



NATIONAL
CENTER for ANALYSIS of LONGITUDINAL DATA in EDUCATION RESEARCH

TRACKING EVERY STUDENT'S LEARNING EVERY YEAR

A program of research by the American Institutes for Research with Duke University, Northwestern University, Stanford University, University of Missouri-Columbia, University of Texas at Dallas, and University of Washington



*Estimating the Effect of
Leaders on Public Sector
Productivity: The Case of
School Principals*

GREGORY F. BRANCH,
ERIC A. HANUSHEK, AND
STEVEN G. RIVKIN

Estimating the Effect of Leaders on Public Sector Productivity: The Case of School Principals

Gregory F. Branch
University of Texas at Dallas

Eric A. Hanushek
Stanford University

Steven G. Rivkin
University of Texas at Dallas

Contents

Acknowledgements	ii
Abstract	iii
Introduction	1
The Productivity of School Leaders	3
The Texas Database	6
Principal Transitions	7
Fixed-effects Estimates of the Variation in Principal Effectiveness	8
Direct Estimation of a Lower Bound on the Variance in Principal Effectiveness	17
Principal Quality and Teacher Turnover	22
The Quality Dimension of Principal Transitions	27
Conclusions	29
References	32
Tables and Figures	35

Acknowledgements

This research has been supported by the Packard Humanities Institute and the Smith Richardson Foundation, the Spencer Foundation, and the Hewlett Foundation. The conclusions of this research do not necessarily reflect the opinions or official position of the Texas Education Agency, the Texas Higher Education Coordinating Board, or the State of Texas.

CALDER • American Institutes for Research
1000 Thomas Jefferson Street N.W., Washington, D.C. 20007
202-403-5796 • www.caldercenter.org

Estimating the Effect of Leaders on Public Sector Productivity: The Case of School Principals

Gregory F. Branch, Eric A. Hanushek, and Steven G. Rivkin

CALDER Working Paper No. 66

January 2012

Abstract

Although much has been written about the importance of leadership in the determination of organizational success, there is little quantitative evidence due to the difficulty of separating the impact of leaders from other organizational components – particularly in the public sector. Schools provide an especially rich environment for studying the impact of public sector management, not only because of the hypothesized importance of leadership but also because of the plentiful achievement data that provide information on institutional outcomes. Outcome-based estimates of principal value-added to student achievement reveal significant variation in principal quality that appears to be larger for high-poverty schools. Alternate lower-bound estimates based on direct estimation of the variance yield smaller estimates of the variation in principal productivity but ones that are still important, particularly for high poverty schools. Patterns of teacher exits by principal quality validate the notion that a primary channel for principal influence is the management of the teacher force. Finally, looking at principal transitions by quality reveals little systematic evidence that more effective leaders have a higher probability of exiting high poverty schools.

Introduction

Leadership quality is often cited as key to organizational success across such diverse places as boardrooms, sports arenas, national legislatures, and schools. Yet it is often quite difficult to distinguish cause and effect, as those anointed as great leaders may simply have been in the right place at the right time. The standard analytical approach to investigating the importance of leaders, developed in the analysis of private firms, has concentrated on the outcomes – typically revenues or profits – associated with the top managers who have varying backgrounds or experiences. But such an outcome-based approach requires an empirical model that effectively isolates the contribution of leadership from a variety of other, perhaps correlated, factors driving outcomes. Consideration of public sector leadership introduces special challenges because of the lack of market discipline and revenue and profit measures of outcomes.

Analysis of the impact of CEOs on corporate results has a lengthy history, where a variety of methods have been used to infer impacts from the pattern of market returns to a company.¹ In an important recent paper, Bertrand and Schoar (2003) provide an innovative new approach. Using semi-parametric methods with panel data to identify the effects of CEOs and top managers, they find that variation in the effectiveness of leaders explains a significant portion of the variance in profits and other outcomes. By simultaneously estimating both firm and manager fixed effects, the authors control for time-invariant differences among firms that could contaminate estimates of the variance in manager effectiveness.²

¹ There is an older literature with event studies on returns with changes in management (e.g., Johnson, Magee, Nagarajan, and Newman (1985)). Other work has looked the impact of management in family owned firms (see Bertrand (2009)) and at changes from firm founders (Bennedsen, Meisner Nielsen, Pérez-González, and Wolfenzon (2007)). Much of research on CEO performance has been motivated by investigations of CEO salaries (see, for example, Englmaier, Illing, and Sadka (2009) and the related symposium).

² However, changes over time in a variety of unobserved factors could inflate the estimated variance of manager behavior or productivity, even if the sample is limited to managers who work in at least two firms. Given the endogenous matching of firms and managers, there is a strong possibility that common unobservables are present in the multiple jobs worked during the sample period. For example, managers who were profitable in their first positions may have more opportunities to gain employment at firms on the upswing in terms of profits due perhaps

A second strand of literature important for our analysis has focused on management practices at the plant level, recognizing that CEOs are far removed from much of the actual management activity. Bloom and Van Reenen (2010) and related articles have documented the importance of differences in business and management practices across firms and countries by relating survey data to productivity differences. Their survey focuses on activities of plant managers in the manufacturing sector and attempts to describe key elements of management (such as the monitoring of employees or the provision of good hiring incentives). They find these surveyed attributes to be correlated not only with firm performance but also with aggregate productivity at the national level. It is at the same time not so much a direct study of the role and importance of managers as it is an investigation of management practices.

In this paper we combine elements of both strands of work in an investigation of variation in the quality of management in education. We begin our investigation of performance at the “plant” level by using the same semi-parametric approach used in Bertrand and Schoar (2003) to estimate the variance in principal effectiveness. We then extend the analysis in order to account more fully for the influences of time-varying unobservables by isolating the additional year-to-year fluctuation in outcomes that accompany changes in the principal. Following estimation of the variance in principal effectiveness, we investigate one widely discussed mechanism through which principals affect achievement: the management of teacher transitions. Importantly, because high teacher turnover can be associated with both improvement and decline in the quality of instruction, the level of turnover per se provides little information on the wisdom of principal personnel decisions. Therefore, we focus on the relationship between the quality of teachers who transition out of a school and the quality of principals.

to a negative random shock in the prior period. Moreover, nonpersistent factors orthogonal to manager quality or behavior will inflate the variance of the estimated fixed effects, even if they do not bias the estimates of the impact of individual CEOs. These are issues we consider below.

Throughout the entire analysis, we are particularly interested in the assignment of principals to schools serving disadvantaged students. From case studies and anecdotal accounts, the importance of principals seems most apparent when considering schools serving disadvantaged populations. It is frequently asserted that high poverty schools are hurt by being unable to attract and retain good principals. The demands of leading such schools, including the need to overcome less desirable working conditions in order to attract and retain high quality teachers, may amplify the importance of having an effective leader.³

The final component of the empirical analysis considers dynamics of the principal labor market. Do the best principals tend to move away from high poverty schools? Since pay and career advancement in the public sector are often insulated from performance in ways different from private sectors, it is not possible simply to track the pay of principals. Instead, we turn to a direct investigation of the quality of movers and stayers within schools, particularly those serving disadvantaged populations.

The productivity of school leaders

A dynamic and skilled school leader is frequently described as the key element of a high-quality school, and stories of the inspirational and effective principal are plentiful and oft-repeated.⁴ The leadership and decision-making provided by a school principal is proximate and tied directly to outcomes in her school, unlike that of a school superintendent of a large district who operates more like

³ Hanushek and Rivkin (2007), Boyd et al. (2011).

⁴ A large qualitative literature focuses on “effective schools” and in that generally places special emphasis on principals and leadership issues. See, for example, Edmonds (1979), Purkey and Smith (1983), or the case studies in Carter (2000).

a CEO in terms of providing broad policy guidance. But there is no clarity from past work about what attributes might lead to success.⁵

A distinctive feature of this analysis is the focus on the public sector.⁶ Unlike private industry, public institutions do not necessarily have a well-defined objective function, complicating the analysis of their behavior. At the same time, with the advent of accountability systems under No Child Left Behind (NCLB) and the complementary state systems, schools do have a direct incentive to maximize achievement on standardized tests.

In contrast to the traditional revenue and return focus of CEO and management studies, our study changes the performance measurement to student achievement. The advantage of this is that it insulates the study from exogenous factors such as prices, aggregate economic fluctuations, or prior capital investment decisions. However, test-measurement error and other test-related issues introduce other complications that must be addressed.

In broad terms, we take a generally agnostic view of the attributes of principals that are important and use school administrative data to concentrate on the more basic question of how important is variation in principal quality. Our administrative data, with very large longitudinal samples of principals and schools, provide a particularly rich source of information for the investigation of leadership effects, because they enable the direct estimation of productivity on the basis of contributions to student achievement.

Most prior analysis of principal effectiveness has been qualitative, though a small number of papers examine the determinants of principal effectiveness and any links with the labor market. Brewer (1993) and Eberts and Stone (1988) find that specific principal characteristics are related to high school student achievement, though the limited set of covariates may fail to account for important factors that

⁵ For some exploratory analysis of this, see Grissom and Loeb (forthcoming).

⁶ One place where management has previously been considered in public employment is in welfare offices, but much of this has concentrated on the narrower question of how managers react to different incentives; see, for example, Courty and Marschke (2004) and Heinrich and Marschke (2010).

could introduce bias. Loeb, Kalogrides, and Bêteille (forthcoming) look at how a school's value-added is related to the movement of teachers and suggest that this relationship is likely to reflect a principal's influence.

Finally, Cullen and Mazzeo (2008) investigate the link between principal salary growth and employment transitions on the one hand and principal effectiveness as measured by state accountability rating and achievement on the other. They find a positive relationship between salary and both accountability rating and student achievement. Nonetheless, the inclusion of just a limited set of student, school, and district controls leaves open the possibility that neither accountability rating nor student achievement provide a meaningful measure of principal effectiveness because of the contributions of unobserved student, family, peer and school factors.

Our central objective is to gain a better understanding of the overall impact and importance of principals, a fundamental issue for education policy. We do this in two different ways. First, following the general approach of Bertrand and Schoar (2003), we estimate models that include both principal and school fixed effects and use the principal fixed effects to generate estimates of the variance in principal effectiveness. Second, we generate lower-bound estimates of the variance in principal effectiveness from the relationship between year-to-year variation in school achievement and principal transitions (following the general approach used by Rivkin, Hanushek, and Kain (2005)). This second approach accounts for some of the year-to-year fluctuations in achievement that inflate variance estimates based on principal fixed effects but that are not caused by principals. Nonetheless, even the second approach may be vulnerable to biases introduced by unobserved factors. Therefore, we take a number of steps to mitigate any bias and examine the robustness of the results in both estimation frameworks. We also validate the estimates by assessing the impact of principal quality on decisions about teachers.

The Texas Database

The administrative data used in this project were constructed as part of the UTD Texas Schools Project. Working with the Texas Education Agency (TEA), this project has combined different data sources to create matched panel data sets of students, teachers, and principals. The panels include all Texas public school teachers, administrators, staff, and students in each year, permitting accurate descriptions of the schools for each principal.

The Public Education Information Management System (PEIMS), TEA's statewide educational database, reports key demographic data including race, ethnicity, and gender for students and school personnel as well as student eligibility for a subsidized lunch. PEIMS also contains detailed annual information on teacher and administrator experience, salary, education, class size, grade, population served, and subject. Importantly, this database can be merged with information on student achievement by campus, grade, and year. Beginning in 1993, the Texas Assessment of Academic Skills (TAAS) was administered each spring to eligible students enrolled in grades three through eight.⁷ These criterion referenced tests, which assess student mastery of grade-specific subject matter, are merged with the student and personnel information. Reading and math tests each contain approximately 50 questions, although the number of questions and average percent correctly answered varies across time and grades. We transform all test results into standardized scores with a mean of zero and variance equal to one for each subject, grade, and year, implying that our achievement measures describe students by their relative position in the overall state performance distribution.

The personnel data combine time as a teacher and as an administrator into total experience, so it is not possible to measure tenure as a principal accurately for those who begin their principal career prior to the initial year of our data (the 1990/91 school year). Therefore, for both the descriptive analysis

⁷ Many special education and limited English proficient students are exempted from the tests. In each year roughly 15 percent of students do not take the tests, either because of an exemption or because of repeated absences on testing days.

and the achievement modeling we concentrate on the period 1995-2001, and we allocate principals to precise experience and tenure categories in the early career while aggregating experience for six or more years.⁸

One of the strengths of our sample is the large number of principals and schools that are observed. Over the 1995-2001 period we observe 7,420 unique principals that yield 28,147 annual observations of principals. The large samples provide the opportunity to investigate differences by the poverty level of the school.

Principal transitions

Our estimation relies on changes in school outcomes that come with a change in principal. Therefore, it is important to understand the extent and character of principal turnover. Moreover, these principal moves are important in their own right. In parallel to concerns about teacher turnover, instability of leadership is often cited as a detriment to school improvement in high poverty and lower performing schools.⁹

We begin by describing the distribution of principals by job tenure and the pattern of transitions with specific attention to differences by school poverty share and mathematics achievement. Schools with a high proportion low income students (those eligible for a subsidized lunch) are more likely to have first year principals and less likely to have principals who have been at the school at least six years than those with a less disadvantaged population (Table 1).¹⁰ Nonetheless, the division of schools by

⁸ Note that the achievement tests in Texas were changed in 2002, thus making comparison of results over time difficult.

⁹ Teachers been shown to have preferences for the income, race, and achievement composition of students along with geographic locations; see Hanushek, Kain, and Rivkin (2004), Boyd, Lankford, Loeb, and Wyckoff (2005), Scafidi, Sjoquist, and Stinebrickner (2007). Some evidence of similar preferences for principals also exists in Loeb, Kalogrides, and Horng (2010).

¹⁰ Nationally in 2009, principal departures from a school are virtually identical for schools that did and did not participate in the free or reduced-price lunch program (Battle (2010)). Gates et al. (2006) find similar higher mobility in disadvantaged schools in North Carolina. As noted earlier the administrative data we use combine years

initial achievement produces much larger differences. The proportion of principals in their initial year with a school is roughly 40 percent higher in schools with the lowest average initial achievement than those with the highest average initial achievement; the proportion of principals that have been at their current school at least six years is roughly 50 percent higher in the highest achievement schools. Differences are also far smaller when ordered by the black or Hispanic enrollment share (not shown).

Table 2 reports annual principal transitions categorized by destination, new role, and tenure. Thirty percent of principals in our sample leave their school each year regardless of tenure in the school (compared to 20 percent nationally in 2009).¹¹ Perhaps more surprising, between 1 in 5 and 1 in 6 principals annually exit the Texas public schools, and the probability of exit does not vary much by tenure.¹² The probability of changing schools and remaining a principal rises from 5.9 percent following the first year at a school to 8.3 percent following the third through fifth years prior to falling back to 5.7 percent for those with tenures of at least six years. By comparison, the probability of transitioning to a non-principal role at a campus declines from 4.5 percent to 2.8 percent as tenure increases from one to at least six years. For those with at least two years of tenure approximately 2 percent make a transition to work as district administrators.

Fixed-effects Estimates of the Variation in Principal Effectiveness

The fundamental identification challenge is the separation of the contribution of principals from other factors that drive student achievement. In this section we describe the value-added model used to estimate principal fixed effects and the results of that estimation. Estimates are produced for a series

as a principal with years spent in other roles including teaching. Therefore, there is no information on tenure as a principal beyond the years observed in the sample.

¹¹ See Battle (2010).

¹² Note that limiting the sample to principals with no more than 25 years of experience in order to lessen the contribution of retirements has little impact on these patterns.

of specifications and samples in order to examine the sensitivity of the variance estimates to potential confounding factors including test measurement error and issues surrounding test construction.

Empirical model

Our basic value-added model relates achievement (A) for student i in school s with principal p in year y to prior achievement, observed student characteristics (X), time varying school and peer characteristics (C), and a vector of principal-by-school fixed effects (θ).¹³ Adding a random error (ε), the basic empirical model is:

$$(1) \quad A_{iy} = \lambda A_{iy-1} + \beta X_{iy} + \delta C_{sy} + \theta_{sp} + \varepsilon_{iy}$$

The concern of course is that other unmeasured factors will be correlated with the principal in each school and with the principal's effectiveness. One fundamental alternative to equation (1) is to add a vector of school fixed effects (γ) – implying that the impact of principals is estimated entirely on the within-school variation in achievement through comparing student performance when different principals run each school. This approach guards against systematic selection of the schools by students and principals, but it also puts extra strain on the data. Below, we consider both equation (1) and the expansion to include the school fixed effects.

In the empirical application, the vector X includes a full set of race and ethnicity indicators and indicators for subsidized lunch eligibility, special education participation, female and English as a second language. It also includes indicators for school changes: a switch to the earliest grade offered in a different school (including structural transitions from elementary to middle school) and a switch to other than the earliest grade offered in a new school. The vector C of time-varying aggregate influences includes average demographic characteristics for students in school s in year y including proportion low

¹³ While the general concept has been used in education for over three decades (see Hanushek (1979)), the recent addition of extensive administrative data bases has led to expansion of both the empirical analysis (Hanushek and Rivkin (2010b)) and the understanding of fundamental underlying estimation and interpretation issues (Meghir and Rivkin (2011)).

income, proportion classified as special needs, proportion that are new entrants to the school and proportion female. All regressions also include a full set of year-by-grade indicators to account for test and other statewide policy changes.¹⁴

As with related prior analyses of value-added for teachers, the lagged achievement term is included to capture past inputs from families, schools, and peers along with overall ability differences. This formulation supports a focus on just the contemporaneous inputs that are related to student learning in year y .

Variance estimates based on three year spells

We begin with estimates of the variation in principal effectiveness from a sample restricted to the first three years a principal leads a school (similar to the approach taken by Bertrand and Schoar (2003)). Because the impact of a principal on school quality likely increases with tenure given the persistence of personnel and other decisions, focusing on a sample of principals with the same tenure profile effectively controls for variation in quality related to differences in length of tenure. Importantly, principal impacts may become more positive or more negative over time, so simple linear or polynomial controls for tenure are not appropriate. On the one hand, over time a principal would be expected to learn about school operations, the effectiveness of various teachers, and other school specific factors, and such learning would presumably improve job performance. On the other hand, however, principal personnel decisions alter the stock of teachers and the school environment, and the impact of a principal increases over time as a principal accounts for more and more of the hiring and retention of the existing stock of teachers.

Table 3 reports the variance and 10th, 25th, 50th, 75th, and 90th percentiles of the principal-by-spell fixed-effect distribution disaggregated by the school poverty rate. From the first column, we find

¹⁴For the actual estimation, the data are aggregated prior to running the regressions to the campus-by-grade-by-year level to reduce the computational burden. All tables report absolute values of t-statistics based on robust standard errors clustered by campus.

that the overall standard deviation of principal effectiveness is estimated to be 0.207. This is a very large effect – perhaps unbelievably large – because it implies that a principal one standard deviation above the mean would, compared to an average principal, move mean achievement from the 50th percentile to the 58th percentile in one academic year. This impact is seen vividly by moving across the given percentiles for all teachers: a principal in the top quarter of the distribution gets average school gains of 0.11 s.d., while one in the bottom quarter has losses of 0.15 s.d..¹⁵

Looking down the first column reveals a monotonic relationship between the school poverty rate and the estimated variance. The variance in principal effectiveness increases noticeably with the share of low income students in the school. An examination of the achievement impacts at various quantiles of the principal quality distribution shows that the increase in dispersion as the poverty level rises is most pronounced at the lower end of the principal quality distribution: The difference in the effectiveness of the principal between the least and most disadvantaged school equals -0.20 standard deviations at the 10th percentile, -0.17 standard deviations at the 25th percentile, but only 0.07 standard deviations at the 90th percentile.

It is tempting to conclude that the best principals are quite evenly distributed across schools while the weakest principals are disproportionately found in high poverty schools. Note, however, that we cannot distinguish between two, quite different explanations for this pattern. High poverty schools may draw a disproportionate share of the overall least effective principals; or, the least effective principals may be evenly distributed but the adverse effect of low-productivity principals is larger in the most disadvantaged schools.

¹⁵ Note that the median principal produces slightly negative average gains. This is a specific subset of principals – those in their first three years of tenure in the school – and they are not far different from what is expected for all principals.

Test Measurement Issues

The interpretation of these estimates is, nonetheless, complicated further by potential test measurement issues. Not only can test measurement error and school differences in the skill distribution of students – the two major issues to be discussed below – contaminate the variance estimates, but these effects may differ systematically by poverty level.

Existing research on teacher quality confirms the special measurement issues arising in the study of achievement test results. Because the measurement of principal effectiveness is based on student assessments, both the structure of tests and errors in test measurement potentially complicate the estimation of principal quality.¹⁶ First, as Kane and Staiger (2002) and Jacob and Lefgren (2005, (2006) point out in other contexts, even in the absence of confounding influences, test-based quality estimates capture both random error and true effects. Consequently variance estimates overstate the actual variation in principal effectiveness, and the magnitude of any upward bias is likely to increase as school size decreases. Following Morris (1983) we utilize a shrinkage estimator to mitigate the impact of the test error. Normalizing average principal quality to zero, the adjusted quality estimate θ_p^a for principal p equals

$$(2) \quad \theta_p^a = \left(1 - \frac{V_p}{V_p + A}\right) \theta_p + \left(\frac{V_p}{V_p + A}\right) \bar{\theta}$$

where θ_p is the estimated fixed effect for principal p, $\bar{\theta}$ is the average fixed effect for all principals, V_p is the estimated variance of the fixed-effect estimate for principal p, and A is the estimate of the overall variance. Essentially, the shrinkage estimator pulls estimates toward the grand mean; the larger the error variance the more the adjusted fixed effect is shrunk toward the grand mean.

¹⁶ Test measurement errors include both issues of reliability and of validity. Reliability encompasses whether a test will consistently yield the same score for a given domain, while validity focuses on how well the test assesses the intended material.

Second, the construction of the test may affect the translation of measured student achievement into principal quality when differences in the achievement distribution exist across schools. For example, in schools where many students score near the top of the test at the beginning of the school year such that test ceiling effects become binding, principal quality might have very little effect on standardized test scores even if the principal is having substantial impact on the overall level of intellectual engagement and quality of instruction in the school. This is particularly relevant for the TAAS scores that we rely upon for our analysis, because the test is focused on lower level skills, and high achieving students could do very well on this test even without attending school. Consequently our test score-based-estimates of principal quality may produce a more compressed distribution for groups of schools with larger shares of initially high achieving students. Because we are particularly interested in the possibility that principals may have larger effects in schools serving predominantly disadvantaged students who tend to have lower initial scores, producing valid comparisons across schools grouped by poverty rate requires that this concern be addressed.

We investigate the sensitivity of the estimates to student composition using two alternative methods. The simplest incorporates a more flexible specification of prior achievement in order to capture differences in expected test score growth by initial score. Preliminary results (not reported) showed that polynomial terms for initial achievement had little effect on the estimates of principal effectiveness. However, the inclusion of quadratic or cubic terms may not fully address the problem given the skewness of the test score distribution.

Our preferred method is to weight performance by fixed weights in order to produce estimates of principal quality not influenced by differences in the test score distribution among schools. In particular, most schools have a mixture of students performing at different achievement levels, and it is the varying proportions that signify “advantaged” or “disadvantaged” in an achievement sense. Equation (3) shows the calculation used to produce weighted school-by-grade-by-year mean test score:

$$(3) \quad \bar{A}_{sgy} = \frac{1}{Nsgy} \sum_{d=1}^{10} \sum_{i=1}^{N_{sgyi}} \frac{F_d}{share_{sgyd}} A_{sgydi}$$

where F_d is the fixed weight share for test decile d , $share_{sgyd}$ is the actual share of students in schools in grade g in year y in decile d , N is enrollment, A is test score, and i indexes child. For this, we normalize scores on schools serving poverty students. The weights come from the distribution of 3rd grade mathematics achievement scores in 1994 for students in the top quartile of schools in terms of proportion of students eligible for a subsidized lunch. F_1 is the share in the bottom test score decile, F_2 in the second decile, and on up to F_{10} in the top decile. The scores of students at the low end of the distribution receive disproportionate weight in schools with small shares of such students relative to the weighting sample of higher poverty schools, while the scores of such students receive less than proportional weight in schools with a high concentration of initially low achieving students.

Table 4 presents three sets of estimates for all schools combined that address these two measurement issues separately and then together. The top panel reproduces the full-sample estimates from Table 3, the second panel contains information based on the estimates summarized in Table 3 but shrunk to the grand mean on the basis of the standard errors; the third panel reports the variance and quantiles for estimates of principal fixed effects based on value-added estimates of reweighted data that eliminate differences in the distribution of initial achievement among principals; finally, the bottom panel reports results based on reweighted data that have been shrunk to remove the effects of test error. Note that separate estimates by poverty quartile (not reported) follow the pattern of change found in the full sample.

The similarity between the results in the top two panels shows that shrinkage has virtually no effect on the estimated distribution of principal quality. Unlike estimates of teacher value-added that often rely on fewer than 15 observations for many teachers, estimated value-added for principals even

in quite small schools typically come from at least several hundred test scores. Consequently the variance of the error is likely to be quite small, and it is not surprising that shrinkage has little effect on the results.

In contrast, there are marked differences among schools in the initial achievement distribution, and this raises the possibility that the greater concentration of test scores at the lower end of the achievement distribution causes the larger variance in principal quality observed for high poverty schools. Re-weighting the scores such that the estimates for all principals are based on the same underlying distribution among student test score deciles lessens this concern. However, it does potentially increase the error variance by placing greater weight on smaller cells, and this may have a particularly large effect in very high and very low poverty schools

Consistent with the notion that re-weighting increases the error variance, the use of the shrinkage estimator has a much larger effect on the re-weighted data. Nonetheless, after shrinking, the variance estimates of the reweighted data (bottom panel) remain larger than those generated from the unweighted data (second panel), suggesting that test limitations may introduce downward bias into estimates of principal effects on learning. Yet because of the possibility that reweighting could exacerbate unobserved influences for students in small cells, we use unweighted data in the remainder of the analysis in order to avoid overstating the variation in principal effectiveness.

School Fixed-Effect Estimates

The prior estimates show substantial variation in principal value-added, but questions remain as to whether unobserved factors confound these estimates. If some salient differences among schools (that are not under the control of the principal) are not accounted for, estimates of θ_{sp} combine true principal effects with unobserved differences in other school or student factors. To understand the potential for such bias, we include school fixed effects (γ) in the specifications. This eliminates the influence of time-invariant school differences at the cost of restricting principal comparisons to others

who lead the same school. But, as a result, the estimated variance of the school fixed-effect estimates of principal quality omits all between school variation in principal effectiveness, potentially underestimating variations in principal quality.

Importantly, the restriction of the sample to the first three years in a school is not feasible in school fixed-effects models that identify principal effectiveness on the basis of within-school achievement differences, because the numbers of schools with two principals observed in their first three years is quite small. Therefore, in the fixed-effects models we consider the full sample and do not control for years at the school, although we do examine the effects of introducing controls for the first year of a spell.

Table 5 reports variance estimates from models with and without school fixed effects (with the latter estimates introduced to allow for potential differences due to sample composition). The first column shows that estimates are very similar to the comparable estimates in Table 4 for the sample restricted to the first three years at a school even though they include the wider range of tenure. But, the second column shows that the inclusion of the school fixed effects substantially reduces the variation in estimated principal quality.

In both the separate poverty categories and for the sample taken as a whole, the inclusion of school fixed effects reduces the estimated variance of principal quality by half. Although some of this reduction likely reflects the elimination of influences from confounding factors, some of it also likely reflects the elimination of real between-school variation in principal effectiveness. Looking down the second column, we see that the variance in effectiveness still increases monotonically with poverty concentration. In fact, the relative increase in the variance of effectiveness is noticeably larger when looking at just within-school variation.

The overall estimate of the variance of principal effectiveness from within-school variance, while smaller than the previous estimates, remains substantial. It indicates that a one standard deviation

increase in principal effectiveness raises school average achievement by slightly more than 0.1 s.d, a magnitude roughly comparable to the estimates of the within-school variance in teacher quality.

Direct Estimation of a Lower Bound on the Variance in Principal Effectiveness

The previous estimates of the distribution of principal quality began with estimation of the impact of individual principals and used these to understand what the overall distribution looks like. Although unobserved factors that are orthogonal to principal quality do not bias estimates of individual principal effectiveness, they do inflate the estimates of the variance in principal quality (and were the subject of various corrections). An alternative approach focuses on how the variance in achievement growth changes across student cohorts as new principals assume leadership in a school. From this, it is possible to estimate the variance in principal effectiveness directly and in so doing to circumvent some of the complications of the prior estimation.¹⁷

We begin with a modified version of the basic achievement equation from equation (1), where we aggregate across the students in a school to get the average achievement gain.¹⁸ Equation (4) relates the average gain in achievement in school s in year y to principal quality (θ), the quality of other school factors including student composition not under the control of the principal (δ), a school fixed effect (γ), and the school average error that includes unobserved student influences:

$$(4) \quad \overline{\Delta A_s^y} = \theta_s^y + \overline{\delta_s^y} + \gamma_s + \overline{u_s^y}$$

Consider the difference between successive years y and y' in average gains in achievement.

This eliminates all school effects that do not vary across the two years, leaving only year-to-year

¹⁷ This approach is similar to that in Rivkin, Hanushek, and Kain (2005), which identified teacher quality impacts through the movement of teachers in and out of schools.

¹⁸ For ease of presentation, this depiction restricts the coefficient on lagged achievement to equal 1; family and other time-varying school inputs are also combined into a single term.

differences in principal quality, other school influences, and other unobserved time-varying factors as determinants of the difference in achievement gain.

$$(5) \quad \left(\overline{\Delta A_s^y} - \overline{\Delta A_s^{y'}} \right) = \left(\theta_s^y - \theta_s^{y'} \right) + \left(\overline{\delta_s^y} - \overline{\delta_s^{y'}} \right) + \left(\overline{\nu_s^y} - \overline{\nu_s^{y'}} \right)$$

Squaring this difference yields a natural characterization of the observed achievement differences between years as a series of terms that reflect the variances and covariances of the principal and other school effects plus a catchall component e that includes all random error and cross-product terms between specific principal and other year-specific effects.

$$(6) \quad \left(\overline{\Delta A_s^y} - \overline{\Delta A_s^{y'}} \right)^2 = \left(\theta_s^y \right)^2 + \left(\theta_s^{y'} \right)^2 - 2\theta_s^y \theta_s^{y'} + \overline{\delta_s^y}^2 + \overline{\delta_s^{y'}}^2 - 2\overline{\delta_s^y} \overline{\delta_s^{y'}} + e_s$$

Taking the expectation of Equation 6 and assuming principals are drawn from common distributions over the restricted time period of the observations yields:

$$(7) \quad E \left(\overline{\Delta A_s^y} - \overline{\Delta A_s^{y'}} \right)^2 = 2 \left(\sigma_\theta^2 - \sigma_{\theta^y \theta^{y'}} \right) + 2 \left(\sigma_\delta^2 - \sigma_{\delta^y \delta^{y'}} \right) + E(e_s)$$

where σ_θ^2 is the variance of principal quality and $\sigma_{\theta^y \theta^{y'}}$ is the covariance in principal quality across the two years. Similarly, σ_δ^2 and $\sigma_{\delta^y \delta^{y'}}$ are the variance and covariance other school influences across years, again assumed to come from a common distribution.

Our focus is on the term $(\sigma_\theta^2 - \sigma_{\theta^y \theta^{y'}})$ that characterizes the influence of principal quality on squared changes in school-average achievement growth. Specifically, we relate principal turnover to the difference between the variance and covariance terms by regressing the squared difference in achievement gain on a dummy variable indicating that a school has a different principal in years y and y' . The parameter on that dummy variable would be equal to two times the variance in principal quality under three key assumptions: 1) the effect of a principal is fixed (no change over time); 2) principals are

assigned randomly to schools; and 3) principal turnover is orthogonal to other school changes that affect achievement gain.

This interpretation comes directly from these three assumptions. In schools with the same principal for years y and y' , the variance and covariance terms are equal and the principal contributes nothing to the variation in achievement gain. In schools with different principals in years y and y' , the covariance term is zero (under assumption 2), and the principal contributes σ_θ^2 to the year-to-year difference in achievement gain as shown in Equation (7).

Violation of the first assumption would almost certainly bias downward the estimate of the variance in principal quality. In our framework any contributions of an incumbent principal to year-to-year differences in school quality would be incorporated into the counterfactual squared difference in achievement gain attributed to factors other than the principal. If ignored, such changes will bias downward estimates of principal effects based on the coefficient for the different principal dummy. Because any such changes are most likely to occur in the initial years in a school, we include a dummy variable indicating that a difference is calculated over the first two years of a principal spell in some specifications.

Similarly, violation of the second assumption would also bias downward the estimate of the variance in principal quality. If individual schools tended to draw principals from a particular part of the quality distribution because of their salary structure, quality of central administration, or other factors, the covariance term will be positive rather than zero in schools with different principals in years y and y' . This would again reduce the difference in the expected value of $(\sigma_\theta^2 - \sigma_{\theta^y\theta^{y'}})$ between schools with different principals and schools with the same principals and thus attenuate the variance estimate. Essentially our approach ignores all between school variation in principal effects, because they cannot be identified.

In contrast, violation of the third assumption could potentially bias upward the coefficient on the different principal dummy variable. A range of shocks including changes in student demographics could increase both the probability of principal turnover and amplify fluctuations in achievement between cohorts. Therefore, it is imperative to account directly for variation over time in student characteristics that could be related to principal turnover. Although unobserved influences cannot be directly accounted for, the sensitivity of the estimate to the inclusion of significant time varying factors provides information on the likely existence and strength of unobserved factors that might confound the estimate.¹⁹

The sensitivity of the turnover coefficient to the period over which we calculate the squared differences also provides information on the validity of the estimates. Given the likely increase over time in the impact of a new principal, we would expect smaller principal turnover coefficients for specifications using squared differences in gains for adjacent years straddling the transition than for specifications using squared differences for non-adjacent years with a one year gap (e.g. 1999 and 1997 versus 1998 and 1997). The opposite finding of larger turnover coefficients for adjacent years would suggest the presence of additional turbulence around principal transitions and raise doubts about the validity of the estimates.

Table 6 reports the different lower-bound estimates of the variation in principal effectiveness, depending on the comparison group and different controls for student demographics. The estimates consider both adjacent and non-adjacent-year variances and the inclusion of first year principal indicators. The different-principal coefficient is positive and highly significant in all specifications, consistent with the existence of significant variation in principal quality. Note that the inclusion of an indicator for observations that span the first two years of a principal spell increases only slightly the estimated variance (compare Columns 2 and 3 and Columns 5 and 6).

¹⁹ Altonji, Elder, and Taber (2005) describe an informal approach to measuring the bias from unobserved factors on the basis of the sensitivity of the results to the inclusion of observed characteristics.

More importantly, the pattern of estimates provides support for the interpretation of the results as capturing the variation in principal effects. First, the non-adjacent-year specifications (col. 4-6) produce estimated standard deviations of principal effects that are roughly 10 percent larger than the comparable adjacent-year specifications. This is consistent with an increase over time in the influence of new principals and not consistent with the hypothesis that other changes or shocks coincident to the principal turnover bias upward the estimate of the variance in principal effects.

Second, the lack of sensitivity of the different-principal coefficient to the exclusion of the highly significant set of demographic characteristics suggests that neither observed nor unobserved factors inflate the estimates of the variation in principal effectiveness. Comparisons of Columns 1 and 2 and Columns 4 and 5 show that the exclusion of these variables has virtually no effect on the magnitude of the different-principal coefficient. As a result, we interpret these estimates as lower bounds on the variance in principal effectiveness.

Although roughly half as large as the fixed-effects estimates, these magnitudes for the lower bound remain educationally significant. A one standard deviation increase in principal quality translates into an increase of roughly 0.05 standard deviations in average student growth. By comparison, within-school estimates suggest that a one standard deviation increase in teacher quality raises achievement by somewhat more than 0.1 standard deviations.²⁰ However, teachers affect only their students, while principals affect all students in a school. Therefore, the overall impact from moving across the principal quality distribution substantially exceeds the benefit from a comparable movement across the teacher quality distribution.

Table 7 reports different principal coefficients for both the adjacent-year and interrupted-year specifications by poverty quartile. All coefficients are positive, though only those for the third poverty

²⁰ Hanushek and Rivkin (2010b) review estimates of the within-school variance in teacher quality and find an average of 0.11 s.d. for reading and 0.15 s.d. for math. The estimated benefits are similar to those of a ten student reduction in class size in early grades (Rivkin, Hanushek and Kain, 2005).

quartile reach significance at the five percent level. Nonetheless, although much noisier, the pattern of estimates mirrors that observed in previous tables: the lower-bound for the variance in principal quality increases monotonically with the poverty share.

Taking both the fixed-effect and the direct-estimation approaches together, three main findings emerge from the analysis: 1) there is significant variation in principal effectiveness; 2) variation increases with the school poverty share; and 3) the magnitude of the estimates falls substantially following the inclusion of school fixed effects and again following the focus on the additional variance associated with a change in principal. Although the latter finding is not definitive evidence that unobserved school differences bias estimates of principal effectiveness, the pattern of estimates is certainly consistent with this belief. Moreover, the findings indicate that sampling error introduces upward bias into the variance estimates.

Finally, the estimates in Tables 6 and 7 ignore any between-school variation in principal effects. This suggests that the findings likely understate the actual variance in principal effects.

Principal Quality and Teacher Turnover

The prior estimates rely on indirect measures of principal impact, namely student learning gains during a principal's tenure in a school. It does not include any observations about what a principal actually does to improve or limit learning. In part to validate the prior estimates of principal effectiveness, we turn to the interactions of principals with teaching staff. This additional investigation, however, also bears directly on a number of current policy discussions.

A primary channel through which principals can be expected to improve the quality of education is by raising the quality of teachers either through improving the instruction by existing teachers or through teacher transitions that improve the caliber of teaching (see, for example, Loeb, Kalogrides, and Bétaille (forthcoming)). Teacher turnover per se has been subject to considerable policy attention, given

the well-known difficulties experienced by new teachers. However, the benefits of reducing turnover clearly depend on both the composition of entering and exiting teachers and the transition rate.

Because it is difficult to separate improvement in the quality of instruction by current teachers from other changes that raise achievement, we focus on teacher turnover.²¹ We expect highly rated principals to be more successful at retaining more effective teachers and moving out less effective teachers. In contrast, we expect less highly rated principals to be less successful in raising the quality of the teaching stock, because of less skill in evaluating teacher quality, less emphasis on teacher effectiveness in personnel decisions, or less success at creating an environment that is attractive for better teachers. Although better principals may also attract and hire more effective teachers, the absence of quality measures for new teachers and the fact that many principals may have little control over new hires lead us to focus specifically on turnover.

Unfortunately, our data do not contain information on personnel decisions that would enable us to separate voluntary and involuntary transitions, and existing evidence suggests that teachers rather than principals initiate the majority of transitions. In addition, the Texas administrative data do not match students and teachers, meaning that we must draw inferences about teacher quality from grade average information. Although this avoids problems introduced by non-random assignment to classrooms that potentially biases estimates of teacher value-added (see Clotfelter, Ladd, and Vigdor (2006) and Rothstein (2010)), it prevents direct comparisons of the quality of individual stayers and leavers. An alternative, however, that we pursue below is to focus on the aggregate mobility patterns.

²¹ The estimation of teacher quality has been an active area of research, particularly in terms of value-added models. See, for example, Hanushek and Rivkin (2006, (2010b), Rothstein (2010), and McCaffrey, Sass, Lockwood, and Mihaly (2009). One result from this research is clarity about the difficulty of estimating teacher effectiveness with small samples of students, and this makes it difficult to identify the impact of specific principals on individual teachers.

The magnitude of teacher turnover

Teacher turnover has been a significant concern, especially in high poverty schools (see Hanushek, Kain, and Rivkin (2004) and Boyd, Lankford, Loeb, and Wyckoff (2005)).²² Prior to examining the relationship between principal quality and the pattern of teacher exits, we describe variation in the magnitude of teacher turnover by principal quality and school poverty.

Figure 1 shows that the teacher transition rate is highest in schools with the least effective principals, regardless of the rate of school poverty. This pattern is consistent with concerns about working conditions being a primary contributor to exits (see, for example, Hanushek, Kain, and Rivkin (2004), Hanushek and Rivkin (2007)). However, in the remaining three categories, there is little difference in overall turnover, with the exception of the highest poverty schools in which turnover in the 2nd quartile of principal quality is closer to the bottom quartile than the others. In fact turnover in the top principal quality schools is uniformly higher than that in the next category.

Prior analyses of teacher turnover suggest that the quality of movers differs systematically by destination. Hanushek, Kain, O'Brien, and Rivkin (2005) and Goldhaber, Gross, and Player (2011) both find that district-switchers tend on average to outperform teachers moving within the district and those exiting the public schools entirely. Consistent with that finding and the importance of leadership, Figure 1 shows that it is the rate of departure to other districts that tends to be much higher in schools with the least effective principals and lower in schools with the most effective principals. In fact district switching decreases monotonically as principal quality rises in all four poverty categories and differences in the rate of district switching is consistent with higher quality principals losing a smaller share of their more effective teachers.

²² For the quality dimension of teacher turnover, see Hanushek and Rivkin (2010a) and Goldhaber, Gross, and Player (2011).

Principal quality and teacher selection

With accurate information on teacher effectiveness and transitions, we could investigate whether better principals are more likely to dismiss the least effective teachers and reduce the probability that the more effective teachers depart voluntarily. In the absence of such information, however, we focus on the relationship within schools between the share of teachers that exit each grade and the grade-average value-added and observe how this varies with our estimates of principal quality. Specifically we investigate whether the relationship between the teacher exit rate in a grade and grade-average value-added is more negative in schools with higher quality principals.²³

This is obviously an imprecise test of whether better principals are dismissing their least effective teachers or retaining their most effective, as small grade average differences in mean value-added provide imperfect information on the probability that there is a very low performing or very high performing teacher in one grade as opposed to another. Nonetheless, while there will be cases in which the lowest performer teaches in the grade with the higher average teacher quality and vice versa, a very low performer is more likely to teach in the grade with the lower average achievement, and a very high performer is more likely to teach in the grade with higher average achievement.

We therefore employ a generalized difference-in-differences approach to investigate the variation in the relationship between grade differences in the rate of teacher exit and average value-added by principal quality. Equation (8) relates teacher turnover (T) in grade g in school s following year y to estimated value-added in grade g, school s, and year y ($\bar{A}_{gsy} - \lambda \bar{A}_{gs,y-1}$),²⁴ quartile of estimated principal quality (Q_i), student characteristics (X), average value-added interacted with quartiles of

²³ Note that it is not necessary for our purposes that the good principals explicitly make decisions to dismiss poor teachers. Teachers could be making these decisions based upon the atmosphere and actions of the principals (Boyd et al. (2011)) as long as the level of an influence over different quality teachers is correlated with principal quality.

²⁴ This measure of average test score gain is an approximation of grade value-added where λ is the production function estimate of the persistence of prior achievement (see equation 1).

estimated principal quality, and a vector of campus-by-year fixed effects. Note that the estimates of principal quality come from the analysis of principals' first three years in a school (Section V.B, above).

$$(8) \quad T_{gsy} = \alpha_1 (\bar{A}_{gsy} - \lambda \bar{A}_{gs,y-1}) + \sum_{p=2}^4 \alpha_{2p} Q_p + \sum_{p=2}^4 \left[\phi_p (\bar{A}_{gsy} - \lambda \bar{A}_{gs,y-1}) Q_p \right] + X\beta + \gamma_{sy} + \nu_{gsy}$$

Our primary interest is the relationship between principal quality and changes in the quality of the teaching force. If higher quality principals are more likely to succeed in retaining effective teachers and moving out ineffective teachers, we would expect $\phi_4 < \phi_3 < \phi_2 < 0$. That is, we expect the grade differences in teacher turnover and average value-added to be more negatively related in schools with a higher quality principal.

The inclusion of school-by-year fixed effects fully accounts for all fixed and time-varying unobserved differences between schools that raised such concern in the estimation of principal effectiveness. In this framework only variation between grades and within schools and years are used to identify the relationship between principal quality and the pattern of teacher exits. Therefore, we are able to use the more precisely estimated principal quality estimates based on the sample of principals in their first three years from a specification without school fixed effects. Note that the direct impacts of principal quality on the magnitude of teacher turnover (the α_{2p} 's) are not separately identified, because principal quality does not vary within school and year.

An additional issue with the estimation of Eq. (8) is that it relies both on achievement measures that contain error but also on estimated parameters from prior stages (λ and the Q_p 's). These measurement errors will tend to attenuate the estimates of interest. Before estimating equation (8), both the grade level value-added and principal quality estimates are shrunk toward the grand means of zero using Bayesian shrinkage estimation described above.

Table 8 reports estimates of our difference-in-differences model from specifications using the sample of all schools and the sample of schools in the highest poverty category.²⁵ As noted, the high poverty schools have the largest variance in principal quality and greatest rate of teacher turnover, elevating the concern about where the exiting teachers are drawn from in the school.

The results in Column 1 reveal a negative relationship between teacher turnover and grade average value-added that increases monotonically as principal quality rises, though only the interaction with the top quartile of principals is statistically significant. This pattern is consistent with the belief that management of the stock of teacher quality constitutes an important pathway through which principals affect school quality. The results in Column 2 are similar for the sample of high poverty schools: all coefficients are negative but only the coefficient on the interaction with the top quartile of principals is statistically significant.

Despite the aggregation to the grade level and reliance solely on the variation within school and year, the results provide strong evidence of more negative teacher selection out of a school in buildings run by highly effective principals, thus validating our prior quality estimates. In the absence of information on principal behavior and the proximate causes of each transition, it is not possible to know the precise actions that underlie the observed relationships. Nonetheless, it reinforces views about the importance of school leadership.

The Quality Dimension of Principal Transitions

Many bemoan high rates of turnover for both teachers and administrators in schools with high concentrations of poverty. Yet, as noted, the magnitude of any turnover problem clearly depends on whether high or low quality personnel are leaving. No prior analysis has been able to describe

²⁵ Estimates for the remaining poverty categories (not reported) are quite imprecise.

systematically any differences that might exist in the mobility patterns of high- and low-quality principals.

In order to focus on this issue, we describe principal transitions following their third year in a school by principal quality and by the share of students eligible for a subsidized lunch. Following the taxonomy in the descriptive section above, principals are observed to follow a variety of career moves: remaining in the same school as principal, remaining in the same school in another capacity, taking a principal position in another school, taking a different position in another school, moving to a central administrative position, or exiting the public schools entirely. Again principals are divided into four quality quartiles on the basis of estimates produced from the sample of the first three years in a school for specifications without school fixed effects. In this section the sample includes only principals with fewer than 25 years of total experience in order to minimize complications introduced by the retirement decision.²⁶

As seen in Table 9, principals in the lowest quartile are least likely to remain in their current position and most likely to leave the public schools entirely.²⁷ With the exception of the lowest poverty schools, however, there is not a monotonic relationship between the probability of remaining in the same position and principal quality. Principals in both the second and third quartiles are substantially more likely to remain than those in the bottom quartile. The most effective principals are always more likely to remain in the same position than those in the bottom quality quartile, but are generally more likely to move than those in the middle quality range. Overall, the probability of exiting the public schools entirely moves as the mirror image of the probability of remaining in the same position.

One troubling aspect of Table 9 from a policy perspective is the substantial share of low performing principals who transition to principal positions at other schools. This is particularly striking in

²⁶ The experience restriction, however, has virtually no effect on the observed transition patterns.

²⁷ Because of the very small shares of principals that transition to other positions, Table 9 focuses on principal transitions to other schools and out of the public schools entirely. (Appendix Table a1 reports shares for these other transitions).

the two highest poverty categories where over 12 percent of poor performers make such a move. In contrast, less than 7 percent of the poorest performers in the two lower poverty categories become principals at other schools. This may reflect the fact that it is difficult to separate the difficulties of the school circumstances from the quality of the principal in high poverty schools – leading both the principals and the district administrators to give principals from high poverty schools another chance at a different school.

The simple conclusion, nonetheless, is that the operation of the principal labor market – constrained in a variety of ways by inertia in salaries, historical absence of good performance measures, and other restrictions in public sector operations – does not appear to screen out the least effective principals. Instead they frequently just move to different schools, perhaps reflecting the bargain necessary to move out an ineffective leader in a public sector organization. Potentially this is where the CEO (superintendent) enters. A good superintendent may be good at decisions on retention and assignment of principals. We cannot directly look at this here, however, and this merits additional study.

Conclusions

An important facet of many school policy discussions is the role of strong leadership, particularly of principals. Leadership is viewed as especially important in revitalizing failing schools. This discussion is, however, largely uninformed by systematic analysis of principals and their impact on student outcomes. More generally, assertions about the importance of leaders in many occupations and circumstances lack empirical backing.

Understanding the impact of principals on learning is a particularly difficult analytical problem. The non-random sorting of principals and students among schools, the varying evolution of principal influence with tenure, and other school changes over time complicate efforts to identify principal effectiveness. The fact that the magnitude of the estimated variance declines substantially following the

inclusion of school fixed effects and again following the focus on the additional variance associated with a change in principal suggests that direct estimates of principal fixed effects combine real value-added with sampling error and other influences.

Nonetheless, the lower bound results based on direct estimation from principal turnover – the most conservative of the alternative estimation approaches – suggest the existence of substantial variation in principal effectiveness. A principal in the top 16 percent of the quality distribution (i.e., one standard deviation above average) will lead *annually* to student gains that are 0.05 s.d. or more higher than average *for all students in the school*. The finding that the variance increases with the school poverty rate is consistent with both the hypothesis that principal skill is more important in the most challenging schools and the hypothesis of larger variation in underlying skills of leaders entering high poverty schools – explanations that need not be mutually exclusive.

There are many channels through which principals influence school quality, though the precise mechanisms likely differ among districts due to variation in the regulatory and institutional structures that define principal authority. Because all principals participate in personnel decisions, we focus on the composition of teacher turnover. For the best principals, teacher turnover is concentrated in the grades within their schools with lower value-added, supporting the belief that improvement in the stock of teacher quality provides an important channel through which principals can raise the quality of education. Moreover, this relationship is strongest in high poverty schools, consistent with the finding of larger variation in principal quality in these schools.

Finally, patterns of principal transitions indicate that it is the least and most effective that tend to leave schools, suggesting some combination of push and pull factors. This pattern is again particularly pronounced in higher poverty schools. A troubling finding on transitions shows that a substantial share of ineffective principals in high poverty schools take principal positions in other schools and districts.

Much more needs to be learned about the dynamics of the principal labor market, just as much more needs to be learned about the underlying factors that contribute to teacher transitions.

We have emphasized the importance of variations in principal quality. But, it is also worthwhile putting this analysis into the context of understanding the importance of managers in the public sector. Recent work has emphasized the role of managers and leaders in profit making firms. Several different approaches have been introduced. One considers how important CEOs might be to differences in firm performance (Bertrand and Schoar (2003)). The other goes into more depth about the specific managerial processes that promote higher performance (Bloom and Van Reenen (2010)).

Borrowing from both perspectives and moving into the public sector, we focus on plant-level (school) management and address the fundamental question of how big is the variation of quality of public managers. In simplest terms, even with the constraints on public sector operations and the diffuse performance incentives, public sector managers appear to have a large impact on outcomes. From a policy viewpoint, added attention to the selection and retention of high quality managers would have a very high pay-off.

References

- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber. 2005. "Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools." *Journal of Political Economy* 113, no. 1 (February): 151-184.
- Battle, Danielle. 2010. *Principal Attrition and Mobility: Results From the 2008-09 Principal Follow-up Survey*. Washington, DC: National Center for Education Statistics (June).
- Bennedsen, Morten, Kasper Meisner Nielsen, Francisco Pérez-González, and Daniel Wolfenzon. 2007. "Inside the Family Firm: the Role of Families in Succession Decisions and Performance." *Quarterly Journal of Economics* 122, no. 2: 647-691.
- Bertrand, Marianne. 2009. "CEOs." *Annual Review of Economics* 1: 121-149.
- Bertrand, Marianne, and Antoinette Schoar. 2003. "Managing with Style: The Effect of Managers on Firm Policies." *The Quarterly Journal of Economics* 118, no. 4 (November): 1169-1208.
- Bloom, Nicholas, and John Van Reenen. 2010. "Why do management practices differ across firms and countries?" *Journal of Economic Perspectives* 24, no. 1 (Winter): 203-224.
- Boyd, Don, Hamilton Lankford, Susanna Loeb, and James Wyckoff. 2005. "The draw of home: How teachers' preferences for proximity disadvantage urban schools." *Journal of Policy Analysis and Management* 24, no. 1 (Winter): 113-132.
- Boyd, Donald, Pam Grossman, Marsha Ing, Hamilton Lankford, Susanna Loeb, and James Wyckoff. 2011. "The influence of school administrators on teacher retention decisions." *American Educational Research Journal* 48, no. 2: 303-333.
- Brewer, Dominic J. 1993. "Principals and student outcomes: Evidence from U.S. high schools." *Economics of Education Review* 12, no. 4 (December): 281-292.
- Carter, Samuel Casey. 2000. *No Excuses: Lessons from 21 High-Performing, High-Poverty Schools*. Washington, D.C.: The Heritage Foundation.
- Clotfelter, Charles T., Helen F. Ladd, and Jacob I. Vigdor. 2006. "Teacher-Student Matching and the Assessment of Teacher Effectiveness." *Journal of Human Resources* 41, no. 4 (Fall): 778-820.
- Courty, Pascal, and Gerald Marschke. 2004. "An Empirical Investigation of Gaming Responses to Explicit Performance Incentives." *Journal of Labor Economics* 22, no. 1 (January): 23-56.
- Cullen, Julie B., and Michael J. Mazzeo. 2008. "Implicit Performance Awards: An Empirical Analysis of the Labor Market for Public School Administrators." University of California, San Diego (December).
- Eberts, Randall W., and Joe A. Stone. 1988. "Student achievement in public schools: Do principals make a difference?" *Economics of Education Review* 7, no. 3: 291-299.
- Edmonds, Ronald R. 1979. "Effective schools for the urban poor." *Educational Leadership* 37, no. 1 (October): 15-24.
- Englmaier, Florian, Gerhard Illing, and Efraim Sadka. 2009. "Introduction to the Symposium on Executive Pay." *CESifo Economic Studies* 55, no. 3/4 (September/December): 399-404.

- Gates, Susan M., Jeanne S. Ringel, Lucrecia Santibañez, Cassandra Guarino, Bonnie Ghosh-Dastidar, and Abigail Brown. 2006. "Mobility and turnover among school principals." *Economics of Education Review* 25, no. 3 (June): 289-302.
- Goldhaber, Dan, Betheny Gross, and Daniel Player. 2011. "Teacher career paths, teacher quality, and persistence in the classroom: Are public schools keeping their best?" *Journal of Policy Analysis and Management* 30, no. 1 (Winter): 57-87.
- Grissom, Jason, and Susanna Loeb. forthcoming. "Triangulating principal effectiveness: How perspectives of parents, teachers, and assistant principals identify the central importance of managerial skills." *American Educational Research Journal*.
- Hanushek, Eric A. 1979. "Conceptual and empirical issues in the estimation of educational production functions." *Journal of Human Resources* 14, no. 3 (Summer): 351-388.
- Hanushek, Eric A., John F. Kain, Daniel M. O'Brien, and Steven G. Rivkin. 2005. "The market for teacher quality." Working Paper No. 11154. National Bureau of Economic Research (February).
- Hanushek, Eric A., John F. Kain, and Steve G. Rivkin. 2004. "Why public schools lose teachers." *Journal of Human Resources* 39, no. 2: 326-354.
- Hanushek, Eric A., and Steven G. Rivkin. 2006. "Teacher quality." In *Handbook of the Economics of Education*, edited by Eric A. Hanushek and Finis Welch. Amsterdam: North Holland: 1051-1078.
- Hanushek, Eric A., and Steven G. Rivkin. 2007. "Pay, working conditions, and teacher quality." *Future of Children* 17, no. 1 (Spring): 69-86.
- Hanushek, Eric A., and Steven G. Rivkin. 2010a. "Constrained Job Matching: Does Teacher Job Search Harm Disadvantaged Urban Schools?" NBER w15816. Cambridge, MA: National Bureau of Economic Research (March).
- Hanushek, Eric A., and Steven G. Rivkin. 2010b. "Generalizations about using value-added measures of teacher quality." *American Economic Review* 100, no. 2 (May): 267-271.
- Heinrich, Carolyn J., and Gerald Marschke. 2010. "Incentives and their dynamics in public sector performance management systems." *Journal of Policy Analysis and Management* 29, no. 1 (Winter): 183-208.
- Jacob, Brian A., and Lars Lefgren. 2005. "Principals as agents: Subjective performance measurement in education." mimeo J.F. Kennedy School of Government (June).
- Jacob, Brian A., and Lars Lefgren. 2006. "When principals rate teachers." *Education Next* 6, no. 2 (Spring): 59-69.
- Johnson, Bruce, Robert P. Magee, Nandu J. Nagarajan, and Harry A. Newman. 1985. "An analysis of the stock price reaction to sudden executive death: Implications for the management labor market." *Journal of Accounting and Economics* 7, no. 1-3: 151-174.
- Kane, Thomas J., and Douglas O. Staiger. 2002. "Volatility in school test scores: Implications for test-based accountability systems." In *Brookings Papers on Education Policy 2002*, edited by Diane Ravitch. Washington, DC: Brookings: 235-269.
- Loeb, Susanna, Demetra Kalogrides, and Tara Béteille. forthcoming. "Effective Schools: Teacher Hiring, Assignment, Development, and Retention." *Education Finance and Policy*.

- Loeb, Susanna, Demetra Kalogrides, and Eileen L. Horng. 2010. "Principal preferences and the uneven distribution of principals across schools." *Educational Evaluation and Policy Analysis* 32, no. 2 (June): 205-229.
- McCaffrey, Daniel F., Tim R. Sass, J. R. Lockwood, and Kata Mihaly. 2009. "The Intertemporal Variability of Teacher Effect Estimates." *Education Finance and Policy* 4, no. 4 (Fall): 572-606.
- Meghir, Costas, and Steven G. Rivkin. 2011. "Econometric methods for research in education." In *Handbook of the Economics of Education, Vol. 3*, edited by Eric A. Hanushek, Stephen Machin, and Ludger Woessmann. Amsterdam: North Holland: 1-87.
- Morris, Carl N. 1983. "Parametric Empirical Bayes Inference: Theory and Applications." *Journal of the American Statistical Association* 78, no. 381 (March): 47-55.
- Purkey, Stewart C., and Marshall S. Smith. 1983. "Effective Schools: A Review." *The Elementary School Journal* 83, no. 4 (March): 427-452.
- Rivkin, Steven G., Eric A. Hanushek, and John F. Kain. 2005. "Teachers, schools, and academic achievement." *Econometrica* 73, no. 2 (March): 417-458.
- Rothstein, Jesse. 2010. "Teacher quality in educational production: Tracking, decay, and student achievement." *Quarterly Journal of Economics* 125, no. 1 (February): 175-214.
- Scafidi, Benjamin, David L. Sjoquist, and Todd R. Stinebrickner. 2007. "Race, poverty, and teacher mobility." *Economics of Education Review* 26, no. 2 (April): 145-159.

Table 1. Distribution of Principals by Tenure at Current School and Student Demographic Characteristics

	Quartile of distribution of student characteristics	Years of Principal Tenure (percent)						All
		1	2	3	4	5	6 or more	
proportion eligible for subsidized lunch	lowest	17.8	15.3	12.6	10.2	8.2	36.1	100.0
	2nd	19.8	15.9	12.3	9.7	7.9	34.5	100.0
	3rd	20.5	17.1	13.7	9.7	7.9	31.1	100.0
	greatest	19.5	17.1	13.2	10.5	8.2	31.6	100.0
mathematics achievement	worst	22.7	19.4	14.3	9.8	7.4	26.3	100.0
	2nd	20.4	16.7	12.7	10.0	8.4	31.9	100.0
	3rd	18.1	15.3	12.5	10.1	7.9	36.3	100.0
	best	16.4	14.0	12.2	10.2	8.5	38.8	100.0

Table 2. Principal Transitions by Tenure

Transition	Years of tenure as principal at school (percent)			
	1	2	3 to 5	6 or more
Remain principal				
Same campus	72.5	68.0	69.2	70.8
New campus, same district	3.9	4.9	6.0	4.7
New district	2.0	2.5	2.3	1.0
Other school job				
Same campus	0.8	0.6	0.3	0.2
New campus, same district	2.0	2.1	2.2	2.1
New district	1.7	1.7	1.1	0.5
Job in Administration				
Same district central office	0.7	1.3	1.6	1.3
New district central office	0.5	0.7	0.7	0.4
Exit Texas public schools	16.0	18.1	16.7	19.0
Total	100.0	100.0	100.0	100.0

Table 3. Distribution of Principal Value-Added by Poverty Quartile

School Low- Income Share (<i>quartile</i>)	Standard Deviation	Average Annual School Achievement Growth by Percentile of Principal Effectiveness				
		10th	25th	50th	75th	90 th
Lowest	0.158	-0.180	-0.063	0.032	0.134	0.215
2 nd	0.176	-0.243	-0.142	-0.032	0.086	0.190
3 rd	0.207	-0.301	-0.162	-0.036	0.103	0.207
Greatest	0.263	-0.383	-0.236	-0.068	0.114	0.285
All	0.207	-0.286	-0.153	-0.019	0.109	0.222

Note: Students are classified as low income if they are eligible for a subsidized lunch. The sample of principals includes all observations for principals in their first three years of tenure at a school. Principal value-added is measured by the principal fixed effect that comes from a regression of mathematics achievement on lagged achievement, principal demographic characteristics, student demographic characteristics, and measures of student mobility. Full grade-year controls are also included

Table 4. Distribution of Principal Value-Added with Alternative Adjustments for Test Measurement Error

	Standard Deviation	Average Annual School Achievement Growth by Percentile of Principal Effectiveness				
		10th	25th	50th	75th	90 th
<i>Unadjusted</i>	0.207	-0.286	-0.153	-0.019	0.109	0.222
<i>Shrunk Estimates^a</i>	0.200	-0.280	-0.150	-0.020	0.106	0.214
<i>Re-weighted Estimates^b</i>	0.270	-0.182	-0.082	0.031	0.174	0.417
<i>Shrunk and Re-weighted Estimates^c</i>	0.241	-0.162	-0.071	0.032	0.164	0.385

Note: The sample of principals includes all observations for principals in their first three years of tenure at a school. Principal value-added is measured by the principal fixed effect that comes from a regression of mathematics achievement on lagged achievement, principal demographic characteristics, student demographic characteristics, and measures of student mobility. Full grade-year controls are also included.

- a. Bayesian shrinkage according to Eq. (2).
- b. Weighting by testing deciles according to Eq. (3).
- c. Both a. and b. adjustments applied.

Table 5. Estimated Standard Deviation of Principal Value-Added by Poverty Quartile and Whether the Specification includes School Fixed Effects

School Low-Income Share (<i>quartile</i>)	Without school fixed effects	With school fixed effects
Lowest	0.179	0.077
2 nd	0.192	0.095
3 rd	0.230	0.118
Greatest	0.277	0.138
All	0.224	0.110

Table 6. Lower-Bound Variance in Principal Quality Estimated from the Year-to-year Squared Difference in School Average Student Gains

	Adjacent-year difference			Interrupted-year difference		
	With student controls ^a			With student controls ^a		
Different principal	0.0052 (3.41)	0.0048 (3.16)	0.0049 (3.24)	0.0058 (4.35)	0.0056 (4.28)	0.0058 (2.87)
Principal in 1st year			0.0036 (2.97)			0.0026 (2.72)
Principal in 1st year*different principal						-0.0002 (0.11)
N	27,767	27,767	27,767	23,232	23,232	23,232
F test for student demographics ^b		29.05	28.21		32.41	32.03
Estimates derived from regression results						
Within-school variance of principal quality	0.0026	0.0024	0.0025	0.0029	0.0028	0.0029
Within-school standard deviation of principal quality	0.051	0.049	0.050	0.054	0.053	0.054

Note: Different-principal coefficients come from regressions of the year-to-year squared differences in school average student gains on an indicator for a different principal in the two comparison years. The absolute value of T-statistics are in parentheses.

- a. Student controls include ethnicity, gender, socioeconomic status, limited English proficiency, special education status and student mobility measures (shares).
- b. F-test that coefficients on all student demographics and mobility measures equal zero.

Table 7. Lower-Bound Variance in Principal Quality Estimated from the Year-to-year Squared Difference in School Average Student Gains by Poverty Quartile

	Poverty Quartile							
	lowest		2nd		3 rd		greatest	
	Adjacent year	Interrupted year	Adjacent year	Interrupted year	Adjacent year	Interrupted year	Adjacent year	Interrupted year
Different Principal	0.0017 (1.75)	0.0015 (1.51)	0.0019 (1.57)	0.0024 (1.68)	0.0048 (2.31)	0.0066 (2.59)	0.0090 (1.37)	0.0081 (0.96)
N	6,809	5,566	6,860	5,684	6,945	5,855	7,153	6,127
Estimates derived from regression results								
Within-school variance of principal quality	0.0009	0.0008	0.0010	0.0012	0.0024	0.0033	0.0045	0.0041
Within-school standard deviation of principal quality	0.029	0.027	0.031	0.035	0.049	0.057	0.067	0.064

Note: The different-principal coefficients come from regressions of year-to-year squared difference in school average student gains on an indicator for a different principal in the two comparison years. All regressions include the student demographic variables listed in Table 6 and indicators for comparisons that include the first year of principal tenure in a school. The absolute value of T-statistics are in parentheses.

Table 8. Estimated Effects of Grade Average Value-added and Principal Quality on Grade-Level Teacher Turnover for All Schools and the Schools in the Top Quartile in Terms of the Poverty Rate

	All schools	Highest poverty schools
Grade-level value-added	-0.023 (1.80)	0.006 (0.33)
Grade-level value-added *2nd quartile principal quality	-0.018 (0.89)	-0.065 (1.79)
Grade-level value-added *3rd quartile principal quality	-0.029 (1.35)	-0.025 (0.65)
Grade-level value-added *4th quartile principal quality	-0.079 (3.68)	-0.102 (3.16)

Note: Principal quality is measured by principal fixed effect for principals in first three years (see Table 3). The campus by year fixed-effect regressions (following Eq. (8)) also include the full set of student demographic variables. Regressions are weighted by enrollment. The absolute value of t-statistics are in parentheses.

Table 9. Principal Transition Distribution for Principals with less than 25 years of experience Quartile of Principal Quality and Poverty Quartile

Poverty Quartile and Principal Transition in Fourth Year	Quartile of principal quality ^a			
	worst	Q2	Q3	best
<i>Lowest quartile low-income</i>				
Same campus, principal	58.6	68.0	73.5	76.3
Moves campus, principal	4.6	4.1	4.6	4.3
Moves district, principal	2.3	1.6	3.8	4.3
Exit Texas Public Schools	27.6	21.3	12.1	8.6
<i>Second quartile low-income</i>				
Same campus, principal	52.4	70.2	81.5	71.7
Moves campus, principal	3.9	2.2	0.8	6.6
Moves district, principal	1.0	6.0	2.4	5.7
Exit Texas Public Schools	34.0	14.9	9.7	15.1
<i>Third quartile low-income</i>				
Same campus, principal	44.4	55.3	63.8	58.1
Moves campus, principal	7.5	5.7	9.5	8.1
Moves district, principal	4.5	4.9	1.9	7.3
Exit Texas Public Schools	35.3	25.2	19.1	23.4
<i>Highest Quartile low-income</i>				
Same campus, principal	62.6	73.2	72.3	67.4
Moves campus, principal	8.6	11.0	6.9	7.3
Moves district, principal	3.6	1.2	3.0	1.5
Exit Texas Public Schools	21.6	12.2	15.9	21.0

Note: a. Principal quality from estimates for principals with three or less years of tenure in school (see Table 3). Transition status refers to the year subsequent to the third year in a school. Complete moves in Appendix Table a1.

Appendix Table a1. Principal Transition Distribution for Principals with less than 25 years of experience, by Quartiles of Principal Quality and Poverty

Poverty Quartile and Principal Transition in Fourth Year	Quartile of principal quality ^a			
	worst	Q2	Q3	best
<i>Bottom quartile low-income</i>				
Same campus, principal	58.6	68.0	73.5	76.3
Same campus, other	1.2	0.0	0.8	0.0
Moves campus, principal	4.6	4.1	4.6	4.3
Moves campus, other	0.0	0.8	1.5	2.2
Same district, distr. Admin	0.0	4.1	1.5	2.2
Moves district, principal	2.3	1.6	3.8	4.3
Move district, other	3.5	0.0	1.6	0.0
Move district, district admin	2.3	0.0	0.8	2.2
Exits	27.6	21.3	12.12	8.6
<i>Second quartile low-income</i>				
Same campus, principal	52.4	70.2	81.5	71.7
Same campus, other	1.0	0.0	0.0	0.0
Moves campus, principal	3.9	2.2	0.8	6.6
Moves campus, other	2.9	0.8	0.8	0.9
Same district, distr. Admin	1.9	0.8	1.6	0.0
Moves district, principal	1.0	6.0	2.4	5.7
Move district, other	1.9	3.7	2.4	0.0
Move district, district admin	1.0	1.5	0.8	0.0
Exits	34.1	14.9	9.7	15.1
<i>Third quartile low-income</i>				
Same campus, principal	44.4	55.3	63.8	58.1
Same campus, other	0.8	0.0	0.0	0.0
Moves campus, principal	7.5	5.7	9.5	8.0
Moves campus, other	0.8	2.4	2.9	1.6
Same district, distr. Admin	3.8	2.4	1.0	0.8
Moves district, principal	4.5	4.9	1.9	7.3
Move district, other	2.3	1.6	1.9	0.8
Move district, district admin	0.8	2.4	0.0	0.0
Exits	35.3	25.2	19.1	23.4
<i>Top Quartile low-income</i>				
Same campus, principal	62.6	73.2	72.3	67.4
Same campus, other	0.0	0.0	0.0	0.0
Moves campus, principal	8.6	11.0	6.9	7.3
Moves campus, other	0.7	1.2	1.0	1.5
Same district, distr. Admin	0.0	0.0	1.0	0.0
Moves district, principal	3.7	1.2	3.0	1.5
Move district, other	0.7	1.2	0.0	0.7
Move district, district admin	2.2	0.0	0.0	0.7
Exits	21.6	12.2	15.8	21.0

Note: a. Principal quality comes from estimates for principals with three or less years of tenure in school (see Table 3). Transition status refers to the year subsequent to the third year in a school.

Figure 1. Teacher Transitions by Principal Effectiveness and Share of Students in a School Eligible for a Subsidized Lunch

