Economics of Education Review 57 (2017) 20–30

Contents lists available at ScienceDirect

Economics of Education Review

journal homepage: www.elsevier.com/locate/econedurev

Workload and teacher absence

Ben Ost[a,*], Jeffrey C. Schiman[b]

[a] Department of Economics, University of Illinois at Chicago, USA
[b] Department of Finance and Economics, Georgia Southern University, USA

A R T I C L E I N F O

Article history:
Received 1 April 2015
Revised 25 October 2016
Accepted 9 January 2017
Available online 12 January 2017

JEL classification:
I20
J45

Keywords:
Teacher absences
School practices

A B S T R A C T

We investigate the determinants of teacher absences both within and across schools. We find that teachers generally respond to increased workload by decreasing their rate of absence. Teachers are less likely to be absent when they are teaching larger classes, have new grade assignments or have fewer years of experience. Moreover, we show that when teachers change schools, their absence rate quickly gravitates towards the mean absence rate of their new school, suggesting that school-level factors are an important determinant of absence rates. Finally, we show that the inverse relationship between workload and absence may lead researchers to underestimate the ceteris paribus effect of certain teacher inputs. We illustrate this point in the context of estimating the effect of teacher experience on test scores and show that controlling for absence rates increases the estimated returns to experience by approximately 10%.

1. Introduction

Although there is a large qualitative literature exploring how working conditions affect teachers' attitudes and self-reported effort (Blase, 1982; Neves de Jesus & Lens, 2005; Reyes & Imber, 1992; Timperley & Robinson, 2000), there is little research relating measurable teacher effort behaviors to their workload. Theoretically, teachers may increase effort when workloads increase, or they may become discouraged and decrease effort. Distinguishing between these possibilities is difficult, largely because few datasets include measures of either teacher effort or workload. In this article, we construct teacher-specific measures of workload to study the extent to which school- and individual-level factors influence teacher absence. We view teacher absences as a proxy for effort, but since we lack direct data on actual effort, the most conservative interpretation of our results is that they document the effect of workload on absences. Although many absences are taken for reasons such as illness, our results strongly suggest that some absences are discretionary.[1]

Using administrative longitudinal data on elementary teachers in North Carolina, we find that teachers are less likely to be absent when they teach larger classes, when they have less experience, and when they are assigned to teach a different grade from the previous year.[2] The result that teacher absences decrease following grade switches is based on a model that controls for teacher

* Corresponding author.
E-mail addresses: bost@uic.edu (B. Ost), jschiman@georgiasouthern.edu (J.C. Schiman).

http://dx.doi.org/10.1016/j.econedurev.2017.01.002

[1] Taking discretionary absences may proxy for low teacher effort, but it could also be correlated with high teacher effort. In particular, it is possible that teachers who put in little effort each day rarely take absences because they are not in need of a break. We cannot rule out this possibility, though we view it as less likely since dedication has been shown to be negatively correlated with absences in other contexts (Gaziel, 2004).

[2] The absences/experience profile has been previously documented in Clotfelter et al. (2009), Miller et al. (2008) and Hansen (2009). The relationship between class size and absences and grade switching and absences are both novel results. Although the absences/experience profile is interesting, the fact that absences rise with experience could be attributable to many different factors and is not necessarily evidence of a workload/absences relationship. For example, teachers may develop relationships with administrators and gain political capital as they gain experience that allows them to take more leave.

experience, so it suggests that teachers who have recently taught their current grade assignment are more likely to be absent than other teachers who have the same level of experience. These results are based on models that include teacher and school fixed effects, so fixed differences across teachers or schools cannot explain the findings.

Before showing the relationship between workload and absences, we first document several interesting patterns that generally suggest that teacher absences are malleable. For example, we show that among teachers who switch schools, teachers' absence rates gravitate to the average rate of absence in their new school. The timing of the change in absence rates coincides perfectly with the movement across schools and there is no evidence of differential trends prior to the school switch. We view this as evidence that school-specific factors play a role in determining absences.

We are unaware of any study that demonstrates the effect of workload on absences, but several studies document the effect of teacher absences on student learning and how absence behavior varies with teacher incentives. For instance, Herrmann and Rockoff (2012) find that teacher absences have negative effects on student performance. Similarly, Clotfelter, Ladd, and Vigdor (2009) find that teacher absences hurt student performance and that teachers are much less likely to take absences when there is a direct financial penalty for doing so. Jacob (2013) finds that teachers take fewer absences following a policy that reduced job security for teachers in Chicago. Finally, in the developing context, there are several randomized experiments that demonstrate that policies that reduce teacher absences increase test scores (Banerjee & Duflo, 2006; Duflo, Hanna, & Ryan, 2012). In addition to the consequences for student performance, Joseph, Waymack, and Zielaski (2014) estimate that teacher absences impose direct financial costs on districts of $1800 per teacher per year.

In addition to providing a better understanding of how teachers allocate effort, this study illustrates a theoretical point noted in Todd and Wolpin (2003) regarding dynamic optimization of inputs in an education production function. In their model, when one input changes exogenously, other inputs either increase or decrease depending on whether they are complements or substitutes in production. Importantly, dynamic optimization implies that the total effect of an input may differ from the ceteris paribus effect of an input. We explore empirically whether there is a difference between the ceteris paribus returns to experience and the total returns to experience. Since we show that teachers are more likely to be absent as they gain experience, past work estimating the returns to experience includes the effect of absences in their estimates of the experience profile (Rivkin, Hanushek, & Kain, 2005; Rockoff, 2004). Researchers are often interested in the total experience effect, including the effect operating through increased absences. However, researchers are also interested in understanding the ceteris paribus effect of experience to identify the educational production function itself.[3] We show that in practice, controlling for the rate of absence increases the estimated returns to experience by approximately 10%.

2. Data and institutional environment

Our data come from the North Carolina Education Research Data Center and contain information on every public school teacher and student in North Carolina from 1995 to 2007. Our focus is on elementary teachers in PK-5. For teachers, we learn years of past teaching experience, whether or not the teacher was just reassigned to teach a new grade, the size of their primary class, and teacher demographics like race, education, and Praxis test score. The Praxis exam is a required test administered to teachers that aims to assess subject and pedagogical skills. We focus on elementary teachers because grade switching is most well defined in the elementary context where teachers generally teach a single grade. We measure absences based on the number of sick or personal days taken during the course of the academic year, as opposed to including absences taken for reasons such as school sanctioned professional development. We trim the top 1% of absences to ensure that outliers do not unduly influence results, but in practice, our estimates are very similar without dropping outliers.

In North Carolina, absence policies are set by law at the state level and specify the maximum number of sick and personal days that teachers are allowed to take.[4] Teachers are allowed to take 10 unpenalized sick days per year and 2 personal days per year (accrued on a monthly basis). If a teacher exhausts her unpenalized absences, she can take up to 20 additional days of extended sick time at a cost of $50 per day.[5] Important for our research, unused days can be carried forward across years with no limit. As such, for teachers who do not take all allotted sick and personal days each year, the number of days that they have available rises with experience. Although this suggests that teachers with many years of experience likely have more allowed absences than teachers with few years of experience, the policy does not actually generate a mechanical relationship between experience and absences. If a teacher takes all allowed absences every year, she will have the same number of absences available in her later years compared to her first year. If a teacher is unconstrained by the absence policy in her first year, this teacher is likely to be unconstrained in future years, and so the higher maximum absences will not necessarily result in more absences taken.

The first column of Table 1 shows characteristics for the sample used to study how workload relates to absence behavior. In the second column, we show these same characteristics for a sample of teachers who switch schools

[3] There are likely other endogenous inputs related to experience, so our approach does not capture a true ceteris paribus effect. We simply show how the estimated returns differ when holding fixed one of the endogenous inputs, namely teacher absences.

[4] In addition to sick and personal days, teachers are granted vacation days, but these cannot be taken while students are present. Furthermore, in our data, we cannot distinguish between vacation days and school holidays and so we focus on just sick and personal days. Results are very similar when we include vacation days in the analysis as well.

[5] Longer leaves such as maternity leave are governed by a separate set of rules.

Table 1.
Descriptive statistics.

	Workload sample		School switching sample	
	Mean	S.D.	Mean	S.D.
Absences	8.55	6.73	8.45	6.51
Peer absences	8.37	1.57	8.17	1.61
Total past experience	13.5	9.37	11.1	9.27
Grade change $t-1$ to t	0.14	0.34	0.35	0.48
Switched school between t and $t+1$	0.15	0.36	0.30	0.46
Teacher has M.A. or higher	0.27	0.44	0.27	0.44
Teacher is female	0.96	0.18	0.95	0.21
Teacher is black	0.13	0.34	0.18	0.39
Teacher is Hispanic	0.0026	0.051	0.0042	0.064
Teacher's Praxis Test score	−0.048	0.85	−0.073	0.84
Class size	21.9	4.43	21.6	5.13
Behaviorally impaired students	0.066	0.29	0.073	0.35
Learnings disabled students	0.79	1.20	0.81	1.20
Share of minority students	0.37	0.26	0.42	0.27
Title 1 status	0.39	0.49	0.40	0.49
School enrollment/100	5.81	2.01	5.93	2.19
Observations	218,860		32,096	

Notes. Here we present descriptive statistics for the sample used to estimate the relationship between workload and teacher absences in Table 5 ("Workload sample") and the sample used in the school-switching analysis in Table 3 ("School switching sample").

during our sample time period. This second sample is used for our analysis of how school factors influence absence-taking behavior. In both samples, the teachers take approximately 8.5 absences per year. Approximately 14% of teachers switch grades across years, but this fraction is 35% for the set of teachers who switch schools. The higher rate of switching grades for the latter sample reflects the fact that teachers are likely to switch grades when they switch schools. In our empirical analysis we are careful to account for this fact when examining the effect of grade switching. Approximately 95% of both samples are female, reflecting the particularly strong occupational segregation of elementary (and in particular early elementary) teachers. Approximately 13% of teachers are black and very few North Carolina teachers are Hispanic. Although teachers who switch schools are slightly more likely to be black, and have slightly fewer years of experience, they do not appear to be dramatically different than the average teacher.

In Table 2, we describe how the number of allowed absences evolve with experience. As previously described, teachers accrue 12 days of absence per year and these can accumulate without limit. Although we have no direct data on each teacher's accrued absences, we can calculate this directly from the data for the subset of teachers who enter our sample with zero years of experience. The first column of Table 2 shows that the number of absences taken rises substantially in the first three years of teaching. The second column of Table 2 shows that average allowed absences rise sharply with experience such that by the third year, teachers have an average of 24 days of absence accrued. Average allowed absences continue to accrue throughout the experience profile, but at a slower rate.

To further explore the relationship between absences and experience, Fig. 1 shows the distribution of absences for the entire sample, split by years of experience. If absences increase with experience strictly because maximum allowed absences increase, we would expect to see a rightward shift of the upper tail of the distribution for more experienced teachers, but the lower tail of the distribution would remain unchanged. Instead, Fig. 1 shows that more experienced teachers take more absences than inexperienced teachers at all points in the distribution, represented by a rightward shift in the *entire* distribution of absences. Importantly, we view the results of Table 2 and Fig. 1 as purely descriptive, since these patterns include any effects from survival bias or unobserved heterogeneity.

Table 2.
Average absences used and remaining by years of experience.

Years of past experience	Absences taken	Allowed absences
0	4.7	12.0
1	7.2	19.3
2	8.2	24.2
3	8.9	28.2
4	9.6	31.4
5	9.7	33.4
6	9.98	36.1
7	10.03	38.1
8	9.7	40.3
9	9.5	43.2
10	9.6	44.7
11	8.97	46.8
12	9.4	52.7

Notes. To calculate allowed absences, we focus on the subset of teachers who enter our sample with zero years of prior of experience. We then follow the absence accrual policy as discussed in the main text.

3. The determinants of teacher absences

3.1. School-level factors

We first examine the extent to which teacher absences are influenced by school-level factors by considering how absence rates change when teachers change schools. If absences primarily reflect teacher fixed characteristics, then one would expect that switching schools would have little effect on the number of absences. For this analysis, we restrict the sample to teachers who switch schools, and examine how a teacher's change in absences from their old to their new school is related to the difference between the average rate of absence in their old and their new school. We refer generically to each teacher's old school as "school 1" and their new school as "school 2". Specifically, for a teacher who switches between school 1 and school 2 from year $t-1$ to year t, we estimate variants of

$$y_{jt} - y_{j(t-1)} = \beta(\overline{y_{s_2}} - \overline{y_{s_1}}) + \epsilon_{jt} \quad (1)$$

where y_{jt} is the absence rate for teacher j in year t, and $\overline{y_{s_k}}$ is the average absence rate in school k. Because the dependent variable in this specification is the *change* in absence taking within a teacher, fixed teacher characteristics will not bias estimates. That said, there are several empirical concerns in estimating the above specification. First, it is possible that teachers who are increasing their absence rate over time choose to move to schools with higher absence rates and teachers who are decreasing their rate of absence over time choose to move to schools with

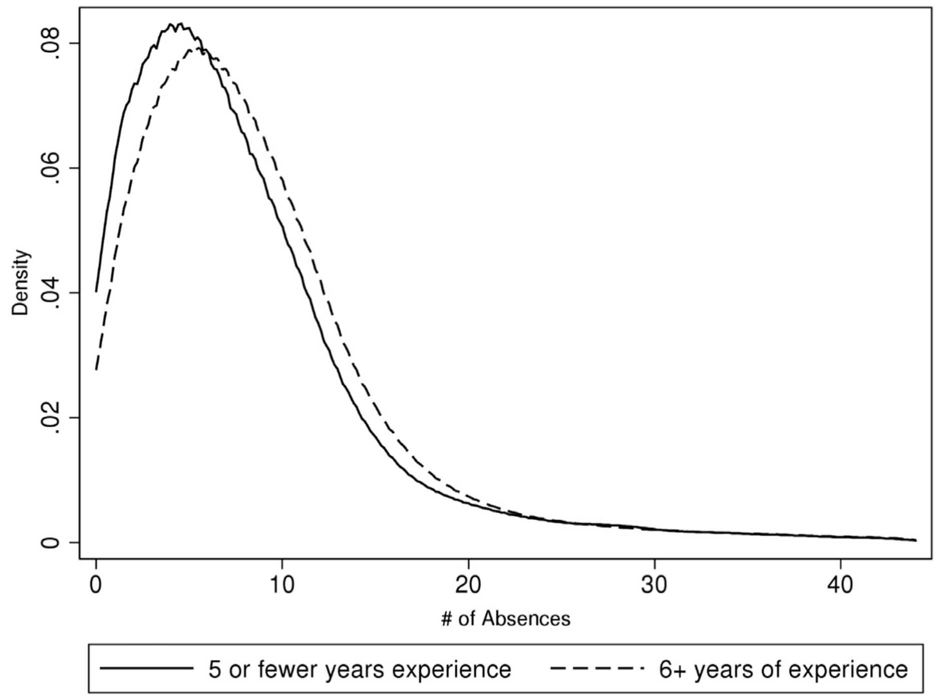

Fig. 1. Distribution of absences by experience.
Notes. We calculate the density using the "Workload sample" described in Table 1.

lower absence rates. Second, it is possible that teacher *j*'s absence-taking behavior could affect the school average absence rate. Third, school-by-time varying factors could bias estimates if they affect both the teacher and the school average absence rates. For example, if there is a contagious illness in school 2 in year *t*, this may lead to an increase in absence for teacher *j* and it will also make it more likely that school 2's average absence rate is high relative to school 1's.

To address the first concern regarding differential trends across teachers, we examine whether the change in school-level average absences predicts changes in the absence rate in the years prior to the move. To address the second and third concerns, we construct a jackknife-style measure of school-level absence in which we exclude any year that teacher *j* was present in a school when calculating the average school absence rate. This implies that for both school 1 and school 2, the average absence rate excludes teacher *j* (addressing the concern that the switching teacher influences the average at each school). The measure also excludes all years when teacher *j* was at the school, ensuring that no school-by-year shock could cause both teacher *j*'s absence and the jackknife school-level absence rate. If there is an illness in year $t+1$ in school 2, this will act to increase teacher *j's* absences, but it will not affect our measure of school 2's absence rate since we exclude years in which teacher *j* is at the school. We use this measure of school average absence both for our main estimation and for all the specification checks. For compactness, in the text that follows we simply refer to "school-level average absence" but this measure always excludes years in which teacher *j* was present in the school.

The first row and column of Table 3 shows our estimate of the effect of the change in school-level average absence on the change in own teacher absences. We find that a 1-day increase in school average absence rates leads to a 0.14-day increase in own absence rates. Columns 2–4 show that the change in school-level absence does not predict changes in absence in the 3 years prior to the move. This suggests that estimates are unlikely to be biased by differential trends. Columns 5–7 show that the change in school-level absence does not predict changes in own absence rates in the years after the school switch. This suggests that the entire adjustment to own absence occurs in the year immediately following the school move. The change in school-average absence could theoretically have been gradual and so we do not view columns 5–7 as a falsification test analogous to columns 2–4. That said, columns 5–7 do provide a partial specification check on our results as it would have been concerning if we had found that the $t-1$ to t effect reversed in $t+1$ or if we had found that the post-switch changes in absences were monotonically increasing each year.

In rows 2–4, we modify our estimation strategy in various ways to address potential sources of bias. Experience has the potential to bias our estimates if the change in school-level average absence *and* the changes in own absence are both related to experience. Although this type of issue would likely have revealed itself in the pre-trend test, we confirm that this issue does not bias our analysis by re-estimating our baseline specification while controlling non-parametrically for experience indicators. Row 2 shows that results are generally robust to this control.

Table 3.
Changes in own absences in response to changes in the absence environment.

	(1) Transition	(2) Before leaving	(3)	(4)	(5) After leaving	(6)	(7)
	$t-1$ to t	$t-4$ to $t-3$	$t-3$ to $t-2$	$t-2$ to $t-1$	t to $t+1$	$t+1$ to $t+2$	$t+2$ to $t+3$
Overall estimate	0.142***	−0.027	−0.017	−0.001	0.013	−0.006	0.004
	(0.023)	(0.033)	(0.029)	(0.026)	(0.025)	(0.029)	(0.032)
Observations	32,096	13,324	17,740	23,437	25,133	19,904	15,557
With experience controls	0.121***	−0.030	−0.017	−0.004	0.007	−0.012	0.002
	(0.023)	(0.033)	(0.028)	(0.025)	(0.025)	(0.029)	(0.032)
Observations	31,898	13,290	17,680	23,338	25,019	19,836	15,521
With balance restriction	0.131***	−0.067	0.027	−0.047	−0.026	−0.005	−0.005
	(0.048)	(0.044)	(0.046)	(0.048)	(0.043)	(0.046)	(0.053)
Observations	5929	5929	5929	5929	5929	5929	5929
Teacher FE	0.089**	−0.003	−0.035	0.011	−0.010	−0.004	−0.023
	(0.041)	(0.058)	(0.050)	(0.045)	(0.047)	(0.056)	(0.070)
Observations	32,096	13,324	17,740	23,437	25,133	19,904	15,557

Notes. The estimates in Table 3 correspond to the approach described in Eq. (1) in Section 3.1. Heteroskedastic-robust standard errors are in parentheses.
* $p < 0.10$.
** $p < 0.05$.
*** $p < 0.01$.

As is apparent from the changing sample sizes across columns, there is a different composition of teachers in each column in Table 3. This is because teachers who switch schools before their third year in a school do not have an absence rate in year $t-3$. Similarly, teachers who leave school 2 shortly after arriving or who switch schools towards the end of our sample period may not have all of the post years defined. The sample differences across columns should not bias any of the individual estimates, but they could make comparisons across columns misleading. For example, if teachers present for just t and $t-1$ were the only type of teacher with responsive absences, we could mistakenly conclude that the year of the switch is more important than other years, when it would actually just reflect compositional differences across columns. To address this issue, we restrict the sample to be a balanced panel of teachers for whom we can observe 3 years before and after their school switch. This is a substantial sample restriction because it removes all school switches that happen near the end or beginning of our sample as well as removing teachers with short tenure at either their first or second school. Row 3 of Table 3 shows that the estimate for absence change from $t-1$ to t is very similar in the restricted sample. Though estimates are somewhat noisier, we do not find statistically significant estimates for either the before or after periods and the signs of the estimates vary across specifications. Broadly, we view the balanced panel analysis as suggesting that our prior estimates are not driven by changing sample composition across columns. The balanced panel is not our preferred specification, given that the estimates are only relevant for a particular type of teacher and are substantially less precise than our baseline specification.

Our estimates cannot be biased by fixed teacher characteristics, but it is possible that certain teachers have a propensity to *change* their absence level following changing schools. Though it seems unlikely that this would occur in such a way that would be completely uncorrelated with the pre-period changes, we are able to provide direct evidence on this possibility by including a teacher fixed effect. The teacher fixed effect allows for the possibility that certain teachers tend to have different *changes* in absences across years. For teachers who only switch schools once, there is no variation in the independent variable within a teacher. As such, the teacher fixed effect specification is identified entirely by teachers who switch schools multiple times. For example, if a teacher switches from school 1 to school 2 to school 3, the teacher fixed effect specification examines whether, within a teacher, the size of the change in school-level average absence between schools is related to the change in own absences across schools. Standard errors are substantially larger after adding the teacher fixed effect, but estimates are quite similar. The teacher fixed effect regression is also not a preferred specification, since it is identified only from the unique subset of teachers who transition twice or more.

Although our measure of school-level absence avoids several issues of endogeneity, one downside of this specification is that absence-taking behavior may change over time within schools and thus the average rate of absence in other years is likely to be an imperfect proxy for the school absence rate actually experienced by the teacher. We address this issue by instrumenting for the actual school-level rate of absence a teacher experienced at a particular school using the average absence rate in years when that teacher was not present. The reduced form from this IV analysis is our baseline specifications shown in row 1 of Table 3. The IV assumption is that the average level of absence at the school when a teacher is not present only affects a teacher's absence taking behavior through its correlation with the average level of absence at the school when the teacher is there. Our first-stage estimate (not shown in a table) is 0.25 with a standard error of 0.01, suggesting that approximately 1/4 of the

Table 4.
Instrumental variable estimates of changes in own absences in response to changes in the absence environment.

	(1) Transition	(2) Before leaving	(3)	(4)	(5) After leaving	(6)	(7)
	$t-1$ to t	$t-4$ to $t-3$	$t-3$ to $t-2$	$t-2$ to $t-1$	t to $t+1$	$t+1$ to $t+2$	$t+2$ to $t+3$
	IV estimates across all switches						
IV estimates	0.542***	−0.107	−0.072	−0.005	0.053	−0.026	0.019
	(0.086)	(0.132)	(0.119)	(0.104)	(0.100)	(0.123)	(0.140)
Observations	32,094	13,323	17,739	23,436	25,132	19,903	15,556
	IV estimates for in-district switches						
IV estimates	0.553***	−0.019	−0.158	0.047	0.019	0.040	0.010
	(0.122)	(0.167)	(0.152)	(0.144)	(0.143)	(0.170)	(0.197)
Observations	22,864	10,188	13,301	17,132	18,185	14,616	11,596
	IV estimates for out-of-district switches						
IV estimates	0.482***	−0.303	0.121	−0.069	0.132	−0.119	0.035
	(0.108)	(0.206)	(0.183)	(0.137)	(0.116)	(0.161)	(0.171)
Observations	9222	3130	4433	6296	6940	5284	3959

Notes. Here we present instrumental variable estimates, where we instrument for the actual rate of absence a teacher experienced at a particular school using the average absence rate in years when that teacher was not present. Our first-stage estimate (not shown in a table) is 0.25 with a standard error of 0.01. The F-statistic on the excluded instrument is over 400. Heteroskedastic-robust standard errors are in parentheses.
* $p < 0.10$, ** $p < 0.05$.
*** $p < 0.01$.

absence rate is persistent across years within a school. The F-statistic on the excluded instrument is over 400.

Row 1 of Table 4 shows the IV estimates that suggest that a 1-day increase in school-level absence increases own absence by 0.54 days. While there are many potential explanations for why teachers respond to school average absences, we suspect that our estimates reflect teachers altering their behavior to match their new school's culture. In rows 2 and 3 of Table 4 we examine whether our estimates differ across teachers switching schools within vs. across districts. Based on the point estimates, teachers are slightly more responsive to school-level absences when they move within district compared to out of district, but the difference is not statistically significant. This finding suggests that even when a teacher remains in the same district, moving to a school with a higher average absence rate increases her rate of absence.

3.2. Workload

To understand how workloads influences teacher absences, we estimate

$$ABSENCE_{jgst} = \boldsymbol{\lambda_{st}}\gamma + \boldsymbol{X_{jgst}}\beta + \delta_g + \eta_t + \phi_s + \theta_j + \varepsilon_{jgst} \quad (2)$$

where $ABSENCE_{jgst}$ records the number of days absent for teacher j teaching in grade g and school s in year t. The vector $\boldsymbol{\lambda_{st}}$ includes the school share of minority students, whether or not the school is Title 1 funded, and school enrollment. The vector of teacher-by-year characteristics $\boldsymbol{X_{jgst}}$ includes a large vector of variables of interest; the exact variables are shown in the table and discussed in the results section. We include a full set of grade and year fixed effects to account for time-varying and grade-specific factors. Our primary specification includes both school fixed effects (ϕ_s) and teacher fixed effects (θ_j). In some specifications we omit these fixed effects and instead include teacher characteristics like demographics, education, Praxis score, sex, and race.

Because the above specification includes teacher fixed effects, teacher experience and the year fixed effects are perfectly collinear for any teacher with continuous employment. We address this issue by implementing the two-step procedure developed by Papay and Kraft (2015). In the first step of the procedure, we estimate the year effects in a model that excludes teacher fixed effects but includes all other covariates. In the second step we impose these year effects on the specification that includes teacher fixed effects. Papay and Kraft show that specifications that cap the experience profile (common in the literature) have the potential to bias *all* estimates of the experience profile.[6] In our sample, some teachers have up to 50 years of experience, and the Papay and Kraft approach technically allows for the inclusion of all 50 experience indicators. For presentational compactness, we instead include indicators for the first 6 years of experience and a final variable indicating 7 or more years of experience. Although this specification seems similar to past work that caps the experience profile, an important difference is that the two-step procedure means that the year fixed effects are identified by all teachers and not just the teachers with above 7 years of experience. As such, our estimates

[6] For example, consider a value-added specification that includes teacher fixed effects, year fixed effects, indicators for the first 6 years of experience and an indicator for having 7 or more years of experience. This specification assumes that the experience profile is flat after 7 years and teachers in this flat region are the only teachers who help to identify the year fixed effects. If teachers with more than 7 years of experience actually continue to improve, this will lead to the later year fixed effects being biased upward. The overstated year fixed effects will lead to understating the returns to experience in the first 6 years so that all estimates of the returns to experience become biased. Papay and Kraft avoid this bias by estimating the year fixed effects in an auxiliary regression.

of the experience gradient in the first 6 years will not be biased even if teacher absences continue to increase with experience above 7 years.

We include several variables that can be viewed as proxies for teacher workload. First, we are interested in the relationship between experience and absences. As teachers gain experience they may be able to reduce work intensity since they can reuse lesson plans and materials, and generally become more efficient. If teachers take absences during particularly difficult workloads, we would expect that teacher absences would be highest during a teacher's first year. However, increased experience can also influence absences through other channels, so the relationship between experience and absences should not be viewed as concrete evidence regarding the influence of workload on teacher absence.

Second, we consider the effect of grade switches on absence rates. When teachers switch grade assignments, they face many of the obstacles faced by novice teachers.[7] Teachers working in a new grade likely have increased workloads since they have to develop new lesson plans and adjust to a new age group. One concern is that being switched may also be correlated with job insecurity. In particular, a teacher may fear that her grade reassignment is a precursor to dismissal. To the extent that teachers view grade-switches as a precursor to dismissal, absences may fall following a switch in order to avoid dismissal. To help shed light on this possibility, we include two additional controls. First, we control for whether the teacher ends up leaving at the end of the year. Second, we interact the grade-switching variable with an indicator for whether the teacher has tenure. Tenured teachers should feel more secure in their jobs and so if there is no differential effect by tenure status, this would provide suggestive evidence that grade-switches do not affect absence by affecting job insecurity.

Finally, we measure teacher-specific class size since larger classes may be more difficult to teach. Since we are focused on self-contained regular elementary classes and include teacher and school fixed effects, the majority of the variation in class size comes from variation in the number of students across cohorts. We also include controls for the number of emotionally/behaviorally disabled students and the number of learning disabled students to further proxy for the difficulty of the teaching assignment.

In Table 5 we present the results from Eq. (1) including various types of fixed effects. Column 1 is shown just as a baseline, since these estimates exclude all of the fixed effects and are unlikely to reflect causal forces. Column 2 shows that adding district fixed effects has little effect on these results, suggesting that district-specific factors are not related to the relationship between the covariates and absence. Column 3 adds teacher fixed effects and Column 4 includes teacher and school fixed effects. Column 5 adds a teacher-by-school fixed effect. Column 6 restricts the sample to teachers who are at their first school in the data. This specification ensures that teachers remain in the same school and so the grade-switching effect cannot be biased by its correlation with school switches.

The estimates are very similar for columns 3 through 6 and quite different than the specifications that exclude the teacher fixed effect. For simplicity of discussion, below we focus on the specification that includes teacher and school fixed effects (column 4). Column 4 shows that absences rise sharply as teachers gain experience and that the profile is steeper than the regression that excludes teacher fixed effects suggests. We suspect that the difference in the experience gradient after adding teacher fixed effects is the result of survival bias. Specifically, the most dedicated teachers are likely to take fewer absences, and these teachers are disproportionately represented at later experience levels.

When teachers teach a new grade, they take approximately 0.5 fewer absences compared to when they teach the same grade as in the previous year. This effect does not appear to vary by tenure status. The result is robust to restricting the sample to teachers who have not switched schools (column 6), suggesting that the grade-switching estimate is not driven by school switches. Column 4 also shows that teachers take fewer absences when they have larger classes, though this effect is relatively small in magnitude. The number of behaviorally disabled or learning disabled students does not affect teacher absences. Together, we view these results as consistent with the notion that teachers take fewer absences during periods in which they might be expected to struggle. Although the absence/experience profile may be attributable to factors such as teacher age, it is difficult to think of a reason why grade switches would relate to absences outside of the workload channel. This is particularly true for tenured teachers who are unlikely to be concerned with their job security.

4. Ceteris paribus effects vs. total effects

As noted in Todd and Wolpin (2003), there are conceptually two different causal effects of an input in a production function. First, there is the ceteris paribus structural effect of the input, based directly on the educational production function. Second, there is the total effect of the input that includes both the structural effect and any endogenous change in other inputs that result. Both of these parameters are of interest. If one is interested in understanding the likely effect of exogenously changing one input through policy, the parameter that includes the endogenous input changes is most relevant. If, on the other hand, one is interested in understanding the educational production function directly, then the ceteris paribus parameter is more relevant. For example, researchers interested in understanding the underlying causes of teacher improvement would be interested in distinguishing between direct effects and endogenous responses.

As an illustrative example, we explore how controlling for absences affects the estimated returns to experience.

[7] Ost (2014) and Ost and Schiman (2015) show that similar to novice teachers, teachers who have just switched grades have lower value added and are more likely to quit teaching compared to similarly experienced teachers who have not been switched. Ost (2014) demonstrates that, conditional on a teacher fixed effect, the timing of teacher switches is essentially random and is unrelated to a variety of time-varying teacher characteristics.

Table 5.
Possible causes of teacher absences.

	(1)	(2)	(3)	(4)	(5)	(6)
Experience = 2	0.918***	0.915***	1.026***	1.014***	0.975***	0.912***
	(0.076)	(0.076)	(0.078)	(0.078)	(0.086)	(0.098)
Experience = 3	1.596***	1.593***	1.901***	1.885***	1.810***	1.757***
	(0.081)	(0.081)	(0.088)	(0.090)	(0.102)	(0.116)
Experience = 4	2.212***	2.206***	2.660***	2.632***	2.573***	2.584***
	(0.087)	(0.087)	(0.099)	(0.101)	(0.116)	(0.133)
Experience = 5	2.277***	2.266***	2.914***	2.890***	2.884***	2.843***
	(0.092)	(0.092)	(0.106)	(0.109)	(0.128)	(0.146)
Experience = 6	2.379***	2.367***	3.165***	3.130***	3.144***	3.074***
	(0.094)	(0.095)	(0.112)	(0.115)	(0.135)	(0.154)
Experience = 7+	1.327***	1.304***	3.087***	3.054***	3.120***	2.977***
	(0.059)	(0.059)	(0.100)	(0.106)	(0.130)	(0.146)
Grade change $t-1$ to t	0.033	0.034	−0.474***	−0.501***	−0.458***	−0.419***
	(0.073)	(0.073)	(0.087)	(0.088)	(0.100)	(0.119)
Grade change $t-1$ to t * Tenured	0.215**	0.218**	0.063	0.078	0.109	0.148
	(0.088)	(0.088)	(0.101)	(0.102)	(0.115)	(0.138)
School switch between t and $t+1$	2.094***	2.136***	1.627***	1.703***	1.871***	2.267***
	(0.045)	(0.046)	(0.049)	(0.051)	(0.061)	(0.078)
Class size	−0.027***	−0.026***	−0.008**	−0.009**	−0.011**	−0.014***
	(0.003)	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)
Behavioral students	0.044	0.058	−0.003	−0.002	−0.033	0.013
	(0.050)	(0.051)	(0.056)	(0.056)	(0.060)	(0.067)
Learning disabled students	−0.022	−0.021	−0.002	0.002	0.003	0.006
	(0.013)	(0.014)	(0.015)	(0.015)	(0.016)	(0.018)
Average peer absences	0.430***	0.354***	0.251***	0.219***	0.224***	0.222***
	(0.009)	(0.010)	(0.011)	(0.012)	(0.013)	(0.014)
Teacher has M.A. or higher	−0.342***	−0.341***				
	(0.032)	(0.033)				
Teacher's Praxis Test score	−0.344***	−0.307***				
	(0.018)	(0.018)				
Teacher is female	2.018***	2.003***				
	(0.065)	(0.066)				
Teacher is black	−0.404***	−0.402***				
	(0.048)	(0.048)				
Teacher is Hispanic	−0.154	−0.106				
	(0.253)	(0.255)				
District effect	N	Y	N	N	N	N
Teacher effect	N	N	Y	Y	N	Y
School effect	N	N	N	Y	N	N
Teacher-by-school effect	N	N	N	N	Y	N
Restricted to first school	N	N	N	N	N	Y
Sample size	218,860	218,860	218,860	218,860	218,860	158,823

Notes. The estimates here correspond to Eq. (2) in Section 3.2. Heteroskedastic-robust standard errors are in parentheses.
* $p < 0.10$.
** $p < 0.05$.
*** $p < 0.01$.

In the preceding section, we document that teachers take more absences as they gain experience. In a traditional value-added model, the total effect of experience includes the direct effect of experience in addition to any effect of changing absence-taking behavior correlated with experience. By controlling for absences in a value-added model, we estimate the returns to experience when teachers do not endogenously alter their absences. Following the recommendation of Kane and Staiger (2008) and Guarino, Reckase, Stacy, and Wooldridge (2014) we estimate the following lagged model,

$$A_{ijgst} = \gamma A_{i(t-1)} + \boldsymbol{X_i}\beta + \boldsymbol{C_{ij}}\delta + f(Exp_j) + \theta_j + \omega_g + \xi_t + \epsilon_{ijgst} \quad (3)$$

where A_{ijgst} is student *i*'s test score with teacher *j* in grade *g* in school *s* at time period *t*. $A_{i(t-1)}$ is lagged test score for student *i*, $\boldsymbol{X_i}$ is a vector of student characteristics, $\boldsymbol{C_{ij}}$ is a vector of classroom characteristics including class size and peer characteristics, and $f(Exp_j)$ is teacher experience, modeled here as a series of dummy variables. Though the lagged model has been shown to outperform other value-added models along a variety of dimensions, in our context, the results are not sensitive to the choice of value-added model. As with the workload analysis, we use the methodology developed by Papay and Kraft (2015) to estimate the year fixed effects.

Columns (1) and (3) of Table 6 provide baseline estimates of the returns to experience for math and reading test scores. In columns (2) and (4) we add a linear control for absences to the value-added model. First, consistent with past work, we find that teacher absences have modest negative effects on student performance (Clotfelter et al., 2009; Miller, Murnane, & Willett, 2008). Second, for both reading and math, the experience profile becomes approx-

Table 6.
Ceteris paribus vs. total returns to experience.

	Math		Reading	
	(1)	(2)	(3)	(4)
Experience = 1	0.0513***	0.0575***	0.0337***	0.0379***
	(0.0059)	(0.0059)	(0.0053)	(0.0053)
Experience = 2	0.0749***	0.0847***	0.0556***	0.0621***
	(0.0064)	(0.0066)	(0.0060)	(0.0062)
Experience = 3	0.0878***	0.0992***	0.0620***	0.0696***
	(0.0073)	(0.0075)	(0.0065)	(0.0068)
Experience = 4	0.0943***	0.1078***	0.0713***	0.0802***
	(0.0081)	(0.0084)	(0.0074)	(0.0077)
Experience = 5	0.0905***	0.1045***	0.0677***	0.0770***
	(0.0090)	(0.0093)	(0.0082)	(0.0084)
Experience = 6	0.1031***	0.1178***	0.0630***	0.0727***
	(0.0101)	(0.0104)	(0.0091)	(0.0094)
Experience = 7+	0.1170***	0.1329***	0.0772***	0.0877***
	(0.0109)	(0.0112)	(0.0095)	(0.0099)
Absence		−0.0020***		−0.0013***
		(0.0003)		(0.0003)
Observations	432,325	432,325	431,754	431,754

Notes. The estimates here correspond to the approach described in Eq. (3) of Section 4. Heteroskedastic-robust standard errors are in parentheses.
* $p < 0.10$, ** $p < 0.05$.
*** $p < 0.01$.

imately 10% steeper, though this varies somewhat across years of experience. The estimated return to the first year of teaching increases by roughly 12% for both reading and math. Taking the difference between the second and first experience dummies and comparing across models shows that the estimated return to the second year of experience increases by approximately 15% for math and 11% for reading. Though our preferred specification includes the set of experience dummies, we have also estimated linear models in order to easily quantify how the returns to experience change when controlling for absences and to more easily statistically test whether the changes in the coefficients are statistically significant. When using the linear specification, we easily reject the null that experience profile is unaffected by the addition of the absence control.[8] The fact that the estimated returns to experience increase after controlling for absences is consistent with teachers taking more absences as they gain experience combined with the fact that teacher absences hurt student performance. In practice, the endogenous absence response tends to make the total effect of experience smaller than the ceteris paribus effect, but not dramatically so.

Table 6 assumes that the effect of absence on student achievement is linear and homogenous. To the extent that the linearity and homogeneity assumptions made in Table 6 are inaccurate, it is possible that controlling for a more flexible metric of teacher absence could differently affect the estimated experience profile. The homogeneity assumption may fail since the effect of teacher absence on student achievement likely depends on the quality of the substitute teacher and also what that substitute is asked to do. In Table 7 we examine whether the experience profile is sensitive to changing the measure of absence in various ways. For convenience, columns 1 and 5 simply repeat columns 2 and 4 from Table 6.

In columns 2 and 6 of Table 7, we split our absence measure according to whether the absence occurred as part of a short spell (fewer than 3 days) or as part of a longer spell (3 or more days). For example, a teacher may have 4 two-day absence spells and 1 four-day spell in a single year. Importantly, both the short-spell variable and the long-spell variable still measure absences in 1-day units so the coefficients should be interpreted as the effect of adding an additional day onto either a short or long spell. The idea behind this measure is that principals may put less effort into finding a high quality substitute for a short absence spell.[9] We find that the estimated experience profile is very similar when allowing the effect of absence to vary by spell-length. Columns 3 and 7 allow the effect of absence to differ according to whether the substitute teacher is known to be certified, known to not be certified, or whether certification status is unknown. Again, the experience profile is fairly similar when allowing the effect of absence to differ according to the certification status of the substitute. The differences in the coefficients across certification types and the difference across spell-lengths are all statistically indistinguishable, suggesting that the simpler model from Table 6 is a reasonable baseline. In columns 4 and 8, we model absences as a series of dummy variables as opposed to assuming that absences affect student achievement linearly. This has little effect on the experience estimates, suggesting that our earlier result is not driven by the particular functional form that we assume for the relationship between teacher absences and student achievement.

5. Conclusion

We show that teacher absences are malleable and respond to aspects such as workload and school-level factors. Theoretically, teachers may take absences as a way to relieve pressure or stress, or they may take absences when their classes are going well. Though we cannot identify the underlying motivation for taking absences, the fact that teachers are more likely to be absent when workload decreases is most consistent with teachers viewing absences as more acceptable when one's class is running smoothly and requires less intensive effort. As noted in the introduction, we cannot rule out the possibility that absences proxy for high effort instead of low effort, so the

[8] We find that in the linear specification (not shown), the returns to experience increase by 8.6% for mathematics (8.2% for reading) after controlling for absences, and this change is highly significant (p-value<0.01). An informal comparison of the differences in the coefficient estimates across columns relative to the standard errors might suggest that the estimates are not significantly different from one another. This informal test fails to account for the fact that because the different columns are estimated on the same samples, the coefficients are not independent. When the dependence of the coefficients is accounted for in the hypothesis test, we strongly reject that the experience profile remains the same after controlling for absences.

[9] In results not shown, we tested whether length of spell predicts whether the replacement substitute is certified. We find no evidence that spell-length predicts certification status of the substitute. That said, certification is a crude measure of substitute quality, so this does not rule out the possibility that principals use higher quality substitutes for longer spells.

Table 7.
Ceteris paribus vs. total returns to experience with additional absence controls.

	Math				Reading			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Experience = 1	0.0575***	0.0576***	0.0583***	0.0555***	0.0379***	0.0370***	0.0378***	0.0366***
	(0.0059)	(0.0060)	(0.0060)	(0.0060)	(0.0053)	(0.0054)	(0.0054)	(0.0054)
Experience = 2	0.0847***	0.0851***	0.0867***	0.0819***	0.0621***	0.0619***	0.0628***	0.0604***
	(0.0066)	(0.0068)	(0.0067)	(0.0066)	(0.0062)	(0.0063)	(0.0063)	(0.0062)
Experience = 3	0.0992***	0.0997***	0.1008***	0.0962***	0.0696***	0.0690***	0.0697***	0.0676***
	(0.0075)	(0.0077)	(0.0075)	(0.0075)	(0.0068)	(0.0069)	(0.0069)	(0.0068)
Experience = 4	0.1078***	0.1101***	0.1104***	0.1041***	0.0802***	0.0794***	0.0805***	0.0778***
	(0.0084)	(0.0086)	(0.0085)	(0.0084)	(0.0077)	(0.0079)	(0.0078)	(0.0077)
Experience = 5	0.1045***	0.1063***	0.1076***	0.1011***	0.0770***	0.0771***	0.0778***	0.0747***
	(0.0093)	(0.0095)	(0.0094)	(0.0093)	(0.0084)	(0.0086)	(0.0085)	(0.0084)
Experience = 6	0.1178***	0.1190***	0.1198***	0.1139***	0.0727***	0.0716***	0.0730***	0.0703***
	(0.0104)	(0.0106)	(0.0105)	(0.0105)	(0.0094)	(0.0096)	(0.0095)	(0.0094)
Experience = 7+	0.1329***	0.1344***	0.1357***	0.1289***	0.0877***	0.0876***	0.0887***	0.0852***
	(0.0112)	(0.0114)	(0.0113)	(0.0113)	(0.0099)	(0.0101)	(0.0101)	(0.0099)
Absences	−0.0020***				−0.0013***			
	(0.0003)				(0.0003)			
Short absence spell		−0.0022***				−0.0013**		
		(0.0006)				(0.0006)		
Long absence spell		−0.0018***				−0.0011***		
		(0.0004)				(0.0004)		
Absence replaced by certified sub			−0.0039***				−0.0011	
			(0.0012)				(0.0010)	
Absence replaced by uncertified sub			−0.0019***				−0.0013***	
			(0.0004)				(0.0004)	
Absence replaced by unknown sub			0.0030				−0.0006	
			(0.0036)				(0.0025)	
Absences 6 to 10				−0.0023				−0.0042
				(0.0041)				(0.0039)
Absences 11 to 15				−0.0229***				−0.0109**
				(0.0055)				(0.0053)
Absences 16+				−0.0371***				−0.0269***
				(0.0070)				(0.0065)
Observations	432,325	430,427	428,991	432,325	431,754	429,858	428,427	431,754

Notes. The estimates here correspond to the approach described in Eq. (3) of Section 4. Table 7 expands on Table 6 by entering a variety of different measures of absences. Heteroskedastic-robust standard errors are in parentheses.
* $p < 0.10$.
** $p < 0.05$.
*** $p < 0.01$.

most conservative interpretation of our findings is that absences fall in response to increased workloads.

The dataset we use has many advantages, but one limitation is that it only represents North Carolina teachers and thus our results may not generalize to other contexts. There is little reason to expect that North Carolina teachers are fundamentally different than teachers in other parts of the country, but unlike in many states, the official absence policy in North Carolina is set at the state level. In areas where absence policies differ across school districts, we expect that the role of districts in determining absence-taking behavior could be more important. That said, the basic structure of North Carolina's absence policy is very similar to the basic structure of absence policies in many districts outside of North Carolina. Though we have no direct evidence on external validity, we see little reason to expect that the relationship between workload and absences would differ greatly in other U.S. contexts.

The fact that teacher absences reduce student performance suggests that policies that reduce absenteeism can improve performance as well as directly cut costs due to lower substitute usage. Our study has several implications for schools and districts that implement policies aimed at reducing absences. First, the general malleability of absences documented in our research suggests that there is likely scope for schools to affect absences. Second, absence-reduction policies could be relatively more effective among teachers with relatively low workloads who we suspect have a higher rate of discretionary absence. Finally, our results suggest that the test-score experience profile would become steeper if absence policies flattened the absence-experience profile.

Although we document the effect of workload on absences, we are unable to empirically investigate the mechanisms behind this relationship. The finding that absences fall with increased workload is consistent with several models of teacher behavior. First, teachers may seek to meet a minimum performance standard that varies with workload so that as workload increases, effort must increase in order to meet the performance standard. Second, as documented in the organizational psychology literature (Gaziel, 2004), teacher absence may be partly determined by one's organizational commitment and to the extent that added responsibilities increase commitment, this could potentially explain our findings. Finally, our findings are also consistent with a model of time allocation in which the marginal productivity of work increases as total workload increases. For example, a teacher with very high workload might accomplish so much each day that her high marginal product outweighs the benefits of absence.[10] Distinguishing between these possibilities remains an important question for future research.

References

Banerjee, A., & Duflo, E. (2006). Addressing absence. *Journal of Economic Perspectives, 20*(1), 117–132. doi:10.1257/089533006776526139.

Blase, J. J. (1982). A social-psychological grounded theory of teacher stress and burnout. *Educational Administration Quarterly, 18*(4), 93–113. doi:10.1177/0013161X82018004008.

Clotfelter, C. T., Ladd, H. F., & Vigdor, J. L. (2009). Are teacher absences worth worrying about in the United States? *Education Finance and Policy, 4*(2), 115–149. doi:10.1162/edfp.2009.4.2.115.

Duflo, E., Hanna, R., & Ryan, S. P. (2012). Incentives work: Getting teachers to come to school. *American Economic Review, 102*(4), 1241–1278. doi:10.1257/aer.102.4.1241.

Finnigan, K. S., & Gross, B. (2007). Do accountability policy sanctions influence teacher motivation? Lessons from Chicago's low-performing schools. *American Educational Research Journal, 44*(3), 594–630. doi:10.3102/0002831207306767.

Gaziel, H. H. (2004). Predictors of absenteeism among primary school teachers. *Social Psychology of Education, 7*(4), 421–434. doi:10.1007/s11218-004-5232-z.

Guarino, C., Reckase, M. D., Stacy, B., & Wooldridge, J. M. (2014). Evaluating specification tests in the context of value-added estimation. *IZA discussion paper 7974*. Institute for the Study of Labor (IZA) *https://ideas.repec.org/p/iza/izadps/dp7974.html.*

Hansen, M. (2009). How career concerns influence public workers' effort: Evidence from the teacher labor market. *Calder Working Paper http://www.urban.org/publications/1001368.html.*

Herrmann, M. A., & Rockoff, J. (2012). Worker absence and productivity: Evidence from teaching. *Journal of Labor Economics*, 749–782 no. October.

Jacob, B. A. (2013). The effect of employment protection on teacher effort. *Journal of Labor Economics, 31*(4), 727–761. doi:10.1086/669942.

Joseph, N., Waymack, N., & Zielaski, D. (2014). *Roll call: The importance of teacher attendance*. National Council on Teacher Quality *http://eric.ed.gov/?id=ED556249.*

Kane, T. J., & Staiger, D. O. (2008). Estimating teacher impacts on student achievement: An experimental evaluation. *NBER working paper 14607*. National Bureau of Economic Research, Inc. *https://ideas.repec.org/p/nbr/nberwo/14607.html.*

Miller, R. T., Murnane, R. J., & Willett, J. B. (2008). Do teacher absences impact student achievement? Longitudinal evidence from one urban school district. *Educational Evaluation and Policy Analysis, 30*(2), 181–200. doi:10.3102/0162373708318019.

Neves de Jesus, S., & Lens, W. (2005). An integrated model for the study of teacher motivation. *Applied Psychology, 54*(1), 119–134. doi:10.1111/j.1464-0597.2005.00199.x.

Ost, B. (2014). How do teachers improve? The relative importance of general and specific human capital. *American Economic Journal: Applied Economics, 6*(2), 127–151.

Ost, B., & Schiman, J. C. (2015). Grade-specific experience, grade reassignments, and teacher turnover. *Economics of Education Review, 46*(June), 112–126. doi:10.1016/j.econedurev.2015.03.004.

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, 130*(October), 105–119. doi:10.1016/j.jpubeco.2015.02.008.

Reyes, P., & Imber, M. (1992). Teachers' perceptions of the fairness of their workload and their commitment, job satisfaction, and morale: Implications for teacher evaluation. *Journal of Personnel Evaluation in Education, 5*(3), 291–302. doi:10.1007/BF00125243.

Rivkin, S. G., Hanushek, E. A., & Kain, J. F. (2005). Teachers, schools, and academic achievement. *Econometrica, 73*(2), 417–458. doi:10.1111/j.1468-0262.2005.00584.x.

Rockoff, J. E. (2004). The impact of individual teachers on student achievement: Evidence from panel data. *American Economic Review, 94*(2), 247–252. doi:10.1257/0002828041302244.

Timperley, H., & Robinson, V. (2000). Workload and the professional culture of teachers. *Educational Management Administration & Leadership, 28*(1), 47–62. doi:10.1177/0263211X000281005.

Todd, P. E., & Wolpin, K. I. (2003). On the specification and estimation of the production eunction for cognitive achievement. *The Economic Journal, 113*(485), F3–33. doi:10.1111/1468-0297.00097.

[10] In most time allocation models, agents get utility from wages – not directly from their marginal products. For the most part, wages for teachers do not vary with workload so this model would only be consistent with our results if teachers obtain utility directly from their marginal product. This is plausible in the case of teachers since there is a literature documenting that many teachers are intrinsically motivated to help students improve (Finnigan & Gross, 2007).