

The Design of Teacher Incentive Pay and Educational Outcomes: Evidence from the New York City Bonus Program

Sarena F. Goodman
Columbia University

Lesley J. Turner
Columbia University

May 2012*

Abstract

Teacher compensation schemes are often criticized for lacking a performance-based component. Proponents argue that teacher incentive pay can raise student achievement and stimulate system-wide innovation. We examine a group-based teacher incentive scheme implemented in New York City and investigate whether specific features of the program contributed to its ineffectiveness. Although overall the program had little effect on student achievement, we show that in schools where incentives to free-ride were weakest, the program led to small increases in math achievement. Our results underscore the importance of carefully considering the design of teacher incentive pay programs.

* Correspondence should be sent to ljt2110@columbia.edu. We are especially grateful to Jonah Rockoff for his thoughtful comments and advice. We also thank Todd Kumler, Bentley MacLeod, Ben Marx, Derek Neal, Petra Persson, Maya Rossin, Jesse Rothstein, Miguel Urquiola, Till Von Wachter, Reed Walker, and seminar participants at the Columbia applied microeconomics colloquium, AEFA annual meeting, Teacher's College Economics of Education Workshop, and the Harvard Kennedy School's Program on Education Policy and Governance's Merit Pay Conference for useful discussions and feedback. We are grateful to the New York City Department of Education for the data used in this paper.

1. Introduction

Teacher compensation schemes are often criticized for their lack of performance pay. In other sectors, incentive pay increases worker effort and output by aligning the interests of workers and employers, providing information about the most valued aspects of an employee's job, and motivating workers to provide costly effort (Gibbons, 1998; Lazear and Oyer, 2010). In this paper, we examine a group-based teacher incentive scheme implemented by the New York City Department of Education (DOE) and investigate whether specific features of the program contributed to its ineffectiveness.

In 2007, close to two hundred schools were randomly selected from a group of high-poverty schools.¹ These schools could earn school-wide bonuses by surpassing goals primarily based on student achievement. Successful schools would earn lump sum payments equal to \$3000 per union teacher (three to seven percent of annual teacher pay). Several independent studies show that the bonus program had little overall effect on either math or reading achievement (Springer and Winters 2009; Goodman and Turner 2010; Fryer 2011). We show that in schools where smaller groups of teachers were responsible for instructing tested students, the program led to small but significant increases in student achievement. Our finding is consistent with predictions that group-based incentives are diluted by the potential for free-riding when payments depend on actions of a large number of workers (Holmstrom 1982).

Several features of the educational sector complicate the design of teacher performance pay. First, performance pay is most effective when employers can measure worker output or when observable effort and productivity are closely aligned. Monitoring teachers is costly and measuring individual teachers' contributions to student achievement is difficult. Second, although education is a complex good and teachers must allocate their effort across several activities, teacher incentive pay is often linked to a single performance measure (e.g., student test scores), which may lead teachers to direct effort away from other beneficial classroom activities (Holmstrom and Milgrom 1991).² Despite these issues, studies from outside the United States demonstrate that teacher incentive pay can increase student achievement (e.g., Lavy 2002; Lavy 2009; Muralidharan and Sundararaman 2011).

¹ This experiment was designed and implemented by the New York City Department of Education and teachers' union, random assignment was conducted by Roland Fryer, and RAND performed the official evaluation.

² Teachers may also be induced to focus on narrow, exam-related basic skills, manipulate test scores, or focus on students whose performance contributes more towards goals (e.g., Jacob and Levitt 2003; Jacob 2005; Cullen and Reback 2006; Neal and Schanzenbach 2010).

Specific features of the NYC bonus program may have limited its effectiveness. First, the program linked incentive pay to school-wide performance goals. In theory, group incentive pay is most effective in the context of a joint production technology (Itoh, 1991). For instance, if an individual teacher's effort has positive impacts on the effort exerted by her peers (e.g., Jackson and Bruegmann 2009), group incentives may outperform individual incentives. Otherwise, relative to individual incentives, group incentives decrease individual returns to effort and will lead to free-riding unless workers monitor each other's effort.

We test for free-riding by allowing the bonus program's impacts to vary by the number of teachers with students who are tested (and therefore contribute to the probability that a school qualifies for the bonus award). To test for the importance of joint production and monitoring, we examine whether program impacts vary by the degree to which teachers report collaborating in lesson planning and instruction using a survey administered prior to program implementation. We show that the bonus program raised math achievement in schools with a small number of teachers with tested students, although these impacts are small (0.08 student-level standard deviations) and only marginally significant in the program's second year. We present suggestive evidence of positive program impacts in schools with a high degree of collaboration.

Second, teachers already faced negative incentives when the bonus program was implemented. In fall 2007, the DOE instituted a district-wide accountability system that imposed sanctions on schools that did not meet the same goals used in determining bonus receipt. Thus, estimated impacts of the bonus program represent the effect of teacher performance pay in schools already under accountability pressure. However, this may be the most appropriate context to examine, since many states have implemented accountability systems and all public school districts face pressure from No Child Left Behind provisions. Finally, we find no differences in the impacts of the bonus program when we compare schools under different degrees of accountability pressure, suggesting that our results are not solely driven by the dilution of incentives due to the accountability system (Goodman and Turner 2010).

Third, teachers' lack of understanding of the bonus program's complex goals may have limited its efficacy. Alternatively, since bonus awards were provided if a school's performance reached a set threshold, if thresholds were set too high or too low, a large number of teachers may have optimally responded by not changing their behavior (Neal 2011). However, the metrics used to determine bonus payments were the same goals used by the district-wide accountability

system and Rockoff and Turner (2010) show that negative incentives provided through this system increased student achievement.³

2. Data and Empirical Framework

Our analyses focus on schools classified as elementary, middle, and kindergarten through grade 8 (K-8) schools eligible for selection into the bonus program. A total of 181 schools were chosen to participate in the bonus program; 128 schools were placed in the control group.⁴ We use publicly available DOE data and measure academic achievement using average math and reading test scores in the 2006-07, 2007-08, and 2008-09 school years (hereafter 2007, 2008, and 2009).

We estimate the main effect of the bonus program using the following model:

$$(1) \quad y_{jt} = \delta D_{jt} + \mathbf{X}_{jt} \boldsymbol{\beta} + \varepsilon_{jt}$$

where y_{jt} is the outcome of interest for school j in year t , D_{jt} is an indicator selection into the bonus program's treatment group (regardless of whether the school ultimately participated), \mathbf{X}_{jt} is a vector of school characteristics, and ε_{jt} is an idiosyncratic error term.⁵ School observations are weighted by the number of tested students. With successful random assignment, D_{jt} is independent of omitted variables and $\hat{\delta}$ represents the casual impact of the bonus program.

3. Results

3.1 Group Bonuses and the Free-Rider Problem

Teachers should respond to the bonus program by increasing their effort until the expected marginal benefit is equal to the expected marginal cost. However, the probability that a treated school reaches its goal and receives a bonus primarily depends on students' performance on math and reading exams. Thus, the impact of an individual's teacher's effort on her expected

³ On a related note, a committee within each school had some discretion over how bonuses would be distributed. However, the distribution scheme was set *ex ante* and most schools chose equal or close to equal distributions.

⁴ A small number of experimental sample schools were excluded *prior* to random assignment. Moreover, two of the 181 schools originally assigned to the treatment group were moved to the control group prior to notification of their assignment; we classify these as treatment group schools. Treatment schools were eligible to earn bonuses if 55 percent full-time United Federal of Teachers staff voted in favor of participation. Twenty-five schools voted not to participate or withdrew from the program after voting. Finally, four schools that were originally assigned to the control group were allowed to vote and participate in the bonus program; we consider these control schools. Ultimately, 158 schools were eligible to earn bonus payments.

⁵ Covariates include the outcome measured in 2007, school type indicators (i.e., elementary, middle, or K-8), the percentage of students that are English Language Learners, special education, Title I free lunch recipients, and minorities, and performance under the NYC accountability system (school accountability scores and peer indices).

bonus is decreasing as the number of teachers with tested students increases.⁶ The diffusion of responsibility for test score gains across many teachers may dilute the incentives of the bonus scheme. Moreover, monitoring may be more difficult in schools with more teachers.

We test for evidence of free-riding by allowing treatment effects on math and reading scores to vary by the number of math and reading teachers, respectively. We only focus on teachers whose students take these exams, rather than the full set of teachers in a school, since only teachers with tested students contribute to the probability that a school earns its bonus.⁷ The first set of regressions in Table 1 show the main effect of the bonus program on math and reading achievement.⁸ We first add an interaction between the number of math/reading teachers (relative to the mean number of such teachers in the sample) and the treatment indicator (columns 2 and 5), and finally, interact treatment status with an indicator for schools in the bottom quartile of the number of teachers with tested students (approximately 10 or fewer teachers in elementary and K-8 schools and 5 or fewer in middle schools). We only present results from specifications that include covariates, however, results are similar when we exclude covariates or instrument for actual treatment with initial assignment.

We find evidence of free-riding. For schools at the bottom of the distribution of the number of teachers with tested students, we estimate a positive effect of the bonus program on math achievement in the first year of the program and a positive, but insignificant effect in the second year, although we cannot reject a test of equality of effects across years. In 2008, the bonus program resulted in a 3.2 point (0.08 student-level standard deviation) increase in math achievement.⁹

Group-based incentive pay may outperform individual incentives in the case of joint production. If the degree to which teachers work together varies across schools, the bonus program may have been effective in schools with a high level of cooperation between teachers.

⁶ Consider two extremes, a school with only one teacher with tested students and a school with an infinite number of these teachers. In the first case, the teacher will either respond to the program by increasing her effort to the expected level necessary to achieve the school's goal or not respond (if the size of the bonus is less than the cost of exerting this level of effort). In the second case, changes in a given teacher's effort do not affect the probability that the school receives the bonus and it will be optimal for teachers to not respond to the program.

⁷ On average, treatment and control group schools have 55 teachers in total, but only 16 teach tested students.

⁸ The small number of middle and K-8 schools that are missing information on the number of teachers with tested subjects are excluded.

⁹ Another implication of this finding is that, in schools with a large number of teachers with tested students, the bonus program had a negative impact on student achievement. One explanation is the bonus program crowded out teachers' intrinsic motivation and only in schools where incentives were not diluted by free-riding did the potential monetary rewards lead to increased teacher effort.

To proxy for the extent of joint production in a school, we construct a measure of school cohesiveness using teachers' answers to a set of five survey questions prior to the announcement of the bonus program.¹⁰ This measure may also incorporate the degree to which teachers are able to monitor their colleagues. We sum responses across survey questions and standardize the index so it has a mean of zero and standard deviation equal to one. Schools with high levels of cohesion are distinct from those with a small number of teachers with tested students.¹¹

Table 2 tests for heterogeneity in the impact of the bonus program by school cohesion. We first interact treatment with the linear index (columns 2 and 5) and then interact treatment with an indicator for schools with above average cohesion (columns 3 and 6). The point estimates for schools with below average cohesion are marginally significant and negative in both subjects and both years, while the interaction of treatment and the indicator for above average cohesion is significant, positive, and of greater magnitude. Results suggest that the bonus program may have had detrimental effects in schools with low levels of cohesion, and small positive effects on achievement in cohesive schools.

3.2 Teacher Effort

A primary motivation for performance-based pay is to provide teachers with incentives to increase effort devoted to raising student achievement. Although we do not directly observe teacher effort, we can measure teacher attendance, which may be correlated with effort decisions and contributes to student achievement (e.g., Miller, Murnane, and Willett 2008; Herrmann and Rockoff forthcoming). We measure teacher absences using aggregate statistics from individual teacher data and estimate models where the dependent variable is the average number of absences taken during the months when schools first learned of their eligibility for the bonus program and when the last exams were taken.¹² If teachers believe that their attendance can affect the probability of bonus receipt by raising student achievement, the program's impacts on absenteeism should be largest over this period.¹³ We only examine absences that teachers likely

¹⁰ These surveys were administered in spring 2007. Questions include: (1) the extent to which teachers report feeling supported by fellow teachers, (2) whether curriculum and instruction is aligned within and across school grades, (3) whether the principal involves teachers in decision making, (4) whether school leaders encourage collaboration, and (5) whether teachers collaborate to improve instruction. We exclude schools with a survey response rate under 10%.

¹¹ This index has a small, negative, and statistically insignificant correlation with the number of math and reading teachers in a school.

¹² We thank Jonah Rockoff for constructing these aggregate statistics for the purpose of this research.

¹³ In the first year of the program, schools learned of their eligibility in November while in the second year, eligibility was known in September. In both years, the last exams occurred in March. Results are robust to alternate definitions of the time period (e.g., November to March in the second year or September to March in the first year).

have some control over – those taken for illness and personal reasons.

Table 3 presents these results; each column within a panel contains the estimates from separate regressions. The first column examines the effect of the bonus program on absences across all teachers within a school and shows no measurable impact on overall attendance. Column 2 focuses on teachers with tested students, while the third and fourth columns follow the same approach as Table 2 and interact the treatment indicator with the number of teachers with tested students (column 3) or an indicator for whether a school falls in the bottom quartile of the number of such teachers (column 4).

Program impacts on attendance are not consistent across years. In the program's first year, for schools with a small number of teachers with tested students, attendance increased.¹⁴ Conversely, in the second year of the program, we find positive but insignificant impacts on absenteeism. Finally, we test whether the bonus program had heterogeneous impacts according to initial teacher effort. For instance, initially low effort (high absence) teachers may be the only group with the ability to respond through increasing attendance. Conversely, if *ex ante* high effort teachers believed that achieving the bonus program goals was a high probability event, they may have responded by reducing their effort. However, we find no evidence teacher absenteeism varies along this dimension (available upon request). In the United States, attendance may not be the dimension along which teachers respond to incentive pay.

4. Conclusions

In many sectors, performance-based pay enhances effort, output, and other desirable outcomes. Evidence from Israel and India suggests that properly structured teacher incentive pay programs can benefit students. However, despite substantial expenditures – over \$40 million in the program's first two years – the NYC bonus program did not raise student achievement. This paper discusses several features of the NYC bonus program that may have contributed to its ineffectiveness. We provide suggestive evidence that the group-based structure of the program may have been detrimental in the majority of schools where the number of teachers responsible for tested students is large. Conversely, the program improved math achievement in schools with fewer teachers responsible for tested students or a more cohesive group of teachers. A lack of monitoring as well as the diffusion of responsibility for test score gains among many teachers

¹⁴ However, impacts are only significant in schools at the 10th percentile in the distribution of number of teachers (results available upon request).

may have diluted the incentives of the opportunity to earn bonuses. Our results are consistent with the long-standing literature in economics on the importance of taking into consideration free-riding, joint production, and monitoring when designing incentive systems and suggest that a one-size-fits-all approach may not be the most effective when implementing incentive pay schemes within a school district.

Given that team-based incentives in other contexts resulted in student achievement gains, other features of the NYC program may have also contributed to its ineffectiveness. Neal (2011) suggests that results from economic theory offer valuable insights into optimal incentive design. For instance, an intervention in India utilized a piece-rate payment scheme: teachers or schools received bonus payments for incremental improvements in student achievement (Muralidharan and Sundararaman 2011). This avoids threshold effects of schemes like the NYC bonus program, which dilute incentives for teachers with a probability of bonus receipt approaches zero or one.

Even so, many challenges in designing effective teacher incentive schemes remain. Incentive pay programs that come about as a compromise between school districts and teachers unions' might contain incentives that are so diluted they are destined to fail. Finally, the extensive margin may be most important margin through which teacher pay can improve student achievement. Small-scale teacher incentive pay experiments cannot provide information concerning the general equilibrium effects of overall increase in teacher pay or movement towards performance-based compensation.

Currently, the U.S. government provides significant funding through the Race to the Top program. Eligibility for Race to the Top funding depends on districts' ability and willingness to link student achievement to individual teachers and use this data in teacher evaluations, but grants districts a great deal of discretion in designing performance pay systems. In 2010, 62 school districts and nonprofit groups received over \$400 million in funding from the federal Teacher Incentive Fund. Our results underscore the importance of the structure of performance pay in education. Policy innovations in this area should be carefully considered, taking into account personnel economics theory and research.

References

Cullen, Julie Berry and Randall Reback. 2006. Tinkering toward accolades: school gaming under a performance accountability system. In *Advances in applied microeconomics volume 14*:

- improving school accountability*, ed. Timothy J. Gronberg and Dennis W. Jansen. Oxford, UK: JAI Press.
- Fryer, Roland G. 2011. Teacher incentives and student achievement: evidence from New York City public schools. Working Paper no. 16850, National Bureau of Economic Research, Cambridge, MA.
- Gibbons, Robert. 1998. Incentives in organizations. *Journal of Economic Perspectives* 12, no. 4: 115–32.
- Goodman, Sarena and Lesley Turner. 2010. Teacher incentive pay and educational outcomes: evidence from the NYC bonus program. Working Paper no. 10-07, Program on Education Policy and Governance Working Paper Series, Cambridge, MA.
- Herrmann, Mariesa A. and Jonah E. Rockoff. Forthcoming. Worker absence and productivity: evidence from teaching. *Journal of Labor Economics*.
- Holmstrom, Bengt. 1982. Moral hazard in teams. *Bell Journal of Economics* 13, no. 2: 324–40.
- Holmstrom, Bengt and Paul Milgrom. 1991. Multitask principal-agent analyses: incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization* 7 (January): 24–52.
- Itoh, Hideshi. 1991. Incentives to help in multi-agent situations. *Econometrica* 59, no. 3: 611–36.
- Jackson, C. Kirabo and Elias Bruegmann. 2009. Teaching students and teaching each other: the importance of peer learning for teachers. *American Economic Journal: Applied Economics* 1, no. 4: 1–27.
- Jacob, Brian A. 2005. Accountability, incentives and behavior: evidence from school reform in Chicago. *Journal of Public Economics* 89: 761–96.
- Jacob, Brian A. and Steven D. Levitt. 2003. Rotten apples: an investigation of the prevalence and predictors of teacher cheating. *Quarterly Journal of Economics* 118, no. 3: 843–77.
- Lavy, Victor. 2002. Evaluating the effect of teachers' group performance incentives on pupil achievement. *The Journal of Political Economy* 110, no. 6: 1286–1317.
- . 2009. Performance pay and teachers' effort, productivity and grading ethics. *American Economic Review* 99, no. 5: 1979–2011.
- Lazear, Edward P. 2003. Teacher incentives. *Swedish Economic Policy Review* 10: 179–214.
- Lazear, Edward P. and Paul Oyer. Forthcoming. Personnel economics. In *Handbook of organizational economics*, ed. Robert Gibbons and D. John Roberts. Princeton, NJ: Princeton University Press.
- Miller, Raegen T. Richard J. Murnane, and John B. Willett. 2008. Do worker absences affect productivity? The case of teachers. *International Labour Review* 147, no. 1: 71–89.
- Muralidharan, Karthik and Venkatesh Sundararaman. 2011. Teacher performance pay: experimental evidence from India. *Journal of Political Economy* 119, no. 1: 39–77.

- Neal, Derek. 2011. The design of performance pay in education. In *Handbook of economics of education, volume 4*, ed. Eric Hanushek, Steve Machin, and Ludger Woessmann. Amsterdam: Elsevier.
- Neal, Derek and Diane W. Schanzenbach. 2010. Left behind by design: proficiency counts and test-based accountability. *Review of Economics and Statistics* 92, no. 2: 263-83.
- Rockoff, Jonah and Lesley J. Turner. 2010. Short-run impacts of accountability on school quality. *American Economic Journal: Economic Policy* 2, no. 4: 119–47.
- Springer, Matthew G. and Marcus A. Winters. 2009. The NYC teacher pay-for-performance program: early evidence from a randomized trial. Manhattan Institute Civic Report No. 56.

Table 1: Free-riding and the Impact of Teacher Incentives on Student Math and Reading Achievement

	Reading			Math		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Year 1: 2007-2008</i>						
Treatment	-0.372 (0.490)	0.046 (0.499)	-0.667 (0.519)	-0.871 (0.530)	-0.536 (0.568)	-1.445 (0.561)*
* Number of teachers (<i>mean = 0</i>)		-0.233 (0.089)**			-0.176 (0.097)+	
* First quartile of number of teachers			2.044 (1.575)			4.670 (1.483)**
Treatment effect: schools in first quartile			1.377 (1.481)			3.225 (1.395)*
Observations	300	300	300	301	301	301
<i>B. Year 2: 2008-2009</i>						
Treatment	-0.579 (0.539)	-0.395 (0.572)	-0.909 (0.556)	-1.297 (0.668)+	-0.979 (0.726)	-1.893 (0.689)**
* Number of teachers (<i>mean = 0</i>)		-0.126 (0.099)			-0.171 (0.144)	
* First quartile of number of teachers			2.122 (2.067)			4.826 (2.579)+
Treatment effect: schools in first quartile			1.213 (1.968)			2.933 (2.461)
Observations	294	294	294	294	294	294

Note: Each column within a panel denotes a separate regression; dependent variable: average math or reading test scores. Robust standard errors in parentheses. The first row in each panel displays the estimated impact of treatment group assignment; in columns 2 and 5, treatment group assignment is interacted with the (demeaned) number of teachers with tested students; in columns 3 and 6, treatment group assignment is interacted with an indicator for being a school in the lowest quartile of teachers with tested students. The number of math teachers for schools in the first quartile is less than or equal to: 10 (elementary and K-8 schools), 5 (middle schools); the number of reading teachers for schools in the first quartile is less than or equal to: 10 (elementary and K-8 schools), 6 (middle schools). The regressions are weighted by number of tested students in each school. Schools with no information on teachers with tested students are dropped. See text for a description of additional controls included in regressions.

+ $p < 0.10$

* $p < 0.05$

** $p < 0.01$

Table 2: School Cohesion and the Impact of Teacher Incentives on Student Math and Reading Achievement

	Reading			Math		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>A. Year 1: 2007-2008</i>						
Treatment	-0.328 (0.494)	-0.091 (0.515)	-0.908 (0.591)	-0.628 (0.530)	-0.270 (0.551)	-1.264 (0.674)+
* Cohesion index		0.283 (0.547)			0.766 (0.619)	
* Above average cohesion index			1.840 (0.975)+			1.962 (1.131)+
Treatment effect: schools with above average cohesion			0.932 (0.789)			0.698 (0.891)
Observations	300	300	300	301	301	301
<i>B. Year 2: 2008-2009</i>						
Treatment	-0.544 (0.540)	-0.328 (0.562)	-1.139 (0.666)+	-1.118 (0.666)+	-0.669 (0.679)	-2.266 (0.869)**
* Cohesion index		0.361 (0.626)			1.105 (0.839)	
* Above average cohesion index			1.850 (1.093)+			3.347 (1.446)*
Treatment effect: schools with above average cohesion			0.710 (0.868)			1.081 (1.105)
Observations	296	296	296	297	297	297

Note: Each column within a panel denotes a separate regression; dependent variable: average math or reading test scores. Robust standard errors in parentheses. The first row in each panel displays the estimated impact of treatment group assignment; in columns 2 and 5, treatment group assignment is interacted with the teacher cohesion index (mean = 0, sd = 1 across all NYC schools); in columns 3 and 6, treatment group assignment is interacted with an indicator for having a cohesion index greater than zero. The regressions are weighted by number of tested students in each school. Schools with teacher survey response rate below 10 percent are dropped. See text for a description of additional controls included in regressions.

+ $p < 0.10$
* $p < 0.05$
** $p < 0.01$

Table 3: The Impact of Teacher Incentives on Teacher Absences Due to Personal and Sick Leave

	<u>All Teachers</u>		<u>Teachers of Tested Students</u>	
	(1)	(2)	(3)	(4)
<i>A. Year 1: 2007-2008</i>				
Treatment	0.001 (0.091)	-0.158 (0.146)	-0.217 (0.148)	-0.156 (0.163)
* Number of teachers (<i>mean = 0</i>)			0.013 (0.022)	
* First quartile of number of teachers				-0.236 (0.390)
Treatment effect: schools in first quartile				-0.391 (0.352)
Observations	301	301	301	301
<i>B. Year 2: 2008-2009</i>				
Treatment	0.045 (0.119)	0.151 (0.175)	0.203 (0.192)	0.161 (0.200)
* Number of teachers (<i>mean = 0</i>)			0.005 (0.032)	
* First quartile of number of teachers				0.158 (0.621)
Treatment effect: schools in first quartile				0.319 (0.576)
Observations	294	294	294	294

Note: Each column within a panel denotes a separate regression; dependent variable: average absences/teacher taken for personal or sick leave between November and March (Panel A) or September and March (Panel B). Robust standard errors in parentheses. The first row in each panel displays the estimated impact of treatment group assignment on absences for all teachers. The second row in each panel displays the estimated impact of treatment group assignment on absences for teachers with tested students. In column 3, treatment group assignment is interacted with the (demeaned) number of teachers with tested students; in column 4, treatment group assignment is interacted with an indicator for being a school in the lowest quartile of teachers with tested students (see Table 1 notes). Schools with no information on teachers with tested students are dropped. See text for a description of additional controls included in regressions.