

NBER WORKING PAPER SERIES

LEARNING JOB SKILLS FROM COLLEAGUES AT WORK:  
EVIDENCE FROM A FIELD EXPERIMENT USING TEACHER PERFORMANCE DATA

John P. Papay  
Eric S. Taylor  
John H. Tyler  
Mary Laski

Working Paper 21986  
<http://www.nber.org/papers/w21986>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
February 2016

We thank Jonah Rockoff as well as seminar participants at the NBER Summer Institute, New York Federal Reserve, and Harvard Graduate School of Education for helpful comments and suggestions. The Bill & Melinda Gates Foundation provided financial support of this research, and we benefited greatly from discussions with our program officer Steve Cantrell. We are equally indebted to the Tennessee Department of Education, and particularly Nate Schwartz, Laura Booker, Tony Pratt, Luke Kohlmoos, and Sara Heyburn for their collaboration throughout. Finally, we thank Verna Ruffin, superintendent in Jackson-Madison County Schools, and the principals and teachers who participated in the program. All opinions and errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by John P. Papay, Eric S. Taylor, John H. Tyler, and Mary Laski. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Learning Job Skills from Colleagues at Work: Evidence from a Field Experiment Using Teacher Performance Data

John P. Papay, Eric S. Taylor, John H. Tyler, and Mary Laski

NBER Working Paper No. 21986

February 2016

JEL No. I2,J24,M53

**ABSTRACT**

We study on-the-job learning among classroom teachers, especially learning skills from coworkers. Using data from a new field experiment, we document meaningful improvements in teacher job performance when high- and low-performing teachers working at the same school are paired and asked to work together on improving the low-performer's skills. In particular, pairs are asked to focus on specific skills identified in the low-performer's prior performance evaluations. In the classrooms of low-performing teachers treated by the intervention, students scored 0.12 standard deviations higher than students in control classrooms. These improvements in teacher performance persisted, and perhaps grew, in the year after treatment. Empirical tests suggest the improvements are likely the result of low-performing teachers learning skills from their partner.

John P. Papay  
Box 1938  
340 Brook Street  
Brown University  
Providence, RI 02912  
john\_papay@brown.edu

Eric S. Taylor  
Harvard Graduate School of Education  
Gutman Library 469  
6 Appian Way  
Cambridge, MA 02138  
eric\_taylor@gse.harvard.edu

John H. Tyler  
Box 1938  
340 Brook Street  
Brown University  
Providence, RI 02912  
and NBER  
john\_tyler@brown.edu

Mary Laski  
Box 1938  
340 Brook Street  
Brown University  
Providence, RI 02912  
Mary\_Laski@brown.edu

*“Some types of knowledge can be mastered better if simultaneously related to a practical problem.” Gary Becker (1962)*

Can employees learn job skills from their coworkers? Whether and how peers contribute to on-the-job learning, and at what costs, are practical questions for personnel management. Economists’ interest in these questions dates to at least Alfred Marshall (1890) and, more recently, Gary Becker (1962) and Robert Lucas (1988). Yet, despite the intuitive role for coworkers in human capital development, empirical evidence of learning from coworkers is scarce.<sup>2</sup> In this paper we present new evidence from a random-assignment field experiment in U.S. public schools: low-performing classroom teachers in treatment schools were each matched to a high-performing colleague in their school, and pairs were encouraged to work together on improving their teaching skills. We report positive treatment effects on teacher productivity, as measured by contributions to student achievement growth, particularly for low-performing teachers. We then test empirical predictions consistent with peer learning and other potential mechanisms.

While there is limited evidence on learning from coworkers specifically, there is a growing literature on productivity spillovers among coworkers generally. Morreti (2004) and Battu, Belfield, and Sloane (2003) document human capital spillovers broadly, using variation between firms, but without insight to mechanisms. Several other papers, each focusing on a specific firm or occupation as we do, also find spillovers; the apparent mechanisms are shared production opportunities or peer influence on effort (Ichino and Maggi 2000, Hamilton, Nickerson and Owan 2003, Bandiera, Barankay and Rasul 2005, Mas and Moretti 2009, Azoulay, Graff Zivin, and Wang 2010). Moreover, these spillovers may be substantial. Lucas

---

<sup>2</sup> We are focused in this paper on coworker peers and learning on-the-job. A large literature examines the role of peers in classroom learning and other formal education settings (for a review see Sacerdote 2010).

(1988) suggests human capital spillovers, broadly speaking, could explain between-country differences in income.

One empirical example of learning from coworkers comes from the study of classroom teachers. Jackson and Bruegmann (2009) find a teacher's productivity, as measured by contribution to her students' test score growth, improves when a new higher-performing colleague arrives at her school; then, consistent with peer learning, the improvements persist after she is no longer working with the same colleague (i.e., teaching the same grade in the same school). The authors estimate that prior coworker effectiveness explains about one-fifth of the variation in teacher performance.

In this paper we also focus on classroom teachers. While we believe the paper makes an important general contribution, a better understanding of on-the-job learning among teachers specifically has sizable potential value for students and economies. Classroom teaching represents a substantial investment of resources: one out of ten college-educated workers in the U.S. is a public school teacher, and public schools spend \$285 billion annually on teacher wages and benefits (U.S. Census Bureau 2015, Table 6).<sup>3</sup> And there is substantial variability in teacher job performance: measured both in the short-run with students' test scores (see Jackson, Rockoff, and Staiger 2014 for a review) and the long-run with students' economic and social success years later as adults (Chetty, Friedman, and Rockoff 2014). One seemingly consistent source of differences in teacher performance is experience on the job (Rockoff 2004, Harris and Sass 2011, Papay and Kraft 2015). Estimated differences due to experience are much larger than differences in formal pre-service or in-service training (Jackson, Rockoff, and Staiger 2014).

---

<sup>3</sup> Authors' calculations of workforce share from Current Population Survey 1990-2010.

We report here on a field experiment in Tennessee designed to study on-the-job, peer learning between teachers who work at the same school. Schools were randomly assigned to either treatment or a business-as-usual control. In treatment schools, low-performing teachers were each matched to a high-performing partner using detailed micro-data from prior performance evaluations. In Tennessee, teachers are observed in the classroom multiple times per year and scored in 19 specific skills (e.g., “questioning,” “lesson structure and pacing,” “managing student behavior”). Each low-performing “target” teacher was identified as such because his prior evaluation scores were particularly low in one or more of the 19 skill areas. Then his high-performing “partner” was chosen because she had high scores in (many of) the same skill areas. Each pair of teachers was encouraged by their principal to work together during the school year on improving teaching skills identified by evaluation data. Thus the topics and skills teachers worked on were specific to each pair and varied between pairs. More generally, pairs were encouraged to examine each other’s evaluation results, observe each other teaching in the classroom, discuss strategies for improvement, and follow-up with each other’s commitments throughout the school year.<sup>4</sup>

We find that treatment—pairing classroom teachers to work together on improving skills—improves teachers’ job performance, as measured by their students’ test score growth. At the end of the school year, the average student in a treatment school scores  $0.06\sigma$  (student standard deviations) higher on standardized math and reading/language arts tests than she would have in a control school, regardless of whether her teacher participated in a partnership. The gains are concentrated among “target” teachers; in target teachers’ classrooms students score  $0.12\sigma$  higher. These are meaningful gains. One standard deviation in teacher performance is

---

<sup>4</sup> The treatment was designed in a collaboration between the research team and several people at the Tennessee Department of Education. TNDOE also played key roles in carrying out the experiment and collecting data.

typically estimated to be  $0.15\text{-}0.20\sigma$  (Hanushek and Rivkin 2010). In other words, a gain of  $0.12\sigma$  is roughly equivalent to the difference between being assigned to a median teacher instead of a bottom quartile teacher. These improvements in teacher performance persist, and perhaps grow, in the school year following treatment; in year two the effect for target teachers is a marginally significant  $0.25\sigma$ .

Interpreting these differences as causal effects of treatment rests mainly on the random assignment of schools. While the “target” and “partner” roles were not randomly assigned, the roles were assigned by algorithm for both treatment and control schools prior to randomization, as we detail in Section 1. The estimates in the previous paragraph are intent-to-treat estimates based on algorithm-assigned roles.

After documenting average treatment effects, we turn to examining mechanisms. In particular we ask: Can the performance improvements be attributed to growth in teachers’ skills from peer learning, or are other changes in behavior or effort behind the estimated effects? Larger effects for target teachers, effects which persist over time, are highly suggestive of skill growth, but could also result if partnering increased target teachers’ motivation or effort, or provided new opportunities to share resources or tasks (Jackson and Bruegmann 2009). In Section 3, we test a number of empirical predictions motivated by these potential mechanisms. If the underlying mechanism is skill growth, we would predict larger treatment effects for target teachers when the high-performing partner’s skill strengths match more of the target teacher’s weak areas. We would also predict larger improvements in classroom observation scores for the skills matched, but smaller or no improvements in the skills left unmatched. Both predictions are true empirically. If, however, the mechanism is shared production or resources, we would predict larger effects when teacher pairs teach similar grade-levels or subjects. If the mechanism is effort

or motivation, we would predict larger effects when there were larger gaps in prior performance between paired teachers, on the assumption that the comparison of performance induces greater effort (as in Mas and Moretti 2009). Neither of these latter predictions is borne out in the data. In short, the available data suggest target teachers learned new skills from their partner.

One contextual feature of the experiment is also important to interpreting these results. The detailed micro-data with which teachers were paired are taken from the state's performance evaluation system for public school teachers, which the Tennessee Department of Education introduced in 2011. Locally the treatment was known as the "Evaluation Partnership Program." These connections to formal evaluation, and its stakes, likely influenced principals' and teachers' willingness to participate and the nature of their participation.<sup>5</sup> The evaluation context also affects the counterfactual behavior of control schools and teachers. This context may partly explain why we find positive effects in this case while other research does not consistently find effects of formal mentoring or formal on-the-job training for teachers (see reviews by Jackson, Rockoff, and Staiger 2014, and Yoon et al. 2007).<sup>6</sup> More generally, this paper also belongs to a small literature on how evaluation programs affect teacher performance (Taylor and Tyler 2012, Steinberg and Sartain 2015, Bergman and Hill 2015, Dee and Wyckoff 2015). Taylor and Tyler (2012) study veteran teachers who were evaluated by and received feedback from experienced, high-performing teachers; the resulting improvements in teacher productivity persisted for years after the peer evaluation ended.

The performance improvements documented in this paper suggest teachers can learn job skills from their colleagues—empirical evidence of the intuitive benefit of skilled coworkers in

---

<sup>5</sup> In one-on-one interviews, some participating teachers said they were willing to participate because teacher pairs were matched based on specific skills and not on a holistic measure of performance.

<sup>6</sup> Exceptions include an example of mentoring studied by Rockoff (2008) and an example of training studied by Angrist and Lavy (2001).

human capital development. The magnitude of those improvements suggests peer learning may be as important as on-the-job experience in teacher skill development (Rockoff 2004, Papay and Kraft forthcoming); indeed, peer learning may be a key contributor to the oft-cited estimates of returns to experience in teaching. Most practically, the treatment and results suggest promising ideas for managing the sizable teacher workforce.

Next, in Section 1, we describe the treatment in detail, along with other features of the experimental setting and data. In Section 2 we describe the average treatment effects and treatment effects by teachers' assigned partnership role. Section 3 discusses potential mechanisms and presents tests of empirical predictions related to those mechanisms. We conclude in Section 4 with some further discussion of the results.

## **1. Treatment, Setting, and Data**

### *1.1 Treatment*

We report on a field experiment designed to study on-the-job, peer learning between teachers who work at the same school. At schools randomly assigned to the treatment condition—known in the schools as the “Evaluation Partnership Program”—low-performing “target” teachers were paired with a high-performing “partner” teacher, and each pair was encouraged to work together on improving each other’s teaching skills over the course of the school year. Importantly, teachers were matched using micro-data from state-mandated performance evaluations. As described further in the next section, these prior evaluations include separate performance ratings for many specific instructional skills (e.g., “questioning,” “lesson structure and pacing,” “managing student behavior”). Each target teacher was identified as such because he had low scores in one or more specific skill areas; his matched partner was selected



because she had high scores in (many of) the same skill areas. Pairs were approached by their school principal and asked to work together for the year focusing on the strength-matched-to-weakness skill areas, with the goal of improving instructional skills. Thus the topics and skills teachers worked on were specific to each pair and varied between pairs. More generally, pairs were encouraged to scrutinize each other's evaluation results, observe each other teaching in the classroom, discuss strategies for improvement, and follow-up with each other's commitments throughout the school year.

While individual teacher pairings were the focus of the intervention, treatment was assigned at the school level. Thus the success of individual pairs may have been influenced by the principal's role or support, or influenced by other teacher pairs in the school working in the same kinds of ways. Certainly the extensive margin of treatment take-up was in the hands of the school principal, as described below.

### *1.1.1 Teacher Evaluation in Tennessee*

All public school teachers in Tennessee are evaluated annually. Beginning in the 2011-12 school year, the state introduced new, more-intensive requirements for teacher evaluation. The new evaluations include both (i) direct assessments of teaching skills in classroom observations, and (ii) measures of teachers' contributions to student achievement. We focus in this section on the classroom observation scores because they are the micro-data used in matching target and partner pairs, and the motivation for each pair's work together.

Each teacher is observed while teaching and scored multiple times during the course of the school year, typically by the school principal or vice-principal. Observations and scores are structured around a rubric known as the "TEAM rubric" which measures 19 different instructional skill areas or "indicators" which are grouped into three categories: instruction,

planning, and classroom environment.<sup>7</sup> The rubric is based in part of the work of Charlotte Danielson (1996). Skill areas include things like “managing student behavior,” “instructional plans,” “teacher content knowledge,” and many others. As an example, Figure 1 reproduces the rubric for “Questioning.” Teachers are scored from 1-5 on each skill area: 1 significantly below expectations, 2 below expectations, 3 at expectations, 4 above expectations, 5 significantly above expectations. As the Figure 1 example suggests, the rubric language describes relatively specific skills and behaviors, not vague general assessments of teaching effectiveness. As described in the next section, we use these micro-data on 19 different skill areas to match high- and low-performing teachers in working pairs.

Table 1 lists all 19 skills and provides summary statistics for teachers’ scores in the year prior to the experiment. Teachers score lowest in “problem solving” and highest in “content knowledge.” Scores for classroom environment skills are most variable, especially “managing student behavior.” Table 1 draws on teachers in the treatment and control schools, but these statistics are similar when calculated using all Tennessee schools.

At the end of the school year, the classroom observation micro-data are aggregated to produce a final, overall, univariate observation evaluation for each teacher. In the 2012-13 school year, the second under the state’s new requirements, teachers in our sample scored quite high in classroom observations: more than four-fifths received an overall average rating of 3 or higher (Table 1 Column 3), while just 1.2 percent scored 2 or lower.<sup>8</sup> Critically for this study, however, there was substantial variation between and within teachers at the 19-skill micro-data level as

---

<sup>7</sup> TEAM stands for Tennessee Educator Acceleration Model. Most Tennessee districts, including the district where this paper’s data were collected, use the TEAM rubric. Some districts use alternative rubrics. The full TEAM rubric is available at: <http://team-tn.org/wp-content/uploads/2013/08/TEAM-General-Educator-Rubric.pdf>.

<sup>8</sup> The small percentage of “low” final overall scores is not a-typical, even after the revisions in teacher evaluation programs in recent years (Wiesberg 2009, New York Times, March 30, 2013)

detailed in Table 1. Two out of five teachers received a score lower than 3 in at least one of the 19 skill areas; and, among that 41 percent, the average number of skills with a score below 3 is 5.8. (The threshold of below 3 was used to identify skills for matching pairs, as described in the next section.) In the state evaluation process, the overall observation rating is combined with student achievement data to produce a summative score for each teacher. It is important to note that, while certainly used in the state’s formal evaluation scores, student achievement data were not used in the matching of teachers or communication with teachers about the goals of the teacher partnerships.

### *1.1.2 Identifying and Matching High- and Low-Performing Teachers*

For the purposes of this experiment, a teacher was identified as a “low-performing” or “target” if he had a score less than 3 in one or more of the 19 skill areas. Similarly, a teacher was identified as a “high-performing” *potential* “partner” if she had a score of 4 or higher in one or more skill area. Both the set of target and the set of potential partners were identified based on pre-experiment evaluation scores: the average of a teacher’s scores from the prior school year 2012-13 and, when available, the first observation of 2013-14.<sup>9</sup>

Our matching algorithm followed these steps and rules: (1) Consider each possible pairing of a target and a partner teacher who work in the same school, and calculate the total number of skill areas (out of 19 possible) where there is a strength-to-weakness skill match for that pairing. A strength-to-weakness match occurs when the target teacher has a score less than 3

---

<sup>9</sup> If a teacher was identified as both a “target” and “partner” by these rules, the teacher was included only on the target list. Additionally, any potential “target” teacher with an *overall* observation rating of “above expectation” (4) or “significantly above expectations” (5) was excluded from the target teacher list.

Two schools had few or no teachers identified as “target” by these rules. In those two cases, after random assignment, we used a threshold of less than or equal to 3, instead of strictly less than 3. One school had many teachers identified as “target.” In that school we limited the set of target teachers to those with 8 or more skill areas (out of 19) with a score less than 3. While these adjustments aided in the practical implementation of the program, our ITT estimates use only the original assignments of teachers, not the assignments after these school-specific adjustments.

in a given skill, and the partner has a score of 4 or higher in the same skill. (2) For each school, list all possible configurations of pairings where each potential partner is matched to just one target teacher. (3) Choose the set of pairings which maximizes the number of strength-to-weakness matches (out of  $19 * T$  possible, where T is the number of target teachers).<sup>10</sup> This algorithm produced a set of “recommended” matches for each school. We created recommended match lists for both treatment and control schools pre-experiment, but the lists were only provided to principals in the treatment schools.<sup>11</sup>

The principal in each treatment school was responsible for introducing each target-partner pair, explaining why the two had been paired, and encouraging the pair in their work together. Each teacher, target or partner, ultimately decided to what extent she would participate and, in that sense, the experiment can be thought of as an encouragement-style design.

Additionally, principals were encouraged to review the recommended list of matches and make changes as they saw fit. This latitude for principals was intended to help improve matches using local knowledge not observed in the evaluation data; for example, a recommended match may have paired two individuals known to not work well together. To encourage data-informed adjustments, we provided principals with a list of additional potential matches for each target teacher. These additional matches were the five potential partners with the highest strength-to-weakness scores from step (1) above, regardless of the optimization and constraints in steps (2) and (3). Principals received a spreadsheet listing the 19 skill areas for each teacher to show where target and potential partner teachers matched. An example is shown in Figure 2. We

---

<sup>10</sup> The core of this matching approach is sometimes called the Hungarian Algorithm or Method.

<sup>11</sup> We can say definitively that no control school received a list of target teachers and proposed partners. However, we discussed the idea of the program with all principals, and we cannot ensure that the kernel of the idea was not adopted by principals in the control schools. Because principals conduct the observations, they had on hand all of the information necessary to conduct such pairings. Our discussions with district officials, though, suggest that principals in control schools did not undertake such a program. Any such activities would bias downward our estimated average treatment effects.

provided the proposed matches and information for participating principals and teachers to principals in early November 2013.<sup>12</sup>

### *1.2 Sample and Random Assignment*

The experiment was conducted at 14 elementary and middle schools (7 treatment, 7 control) in a medium-sized district in Tennessee. The district, Jackson-Madison County School System, is the 12th largest in the state enrolling approximately 13,000 students. Across the district, 77 percent of students are economically disadvantaged, 61 percent are African-American, 32 percent are white, and 7 percent are Hispanic. The district spends about \$9,750 per pupil annually. In 2013-14, the state of Tennessee's measure of student test score growth ranked Jackson-Madison as a Level 1 district, the lowest-performing category in the state.

In the summer of 2013, principals from all 21 of the district's elementary and middle schools were briefed on the Evaluation Partnership Program and 14 agreed to participate in the study. We blocked the 14 schools in seven pairs based on level (elementary or middle) and student enrollment, and randomly assigned one to treatment within each pair.<sup>13</sup>

During the experiment year 2013-14, the 14 schools enrolled nearly 3,000 students in grades 4 through 8 with at least 136 teachers in mathematics and reading/language arts. We focus on grades 4 through 8 because these are the grades where we observe pre- and post-experiment student achievement measured by state tests. Descriptive information on the students and teachers is provided in Table 2.

---

<sup>12</sup> Participating principals and teachers mainly communicated directly with the research team. The exception is that lists of proposed matches were emailed to principals by the TNDOE. The research team prepared the match reports using de-identified data, where teachers were only known by randomly generated ID numbers. The TNDOE then replaced the random IDs with actual names and sent the reports to principals.

<sup>13</sup> This paper focuses on elementary and middle schools where we have data on student achievement test scores. We also recruited the district's high schools and pre-kindergarten schools. Two high schools and one pre-k school agreed to participate, were randomized, and participated in the program.

Interpreting the results of this experiment as causal effects rests largely on the success of randomly assigning schools to treatment and control conditions. Table 2 reports the traditional test of randomization, comparing pre-treatment characteristics of students, teachers, and teacher pairs. These tests include fixed effects for the randomization block pairs. We read the results as evidence of successful randomization. Most differences are substantively quite small.<sup>14</sup> Only 2 of the 20 characteristics show a statistically significant difference between treatment and control means: the proportion of English language learner students, and the difference in baseline observation scores for teacher pairs. Additionally, in Appendix Table A1, we check for covariate balance separately for teachers in each assigned program role: low-performing target teachers, high-performing partners, and teachers not assigned a role. The results are similar.

### *1.3 Data*

The Tennessee Department of Education provided two primary sources of data for this paper. First, we use the teacher evaluation micro-data described above for the years 2012-13 through 2014-15. The pre-randomization evaluation data are used in matching target and partner teachers; the post-randomization evaluation data are used as outcome measures of observed teacher job performance. Second, we use state administrative records for the years 2012-13 through 2014-15 that include (i) student scores from annual state standardized tests in math and reading/language-arts in grades 3 through 8, (ii) information on student demographics and special educational programs, (iii) records linking each student to her assigned teacher(s) for each subject each year, (iv) and information about teacher experience and prior performance. We standardize all test scores (mean zero, standard deviation one) within year-grade-subject cells

---

<sup>14</sup> Note that the teacher value-added scores are in teacher standard deviation units, not student standard deviations as is often the case. In student standard deviations the differences are roughly one-tenth to one-fifth the magnitudes in Tables 2 and A1.

using the statewide distribution (as opposed to the district specific mean and standard deviation). Additionally, TDOE provided data on where each teacher was working at the beginning of the 2015-16 school year.

## 2. Effects on Student Achievement

In this paper, we ask two primary empirical questions: First, did the treatment—pairing classroom teachers to work together on improving skills—benefit (harm) teacher performance? Second, if there were improvements, is there evidence that those improvements are the result of growth in teachers’ skills from peer learning?

Our primary measure of teacher performance is growth in student achievement test scores. Student learning that is measurable in standardized tests is, to be certain, only one aspect of a teacher’s job responsibilities. Nevertheless, existing empirical evidence suggests student test scores capture important variation in teacher performance (Jackson, Staiger, and Rockoff 2014 provide a review of the literature).

### 2.1 Average Treatment Effects on Student Achievement

Our most straightforward treatment effect estimates simply compare mean test scores in treatment and control schools. The treatment-control difference in means,  $\delta$ , is estimated by fitting the regression specification

$$A_{ijkt} = \delta EPP_{s(j)} + X_i\beta + \pi_{b(s)} + \varepsilon_{ijkt} \tag{1}$$

where  $A_{ijkt}$  is the end-of-year  $t$  (the experiment year) test score in subject  $k$  (math or reading) for student  $i$  assigned to teacher  $j$  in school  $s$ .<sup>15</sup> The treatment indicator,  $EPP_{s(j)}$ , varies only at the school-level  $s$ . All estimates in the paper include randomization block fixed effects,  $\pi_{b(s)}$ .

Throughout the paper, all statistical inferences account for error clustering within schools. We report p-values obtained by the wild cluster bootstrap-t method suggested by Cameron, Gelbach, and Miller (2008). This approach provides asymptotic refinement when the number of clusters is small, as in our setting with 14 clusters. Cameron and coauthors show that rejection rates can be as high as 10 percent for a nominally  $\alpha = 0.05$  test using conventional clustering methods; using the wild cluster bootstrap-t method rejection rates approach 5 percent. However, in this setting inference using conventional cluster-robust standard errors (as implemented in Stata) turns out to be quite similar.<sup>16</sup>

We report  $\hat{\delta}$  without and with additional covariates  $X_i$ , which are included to improve precision. The additional covariates include student  $i$ 's prior achievement—measured with the average of student  $i$ 's math and reading scores from the prior school year ( $t - 1$ )—as well as her gender, race/ethnicity, English language learner status, and special education status.<sup>17</sup> The vector  $X_i$  also includes a pre-experiment “value added” measure of teacher  $j$ 's contributions to student test scores in subject  $k$ ; this measure comes from that state's TVAAS system for 2011-12 and

---

<sup>15</sup> The student-teacher link records allow students to be linked to more than one teacher for a given subject, though three-quarters in our sample are linked to just one teacher. When a student has more than one teacher, the state assigns a “percent responsibility” to each teacher. When a student has two or more teachers, we include one observation for each student-by-teacher pairing and weight by the “percent responsibility.” But our results are robust to assigning students to the one teacher with the highest weight.

<sup>16</sup> Errors may also be correlated within a group of students taught by the same teacher. In our setting teachers are nested within schools. Thus, clustering errors at the school level is equivalent to clustering at both the school and teacher levels simultaneously (Cameron, Gelbach, and Miller 2011).

<sup>17</sup> While the field experiment occurred only in one district, we observe student test scores throughout the state. As a result we have very little missing data in  $X_i$ , for example, less than 4 percent of students are missing baseline achievement. When baseline achievement or another given covariate is missing, we replace it with a value of zero and include an indicator = 1 for all students missing the given covariate. Our results are robust to excluding these approximately 4 percent of students.



2012-13. Last,  $X_i$  includes grade-by-subject fixed effects and we allow the slope on prior achievement score to differ by subject and grade.

Estimates of the average treatment effect on student achievement,  $\delta$  in Equation 1, are reported in Table 3 Panel A. These are school-level intent-to-treat effects, and do not use any variation in assigned teacher roles or treatment take-up.<sup>18</sup> Students in treatment schools score about  $0.06\sigma$  (student standard deviations) higher than their peers in control schools, on average pooling math and reading outcomes. The difference is marginally statistically significant ( $p = 0.08$ ) when we control for pre-treatment covariates.

The estimated average treatment effect is somewhat larger when we focus on math achievement alone ( $0.07\sigma$ ;  $p = 0.040$ ) and somewhat smaller for reading and language arts alone ( $0.05\sigma$ ;  $p = 0.072$ ). This is consistent with the typical pattern in empirical research on schooling: most general interventions affect reading achievement less than math achievement. Additionally, the reading estimates are more sensitive to the inclusion of pre-treatment covariates, though we cannot reject that the reading estimates in Column 1 and 2 are equivalent statistically.

These average differences in student achievement can be interpreted as causal effects of treatment under the traditional experimental design assumption: At the beginning of the experiment, there was no difference in potential outcomes—student achievement growth, teacher or school performance, etc.—between treatment and control samples. This assumption rests

---

<sup>18</sup> We focus the paper's discussion on ITT estimates. Only one of the seven treatment schools did not participate at all in the Evaluation Partnership Program. The non-participating treatment school received the partnership match lists and program materials just as the other six schools did, but chose not to move forward. Thus the implied first-stage for a school-level TOT estimate would be about 0.86, suggesting the TOT estimates would be about 15 percent larger than the numbers reported in Table 3 Panel A.

largely on the success of random assignment; evidence in support of successful random assignment is presented in Table 2.

These positive average treatment effects are educationally and economically meaningful. Gains of  $0.06\sigma$  represent roughly one-third of a standard deviation in teacher performance, which is typically estimated at  $0.15\text{-}0.20\sigma$  in math and somewhat smaller in reading (Hanushek and Rivkin 2010, Jackson, Staiger, and Rockoff 2014). Put differently, the  $0.06\sigma$  difference is roughly equivalent to the difference between being assigned to a median teacher and a 63rd percentile teacher. Additionally, these average treatment effects are also roughly one-quarter the estimated gain from reducing class size by 30 percent in elementary grades (Kruger 1999), or the estimated gain from doubling the amount of class time middle and high school students spend in math (Taylor 2014, Cortes, Goodman, and Nomi 2015). However, unlike reducing class size or increasing class time, the current treatment—pairing classroom teachers to work together on improving skills—does not require any meaningful increase in teacher salary expenditures.

## *2.2 Treatment Effects for Target Teachers and Other Teachers*

Next we estimate treatment effects separately for teachers with different roles in the partnership program. The experiment was designed first to improve the job performance of low-performing “target” teachers; thus, the estimates in the top panel of Table 3 may mask important heterogeneity by role. Teachers were assigned to one of three roles: (i) low-performing target teachers, (ii) high-performing potential partner teachers, and (iii) all other teachers who were not assigned a role in partnerships.

Building on Specification 1, we estimate the following regression to test for differences by assigned teacher role:

$$\begin{aligned}
A_{ijkt} = & \delta^T (EPP_{s(j)} * Target_{j(i)}) + \delta^P (EPP_{s(j)} * Partner_{j(i)}) + \delta^N (EPP_{s(j)} * NoRole_{j(i)}) \\
& + \alpha^P Partner_{j(i)} + \alpha^N NoRole_{j(i)} + X_i \beta + \pi_{b(s)} + \varepsilon_{ijkt}
\end{aligned}
\tag{2}$$

where  $Target_j$ ,  $Partner_j$ , and  $NoRole_j$  are a set of mutually-exclusive and exhaustive indicator variables varying between teachers. This specification is algebraically equivalent to a more conventional specification with a main effect of treatment and interactions with two of the three teacher roles.  $Target_j = 1$  if teacher  $j$  was listed as a low-performing target teacher on the principal's Evaluation Partnership Program report. Recall that the report was created for both treatment and control schools, but only provided to treatment schools. Similarly,  $Partner_j = 1$  if teacher  $j$  was listed as a high-performing potential partner (either in the recommended pairings list or list of other potential partners). All other teachers have  $NoRole_j = 1$ , and were not listed on the principal's report. The estimates from Specification 2 are best interpreted as intent-to-treat because the role indicators are based on the original reports created by the research team, and not based on any post-randomization endogenous decisions by principals or teachers. All other details of estimation are the same as for Specification 1.

Treatment effects are largest for low-performing target teachers. As reported in the bottom panel of Table 3, treatment leads to test-score gains of  $0.12\sigma$  in target teachers' classrooms (compared to students of teachers who would have been target teachers in control schools, pooling math and reading). The estimates for high-performing partner teachers are positive, but much smaller and not statistically significant. Indeed the estimates for partner teachers are quite similar to the estimates for teachers who were not assigned a role in the program.

Again, these improvements are meaningful. A gain of  $0.12\sigma$  is roughly equivalent to the difference between being assigned to a median teacher instead of a bottom quartile teacher. A gain of  $0.12\sigma$  is also at least as large as the difference in performance between a novice teacher and a 5 to 10 year veteran (Rockoff 2004, Papay and Kraft 2015). However, the gains for target teachers are not necessarily substituting for experience gains. Later in the paper we examine heterogeneity of treatment effects by experience.

Moreover,  $0.12\sigma$  likely underestimates the effect of treatment on teachers who actually participate in the program. Table 4 reports estimates of program take-up by teacher role (the first stage results from a traditional 2SLS estimate of TOT).<sup>19</sup> Among treatment teachers assigned by the research team to the target role, 61 percent participated in the program in the target role, suggesting a treatment-on-the-treated estimate of about  $0.20\sigma$  ( $= 0.12/0.61$ ). These improvements are large but similar to the gains documented by Taylor and Tyler (2012) studying a program of evaluation and feedback, especially the gains Taylor and Tyler estimate for the ex-ante lowest performing teachers.

### *2.3 Treatment Effects After Two Years*

Last we estimate treatment effects one year after treatment. Table 5 shows estimates of Specifications 1 and 2 using test score data for students taught by treatment and control teachers in 2014-15.<sup>20</sup> The improvements in teacher performance persist, and perhaps grow, in the year after treatment. The estimated effect for low-performing target teachers is  $0.25\sigma$ , and marginally

---

<sup>19</sup> The sample and specification are identical to Table 3 Panel B Column 1 following Equation 2, except that the dependent variables are indicators for participation in a specific role.

<sup>20</sup> Some of the 136 teachers in the sample switched schools between 2013-14 and 2014-15. As long as a teacher is working in a Tennessee public school teaching grades 4-8 we include her in the sample. In these cases “school”,  $s$ , in Specifications 1 and 2 for purposes of estimation is her school during 2013-14. For all teachers “role” is her program role during 2013-14.

significant ( $p = 0.068$ , Column 1 Panel B). The estimated average treatment effect is also larger at  $0.11\sigma$ , though the p-value is 0.22 (Column 1 Panel A).

Two considerations are important in interpreting these second-year results. The first consideration is attrition. The sample for Table 5 Column 1 includes 96 of the original 136 treatment and control teachers. Some teachers were not working in Tennessee public schools; others switched to teaching non-tested grades or subjects.<sup>21</sup> To assess the potential influence of this attrition, Columns 2 and 3 provide bounds for the treatment effect following the approach described by Lee (2011).<sup>22</sup> Bounds for the average treatment effect are  $0.05\sigma$  and  $0.14\sigma$ . Bounds for the effect among low-performing teacher teachers are  $0.21\sigma$  and  $0.34\sigma$  (with p-values of 0.052 and 0.040 respectively). In short, the continued performance improvement among target teachers is robust to attrition, in part, perhaps, because attrition was relatively well balanced.

The second consideration is teachers who were target teachers in 2013-14 and again in 2014-15. In 2014-15 the treatment schools continued to use the Evaluation Partnership Program, and teacher roles and partnerships were re-assigned based on evaluations from 2013-14. Two teachers were recommended as target teachers in both years. These repeat target teachers may have an outsized influence on the estimated year-two effect if the effect is increasing in the dose of treatment, or if the improvement in performance only occurs in years when a teacher is

---

<sup>21</sup> We find no statistically significant treatment effects on the probability that teachers leaving their school or stop teaching altogether (at least in Tennessee public schools, results in Appendix Table A3).

<sup>22</sup> Our setting requires one modification to the standard Lee bounds procedure. We are concerned about attrition at the teacher level, but have outcome data at the student level. Thus determining which teachers to “trim” is not a simple function of an observed variable. To determine which teachers to trim we follow these steps: (i) Estimate the number of teachers who should be trimmed separately by role. The answer turns out to be trim one control target, one treatment partner, and three control teachers with no role. (ii) Repeatedly estimate  $\hat{\delta}^T$  in Specification 2 excluding one control target teacher in each iteration. Identify the control target teacher  $j$  whose exclusion minimizes  $\hat{\delta}^T$ . (iii) Repeat step (ii) excluding partner teachers to find the exclusion which minimizes  $\hat{\delta}^P$ . (iv) Repeat step (ii) excluding no-role teachers to find the exclusions which minimize  $\hat{\delta}^N$ . In this case each iteration excludes a unique combination of three control no-role teachers. (v) The five teachers identified in steps (ii)-(iv) are the teachers we trim to estimate the lower bound estimate. (vi) Repeat steps (ii)-(iv), this time maximizing each  $\hat{\delta}$ , to identify the five upper-bound teachers.

actively participating in a partnership. In Table 5 Column 4 we report estimates from a specification identical to Column 1, except that treatment effect for target teachers is interacted with an indicator = 1 if teacher  $j$  was a target teacher in both years. Column 4 does not provide causal estimates of one versus two years of treatment; being treated a second year is endogenous. Column 4 does provide evidence consistent with improvements in job performance for target teachers persist, perhaps grow, after treatment ends. The  $0.25\sigma$  improvement holds for teachers who were targets just one year, though it is not statistically significant.

### **3. Growth in Teachers' Skills and Other Potential Mechanisms**

While the previous section documents educationally meaningful and economically significant impacts, the average effect estimates do not shed light on the mechanisms through which the peer pairings influence productivity. We now move on to our second empirical question: Can the improvements in student learning—the positive average treatment effects—be attributed to growth in teachers' skills from peer learning? Or are other changes in behavior or effort behind the treatment effects?

In treatment schools, low-performing teachers were paired with a high-performing partner, and each pair was explicitly asked to work together on improving teaching skills. Thus our first hypothesized mechanism is teacher skill growth. Nevertheless, there are at least two other potential mechanisms that could contribute to the treatment effects: changes in teachers' motivation or effort, and changes in shared tasks (joint production) or resources. These three categories of mechanism are not mutually exclusive; all three could be contributing, to varying extents, to the average treatment effects. Jackson and Bruegmann (2009) describe how these three categories of teacher spillovers likely affect performance in a typical school context. Our

discussion of these three mechanisms focuses on how the treatment's pairing of teachers may have changed that typical context. The experimental setting and control counterfactual rule out many first-order features of these mechanisms as we highlight in the next paragraphs. Later we present empirical tests of predictions from these three hypothesized mechanisms. We cannot test all predictions empirically, especially in the case of the effort or motivation hypothesis, and thus the analysis is not definitive. However, the tests we can conduct are a step toward sorting out the relative contributions of different mechanisms.

While the stated purpose (to participating teachers) of the intervention was to improve teacher instructional skills, a second potential mechanism is that program participation could result in changes in teachers' motivation or effort. Asking a low-performing teacher to spend more time with a high-performing colleague, and talk together about performance explicitly, may have made her more optimistic or enthusiastic about work or made her more embarrassed about her poor performance. Similarly, treatment teachers may have felt more accountability to their new partner. These interactions may, in turn, lead to increased effort—either transitory increases in effort (e.g., motivated by specific accountability to ones' partner or direct monitoring by ones' partner) or lasting increases in effort (e.g., finding a new preferred equilibrium level of effort as a result of interacting with ones' partner). There is evidence for coworker effects on effort outside the education sector (for example, Mas and Moretti 2009). However, the scope for changes in optimism, embarrassment, or accountability is limited in the current experimental comparison: the treatment likely increased the degree of interaction with one coworker, but the mix of coworkers and typical coworker interactions were the same in treatment and control schools. Furthermore, teachers (and particularly low-performing teachers) in Tennessee already face fairly strong extrinsic incentives to increase their performance, as their schools are under

substantial test-based accountability pressures and as individuals they are at risk of losing their jobs for low evaluation ratings. The effect of any marginal accountability to one's partner is likely to be small.

The third potential mechanism to consider is changes in teachers' opportunities to share resources or production tasks. Teacher partnerships formed by the treatment program may have expanded to activities outside the original program scope. For example, teachers paired by the treatment may have been more likely to share existing lesson plans or cooperate in creating new lessons in ways that benefited productivity. Again, the control counterfactual rules out the typical, first-order resource or task sharing among teachers within a school. For example, teachers in treatment and control schools could exchange lessons and have group collaboration time; thus, to explain treatment effects, the shared production must be a specific result of the new partnership.

Our objective in this section is to present empirical evidence to help discriminate among these three potential mechanisms. We have already seen two important empirical results that are consistent with teacher skill growth. First, improvements in performance are largest for the low-performing target teachers; the average effect was perhaps entirely driven by target teacher improvements. Second, the performance improvements persisted, and perhaps even grew, in the year after treatment ended. However, the pattern of larger effects for target teachers could result from an asymmetric change in teachers' motivation or effort, or an asymmetric change in resource or tasks sharing as a result of treatment. Such changes might then persist after the partnership has formally ended. In the remainder of this section, we present several additional empirical tests. In short, we examine whether treatment effects vary with the characteristics of



teacher pairings in ways that are most consistent with skill growth or with either of the other two potential mechanisms.

### *3.1 Learning Skills*

First, if low-performing target teachers did learn new skills from their high-performing partner, we would expect larger treatment effects when the partner teacher's specific skill strengths matched the target teacher's specific weaknesses. In Table 6 Columns 2 through 5 we test this prediction. (For ease of comparison, we reproduce our main results from Table 3 in Column 1.) In Column 2, for the target teachers, we interact the treatment indicator with the proportion of teacher  $j$ 's weak skill areas matched by her recommended partner's strong skill areas.<sup>23</sup> The "proportion skills matched" measure is based on the one-to-one pairings recommended in the original principal reports (in the spirit, again, of intent-to-treat estimates). We have standardized the "proportion skills matched" (mean 0, s.d. 1) for comparison with other measures of pair characteristics in Table 6. Consistent with peer learning, the coefficient is positive and statistically significant ( $p = 0.052$ ). Student achievement gains are larger in pairs where the high-performing teacher is better suited to teach new skills to her low-performing partner.

Moreover, the result for "proportion skills matched" is not driven by target teachers who simply have more weak skills to match on. In Column 3 we replace "proportion skills matched" with "number of skills attempted to match", and in Column 4 we include both simultaneously. The target teacher's number of weak skills does not itself predict a larger treatment effect; the

---

<sup>23</sup> Recall that proposed pairings were determined algorithmically based on matches in 19 specific skill areas measured in each teacher's prior evaluation micro-data. We count up the number of skill areas in which there is a match: the target teacher has a score less than 3 and the recommended partner has a score of 4 or greater (see Section 1). Then, we divide the number of matches by the number of areas in which the target teacher scored less than 3.

effect is larger only when weak skills are matched. By contrast, if attention to low skill scores, by the program or by the partnership, generated a change in teacher effort, we would have expected a positive coefficient on “number of skills attempted to match.”

In Column 5 we replace the continuous proportion matched with an indicator = 1 if the proportion skills matched is above the pair median (about 0.5). This less-parametric approach also shows larger performance improvements when target and partner teachers are better matched on skills, and the estimates are more precise. In fact, the treatment effects appear concentrated among target teachers who were better matched to partners with relevant skills to share. By contrast, if low-performing target teachers’ motivation, effort, or joint production behavior changed, we would likely see positive treatment effects even when there are few or no strength-to-weakness skill matches. This is apparently not the case. The estimated treatment effect when the proportion skills matched is below median (Column 5 Row 1) is positive ( $0.036\sigma$ ;  $p = 0.340$ ), but not statistically significant and similar in size to the point estimates for partner and no assigned role teachers.

Evidence for skill growth is also apparent when we examine other aspects of teacher performance. To this point we have focused on performance as measured by student test score growth. Table 7 reports treatment effects for a second measure of teacher performance: evaluation scores from direct, in-class observations of teaching practices. (These are observation scores from teachers’ formal evaluations described in Section 1.) In Column 1 the outcome is teacher  $j$ ’s average score across all 19 skills and all classroom observations.<sup>24</sup> All outcome variables are standardized (mean 0, s.d. 1). As measured in direct observations, target teachers’

---

<sup>24</sup> As described in Section 1, teachers are scored on 19 skills multiple times per year. We first calculate an average score for each skill, standardize the skill scores, and then average the skill scores to obtain the overall average. We use all post-randomization scores from 2013-14 and 2014-15. We regress each outcome on a treatment indicator, randomization block fixed effects, and controls for teacher experience to improve precision.

skills improve on average by a marginally significant 0.05 teacher standard deviations. The coefficients for partner teachers are similar in magnitude, but estimates for both partner and no-role teachers have very wide (implicit) confidence intervals.

Consistent with the skill growth mechanism, notably, target teachers' improvements are concentrated in skills where there was a match—a match between the target's weak skills and her partner's strong skills. We examine treatment effects separately (Column 2) for the subset of skills where there was a match—the target scored less than 3 and her partner scored 4 or higher—and (Column 3) the subset of skills where the target was weak—she scored less than 3—but there was no match. We also report (Column 4) the effect for skill areas where we did not attempt to match because the target scored 3 or higher.<sup>25</sup> In matched skills target teachers improved 0.27 teacher standard deviations. But, by contrast, target teachers improved much less, if at all, in weak skill areas where there was no match with a partner strength. In summary, performance gains are larger in skill areas where the high-performing teacher is better suited to teach new skills to her low-performing partner.<sup>26</sup>

### *3.2 Motivation or Effort*

The evidence presented above is largely consistent with teachers building skills, rather than simply increasing motivation or effort. Here, we present additional evidence to differentiate between these mechanisms. If changes in effort contributed to the treatment effects, we might expect the treatment effect to be positively correlated with the size of the gap between the target

---

<sup>25</sup> For any individual teacher these three categories are mutually exclusive. Each of the 19 skills must belong to one and only one of matched, no matched, or not attempted to match.

<sup>26</sup> One note of caution in interpreting these classroom observation results. Many of the observations are conducted by the school principal who, in treatment schools, was certainly aware of the program and its goals of improved practices and improved evaluation scores. Treatment principals may have, consciously or unconsciously, inflated observations scores to recognize participation in the program; or, alternatively, principals may have been more critical or more aware of low-performance as a result of the program. However, to produce the results in Table 7, the principals would have had to differentially score target teachers on the matched skills specifically.

teacher's baseline job performance and her partner's baseline performance. For example, Mas and Moretti (2009) find that low-performing grocery check-out clerks increase their work effort when a higher-performing peer works at the same time and can observe the low-performer's effort; but high-performers are not affected by other high-performers. We test this prediction in Table 6 Columns 6-9, interacting the treatment indicator with the difference in prior-year skill measures. Note that while a positive correlation would be consistent with motivation, it might also be consistent with peer learning. A larger gap in skills may indicate a pair where the high-performer has more things to teach her partner.

Looking across the estimates in these columns, we do not find evidence of such a positive correlation. In fact, the imprecise point estimates are all negative. In Column 6, we examine differences in prior-year classroom observation scores (partner teacher minus target teacher). These observation scores are the rubric scores gathered in the formal evaluation process described in Section 1. The point estimate is small, negative, and not statistically significant. In Column 7 we include interactions for both difference in observation scores (possibly indicative of effort) and proportion skills matched (indicative of learning). Both measures have been standardized (mean 0, s.d. 1 throughout Table 6) to facilitate comparisons like Column 7. If anything, the coefficient on proportion skills matched is more positive than in Column 2, and the coefficient on observation scores more negative than in Column 6.

In Columns 8 and 9 we use an alternative measure of the gap between the target and partner baseline performance: the difference in prior "value added scores".<sup>27</sup> Value added scores are designed to measure each individual teacher's contribution to student test score growth. The

---

<sup>27</sup> We use value added scores provided by the state's TVAAS evaluation system. We calculate an overall value added score for each teacher by averaging all her subject-by-year value added scores from 2011-12 and 2012-13.

pattern for difference in value added is quite similar to the pattern for difference in observation scores.

### *3.3 Sharing Resources or Tasks*

The third mechanism category concerns new opportunities to share productive resources or job tasks. Being paired with another teacher may foster willingness or opportunities to collaborate at work in ways that improve performance even if the target teacher's skills do not change. Under this hypothesis we would expect the treatment effect to be greater when teacher pairs teach the same grade-level or subject area. The assumption motivating this test is that, even absent the treatment, shared production activities are easier or higher-return when teachers teach the same grade-level or subject. Sharing lesson plans is a simple concrete example: for a 6th grade mathematics teacher, sharing plans with another 6th grade math teacher would be more helpful than with an 8th grade teacher or a social studies teacher. To test this mechanism, in Column 10, we interact the target treatment effect with an indicator = 1 if the target and partner teachers are both teaching the subject  $k$ . In Column 11, we interact the effect with an indicator = 1 if the target and partner are both teaching the same grade level(s).<sup>28</sup> In both cases the point estimate is positive, which would be consistent with this third mechanism. But in both cases the estimates are small and (implicit) confidence intervals are quite wide—we could not reject large negative effects inconsistent with this mechanism. Furthermore, the estimated treatment effects for teachers in non-matched grades and subjects remains large, statistically significant, and of approximately the same magnitude as in Column 1. And, as above, this effect could derive in

---

<sup>28</sup> In practice teachers often teach more than one grade level, especially in middle and high school. We calculate the student-weighted average of grade level for each teacher. The indicator is = 1 if the difference in that average grade level is less than one.

part because peer learning is facilitated by shared subject or grade level. In short, we do not find strong evidence for a shared production mechanism.

To summarize the results across the three mechanism categories, first, we find evidence consistent with the hypothesis that low-performing target teachers learn new skills from their high-performing partner. Namely, improvements in performance, as measured by student test scores, are concentrated among target teachers; improvements are larger when more deficient skills are matched by the partner's skills; and improvements persist, perhaps grow, in the year following treatment. Additionally, improvements, as measured in classroom observations, are concentrated in matched skills. By contrast, predictions for other mechanisms are not borne out in our data. We do not find evidence consistent with a motivation or effort hypothesis, nor a shared resources or task hypothesis. Our empirical tests are partial and the three mechanisms are not mutually exclusive, so we cannot rule out any of these mechanisms. However, the available evidence suggests skill growth accounts for part of, perhaps much of, the average treatment effect.

#### **4. Discussion and Conclusion**

In this paper we study on-the-job learning among classroom teachers, especially learning skills from coworkers, in a field experiment. We document meaningful improvements in job performance among treatment teachers, and the patterns of performance gains suggest teachers learned job skills from their colleagues. The gains are empirical evidence of the intuitive, long-theorized benefits of coworkers in human capital development.

The relatively low-performing teachers targeted by our intervention—and ultimately their students—benefited substantially from partnering with a higher-performing colleague at their

school. Target teachers' performance improved 0.12 student standard deviations in the year of treatment, and perhaps double that much in the following year. These performance improvements were larger when teacher partnerships were better matched on strong and weak skills, but otherwise we find little evidence of heterogeneity.

Our estimates are consistent with prior evidence of learning from coworkers among teachers. Jackson and Bruegmann (2009) find that when a teacher begins working with higher-performing colleagues her own performance improves as a result. A one standard deviation increase in peer performance, as measured by prior contributions to student test score growth, generates a  $0.03\text{-}0.04\sigma$  improvement in own performance, also measured with current student test scores. Importantly, Jackson and Bruegmann define "working with" as teaching in the same grade and school as the peer; the peer work we study is more direct, and the effects are larger. Taylor and Tyler (2012) find that teacher performance improves  $0.05\sigma$  during a school year in which a teacher from a different school and the school principal conduct classroom-observation-based evaluations and subsequently provides feedback. Both Jackson and Bruegmann (2009) and Taylor and Tyler (2012) report that the gains in performance are sustained into the future; indeed in the latter case the effects grow from  $0.05\sigma$  during the peer evaluation year to  $0.10\sigma$  in the years after the peer evaluation year. We similarly find that performance improvements persist, perhaps grow, in the year following treatment.

Particularly notable, the treatment effects for target teachers appear to hold for both experienced and inexperienced teachers. In other words, the treatment is not simply substituting for the returns to experience. If anything, the treatment effect was  $0.03\sigma$  larger for target teachers with more than ten years of experience, though the result is far from statistically significant (Appendix Table A2). Performance growth for experienced teachers is rare in the literature, but

this result is consistent with results in Taylor and Tyler (2012). In Appendix Table A2 we examine several potential sources of heterogeneity, including teaching experience, prior performance, school level, and the difference in experience between target and partner. None show a statistically significant relationship to the treatment effect.

The experiment and results suggest practical alternatives to formal on-the-job training, especially in professional occupations. The contrast in approaches is particularly strong in the case of teachers. Formal courses, called “professional development,” are today the primary approach to on-the-job training for public school teachers. Collectively K-12 schools spend about \$18 billion per year on professional development courses, of which \$3 billion is paid to external providers (Gates Foundation 2014); the average teacher spends at least 20 hours each year in “professional development”.<sup>29</sup> Despite the substantial commitment of resources, the empirical evidence suggests little effect on teacher performance (see reviews by Jackson, Rockoff, and Staiger 2014, and Yoon et al. 2007). Similarly, public school systems spend tremendous resources paying for teachers’ graduate tuition, and paying higher salaries once teachers’ obtain their graduate degree. There is limited evidence that such degrees significantly improve teacher effectiveness (Jackson, Rockoff, and Staiger 2014).

By contrast, the one-on-one personalized approach to on-the-job training we study in this paper is apparently much more successful and much less costly. The primary marginal cost for treatment schools was the time that target and partner teachers allocated to working with each other. The opportunity cost of that time is important, although it could plausibly substitute for the time teachers would have spent in formal professional development courses. The cost is low in part because the high-performing coworker provides the teaching expertise. There could be

---

<sup>29</sup> Author’s calculation from the Schools and Staffing Survey 2011-12.



additional costs if higher-performing partners substitute away from other activities, especially attention to their own students. However, we do not find evidence of reduced performance among the higher-performing partners. If anything, the high-performers may have also benefited from participating in the pairings.

One important contextual feature of the experiment is the formal teacher evaluation system. All teachers in our study—treatment and control, target and partner and no role—are subject to Tennessee’s new formal performance evaluation system. Teacher pairs were identified based on prior evaluation results, and teacher pairs were encouraged, in part, to work on improving evaluation results. These connections to formal evaluation likely influenced principals’ and teachers’ willingness to participate, and the nature of their participation. For these reasons we think this study has contributions for the small, still-developing literature on how evaluation affects teacher performance (Taylor and Tyler 2012, Steinberg and Sartain 2015, Bergman and Hill 2015). One additional result on this subject comes from a survey of teachers at the end of the experiment. Teachers were asked a series of questions to measure their attitude toward formal evaluation, for example, “I have a favorable impression of the teacher evaluation system” rated on a six point agree/disagree scale.<sup>30</sup> Judging from survey responses, teachers in treatment schools left with more favorable opinions of evaluation: attitudes about evaluation were 0.23 standard deviations more positive, as measured by a composite of the four survey questions. However, survey response rates were lower in treatment schools (approximately 45

---

<sup>30</sup> The other three questions on this topic were: “In general, my colleagues have a favorable impression of the teacher evaluation system” and “I receive valuable feedback and guidance through teacher evaluation that helps me improve” both rated in the six point agree/disagree scale. And “What do you feel is the primary purpose of the teacher evaluation system? To help teachers improve. To rate teachers. Some of both.” Surveys were collected by the authors working directly with participating schools.

To create a composite score for evaluation attitudes we conducted a factor analysis of these questions and use the predicted first factor as our dependent variable. The factor analysis inputs were the three agree/disagree responses and separate binary indicators for “To help teachers improve” and “To rate teachers”. The first factor explains nearly 100 percent of the variation in responses to these four questions.

percent versus 66 percent), and thus this result should be interpreted with caution. If treatment suppressed responses from teachers with negative opinions, then the treatment effect on attitudes could easily be negative, but the empirical direction of any non-response bias is not clear.

The teacher job performance improvements documented in this paper suggest learning from colleagues is at least as valuable as formal on-the-job training or the gains from experience in developing teaching skills. Indeed peer learning may be a key contributor to the oft-cited estimates of returns to experience in teaching. Most practically, the treatment and results suggest promising ideas for managing the sizable teacher workforce.

## References

- Angrist, J. & Lavy, V. (2001). Does Teacher Training Affect Pupil Learning? Evidence from Matched Comparisons in Jerusalem Public Schools. *Journal of Labor Economics*, 19(2), 343–69.
- Azoulay, P., Zivin, J. S., & Wang, J. (2010). Superstar Extinction. *Quarterly Journal of Economics*, 125(2), 549-589.
- Bandiera, O., Barankay, I., & Rasul, I. (2005). Social Preferences and the Response to Incentives: Evidence from Personnel Data. *Quarterly Journal of Economics*, 120(3), 917-962.
- Battu, H., Belfield, C. R., & Sloane, P. J. (2003). Human Capital Spillovers within the Workplace: Evidence for Great Britain. *Oxford Bull Econ & Stats Oxford Bulletin of Economics and Statistics*, 65(5), 575-594.
- Becker, G. S. (1962). Investment in Human Capital: A Theoretical Analysis. *Journal of Political Economy*, 70(5 part 2), 9-49.
- Bergman, P., & Hill, M. J. (2015). The Effects of Making Performance Information Public: Evidence From Los Angeles Teachers and a Regression Discontinuity Design. CESifo Working Paper 5383.
- Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-Based Improvements for Inference with Clustered Errors. *Review of Economics and Statistics*, 90(3), 414-427.
- Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2011). Robust Inference with Multiway Clustering. *Journal of Business and Economic Statistics*, 29(2), 238-249.
- Chetty, R., Friedman, J. N., & Rockoff, J. E. (2014). Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood. *American Economic Review*, 104(9), 2633-2679.
- Cortes, K. E., Goodman, J. S., & Nomi, T. (2015). Intensive Math Instruction and Educational Attainment: Long-Run Impacts of Double-Dose Algebra. *Journal of Human Resources*, 50(1), 108-158.
- Danielson, C. (1996). *Enhancing Professional Practice: A Framework for Teaching*. Alexandria, VA: Association for Supervision and Curriculum Development.
- Dee, T. S., & Wyckoff, J. (2015). Incentives, Selection, and Teacher Performance: Evidence from IMPACT. *Journal of Policy Analysis and Management*, 34(2), 267-297.
- Gates Foundation. (2014). *Teachers Know Best: Teachers Views on Professional Development*. Seattle: Bill & Melinda Gates Foundation.

- Hamilton, B. H., Nickerson, J. A., & Owan, H. (2003). Team Incentives and Worker Heterogeneity: An Empirical Analysis of the Impact of Teams on Productivity and Participation. *Journal of Political Economy*, 111(3), 465-497.
- Hanushek, E. A., & Rivkin, S. G. (2010). Generalizations about Using Value-Added Measures of Teacher Quality. *American Economic Review*, 100(2), 267-271.
- Harris, D., & Sass, T. (2011). Teacher Training, Teacher Quality, and Student Achievement. *Journal of Public Economics*, 95, 798-812.
- Ichino, A., & Maggi, G. (2000). Work Environment and Individual Background: Explaining Regional Shirking Differentials in a Large Italian Firm. *Quarterly Journal of Economics*, 115(3), 1057-1090.
- Jackson, C. K., & Bruegmann, E. (2009). Teaching Students and Teaching Each Other: The Importance of Peer Learning for Teachers. *American Economic Journal: Applied Economics*, 1(4), 85-108.
- Jackson, C. K., Rockoff, J. E., & Staiger, D. O. (2014). Teacher Effects and Teacher-Related Policies. *Annu. Rev. Econ. Annual Review of Economics*, 6(1), 801-825.
- Krueger, J. (1999). Experimental Estimates of Education Production Functions. *Quarterly Journal of Economics*, 114(2), 497-532.
- Lucas, R. E. (2988). On the Mechanics of Economic Development. *Journal of Monetary Economics*, 22(1), 3-42.
- Marshall, A. (1890). *Principles of Economics*. London: Macmillan.
- Mas, A., & Moretti, E. (2009). Peers at Work. *American Economic Review*, 99(1), 112-145.
- Moretti, E. (2004). Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions. *American Economic Review*, 94(3), 656-690.
- Kraft, M.A. & Papay, J.P. (2014). Do Supportive Professional Environments Promote Teacher Development? Explaining Heterogeneity in Returns to Teaching Experience. *Educational Evaluation and Policy Analysis*, 36(4), 476-500.
- Papay, J. P. & Kraft, M. A. (2015). Productivity Returns to Experience in the Teacher Labor Market: Methodological Challenges and New Evidence on Long-Term Career Improvement. *Journal of Public Economics*, forthcoming.
- Rockoff, J. E. (2004). The Impact of Individual Teachers on Student Achievement: Evidence from Panel Data. *American Economic Review*, 94(2), 247-252.

- Rockoff, J. E. (2008). Does Mentoring Reduce Turnover and Improve Skills of New Employees? Evidence from Teachers in New York City. Columbia Business School Working Paper.
- Sacerdote, B. I. (2011). Peer Effects in Education: How Might They Work, How Big Are They and How Much Do We Know Thus Far? In *Handbook of Economics of Education Volume 3*, Hanushek, E. A., Machin, S., & Woessmann, L. eds. Amsterdam: North Holland.
- Steinberg, M. P., & Sartain, L. S. (2015). Does Teacher Evaluation Improve School Performance? Experimental Evidence from Chicago's Excellence in Teaching Pilot. *Education Finance and Policy*, forthcoming Winter 2015.
- Taylor, E. S. (2014). Spending More of the School Day in Math Class: Evidence from a Regression Discontinuity in Middle School. *Journal of Public Economics*, 117, 162-181.
- Taylor, E. S., & Tyler, J. H. (2012). The Effect of Evaluation on Teacher Performance. *American Economic Review*, 102(7), 3628-3651.
- U.S. Census Bureau. (2015). *Public Education Finances: 2013*. Educational Finance Branch G13-ASPEF. Washington, D.C.: United States Census Bureau
- Weisberg, D., Sexton, S., Mulhern, J., & Keeling, D. (2009). *The Widget Effect: Our National Failure to Acknowledge and Act on Teacher Effectiveness*. New York: The New Teacher Project.
- Yoon, K. S., Duncan, T., Lee, S. W., Scarloss, B., & Shapley, K. (2007). Reviewing the Evidence on How Teacher Professional Development Affects Student Achievement: Issues & Answers Report, REL 2007–No. 033. Washington, DC: US Department of Education, Institute of Education Sciences.

<b>Questioning</b>		
<b>Significantly Above Expectations (5)</b>	<b>At Expectations (3)</b>	<b>Significantly Below Expectations (1)</b>
<ul style="list-style-type: none"> <li>• Teacher questions are varied and high-quality, providing a balanced mix of question types:               <ul style="list-style-type: none"> <li>○ knowledge and comprehension;</li> <li>○ application and analysis; and</li> <li>○ creation and evaluation.</li> </ul> </li> <li>• Questions require students to regularly cite evidence throughout lesson.</li> <li>• Questions are consistently purposeful and coherent.</li> <li>• A high frequency of questions is asked.</li> <li>• Questions are consistently sequenced with attention to the instructional goals.</li> <li>• Questions regularly require active responses (e.g., whole class signaling, choral responses, written and shared responses, or group and individual answers).</li> <li>• Wait time (3-5 seconds) is consistently provided.</li> <li>• The teacher calls on volunteers and non-volunteers, and a balance of students based on ability and sex.</li> <li>• Students generate questions that lead to further inquiry and self-directed learning.</li> <li>• Questions regularly assess and advance student understanding.</li> <li>• When text is involved, majority of questions are text based.</li> </ul>	<ul style="list-style-type: none"> <li>• Teacher questions are varied and high-quality providing for some, but not all, question types:               <ul style="list-style-type: none"> <li>○ knowledge and comprehension;</li> <li>○ application and analysis; and</li> <li>○ creation and evaluation.</li> </ul> </li> <li>• Questions usually require students to cite evidence.</li> <li>• Questions are usually purposeful and coherent.</li> <li>• A moderate frequency of questions asked.</li> <li>• Questions are sometimes sequenced with attention to the instructional goals.</li> <li>• Questions sometimes require active responses (e.g., whole class signaling, choral responses, or group and individual answers).</li> <li>• Wait time is sometimes provided.</li> <li>• The teacher calls on volunteers and non-volunteers, and a balance of students based on ability and sex.</li> <li>• When text is involved, majority of questions are text based</li> </ul>	<ul style="list-style-type: none"> <li>• Teacher questions are inconsistent in quality and include few question types:               <ul style="list-style-type: none"> <li>○ knowledge and comprehension;</li> <li>○ application and analysis; and</li> <li>○ creation and evaluation.</li> </ul> </li> <li>• Questions are random and lack coherence.</li> <li>• A low frequency of questions is asked.</li> <li>• Questions are rarely sequenced with attention to the instructional goals.</li> <li>• Questions rarely require active responses (e.g., whole class signaling, choral responses, or group and individual answers).</li> <li>• Wait time is inconsistently provided.</li> <li>• The teacher mostly calls on volunteers and high-ability students.</li> </ul>

Figure 1—Example from TEAM rubric, “Questioning” skills

Teacher Name	Potential Match Teacher Name	Instructional Plans (IP)	Student Work (SW)	Assessment (AS)	Expectations (EX)	Managing Student Behavior (MSB)	Environment (ENV)	Respectful Culture (RC)	Standards and Objectives (SO)	Motivating Students (MS)	Presenting Instructional Content (PIC)	Lesson Structure and Pacing (LS)	Activities and Materials (ACT)	Questioning (QU)	Academic Feedback (FEED)	Grouping Students (GRP)	Teacher Content Knowledge (TCK)	Teacher Knowledge of Students (TKS)	Thinking (TH)	Problem Solving (PS)
Jane Blue				o					o		o			o						o
	Jane Brown	x	x	x	x	x	x	x	x	x	x	x	x	x	x		x	x	x	x
	Jane Yellow	x	x	x	x	x	x	x	x	x	x	x	x	x		x	x	x		
	John Red	x	x	x	x	x	x	x	x	x			x				x	x		
	Jane Orange	x	x	x	x	x	x	x		x	x	x	x				x	x	x	
	John Pink	x		x	x	x	x	x		x	x	x	x		x		x	x		
John Green									o	o										o
	John Black	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
	Jane Yellow	x	x	x	x	x	x	x	x	x	x	x	x	x		x	x	x		
	Jane Orange	x	x	x	x	x	x	x		x	x	x	x				x	x	x	
	John Pink	x		x	x	x	x	x		x	x	x	x		x		x	x		
	Jane White	x	x		x	x	x	x		x	x	x			x		x	x		

Note: An 'x' indicates that the teacher had an average score of 4 or higher on that element, an 'o' indicates an average score of less than 3.

Figure 2—Sample report for school principals showing potential partner matches for target teachers

Table 1—Classroom observation evaluation scores characteristics,  
year prior to experiment

(A) Scores in 19 skill areas, and scores averaging across skill areas			
	Mean	St. Dev.	Proportion scoring < 3
	(1)	(2)	(3)
Overall average	3.66	(0.68)	0.18
Instruction average	3.56	(0.69)	0.21
Problem Solving	3.19	(0.78)	0.23
Thinking	3.34	(0.75)	0.16
Questioning	3.39	(0.74)	0.17
Grouping Students	3.43	(0.79)	0.16
Lesson Structure and Pacing	3.43	(0.89)	0.22
Academic Feedback	3.49	(0.79)	0.14
Activities and Materials	3.56	(0.80)	0.13
Standards and Objectives	3.59	(0.81)	0.15
Presenting Instructional Content	3.64	(0.86)	0.15
Teacher Knowledge of Students	3.75	(0.88)	0.13
Motivating Students	3.85	(0.81)	0.08
Teacher Content Knowledge	4.12	(0.78)	0.05
Planning average	3.62	(0.80)	0.17
Assessment	3.38	(0.91)	0.18
Student Work	3.66	(0.89)	0.10
Instructional Plans	3.83	(0.88)	0.09
Environment average	3.97	(0.82)	0.11
Expectations	3.81	(0.92)	0.10
Managing Student Behavior	3.94	(0.96)	0.09
Respectful Culture	4.06	(0.89)	0.07
Environment	4.07	(0.92)	0.06
(B) Teachers with skill scores < 3			
	Proportion	Mean	
	(4)	(5)	
One or more skill areas < 3	0.41		
Number of skills with score < 3		2.37	
Number of skills with score < 3   at least one < 3		5.84	

Note: Authors' calculations using 2012-13 data for treatment and control schools. Data for Tennessee as a whole are similar. Observation scores in natural units (1-5 scale).



Table 2—Student, teacher, and pair characteristics,  
and pre-treatment balance

	Cont. mean (st. dev.) (1)	Treat. mean (st. dev.) (2)	Diff = 0 p-value (3)
Student characteristics			
Baseline test scores			
Mathematics	0.078 (0.531)	0.036 (0.658)	0.356
Reading/language arts	0.065 (0.562)	0.038 (0.668)	0.576
Average	0.066 (0.541)	0.040 (0.646)	0.480
Female	0.490	0.488	0.952
Race/ethnicity			
White	0.333	0.300	0.724
African-American	0.610	0.594	0.828
Latino(a)	0.047	0.087	0.176
Other	0.010	0.018	0.160
English language learner	0.015	0.038	0.004
Special education	0.118	0.126	0.704
Retained in grade	0.001	0.002	0.300
Teacher characteristics			
Years of experience	12.151 (11.92)	13.800 (10.18)	0.220
Baseline job performance			
Value-added	-0.023 (0.720)	0.112 (0.783)	0.472
Classroom observation score	3.656 (0.595)	3.846 (0.551)	0.364
Teacher pair characteristics			
Proportion skills matched	0.566	0.599	0.756
Difference in			
Value-added	0.123 (0.561)	0.688 (0.840)	0.172
Observation score	1.157 (0.412)	1.011 (0.440)	0.040
Years of experience	2.319 (14.79)	3.764 (12.11)	0.708
Teach same grade(s)	0.386	0.487	0.436
Teach same subject(s)	0.794	0.938	0.388

Note: Means and standard deviations net of randomization block fixed effects. Baseline student test scores and baseline teacher value-added standardized (mean 0 s.d. 1) using the statewide distributions for students and teachers respectively. Observation scores in natural units (1-5 scale). Column 3 reports wild cluster bootstrap-t p-values (Cameron, Gelbach, and Miller 2008) with 500 replications.

Table 3—Treatment effect on student achievement (intent to treat)

	Treatment effect		Observations	
	(1)	(2)	Student (3)	Teacher (4)
<b>(A) Average treatment effect</b>				
Math and reading	0.055 [0.224]	0.056 [0.080]	2947	136
Math	0.092 [0.032]	0.067 [0.040]	2874	86
Reading	0.018 [0.648]	0.052 [0.072]	2637	88
<b>(B) Treatment effect by teacher role</b>				
Math and reading				
Low-performing target teachers	0.082 [0.024]	0.123 [0.000]	2947	136
High-performing partner teachers	0.039 [0.532]	0.029 [0.252]		
No assigned role	0.013 [0.824]	0.029 [0.468]		
Pre-experiment covariates		√		

Note: Panel A: Each cell is an estimate from a separate regression. The dependent variable is a student test score, standardized (mean 0, standard deviation 1) within subject by grade-level cells using the statewide student distribution. All regressions include randomization block fixed effects. The vector of pre-experiment covariates includes: (i) Baseline student achievement: the average of each student's prior year math and reading scores. Prior test-score slope is allowed to vary by outcome subject and grade-level. (ii) Teacher pre-experiment value-added in the outcome subject: the average of 2012-13 and 2011-12 TVASS scores. (iii) Indicators for student gender, race/ethnicity, English language learner status, special education status, and whether the student is repeating the grade. When baseline test scores or value-added are missing we set the value to zero and include an indicator = 1 for missing. If the student had two or more teachers in a given subject, we include one observation per teacher and weight each observation by the proportion of responsibility allocated by the state to the teacher. Three quarters of students had one teacher in a given subject.

Panel B: Each column reports estimates from a separate regression. Estimation is identical to Panel A, except that the single treatment indicator is replaced with three indicators: treatment \* target teacher, treatment \* partner teacher, and treatment \* no assigned role. The specification also includes main effects for teacher role (i.e., "partner" and "no assignment" with "target" the omitted category).

P-values in brackets for the test that the coefficient equals zero. P-values estimated using wild cluster (school) bootstrap-t methods (Cameron, Gelbach, and Miller 2008) with 500 replications

Table 4—Teacher participation

	Dep. var. = 1 if <i>participated as a...</i>	
	Target (1)	Partner (2)
Treatment * <i>assigned</i> role:		
low-performing target	0.608 [0.000]	0.154 [0.344]
high-performing partner	0.070 [0.140]	0.366 [0.000]
no assignment	0.073 [0.176]	0.014 [0.744]
F-statistic excluded instruments jointly zero	18.522	31.520

Note: Each column reports estimates from a separate LPM regression; specifically, the first stage regressions from 2SLS estimation where actual role is instrumented with assigned role. Estimation is identical to Table 3 Panel B Column 1, except that the dependent variables are indicators = 1 if we observe participation in the target or partner roles respectively. The sample includes 14 schools, 2,947 students, and 136 teachers.

P-values in brackets for the test that the coefficient equals zero. P-values estimated using wild cluster (school) bootstrap-t methods (Cameron, Gelbach, and Miller 2008) with 500 replications.

Table 5—Treatment effects on student achievement in year two  
 Dep. var. = student math and reading test scores

	(1)	Treatment effect bounds (Lee bounds)		(4)
		lower (2)	upper (3)	
<b>(A) Average treatment effect</b>				
Math and reading	0.106 [0.220]	0.051 [0.556]	0.141 [0.212]	
<b>(B) Treatment effect by teacher role</b>				
Math and reading				
Low-performing target teachers	0.252 [0.068]	0.207 [0.052]	0.340 [0.040]	0.252 [0.212]
Low-performing target teachers * target again in year two				-0.000 [0.988]
High-performing partner teachers	0.056 [0.684]	-0.025 [0.776]	0.072 [0.624]	0.056 [0.688]
No assigned role	0.013 [0.908]	0.008 [0.864]	0.047 [0.536]	0.013 [0.908]

Note: Each column within panels reports estimates from a separate regression. The dependent variable is a student test score from spring 2015, standardized (mean 0, s.d. 1) within subject by grade-level cells using the statewide distribution. All regressions include randomization block fixed effects, and the vector of pre-experiment covariates described in Table 3. Additionally, Panel B regressions include main effects for teacher role (i.e., "partner" and "no assignment" with "target" the omitted category). Columns 2 and 3 report lower and upper bounds on the treatment effect estimate following the approach suggested by Lee (2009). We trim at the teacher level (dropping all students assigned to the trimmed teachers) within groups defined by teacher role (target, partner, no role) as described in the text. Column 4 repeats the specification and sample in Column 1, but adds an interaction between the target treatment effect and an indicator = 1 if the teacher was a target in year two. Regressions are weighted as described in Table 3. The sample includes 14 schools, 2,876 students, and 96 teachers.

P-values in brackets for the test that the coefficient equals zero. P-values estimated using wild cluster (school) bootstrap-t methods (Cameron, Gelbach, and Miller 2008) with 500 replications.

Table 6—Treatment effect heterogeneity by assigned teacher role and teacher pair characteristics  
 Dep. var. = student math and reading test scores

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Treatment * low-performing target	0.123 [0.000]	0.133 [0.000]	0.121 [0.000]	0.128 [0.000]	0.036 [0.340]	0.124 [0.000]	0.138 [0.000]	0.124 [0.000]	0.139 [0.000]	0.112 [0.024]	0.113 [0.000]
Treatment * high-performing partner	0.029 [0.252]	0.030 [0.216]	0.029 [0.268]	0.030 [0.220]	0.029 [0.208]	0.029 [0.264]	0.031 [0.196]	0.029 [0.252]	0.030 [0.200]	0.028 [0.252]	0.027 [0.272]
Treatment * no assignment	0.029 [0.468]	0.027 [0.496]	0.028 [0.492]	0.026 [0.552]	0.028 [0.476]	0.029 [0.468]	0.027 [0.512]	0.029 [0.464]	0.028 [0.484]	0.029 [0.468]	0.030 [0.444]
Treatment * target											
* proportion skills matched (std)		0.056 [0.052]		0.066 [0.020]			0.061 [0.028]		0.067 [0.044]		
* number of skills attempted to match (std)			-0.016 [0.632]	-0.028 [0.344]							
* proportion skills matched, above median (binary)					0.156 [0.000]						
* difference in prior observation scores (std)						-0.015 [0.816]	-0.032 [0.480]				
* difference in prior value-added scores (std)								-0.014 [0.636]	-0.036 [0.456]		
* both teaching the same subject (binary)										0.036 [0.612]	
* both teaching the same grade (binary)											0.045 [0.308]

Note: Each column reports estimates from a separate regression. Column 1 is identical to Table 3 Panel B Column 2. For Columns 2-11, all details of estimation are the same as Column 1 (see notes for Table 3), except that we add various characteristics of the target teacher and teacher pair as additional right-hand-side variables, and interact those characteristics with the treatment indicator for target teachers. The interaction coefficients are shown above. In each case the regression also includes a main effect for characteristic shown above. A "skill match" occurs when the target teacher has a score below 3 in the skill and the assigned partner has a score of 4 or higher (19 skills possible). The denominator in "proportion skills matched" is the number of skills where the target has a low score. All "difference" measures are partner score minus target score. The "proportion", "difference", and "number" measures have been standardized (mean 0, s.d. 1) within the sample. "Both teaching the same subject" is an indicator = 1 if assigned partner and target both teach the subject of the outcome score (math or reading). "Both teaching the same grade" is an indicator = 1 if the difference in average grade level of the assigned partner's students and the average grade level of the target's students is less than one. When a pair characteristic is missing we set the value to zero and include an indicator = 1 for missing. Regressions are weighted as described in Table 3. The sample includes 14 schools, 2,947 students, and 136 teachers.

P-values in brackets for the test that the coefficient equals zero. P-values estimated using wild cluster (school) bootstrap-t methods (Cameron, Gelbach, and Miller 2008) with 500 replications.

Table 7—Treatment effect on teaching skills scored in classroom observations

	Dep. var. = Average of evaluation scores on...			
	All 19 skills	Skills matched	Skills <i>not</i> matched	Skills where no attempt was made to match
	(1)	(2)	(3)	(4)
Low-performing target teachers	0.050	0.267	0.034	-0.071
	[0.060]	[0.000]	[0.500]	[0.348]
High-performing partner teachers	0.048			
	[0.836]			
No assigned role	-0.062			
	[0.892]			

Note: Each cell reports an estimate from a separate regression. The dependent variable is a measure of observed teaching practices from formal classroom observations conducted as part of the teacher's performance evaluation (see text for more details). In Column 1 the outcome is the teacher's average score across all 19 skills. In Columns 2 through 4 the outcome is the target teacher's average score across a specific subset of skills described in the column headers. The skills which contribute to each subset differ from teacher to teacher. But for any given teacher the three columns are mutually exclusive and exhaustive. Evaluation scores contributing to the dependent variables are from the 2013-14 and 2014-15 school years. Dependent variables are standardized (mean 0, standard deviation 1), and each of the 19 skill scores was standardized before taking averages. All regressions include randomization block fixed effects, and controls for teacher experience (indicators for quartiles of teacher experience).

P-values in brackets for the test that the coefficient equals zero. P-values estimated using wild cluster (school) bootstrap-t methods (Cameron, Gelbach, and Miller 2008) with 500 replications.

Appendix Table A1—Pre-treatment balance by teacher role

	Treat. - cont. difference [p-value] by teacher's assigned role		
	Target	Partner	No role
	(1)	(2)	(3)
Teacher characteristics			
Years of experience	-0.073 [0.852]	1.497 [0.528]	2.462 [0.468]
Baseline job performance			
Value-added	-0.355 [0.220]	-0.025 [0.852]	0.472 [0.088]
Classroom observation score	0.189 [0.216]	0.243 [0.268]	-0.072 [0.900]

Note: Each cell reports a treatment minus control difference in means. The three estimates in each row come from a single regression. The dependent variable described by the row label. All regressions include randomization block fixed effects, and main effects for teacher role (i.e., "partner" and "no assignment" with "target" the omitted category). P-values in brackets for the test that the difference equals zero. P-values estimated using wild cluster (school) bootstrap-t methods (Cameron, Gelbach, and Miller 2008) with 500 replications.

Appendix Table A2—Treatment effect heterogeneity, additional teacher, partner, and pair characteristics  
 Dep. var. = student math and reading test scores

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment * low-performing target	0.057 [0.220]	0.126 [0.344]	0.136 [0.000]	0.116 [0.000]	0.187 [0.008]	0.118 [0.000]	0.130 [0.000]	0.095 [0.060]
Treatment * high-performing partner	0.028 [0.276]	0.027 [0.316]	0.029 [0.280]	0.029 [0.252]	0.028 [0.276]	0.030 [0.256]	0.050 [0.064]	0.032 [0.204]
Treatment * no assignment	0.030 [0.456]	0.033 [0.400]	0.029 [0.436]	0.031 [0.448]	0.022 [0.596]	0.028 [0.464]	0.040 [0.488]	0.025 [0.536]
Treatment * target								
* target's years of experience	0.005 [0.224]	-0.009 [0.728]						
* target's years of experience ^ 2		0.000 [0.428]						
* target has fewer than 10 years experience (binary)			-0.029 [0.692]					
* target's prior value-added scores (std)				-0.022 [0.532]				
* target's prior observation scores (std)					0.046 [0.464]			
* target is an elementary school teacher (binary, middle school omitted)						0.010 [0.864]		
* difference in years of experience (partner - target)							0.012 [0.848]	
* partner's years of experience								0.003 [0.340]

Note: Each column reports estimates from a separate regression. The dependent variable is a student test score, standardized (mean 0, s.d. 1) within subject by grade-level cells using the statewide distribution. All regressions include randomization block fixed effects, and the vector of pre-experiment covariates described in Table 3. Additionally, all regressions include main effects for teacher role (i.e., "partner" and "no assignment" with "target" the omitted category). In each specification we interact the treatment indicator with a characteristic of the target teacher, pair, or partner teacher. When a characteristic is missing we set the value to zero and include an indicator = 1 for missing. Regressions are weighted as described in Table 3. The sample includes 14 schools, 2,947 students, and 136 teachers.

P-values in brackets for the test that the coefficient equals zero. P-values estimated using wild cluster (school) bootstrap-t methods (Cameron, Gelbach, and Miller 2008) with 500 replications.



Appendix Table A3—Treatment effects on teacher job changes

	No longer teaching in the same school		No longer teaching in Tennessee	
	one year later	two years later	one year later	two years later
	(1)	(2)	(3)	(4)
<b>(A) Average treatment effects</b>				
All teachers	0.084 [0.184]	0.067 [0.384]	0.089 [0.012]	0.036 [0.556]
Control mean	0.277	0.415	0.092	0.246
<b>(B) Treatment effects by teacher role</b>				
Low-performing target teachers	0.038 [0.776]	0.010 [0.956]	0.073 [0.608]	0.028 [0.856]
Control mean	0.375	0.438	0.125	0.250
High-performing partner teachers	0.070 [0.668]	0.079 [0.652]	0.109 [0.324]	0.188 [0.216]
Control mean	0.240	0.400	0.080	0.200
No assigned role	0.144 [0.272]	0.087 [0.400]	0.085 [0.176]	-0.118 [0.288]
Control mean	0.250	0.417	0.083	0.292

Note: Each column within panels reports estimates from a separate LPM regression. The dependent variable is an indicator = 1 if the teacher is no longer teaching in the same school (or in the state of Tennessee) in 2014-15 the year after treatment in 2013-14 (or 2015-16 two years after treatment). All regressions include randomization block fixed effects. Additionally, all regressions in Panel B include main effects for teacher role (i.e., "partner" and "no assignment" with "target" the omitted category). The sample includes 14 schools, and 136 teachers.

P-values in brackets for the test that the coefficient equals zero. P-values estimated using wild cluster (school) bootstrap-t methods (Cameron, Gelbach, and Miller 2008) with 500 replications.