Who Leaves? The Implications of Teacher Attrition for Student Achievement

Don Boyd
University at Albany

Pam Grossman
Stanford University

Hamp Lankford
University at Albany

Susanna Loeb
Stanford University

Jim Wyckoff
University at Albany

We are grateful to the New York City Department of Education and the New York State Education Department for the data employed in this paper. The research is supported by funding from the Carnegie Corporation of New York, the National Science Foundation, the U.S. Department of Education and the Spencer Foundation. The views expressed in the paper are solely those of the authors and may not reflect those of the funders. Any errors are attributable to the authors.

Student achievement in many schools in the U.S. is disturbingly low. For example, 38 percent of middle schools in New York States had at least half of their students score below proficiency on the 8th grade math exam in 2005. Such results have led to substantial interest in developing policies to improve academic achievement, especially among low-achieving students. Naturally much of the policy interest focuses on the role that can be played by teachers and policies affecting the teacher workforce. Improving the quality of teaching has many dimensions, one of which is teacher retention. There is much consternation over high attrition rates, especially during the first few years of teachers' careers. The implicit assumption, that is not well documented, is that those leaving low-performing schools are the more able teachers. This leads many policymakers to the conclusion that to improve student achievement, policies must focus on improving teacher retention. However, there is remarkably little evidence that documents the effectiveness of teachers who leave low-performing schools.

This paper intends to address this issue by examining the value-added of novice teachers by their retention status, with particular attention paid to retention patterns in low-performing schools. The paper analyzes data on teacher attrition for New York City schools from 2000–2005 We find that teacher attrition rates do tend to be significantly higher in schools having lower student achievement, but find that the relationship between teacher effectiveness and teacher retention is less clear. In particular, first-year teachers that we identified as less effective in improving student test scores tended to have higher attrition rates than did those identified as being more effective teachers. The first-year differences are meaningful in size, however, the pattern is less consistent for teachers in their second and third years.

Teacher retention is potentially important to student learning for several reasons. For example, Ingersoll (2001) argues that major school staffing problems result from the "'revolving door' – where large numbers of qualified teachers depart their jobs for reasons other than retirement." If better teachers are more likely to leave teaching, especially in difficult-to-staff schools, students in these schools are more likely to have inexperienced, less qualified teachers, both of which have been shown to be associated with lower student achievement. Teacher retention also is an important factor in determining a school's learning environment for students. School administrators may find it more difficult to implement new policies, make necessary changes or meet higher standards when the teaching workforce is in constant flux. Poor retention of teachers also may indicate that other aspects of the school are not conducive to student learning. For example, teachers may be less likely to remain in a school that has ineffective leadership, which likely indicates that student achievement suffers as well.

Recent research by Richard Ingersoll (Ingersoll, 2001, 2004; Ingersoll and Kralik, 2004; Ingersoll & Smith, 2003 and 2004) and others (e.g., Hanushek, Kain and Rivkin, 2004; Johnson, 2004; Boyd, Lankford, Loeb, & Wyckoff, 2005; Loeb, Darling-Hammond, & Luczak, 2005) have dramatically increased our understanding of teacher retention. These studies have shown that both student and teacher characteristics affect teacher mobility. Teachers are more likely to stay in schools in which student achievement is higher and teachers – especially white teachers – are more likely to stay in schools with higher proportions of white students. Teachers who have scored higher on tests of academic achievement are more likely to leave, as are teachers whose home town is farther from the school in which they teach. Attributes of teachers and the students they teach appear to interact in important ways. In particular, teachers having stronger qualifications (as measured by scores on a general knowledge certification exam) are more likely to quit or transfer than are less-qualified teachers, especially if they teach in low-achieving schools.[1]

Research and other discussions pertaining to teacher attrition have often suggested that reducing teacher attrition would help improve the teacher workforce and the educational outcomes of students. How teacher attrition affects the overall quality of the teacher workforce depends upon a number of factors, including the typical gains in effectiveness teachers realize from additional years of experience as well as how the average quality of entering cohorts of teachers differ from those who entered the profession earlier. A crucial element is whether those teachers who leave teaching are more or less effective compared to their peers who remain. There is almost no research comparing the effectiveness of teachers by retention status. In the first such analysis, Hanushek, Kain, O'Brien and Rivkin (2005) find that the teachers leaving schools in an urban Texas district on average have lower student achievement gains than do the teachers who remain. This is true for those transferring within the district as well as those leaving. They find that the differences in teacher effectiveness is greater for teachers making intra-district transfers following their second and third years of teaching. Using NYC data, Boyd, Lankford, Loeb, Rockoff and Wyckoff (2006) contrast the effectiveness of those who stay in the same school, transfer within NYC or leave teaching in the City. They find little difference in the effectiveness between those who leave NYC and those who stay in the same school. However, they find that those who transfer within NYC on average are less effective than teachers who remained in the same school as well as less effective than the other teachers in their new schools.

[1] Boyd, Lankford, Loeb and Wyckoff (2005).

The goal of this paper is to provide additional evidence regarding whether and, if so, how the effectiveness of teachers differ by their retention status. Whereas the above two studies looked at the attrition of all teachers, we here focus on teachers in their first three years. In addition to contrasting differences in effectiveness by retention status, we analyze how attrition rates differ with the effectiveness of those individuals and whether such patterns differ across grade levels, subjects taught and schools grouped by the academic performance of students.

Teacher Attrition and Student Performance

To see the potential importance of teacher attrition, consider Table 1 which shows the cumulative transition rates for entering cohorts of NYC teachers in grades four through eight who started teaching between 2000 and 2003.[2] Here we track whether the individuals remained in the same school, transferred to another school within NYC, transferred to another public school district in New York or left the New York State public school system. We will characterize this last category as individuals having quit teaching, although a few individuals may actually have taken teaching jobs in private schools or public schools in other states. The table shows that after three years, roughly a quarter of the teachers had left the NYS public school system, another eight percent were teaching in other school districts in New York State. Twenty percent of the entering teachers were teaching in another school within NYC.[3] The net result was that by the start of the fourth year only 46 percent of the teachers continued to teach in the same school where they began their careers.

Much of the discussion regarding teacher retention is from a system-wide perspective; what percentage of teachers leave the profession and how do their attributes compare to those who remain in teaching. However, it is understandable that those specifically interested in educational outcomes in New York City would take a narrower view, being concerned about those who transfer out of NYC as well as those who leave the profession. From an even more local perspective, a school principal, parents and others likely will have a school-level perspective, primarily being concerned with teacher turnover in particular schools. Furthermore, a school-level perspective can be seen as having broad policy interest given the systematic sorting of students and teachers that exists across schools. In particular, there are large differences in student achievement across schools with the lower-achieving students, particularly those in urban areas, often taught by the least skilled teachers.

[2] Here we stop with the 2003 cohort so that we can follow the teachers for three years.
[3] The 20 percent figure does not include the additional roughly one percent of teachers who transferred within NYC but subsequently transferred to another district or quit teaching.

The importance of focusing on particular schools or groups of schools can be illustrated using New York State data. Historically, New York has tested students in 4th and 8th grade in math and English Language Arts (English and reading). Based on their performance, students are grouped into four levels. Proficiency is defined to be at level 3 or above. Level 1 indicates the student has not mastered any of the areas being tested. A large portion of New York's middle-school students score below proficiency and the vast majority of these low-performers are concentrated in relatively few schools. In a typical middle school in 2005, 12 percent of the students performed at level 1 and 33 percent performed at level 2 on the 8th grade math test As shown in Figure 1, 45 percent of the students scoring at level 1 were concentrated in 10 percent of New York's middle schools. Sixty-five percent of level 1 students were in 20 percent of the schools. In large part, this concentration reflects the fact that student failure rates are exceptionally high in these schools. In a typical school drawn from the 20 percent of the schools having the most level 1 students, 25 percent of the students scored at level 1 and 42 were at level 2; two-thirds of the students were not proficient. Unfortunately, the schools with the greatest number of students scoring below proficiency are remarkably constant over time. For example, the 20 percent of schools that enrolled 65 percent of the level 1 students in 2005 also contained at least 60 percent of the level 1 students in every year since 1999. Although many other schools have students who are below proficiency, the problem is most striking in a relatively few schools.[4]

The schools having high failure rates are concentrated in a small number of districts, with 66 percent of these schools in New York City alone. However, all New York City schools do not have high percentages of their students scoring at levels 1 or 2. Ordering NYC schools similar to that in Figure 1, one finds that 54 percent of all NYC students scoring at Level 1 are concentrated in 20 percent of the City's middle schools. Even within New York City there are big differences in student performance across schools.

How can this low achievement be addressed? The effectiveness of teachers accounts for a large share of the variance in student performance after students' own attributes have been taken into account. Thus policies must insure that low-performing schools have effective teachers. Unfortunately, teachers in the lowest performing schools are much less qualified than those in better performing schools (Lankford, Loeb and Wyckoff, 2002.) However, teacher qualifications, such as scores on certification exams and educational status, are not synonymous with teacher effectiveness. This paper explores the relationship between teacher

[4] A similar pattern is found for performance on the 8th grade ELA exam as well as the 4th grade math and ELA exams. However, the overall failure rates on the 4th grade exams are lower, in part because of improving scores in recent years – a trend that has not carried over to the 8th grade exams.

retention and teacher effectiveness to better understand how student achievement can be improved, especially in schools with historically low performance. In particular, we focus on how teacher attrition – to other NYC schools, other public schools and leaving the profession – is related to school-level differences in the effectiveness of teachers at improving the educational outcomes of students.

Measuring Teacher Effectiveness

A student's acquisition of skills and knowledge is a complex social enterprise. Thus, disentangling the contributions of a particular teacher from the contributions of other school inputs and the range of other determinants (e.g., home influences) is a challenge. We estimate the effectiveness of teachers in improving the educational outcomes of students using the fairly typical value-added model shown in (1). Y_{ijst} is the test score of the ith student taught by the jth teacher in school s in year t. We separately analyze scores in math and English language arts (ELA).[5] To net out the effects of a student's own background and past academic achievement,

$$Y_{ijst} = X_{it}\alpha + C_{jt}\beta + \delta_s + \gamma_{jt} + \varepsilon_{ijst} \quad (1)$$

the vector of student attributes, X_{it}, includes measures indicating the student's poverty status, whether the student is an English language learner, the student's race, school absences and suspensions in the prior year as well as the students' scores on both the math and ELA exams in the prior grade and those scores squared and cubed. Averages of these same variables for the student's classmates are included in C_{jt}.To capture the wide range of school-level factors affecting student outcomes, we include school fixed-effects, δ_s. We account for these various factors in an effort to isolate the value-added by a student's teacher, here measured by γ_{jt}. Note that this is a teacher-year fixed effect which is estimated separately for each year a person teaches and that estimates of the γ_{jt} only provide information about a teacher's effectiveness relative to other teachers in the same school, as a result of the model also including school fixed effects. Finally, ε_{ijst} is an error term capturing other factors affecting the student's score (e.g., test measurement error).

Given the inclusion of the various controls for student and class attributes along with the school fixed-effects, the estimated teacher-year effects are useful measures of within-school

[5] Test scores are normalized by grade and year to have zero means and standard deviations of one.

differences in teacher effectiveness.[6] Even so, the estimated teacher fixed-effects are subject to statistical error so that part of the observed differences in these estimates is measurement error, not actual differences in the effectiveness of teachers. We employ the empirical Bayes approach to adjust the teacher-effect estimates for such estimation error.[7] We use with-in school differences in these adjusted fixed-effects by level of experience as our measure of teacher effectiveness.[8] Teacher fixed effects were estimated separately for math and ELA as well as separately for grades 4-5 and 6-8. From the empirical-Bayes-adjusted estimate of a teacher's effectiveness in a particular year, we subtracted the mean adjusted estimate for all teachers in the school having the same level of experience. For example, our measure of a 6-8 grade, second-year teacher in math is calculated as the difference between the teacher's adjusted fixed effect and the average of the adjusted fixed-effects for all the second-year math teachers who taught in that same school during the period 2000-2005. We view such an approach as yielding fairly clean estimates of how effective each teacher is relative to their equally experienced peers in the same school.

Before turning to the discussion of the relationship between our measures of teacher effectiveness and teacher attrition, it will prove useful having a sense of how the estimated teacher effects are distributed. For teachers in grades 4 and 5 we have estimates of their effectiveness in both math and ELA. For first-year teachers in these grades, the standard deviation of our measure of effectiveness in math is 0.132 and 0.100 in ELA. The correlation between the effectiveness measures for the two subjects in 0.473. The standard deviation of our measure of effectiveness in grades 6-8 math is 0.101 and 0.081 in grades 6-8 ELA. Because of course specialization at the middle-school level, only 20 percent of the 6-8 grade teachers in our sample taught both subjects. For them, the effectiveness measures for the two subjects have a correlation of 0.379.[9]

---

[6] An alternative would have been to explicitly model the gains from experience and estimate teacher effects, instead of teacher-year effects. We chose the latter because of our focus on teachers in the first few years of their careers and our decision to compare teachers only to other teachers in their same school having the same level of experience.

[7] See Jacob and Lefgren (2005) for a clear summary of the approach we use as well as Kane and Staiger (2002) and Morris (1983) for earlier references.

[8] From the empirical-Bayes-adjusted estimate of a teacher's effectiveness in a particular year we subtracted the mean adjusted estimate for all teachers in the school having the same level of experience. It is these normalized, adjusted estimates of within school (and experience level) differences in teacher effectiveness we will utilize in the remainder of the paper.

[9] The standard errors for the test groupings are roughly comparable to those reported by Rockoff (2004) and smaller than those reported by Hanushek, Kain, O'Brien and Rivkin (2005) as well as Jacob and Lefgren (2005). Exploring the reasons for these differences go beyond the scope of this paper. However, one factor might be that our measures based on comparisons within school and for the same level of experience might do a better job of netting our differences associated with gains from experience.

Teacher Attrition in New York City

Teacher attrition for novice teachers in New York City is marked by two dominant themes. Teachers of low-performing students are much more likely to leave their current schools during their first two years of teaching than are teachers of high performing students. Within each of these student groups, teachers who are less effective in raising student achievement are more likely to leave their current school than are more effective teachers. These results hold up across performance on math or ELA and across grade levels. However, where these teachers go when they leave their current school does appear to differ.

Table 2 summarizes attrition rates for fourth and fifth grade teachers in their first two years.[10] Overall, school-level attrition for first-year teachers is somewhat greater than that for second-year teachers (20 v. 17 percent), but of greater interest is the destination of the teachers who leave. More than half of the teachers who leave following their first year transfer to another school within New York City, a third leave the New York State system (e.g., quit teaching) and only 15 percent transfer to the suburbs. Contrast this to second-year attrition where only 27 percent of those leaving transfer within New York City and more than half leave the New York State system. Attrition patterns differ somewhat when schools are disaggregated by the performance of their students based on the proportion who score at the lowest of four levels on the 4th grade math exam, indicating no proficiency across the three tested areas. In both the first and second year, we observe a greater proportion of teachers leaving the lowest performing schools than the highest performing schools such that by the end of the second year the difference between these groupings is 7 percentage points (71 v. 64 percent).

We are particularly interested in how teacher retention varies by the relative effectiveness of teachers within schools. In Table 2 teachers are grouped according to whether their school-level differences in effectiveness are in the top or bottom quartile of values or are in the inter-quartile range. Teachers whose math value-added places them in the top quartile of effectiveness have a cumulative two-year retention rate of 70 percent (= 100 · 0.848 · 0.823), while only 62 percent of those in the lowest-quartile remain after two years. These differences largely result from differences in within-district transfers and leaving the NYS system during the first year. Similar patterns result when teacher effectiveness is measured using student ELA value-added rather than math value-added.

[10] The analysis in the remainder of the paper employs data for those who started out as novice teachers in NYC between 1999-2000 and 2004-2005.

Given our measure of effectiveness is a within-school comparison, it is instructive to contrast retention rates for teachers within the same schools or for teachers working in similar schools. In terms of the latter comparison, the last panel in Table 2 shows that there are substantial differences in retention and patterns of attrition between more and less effective teachers in higher- and lower-performing schools. For example examine the math panel. In the quartile of schools having more low performing students, the school-level first-year retention rates for the teachers in the top quartile of math effectiveness is ten percentage point higher than for teachers in the bottom quartile (84 vs. 74 percent), with a similar difference in the second year. The result is that after two years, only 56 percent of the less effective math teachers remained in the same low-scoring school, compared to 71 percent of their relatively more effective peers. In higher scoring schools the first-year differences for more and less effective teachers are similar. However, in the second year the less effective math teachers have a slightly higher retention rate. The net effect is that the two-year cumulative retention rate in the higher scoring schools is six percentage points higher for the more effective teachers compared to their less effective peers (73 vs. 67). It is interesting that the cumulative school-level retention rates of the more effective teachers are roughly the same in the higher- and lower-scoring schools (73 vs 71 percent). However, for the less effective teachers, retention is 11 percentage points lower in the schools having lower scoring students than in those schools having higher scoring students (56 vs. 67 percent). When these same fourth and fifth grade teachers are grouped according to their effectiveness in teaching ELA, the differences in retention rates are in most cases similar, but somewhat smaller.

The more effective math teachers (top quartile) in high-scoring schools (top quartile) have two-year cumulative retention rates that are more than 17 percentage points greater than less effective math teachers in low-scoring schools (73 v. 56 percent). This overall difference aggregates two important components. Among first-year teachers working in lower-scoring schools, the percent returning to the same school is ten percentage points higher for the more effective math teachers, compared to those who are less effective (84 vs. 74). A smaller percentage of the more effective teachers transfer within NYC (9 vs. 15 percent). The more effective teachers are also four percentage points less likely to quit teaching. The pattern for first-year teachers working in the high-scoring quartile of schools is similar. The second factor contributing to the 17 percentage point difference is the lower attrition rates for both more and less effective teachers who teach in schools having higher scoring students. The attrition pattern differs somewhat for second year teachers. Less effective teachers in low-performing schools continue to transfer within the district at higher rates, but there is only a small difference in quit

rates between the two groups. The pattern for second-year teachers in the higher scoring schools is different. In particular, a higher percentage of the more effective teachers quit (12 vs 8 percent).

Table 3 summarizes attrition rates for grades 6-8 math and ELA teachers in their first two years of teaching.[11] The patterns for math teachers are shown in the top half of the table, with ELA teachers shown below. For both groups, those teaching in schools where students' academic performance is lower are substantially more likely to transfer to other schools in New York City and to leave teaching in the New York State public system. Among schools having relatively few students failing the exam, 80 percent of first-year math teachers returned to the same school the following year. In the quartile of schools having the highest failure rates, only 68 percent of first year math teachers returned. The largest component of this difference is the nine percentage point difference in NYC transfers (21 vs. 12 percent). The differences for second-year math teachers are even larger in magnitude. The patterns of attrition linked to the school groupings are similar for grades 6-8 ELA teachers, although the differences in first-year attrition rates are not as large.

Again, it appears that the interaction of student performance and teacher effectiveness provides insights into patterns of teacher retention. Seventy two percent of the more effective teachers in high-performing schools remain after two years (0.877*0.824), while only 40 percent (0.669*0.608) of less effective teachers in low-performing schools return for a third year. In the higher scoring schools, there is a ten percentage point difference in first-year retention rates between more and less effective teachers, reflecting the less effective teachers having both higher NYC transfer and quit rates. In schools with lower-scoring students, more effective teachers had higher retention rates both years. Whereas for first-year teachers this reflected a ten percentage point difference in percentage of teachers quitting, for second-year teachers the major factor was a large difference in within-NYC transfer rates.

The patterns for ELA teachers in these same grades are different. In both higher- and lower-scoring scoring schools, less effective teachers have higher school-level retention rates, in both the first and second year. A key element of this result is that the more effective teachers have higher within-NYC transfer rates.

Within-School Differences in Attrition

[11] Whereas the 4th and 5th grade teachers in our analysis taught both math and ELA, the vast majority of teachers in grades six through eight only taught one of the two subjects. However, those teaching both are included in the separate retention analyses for math and ELA teachers.

The above analysis explores how general retention patterns, and retention patterns within broad groupings of schools, differ with the effectiveness of teachers measured relative to that of their same-school peers. However, from a policy perspective a key question concerns how within-school differences in teacher retention are related to within-school differences in the effectiveness of teachers. We analyze within school differences using a multivariate analysis to model the probabilities of remaining in the same school (h=1), transferring to another NYC school (h=2), transferring to another NYS district (h=3) and leaving the NYS public system (h=4) using the fixed-effect multinomial logit model shown in (2) where E_j is either a scalar or vector measure of teacher effectiveness. The α_s^h are school-specific constant terms that capture the general attrition pattern in each school. The logit specification is attractive in that we do not have to estimate these school fixed effects. Rather they can be swept from the model, allowing us to focus on estimating how within-school differences in teacher effectiveness explain which teachers in a school leave conditional on the schools' overall attrition rates. Without loss of generality, we employ the normalizations $\alpha_s^1 = \beta^1 = 0$. Thus, by estimating the parameters

$$P_j^h = \frac{\exp(\alpha_s^h + \beta^h E_j)}{\sum_g \exp(\alpha_s^g + \beta^g E_j)}, \; h = 1, 2, 3, 4 \qquad (2)$$

β^2, β^3 and β^4, we can make inferences regarding how the transition probabilities within a school vary with the relative effectiveness of teachers. Note: a word of caution is warranted here in that nonzero values of the $\beta's$ do not necessarily imply a causal relationship. We are merely interested in obtaining a clearer empirical description of how attrition varies with teachers' own effectiveness.[12]

Parameter estimates for several different specifications are shown in Table 4 for 4th and 5th grade teachers.[13] The first specification (model A), examining teachers having up to three

[12] Because all the $\beta's$ enter the formula for each of the transition probabilities, interpreting the estimated coefficients is somewhat complicated. However, given the normalization we employ, it follows that $P_j^h / P_j^1 = \exp(\alpha_s^h)\exp(\beta^h E_j)$. Thus, a negative value of β^h would indicate that an increase in teacher effectiveness is associated a reduction in the probability of transition h relative to the probability of remaining in the same school. Note that the magnitude of the effect depends upon β^h as well as the school's baseline pattern of transition, captured by the school fixed effect, α_s^h; in particular, the magnitude of the effect of an increase in E_j will be larger as $\exp(\alpha_s^h)$ is larger.

[13] The models estimated are somewhat more complicated than the specification shown in (1). In particular, we included alternative-specific dummy variables indicating one's years of experience beyond

years of experience, indicates that a teacher's probability of making a NYC transfer, relative to that of remaining in the same school, is lower as the teacher's relative effectiveness in teaching math is higher. There is a similar inverse relationship for the probability of leaving the NYS system. These results are consistent with our findings above.

In contrast to the results for effectiveness in teaching math, we find no systematic relationship between the transition probabilities and individuals' effectiveness in teaching ELA. The estimated coefficients are small in magnitude relative to those estimated for math and their standard errors are somewhat larger. We estimated a range of different specifications, but never found evidence of a systematic relationship between the transition probabilities and 4th and 5th grade teachers' effectiveness in ELA. (In light of this finding, ELA effectiveness is not included in the other specifications reported in Table 4.)

A pertinent question is whether the systematic pattern found for math effectiveness holds equally for first, second and third year teachers. The results for model B indicates that it does not. Here the measure of teacher effectiveness in math is entered for second- and third-year teachers separately from that for first-year teachers. The qualitative pattern for first-year teachers is as described above. However, there is no evidence of a systematic relationship between math effectiveness and the transition probabilities for second and third year teachers. We estimated a range of models and found this result to be robust. For example, for second- and third-year teachers we estimated models (not shown) that included variables reflecting their current and/or first-year effectiveness as well as models including the average of teacher effectiveness in the current and prior-years. In none of the cases did we find within-school differences in math effectiveness of second- and third-year teachers systematically varying with within-school differences in retention patterns.

Does the retention relationship for first-year teachers hold across schools with low and high performing students? This is explored in model C where the variable measuring effectiveness in math is interacted with dummy variables indicating whether a school is in the quartile of NYC schools having the lowest student performance on the 4th grade math exam (lower-scoring school) and a dummy variable indicating the school is in the quartile having the highest student achievement (higher-scoring school). The results indicate there is no statistically significant difference in the math effectiveness coefficients for NYC transfers and quitting across these school groupings. However, for the middle-range schools (i.e., neither high nor low scoring) more effective teachers are estimated to be more likely to transfer to another district

---

the first year in order to account for the fact that transition patters for teachers in their second and third years might well differ systematically from those for first-year teachers.

compared to remaining in the same school. The pattern is different in both the lower- and higher-scoring schools for which one cannot reject the hypothesis that the probability of a NYS transfer is unrelated to a teacher's relative effectiveness.

Although the within-school relationship between retention and teacher effectiveness in math is statistically significant, is it of practical importance? Figures 2a and 2b show estimated retention probabilities of first-year teachers for typical lower- and higher-performing schools.[14] In the lower-performing school, a first-year teacher whose effectiveness in math is one standard deviation above the mean (+0.13) has an estimated probability of transferring that is four percentage points lower than that for a teacher one standard deviation below the mean (10.4 vs 14.5 percent). There is a five percentage point difference in the estimated probability of quitting (7.4 vs. 12.2 percent). In total, a two standard deviation improvement in teacher effectiveness reduces first-year attrition from the school by nine percentage points or by more than 25 percent.

Table 5 shows results for 6-8 grade math teachers. Model D estimated for first, second and third year teachers indicates that the relative likelihood of a first-year teacher in a school quitting is lower as a teacher's effectiveness is higher relative to other math teachers in the school. The coefficient for within NYC transfers is negative, but not statistically significant. The differences by level of experience are similar to what we found for grades four and five; the transition probabilities for second and third year teachers are not systematically related to the effectiveness of second- and third-year math teachers in grades 6-8. Finally, there are a few systematic differences across schools grouped by student performance on the 8th grade math exam (model E). In the schools other than those having high failure rates, it appears relatively more effective teachers are less likely to transfer within NYC. This is not the case in the low-scoring schools. However, in these schools the relatively more effective math teachers are less likely to leave teaching.

Results for ELA teachers in grades six through eight are shown in Table 6. The pattern of more effective first-year teachers being less likely to leave teaching holds here as well.

---

[14] In the two graphs the school-specific constants for the three transition, i.e., the α_s^h, were set so that the estimated transition probabilities equal the observed transition rates for the top and bottom quartiles of schools grouped by the percent of low-performing students where these probabilities are evaluated at $E_j = 0$. Some such normalization is needed since the school-fixed effects were not directly estimated.

As noted in footnote 5, even with a constant value of the coefficients for math effectiveness, the magnitude of the changes in the relative likelihoods associated with a change in effectiveness will differ depending upon the reference probability levels. Thus, as the graphs show, there is a larger difference in the probabilities associated with a given difference in effectiveness in a low performing school since transition probabilities are higher overall.

However, different from the other cases, it appears that more effective teachers are more likely to transfer within NYC. Finally, there are no meaningful differences across the school groupings.

Our findings for math and ELA in grades four and five raise intriguing questions as to what could explain the lack of a relationship between retention and effectiveness in ELA. One possibility relates to a finding by Jacobs and Lefgren (2005) who compare estimates of teacher's effectiveness, obtained using an approach quite similar to that we used, with subjective evaluations provided by principals. In particular, principals were asked to assess how effective teachers were "at raising student math (reading) achievement." They find that "principals appear good at identifying those teachers who produce the largest and smallest standardized achievement gains," and that principals were less successful at distinguishing teachers' effective in reading, compared to math. For example, Jacobs and Lefgren estimated that 69 percent of the teachers principals gave the lowest rating in math actually were correctly classified. In contrast, they estimated that principals only correctly identified 42 percent of the "bottom" teachers in reading. To the extent that principals "counsel out" teachers they perceive as being less effective, a difference in their ability to distinguish effective teaching in math vs. ELA could be a factor in explaining the difference in attrition results. However, two cautions are in order. First, as noted above, the observed pattern that more effective math teaches have lower attrition rates does not establish causation. Second, even if it were true that lower effectiveness were a cause for higher attrition, this would only raise questions as to why. When we observe teachers working (or not working) in particular jobs or teachers making job and career changes, separating out supply and demand factors is most difficult. For example, a teacher leaving a teaching job could be the result of "counseling out" or a more direct action taken by the employer, or a choice made by the teacher that some other opportunity was preferable to continuing in the same job. From this perspective, our findings regarding the difference between math and ELA effectiveness could reflect less effective teacher wanting to leave and teachers themselves being better judges of their effectiveness in math.

Differences in Teacher Effectiveness

Even in the cases where we find systematic within-school differences in retention patterns linked to the relative effectiveness of teachers, questions remains as the differences between those who remain in a school and those who make each of the three transitions. Consider first-year teachers in grades four and five who began teaching in schools grouped in the bottom quartile of student performance. Figure 3 shows the distributions of teacher

effectiveness in math with separate distributions included for those who returned to the same school the following year, transferred within NYC, transferred to another NYS district and left NYS teaching. Two differences stand out. First, comparing the distribution of effectiveness for those transferring within NYC with that for teachers remaining in the same school, there are quite meaningful differences in the left tails of the distributions. Teacher effectiveness in the middle of the distribution do not differ much between those making NYC transfers and those remaining in the same school. There are modest differences in the upper tails of the two distributions. Overall, it is the meaningful numbers of relatively ineffective teachers who transfer within the NYC system that stands out.

Consider the somewhat arbitrary thresholds of plus or minus one standard deviation in math effectiveness (0.13). Of those first-year teachers who remained in the same school, roughly 15 percent had efficiency estimates below this threshold. Fully one quarter of those transferring within NYC were below this threshold. In the top tails of the distributions, 17.5 percent of those who remained in the same school had effectiveness estimates that were at least one standard deviation above the mean. This contrasts to nine percent for those who transferred within NYC.

The second notable pattern in effectiveness is between those who leave teaching in New York State and those remaining in the same school. The distribution of effectiveness for those quitting is substantially to the left of the distribution for those remaining in the same school over most of the range of effectiveness. Here it is pertinent that almost two-thirds of the first-year teaching leaving the NYS system were less effective than the overall average for first-year teachers in their schools.

Even though such a statistic can mask important considerations, it is pertinent to consider the average difference in effectiveness between those who leave and those who stay in the same school after the first year. Table 7 shows the average within-school differences in teacher effectiveness for those making each of the transitions compared to those who remain in the same school. Those transferring within NYC on average are less effective in teaching math by -0.046 compared to those first-year teachers who remain in the same schools. The average difference for those who quit teaching is -.044. To put these numbers into perspective, 0.050 is the average difference in effectiveness between the second year of teaching and the first for those individuals who remain in the same school for a second year. For those teaching math in grades 6-8, the -0.043 average within-school difference for those leaving the NYS system, compared to those returning to the same school, is slightly bigger than the 0.038 gain in effectiveness from those returning having an extra year of experience. The average within-

school differences for ELA 4-5 are smaller, however they are comparable to the average difference in teachers' effectiveness between their first and second year (0.019).

Generally, our findings are consistent with those of Hanushek, Kain, O'Brien and Rivkin (2005), although we find more systematic differences for first-year teachers whereas they report bigger differences for teachers transferring after their second and third years. The above results differ from the finding in Boyd, Lankford, Loeb, Rockoff and Wyckoff (2006) that those teachers leaving NYC differed little in effectiveness compared to those remaining in the same school. Some of this difference is due to that analysis grouping together those leaving teaching and those transferring to other districts. However, much of the difference reflects our focus here on teachers in the first year or two of teaching. As discussed above, the pattern for the first year is quite different.

Characterizing NYC Transfers

There is substantial movement of teachers between schools in New York City. Whether these transfers on net hurt or improve student achievement is important in considering whether policies should be put in place to discourage principles from facilitating a game of musical chairs in which less effective teachers move from one school to another. However, before considering such issues, one need to have a better understanding of the nature of those transfers.

The top panel of Table 8 compares the new schools of fourth and fifth grade teachers who move to the schools where they initially taught. Because the nature and pattern of moves might be different for more and less effective first-year teachers, results are shown separately for individuals who were in the quartile of most effective teachers, the quartile of least effective teachers and the middle group. It is clear that the transitions for teachers in these groups were quite different, on average. The new schools where the less effective teachers moved on average had seven percentage points more minority students compared to the schools where they initially taught. Contrast this to the six percentage point reduction for the middle group. The new schools for the less effective teachers on average had three percentage points higher failure rate compared to the 1.5 percentage point reduction for other teachers. Thus, we see a systematic sorting of teachers linked to their ability to improve student outcomes.

As found in earlier work, those teachers who transfer on average are less effective than their peers in the new schools. In our case, the difference is -0.035, with a standard error of 0.017. However this average masks important differences related to the heterogeneity of those who transfer. As shown in the bottom panel of Table 8, those teachers who were in the top quartile of teacher effectiveness in their first school on average exceeded the effectiveness of

their peers in their new schools by 0.081. Similarly, those teachers who where in the bottom quartile of teacher effectiveness in their first year were less effective by -0.111 compared to their new peers. Note, however, that the pattern of moves is such that the second-year differences in effectiveness are somewhat smaller than they were in the first year.

Conclusion

Is teacher attrition a problem? Would a general reduction in teacher attrition help us attain the goal of improve the quality of the teacher workforce? Even though much more research will be needed to answer these questions, our findings as well as those of Hanushek, Kain, O'Brien and Rivkin (2005) make clear that to be effective, policies seeking to reduce retention rates need to be well targeted. Patterns of teacher attrition are such that individuals who leave teaching or make intra-district transfers on average are less effective than those teachers who remain. Even so, it is clear that many of those leaving are very effective. Our finding that almost two-thirds of the fourth and fifth grade teachers who transfer are "below average" illustrates the point. It follows that a third of those transferring are better than average, with many of these individuals being quite effective in improving the educational outcomes of students.

Those interested in Implementing targeted retention efforts face several challenges. First, it will be necessary to identify the effective teachers to be targeted. Second, strategies for targeting incentives and other policies in an effort to retain these individuals would need to be designed and implemented. There are numerous challenges associated with Introducing value-added assessments of teacher effectiveness and questions regarding the precision with which the effectiveness of individual teachers can be measured. However, the bigger challenge is likely to be the political and administrative hurdles that would have to be faced in implementing the types of targeted incentives needed to retain particular individuals.

Our finding that large numbers of relatively ineffective teachers are transferring between schools raises important questions regarding district policies that allow ineffective teachers to transfer. Even though "passing the buck" may be a rational strategy from the perspective of a local school official, doing so does not help improve the overall quality of the teacher workforce.

Table 1

Cumulative Teacher Turnover for Entering NYC Teacher, 2000-2003

	Remain in Same School	Transfer Within NYC	Transfer Out of NYC	Quit Teaching
After first year	76.8	13.1	3.1	7.0
After second year	59.2	18.0	6.3	16.5
After third year	46.4	20.0	8.1	25.5

Table 2
Attrition for First and Second-Year NYC Teachers in Grades 4-5, by Student Test Performance and Estimated Teacher Effectiveness, 2000-2005

		First Year Teachers				Second Year Teachers				
		N	Remain in same school	Transfer within NYC	Transfer outside NYC	Leave NYS system	N	Remain in same school	Transfer within NYC	Transfer outside NYC	Leave NYS system
All first-year teachers in grades 4-5		3556	0.804	0.102	0.029	0.065	2054	0.830	0.047	0.035	0.089
Schools ranked by student performance											
Higher scoring schools (top quartile)		599	0.823	0.107	0.022	0.048	357	0.863	0.025	0.028	0.084
Middle scoring schools (middle 50%)		2060	0.804	0.092	0.035	0.068	1200	0.828	0.042	0.038	0.093
Lower scoring schools (bottom quartile)		897	0.790	0.119	0.021	0.069	497	0.811	0.074	0.032	0.082
Teachers grouped by estimated effectiveness, Math											
Relatively more effective (top quartile)		956	0.848	0.074	0.031	0.046	521	0.823	0.042	0.036	0.098
Middle group (middle 50%)		1654	0.799	0.106	0.027	0.067	997	0.846	0.048	0.028	0.078
Relatively less effective (bottom quartile)		946	0.767	0.121	0.031	0.081	536	0.808	0.049	0.045	0.099
Teachers grouped by estimated effectiveness, ELA											
Relatively more effective (top quartile)		959	0.825	0.084	0.036	0.054	532	0.844	0.039	0.030	0.086
Middle group (middle 50%)		1641	0.796	0.108	0.026	0.070	997	0.826	0.046	0.040	0.087
Relatively less effective (bottom quartile)		956	0.797	0.108	0.027	0.068	525	0.823	0.055	0.029	0.093
Elementary Schools	**Teacher Effectiveness - Math**										
higher scoring students	Relatively more effective	163	0.883	0.067	0.018	0.031	97	0.825	0.021	0.031	0.124
	Relatively less effective	157	0.777	0.134	0.032	0.057	97	0.866	0.031	0.021	0.082
lower scoring students	Relatively more effective	247	0.842	0.089	0.024	0.045	122	0.844	0.049	0.025	0.082
	Relatively less effective	249	0.739	0.149	0.024	0.088	136	0.757	0.103	0.044	0.096
Elementary Schools	**Teacher Effectiveness - ELA**										
higher scoring students	Relatively more effective	159	0.830	0.094	0.025	0.050	91	0.868	0.011	0.044	0.077
	Relatively less effective	176	0.807	0.125	0.023	0.045	88	0.807	0.057	0.023	0.114
lower scoring students	Relatively more effective	245	0.812	0.094	0.037	0.057	125	0.888	0.048	0.008	0.056
	Relatively less effective	240	0.783	0.129	0.017	0.071	128	0.789	0.109	0.016	0.086

Table 3
Attrition for First and Second-Year NYC Math and ELA Teachers in Grades 6-8, by Student Test Performance and Estimated Teacher Effectiveness, 2000-2005

		First Year Teachers					Second Year Teachers				
		N	Remain in same school	Transfer within NYC	Transfer outside NYC	Leave NYS system	N	Remain in same school	Transfer within NYC	Transfer outside NYC	Leave NYS system
All first-year math teachers in grades 6-8		1579	0.742	0.167	0.023	0.069	807	0.690	0.145	0.031	0.134
Schools ranked by student performance											
Higher scoring schools (top quartile)		250	0.804	0.12	0.004	0.072	85	0.847	0.059	0.035	0.059
Middle scoring schools (middle 50%)		839	0.762	0.154	0.029	0.055	562	0.687	0.133	0.032	0.148
Lower scoring schools (bottom quartile)		490	0.678	0.208	0.022	0.092	160	0.619	0.231	0.025	0.125
Teachers grouped by estimated effectiveness, Math											
Relatively more effective (top quartile)		438	0.776	0.153	0.030	0.041	203	0.695	0.148	0.010	0.148
Middle group (middle 50%)		745	0.733	0.173	0.020	0.074	396	0.682	0.144	0.040	0.134
Relatively less effective (bottom quartile)		396	0.720	0.169	0.020	0.091	208	0.702	0.144	0.034	0.120
Middle Schools	**Teacher Effectiveness - Math**										
higher scoring students	Relatively more effective	65	0.877	0.108	0	0.015	34	0.824	0.088	0.029	0.059
	Relatively less effective	60	0.783	0.15	0	0.067	29	0.862	0	0.034	0.103
lower scoring students	Relatively more effective	143	0.706	0.21	0.035	0.049	72	0.667	0.153	0.014	0.167
	Relatively less effective	133	0.669	0.158	0.023	0.15	74	0.608	0.243	0.027	0.122
All first-year ELA teachers in grades 6-8		1840	0.753	0.147	0.026	0.074	919	0.724	0.127	0.034	0.115
Schools ranked by student performance											
Higher scoring schools (top quartile)		167	0.778	0.114	0.030	0.078	90	0.833	0.033	0.056	0.078
Middle scoring schools (middle 50%)		1261	0.762	0.137	0.028	0.073	641	0.749	0.109	0.027	0.115
Lower scoring schools (bottom quartile)		412	0.714	0.189	0.019	0.078	188	0.585	0.234	0.048	0.133
Teachers grouped by estimated effectiveness, ELA											
Relatively more effective (top quartile)		505	0.733	0.174	0.026	0.067	232	0.681	0.164	0.039	0.116
Middle group (middle 50%)		859	0.773	0.123	0.033	0.071	450	0.751	0.113	0.029	0.107
Relatively less effective (bottom quartile)		476	0.737	0.160	0.015	0.088	237	0.713	0.118	0.038	0.131
Middle Schools	**Teacher Effectiveness ELA**										
higher scoring students	Relatively more effective	36	0.806	0.083	0.028	0.083	24	0.750	0.000	0.083	0.167
	Relatively less effective	35	0.857	0.057	0.000	0.086	19	0.895	0.053	0.000	0.053
lower scoring students	Relatively more effective	128	0.656	0.234	0.031	0.078	53	0.585	0.264	0.019	0.132
	Relatively less effective	118	0.754	0.178	0.000	0.068	55	0.60	0.182	0.091	0.127

Table 4
Estimates of Parameters Reflecting How Transition Probabilities Vary with Teacher Effectiveness, Fixed-Effect Logit Models for Teachers in Grades 4-5

	NYC transfer		NYS transfer		Leave NYS system	
	β_2	s.e.	β_3	s.e.	β_4	s.e.
A: Model including measures of math and ELA effectiveness, 1st – 3rd year teachers						
Math effectiveness	**-1.089***	(0.402)	0.163	(0.628)	**-0.937***	(0.406)
ELA effectiveness	-0.223	(0.547)	0.151	(0.871)	0.127	(0.562)
B: Model entering math effectiveness separately for 2nd and 3rd year teachers						
Effectiveness - 1st year teachers	**-1.653***	(0.429)	0.703	(0.793)	**-2.295***	(0.530)
Effectiveness – 2nd & 3rd years teachers	-0.133	(0.623)	-0.226	(0.761)	0.244	(0.473)
C: Model for 1st year teachers with interactions for schools having lower and higher scoring students						
Effectiveness in math	-1.734*	(0.595)	1.806***	(1.064)	-2.172*	(0.734)
x high student failure dummy	-0.264	(0.960)	-1.676	(1.8134)	-0.664	(1.154)
X low student failure dummy	0.886	(1.152)	-3.482***	(2.099)	0.467	(1.412)

Table 5
Estimates of Parameters Reflecting How Transition Probabilities Vary with Teacher Effectiveness, Fixed-Effect Logit Models for Math Teachers in Grades 6-8

	NYC transfer		NYS transfer		Leave NYS system	
	β_2	s.e.	β_3	s.e.	β_4	s.e.
D: Model entering math effectiveness separately for 2nd and 3rd year teachers						
Math effectiveness - 1st year	-0.368	(0.672)	0.010	(1.732)	**-3.208***	(0.935)
Math effectiveness – 2nd & 3rd years	-0.409	(0.888)	-1.013	(1.830)	0.280	(0.897)
E: Model for 1st year teachers with interactions for schools having lower and higher scoring students						
Math effectiveness	-1.871**	(0.958)	-0.521	(2.474)	-1.137	(1.435)
Math effect. * high student failure	3.726*	(1.455)	0.983	(3.706)	-4.212**	(1.947)
Math effect. * low student failure	0.401	(2.123)	1.189	(5.613)	-1.126	(3.158)

Table 6
Estimates of Parameters Reflecting How Transition Probabilities Vary with Teacher Effectiveness, Fixed-Effect Logit Models for ELA Teachers in Grades 6-8

	NYC transfer		NYS transfer		Leave NYS system	
	β_2	s.e.	β_3	s.e.	β_4	s.e.
F: Model entering ELA effectiveness separately for 2nd and 3rd year teachers						
Effectiveness in ELA - 1st year	4.226**	(1.932)	-2.082	(5.947)	-7.731*	(2.633)
Effectiveness in ELA - 2nd & 3rd years	2.969	(2.338)	-2.436	(3.768)	-0.505	(2.223)
G: Model for 1st year teachers with interactions for schools having lower and higher scoring students						
Effectiveness in ELA	3.485***	(1.993)	-2.207	(6.107)	-6.391**	(2.731)
Effect. In ELA * high student failure	2.704	(2.099)	0.793	(6.731)	-4.227	(2.676)
Effect. In ELA * low student failure	0.447	(4.350)	-1.600	(12.885)	0.089	(6.047)

Table 7
Average Within-School Differences in Teacher Effectiveness Comparing Those Making Transitions to Teachers Remaining in the Same School

NYC transfer	NYS transfer	Leave NYC
Math Grades 4-5		
-0.046	0.007	-0.044
(0.009)	(0.015)	(0.009)
ELA Grades 4-5		
-0.019	0.012	-0.015
(0.007)	(0.012)	(0.007)
Math Grades 6-8		
-0.011	0.003	-0.043
(0.008)	(0.020)	(0.010)
ELA Grades 6-8		
0.001	0.015	-0.010
(0.006)	(0.014)	(0.007)

Table 8
Characterizing the Moves of Teachers Transferring Within New York City, First-Year Teachers in Grades 4-5

	Relatively more effective teachers	Middle group of teachers	Relatively less effective teachers
Change in the percent of students in the school who are minority	0.0	-6.3	7.4
Change in the percent of students in school failing the 4th grade math exam	-1.5	-1.6	3.0
Effectiveness in 2nd year for those moving compared to nonmoving peers in new schools	0.081**	-0.030	-0.111*
	(0.034)	(0.026)	(0.037)
Effectiveness in 1st year for those who moved compared to nonmoving peers in first schools	0.174*	0.001	-0.180
	(0.025)	(0.018)	(0.026)

Figure 1

Concentration of New York State Students Failing the 8th Grade Math Exam

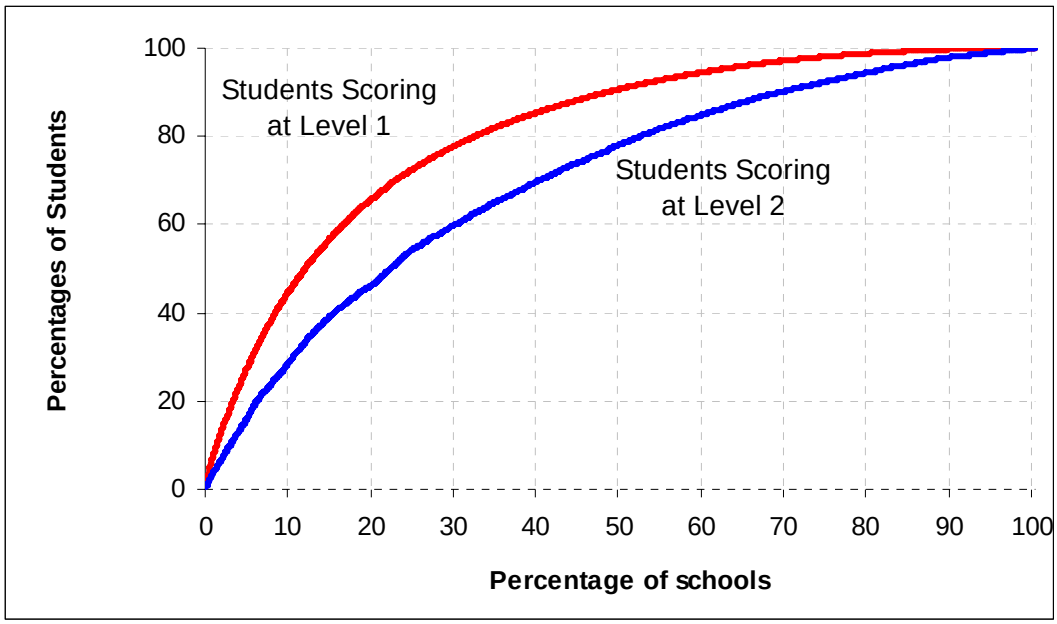

Figure 2a
Estimated Transition Probabilities for a Typical Lower Performing School by Teachers' Effectiveness in Math, Grades 4-5

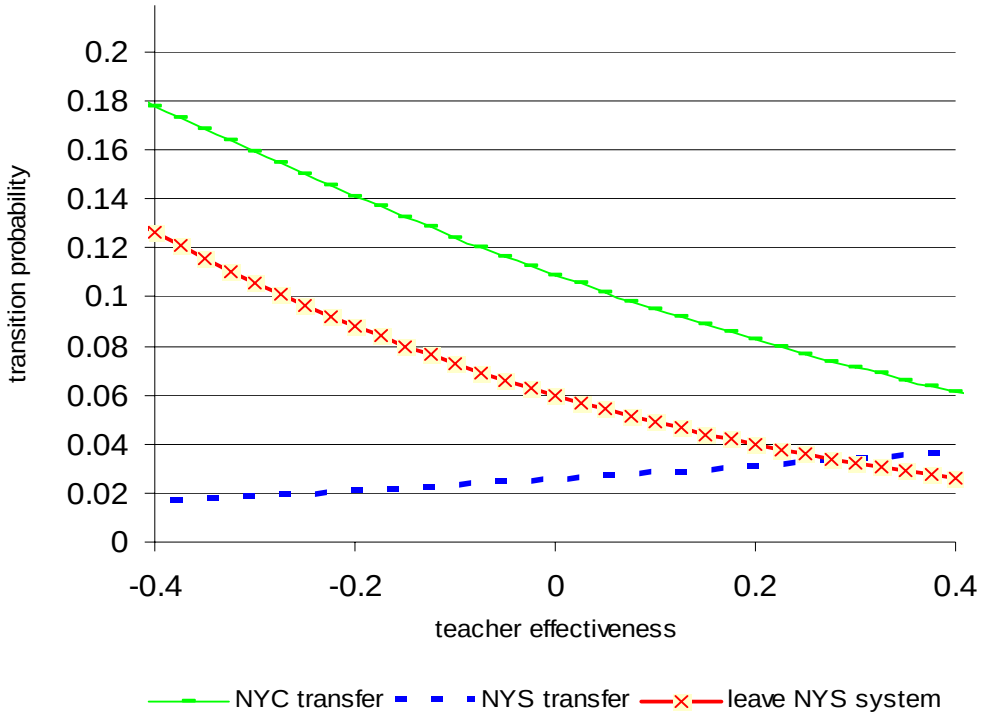

Figure 2b
Estimated Transition Probabilities for a Typical Higher Performing School by Teachers' Effectiveness in Math, Grades 4-

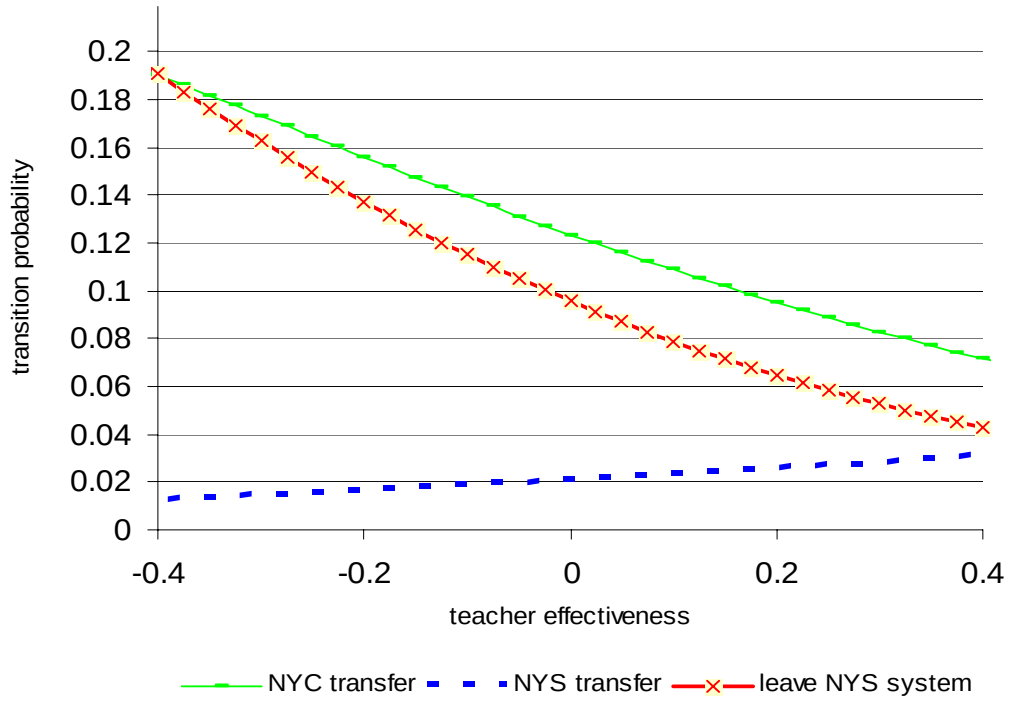

Figure 3
Distributions of 4th and 5th Grade Teacher Effectiveness in Math by Retention Status, Schools in Lower-Scoring Quartile

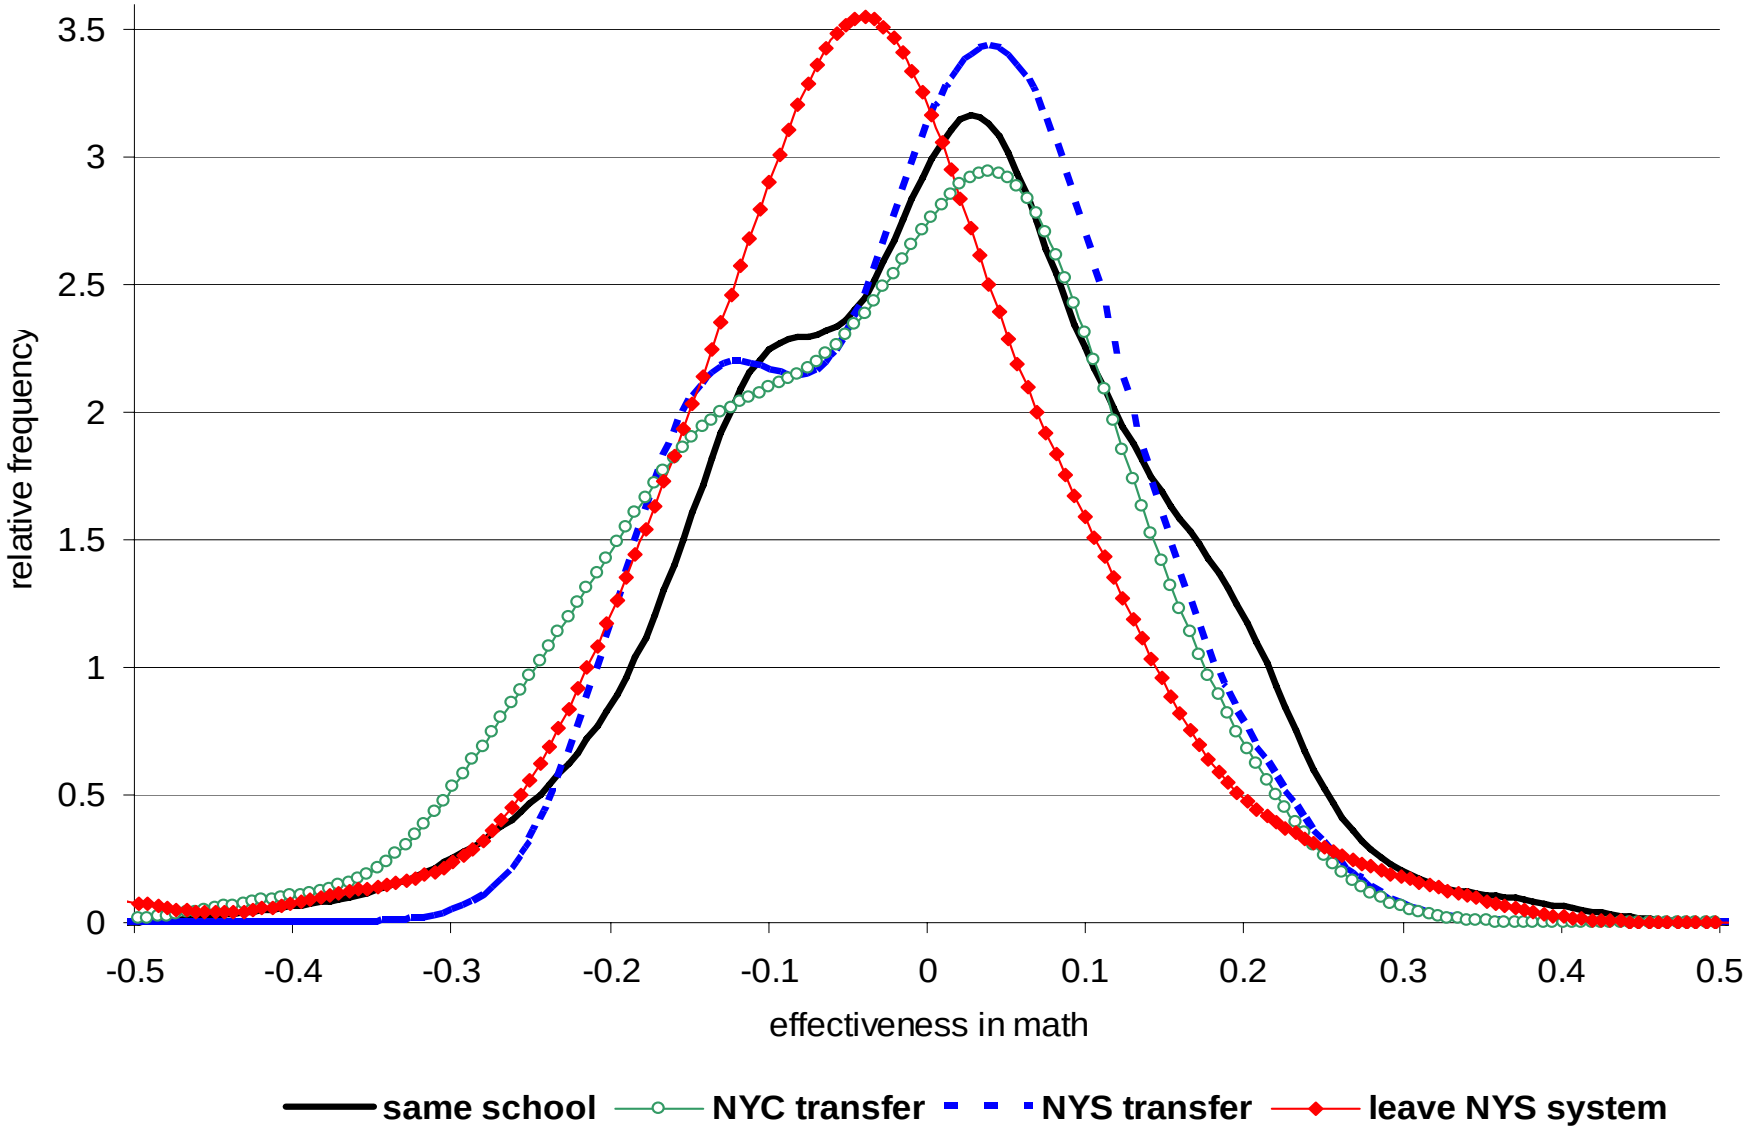

References

Boyd, D., H. Lankford, S. Loeb, J. Rockoff and J. Wyckoff (2006) "The Narrowing Gap in New York City Teacher Qualifications and Its Implications for Student Achievement in High-Poverty Schools," working paper.

Boyd, D., H. Lankford, S. Loeb and J. Wyckoff (2005) "Explaining the Short Careers of High-Achieving Teachers in Schools with Low-Performing Students," American Economic Review Proceedings 95(2), 166-171. An expanded version of this paper is available at www.teacherpolicyresearch.org .)

Boyd, D., H. Lankford, S. Loeb and J. Wyckoff (forthcoming) "Do High-Stakes Tests Affect Teachers' Exit and Transfer Decisions? The Case of the 4th Grade Test in New York State," forthcoming in Public Finance Review

Falch, T. and M Ronning (2005) "The Influence of Student Achievement on Teacher Turnover," CES-IFO Working Paper No. 1469.

Hanushek, E., J. Kain, D. O'Brien and S. Rivkin (2005) "The Market for Teacher Quality," NBER working paper.

Hanushek, E., J. Kain and S. Rivkin (2004) "Why Public Schools Lose Teachers," Journal of Human Resources 39(2) 326-354.

Ingersoll, R. (2001) "Teacher Turnover and Teacher Shortages: An Organizational Analysis," American Educational Research Journal 38(3) 499-534.

Johnson, S (2004) "Finders and Keepers: Helping New Teachers Survive and Thrive in our Schools, The Project on The Next Generation of Teachers, San Francisco: Jossey-Bass.

Ingersoll, R. (2004) Why do High Poverty Schools Have Difficulty Staffing Their Classrooms with Qualified Teachers? Report for Renewing Our Schools, Securing Our Future.

Ingersoll , R. and Kralik J.M. (2004). The Impact of Mentoring on Teacher Retention: What the Research Says. *Research Review*. Education Commission of the States.

Ingersoll, R. and T. Smith (2003). "The Wrong Solution to the Teacher Shortage." *Educational Leadership*, 60(8): 30-33.

Ingersoll, R. and T. Smith (2004) What are the Effects of Mentoring and Induction on Beginning Teacher Turnover, *American Education Research Journal* 41(3) 681-714.

Jacobs, Brian and Lars Lefgren (2005) "Principals as Agents: Subjective Performance Measurement in Education", working paper.

Kane, T., J. Rockoff and D. Staiger (2006) "What Does Certification Tell Us about Teacher Effectiveness? Evidence from New York City" NBER Workingn Paper 12155, April 2006.

Kane, T., and D. Staiger (2002) "Volatility in School Test Scores: Implications for Test-Based Accountability Systems," in Ravitch, D. (ed.) Brookings Papers on Education Policy 2002, Brookings Institution, Washington, D.C.

Loeb, S., L. Darling-Hammond, and J. Luczak (2005) “How Teaching Conditions Predict Teacher Turnover in California Schools,” *Peabody Journal of Education*, 80(3) pp. 44-70

Morris, C. (1983) “Parametric Empirical Bayes Inference: Theory and Application,” Journal of the American Statistical Association 78(381) 47-55.

Podgursky, M. R. Monroe and D. Watson (2004) “ The Academic Quality of Public School Teachers: An Analysis of Entry and Exit Behavior,” Economics of Education Review 23(5) 507-518.