification. If classrooms progress through school intact there is no way to identify the separate effects of teachers in different grades because of perfect multicollinearity. Schools with only one classroom per grade or one teacher per subject per grade therefore provide no information unless some students attend one of the grades in another school

in the sample. In schools with multiple classrooms per grade identification comes from the remixing of classrooms after each grade and school transfers. Time-varying unobservables that complicate identification of current year teacher effects can also bias persistence estimates.

Second, test measurement error that inflates estimates of the variance in teacher quality can attenuate persistence estimates if not appropriately addressed. Because fixed effects tend to exacerbate measurement error induced attenuation bias, the problem is likely to be particularly problematic in this context given the presence of current classroom or teacher fixed effects. Empirical Bayes shrinkage estimators or other method that addresses this issue can mitigate the bias.

Third, as Jacob et al (2010) point out, statistical properties and coverage of tests introduce uncertainty into the meaning of any persistence estimate. The absence of strict vertical scaling means that the amount of learning reflected in a given magnitude of test score growth varies across the achievement distribution. Monotonic transformations of a given scale will alter the estimated degree of persistence, and there is no strong argument in favor of a given scaling. Test coverage and curriculum structure also influence the interpretation of persistence estimates. If knowledge gained from a teacher in grade g relates to only a small portion of the material tested in grade g+1, a finding of little or no persistence contains little information on the actual degree to which the student retains knowledge learned in the prior grade. This may occur for a number of reasons including the focus of tests on rudimentary material that is barely covered in some schools or lack of cumulative structure in certain subjects, e.g. biology followed by chemistry.

Fourth, within classroom variation in current teacher value-added could also affect estimates of persistence. If teachers tend to focus on lower achievers who received inferior instruction during the prior grade, classroom average value-added would overstate teacher value-

added to higher achievers and understate current teacher value-added to lower achievers. The failure to account for compensatory behavior by the current teacher would tend to bias downward estimates of persistence.

These general issues affect the interpretation of persistence estimates in the work by Jacob et al (2010) and in the other papers referenced above to which we now turn. The one that stands in sharpest contrast to Jacob, Lefgren, and Sims is Konstantopoulos (2011), which finds a much higher level of persistence using data from the Tennessee STAR experiment. Although the estimated variances of contemporaneous year teacher effects are not reported, the finding that a one standard deviation increase in teacher quality raises subsequent year achievement by roughly 0.1 standard deviations suggests substantial persistence in teacher quality effects.

The use of the Tennessee STAR experimental data imparts both advantages and challenges onto this study. Ideally, selection bias would not affect a study using random assignment data, and evidence suggests that the initial kindergarten randomization was done well. However, research by Hanushek (1999b) shows the existence of substantial non-random attrition in the subsequent years, and this can introduce bias in the estimation of persistence of initial teacher quality. For example, if more educationally committed families were more sensitive to teacher quality in the decision to remain in the experimental classes, one would expect students taught by the less effective teacher to become more negatively selected over time. This would tend to increase the estimate of persistence above its true value.

In addition to the issue of selection bias, the structure of the STAR experiment complicates estimation of persistence. Students initially assigned to a particular class type are scheduled to remain in that class size category for the duration of the experiment. This means that in schools with only one teacher per class type classes would remain intact from grade-to-

grade in the absence of entry or exit. In this case the effects of prior year teachers could not be

separate from those of current teachers, similar to the one teacher per grade example discussed

above. This requires movement and transfers among classrooms within the school or between

schools, and such movements often introduce selection bias since they are often not part of the

initial or even follow-up randomizations. The character of any such biases is complicated by the

fact that the equations include both indicators for current year teachers and estimated teacher

quality for the prior grade.

McCaffrey et al (2009) takes a Bayesian approach and examines persistence using data

for a large urban district. They find a rate of persistence significantly above zero but far closer to

zero than to one (no depreciation). The teacher effect estimates do not appear to be very

sensitive to the specification of the prior distribution, arbitrary rescaling of the data, or to the use

of marginal models as opposed to joint models that pool mathematics and reading achievement.

However, these estimates do not include any student background variables and thus unobserved

heterogeneity may contaminate the estimates. Moreover, although the estimates are not sensitive

to the method used to deal with missing data, the absence of such controls raises the probability

that non-random attrition does introduce substantial bias not accounted for by any of the missing

data methods used in the paper.

Chetty et al. (2011) provide a different perspective on the issue of persistence. They

relate value-added measures for teachers in grades 4-8 to student achievement in subsequent

years, finding that about 30 percent still shows up in achievement after three years. But, perhaps

more importantly, they link teacher value-added estimates to students' subsequent outcomes

including earnings in the labor market. They find that on average a one standard deviation

improvement in teacher value-added in a single grade raises earnings by about one percent at age

28. In other words, the impact of individual teachers carries through to the long run success of students.

Clearly these estimates of persistence fail to find common ground, and this is not surprising given differences among studies in both methods and material tested. It should be noted that Rothstein (2010) produces estimates similar to those of Jacob et al (2010), while Kane & Staiger (2008) produce somewhat higher estimates of persistence but ones that are sensitive to the specification. Chetty et al. (2011) show that depreciation in the impacts over time stabilizes after about three years and that the impact remains in life experiences after school. Many of the issues raised for the four papers discussed in this section are also relevant to these two studies. Finally, Carrell & West (2010) find that lasting teacher effects appear to be orthogonal to or even negatively correlated with short term effects for a sample of Air Force Academy students. This pattern is consistent with the model introduced above that distinguishes between short and longer term effects.

6. Selected Policy Issues

As noted, one of the reasons for broad interest in the analyses of teacher quality is the very direct relationship with policy deliberations. In this section we consider three areas of policy that have received considerable attention in recent years. The first is the desirability of test-based accountability as an important component of school regulations. Advocates of accountability implicitly assume that test scores are related to important long term outcomes including academic attainment and labor market performance, but some have questioned this assumption. A second discussion uses information about the impact of achievement on economic outcomes to place a monetary value on differences in teacher quality. Finally, the third discusses the method used to reduce the size of the teaching force in response to falling enrollment,

declining budgets, or other factors. Regulations and practices in this area have long-term effects

on the distribution of teacher quality, and we consider the implications of the widely used last in,

first out policy.

Achievement and Longer-term outcomes

Economists' interest in human capital has largely centered on the quantity of schooling,

but school quality has received increasing consideration over time. Mincer (1970) and Mincer

(1974) pioneered the study of the economic returns to additional years in school using readily

available census data A large number of researchers have subsequently estimated "Mincer

earnings functions" that relate the log of individual earnings to years of schooling and labor

market experience.   The Mincer earnings work has been extended in a number of studies to

incorporate the returns to achievement or school quality. This work is highly relevant to the

current debate over test-based accountability and value-added, because it provides evidence on

the degree to which achievement differences translate into future differences in earnings.

The most common estimation begins with a standard Mincer earnings model with the

addition of a measure of achievement (A_i) such as:

$$(8) \qquad \ln Y_i = \alpha_0 + rS_i + \alpha_1 Exper_i + \alpha_2 Exper_i^2 + \phi A_i + \varepsilon_i$$

where Y_i is earnings of individual i, S is school attainment, *Exper* is potential labor market

experience, and ε is a random error.   When cognitive skills are standardized to mean zero and

standard deviation of one, ϕ is interpreted simply as the percentage increase in annual earnings

that can be attributable to a one standard deviation increase in achievement.[27]

---

[27] Much of the labor economics consideration of earnings determination has focused on the estimates of r,
interpreted as the rate of return on schooling (see Card (2001) but also the re-interpretation in Heckman et al
(2006)). An underlying theme has been the influence of unmeasured ability on both the decision to attend further
schooling and subsequent earnings, and one approach has been to add measured achievement to reflect ability. The
interpretation that we give is that there is no conceptual difference between ability and achievement, that they can be
measured with the same instruments, and that – borrowing from the educational production function literature –
achievement can be affected by schools.

Table 2 summarizes different estimates of the impact of test performance on earnings (ϕ). The first three consider young workers (Mulligan (1999); Murnane et al (2000); Lazear (2003)). These studies employ different nationally representative data sets that follow students after they leave school and enter the labor force. When scores are standardized, they suggest that one standard deviation increase in mathematics performance at the end of high school translates into 10-15 percent higher annual earnings.[28]. The estimates do, however, come early in the worker's career, suggesting the impact may actually rise with experience.[29]

In a different set of estimates using data on a sample of workers for the U.S., Hanushek & Zhang (2009) provide estimates of returns (ϕ) of 20 percent per standard deviation. One distinguishing feature of these estimates is that they come for a sample of workers throughout the career, as opposed to the prior estimates that all come from early-career earnings.

Using yet another methodology that relies upon international test scores and immigrants into the U.S., Hanushek & Woessmann (2009) obtain an estimate of 14 percent per standard deviation. These estimates come from a difference-in-differences formulation based on whether the immigrant was educated in the home country or in the United States. They find that skills measured by international math and science tests from each immigrant's home country are significant in explaining earnings within the United States.

Finally, Chetty et al (2010) look at how kindergarten test scores affect earnings at age 25-27 and find an increase of 18 percent per standard deviation. These estimates do not control for

---

[28] It is convenient to convert test scores into measures of the distribution of achievement across the population. A separate review of the normalized impact of measured cognitive skills on earnings by Bowles et al (2001) finds that the mean estimate is only 0.07, or slightly over half of the specific studies here. More details on the individual studies shown here can be found in Hanushek (2011).

[29] These estimates are derived from observations at a point in time. Over the past few decades, the returns to skill have risen. If these trends continue, the estimates may understate the lifetime value of skills to individuals. On the other hand, the trends themselves could change in the opposite direction. For an indication of the competing forces over a long period, see Goldin & Katz (2008).

any intervening school attainment differences but do control for a rich set of parental characteristics.

The finding that moving a standard deviation in cognitive skills yields 10-20 percent higher income may sound small, but these increments apply throughout the lifetime. In 2010, the average present value of income for fulltime, full-year workers age 25-70 is $1.16 million.[30] Thus, one standard deviation higher performance even at a low return of 13 percent per standard deviation amounts to over $150,000.

This approach will understate the full impact of achievement to the extent that higher achievement leads to higher levels of schooling, but that is generally not considered.[31] Thus, these estimates can be considered lower bounds on the earnings impacts of higher achievement.

Economic Value of Increased Teacher Quality

The prior parameter estimates provide the building blocks for assessing the implication of differences in teacher quality for lifetime income. Specifically, we can combine information about the distribution of teacher effectiveness and the labor market impact of higher achievement to estimate the economic impact of a given change in teacher value-added to achievement,

First, we use evidence on the distribution of teacher effectiveness to determine the magnitude of a one standard deviation difference in teacher quality, which constitutes a difference that is well within the range of the sample Table 1 provides a set of estimates based on within school variations, but for this we want the total variation that includes any between-school variation. Because of difficulties in identifying the between-school variance in quality, this

---

[30] Calculations use average income by age for all fulltime, full-year workers in the labor force in the first quarter of 2010. It is assumed that incomes rise 1 percent per year because of overall productivity improvements in the economy and that future incomes are discounted at 3 percent.
[31] Murnane et al (2000) is an exception for tracing through the indirect effects. See also the discussion of the form of estimation in Hanushek & Zhang (2009).

analysis considers two magnitudes of total variance in teacher quality (σ_T), a lower estimate of 0.2 standard deviations (s.d.) and a larger estimate of 0.3 s.d. Thus, a teacher who is one standard deviation above the mean to the distribution of teachers in terms of quality (i.e., roughly comparing the 84[th] percentile teacher to the 50[th] percentile teacher) is estimated to produce marginal learning gains of 0.2-0.3 standard deviations of student achievement compared to the average teacher.[32]

The implication for earnings depends on the persistence of this learning into the future and how this increased learning in any given year carries through into the labor market, i.e., the persistence discussed above. The baseline calculations presume that 70 percent of the added learning persists into the future, i.e., that θ in Equation (1) is 0.3. As a result of the imprecision of prior estimates, however, the impact of depreciation that is twice as large (i.e., $\theta = 0.6$) is also investigated.

It is now possible to calculate the value of an above average teacher in terms of effectiveness. As a base case, consider the effect of having a teacher who is 0.5 s.d. above average in the teacher effectiveness distribution (i.e., at the 69[th] percentile) as opposed to a teacher at the mean of the distribution. She would, according to the above estimates, annually produce a 0.1 s.d. average improvement in cognitive skills of her students (assuming that the standard deviation of teacher effectiveness in units of student achievement is 0.2 s.d.) above what would have been produced by the mean teacher.

Combining the improvement in scores for an individual with a conservative estimate of an impact on future individual earnings of 13 percent per standard deviation of achievement

---

[32] In terms of the student achievement distribution, this would move a student from the 50[th] percentile to the 58[th] to 62[nd] percentile.

(from Table 2 above) and $\theta = 0.3$ yields a present value of $10,600 over a lifetime of work for the average worker.

But this is not yet the full impact of the above average teacher. The impact on one student is replicated across all of the other students in the class. Thus, calculation of the impact of a teacher depends directly on class size, the full annual economic value of teachers at different points in the distribution can be found simply by multiplying the individual value time the class size. Figure 1 displays the impact of different quality teachers according to class sizes at varying percentiles of the distribution. A teacher who is at the 60^{th} percentile (0.25 s.d. above average) raises individual earnings by $5,292, and this translates into a present value of $105,830 for a class size of 20 students. A teacher who is one standard deviation above the mean (84^{th} percentile) produces over $400,000 in added earnings for her class of twenty.[33]

The first thing to note is that this is an *annual* increment by the teacher. Any teacher who stays at the given level of performance produces such an amount each year.

The second thing to note from the bottom half of Figure 1 is that a below average teacher leads to a similar decrease in lifetime earnings.[34] Thus, having an effective teacher followed by an equally ineffective teacher will cancel out the gains.

The precise marginal economic value depends crucially on the three parameters of the teacher distribution and of how achievement evolves over time and affects earnings: $\sigma_T = 0.2$; $\phi = 0.13$; and $\theta = 0.3$. The impact of the different parameters is straightforward. A lower depreciation rate (higher persistence of achievement), a wider distribution of the teacher

---

[33] Chetty et al (2010), extrapolating from their data on early career earnings, estimate the impact of a high quality teacher at about $214,000 per class of 20 for a teacher one s.d. above the mean. This is very close to the lower bound estimate in Table 3, below.

[34] The decrease is slightly different because the estimates come from Mincer earnings functions which relate the logarithm of earnings to the level of cognitive skills and thus to a slight different percentage change when evaluated at a different place in the distribution.

effectiveness distribution, and a larger labor market payoff to skill lead to a larger economic

value of teacher effectiveness. Additionally, the return to skill of $\phi = 0.13$ most closely mirrors

the labor market estimates for young workers and for time periods in the past when the demand

for skill was less. More recent estimates and consideration of the full age range of workers

yields larger estimates, suggesting that $\phi = 0.2$ is a plausible upper bound on the estimates. The

baseline estimates do use a depreciation rate of 0.3, whereas other studies described above

suggest larger depreciation, particularly of achievement gains induced by the teacher. We thus

also look at $\theta = 0.6$, or a depreciation rate that is twice as large.

[Sidebar: Economic Growth about here]

      Table 3 presents alternative estimates of marginal impacts evaluated at one point in the

teacher distribution – one standard deviation above the mean, or the 84[th] percentile. Even at the

lower bound in column (1) of the table defined by the previous quality and earnings parameters

($\sigma_T = 0.2$ and $\phi = 0.13$) but higher depreciation ($\theta = 0.6$), a good teacher with a class of 15

annually produces \$182,000 more in present value than the average teacher. The marginal

annual economic value of a good teacher (compared to the average) evaluated at a given class

size – say 20 students per class –do make a large difference in the estimated impact. The annual

economic value with class size of 20 ranges from a quarter of a million dollars to a million

dollars at the top of the range for the three parameters together. (The final column is an upper

bound on estimates based on current empirical work).

      While the difference in estimates across the parameters is large, the more striking feature

of the table is the magnitude of the lower bound. A teacher in the top 15 percent with a class of

20 or more students yields at least \$240,000 in economic surplus each and every year compared

to an average teacher.

A host of unknown factors – including compensatory behavior of parents and schools, the cumulative nature of skills, the specific attributes valued in the labor market, and the nature of peer-classroom interactions come into play in determining the long run impact of specific teachers. But even twice the depreciation of achievement that was used in the baseline yields very large estimates of the value of an effective teacher – say, $150,000 per year present value for a 75th percentile teacher with a class of 20 students.

A final caveat is that none of the prior estimates in Table 3 address any possible offset from direct class-size increases that would result if teachers were simply laid off instead of being replaced by average teachers (see the next section, however). The magnitude of any such effects has been controversial and does not readily permit explicit analysis; see, for example, Hanushek (1999a), Krueger (1999), Mishel & Rothstein (2002). At large class sizes, the estimates might be interpreted as an upper bound.[35] Of course, that would also imply that those at lower class size are lower bounds.

These estimates of the economic impact of teachers – which come from chaining together parameters estimated across different studies – are validated by the direct estimates of Chetty et al. (2010). They consider how the least effective teachers affect the earnings of students, and they also find present values of teacher impacts for the bottom five percent of teachers that are greater than $250.000 per class. While it is difficult to compare their estimates and those here directly because of varying discount rates, class sizes, and the like, it is clear that they confirm the orders of magnitude described here.

---

[35] Among the various class size estimates, those in Krueger (1999) are among the largest – suggesting that a decrease in class size of eight students for four years might yield an achievement gain of 0.2 standard deviations.

LIFO (Last In, First Out)

Many union contracts for teachers specify that any layoffs should be done according to inverse seniority rules. The least senior teachers must be dismissed before any tenured teachers are eliminated. But, as noted, there is little relationship between experience and effectiveness in the classroom, except perhaps for teachers in their first few years. A layoff policy based on seniority thus would be expected to bear little relationship to the impact on student achievement. Until very recently, few teachers were laid off, so the rules were not binding. However, the fiscal distress of states and localities in the aftermath of the 2008 recession led to some layoffs of teachers. Recent analyses have linked information about teacher value-added to lay off policies, and they indicate that the impact of reductions depend importantly on how the layoffs are implemented (Glazerman et al (2010), Goldhaber & Theobald (2011), Boyd et al (2011)).

Goldhaber & Theobald (2011) evaluate the pattern of layoff notices sent to teachers in the State of Washington over two school years (2008-09 and 2009-10).[36] They found that, although teachers with master's degrees and some specialties had slightly lower probability of getting a notice, the overwhelmingly most important factor in the layoff decision is simply seniority. Because the less senior teachers have lower salaries, in order to obtain the same fiscal savings, districts would have to lay off 20 percent fewer teachers if those laid off fell at the district average salary. More interesting, however, is the impact of layoffs on student outcomes that is estimated using the value-added of teachers.

Their analysis incorporates teacher effectiveness measures in the consideration of who gets layoff notices (for the subset of teachers where value-added estimation is possible). The teacher value-added measure is never close to being a statistically significant determinant of

---

[36] Note that, because of advance notice provisions in law and contracts, school districts send out many more layoff notices than they ultimately exercise.

receiving a layoff notice. Moreover, the differential in average effectiveness (value-added) between teachers separated on the basis of effectiveness as opposed to the largely seniority basis employed in the notices is 0.2 standard deviations of student achievement.

Boyd et al (2011) pursue a very similar simulation in trying to investigate the implications of a hypothetical five percent reduction in teacher salaries for New York City. Their estimates are similar but somewhat larger than those in Goldhaber & Theobald (2011). In their analysis, average teacher who is laid off under a value-added system is 26 percent of a standard deviation in student achievement less effective than the typical seniority-based layoff. As they note, this is more than twice the size of the "first-year teaching" effect.

Part of the importance of these findings, however, depends upon decision makers being able to identify differences in teacher effectiveness, since only a majority of teachers can participate in formal value-added analysis. For example, in the evaluation system used in Washington, DC, less than one-fifth of the teachers were assessed by value-added measures – implying that other elements of evaluation must necessarily be incorporated.

Several studies have related estimates of teacher value-added to the evaluations of principals (for example, Murnane (1975), Armor et al (1976), Jacob & Lefgren (2008)). The recent study by Jacob & Lefgren (2008) finds that principals can do quite well at identifying which teachers are in the extreme of the distribution, although they do worse at differentiating teachers in the middle of the distribution. Identifying the extremes, however, is probably most relevant for teacher personnel and evaluation systems, particularly as related to any layoff policies. Therefore greater flexibility in lay-off decisions could lead to substantial changes in the distribution of teacher quality following a reduction in the size of the teaching force.

Of course, major changes in the terms of the standard teacher contract could have implications (positive or negative) for the teacher labor market. They could affect not only career decisions of existing teachers but the supply of potential teachers into the profession. The first order effects on achievement from selective decision making about layoffs are, nonetheless, large.

7. Some Concluding Thoughts

This discussion of course falls far short of describing all of the relevant policy issues related to teacher quality. Experience on the impact of employing value-added measures in personnel decisions, contracts, and overall policy is currently quite limited, implying that some of the gains (and losses) from changed policies remain speculative at this time. At the same time, policies are moving rapidly, suggesting that they will soon be subject to empirical analysis.

A big issue in all of the policy considerations is simply what drives the observed differences in teacher quality. Without knowing what leads to better or worse performance, it is hard to know what should be done to train teachers. It is hard to know how to hire teachers who have no observed performance to go with them. And, it is hard to decide on such issues as mentoring new teachers or providing professional development.

Clearly research is continuing to fill in these missing elements. At the same time, it seems sensible to develop systems that recognize the large differences and act upon them rather than delaying decisions until all of the answers are available.

References

Aaronson D, Barrow L, Sander W. 2007. Teachers and Student Achievement in the Chicago Public High Schools. *Journal of Labor Economics* 25: 95-135

Armor DJ, Conry-Oseguera P, Cox M, King N, McDonnell L, et al. 1976. *Analysis of the school preferred reading program in selected Los Angeles minority schools*. Santa Monica, CA: Rand Corp.

Ballou D, Sanders W, Wright P. 2004. Controlling for student background in value-added assessment of teachers. *Journal of Educational and Behavioral Statistics* 29: 37-65

Black SE, Machin S. 2011. Housing Valuations of School Performance. In *Handbook of the Economics of Education, Vol. 3*, ed. EA Hanushek, S Machin, L Woessmann, pp. 485-519. Amsterdam: North Holland

Boardman AE, Murnane RJ. 1979. Using panel data to improve estimates of the determinants of educational achievement. *Sociology of Education* 52: 113-21

Bowles S, Gintis H, Osborne M. 2001. The determinants of earnings: A behavioral approach. *Journal of Economic Literature* 39: 1137-76

Boyd D, Lankford H, Loeb S, Wyckoff J. 2005. The draw of home: How teachers' preferences for proximity disadvantage urban schools. *Journal of Policy Analysis and Management* 24: 113-32

Boyd D, Lankford H, Loeb S, Wyckoff J. 2011. Teacher Layoffs: An Empirical Illustration of Seniority versus Measures of Effectiveness. *Education Finance and Policy* 6: 439-54

Braun H, Chudowsky N, Koenig J, eds. 2010. *Getting value out of value-added: Report of a workshop*. Washington, DC: National Academies Press

Card D. 2001. Estimating the return to schooling: Progress on some persistent econometric problems. *Econometrica* 69: 1127-60

Carrell Scott E, West James E. 2010. Does Professor Quality Matter? Evidence from Random Assignment of Students to Professors. *Journal of Political Economy* 118: 409-32

Chetty R, Friedman JN, Hilger N, Saez E, Schanzenbach DW, Yagan D. 2010. How Does Your Kindergarten Classroom Affect Your Earnings? Evidence From Project STAR. In *NBER*. Cambridge, MA: National Bureau of Economic Research

Chetty R, Friedman JN, Rockoff JE. 2011. The Long-Term Impacts of Teachers: Teacher Value-Added and Student Outcomes in Adulthood. In *NBER*. Cambridge, MA: National Bureau of Economic Research

Chubb JE. 2009. *Learning from No Child Left Behind: How and Why the Nation's Most Important but Controversial Education Law Should be Renewed*. Stanford, CA: Education Next Books, Hoover Institution Press

Coleman JS, Campbell EQ, Hobson CJ, McPartland J, Mood AM, et al. 1966. *Equality of educational opportunity*. Washington, D.C.: U.S. Government Printing Office

Dee TS. 2005. A teacher like me: Does race, ethnicity, or gender matter? *American Economic Review* 95: 158-65

Glazerman S, Loeb S, Goldhaber D, Staiger D, Raudenbush S, Whitehurst G. 2010. *Evaluating Teachers: The Important Role of Value-Added*, Brookings Institution, Washington, DC

Goldhaber D, Gross B, Player D. 2011. Teacher career paths, teacher quality, and persistence in the classroom: Are public schools keeping their best? *Journal of Policy Analysis and Management* 30: 57-87

Goldhaber D, Theobald R. 2011. Managing The Teacher Workforce in Austere Times: The Implications of Teacher Layoffs. In *Center for Education Data and Research*. Bothell, WA: University of Washington Bothell

Goldin C, Katz LF. 2008. *The race between education and technology*. Cambridge, MA: Harvard University Press

Guarino CM, Reckase MD, Wooldridge JM. 2011. Can Value-Added Measures of Teacher Performance be Trusted? In *Education Policy Center*. East Lansing, MI: Michigan State University

Hanushek EA. 1971. Teacher characteristics and gains in student achievement: Estimation using micro data. *American Economic Review* 60: 280-88

Hanushek EA. 1979. Conceptual and empirical issues in the estimation of educational production functions. *Journal of Human Resources* 14: 351-88

Hanushek EA. 1999a. The evidence on class size. In *Earning and learning: How schools matter*, ed. SE Mayer, PE Peterson, pp. 131-68. Washington, DC: Brookings Institution

Hanushek EA. 1999b. Some findings from an independent investigation of the Tennessee STAR experiment and from other investigations of class size effects. *Educational Evaluation and Policy Analysis* 21: 143-63

Hanushek EA. 2009. Building on No Child Left Behind. *Science* 326: 802-02

Hanushek EA. 2011. The economic value of higher teacher quality. *Economics of Education Review* 30: 466-79

Hanushek EA, Kain JF. 1972. On the value of `equality of educational opportunity' as a guide to public policy. In *On equality of educational opportunity*, ed. F Mosteller, DP Moynihan, pp. 116-45. New York: Random House

Hanushek EA, Kain JF, O'Brien DM, Rivkin SG. 2005. The market for teacher quality. National Bureau of Economic Research

Hanushek EA, Kain JF, Rivkin SG. 2004a. Disruption versus Tiebout improvement: The costs and benefits of switching schools. *Journal of Public Economics* Vol 88/9-10: 1721-46

Hanushek EA, Kain JF, Rivkin SG. 2004b. Why public schools lose teachers. *Journal of Human Resources* 39: 326-54

Hanushek EA, Kain JF, Rivkin SG. 2009. New evidence about *Brown v. Board of Education*: The complex effects of school racial composition on achievement. *Journal of Labor Economics* 27: 349-83

Hanushek EA, Rivkin SG. 2006. Teacher quality. In *Handbook of the Economics of Education*, ed. EA Hanushek, F Welch, pp. 1051-78. Amsterdam: North Holland

Hanushek EA, Rivkin SG. 2007. Pay, working conditions, and teacher quality. *Future of Children* 17: 69-86

Hanushek EA, Rivkin SG. 2010a. Constrained Job Matching: Does Teacher Job Search Harm Disadvantaged Urban Schools? : Hoover Institution, Stanford University

Hanushek EA, Rivkin SG. 2010b. Constrained Job Matching: Does Teacher Job Search Harm Disadvantaged Urban Schools? In *NBER*. Cambridge, MA: National Bureau of Economic Research

Hanushek EA, Rivkin SG. 2010c. Generalizations about using value-added measures of teacher quality. *American Economic Review* 100: 267-71

Hanushek EA, Sarpça S, Yilmaz K. 2011. Private Schools and Residential Choices: Accessibility, Mobility, and Welfare. *B.E. Journal of Economic Analysis & Policy: Contributions* 11: article 44

Hanushek EA, Woessmann L. 2008. The role of cognitive skills in economic development. *Journal of Economic Literature* 46: 607-68

Hanushek EA, Woessmann L. 2009. Do better schools lead to more growth? Cognitive skills, economic outcomes, and causation. In *NBER Working Paper*. Cambridge, MA: National Bureau of Economic Research

Hanushek EA, Woessmann L. 2011. How much do educational outcomes matter in OECD countries? *Economic Policy* 26: 427-91

Hanushek EA, Yilmaz K. 2007. The complementarity of Tiebout and Alonso. *Journal of Housing Economics* 16(2): 243-61

Hanushek EA, Zhang L. 2009. Quality-consistent estimates of international schooling and skill gradients. *Journal of Human Capital* 3: 107-43

Heckman JJ, Lochner LJ, Todd PE. 2006. Earnings functions, rates of return and treatment effects: The Mincer equation and beyond. In *Handbook of the Economics of Education*, ed. EA Hanushek, F Welch, pp. 307-458. Amsterdam: North Holland

Isenberg E, Hock H. 2010. *Measuring School and Teacher Value Added for IMPACT and TEAM in DC Public Schools: Final Report*, Mathematica Policy Research, Princeton, NJ

Ishii J, Rivkin SG. 2009. Impediments to the Estimation of Teacher Value Added. *Education Finance and Policy* 4: 520-36

Jackson CK, Bruegmann E. 2009a. *Teaching Students and Teaching Each Other: The Importance of Peer Learning for Teachers*

Jackson CK, Bruegmann E. 2009b. Teaching Students and Teaching Each Other: The Importance of Peer Learning for Teachers. *American Economic Journal: Applied Economics* 1: 85-108

Jacob Brian A, Lefgren L. 2008. Can Principals Identify Effective Teachers? Evidence on Subjective Performance Evaluation in Education. *Journal of Labor Economics* 26: 101-36

Jacob BA, Lefgren L, Sims D. 2010. The persistence of teacher-induced learning. *Journal of Human Resources* 45: 915-43

Kane TJ, Rockoff JE, Staiger DO. 2008. What does certification tell us about teacher effectiveness? Evidence from New York City. *Economics of Education Review* 27: 615-31

Kane TJ, Staiger DO. 2002. Volatility in school test scores: Implications for test-based accountability systems. In *Brookings Papers on Education Policy 2002*, ed. D Ravitch, pp. 235-69. Washington, DC: Brookings

Kane TJ, Staiger DO. 2008. Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation. In *NBER*. Cambridge, MA: National Bureau of Economic Research

Kim H-K. 2001. Is there a crowding-out effect between school expenditure and mother's child care time? *Economics of Education Review* 20: 71-80

Koedel C, Betts J. 2011. Does Student Sorting Invalidate Value-Added Models of Teacher Effectiveness? An Extended Analysis of the Rothstein Critique. *Education Finance and Policy* 6: 18-42

Konstantopoulos S. 2011. Teacher effects in early grades: Evidence from a randomized study. *Teachers College Record* 113

Koretz D. 2008. *Measuring up: What educational testing really tells us*. Cambridge, MA: Harvard University Press

Krueger AB. 1999. Experimental estimates of education production functions. *Quarterly Journal of Economics* 114: 497-532

Lazear EP. 2003. Teacher incentives. *Swedish Economic Policy Review* 10: 179-214.

Lockwood JR, McCaffrey DF. 2009. Exploring Student-Teacher Interactions in Longitudinal Achievement Data. *Education Finance and Policy* 4: 439-67

Loeb S, Darling-Hammond L, Luczak J. 2005. How Teaching Conditions Predict Teacher Turnover in California Schools. *Peabody Journal of Education* 80: 44-70

McCaffrey DF, Koretz D, Lockwood JR, Hamilton LS. 2004. *Evaluating Value-Added Models for Teacher Accountability*, RAND

McCaffrey DF, Sass TR, Lockwood JR, Mihaly K. 2009. The Intertemporal Variability of Teacher Effect Estimates. *Education Finance and Policy* 4: 572-606

Meghir C, Palme M. 2005. Educational Reform, Ability, and Family Background. *American Economic Review* 95: 414-23

Meghir C, Rivkin SG. 2011. Econometric methods for research in education. In *Handbook of the Economics of Education, Vol. 3*, ed. EA Hanushek, S Machin, L Woessmann, pp. 1-87. Amsterdam: North Holland

Mincer J. 1970. The distribution of labor incomes: a survey with special reference to the human capital approach. *Journal of Economic Literature* 8: 1-26

Mincer J. 1974. *Schooling, experience, and earnings*. New York: NBER

Mishel L, Rothstein R, eds. 2002. *The class size debate*. Washington, DC: Economic Policy Institute

Mulligan CB. 1999. Galton versus the human capital approach to inheritance. *Journal of Political Economy* 107: S184-S224

Murnane RJ. 1975. *Impact of school resources on the learning of inner city children*. Cambridge, MA: Ballinger

Murnane RJ, Willett JB, Duhaldeborde Y, Tyler JH. 2000. How important are the cognitive skills of teenagers in predicting subsequent earnings? *Journal of Policy Analysis and Management* 19: 547-68

Neal D. 2011. The design of performance pay in education. In *Handbook of the Economics of Education, Volume 4*, ed. EA Hanushek, S Machin, L Woessmann, pp. 495-550. Amsterdam: North Holland

Nye B, Konstantopoulos S, Hedges LV. 2004. How Large Are Teacher Effects? *Educational Evaluation and Policy Analysis* 26: 237-57

Rivkin SG, Hanushek EA, Kain JF. 2005. Teachers, schools, and academic achievement. *Econometrica* 73: 417-58

Rockoff JE. 2004. The impact of individual teachers on student achievement: Evidence from panel data. *American Economic Review* 94: 247-52

Rothstein J. 2009. Student sorting and bias in value-added estimation: Selection on observables and unobservables. *Education Finance and Policy* 4: 537-71

Rothstein J. 2010. Teacher quality in educational production: Tracking, decay, and student achievement. *Quarterly Journal of Economics* 125: 175-214

Sacerdote B. 2011. Peer effects in education: How might they work, how big are they and how much do we know thus far? In *Handbook of the Economics of Education*, ed. EA Hanushek, S Machin, L Woessmann, pp. 249-77. Amsterdam: North Holland

Sanders WL, Horn SP. 1994. The Tennessee value-added assessment system (TVAAS): Mixed-model methodology in educational assessment. *Journal of Personnel Evaluation in Education* 8: 299-311

Song J, Felch J. 2011. Times updates and expands value-added ratings for Los Angeles elementary school teachers. In *Los Angeles TImes*. Los Angeles

Todd PE, Wolpin KI. 2003. On the specification and estimation of the production function for cognitive achievement. *Economic Journal* 113: F3-33

Table 1. The Distribution of Teacher Effectiveness (standard deviations of student achievement)

Study	Location	teacher effectiveness (σ_τ)	
		reading	math
Rockoff (2004)	New Jersey	0.10	0.11
Nye et al (2004)	Tennessee	0.07	0.13
Rivkin et al (2005)	Texas	0.15	0.11
Aaronson et al (2007)	Chicago		0.13
Kane et al (2008)	New York City	0.08	0.11
Jacob & Lefgren (2008)	Midwest city?	0.12	0.26
Kane & Staiger (2008)	Los Angeles	0.18	0.22
Koedel & Betts (2011)	San Diego		0.23
Rothstein (2009)	North Carolina	0.11	
Hanushek & Rivkin (2010a)	Texas city?		0.11
Average		0.13	0.17

Note: All estimates indicate the standard deviation of teacher effectiveness in terms of student achievement standardized to mean zero and variance one. All are corrected for test measurement error. All except Kane & Staiger (2008) use within-school estimators.

Source: Hanushek & Rivkin (2010c)

Table 2. Estimates of the Value of Cognitive Skills from Mincer Earnings Functions

| Study	Earnings coefficient (ϕ)
Mulligan (1999)	0.11
Murnane et al (2000)	0.10-0.15
Lazear (2003)	0.12
Hanushek & Zhang (2009)	0.20
Hanushek & Woessmann (2009)	0.14
Chetty et al (2010)	0.18

Source: Hanushek (2011)

Table 3. Sensitivity of Demand Based on Earnings to Key Parameters

(marginal annual economic value of teacher one standard deviation above mean)

class size	$\theta = 0.6$				$\theta = 0.3$			
	$\sigma_T = 0.2$		$\sigma_T = 0.3$		$\sigma_T = 0.2$		$\sigma_T = 0.3$	
	$\phi = 0.13$	$\phi = 0.2$	$\phi = 0.13$	$\phi = 0.2$	$\phi = 0.13$	$\phi = 0.2$	$\phi = 0.13$	$\phi = 0.2$
5	$60,652	$93,573	$91,215	$140,923	$106,556	$164,741	$160,566	$248,858
10	$121,303	$187,145	$182,430	$281,847	$213,113	$329,482	$321,132	$497,715
15	$181,955	$280,718	$273,645	$422,770	$319,669	$494,223	$481,698	$746,573
20	$242,607	$374,290	$364,860	$563,693	$426,225	$658,964	$642,264	$995,431
25	$303,259	$467,863	$456,075	$704,617	$532,781	$823,706	$802,831	$1,244,288
30	$363,910	$561,435	$547,290	$845,540	$639,338	$988,447	$963,397	$1,493,146

Note: θ=depreciation rate ; σ_T=standard deviation of teacher quality; ϕ=labor market return to one standard deviation higher achievement.

Figure 1. Impact on Student Lifetime Incomes by Class Size and Teacher Effectiveness (compared to average teacher)
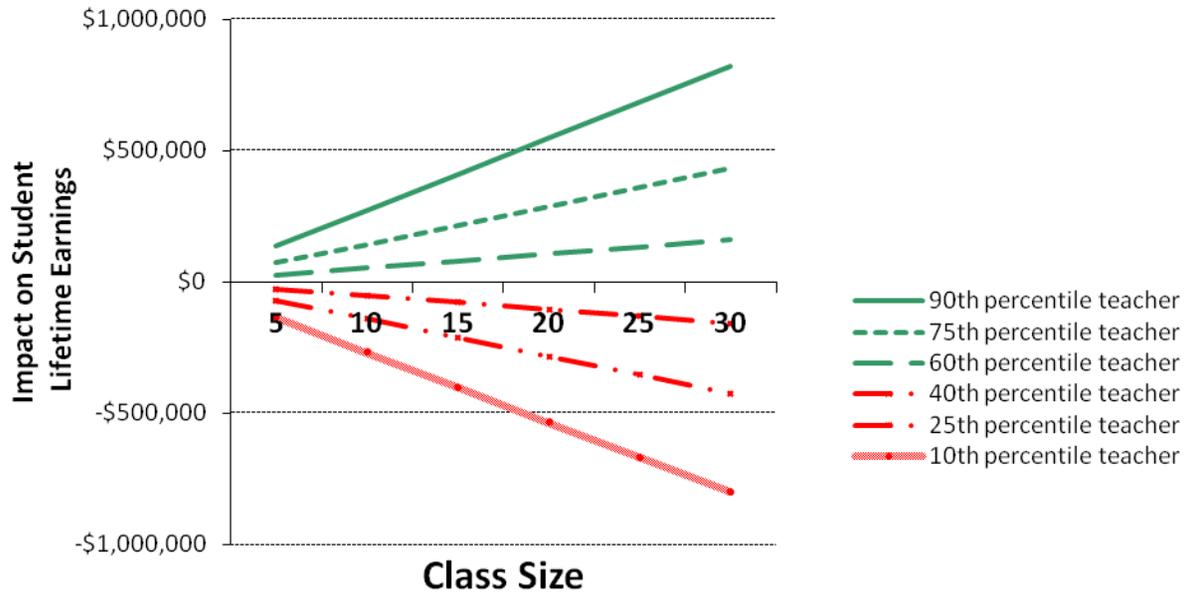
Additional Material

Sidebar: Economic growth

      An alternative estimate of the value of effective teachers focuses on the impact of student

performance on economic growth. Recent analysis has demonstrated a very close tie between

international assessments of achievement and a the country's economic growth rate (Hanushek &

Woessmann (2008)).   Hanushek (2011) links teacher quality to potential implications for performance

on international tests and, through that linkage, to effects on the aggregate U.S. economy.

Specifically, consider eliminating the bottom end of the teacher quality distribution and replacing these

teachers with average teachers.   Using the prior bound on variations in teacher effectiveness as

measured by achievement growth—specifically, 0.20 to 0.30 sd.—it is possible to see the impact of the

least effective teachers. Eliminating the least effective 5-8 percent of teachers would bring student

achievement up by 0.4 s.d. or higher on the international tests . Using prior estimates of the

achievement-growth linkage, Hanushek & Woessmann (2011) suggest that 0.4 s.d. in gains could

translate into a present value of added GDP of some $70 trillion (compared to a U.S. GDP of $15

trillion in 2011.

Key terms:

Value-added of teachers: the separate contribution of teachers to learning, independent of family, peer, and other influences

Persistence: the amount of learning in one year that carries over into the future.

Test measurement error: the portion of test scores that does not reflect the true knowledge of a student in a given domain.

Mincer earnings function: The common statistical model, named after Jacob Mincer, that explains individual earnings in terms of school attainment and experience.

LIFO: In education, "last in, first out" refers to the common practice of eliminating the teachers with least seniority if reductions are necessary.

Within school variance: The variation in achievement that is found inside schools as opposed to between schools; often the focus of estimates of teacher value-added.

Summary of Points

- The differences in effectiveness of teachers are not captured by common observable characteristics such as teacher experience or teacher degrees.
- Identification of teacher effectiveness on the basis of teacher value-added has shown that teacher quality varies dramatically even within individual schools.
- Teacher value-added has entered into policy discussions, leading to considerable current research to understand how estimation approaches affect the statistical estimates.
- Test measurement error influences the estimates of teacher value-added but its impact can be ameliorated with a variety of techniques.
- It is important to consider how family choices of schools and decisions of schools about classroom assignments influence estimates of teacher value-added.
- One topic of current investigation is the persistence of learning across grades, because this determines the ultimate influence of good or bad teachers on student outcomes.
- Elements of teacher value-added have entered directly into a variety of policy discussions including such things as which teachers should be let go when there is a reduction in the teacher force.

Future Issues

- A variety of technical issues continue to be topics for additional research: the best ways to deal with test measurement error; the amount of persistence in learning; and the interaction of teachers with classroom composition of students.
- On the policy side, further consideration of how estimates of individual teacher value-added should be incorporated into personnel policies including the retention of teachers and their pay.