



Dynamic effects of teacher turnover on the quality of instruction[☆]



Eric A. Hanushek^a, Steven G. Rivkin^{b,*}, Jeffrey C. Schiman^c

^a Hoover Institution/Stanford University, University of Texas at Dallas, and National Bureau of Economic Research, USA

^b Department of Economics/University of Illinois at Chicago, University of Texas at Dallas, and National Bureau of Economic Research, USA

^c Department of Economics/Georgia Southern University, USA

ARTICLE INFO

Article history:

Received 11 August 2015

Revised 1 August 2016

Accepted 3 August 2016

Available online 3 October 2016

Keywords:

Teacher turnover

Teacher quality

Estimation of teacher effects

ABSTRACT

It is widely believed that teacher turnover adversely affects the quality of instruction in urban schools serving predominantly disadvantaged children, and a growing body of research investigates various components of turnover effects. The evidence at first seems contradictory, as the quality of instruction appears to decline following turnover despite the fact that most work shows higher attrition for less effective teachers. This raises concerns that confounding factors bias estimates of transition differences in teacher effectiveness, the adverse effects of turnover or both. After taking more extensive steps to account for nonrandom sorting of students into classrooms and endogenous teacher exits and grade-switching, we replicate existing findings of adverse selection out of schools and negative effects of turnover in lower-achievement schools. But we find that these turnover effects can be fully accounted for by the resulting loss in experience and productivity loss following the reallocation of some incumbent teachers to different grades.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

There are conflicting views about the impact of teacher turnover on student outcomes. In policy discussions, the disruption from high turnover is viewed as a key impediment to high-quality instruction, and it is seen as doubly bad for high-poverty schools where seniority-determined job priority and the absence of substantial compensating differentials amplify the adverse impacts. These policy debates are often predicated on an assumption that schools

serving disadvantaged children tend to lose their better teachers year after year. However, recent research tends to find negative quality selection for departing teachers, contradicting this assumption and raising doubts about the prevailing negative view of turnover. Taking extensive steps to address problems introduced by non-random sorting of students into classrooms and endogenous teacher grade-switching and exits, this paper integrates analysis of aggregate disruptions from turnover with a parallel analysis about the effectiveness of school leavers in order to assess the net impact on the quality of instruction.

The high rate of turnover by teachers does not distinguish teaching from other occupations, because exit rates from teaching mirror those in non-teaching occupations (Ballou & Podgursky, 2002; Stinebrickner, 2002). Consequently, disproportionate harm from the turnover of teachers must come from the character of that turnover on education production, rather than simply the level. Indeed, following the general conclusion from other industries where early career moves are viewed as improving job

[☆] We thank Dan O'Brien for help with the data development and early analytical work and anonymous referees and the editor for thoughtful comments. Greg Branch provided superb research assistance. This research has received support from the Spencer Foundation, the Hewlett Foundation, and the Packard Humanities Institute. 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.

* Corresponding author.

E-mail address: sgrivkin@uic.edu (S.G. Rivkin).

matches, this larger perspective even raises the possibility that concerns about teacher turnover may be misplaced and that turnover may have little adverse effect on the quality of instruction.¹

The unique structure of teacher labor markets with barriers to entry, seniority-influenced teacher assignments, tenure restrictions, and rigid salary schedules is, however, often identified as fostering turnover-related productivity problems, particularly in high-poverty schools. [Dolton and van der Klaauw \(1999\)](#) find that teachers with better alternative earnings opportunities are more likely to exit teaching. The findings in [Chingos and West \(2012\)](#) reinforce this, showing that the relationship between estimated teacher effectiveness and earnings is stronger outside of teaching than in teaching for groups of teachers who leave for other industries.

In the case of schools, discussions frequently associate high turnover with lower average teacher quality, reduced school-specific human capital, disrupted school programs, and lessened teacher collaboration. Although unable to distinguish the mechanisms definitively, [Ronfeldt, Loeb, and Wyckoff \(2013\)](#) find high aggregate turnover adversely affects the quality of instruction in the subsequent year, particularly in lower-achievement schools. Moreover, they present evidence that turnover harms students even in classes with teachers who remain in the school.

Balanced against these aggregate findings are results suggesting that the composition of departing teachers may positively affect school productivity. With the recent expansion in the ability to judge teacher effectiveness has come evidence on productivity differences between teachers who transition out of a school and those who remain. The pattern of adverse selection out of teaching found in preliminary work on Texas ([Hanushek, Kain, O'Brien, & Rivkin, 2005](#)) was reinforced by subsequent work providing similar evidence for North Carolina and Florida ([Goldhaber, Gross, & Player, 2011; Feng & Sass, 2011](#)), though [Feng and Sass \(2011\)](#) also find higher exit rates for teachers from the upper tail of the distribution.

Nevertheless, identifying the net effects of teacher turnover is difficult. First, unobserved shocks may affect both the probability of teacher transitions and the quality of instruction. Second, schools may fill vacancies with a teacher new to the school or with a teacher from another grade, and any association between the fraction of teachers who switched grades and the fraction who are new to the school would likely introduce omitted variables bias given the evidence that switching grades adversely affects the quality of instruction ([Ost, 2014](#)). Finally, the decision to fill a vacancy with a teacher from another grade rather than a teacher new to the school is unlikely to be random. Rather the desirability of working with other teachers in a specific grade, perceptions of the incoming cohort of students, and principal discretion likely influence the allocation of incumbent teachers among grades.

A number of steps can be taken to address these concerns, though some involve analytical tradeoffs. First consider the level of aggregation. The measurement of turnover at the school-grade-year level enables the inclusion of school-by-year and school-by-grade fixed effects to account for time-varying school-wide shocks and persistent grade differences that may be associated with turnover. This approach, adopted by [Ronfeldt et al. \(2013\)](#), uses within-school variation in turnover among grades and years to identify the turnover effect. Although appealing, a focus on school-grade-year variation potentially amplifies any biases resulting from purposeful grade assignments of new and incumbent teachers in a given year. Specifically, if vacancies in “more desirable” grades are more likely to be filled with teachers from another grade while vacancies in “less desirable” grades are more likely to be filled by new teachers to the school, the fraction of teachers new to the school will be correlated with unobserved determinants of the work environment and achievement. The inclusion of the fraction of teachers who switched into the grade as an additional variable, a step not taken in [Ronfeldt et al. \(2013\)](#), can mitigate but not eliminate the biases introduced by such purposeful assignments. Importantly, neither the inclusion of both school-by-year and school-by-grade fixed effects nor the measurement of turnover as the fraction in a grade who exit the school rather than the fraction who are new arrivals fully resolves these issues.

Alternatively, turnover can be aggregated to the school-year level, addressing problems introduced by the non-random sorting of teachers among grades. In this case year-to-year variation in the share of all teachers at a school who are new identifies the total turnover effect, and the fraction of all teachers who switch grades can be included as a control. The primary drawback of this more aggregate measure is the deficiency of controls for time-varying, school wide shocks, as school but not school-by-year fixed effects can be included. Nonetheless, the investigation of alternative specifications provides information on the magnitudes of potential biases.

An additional benefit of the more aggregate measure is that it captures negative effects of turnover on the quality of instruction throughout the school and not just in the grades that experience transitions. [Ronfeldt et al. \(2013\)](#) emphasize the possibility that transitions may lessen the benefits of teacher cooperation within grades, but transitions in one grade could affect other grades, even those without staffing changes, through harm to coordination of curriculum and instruction across grades.

Recent research also opens more general questions about estimates of the quality distribution of teacher transitions. [Guarino, Reckase, and Wooldridge \(2015\)](#), [Jacob, Lefgren, and Sims \(2010\)](#), [Meghir and Rivkin \(2011\)](#), [Rothstein \(2010\)](#), and [Dieterle, Guarino, Reckase, and Wooldridge \(2015\)](#) focus both on potential biases introduced by the purposeful sorting of students and teachers into schools and classrooms and on the dynamics of teacher effects. [Kane and Staiger \(2008\)](#), [Kane, McCaffrey, Miller, and Staiger \(2013\)](#), and [Chetty, Friedman, and Rockoff \(2014b\)](#) provide some evidence in support of standard value-added models, but overall questions remain about the influence of unobserved heterogeneity and of

¹ An alternative perspective, not analyzed here, is that turnover in other settings may not be so beneficial. There may be greater adverse effects that are concealed by the absence of accurate labor productivity measures.

purposeful classroom assignment and must be considered in the estimation of the quality distributions of movers (Rothstein, 2014).

This paper examines the combined effects of overall turnover and the quality distribution of teacher transitions for a large, urban district in Texas with special emphasis on nonrandom sorting of students into classrooms, endogenous teacher exits and grade-switching. The analysis underscores the importance of a broader consideration of teacher turnover and the measurement of teacher effectiveness. Results are consistent with existing findings of adverse selection out of schools and negative effects of turnover in lower-achievement schools. However, we find that the decline in experience with teacher transitions and productivity loss following the reallocation of some incumbent teachers to different grades account fully for the negative turnover effect once we also consider the nonrandom assignment of teachers to grades.

The next section describes the administrative data. Section 3 develops the empirical model for the analysis of the quality distribution of transitions and presents results, highlighting the additional complications introduced by classroom selection and nonrandom sorting of students and teachers into schools and classrooms. Section 4 modifies the model to examine the effects of teacher turnover and describes the methods used to address endogenous teacher exits and movements across grades. Throughout, the analysis considers heterogeneity by school average achievement, both because policy is concerned with the impacts of turnover in low-performing schools and because achievement levels appear to be related to working conditions and teacher transitions (Hanushek, Kain, & Rivkin, 2004). Section 5 synthesizes the aggregate and compositional components in order to assess the various channels of turnover effects. Finally, Section 6 discusses implications for public policy given the institutional structures within which public schools currently operate.

2. Texas Schools Project data

The stacked panel data sets constructed by the Texas Schools Project contain administrative records on students and teachers collected by the Texas Education Agency (TEA) beginning in 1989.² These data permit the linkage of students over time and of students and teachers in the same school, grade, and year, but the statewide data have historically not matched students and classroom teachers. This shortcoming is overcome by merging in teacher-student matched data provided by one large Texas urban district, known henceforth as the “Lone Star” District. Typically, this match identifies a subject specialist in middle school and a general teacher in elementary school. Only regular classroom teachers are considered in the analysis.

The student background data contain a number of student, family, and program characteristics including race, ethnicity, gender, and eligibility for a free- or reduced-price

Table 1

Average annual transition shares for Lone Star district teachers, 1997–2001.

	All (1)	New teachers (2)
Stay at same campus	0.806	0.689
Change campus	0.093	0.113
Change district	0.042	0.074
Exit Texas public schools	0.060	0.124
Sample Size	6241	646

Notes. Turnover by type of transition is averaged across years. New teachers (col. 2), have 0 years of past experience.

lunch (the measure of economic disadvantage), classification as special needs, and classification as limited English proficient. During the study period, students were annually tested in a number of subjects using the Texas Assessment of Academic Skills (TAAS), which was administered each spring to eligible students enrolled in grades three through eight. These criterion-referenced tests evaluate student mastery of grade-specific subject matter, and this paper presents results for mathematics.³ Test scores are converted to z-scores using the mean and standard deviation for the entire state separately for each grade and year to account for effects of test score inflation and other changes to the tests.

In this paper we study students and teachers in grades 4 through 8 for the school years 1996–1997 to 2000–2001, the period with data for the Lone Star district. We eliminate any student without valid test scores or other missing data and classrooms with fewer than five students with non-missing data.

3. Teacher transitions and productivity

Teacher movement is substantial within the Lone Star district. Table 1 shows that among teachers with 0–1 years of prior experience less than 70 percent on average remain at their campus for a second year: 12 percent exit the Texas public schools entirely, 11 percent change campuses within the Lone Star District, and another 7 percent move to another Texas district each year. Even among more experienced teachers, some 19 percent leave their current school each year.

Although considerable attention has been given to the amount of teacher turnover in U.S. schools, the policy implications depend crucially on how turnover affects the distribution of teacher quality. In particular, longer-term changes in the quality distribution may exert greater effects over time than the shorter-term disruptions resulting from teacher transitions. Understanding the character and implications of job mobility naturally begins with the development of measures of teacher quality.

3.1. Key issues in estimating teacher productivity

Estimation of teacher effectiveness has been controversial in part because of its potential use in personnel

² The underlying data were developed at the Texas Schools Project at the University of Texas at Dallas. For more detail, see <http://www.utdallas.edu/research/tsp-erc/>.

³ There is some evidence that teachers differ in their value-added across subjects, something that we do not consider here; see Goldhaber, Cowan, and Walch (2013).

decisions, and this has spurred a vast amount of research. The primary analytical task in estimating effectiveness is separating teacher contributions to achievement from other student, family, school, and community factors, and here there is an ongoing debate over the best approach.⁴ It is important to locate our investigation of teacher mobility in the existing research on teacher effectiveness – a research area that has simultaneously gone in several different directions.

One issue in the debates has been whether lagged student achievement or alternatively student fixed effects in value-added models are the best way to account for heterogeneity in non-teacher factors. Historically, some form of lagged achievement has been the most common approach to estimating value-added, but student fixed effects – particularly with new data and larger computing power – provide an appealing alternative to account for unobserved heterogeneity that is time invariant. [Sass, Semykina, and Harris \(2014\)](#) offer evidence that supports student fixed effects as the preferred model. But, other analyses that examine the sensitivity of a number of value-added specifications to various shocks and confounding factors (including some not tested directly by [Sass et al. \(2014\)](#)) raise concerns about this conclusion. Using Monte Carlo analysis, [Guarino et al. \(2015\)](#) examine the sensitivity of a set of models to various shocks and confounding influences. Their results indicate that the lagged achievement specification is quite robust to commonly considered shocks and outperforms models that might have more conceptual appeal including those with student fixed effects. Our initial analysis of Lone Star district data showed that the inclusion of student fixed effects more than doubled estimates of the variance in teacher effectiveness, suggesting the presence of strong inverse matching of weaker students and stronger teachers, substantial error in measurement exacerbated by the fixed effects, or some combination of the two.

The lack of consensus regarding model specification and the sensitivity of estimates to specification differences do raise questions about the interpretation of existing evidence about teacher effectiveness and teacher mobility. [Feng and Sass \(2011\)](#) estimate a number of alternative specifications that differ in the treatment of depreciation and in the approach for accounting for student heterogeneity. The use of specifications with either student fixed effects with no lagged achievement variable or the achievement gain as the dependent variable lack conceptual appeal given the evidence on knowledge persistence and the relatively poorer performance of such models in the presence of common shocks ([Meghir & Rivkin, 2011](#)). Importantly, their estimates from the arguably more appealing models that control for unobserved heterogeneity with prior achievement are quite sensitive to whether

teacher quality is measured at a single point in time or using all years of data for a teacher regardless of whether they follow or precede a school transition. In fact, it is only the estimates based on all years of data that show the bimodal distribution of the effectiveness of exiting teachers emphasized by the authors.

Though [Goldhaber et al. \(2011\)](#) undertakes more extensive efforts to examine the potential effects of non-random student sorting into classrooms and unobserved school influences, differences across the inclusion of student fixed effects in virtually all value-added specifications and sensitivity of the estimates to the treatment of unobserved student and school heterogeneity again raises concerns. In addition, the use of a continuous variable to account for experience is not consistent with the non-linear pattern of diminishing returns generally found in previous studies ([Hanushek & Rivkin, 2012](#)), and the inclusion of average mathematics achievement as a control raises questions given that it is determined partly by teacher productivity.

A second major issue arises from the endogenous allocation of teachers among schools and grades. Because the school and classroom distribution results from the purposeful choices of teachers and administrators based in part on working conditions, it is imperative to account for school factors that affect both learning and teacher transitions. [Feng and Sass \(2011\)](#) rely exclusively on a limited set of school characteristics to account for both student sorting among schools and the effects of other school factors on achievement. [Goldhaber et al. \(2011\)](#) do include school fixed effects in some specifications. Nevertheless, the majority of these specifications, including most of the robustness analysis, does not include school fixed effects, and those that do will not capture any systematic differences by grade in a school or district that may be related to both achievement and teacher turnover.

Based on the balance of the evidence and initial work, we specify a cubic lagged achievement value-added regression that estimates average quality differences by transition status.⁵ In order to account for time- and grade varying differences among schools in cohort quality, working conditions, curricula and other factors, we also include school-by-grade-by-year fixed effects. These account comprehensively for all differences among schools, grades, and years factors that affect achievement and transitions at the cost of neglecting any true differences in teacher effectiveness across the same dimensions.

Nevertheless, nonrandom sorting of students among classrooms could still bias estimates of teacher effectiveness. [Clotfelter, Ladd, and Vigdor \(2006\)](#) and [Rothstein \(2010\)](#) document non-randomness across classrooms on the basis of student characteristics and prior performance in North Carolina, but whether such sorting introduces substantial bias in commonly used teacher value-added models is a topic of considerable debate. In an influential article, [Rothstein \(2010\)](#) argues that much of the sorting occurs on the basis of time-varying student heterogeneity, and therefore lagged test scores and even student fixed

⁴ For earlier estimates of teacher value-added, see, for example, [Hanushek \(1971, 1992\)](#), [Armor et al. \(1976\)](#), [Murnane \(1975\)](#), [Murnane and Phillips \(1981\)](#), [Aaronson, Barrow, and Sander \(2007\)](#), [Rockoff \(2004\)](#), [Rivkin, Hanushek, and Kain \(2005\)](#), [Boyd, Grossman, Lankford, Loeb, and Wyckoff \(2006\)](#), [Kane, Rockoff, and Staiger \(2008\)](#) and [Chetty, Friedman, and Rockoff \(2014a, 2014b\)](#). These and other studies are reviewed and critiqued in [Hanushek and Rivkin \(2010, 2012\)](#).

⁵ See [Hanushek \(1979, 1986\)](#) for a discussion of value-added models.

effects may fail to capture important determinants of classroom allocation.⁶

A series of papers investigates the magnitude of any such bias. Kane and Staiger (2008) develop a specification test of the validity of non-experimental estimates for a small sample of Los Angeles teachers and cannot reject unbiasedness of various standard estimators. A follow-on study expands the sample and finds similar results, thus providing additional evidence against the importance of such sorting effect (Kane et al., 2013). Finally, the specification checks in Chetty, Friedman, and Rockoff (2014a) provide additional evidence against the presence of substantial bias in value-added models with cubic polynomials in prior achievement caused by extensive sorting on unobserved characteristics. Nonetheless, the exchanges between Rothstein (2014) and Chetty, Friedman, and Rockoff (2016) illustrate that this debate remains unresolved.⁷

Consequently, we take additional steps to understand the likely influences of any sorting among classrooms conditional on prior achievement and the other included variables. We do this by dividing school, grade, and year cells into “sorted” and “not-sorted” categories on the basis of two different dimensions of the classroom allocation process. The first approach follows in the spirit of Clotfelter et al. (2006) and is based on an F-test of the equality of mean prior year test score; the second approach uses a chi-square test to examine the independence of transitions of students who remain in the same school from grade $G-1$ to grade G .⁸ The school observations where we reject the null hypothesis of no significant differences in mean prior test score among classrooms (first approach) or the null hypothesis that the allocation of students across classrooms in grade g is independent of the allocation in grade $g-1$ based on a chi-squared test of the transition matrix (second approach) are considered observations potentially affected by purposeful sorting (“sorted”). The remaining schools are classified as “not-sorted”.⁹ We then estimate the transition regressions separately for the two sets of “sorted” and “not-sorted” samples and compare the results to estimates based on the entire sample.¹⁰

⁶ Some specification issues in Rothstein (2010) do raise questions about the strength of this critique. The evidence that time varying classroom heterogeneity is important is based on findings from models with student fixed effects and test score gain as the dependent variable. If the assumption of no knowledge depreciation is incorrect, it could appear that much student heterogeneity was time varying even if that were not the case.

⁷ Within the class of lagged achievement models, there does appear to be a growing body of evidence that unobserved influences does not introduce substantial bias into estimates produced by models that include cubic polynomials of prior achievement. See the empirical evidence provided by Kane and Staiger (2008), Koedel and Betts (2011), Kane et al. (2013), and Chetty, Friedman, and Rockoff (2014a).

⁸ Dieterle et al. (2015) develop a similar approach to that used in Clotfelter, Ladd, and Vigdor (2006) in which they examine the characteristics jointly.

⁹ Appendix Table A1 describes the joint distribution of the distribution of schools into “sorted” and “not-sorted” categories on the basis of these two criteria and illustrates that there is substantial but incomplete agreement in the divisions. Approximately three-quarters of the common schools fall into the same category of sorted or not-sorted, but student sample losses from the tests of placements across years are large.

¹⁰ Goldhaber, Gross, and Player (2011) follow a variant of the first approach, though they classify schools on the basis of sorting of demographic characteristics rather than prior achievement.

Note these tests are weak in the sense that, given the sample sizes used in the analysis, the failure to reject the hypotheses of independence or no significant difference in mean prior test score at the five percent level does not provide strong evidence that a school randomly assigns students among classrooms. Nonetheless, preliminary work (not reported) shows that teachers in the “not-sorted” samples appear to have little or no effect on the learning of students in the prior grade but substantial variation in the estimated effect on current achievement. This supports the belief that transition estimates for the “not-sorted” samples provide useful robustness checks.

3.2. Empirical model

Eq. (1) models achievement of student i in school s in with teacher j in grade G and year y as a cubic function of lagged mathematics achievement ($f(A)$), family background and other influences outside of schools (X), peer composition (P), school factors (S), a set of teacher transition indicators (T), a school-by-grade-by-year effect (η) and a random error (e):

$$A_{isGjy} = \theta f(A_{iG-1}) + \beta X_{iGy} + \lambda P_{isGy} + \delta S_{isGy} + \sum_k \kappa_k T_{kjsGy} + \eta_{sGy} + e_{isGjy} \quad (1)$$

where $T_{kjsGy} = 1$ if teacher j undergoes transition type k at the end of year y and $=0$ otherwise.¹¹ We focus on three types of moves: (1) move to a different school in the same district; (2) move to a school in a different district; and (3) exit from public schools in Texas. The comparison group is the set of teachers who remain in the same school for the following year. κ_k is the conditional mean effectiveness of teachers in each transition type compared to those who remain in the school. Individual controls (X) include indicators for female, race-ethnicity, low income, special needs, limited English proficient, first year in middle school, and family initiated move. Peer composition (P) likely also affects achievement, and the school-grade-year effects account for all peer differences by school, grade, and year. School controls (S) include a full set of teacher experience dummy variables and a full set of year-by-grade effects. The estimation presumes that there are at least two observations of achievement for each student and that there are multiple students with each teacher.

3.3. Effectiveness of departing teachers

The evidence of sizeable variation in teacher effectiveness motivates the attention to teacher transitions and how transitions alter the distribution of teacher quality. Given that teachers initiate a majority of job separations, teacher movement would arguably improve the well-being of most school leavers, even those who choose to leave the profession entirely, but the impact on students is less clear.

¹¹ A much more flexible specification of prior achievement that included indicators for each decile was used in preliminary work on teacher transitions, and the results were virtually identical to the parsimonious specification used in this paper.

Table 2

Differences in average teacher quality by transition compared to teachers who do not move, for all schools and by the allocation of students across classrooms.

	All schools (1)	Sorting tested by pretest mean achievement		Sorting tested by classroom assignment patterns	
		Sorted (2)	Not-sorted (3)	Sorted (4)	Not-sorted (5)
Change campus	–0.024* (0.013)	–0.029* (0.016)	–0.010 (0.021)	–0.030 (0.023)	–0.056* (0.029)
Change district	–0.008 (0.017)	–0.011 (0.020)	0.013 (0.034)	0.015 (0.029)	–0.056 (0.039)
Exit Texas public schools	–0.056** (0.013)	–0.041** (0.016)	–0.085*** (0.022)	–0.046* (0.020)	–0.121*** (0.037)
Sample Size	205,711	127,711	64,383	61,713	27,814

Notes. All regressions include school-grade-year fixed effects. Coefficients on teacher transition variables come from regressions of math score on the transition variables plus lagged test score, indicators for female, race-ethnicity, low income, special needs, limited English proficient, first year in middle school, family initiated move, shares of students in campus, grade, and year who are female, black, Hispanic, Asian, Native American, low income, special needs, limited English proficient, movers, peer average lagged achievement, a full set of teacher experience dummies, and a full set of year-by-grade dummies. No move is the omitted category. Sorted and not-sorted schools are based on statistical tests related to the entering achievement patterns; see text. Standard errors clustered by teacher-year are in parentheses.

* $p < 0.10$,** $p < 0.05$,*** $p < 0.001$.

The constrained labor markets for teachers – with strict district salary schedules that vary only modestly across districts – may lead the dynamics of the teacher labor market to diverge sharply from those of less fettered markets. The fact that much of the movement involves changes across the “establishments” of a single firm (district) in a context in which teachers typically maintain significant control over assignment to open positions introduces another dimension through which the choice process can lead to substantial inequality in teacher effectiveness among both districts and schools.

In the empirical analysis, we use a number of variants of the basic models to describe differences in teacher effectiveness by transition status. The specifications differ by the steps taken to account for potential confounding factors and time- and grade-varying school shocks and the timing of the measurement of teacher effectiveness. Finally, the pattern of teacher transitions is permitted to differ by school average student achievement, a factor previously identified as related both to the likelihood of teacher exits and the degree of harm caused by teacher departures (Hanushek et al., 2004) and also to the magnitude of aggregate turnover effects (Ronfeldt et al., 2013).

The baseline estimates of mean differences in value-added to mathematics achievement by transition type, contrary to much popular discussion, provide no evidence that more effective teachers have higher probabilities of exiting schools in the Lone Star District. As shown in the first column of Table 2, regardless of the destination of the departing teachers, all coefficients are negative, implying that the average leaver is less effective than the average stayer in each school. Models without school-by-grade-by-year fixed effects (not reported) produce the same pattern of estimates.

Those who exit the Texas public schools entirely are significantly less effective on average than those who stay. In the school year immediately prior to exiting, the average value-added of a teacher who left the Texas public schools entirely was 0.056 standard deviations (of student achieve-

ment) below the average for a teacher remaining in the same school.¹² Note that estimates for Texas reported in Rivkin, Hanushek, and Kain (2005) show that a standard deviation of teacher effectiveness in terms of the student achievement distribution is 0.11 s.d. Therefore, on average those exiting from Texas schools are roughly 50 percent of a teacher standard deviation less effective than their colleagues who remain in the same school. Moreover, those who switch campuses within the same district are also significantly less effective than stayers, though the deficit is smaller than that observed for those exiting the Texas public schools. In contrast, those switching to another Texas school district are not significantly different on average from teachers who remain in the same school.

Because purposeful classroom assignments on the basis of unobserved factors potentially contaminate the estimates, the results in the center and right panels of Table 2 include estimates from the separate samples of “sorted” and “not-sorted” observations. These columns provide little or no evidence that such classroom sorting drives the results. Rather the estimated differences between stayers and teachers switching to a different campus or exiting from Texas public schools are mostly larger in magnitude and more significant in the “not-sorted” samples, where any biases are almost certainly smaller. Exiting teachers in the “not-sorted” samples have value-added estimates ranging from 0.085 to 0.12 standard deviations (of student achievement) below teachers who stay. These estimates are statistically significant regardless of the method used to divide the schools.

The mean differences offer a limited view of the character of transitions, because there is substantial quality variation within each of the streams. Figs. 1 and 2 provide kernel density plots of teacher value-added that illustrate both the mean differences and dispersion of each of the streams

¹² This lower achievement of exiting teachers contrasts with the findings of Chingos and West (2011) who find both that moving into administration and exiting tend to be related to higher quality teachers in Florida.




Fig. 1. Distribution of teacher quality across schools.




Fig. 2. Distribution of teacher quality within school-grade-years.

in terms of teacher quality. Although non-persistent factors, including simple measurement error in the tests, certainly inflate the dispersion for all streams, the magnitude of the observed variation clearly indicates substantial productivity differences among stayers, school changers, and those who exit the public schools. Consistent with [Sass, Hannaway, Xu, Figlio, and Feng \(2012\)](#), there is some evidence that the relatively small number of district switch-

ers also contains a disproportionate number of the most effective teachers. However, differences between the quality distributions of stayers on the one hand and campus switchers and exiters from the Texas public schools on the other emerge across the entire distributions.

The high transition rates of teachers early in the career magnify the importance of the pattern of movement for this group, and the estimates in [Table 3](#) reveal a

Table 3
Differences in average teacher quality by transition type and experience at time of move.

	Teacher experience at time of move			
	One year (1)	Two years (2)	Three years (3)	Four of more years (4)
Change campus	0.091** (0.039)	0.011 (0.056)	−0.056 (0.038)	−0.042*** (0.014)
Change district	0.102** (0.043)	−0.078* (0.044)	−0.007 (0.043)	−0.024 (0.022)
Exit public schools	−0.041 (0.034)	0.000 (0.036)	−0.082** (0.039)	−0.062*** (0.016)

Notes: The table reports within school-grade-year comparisons. No move is the omitted category. Models include the same controls as Table 2 except here they include interactions between one, two, three, and four+ years of experience and the transition variables. The regression is based on the same sample of 205,711 observations. Standard errors clustered by teacher-year are in parentheses.

* $p < 0.10$,

** $p < 0.05$,

*** $p < 0.001$.

sharp divergence between patterns for first year and those for more experienced teachers. Specifically, the first year teachers who change schools or districts are significantly more effective on average than stayers; this finding does not hold for any other level of experience. In contrast, first-year teachers who exit the Texas public schools are less effective on average, though the estimate is not significant.

The concentration of newer teachers in higher poverty and lower achievement schools raises the possibility that the quality distribution of transitions differs by school characteristics. Lower average achievement is arguably the best available measure of academic disadvantage, as prior research highlights the pronounced movement of more experienced teachers toward higher achieving schools (Hanushek et al., 2004).¹³

To examine transition outcomes by school average achievement, we interact the transition indicators with indicators for high-achievement and low-achievement schools (i.e., for schools with mean achievement above or below the median for schools in the district). The estimates in the first column of Table 4 provide little support for the view that schools with lower achievement are more likely to lose their higher-performing teachers. To the contrary, the estimates indicate that teachers who depart lower-achievement schools regardless of destination are even less effective on average relative to stayers.

Column 2 of Table 4 reports average effectiveness for departing first year teachers, and the estimates show that positive selection of campus and district switchers holds regardless of school average achievement. (Estimates for second, third, and fourth year teachers show little pattern and are rarely significant). The value-added differential be-

tween first-year campus switchers and all stayers equals 0.22 for higher-achievement schools and 0.04 for lower-achievement schools; a smaller differential by school average achievement emerges for the positive selection of district switchers. In contrast, significant negative selection for those who leave the Texas public schools appears only for the lower-achievement schools; the corresponding estimate for first-year teachers exiting from higher-achievement schools is small and not significant.

A related question considers the destination of within district campus switchers: are more effective teachers more likely to gravitate toward higher achieving schools? Specifications that classify transitions by both origin and destination school characteristics (not shown) reveal little or no evidence of significant differences by destination school type regardless of the student characteristics in the origin school.

The basic results about the effectiveness of movers paint a clear picture that teachers leaving lower-achievement schools are less effective than the average teacher who does not move. This holds for new teachers as well as experienced teachers, though the effectiveness of teachers who switch campuses and districts relative to those with comparable levels of experience appears to be much higher for first year teachers. Finally, teachers who leave the Texas public schools tend to be noticeably less effective than the average teacher who does not move, and this holds regardless of experience.

The timing of the measurement of effectiveness is the final element of the estimation that we examine. To this point estimates of teacher effectiveness are based on the academic year immediately prior to any transition, but this chronology potentially complicates interpretation of the results. There are several reasons why movers may appear less effective in their transition year. One possibility is that movers may put forth less effort once they decide to leave the school. A second possibility is that a negative shock such as an unruly class or a bad relationship with a principal may simultaneously induce a transition and degrade instructional effectiveness.

To isolate persistent productivity differences, we generate estimates based on value-added in the year prior to the transition year and compare those with estimates based on the transition year (immediately preceding the move). For example, we describe the distribution of quality for transitions following the 1999 school year with value-added based on average student achievement during the 1998 school year, implying that any shocks or changes in effort related to the transition do not affect the estimates of teacher effectiveness. (Note, however, that this approach does reduce the sample size by eliminating all teacher-year observations for which the teacher is not in the sample in the prior year).

The final two columns of Table 4 present estimates of the effectiveness of departing teachers based alternately on achievement in the transition year and the previous year (separately by school achievement category) using the subsample of teachers in the data in successive years.¹⁴

¹³ We repeated the same distributional analysis with schools divided by percentage of black students (not reported), and obtained qualitatively similar results. Information on eligibility for a subsidized lunch provides a crude measure of income, and our earlier work on teacher mobility suggested that black concentration and not overall minority concentration was most salient for teacher moves (Hanushek, Kain, & Rivkin, 2004).

¹⁴ Note that the point estimates in Column 6 differ somewhat from the point estimates in Column 1 based on the entire sample because the sub-

Table 4

Differences in average teacher quality by transition type and school average prior achievement.

	By teacher experience		Sorted (3)	Not sorted (4)	By timing of teacher effectiveness	
	All (1)	1st Year (2)			Prior year (5)	Current year (6)
<i>Change campus</i>						
From low achieving school	−0.050*** (0.018)	0.039 (0.047)	−0.057** (0.023)	−0.034 (0.026)	0.026 (0.023)	−0.063*** (0.023)
From high achieving school	0.006 (0.019)	0.218*** (0.062)	0.001 (0.022)	0.024 (0.035)	−0.018 (0.029)	−0.049** (0.021)
<i>Change district</i>						
From low achieving school	−0.004 (0.023)	0.067 (0.049)	−0.018 (0.026)	0.069 (0.047)	0.006 (0.037)	0.040 (0.035)
From high achieving school	−0.014 (0.026)	0.123* (0.064)	−0.004 (0.030)	−0.051 (0.048)	−0.034 (0.031)	−0.047 (0.035)
<i>Exit public schools</i>						
From low achieving school	−0.073*** (0.018)	−0.120*** (0.040)	−0.061*** (0.022)	−0.096*** (0.032)	−0.035 (0.024)	−0.018 (0.023)
From high achieving school	−0.040** (0.019)	0.030 (0.053)	−0.022 (0.024)	−0.076** (0.031)	−0.066*** (0.025)	−0.061** (0.024)

Notes. For each type of departure, schools are divided into low achieving or high achieving by being below or above median school achievement. The omitted transition category is no move. Each regression includes the same variables as in Table 2 specifications. Sample sizes are 205,711 in Columns 1–4, 125,843 in Column 5 and 158,343 in Column 6. Standard errors clustered by teacher-year are in parentheses.

* $p < 0.10$,** $p < 0.05$,*** $p < 0.001$.

Two findings stand out in the comparison of performance in the transition year and the year prior. First, the apparent negative selection of campus switchers based on the transition year measure does not emerge in the estimates based on the penultimate year regardless of school achievement category. It appears that these teachers experience a temporary effectiveness decline in their final year in the school. Second, the use of the penultimate year to measure effectiveness produces similar estimates of the negative selection of those exiting the Texas public schools for both achievement groups. Taken together the results suggest that the lower performance of those leaving the low-achievement public schools reflects actual skill differences, while interpretation of the lower effectiveness of campus switchers in the year of the move is less clear.

The absence of information on contract offers prevents us from determining whether departures from the public schools result from voluntary decisions on the part of poorly-performing teachers, district decisions not to renew contracts, or principal pressure to quit.¹⁵ These alternative channels carry different implications for policy, but within our data it is impossible to distinguish among them.

4. Aggregate disruption

Even in the presence of negative selection out of schools, teacher departures may adversely affect the quality of instruction through a number of channels. First, turnover may reduce the amount of accumulated general and specific human capital: in the Lone Star Dis-

trict, roughly one third of teachers new to a school have no prior teaching experience. Second, many new hires may come from the lower portion of the quality distribution. Third, turnover may lead to shuffling of teachers among grades, and Ost (2014) finds that movement into a grade not taught in the prior year tends to lower value-added. Fourth, the composition of peer teachers may affect productivity.¹⁶ Fifth, the disruption associated with turnover may reduce productive interactions among teachers.

4.1. Modeling impacts of teacher turnover

Identification of the impacts of teacher turnover is challenging because of unobserved shocks associated with exits, the purposeful movements of teachers among grades, and spillovers across grades. Previous work by Ronfeldt et al. (2013) measured turnover at the grade level and used school-by-year or school-by-grade fixed effects to account for unobserved influences. However, these fixed effects may fail to account for important confounding factors related to the endogenous sorting of teachers among grades. This leads us to take additional steps in an effort to isolate the causal effects of turnover.

First, we include the fraction of teachers who were in another grade in the same school in the prior year as an additional regressor.¹⁷ If the fraction of teachers in a grade who are new to the school or the fraction of teachers in

¹⁶ See for example Jackson (2012).

¹⁷ If turnover in one grade adversely affects the quality of instruction in other grades due to grade reassignments and hiring, the strict exogeneity assumption underlying in the school-by-year fixed effect regressions will be violated; i.e., the fixed effects will necessarily not be fixed. For discussion of the strict exogeneity requirement in these fixed effects models, see Wooldridge (2002).

sample required to analyze the effects of timing excludes all first-year teachers and those without two successive years of value-added results.

¹⁵ See the suggestive evidence in Branch, Hanushek, and Rivkin (2012) on principal decision making.

a grade who exited following the previous year is related to the fraction that moved from another grade in the same school, potentially important omitted variables bias is introduced. Appendix Table A2 shows that the correlations between the share that moved from another grade and the share that exited the school following the prior year equals 0.14, and the correlation between the share that moved from another grade and the share new to the school equals -0.17 . This comparison highlights the importance of controlling for grade switching.

Following Ost and Schiman (2015), Appendix Table A3 reports linear probability models of the determinants of grade switching. These reveal that grade switching is strongly related to the probability a teacher leaves any grade. Moreover, regardless of the structure of the fixed effects, the estimates show that more effective teachers in the prior year are far less likely to switch grades. Thus, neither the incidence nor the composition of grade switching is random, and the absence of controls for it likely compromises the estimation of the aggregate effects of teacher turnover.

Importantly, just controlling for the share of newly re-assigned teachers does not fully account for unobserved factors related to teacher transitions. Consider the possibility that a higher rate of transitions out of one grade relative to others may result from the presence of a problematic teacher or impending arrival of a disruptive student cohort. A principal who uses vacancies to reallocate teachers among grades based on cohort characteristics, teacher requests, or other factors might allocate new teachers to particularly difficult classrooms and grades. Given the absence of controls for climate and disruption, this practice would tend to amplify the magnitude of any effect of the share of new teachers estimated at the grade level. Of course other principal decision rules might introduce bias in the opposite direction.

Neither models with school-by-grade nor school-by-year fixed effects account for time-varying, grade specific factors that may introduce bias. In the case of school-by-grade fixed effects, the specification accounts for fixed differences among grades including a particularly strong or weak incumbent teacher but not time varying differences. In the case of school-by-year fixed effects, any grade specific unobservables related to turnover would be problematic. Moreover, school-by-year fixed effects absorb any negative effects on cooperation, planning and curriculum implementation that involves all grades.

Consequently, we not only estimate grade-level models with school-by-year and school-by-grade fixed effects but also an additional set of specifications that aggregate turnover and grade reassignments to the school-by-year level rather than school-by-grade-by-year level. This extension circumvents problems introduced by the purposeful assignment of teachers to grades based on time-varying factors and captures any school-wide disruption effects. Although aggregation to the school-by-year level reduces the variation in turnover, it eliminates biases resulting from the purposeful assignment of teachers across grades. Moreover, the approach also accounts for grade-specific factors that induce exits including cohort quality, preferences re-

garding peer teachers, or purposeful placement of teachers for reasons unrelated to their preferences. Comparisons of the magnitudes of estimates based on different levels of aggregation for models with school fixed effects provides information on the relative importance of spillovers versus specification error induced by teacher movements among grades.

Importantly, this aggregation precludes the inclusion of school-by-year effects to account for correlated time-varying factors. The sensitivity of coefficients in the grade-level models to the inclusion of school-by-year effects will provide some information on the importance of time-varying factors, though it will not be possible to determine whether any reduction in the magnitude of the turnover coefficient following the replacement of school with school-by-year fixed effects results from the elimination of confounding influences, from the absorption of school-wide turnover effects, or from some combination of the two. Nonetheless, assuming that school-wide spillover effects are unlikely to be large relative to those that are grade specific, evidence of much larger coefficients in the absence of school-by-year fixed effects would raise concerns about any specification that does not include school-by-year fixed effects.

A final issue concerns the timing of the measurement of turnover. Ronfeldt et al. (2013) find very similar results when turnover is measured as the share of departing teachers from the prior year or the share of new teachers in the current year. Nonetheless, the timing determines the character of potential confounding influences. On the one hand, the transition rate out of a grade may be related to grade-specific experiences during that year or expectations for that grade in the subsequent year but not to any movement of teachers in response to vacancies created by the departures. On the other hand, the proportion of new teachers in the current year has a weaker relationship with grade-specific prior year factors and shocks but potentially a stronger relationship with influences related to purposeful teacher movements among grades. Consequently, we investigate both measures in our analysis.

4.2. Empirical models of the importance of aggregate teacher turnover

In our first approach we estimate a value-added model that includes average teacher turnover (\bar{T}_{iGy}) and grade reassignments (\bar{R}_{sGy}) measured at the school-grade-year level, the same student and school covariates included in the analysis of teacher effectiveness and school (ϕ_s), school-by-grade (μ_{sG}), and school-by-year (η_{sy}) fixed effects:

$$A_{isGy} = \theta f(A_{iG-1}) + \kappa \bar{T}_{sGy} + \rho \bar{R}_{sGy} + \beta X_{iGy} + \lambda P_{isGy} + \delta S_{isGy} + \mu_{sg} + \eta_{sy} + e_{isGy} \quad (2)$$

The exclusion of controls for teacher experience in some specifications highlights the contribution of the loss of experience to the turnover effects, while the exclusion of the share of teachers new to the grade illuminates the consequences of not accounting for such movement.

Table 5Estimated effects of proportion of teachers new to the grade on achievement gains ($y-1$ to y).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel A. Without experience controls</i>								
Share of teachers new to the school in grade g	-0.066*** (0.016)	-0.086*** (0.017)	-0.047*** (0.016)	-0.044*** (0.016)	-0.111*** (0.019)	-0.109*** (0.018)	-0.044** (0.018)	-0.047*** (0.018)
Share of teachers who switched into grade g	-0.018 (0.025)	-0.041* (0.023)	-0.021 (0.024)	-	-0.020 (0.025)	-	0.024 (0.026)	-
<i>Panel B. With experience controls</i>								
Share of teachers new to the school in grade g	-0.026 (0.017)	-0.036** (0.017)	-0.001 (0.017)	0.001 (0.017)	-0.069*** (0.019)	-0.067*** (0.018)	-0.007 (0.019)	-0.011 (0.018)
Share of teachers who switched into grade g	-0.013 (0.025)	-0.032 (0.023)	-0.012 (0.024)	-	-0.014 (0.025)	-	0.027 (0.026)	-
School fixed effects	N	Y	N	N	N	N	N	N
School-by-grade fixed effects	N	N	Y	Y	N	N	Y	Y
School-by-year fixed effects	N	N	N	N	Y	Y	Y	Y

Notes. Coefficients on the teacher turnover variable come from regressions of math score on the turnover variable plus a cubic in lagged test score, indicators for female, race-ethnicity, low income, special needs, limited English proficient, first year in middle school, family initiated move, shares of students in campus, grade, and year who are female, black, Hispanic, Asian, Native American, low income, special needs, limited English proficient, movers, peer average lagged achievement, a full set of teacher experience indicators (in the bottom panel), and a full set of year-by-grade dummies. All regressions come from a consistent sample of 205,711 observations. Standard errors clustered by teacher-year are in parentheses.

* $p < 0.10$.** $p < 0.05$.*** $p < 0.001$.

In our second approach, we highlight the effects of non-random teacher assignments to grades by calculating average turnover (\bar{T}_{sGy}) and grade reassignments (\bar{R}_{sGy}) at the school-year level. In Eq. (3) we are no longer able to include school-year fixed effects and therefore include only school fixed effects.

$$A_{isGy} = \theta f(A_{iG-1}) + \kappa \bar{T}_{sy} + \rho \bar{R}_{sy} + \beta X_{iGy} + \lambda P_{isGy} + \delta S_{isGy} + \phi_s + e_{isGy} \quad (3)$$

In each approach we investigate measuring turnover alternately as either the share of new entrants in year y or the share of teachers who exited the school between years $y-1$ and y . In models with turnover measured at the grade level these shares are grade specific while in the models with turnover aggregated to the school level the shares reflect the turnover rate for all grades included in the sample combined.

4.3. Impacts of aggregate teacher turnover

Tables 5 and 6 present estimates for alternative specifications of Eq. (2); turnover is measured as proportion new to the grade in Table 5 and proportion who left the grade prior to the current year in Table 6. The top panels do not include measures of teacher experience, while the bottom panels include a full set of teacher experience indicators for each year of experience of the current teacher.

The estimates in Tables 5 suggest that teacher turnover adversely affects the quality of instruction, that changes in the experience distribution account for a portion of the turnover effect, and that the failure to account for teacher grade switching has little or no effect on the turnover coefficients. Even in the specification that includes both school-by-year and school-by-grade fixed effects, proportion new to the school is significant at the 5 percent level in the absence of experience controls.

Estimated effect magnitude is quite sensitive to model specification, raising concerns about the influences of

confounding factors. More specifically, the inclusion of the school-by-grade fixed effects substantially reduces the magnitude of the proportion new coefficient even in specifications that already include school-by-year fixed effects (Column 7 v. 5), while the addition of school-by-year fixed effects to a specification that already includes school-by-grade fixed effects has virtually no effect on the estimate (Column 7 v. 3). This suggests that new entrants tend to be concentrated in lower-achieving grades and that the failure to account for grade-specific differences would likely bias the results. This pattern also tempers concerns that the inability to account for time-varying school factors will inflate the magnitude of the turnover coefficient from the models that aggregate turnover to the school-year level. Finally, inclusion of controls for experience (bottom panel) cuts the magnitude of the proportion new coefficient by roughly 0.04 s.d. in all specifications. This is consistent with a loss of general or school-specific experience reducing the quality of instruction, an issue to which we return below.

Although noticeably smaller in magnitude, a similar pattern emerges in Table 6 when we measure turnover as the share of teachers exiting between $y-1$ and y . This suggests that the share who exits a grade following the previous academic year provides a noisier measure of turnover-related disruptions experienced by a particular grade than the direct measure of the share of new teachers. Some departures will not be replaced, others will be replaced by grade switchers, and only a subset will be replaced by new entrants. Moreover, the exit rate will not capture increases in the number of teachers in a grade. Therefore, in the remainder of the paper we focus on the fraction of teachers new to the grade and rely upon the fixed effects and aggregation to the school level to account for confounding influences.

It should be noted that the sensitivity to the structure of the school fixed effects contrasts the stability of the findings in Ronfeldt et al. (2013), though that work does

Table 6Estimated effects of teacher turnover following the prior year on achievement gains ($y-1$ to y).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel A. Without experience controls</i>								
Share of teachers exiting between $y-1$ and y	-0.036*	-0.047**	-0.022	-0.022	-0.077***	-0.070***	-0.046**	-0.037*
	(0.019)	(0.019)	(0.019)	(0.019)	(0.020)	(0.020)	(0.021)	(0.020)
Share of teachers who switched into grade g	0.009	-0.003	0.000	-	0.041	-	0.056**	-
	(0.024)	(0.023)	(0.025)		(0.025)		(0.027)	
<i>Panel B. With Experience Controls</i>								
Share of teachers exiting between $y-1$ and y	0.002	-0.010	0.011	0.010	-0.041**	-0.038*	-0.017	-0.010
	(0.019)	(0.019)	(0.019)	(0.019)	(0.020)	(0.020)	(0.020)	(0.020)
Share of teachers who switched into grade g	-0.003	-0.018	-0.011	-	0.023	-	0.038	-
	(0.024)	(0.023)	(0.024)		(0.026)		(0.027)	
School fixed effects	N	Y	N	N	N	N	N	N
School-by-grade fixed effects	N	N	Y	Y	N	N	Y	Y
School-by-year fixed effects	N	N	N	N	Y	Y	Y	Y

Notes. See notes to Table 5.

* $p < 0.10$,** $p < 0.05$,*** $p < 0.001$.**Table 7**

Estimated effects of proportion of teachers new to the grade on achievement gains, by level of aggregation.

<i>Shares aggregated to the school-grade-year level</i>				
<i>Without experience controls</i>				
Share new teachers in y	-0.065***		-0.066***	-0.080***
	(0.016)		(0.017)	(0.017)
Share grade switchers $y-1$ to y			-0.018	-0.041*
			(0.025)	(0.023)
<i>With experience controls</i>				
Share new teachers in y	-0.025		-0.026	-0.032*
	(0.017)		(0.017)	(0.017)
Share grade switchers $y-1$ to y			-0.013	-0.032
			(0.025)	(0.023)
<i>Shares aggregated to the school-year level</i>				
<i>Without experience controls</i>				
Share new teachers in y	-0.035		-0.037	-0.043*
	(0.023)		(0.023)	(0.026)
Share grade switchers $y-1$ to y			-0.029	-0.071**
			(0.032)	(0.033)
<i>With experience controls</i>				
Share new teachers in y	0.009		0.007	0.010
	(0.024)		(0.024)	(0.026)
Share grade switchers $y-1$ to y			-0.021	-0.060*
			(0.032)	(0.033)
School effect	N		N	Y
				Y

Notes. See notes to Table 5.

* $p < 0.10$,** $p < 0.05$,*** $p < 0.001$.

not report estimates from the full model that includes both school-by-grade and school-by-year fixed effects. A number of factors could contribute to the divergent findings, but differences in the time periods of the studies raise the possibility that the passage of federal accountability under the No Child Left Behind Act (NCLB) discouraged any systematic placement of new teachers to difficult grades. This would reduce the influence of unobserved, grade-specific confounding factors in Ronfeldt et al. (2013) study that uses data from the post-NCLB period.

With the possibility of nonrandom grade assignments of new entrants in mind, we move to Table 7 where we compare estimates based on the turnover and grade-

switching shares aggregated to the school-grade-year level (top panel) and the school-year level (bottom panel). Regardless of whether school fixed effects are included, aggregation produces virtually the same impact on the proportion new coefficient as does the inclusion of school-by-grade fixed effects in the grade-level specifications previously reported in Table 5. Each method reduces the coefficient magnitude by roughly 0.03 to 0.04, and in specifications with experience controls each method produces estimates that are close to zero and not significant. Given that both aggregation and school-by-grade fixed effects eliminate the influences of the systematic placement of new teachers to particular grades, the findings strongly

Table 8

Estimated effects of proportion of teachers new to the grade on achievement gains, by average achievement and the level of aggregation.

	School average prior achievement			
	Below median achievement		Above median achievement	
<i>Without experience controls</i>				
Share new teachers in t	-0.038 (0.028)	-0.073* (0.040)	-0.043* (0.024)	-0.010 (0.032)
Share grade switchers $t-1$ to t	0.033 (0.036)	-0.123*** (0.046)	0.019 (0.039)	-0.014 (0.044)
<i>With experience controls</i>				
Share new teachers in t	0.003 (0.028)	0.006 (0.040)	-0.008 (0.024)	0.025 (0.033)
Share grade switchers $t-1$ to t	0.055 (0.037)	-0.104** (0.045)	0.003 (0.038)	-0.012 (0.043)
School-grade-year level aggregation	Y	N	Y	N
School-year level	N	Y	N	Y
School effect	N	Y	N	Y
School-grade effect	Y	N	Y	N
School-year effect	Y	N	Y	N

Notes. See notes to Table 7. Turnover is aggregated to either the school-grade-year level or school-year level as noted in the table.

support the presence of such purposeful assignments during this period.

In contrast to the case for the proportion new coefficient, aggregation has a different effect than the inclusion of school-by-grade fixed effects on the share of grade-switchers coefficient. While the addition of school-by-grade fixed effects has little impact, aggregation amplifies the estimated negative effect of grade switching. This suggests that incumbent teachers use time- and grade-varying information about student cohort quality in their classroom requests rather than gravitating toward a particular grade in the school.

We now return to possible heterogeneity in the pattern of turnover effects by school average achievement, similar to the focus in Ronfeldt et al. (2013).¹⁸ In Table 8 we present estimates using turnover aggregated to the school-grade-year level with school-year and school-grade fixed effects or school-year level with school fixed effects (indicated in table). The estimates reveal little difference by school achievement level regardless of whether experience controls are included. The school-level proportion new estimates are somewhat noisier, though once again there is little or no evidence of a negative proportion new effect other than through a reduction in experience.

Finally, aggregation to the school-year level inflates the magnitude of the proportion of teachers who switch grades coefficient for the lower-achievement schools. Note that this is precisely the pattern that would be expected if incumbent teachers tend to switch to grades that are unobservably better for raising achievement and teachers new to the school tend to receive the residual assignments to the more difficult grades. Note that there is little or no evidence of such selective assignment of grade switchers in the higher-achievement schools, consistent with the notion

that such practices occur much more frequently in lower-achievement schools.

5. Reconciling selection and turnover effects in low-achievement schools

Taken as a whole, the evidence on turnover in Lone Star District paints a complicated picture in which teacher turnover does not benefit children attending low-achievement schools despite the tendency for ineffective teachers to leave these schools. The estimates in Table 8 support the notion that the loss of experience and grade switching offset the potential gains from the departure of less effective teachers, though turnover-induced disruption and the ineffectiveness of new hires may also play a role.¹⁹ Given the obstacles to the measurement of disruption effects, we focus here on the effectiveness of replacements for the departing teachers.

Estimating the effectiveness of replacement teachers is challenging because leavers and their replacements teach in different years. Note that given the large number of replacements in their first year of experience, it is not feasible to use prior performance to estimate the effectiveness of entrants. Therefore, in order to provide a comparable benchmark and include new teachers, we measure effectiveness in the final year for departing teachers and the first year for new teachers. We adopt a two-step procedure to compare leavers, grade-switchers, and new entrants. First, we estimate value added for each teacher separately for each year. Second, we regress these value-added estimates on indicators for new entrant, grade switcher, and leaver in a specification that includes school-by-grade-by-year fixed effects. Importantly, in or-

¹⁸ In this section schools are divided into achievement categories on the basis of being above or below median school-average achievement for the entire period. This ensures that schools do not switch samples following changes in achievement and introduce selection bias. Note, however, that a division of schools on the basis of achievement in the first year in the sample produces almost identical estimates.

¹⁹ Consistent with existing research, we find that the return to experience is concentrated in the initial years of teaching. In comparison to teachers with more than three years of experience, first year teachers are on average 0.10 standard deviations less effective and second year teachers are on average 0.03 standard deviations less effective. There appears to be little additional return to experience following the second year of teaching.

Table 9
Average differences in teacher quality by transition type.

	Overall (1)	School average prior achievement	
		Below median (2)	Above median (3)
<i>Panel A. No experience controls</i>			
New to school	−0.053** (0.022)	−0.083** (0.033)	−0.025 (0.027)
Switched into grade	−0.068** (0.029)	−0.124*** (0.041)	−0.013 (0.036)
Exited	−0.044** (0.022)	−0.075** (0.031)	−0.020 (0.030)
<i>Panel B. With experience controls</i>			
New to school	−0.011 (0.023)	−0.021 (0.036)	−0.002 (0.027)
Switched into grade	−0.062** (0.028)	−0.115*** (0.039)	−0.002 (0.038)
Exited	−0.033 (0.021)	−0.059** (0.029)	−0.012 (0.029)

Notes. Estimates are relative to teachers who stay. Explanatory variables are described in note to Table 2. Standard errors clustered by school are in parentheses.

* $p < 0.10$,

** $p < 0.05$,

*** $p < 0.001$.

der to compare leavers with their replacements we move leavers forward by one year in the second stage regression. In other words, year is reset to y for teachers who exit a school following year $y-1$, even though value-added is based on performance in year $y-1$.

Eq. (4) models achievement for student i in school s in grade G with teacher j and year y as a cubic function of lagged mathematics achievement, family background and other influences outside of schools (X), peer composition (P), school factors (S), a teacher-school-grade-year fixed effect (VA), and a random error (e):

$$A_{isGjy} = \theta f(A_{iG-1}) + \beta X_{iGy} + \lambda P_{isGy} + \delta S_{isGy} VA_{sGjy} + e_{isGy} \quad (4)$$

Having obtained a measure of value-added (VA_{sGjy}) for each teacher-school-grade-year combination, we use these estimates to compare the effectiveness of leavers, new entrants and grade switchers:

$$VA_{sGjy} = \sum_k \kappa_k T_{ksGjy} + \eta_{sGy} + e_{isGjy} \quad (5)$$

where $T_{ktsGy} = 1$ if teacher t undergoes transition type k and =0 otherwise. Here we focus on a different set of three types of moves than employed previously: (1) enter grade G and school s in year y ; (2) switch into grade G in year y after teaching in a different grade in the same school in year $y-1$; and (3) exit grade G and school S after $y-1$. The omitted category includes teachers who neither enter the grade in year y nor leave the school following year y . We weight the regressions by enrollment in teacher-school-grade-year cells. All specifications include school-grade-year fixed effects, and we estimate the models with and without experience controls.

The estimates Table 9 show that, despite the negative selection out of schools, new entrants tend to be even less

effective than the teachers who departed. However, the lower experience of new entrants appears to fully account for the differential. In specifications that do not include experience controls (Panel A) the new-entrant coefficient is slightly more negative than the exiting teacher coefficient, but in specifications that do include experience controls (Panel B) the new entrant coefficient is small and insignificant regardless of school-average achievement. By comparison, controlling for experience reduces the coefficient on school leavers by roughly 25 percent; it remains statistically significant for low-achievement schools. Roughly one third of new entrants but only 17 percent of teachers leaving after the school year have no prior experience, consistent with the larger impact of experience on the entrant coefficient.

Consistent with the aggregate estimates in Tables 7 and 8, Table 9 also shows the lower effectiveness of grade switchers in low-achievement schools, even conditional on experience. Although reassignment to a different grade can occur even in the absence of any exits, grade reassignments appear to contribute to the negative effect of turnover. The fraction of teachers in a grade who taught in a different grade in the prior year is roughly 33 percent (1.3 percentage points) higher in grades that had at least one teacher leave the school following the prior year than in grades that did not have any teachers leave the school.

There is no direct evidence on disruption effects. But, the findings on experience and grade-switching suggest that any turnover-induced disruption to school operations has at most a small effect on the quality of instruction.

6. Conclusions

This analysis shows that on net turnover adversely affects the quality of instruction in lower-achievement schools. This result is due to a turnover-induced loss of general and grade-specific experience that is sufficient to offset the potential gains that come from the departure of teachers who on average are less effective than stayers. In higher-achievement schools, there is little evidence of adverse turnover effects. Importantly, turnover is not accurately characterized by assertions that high-poverty schools regularly lose their most effective teachers – which would imply that the reduction of turnover per se would substantially improve the quality of schools.

Our analysis of teacher turnover and student achievement in the Lone Star District of Texas takes special care in measuring teacher effectiveness to account for potential confounding factors including the endogeneity of school choice and classroom assignment. The results show that failure to account adequately for the grade-assignment of teachers (here through school-by-grade fixed effects or the measurement of the turnover variables at the school level) inflates estimates of teacher turnover effects on achievement.

Teachers who exit Texas public schools (6 percent annually of all teachers and 12 percent of new teachers) are on average less effective than those who remain. This finding makes clear the inadequacy of unfocused teacher policies, including universal pay increases, designed to reduce overall turnover without consideration of quality. The

detailed analysis shows this conclusion is particularly true for policies designed to reduce turnover among experienced teachers. In contrast, the findings provide support for evaluation and compensation systems such as those implemented in Washington, DC and Dallas, Texas, that link pay increases with performance in an effort to retain, support and attract more effective teachers. The additional compensation payed to highly-effective teachers in disadvantaged schools in Washington seems particularly promising given the larger turnover costs for these schools.

The estimates are also consistent with the belief that a lack of success leads many teachers to exit, particularly from low-achieving schools. This finding reinforces general policy prescriptions about improving teacher quality but adds a complementary impact through the effect on teacher turnover. General professional development programs have not proved very successful,²⁰ but there are a few suggestive studies that indicate a focus on feedback and mentoring might be merited.²¹ The turnover results would also support increased efforts to improve the selection of entering teachers, although the existing evidence again is not clear about methods to predict accurately future effectiveness in the classroom.²²

²⁰ See [Garet et al. \(2008\)](#), [Garet et al. \(2011\)](#), and [TNTF \(2015\)](#).

²¹ [Taylor and Tyler \(2012\)](#) find that feedback through high-quality teacher evaluations leads to improvements even among more experienced teachers. Another recent study by [Papay, Taylor, Tyler, and Laski \(2016\)](#) provides evidence from a random assignment field experiment that mentoring greatly improves teacher effectiveness, raising test scores by 0.12 standard deviations. Nonetheless, there is little evidence to date about the ability to implement such programs on a large scale.

²² In their review, [Staiger and Rockoff \(2010\)](#) find little support for the effectiveness of current hiring processes of schools. On the other hand, [Jacob, Lefgren, and Sims \(2010\)](#) find some evidence that a combination of instruments has some predictive power in terms of future effectiveness, and [Jacob, Rockoff, Taylor, Lindy, and Rosen \(2016\)](#) suggest that existing applicant information could produce better results.

Finally, our estimates of the average ineffectiveness of teachers exiting from Texas public schools undoubtedly reflects forced departures of poorly performing teachers along with voluntary choices of teachers who recognize that they are not effective. In the absence of information on the circumstances of the separation, it is not possible to quantify the relative quality of voluntary leavers versus those forced out.

Appendix

Table A1

Correspondence between assignment into sorted and not-sorted categories by test statistic.

Division by Chi-square test	Division by test of mean classroom achievement			
	Sorted	Unsorted	Missing	Total
<i>Panel A. Proportions</i>				
Sorted	0.74	0.26	0.00	0.17
Unsorted	0.26	0.74	0.00	0.14
Missing	0.34	0.36	0.30	0.68
Total	0.40	0.40	0.21	1.00
<i>Panel B. Numbers (School/grade/year)</i>				
Sorted	324	114	0	438
Unsorted	94	273	0	367
Missing	588	624	527	1739
Total	1006	1011	527	2544

Notes. The Chi Square test of the independence of transitions across grades requires an additional year of data with at least two teachers in the grade, leading to the classification of missing for a number of observations.

Table A2

Correlation coefficients between transitions and aggregate share new to grade in subsequent year.

	Stay in campus	Move campus	Move district	Exit texas	Share new hires in grade	Share exiting in prior year	Share current teachers switched to new grade
Stay in campus	1.0000						
Move Campus	-0.6277	1.000					
Move district	-0.4165	-0.0628	1.000				
Exit texas	-0.5478	-0.0826	-0.0548	1.000			
Share new hires in grade	-0.0886	0.0397	0.0379	0.0655	1.000		
Share exiting prior year	-0.0693	0.0141	0.0396	0.0626	0.5196	1.0000	
Share current teachers switched to new grade	-0.0040	-0.0028	0.0159	-0.0030	-0.1752	0.1390	1.000

Notes. Estimated from a sample of 215,166 student-year observations.

Table A3
Predictors of grade reassignments.

	Switched between t and $t-1$					
	Without experience controls			With experience controls		
	(1)	(2)	(3)	(4)	(5)	(6)
Share of teachers exiting in $t-1$	0.765*** (0.025)	0.741*** (0.028)	0.764*** (0.034)	0.751*** (0.025)	0.733*** (0.028)	0.749*** (0.034)
Math gains $t-2$ to $t-1$	-0.160*** (0.021)	-0.171*** (0.021)	-0.202*** (0.023)	-0.149*** (0.021)	-0.159*** (0.021)	-0.185*** (0.023)
Average gains in $t-1$ produced by new teachers in $t-1$	0.066 (0.043)	0.061 (0.044)	0.048 (0.053)	0.069 (0.043)	0.058 (0.044)	0.044 (0.052)
Current Salary	-0.000*** (0.000)	-0.000*** (0.000)	-0.000*** (0.000)	-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
Teacher is male	-0.024 (0.018)	-0.016 (0.019)	-0.015 (0.018)	-0.024 (0.018)	-0.016 (0.019)	-0.015 (0.018)
Teacher has advanced degree	0.040** (0.018)	0.048** (0.019)	0.053*** (0.019)	0.039* (0.021)	0.046** (0.021)	0.053** (0.022)
Teacher is black	-0.041** (0.018)	-0.034* (0.019)	-0.029 (0.018)	-0.046*** (0.017)	-0.040** (0.019)	-0.035* (0.018)
Teacher is Hispanic	0.031 (0.036)	0.019 (0.038)	0.027 (0.037)	0.037 (0.035)	0.026 (0.037)	0.032 (0.036)
Teacher is other race	-0.051 (0.049)	-0.035 (0.054)	-0.039 (0.055)	-0.056 (0.048)	-0.044 (0.054)	-0.049 (0.054)
School FE	N	Y	N	N	Y	N
School-year fe	N	N	Y	N	N	Y
Sample Size	202,882	202,882	202,882	202,882	202,882	202,882

References

- Aaronson, D., Barrow, L., & Sander, W. (2007). Teachers and Student Achievement in the Chicago Public High Schools. *Journal of Labor Economics*, 25(1), 95–135.
- Armor, D. J., Conry-Oseguera, P., Cox, M., King, N., McDonnell, L., Pascal, A., et al. (1976). *Analysis of the school preferred reading program in selected Los Angeles minority schools*. Santa Monica, CA: Rand Corp.
- Ballou, D., & Podgursky, M. (2002). Returns to seniority among public school teachers. *Journal of Human Resources*, 37(4), 892–912 (Fall).
- Boyd, D., Grossman, P., Lankford, H., Loeb, S., & Wyckoff, J. (2006). How changes in entry requirements alter the teacher workforce and affect student achievement. *Education Finance and Policy*, 1(2), 176–216 (Spring).
- Branch, G. F., Hanushek, E. A., & Rivkin, S. G. (2012). *Estimating the effect of leaders on public sector productivity: the case of school principals*. Cambridge, MA: National Bureau of Economic Research NBER Working Paper W17803.
- Chetty, R., Friedman, J. N., & Rockoff, J. (2014a). Measuring the impacts of teachers I: evaluating bias in teacher value-added estimates. *American Economic Review*, 104(9), 2593–2632 (September).
- Chetty, R., Friedman, J. N., & Rockoff, J. (2014b). Measuring the impacts of teachers II: teacher value-added and student outcomes in adulthood. *American Economic Review*, 104(9), 2633–2679 (September).
- Chetty, R., Friedman, J. N., & Rockoff, J. (2016). Using lagged outcomes to evaluate bias in value-added models. *American Economic Review*, 106(5), 393–399.
- Chingos, M. M., & West, M. R. (2011). Promotion and reassignment in public school districts: how do schools respond to differences in teacher effectiveness? *Economics of Education Review*, 30(3) (June).
- Chingos, M. M., & West, M. R. (2012). Do more effective teachers earn more outside the classroom? *Education Finance and Policy*, 7(1), 8–43 (Winter).
- Clotfelter, C. T., Ladd, H. F., & Vigdor, J. I. (2006). Teacher-student matching and the assessment of teacher effectiveness. *Journal of Human Resources*, 41(4), 778–820 (Fall).
- Dieterle, S., Guarino, C. M., Reckase, M. D., & Wooldridge, J. M. (2015). How do principals assign students to teachers? Finding evidence in administrative data and the implications for value added. *Journal of Policy Analysis and Management*, 34(1), 32–58 (Winter).
- Dolton, P. J., & van der Klaauw, W. (1999). The turnover of teachers: a competing risks explanation. *Review of Economics and Statistics*, 81(3), 543–552 (August).
- Feng, L., & Sass, T. (2011). *Teacher quality and teacher mobility*. Washington, DC: National Center for Analysis of Longitudinal Data in Education Research CALDER Working Paper 57 (January).
- Garet M. S., Cronen S., Eaton M., Kurki A., Ludwig M., Jones W., et al. (2008). The impact of two professional development interventions on early reading instruction and achievement. National Center for Education Evaluation and Regional Assistance, Institute of Education Sciences. Washington, DC: U.S. Department of Education (September).
- Garet, M. S., Wayne, A. J., Stancavage, F., Taylor, J., Eaton, M., Walters, K., et al. (2011). *Middle school mathematics professional development impact study: findings after the second year of implementation*. Washington, DC: Institute of Education Sciences NCEE 2011-4024 (April).
- Goldhaber, D., Cowan, J., & Walch, J. (2013). Is a good elementary teacher always good? Assessing teacher performance estimates across subjects. *Economics of Education Review*, 36(October), 216–228.
- Goldhaber, D., Gross, B., & Player, D. (2011). Teacher career paths, teacher quality, and persistence in the classroom: are public schools keeping their best? *Journal of Policy Analysis and Management*, 30(1), 57–87 (Winter).
- Guarino, C. M., Reckase, M. D., & Wooldridge, J. M. (2015). Can value-added measures of teacher performance be trusted? *Education Finance and Policy*, 10(1), 117–156.
- Hanushek, E. A. (1971). Teacher characteristics and gains in student achievement: estimation using micro data. *American Economic Review*, 60(2), 280–288 (May).
- Hanushek, E. A. (1979). Conceptual and empirical issues in the estimation of educational production functions. *Journal of Human Resources*, 14(3), 351–388 (Summer).
- Hanushek, E. A. (1986). The economics of schooling: production and efficiency in public schools. *Journal of Economic Literature*, 24(3), 1141–1177 (September).
- Hanushek, E. A. (1992). The trade-off between child quantity and quality. *Journal of Political Economy*, 100(1), 84–117 (February).
- Hanushek, E. A., Kain, J. F., O'Brien, D. M., & Rivkin, S. G. (2005). *The market for teacher quality*. Cambridge, MA: National Bureau of Economic Research NBER Working Paper No. 11154 (February).
- Hanushek, E. A., Kain, J. F., & Rivkin, S. G. (2004). Why public schools lose teachers. *Journal of Human Resources*, 39(2), 326–354 (Spring).
- Hanushek, E. A., & Rivkin, S. G. (2010). Generalizations about using value-added measures of teacher quality. *American Economic Review*, 100(2), 267–271 (May).
- Hanushek, E. A., & Rivkin, S. G. (2012). The distribution of teacher quality and implications for policy. *Annual Review of Economics*, 4, 131–157.
- Jackson, C. K. (2012). Match quality, worker productivity, and worker mobility: direct evidence from teachers. *Review of Economics and Statistics*, 95(4), 1096–1116 (October).
- Jacob, B. A., Lefgren, L., & Sims, D. (2010). The persistence of teacher-induced learning. *Journal of Human Resources*, 45(4), 915–943 (Fall).

- Jacob, B., Rockoff, J. E., Taylor, E. S., Lindy, B., & Rosen, R. (2016). *Teacher applicant hiring and teacher performance: evidence from DC public schools*. Cambridge, MA: National Bureau of Economic Research NBER Working Paper No. 22054(March).
- Kane, T. J., McCaffrey, D. F., Miller, T., & Staiger, D. O. (2013). *Have we identified effective teachers? Validating measures of effective teaching using random assignment*. MET Project: Bill and Melinda Gates Foundation (January).
- Kane, T. J., Rockoff, J. E., & Staiger, D. O. (2008). What does certification tell us about teacher effectiveness? Evidence from New York City. *Economics of Education Review*, 27(6), 615–631 (December).
- Kane, T. J., & Staiger, D. O. (2008). *Estimating teacher impacts on student achievement: an experimental evaluation*. Cambridge, MA: National Bureau of Economic Research NBER Working Paper 14607(December).
- Koedel, C., & Betts, J. (2011). Does student sorting invalidate value-added models of teacher effectiveness? An extended analysis of the Rothstein critique. *Education Finance and Policy*, 6(1), 18–42 Winter.
- Meghir, C., & Rivkin, S. G. (2011). Econometric methods for research in education. In E. A. Hanushek, S. Machin, & L. Woessmann (Eds.), *Handbook of the economics of education: Vol. 3* (pp. 1–87). Amsterdam, North Holland.
- Murnane, R. J. (1975). *Impact of school resources on the learning of inner city children*. Cambridge, MA: Ballinger.
- Murnane, R. J., & Phillips, B. (1981). What do effective teachers of inner-city children have in common? *Social Science Research*, 10(1), 83–100 March.
- Ost, B. (2014). How do teachers improve? The relative importance of specific and general human capital. *American Economic Journal: Applied Economics*, 6(2), 127–151 April.
- Ost, B., & Schiman, J. (2015). Grade-specific experience, grade reassignments, and teacher turnover. *Economics of Education Review*, 46(June), 112–126.
- Papay, J. P., Taylor, E. S., Tyler, J. H., & Laski, M. (2016). *Learning job skills from colleagues at work: evidence from a field experiment using teacher performance data*. Cambridge, MA: National Bureau of Economic Research NBER Working Paper No. 21986(February).
- Rivkin, S. G., Hanushek, E. A., & Kain, J. F. (2005). Teachers, schools, and academic achievement. *Econometrica*, 73(2), 417–458 (March).
- Rockoff, J. E. (2004). The impact of individual teachers on student achievement: evidence from panel data. *American Economic Review*, 94(2), 247–252 May.
- Ronfeldt, M., Loeb, S., & Wyckoff, J. (2013). How teacher turnover harms student achievement. *American Educational Research Journal*, 50(1), 4–36 February.
- Rothstein, J. (2010). Teacher quality in educational production: tracking, decay, and student achievement. *Quarterly Journal of Economics*, 125(1), 175–214 February.
- Rothstein, J. (2014). *Revisiting the impacts of teachers*. Berkeley: Department of Economics: University of California (mimeo)(October).
- Sass, T. R., Hannaway, J., Xu, Z., Figlio, D. N., & Feng, L. (2012). Value added of teachers in high-poverty schools and lower poverty schools. *Journal of Urban Economics*, 72(2–3), 104–122 (September–November).
- Sass, T. R., Semykina, A., & Harris, D. N. (2014). Value-added models and the measurement of teacher productivity. *Economics of Education Review*, 38, 9–23.
- Staiger, D. O., & Rockoff, J. E. (2010). Searching for effective teachers with imperfect information. *Journal of Economic Perspectives*, 24(3), 97–118 (Summer).
- Stinebrickner, T. R. (2002). An analysis of occupational change and departure from the labor force. *Journal of Human Resources*, 37(1), 192–216 (Winter).
- Taylor, E. S., & Tyler, J. H. (2012). The effect of evaluation on teacher performance. *American Economic Review*, 102(7), 3628–3651.
- TNTP. (2015). *The mirage: Confronting the hard truth about our quest for teacher development: The New Teachers Project*.
- Wooldridge, J. M. (2002). *Econometric analysis of cross section and panel data*. Cambridge, MA: MIT Press.