

Teacher-to-classroom assignment and student achievement

Bryan S. Graham*, Geert Ridder†, Petra Thiemann‡, Gema Zamorro§

March 8, 2021

Abstract

We study the effects of counterfactual teacher-to-classroom assignments on average student achievement in elementary and middle schools in the US. We use the Measures of Effective Teaching (MET) experiment to semiparametrically identify the average reallocation effects (AREs) of such assignments. Our findings suggest that changes in within-district teacher assignments could have appreciable effects on student achievement. Unlike policies which require hiring additional teachers (e.g., class-size reduction measures), or those aimed at changing the stock of teachers (e.g., VAM-guided teacher tenure policies), alternative teacher-to-classroom assignments are resource neutral; they raise student achievement through a more efficient deployment of *existing* teachers.

JEL codes: I20, I21, I24

Keywords: teacher quality, teacher assignment, education production, student achievement, average reallocation effects, K–12

*Professor of Economics, Department of Economics, University of California - Berkeley, CA 94720-3880, United States, and National Bureau of Economic Research, e-mail: bgraham@econ.berkeley.edu.

†Professor of Economics, Department of Economics, University of Southern California, Los Angeles, CA 90089-0253, United States, e-mail: ridder@usc.edu.

‡Corresponding author. Assistant Professor of Economics, Department of Economics, Lund University, P.O. Box 7080, 22007 Lund, Sweden, and IZA, e-mail: petra.thiemann@nek.lu.se, phone: +46 (0)46 222 8686.

§Professor and 21st Century Endowed Chair in Teacher Quality, Department of Education Reform, University of Arkansas, Fayetteville, AR 72701, United States, and CESR, e-mail: gzamarro@uark.edu.

1 Introduction

Approximately four million teachers work in the public elementary and secondary education system in the United States. These teachers provide instruction to almost fifty million students, enrolled in nearly one hundred thousand schools, across more than thirteen thousand school districts (McFarland et al., 2019, Snyder et al., 2017). Differences in measured student achievement are substantial across US schools as well as across classrooms within these schools. Beginning with Hanushek (1971), a large economics of education literature attributes cross-classroom variation in student achievement to corresponding variation in (largely) latent teacher attributes. These latent attributes are sometimes referred to as teacher quality or teacher value-added.

The implications of value-added measures (VAM) for education policy are controversial both within the academy and outside it.¹ The most contentious applications of VAM involve their use in teacher tenure and termination decisions (cf., Chetty et al., 2012, Darling-Hammond, 2015). The premise of such applications is that changes in the stock of existing teachers – specifically rooting out teachers with low VAMs and retaining those with high ones – could lead to large increases in student achievement and other life outcomes.

In this paper we pose an entirely different question: is it possible to raise student achievement, without changes to the existing pool of teachers, by changing who teaches whom? Schools and school districts are the loci of teacher employment. To keep our analysis policy-relevant, we therefore focus on the achievement effects of different within-school and within-district teacher-to-classroom assignment policies.

For teacher assignment policies to matter, teachers must vary in their effectiveness in teaching different types of students. For example, some teachers may be especially good at teaching English language learners, minority students, or accelerated learners (c.f., Dee, 2004, Loeb et al., 2014a). Formally educational production must be non-separable in some teacher and student attributes (Graham et al., 2007, 2014, 2020).

Using data collected in conjunction with the Measures of Effective Teaching (MET) project, we present experimental evidence of such non-separabilities. The MET study was run in six urban public school districts in the US during two school years (AY 2009/10 and 2019/11) in grades 4–10. The goal of the study was to evaluate different measures teacher effectiveness; to this end, MET researchers randomly assigned teachers to classrooms within schools. We exploit the experimental design in combination with rich data on teaching practices collected throughout the experiment. Specifically, we study complementarities between (i) an observation-based, pre-experiment measure of teaching practice – Danielson’s (2011) Framework for Teaching (FFT) instrument – and (ii) students’ and classroom peers’ baseline

¹See, for example, the March 2015 special issue of *Educational Researcher* on value-added research and policy or the American Statistical Association’s statement on value-added measures (Morganstein and Wasserstein, 2014).

test scores, also measured pre-experiment. Observational measures of teaching practices, such as the FFT, are routinely collected by school principals in nearly all US public schools (U.S. Department of Education, 2020). They are used to inform teaching assignments, as well as to promote, recognize, and retain good teachers. It seems reasonable, therefore, to explore whether these measures might also aid in determining classroom assignments.²

Alternative teacher assignments change the joint distribution of teacher FFT and student baseline achievement. When the educational production function is non-separable in these variables, these changes can affect average achievement.

To quantify the potential achievement effects of alternative teacher-to-classroom assignments, we estimate the gains associated with implementing an allocation that maximizes average achievement. Because our model allows for both complementarities and non-linearity in the underlying variables (FFT and student baseline achievement), determining an outcome-maximizing assignment is non-trivial. We compute the optimal assignment using linear programming methods under the constraint that the existing pool of teachers remains unchanged.

We find that within-district changes in teacher-to-classroom assignments could increase average classroom student achievement by as much as 0.04 standard deviations. This effect corresponds to the estimated difference in average achievement across the teacher-to-classroom assignment which maximizes aggregate achievement versus the one which minimizes it. Perhaps more realistically, comparing the status quo assignment in MET schools – which was generated by random assignment of teachers to classrooms³ – with the optimal one generates an estimated increase in average test scores of 0.02 standard deviations.

To benchmark these effects consider a policy which removes the bottom $\tau \times 100$ percent of teachers from classrooms – sorted according to their VAM – and replaces them with average teachers (i.e., teachers with VAMs of zero). Assuming a Gaussian distribution for teacher value-added, the effect of such an intervention would be to increase the mean of the student test score distribution by

$$(1 - \tau) \sigma \frac{\phi\left(\frac{q_\tau}{\sigma}\right)}{1 - \Phi\left(\frac{q_\tau}{\sigma}\right)}$$

standard deviations.⁴ Here σ corresponds to the standard deviation of teacher value-added and q_τ to its τ^{th} quantile.

Rockoff (2004, Table 2) and Rothstein (2010, Table 6) estimate a standard deviation

²Observational measures of teacher teaching practices are also increasingly used in research (e.g. Araujo et al., 2016, Aucejo et al., 2019).

³The joint distribution of FFT and student baseline achievement looks very similar under the random allocation, compared to the allocation under non-experimental conditions in the prior school year. See Section 3 for details.

⁴As is common practice, we abstract from the effects of possible behavioral responses to VAM-guided teacher retention policies. For example, teachers might become demoralized by VAM retention systems and teach less effectively as a result; or they could be motivated to teach more effectively.

of teacher value-added of between 0.10 and 0.15. Taking the larger estimate and setting $\tau = 0.05$ (0.10) generates an expected increase in student test scores of 0.015 (0.026) standard deviations; smaller than, albeit comparable to, our reallocation effects. We note that replacing five (ten) percent of teachers would be difficult to do in practice. It would be even more difficult to correctly identify the bottom five (ten) percent of teachers (according to VA) and replace them with average ones. We conclude that the achievement effects of teacher assignment policies are meaningful.

In contrast to policies which replace teachers, and/or require the hiring of new ones, the within-school and within-district teacher reassignment policies we explore in this paper are potentially resource neutral. Districts require no new teachers or financial resources to implement them. Of course, reassigning teachers across classrooms, and especially across schools within a district, may involve other types of costs. For example, many school districts operate under collective bargaining agreements which give senior teachers partial control over their school assignment (e.g., [Cohen-Vogel et al., 2013](#)).⁵

Work by Susanna Loeb and co-authors suggests that US school districts have *de facto* teacher-to-classroom assignment policies ([Kalogrides et al., 2011](#), [Grissom et al., 2015](#)). For example, they find that less experienced, minority, and female teachers are more likely to be assigned to predominantly minority classrooms. They also present evidence that principals use teacher assignments as mechanisms for retaining teachers – as well as for encouraging less effective teachers to leave – and that more experienced teachers exert more influence on classroom assignment decisions.

Our work helps researchers and policy-makers understand the achievement effects of such policies and the potential benefits of alternative ones. The findings presented below suggest that teacher-to-classroom assignment policies are consequential and that changes to them could meaningfully increase average student achievement.

In addition to our substantive results, we present new identification results for average reallocation effects (AREs). Identification and estimation of AREs under (conditional) exogeneity is considered by [Graham et al. \(2014, 2020\)](#). Unfortunately these results do not apply directly here. Although teachers were randomly assigned to classrooms as part of the MET experiment, compliance was imperfect. Furthermore some students moved across classrooms after the random assignment of teachers, which raises concerns about bias due to endogenous student sorting. We develop a semiparametric instrumental variables estimator (e.g., [Ai and Chen, 2003](#)) which corrects for student and teacher non-compliance. Our analysis highlights how complex identification can be in the context of multi-population matching models where agents sort endogenously.

⁵Moreover, many teachers have preferences over school and student characteristics. Thus, one might need to pay bonuses to move teachers to certain schools or classrooms (e.g., classrooms with high fractions of disadvantaged students).

Our study builds on a series of papers that document non-separability of teacher and student inputs along different characteristics. Starting with [Dee \(2004, 2005, 2007\)](#), these studies find non-separabilities along students’ and teachers’ race and ethnicity (e.g., [Dee, 2004, 2005](#), [Lusher et al., 2018](#), [Egalite and Kisida, 2018](#), [Joshi et al., 2018](#)), gender (e.g., [Dee, 2005, 2007](#), [Lim and Meer, 2017](#), [Egalite and Kisida, 2018](#), [Gong et al., 2018](#)), student background and teacher quality/value-added (e.g., [Jackson, 2013](#), [Loeb et al., 2014b](#), [Condie et al., 2014](#)), and student background and teacher experience (e.g., [Clotfelter et al., 2006](#)).

In this paper we consider non-separabilities in teaching practice and prior student achievement. The FFT relies on classroom observations that are evaluated based on a professional teaching rubric. Similar measures are often used in practice: nearly all principals in US public schools – 98 percent – collect observational measures that are based on professional rubrics to evaluate teachers; these evaluations in turn are used to determine teaching assignments ([U.S. Department of Education, 2020](#)).

In independent work, [Aucejo et al. \(2019\)](#) also explore non-separability based on the FFT in the MET data. They present an interesting new approach which emphasizes match effects between within-classroom dispersion in student background and the FFT in elementary school classrooms. Moreover, they investigate different sub-components of the FFT, something not done here. Because our focus is on the achievement effects of explicit alternative teacher-to-classroom assignments, our research design emphasizes the identification of structural features of the education production function. In contrast, [Aucejo et al. \(2019\)](#) implement a reduced form intent-to-treat design. Both approaches have merit, depending on the questions being asked. Our focus on counterfactual assignments potentially distant from the status quo one requires the identification, as we detail below, of at least some structural features of education production.

To our knowledge, this paper is the first to derive and quantify reallocation effects from alternative teacher-to-classroom assignments. Specifically, we focus on an “optimal” allocation that maximizes aggregate outcomes. The paper thus also relates to the literature on optimal assignments under resource constraints, as explored in other contexts, for example the assignments of students to dorm rooms ([Bhattacharya, 2009](#)), of bed nets to households ([Bhattacharya and Dupas, 2012](#)), and of refugees to geographic locations ([Trapp et al., 2018](#)). See [Bonhomme \(2020\)](#) for related discussion and references.

Our work naturally has several limitations that should be addressed in future research. First, it seems likely that assignments based upon a *different* combination of teacher and classroom attributes could lead to even greater achievement gains.⁶ We leave explorations along these lines to further study. Second, readers may reasonably raise questions about the external validity of the results reported below. Indeed, our results need to be interpreted with caution, since our data comes from an experiment in urban school districts, and

⁶That is, we did not consider the universe of possible assignment variables.

schools participated voluntarily. We do, however, hope our results are sufficiently compelling to motivate additional research and experimentation with teacher-to-classroom assignment policies.

2 Model and identification

Our goal is to identify the average achievement effects of alternative assignments of teachers to MET classrooms. These are average reallocation effects (AREs), as introduced by [Graham et al. \(2007, 2014\)](#). The identification challenge is to use the observed MET teacher-to-classroom assignments and outcomes to recover these AREs.

Our analysis is based upon experimentally generated combinations of student and teacher attributes, that is, it exploits the random assignment of teachers to classrooms in the MET experiment.⁷ Like in many other field experiments, various deviations from MET’s intended protocol complicate our analysis. In this section we outline a semiparametric model of educational production and consider its identification based upon the MET project *as implemented*, using the MET data *as collected*.

It is useful, however, to first explore nonparametric identification of reallocation effects under an ideal implementation of the MET project (henceforth MET *as designed*). Such an approach clarifies how the extra restrictions introduced below allow for the identification of reallocation effects despite non-compliance, attrition, and other deviations from the intended experimental protocol.

Nonparametric identification under ideal circumstances

Our setting features two populations, one of students and the other of teachers. Each student is distinguished by an observed attribute X_i , in our case a measure of baseline academic achievement, and an unobserved attribute, say “student ability,” V_i . Similarly, each teacher is characterized by an observed attribute W_i , in our case an observation-based measure of teaching pedagogy, and an unobserved attribute, say “teacher quality,” U_i .⁸

Let $i = 1, \dots, N$ index students. Let C be the total number of MET classrooms or equivalently teachers. We define G_i to be a $C \times 1$ vector of classroom assignment indicators. The c^{th} element of G_i equals one if student i is in classroom $c \in \{1, \dots, C\}$ and zero otherwise. The indices of student i ’s peers or classmates are therefore given by the index set

$$p(i) = \{j : G_i = G_j, i \neq j\}.$$

⁷To be precise, the randomization was carried out within schools. See Section 3 for details.

⁸We use “ability” as a shorthand for latent student attributes associated with higher test scores; likewise we use “quality” as a shorthand for latent teacher attributes associated with higher test scores.

Next we define the peer average attribute as $\bar{X}_{p(i)} = \frac{1}{|p(i)|} \sum_{j \in p(i)} X_j$ (i.e., the average of the characteristic X across student i 's peers). We define $\bar{V}_{p(i)}$ similarly.

The MET project protocol did not impose any requirements on how students, in a given school-by-grade cell, were divided into classrooms. Evidently schools followed their existing procedures for dividing students within a grade into separate classrooms. An implication of this observation is that the MET experiment implies no restrictions on the joint density

$$f_{X_i, V_i, \bar{X}_{p(i)}, \bar{V}_{p(i)}}(x, v, \bar{x}, \bar{v}), \quad (1)$$

beyond the requirement that the density be feasible.⁹ For example, if most schools tracked students by prior test scores, then we would expect X_i and $\bar{X}_{p(i)}$ to positively covary. If, instead, students were randomly assigned to classrooms and hence peers, we would have, ignoring finite population issues, the factorization

$$f_{X_i, V_i, \bar{X}_{p(i)}, \bar{V}_{p(i)}}(x, v, \bar{x}, \bar{v}) = f_{X_i, V_i}(x, v) f_{\bar{X}_{p(i)}, \bar{V}_{p(i)}}(\bar{x}, \bar{v}).^{10}$$

Our analysis allows for arbitrary dependence between own and peer attributes, both observed and unobserved, and consequently is agnostic regarding the protocol used to group students into classrooms.

Two implications of this agnosticism are (i) our analysis is necessarily silent about the presence and nature of any peer group effects, and (ii) it is likely that more complicated policies, involving *simultaneously* regrouping students into new classes *and* reassigning teachers to them, could raise achievement by more than what is feasible via reassignments of teachers to *existing* classrooms alone, which is the class of policies we consider.¹¹

Although nothing about the MET protocol generates restrictions on the joint density (1), random assignment of teachers to classrooms – however formed – ensures that

$$f_{X_i, V_i, \bar{X}_{p(i)}, \bar{V}_{p(i)}, W_i, U_i}(x, v, \bar{x}, \bar{v}, w, u) = f_{X_i, V_i, \bar{X}_{p(i)}, \bar{V}_{p(i)}}(x, v, \bar{x}, \bar{v}) f_{W_i, U_i}(w, u). \quad (2)$$

Here W_i and U_i denote the observed and unobserved attributes of the teacher assigned to the classroom of student i . A perfect implementation of MET as designed would ensure that student and teacher attributes vary independently of each other. Our research design is fundamentally based upon restriction (2).

⁹We cannot, for example, have an assignment which implies all students have above-average peers.

¹⁰Indeed additional restrictions on the second density to the right of the equality above would also hold; see [Graham et al., 2010](#), and [Graham, 2011](#), for additional discussion and details.

¹¹Learning about the effects of policies which simultaneously regroup students and reassign teachers would require double randomization. See [Graham \(2008\)](#) for an empirical example and [Graham et al. \(2010, 2020\)](#) for a formal development.

Let Y_i be an end-of-year measure of student achievement, generated according to

$$Y_i = g(X_i, \bar{X}_{p(i)}, W_i, V_i, \bar{V}_{p(i)}, U_i). \quad (3)$$

Other than the restriction that observed and unobserved peer attributes enter as means, equation (3) imposes no restrictions on educational production.¹² Under restriction (2) the conditional mean of the outcome given observed own, peer, and teacher attributes equals

$$\begin{aligned} \mathbb{E}[Y_i | X_i = x, \bar{X}_{p(i)} = \bar{x}, W_i = w] &= \iiint \left[g(x, \bar{x}, w, v, \bar{v}, u) f_{V_i, \bar{V}_{p(i)} | X_i, X_{p(i)}}(v, \bar{v} | x, \bar{x}) \right. \\ &\quad \left. \times f_{U_i | W_i}(u_i | w_i) \right] dv d\bar{v} du \\ &= m^{\text{amf}}(x, \bar{x}, w). \end{aligned} \quad (4)$$

Equation (4) coincides with (a variant of) the Average Match Function (AMF) estimand discussed by [Graham et al. \(2014, 2020\)](#). The AMF can be used to identify AREs. Our setting – which involves multiple students being matched to a single teacher – is somewhat more complicated than the one-to-one matching settings considered by [Graham et al. \(2014, 2020\)](#). One solution to this “problem” would be to average equation (3) across all students in the same classroom and work directly with those averages. As will become apparent below, however, working with a student-level model makes it easier to deal with non-compliance and attrition, which have distinctly student-level features. It also connects our results more directly with existing empirical work in the economics of K-to-12 education, where student-level modelling predominates, and results in greater statistical power.

The decision to model outcomes at the student level makes the analysis of teacher re-assignments a bit more complicated, at least superficially. To clarify the issues involved it is helpful to consider an extended example. Assume there are two *types* of students, $X_i \in \{0, 1\}$, and two *types* of teachers, $W_i \in \{0, 1\}$. For simplicity assume that the population fractions of type $X_i = 1$ students and type $W_i = 1$ teachers both equal one-half, that is, half of the students are of type 1 and half of the students are taught by a teacher of type 1. Assume, again to keep things simple, that classrooms consist of three students each.

Table 1 summarizes this basic set-up. Column 1 lists classroom types. For example, a 000 classroom consists of all type-0 students. There are four possible classroom types, each assumed to occur with a frequency of one-fourth.¹³ The status quo mechanism for grouping students into classrooms induces a joint distribution of own and peer average attributes. This joint distribution is given in the right-most column of Table 1. For instance, $\frac{1}{4}$ (3 out of 12) of the students are in a classroom with two type-0 peers, so that $f_{X_i, \bar{X}_{p(i)}}(0, 0) = \frac{1}{4}$,

¹²A fully nonparametric model would allow achievement to vary with any exchangeable function of peer attributes; see [Graham et al. \(2010\)](#) for more details.

¹³Within classrooms students are exchangeable.

Table 1: Feasible teacher reassignments

Classroom Type	Status quo	Counterfactual	X_i	$\bar{X}_{p(i)}$	$f(X_i, \bar{X}_{p(i)})$
	$\Pr(W_i = 1 X_i, \bar{X}_{p(i)})$	$\Pr(\tilde{W}_i = 1 X_i, \bar{X}_{p(i)})$			
000	$\frac{1}{2}$	$\frac{1}{3}$	0	0	$\frac{1}{4}$
			0	0	$\frac{1}{4}$
			0	0	$\frac{1}{4}$
001	$\frac{1}{2}$	$\frac{1}{2}$	0	$\frac{1}{2}$	$\frac{1}{6}$
			0	$\frac{1}{2}$	$\frac{1}{6}$
			1	0	$\frac{1}{12}$
011	$\frac{1}{2}$	$\frac{1}{2}$	0	1	$\frac{1}{12}$
			1	$\frac{1}{2}$	$\frac{1}{6}$
			1	$\frac{1}{2}$	$\frac{1}{6}$
111	$\frac{1}{2}$	$\frac{2}{3}$	1	1	$\frac{1}{4}$
			1	1	$\frac{1}{4}$
			1	1	$\frac{1}{4}$

Note: The population fraction of type $X_i = 1$ students is $\frac{1}{2}$ and that of type $W_i = 1$ teachers is also $\frac{1}{2}$. Classrooms of three students each are formed, such that the frequency of each of the four possible classroom configurations is $\frac{1}{4}$ in the population of classrooms of size 3. Under the status quo teachers are assigned to classrooms at random; in the counterfactual teachers are assigned more assortatively. See the main text for more information.

and $\frac{1}{6}$ (2 out of 12) of the students are in a classroom with one type-0 and one type-1 peer, so that $f_{X_i, \bar{X}_{p(i)}}(0, \frac{1}{2}) = \frac{1}{6}$. The MET experiment implies no restrictions on the joint density $f_{X_i, \bar{X}_{p(i)}}(x, \bar{x})$, consequently we only consider policies which leave it unchanged.

Next assume, as was the case in the MET experiment, that under the status quo teachers are randomly assigned to classrooms. This induces the conditional distribution of W_i given X_i and $\bar{X}_{p(i)}$ reported in column 2 of Table 1. Of course, from this conditional distribution, and the marginal for X_i and $\bar{X}_{p(i)}$, we can recover the joint distribution of own type, peer average type, and teacher type (i.e., of X_i , $\bar{X}_{p(i)}$ and W_i).

Now consider the AMF: $m^{\text{amf}}(x, \bar{x}, w)$. Consider the subpopulation of students with $X_i = 1$ and $\bar{X}_{p(i)} = \frac{1}{2}$. Inspecting Table 1, this subpopulation represents $\frac{1}{6}$ of all students (right-most column of Table 1). If we assign to students in this subpopulation a teacher of type $W_i = 1$, then the expected outcome coincides with $m^{\text{amf}}(1, \frac{1}{2}, 1)$. Under random assignment of teachers the probability of assigning a type-1 teacher is the same for all subpopulations of students.

Finally consider a counterfactual assignment of teachers to classrooms. Since we leave the composition of classrooms unchanged, $f_{X_i, \bar{X}_{p(i)}}(x, \bar{x})$ is left unmodified. The counterfactual assignment therefore corresponds to a conditional distribution for teacher type, $\tilde{f}_{\tilde{W}_i | X_i, \bar{X}_{p(i)}}(w | x, \bar{x})$

which satisfies the feasibility condition:

$$\iint \tilde{f}_{\tilde{W}_i|X_i, \bar{X}_{p(i)}}(w|x, \bar{x}) f_{X_i, \bar{X}_{p(i)}}(x, \bar{x}) dx d\bar{x} = f(w) \quad (5)$$

for all $w \in \mathbb{W}$. Here \tilde{f} denotes a counterfactual distribution, while f denotes a status quo one. We use \tilde{W}_i to denote an assignment from the counterfactual distribution. Note that by feasibility of an assignment $\tilde{W}_i \stackrel{D}{=} W_i$ marginally, but will differ conditional on student attributes. Condition (5), as discussed by [Graham et al. \(2014\)](#), allows for degenerate conditional distributions, as might occur under a perfectly positive assortative matching.

Average achievement under a counterfactual teacher-to-classroom assignment equals:

$$\beta^{\text{are}}(\tilde{f}) = \iint \left[\int m^{\text{amf}}(x, \bar{x}, w) \tilde{f}_{\tilde{W}_i|X_i, \bar{X}_{p(i)}}(w|x, \bar{x}) dw \right] f_{X_i, \bar{X}_{p(i)}}(x, \bar{x}) dx d\bar{x}. \quad (6)$$

Since all the terms to the right of the equality are identified, so too is the ARE. Conceptually we first – see the inner integral in equation (6) – compute the expected outcome in each type of classroom (e.g., $X_i = x$ and $\bar{X}_{p(i)} = \bar{x}$) given its new teacher assignment (e.g., to type $\tilde{W}_i = w$). We then – see the outer two integrals in equation (6) – average over the status quo distribution of $X_i, \bar{X}_{p(i)}$, which is left unchanged. This yields average student achievement under the new assignment of teachers to classrooms.

In addition to the feasibility condition (5) we need to also rule out allocations that assign different teachers to students in the same classrooms. Note that $m^{\text{amf}}(x, \bar{x}, w)$ is the average outcome for the subpopulation of students of type $X_i = x$ with peers $\bar{X}_{p(i)} = \bar{x}$. For example, in [Table 1](#) classroom 001 has students from two subpopulations so defined. Assignment of teachers to subpopulations of students opens up the possibility that a classroom is assigned to teachers of different types for its constituent subgroups of students. If, as indicated in [Table 1](#), the teacher-type assignment probability is the same for all subpopulations of students represented in a classroom, then the ARE in equation (6) coincides with one based on direct assignment of teachers to classrooms. This implicit restriction on teacher assignments provides a link between models for individual outcomes and classroom-level reallocations.

Semiparametric identification under MET as implemented

In the MET experiment as implemented not all teachers and students appear in their assigned classrooms. This occurs both due to attrition (e.g., when a student changes schools prior to follow-up) as well as actual non-compliance (e.g., when a teacher teaches in a classroom different from their randomly assigned one).

In this section we describe our approach to identifying AREs in MET as implemented. Relative to the idealized analysis of the previous subsection we impose two types of additional

restrictions. First, we work with a semiparametric, as opposed to a nonparametric, educational production function. Second, we make behavioral assumptions regarding the nature of non-compliance. Both sets of assumptions are (partially) testable.

Educational production function

Our first set of restrictions involve the form of the educational production function. A key restriction we impose is that *unobserved* student, peer, and teacher attributes enter separably. Although this assumption features in the majority of economics of education empirical work (e.g., Chetty et al., 2014a,b), it is restrictive. We also discretize the observed student and teacher attribute. This allows us to work with a parsimoniously parameterized educational production function that nevertheless accommodates complex patterns of complementarity between student and teacher attributes. Discretization also allows us to apply linear programming methods to study counterfactual assignments (c.f., Graham et al., 2007, Bhattacharya, 2009).

Specifically we let X_i be a vector of indicators for each of K “types” of students. Types correspond to intervals of baseline test scores. Our preferred specification works with $K = 3$ types of students: those with low, medium, and high baseline test scores.¹⁴ In this case X_i is a 2×1 vector of dummies for whether student i ’s baseline test score was in the medium or high range (with the low range being the omitted group). This definition of X_i means that $\bar{X}_{p(i)}$ equals the 2×1 vector of fractions of peers in the medium and high baseline categories (with the fraction low range omitted).

We discretize the distribution of the teacher attribute similarly: W_i is a vector of indicators for L different ranges of FFT scores. In our preferred specification we also work with $L = 3$ types of teachers: those with low, medium, and high FFT scores. Hence W_i is again a 2×1 vector of dummies for whether the teacher of student i ’s FFT score was in the medium or high range (with the low range again being the omitted group).

We assess the sensitivity of our results to coarser and finer discretizations of the baseline test score and FFT distributions. Specifically we look at $K = L = 2$ and $K = L = 4$ discretizations.

We posit that end-of-school year achievement for student i is generated according to

$$\begin{aligned}
Y_i = & \underbrace{\alpha + X_i' \beta + V_i}_{\text{Student Ability}} + \underbrace{\bar{X}_{p(i)}' \gamma + \rho \bar{V}_{p(i)}}_{\text{Peer Effect}} + \underbrace{W_i' \delta + U_i}_{\text{Teacher Quality}} \\
& + \underbrace{(X_i \otimes \bar{X}_{p(i)})' \zeta}_{\text{Student-Peer Complementarity}} + \underbrace{(X_i \otimes W_i)' \eta + (W_i \otimes \bar{X}_{p(i)})' \lambda}_{\text{Student-Teacher Complementarity}}
\end{aligned} \tag{7}$$

Observe that – as noted above – own, V_i , peer, $\bar{V}_{p(i)}$, and teacher, U_i , unobservables enter

¹⁴Precise variable definitions are given below.

equation (7) additively. The labelled grouping of terms in equation (7) highlights the flexibility of our model relative to those typically employed by researchers. As in traditional models, observed and unobserved student and teacher attributes are posited to directly influence achievement. We add to this standard set-up the possibility of complementarity between own and peer attributes, and complementarity between own and teacher attributes. Additionally our model allows for both observed and *unobserved* peer attributes to influence achievement.

Conditional on working with a discrete student and teacher type space, equation (7) is unrestrictive in how own and teacher attributes interact to generate achievement. In contrast, equation (7) restricts the effect of peers’ observed composition on the outcome. Partition $\zeta = (\zeta_1, \dots, \zeta_{K-1})$ and similarly partition $\lambda = (\lambda_1, \dots, \lambda_{L-1})$. The $(K - 1) \times 1$ gradient of student i ’s outcome with respect to peer composition is

$$\frac{\partial Y_i}{\partial \bar{X}_{p(i)}} = \gamma + \sum_{k=1}^{K-1} X_{ki} \zeta_k + \sum_{l=1}^{L-1} W_{li} \lambda_l, \quad (8)$$

which is constant in $\bar{X}_{p(i)}$, although varying heterogeneously with student and teacher type. Put differently convexity/concavity in $\bar{X}_{p(i)}$ is ruled out by equation (7). It should be noted that the MET data, in which the assignment of peers is not random, are not suitable for estimating peer effects (non-linear or otherwise).

For completeness we also include the interaction of teacher type with peer composition – the $(W_i \otimes \bar{X}_{p(i)})$ regressor in equation (7) – although λ is poorly identified in practice.¹⁵ Due to our limited sample size, we do not include the third order interactions of own, peer, and teacher types.¹⁶

Relative to a standard “linear-in-means” type model typically fitted to datasets like ours (e.g., Hanushek et al., 2004):

$$Y_i = \alpha + X_i' \beta + \bar{X}_{p(i)}' \gamma + W_i' \delta + V_i + U_i, \quad (9)$$

equation (7) is rather flexible. It allows for rich interactions in observed own, peer, and teacher attributes and is explicit in that both observed *and* unobserved peer attributes may influence own achievement. The “linear-in-means” model (9) presumes homogenous effects

¹⁵Admittedly, the $(W_i \otimes \bar{X}_{p(i)})$ term is not entirely straightforward to interpret; however it seemed ad hoc to include some second-order interactions while a priori excluding others. We also report specifications which exclude this term below (as its coefficient is always insignificantly different from zero).

¹⁶The educational production function in equation (7) includes $J = \dim(\alpha) + \dim(\beta) + \dim(\gamma) + \dim(\delta) + \dim(\zeta) + \dim(\eta) + \dim(\lambda) = 1 + 2 + 2 + 2 + 4 + 4 + 4 = 19$ parameters. A fully interacted model would introduce 8 additional parameters, for $3 \times 3 \times 3 = 27$ in total. Even this model would not be fully flexible since, although constructed from averages of binary indicators, $\bar{X}_{p(i)}$ is not binary-valued. A more flexible model would therefore also include, for example, interactions of X_i with the squares of the elements of $\bar{X}_{p(i)}$ (and so on.)

and does not explicitly incorporate unobserved peer attributes.¹⁷

As mentioned above, a student’s assigned teacher and peers may deviate from her realized ones due to attrition and non-compliance. To coherently discuss our assumptions about these issues we require some additional notation. Let W_i^* and $\bar{X}_{p^*(i)}$ denote student i ’s assigned teacher and peer attribute (here $p^*(i)$ is the index set of i ’s assigned classmates). Random assignment of teachers to classrooms ensures that a student’s assigned teacher’s attributes are independent of her own unobservables:

$$\mathbb{E} [V_i | X_i, \bar{X}_{p^*(i)}, W_i^*] = \mathbb{E} [V_i | X_i, \bar{X}_{p^*(i)}] \stackrel{def}{=} g_1 (X_i, \bar{X}_{p^*(i)}). \quad (10)$$

Here $g_1(x, \bar{x})$ is unrestricted. Under double randomization, with students additionally grouped into classes at random, we would have the further restriction

$$\mathbb{E} [V_i | X_i, \bar{X}_{p^*(i)}] = \mathbb{E} [V_i | X_i].$$

However, since the MET experiment placed no restrictions on how students were grouped into classrooms, we cannot rule out the possibility that a student’s peer characteristics, $\bar{X}_{p^*(i)}$, predict her own unobserved ability, V_i . Consequently our data are necessarily silent about the presence and nature of any peer group effects in learning. This limitation does not limit our ability to study the effects of teacher reallocations, because we leave the student composition of classrooms – and hence the “peer effect” – fixed in our counterfactual experiments.

Finally, even with double randomization, we would still have $\mathbb{E} [V_i | X_i] \neq 0$. Observed and unobserved attributes may naturally covary in any population (e.g., for example, average hours of sleep, which is latent in our setting, plausibly covaries with baseline achievement and also influences the outcome). Such covariance is only a problem if, as is true in the traditional program evaluation setting, the policies of interest induce changes in the marginal distribution of X_i – and hence the joint distribution of X_i and V_i .¹⁸ This is not the case here: any reallocations leave the joint distribution of X_i and V_i unchanged.

The MET protocol also ensures that assigned peer unobservables, $\bar{V}_{p^*(i)}$, are independent of the observed attributes of one’s assigned teacher:

$$\mathbb{E} [\bar{V}_{p^*(i)} | X_i, \bar{X}_{p^*(i)}, W_i^*] = \mathbb{E} [\bar{V}_{p^*(i)} | X_i, \bar{X}_{p^*(i)}] \stackrel{def}{=} g_2 (X_i, \bar{X}_{p^*(i)}), \quad (11)$$

with $g_2 (X_i, \bar{X}_{p^*(i)})$, like $g_1 (X_i, \bar{X}_{p^*(i)})$, unrestricted.

Random assignment of teachers to classrooms also ensures independence of the unobserved

¹⁷Manski (1993) did allow for unobserved peer attributes as did Graham (2008); but these cases are exceptional.

¹⁸A prototypical example is a policy which increases years of completed schooling. Such a policy necessarily changes the joint distribution of schooling and unobserved labor market ability relative to its status quo distribution.

attribute of a student's *assigned* teacher and observed student and peer characteristics:

$$\mathbb{E} [U_i^* | X_i, \bar{X}_{p^*(i)}, W_i^*] = \mathbb{E} [U_i^* | W_i^*] = 0. \quad (12)$$

The second equality is just a normalization.¹⁹

Under MET *as designed* we could identify AREs using equations (10), (11), and (12). To see this let, as would be true under perfect compliance, $W_i = W_i^*$ and $\bar{X}_{p(i)} = \bar{X}_{p^*(i)}$ for all $i = 1, \dots, N$. Using equations (10), (11), and (12) yields, after some manipulation, the partially linear regression model (e.g., [Robinson, 1988](#)):

$$Y_i = W_i' \delta + (X_i \otimes W_i)' \eta + (W_i \otimes \bar{X}_{p(i)})' \lambda + h(X_i, \bar{X}_{p(i)}) + A_i \quad (13)$$

with $\mathbb{E} [A_i | X_i, \bar{X}_{p(i)}, W_i] = 0$ for

$$A_i \stackrel{def}{=} [V_i - g_1(X_i, \bar{X}_{p(i)})] + \rho [\bar{V}_{p(i)} - g_2(X_i, \bar{X}_{p(i)})] + U_i, \quad (14)$$

and where the nonparametric regression component equals

$$h(X_i, \bar{X}_{p(i)}) \stackrel{def}{=} \alpha + X_i' \beta + g_1(X_i, \bar{X}_{p(i)}) + \bar{X}_{p(i)}' \gamma + \rho g_2(X_i, \bar{X}_{p(i)}) + (X_i \otimes \bar{X}_{p(i)})' \zeta. \quad (15)$$

Note, even under this perfect experiment, we cannot identify β , γ , and ζ ; these terms are confounded by $g_1(X_i, \bar{X}_{p(i)})$ and $g_2(X_i, \bar{X}_{p(i)})$ and hence absorbed into the nonparametric component of the regression model. This lack of identification reflects the inherent inability of the MET experiment to tell us anything about peer group effects. Nevertheless, as detailed below, knowledge of δ , η , and λ is sufficient to identify the class of reallocation effects we focus upon.

Patterns of non-compliance

Unfortunately, we do not observe student outcomes under full compliance. Non-compliance may induce correlation between A_i and X_i , $\bar{X}_{p(i)}$ and W_i in regression model (13). Our solution to this problem is to construct instrumental variables for *observed* teacher and peer attributes, W_i and $\bar{X}_{p(i)}$ – which necessarily reflect any non-compliance and attrition on the part of teachers and students – from the *assigned* values, W_i^* and $\bar{X}_{p^*(i)}$.

Rigorously justifying this approach requires imposing restrictions on how, for example, realized and assigned teacher quality relate to one another. The first assumption we make along these lines is:

¹⁹Reallocations leave the joint distribution of U_i and W_i unchanged, so we are free to normalize this mean to zero.

Assumption 1. (IDIOSYNCRATIC TEACHER DEVIATIONS)

$$\mathbb{E} [U_i - U_i^* | X_i, \bar{X}_{p^*(i)}, W_i^*] = 0. \quad (16)$$

Assumption 1 implies that the *difference* between *realized* and *assigned* (unobserved) teacher “quality” cannot be predicted by own and assigned peer and teacher observables. While Assumption 1 is not directly testable, we can perform the following plausibility test. Let $R_i - R_i^*$ be the difference between the realized and assigned value of some observed teacher attribute other than W_i (e.g., years of teaching experience). Under equation (16), if we compute the OLS fit of this difference onto 1, X_i , W_i^* , and $\bar{X}_{p^*(i)}$, a test for the joint significance of the non-constant regressors should accept the null of no effect. Finding that, for example, students *assigned* to classrooms with low average peer prior-year achievement, tend to *move into* classrooms with more experienced teachers suggests that Assumption 1 may be implausible.

Assumption 1 and equation (12) above yield the mean independence restriction

$$\mathbb{E} [U_i | X_i, W_i^*, \bar{X}_{p^*(i)}] = 0. \quad (17)$$

This equation imposes restrictions on the unobserved attribute of student i ’s *realized* teacher. It is this latent variable which drives the student outcome actually observed.

Our second assumption involves the relationship between the unobserved attributes of a student’s assigned peers and those of her realized peers. These two variables will differ if some students switch out of their assigned classrooms.

Assumption 2. (CONDITIONALLY IDIOSYNCRATIC PEER DEVIATIONS)

$$\mathbb{E} [\bar{V}_{p(i)} - \bar{V}_{p^*(i)} | X_i, \bar{X}_{p^*(i)}, W_i^*] = \mathbb{E} [\bar{V}_{p(i)} - \bar{V}_{p^*(i)} | X_i, \bar{X}_{p^*(i)}] \quad (18)$$

Assumption 2 implies that the *difference* between *realized* and *assigned* unobserved peer “quality” cannot be predicted by assigned teacher observables. We do allow for these deviations to covary with a student’s type and the assigned composition of her peers. Assumption 2 and equation (11) yield a second mean independence restriction of

$$\mathbb{E} [\bar{V}_{p(i)} | W_i^*, X_i, \bar{X}_{p^*(i)}] = g_2^* (X_i, \bar{X}_{p^*(i)}), \quad (19)$$

where $g_2^* (X_i, \bar{X}_{p^*(i)}) \stackrel{def}{=} \mathbb{E} [\bar{V}_{p(i)} - \bar{V}_{p^*(i)} | X_i, \bar{X}_{p^*(i)}] + g_2 (X_i, \bar{X}_{p^*(i)})$ is unrestricted. We can also pseudo-test Assumption 2 using observed peer attributes. Finding, for example, that – conditional on own type, X_i , and assigned peers’ average type, $\bar{X}_{p^*(i)}$ – assigned teacher quality, W_i^* , predicts differences between the realized and assigned values of other observed peer attributes provides evidence against Assumption 2.

The experiment-generated restrictions – equations (10), (11), and (12) above – in conjunction with our two (informally testable) assumptions about deviations from the experiment protocol – Assumptions 1 and 2 above – together imply the following conditional moment restriction:

$$\mathbb{E} [U_i + V_i + \rho \bar{V}_{p(i)} | W_i^*, X_i, \bar{X}_{p^*(i)}] = g_1 (X_i, \bar{X}_{p^*(i)}) + \rho g_2^* (X_i, \bar{X}_{p^*(i)}). \quad (20)$$

We wish to emphasize two features of restriction (20). First, the conditioning variables are *assigned* peer and teacher attributes, not their realized counterparts. This reflects our strategy of using assignment constructs as instruments. Second, any function of W_i^* , as well as interactions of such functions with functions of X_i and $\bar{X}_{p^*(i)}$ do not predict the composite error $U_i + V_i + \rho \bar{V}_{p(i)}$ conditional on X_i and $\bar{X}_{p^*(i)}$; hence such terms are valid instrumental variables.

More specifically we redefine h to equal

$$\begin{aligned} h (X_i, \bar{X}_{p^*(i)}, \bar{X}_{p(i)}) &\stackrel{def}{=} \alpha + X_i' \beta + g_1 (X_i, \bar{X}_{p^*(i)}) + \bar{X}_{p(i)}' \gamma \\ &\quad + \rho g_2^* (X_i, \bar{X}_{p^*(i)}) + (X_i \otimes \bar{X}_{p(i)})' \zeta \end{aligned} \quad (21)$$

and A_i to equal

$$A_i \stackrel{def}{=} (V_i - g_1 (X_i, \bar{X}_{p^*(i)})) + \rho (\bar{V}_{p(i)} - g_2^* (X_i, \bar{X}_{p^*(i)})) + U_i. \quad (22)$$

Equations (7), (21), and (22) yield an outcome equation of

$$Y_i = W_i' \delta + (X_i \otimes W_i)' \eta + (W_i \otimes \bar{X}_{p(i)})' \lambda + h (X_i, \bar{X}_{p^*(i)}, \bar{X}_{p(i)}) + A_i. \quad (23)$$

Condition (20) implies that A_i is conditionally mean zero given X_i , $\bar{X}_{p^*(i)}$ and W_i^* .

Summarizing, the experimentally-induced restrictions (10), (11), and (12), and our Assumptions 1 and 2 together imply that:

$$\mathbb{E} [A_i | X_i, \bar{X}_{p^*(i)}, W_i^*] = 0. \quad (24)$$

The estimation simplifies if we impose a restriction on the peer attrition/non-compliance that is similar to Assumption 2, but is on the observable peer average:

$$\mathbb{E} [\bar{X}_{p(i)} - \bar{X}_{p^*(i)} | X_i, \bar{X}_{p^*(i)}, W_i^*] = \mathbb{E} [\bar{X}_{p(i)} - \bar{X}_{p^*(i)} | X_i, \bar{X}_{p^*(i)}]. \quad (25)$$

This restriction is directly testable. By equation (25),

$$\mathbb{E} [\bar{X}_{p(i)} | X_i, \bar{X}_{p^*(i)}, W_i^*] = \mathbb{E} [\bar{X}_{p(i)} | X_i, \bar{X}_{p^*(i)}].$$

If this restriction holds the outcome equation is as in (23), but with a redefined nonparametric h that is a function of X_i and $\bar{X}_{p^*(i)}$ only,

$$h(X_i, \bar{X}_{p^*(i)}) \stackrel{def}{=} \alpha + X_i' \beta + g_1(X_i, \bar{X}_{p^*(i)}) + \mathbb{E}[\bar{X}_{p(i)} | X_i, \bar{X}_{p^*(i)}]' \gamma + \rho g_2^*(X_i, \bar{X}_{p^*(i)}) + (X_i \otimes \mathbb{E}[\bar{X}_{p(i)} | X_i, \bar{X}_{p^*(i)}])' \zeta, \quad (26)$$

and A_i is

$$A_i \stackrel{def}{=} (V_i - g_1(X_i, \bar{X}_{p^*(i)})) + \rho (\bar{V}_{p(i)} - g_2^*(X_i, \bar{X}_{p^*(i)})) + U_i + (\bar{X}_{p(i)} - \mathbb{E}[\bar{X}_{p(i)} | X_i, \bar{X}_{p^*(i)}])' \gamma + (X_i \otimes (\bar{X}_{p(i)} - \mathbb{E}[\bar{X}_{p(i)} | X_i, \bar{X}_{p^*(i)}]))' \zeta. \quad (27)$$

The conditional moment restriction in equation (24) also holds for this error.

Equations (23) and (24) jointly define a partially linear model with an endogenous parametric component. This is a well-studied semiparametric model (see, for example, [Chen et al., 2003](#)). The parameters δ , η , and λ are identified; $h(X_i, \bar{X}_{p^*(i)})$ is a nonparametric nuisance function.

We implement the partial linear IV estimator using the following approximation for $h(x, \bar{x})$:

$$h(X_i, \bar{X}_{p^*(i)}) \approx X_i' b + \bar{X}_{p^*(i)}' d + (X_i \otimes \bar{X}_{p^*(i)})' f.$$

For this approximation we estimate δ , η , and λ by linear IV fit of Y_i onto a constant, X_i , $\bar{X}_{p^*(i)}$, $(X_i \otimes \bar{X}_{p^*(i)})$, W_i , $(X_i \otimes W_i)$, and $(W_i \otimes \bar{X}_{p(i)})$ using the excluded instruments W_i^* , $(X_i \otimes W_i^*)$, and $(W_i^* \otimes \bar{X}_{p^*(i)})$. Note that both assigned and realized peer groups enter the main equation.

As in the case with perfect compliance, we do not identify β , γ , and ζ , again reflecting the inherent inability of the MET experiment to tell us anything about peer group effects. Nevertheless knowledge of δ , η , and λ is sufficient to identify the class of reallocation effects we focus upon.

3 Data and tests of identifying assumptions

The Measures of Effective Teaching (MET) study

The MET study was conducted during the 2009/10 and 2010/11 school years in elementary, middle, and high schools located in six large urban school districts in the United States.²⁰

²⁰The school districts are: Charlotte-Mecklenburg (North Carolina), Dallas Independent School District (Texas), Denver Public Schools (Colorado), Hillsborough County Public Schools (Florida), Memphis City Schools (Tennessee), and the New York City Department of Education (New York). The participation of school districts and schools in the MET study was voluntary.

The goal of the study was to examine determinants and consequences of teacher quality and teaching practices. In the first year of the study, the MET researchers collected detailed background information on teaching practices for each teacher as well as background performance measures for each student. In the second year of the study, MET researchers conducted a field experiment, randomly assigning teachers to classrooms within school-by-grade-by-subject cells (“randomization blocks”). The randomization blocks typically consisted of two to three classrooms each. Classroom composition was not manipulated as part of the study. For more details about the study design, see [White et al. \(2019\)](#) and [Kane et al. \(2013\)](#).

The dataset contains detailed information on the students, their classroom teachers, and their classroom peers. As our key measure of teaching practices, we use Danielson’s (2011) “Framework for Teaching” (FFT) measure, collected at baseline (school year 2009/10). The FFT is designed to provide feedback to teachers, but has also been used to study the impact of teaching practices on student outcomes (c.f., [Garrett and Steinberg, 2015](#), [Aucejo et al., 2019](#)). The FFT is an observational measure: teachers are video-taped several times during a school year and subsequently rated – using a specially designed rubric – by trained raters (often former teachers).

The FFT is positively correlated with a range of teacher quality measures. Using the teacher variables available in the MET data, we find that the FFT is positively and statistically significantly correlated with three different measures of teaching practices (“Classroom Assessment Scoring System (CLASS)”, “Mathematical Quality of Instruction (MQI)”, “Protocol for Language Arts Teaching Observations (PLATO)”, see [White et al. \(2019\)](#) for details), principal ratings, teacher value added, and content knowledge in ELA. By contrast, it is not significantly correlated with teacher experience, teacher qualification (whether the teacher holds a Master’s degree), and content knowledge in math (see Appendix Table B.2).

For our main outcome measure we use students’ 2010/11 end-of-year standardized state test scores in the subjects mathematics and English language arts (ELA). The dataset also includes background characteristics from school district records for students (age, gender, race/ethnicity, special-education status, free/reduced-price lunch eligibility, gifted status, and whether a student is an English language learner) and for teachers (education, teaching experience in the district). Through a section identifier, we can match a student to her classroom peers.

Sample and summary statistics

In constructing our estimation sample, we closely follow [Garrett and Steinberg \(2015\)](#), who investigate the impact of FFT on students’ test score outcomes in the MET study. We restrict our sample to all elementary- and middle-school students (grades 4–8) who took part in the randomization. Furthermore, we only include students with non-missing information

on baseline and final test score outcomes as well as students with non-missing information on the characteristics of the assigned and realized teacher, and the characteristics of the assigned and realized classroom peers. Details of the data construction from the MET data files are provided in Appendix B.

Our final sample consists of about 8,500 students and 614 teachers in math and of about 9,600 students and 649 teachers in ELA.²¹ The majority of the students (60 percent) are elementary school students (grades 4–5), and the remaining students are middle school students (grades 6–8). Nearly 60 percent of the students are economically disadvantaged, that is, eligible for free/reduced-price lunch.²² The student sample is diverse, with equal proportions of white, black, and Hispanic students.

Student test scores are centered at the relevant district-level mean and standardized by the relevant district-level standard deviation. MET students exceed their district average test score by about 0.15 standard deviations on average, perhaps because the participation of schools in the MET study was voluntary (see Kane et al., 2013). The majority of the teachers are female (about 85 percent), and about two-thirds of them are white. Most of the teachers have substantial teaching experience – seven years on average – and about 40 percent of the teachers have graduated from a Master’s program.²³

Teaching practices

The MET data includes ratings of teaching practices for two domains of the FFT, “Classroom Environment” and “Instruction.” These domains are divided into four components each,²⁴ and each component is rated on a four-point scale by each rater (unsatisfactory, basic, proficient, distinguished), with high scores in a category indicating that a teacher is closer to an ideal teaching practice according to the FFT. The teachers are video-taped at least twice during the school year during lessons, and each video is rated independently by at least two raters. In our analysis, we average the scores across all videos, raters, components, and domains. This aggregation of the scores is the most common one in the literature (cf., Garrett and Steinberg, 2015, Aucejo et al., 2019).

The FFT is on average 2.5 in math and 2.6 in ELA, which corresponds to a rating between “basic” and “proficient”; for the purpose of our analysis, we create three categories of FFT (see Appendix Figure A.1). We set the cutoffs at FFTs of 2.25 and 2.75. In our sample, low-FFT teachers have “basic” teaching practices on average (average FFT of 2.1 in both math

²¹Appendix Table A.1 summarizes all student and teacher variables.

²²This variable is missing for 30 percent of the students in the math sample and for 25 percent of students in the ELA sample.

²³Teaching experience and teacher education are missing for about 30 percent of teachers.

²⁴The components of the domain “Classroom Environment” are: creating an environment of respect and rapport; establishing a culture for learning; managing classroom procedures; managing student behavior. The components of the domain “Instruction” are: communicating with students; using questioning and discussion techniques; engaging students in learning; using assessment in instruction.

and ELA), and high-FFT teachers have “proficient” teaching practices on average (average FFT of 2.9 in both math and ELA); 18 percent of the teachers in math and 14 percent of the teachers in ELA are classified as having a low FFT, 62 percent of the teachers in math and 58 percent of the teachers in ELA are classified as having a middle FFT, and 20 percent of the teachers in math and 28 percent of the teachers in ELA are classified as having a high FFT. Our results are not sensitive to the exact position of the cutoffs, and we investigate a model with two or four, instead of three, FFT categories in our sensitivity checks below. This categorization also creates variation in FFT levels within randomization blocks: in our sample, 65 percent of randomization blocks in math, and 70 percent of randomization blocks in ELA include teachers with different levels of FFT (for example, both a low- and a middle-FFT teacher).

Test score outcomes

The end-of-year test scores provided in the MET data are z -scores, i.e., they are standardized such that district-wide, they will be mean zero with unit standard deviation. Since MET schools are not representative at the district level, the mean test scores may – and do – deviate from zero. We use the 2010/11 z -score as our outcome variable.

To identify reallocation effects, we split 2009/10 baseline test scores into three bins, corresponding to terciles of the within-district z -score distribution. In our estimation sample, 27 percent of the students in math, and 26 percent in ELA, have “low” baseline test scores, 36 percent of the students in math, and 35 percent of the students in ELA, have “middle” baseline test scores, and 36 percent of the students in math, and 39 percent of the students in ELA, have “high” baseline test scores.²⁵ To include classroom peers into the analysis, we compute the fraction of each student’s classmates with high, middle, and low baseline test scores (leave-own-out means).

Non-compliance

As outlined in Section 2, some teachers and students switched classrooms or schools before the start of the school year, which we take into account in our identification strategy. In our sample, 69 percent of the students in math, and 73 percent of the students in ELA, are actually taught by their randomly assigned teachers. This level of non-compliance is high enough to make analyses which ignore it potentially problematic.

We also observe changes in classroom peers after the randomization but before the beginning of the school year. Changes in assigned peers were driven by students who leave schools, repeat a grade or, in some cases, by the need to adjust their schedule. On average,

²⁵These numbers are not exactly 33 percent in each category, because the students in our sample have on average higher baseline test scores than the full student population in a district.

students’ realized peers, however, are not appreciably different from their assigned ones. The difference in baseline z -scores between the assigned and the realized peers amounts to just 0.02 standard deviations on average both in math and in ELA.

Tests of identifying assumptions and restrictions

In this section, we report the results of the series of tests discussed earlier; specifically tests designed to assess the plausibility of Assumptions 1 and 2. These two assumptions impose restrictions on the nature of non-compliance by students and teachers. We further *directly* test restriction (25), which is a restriction on the nature of non-compliance by peers. We also assess the quality of the initial MET randomization of teachers to classrooms.

To test whether the randomization was successful in balancing student characteristics across teachers with different levels of FFT we regress the FFT of a student’s assigned teacher on the student’s own characteristics, controlling for randomization block fixed effects. None of the student characteristics predict assigned teacher’s FFT, individually or jointly, which confirms that the randomization indeed “worked” (see Appendix Table A.2).

Covariate balance, however, is not a sufficient condition to identify reallocation effects under non-compliance by both students and teachers (see Section 2). Assumption 1 requires that own and assigned peer and teacher observables should not predict the difference between realized and assigned unobserved teacher quality. Since this assumption involves a statement about unobserved variables, we cannot test it directly. Instead we “test” it indirectly as described in Section 2. Specifically we use those teacher background characteristics that are not part of the the model as replacements for the unobserved quality of a teacher: a teacher’s demographics, experience, and education. We regress the difference between realized and assigned teacher characteristics on the student’s baseline test score, the FFT of the assigned teacher, and the average baseline test score of the assigned peers. Consistent with Assumption 1, these variables do not jointly predict differences between the characteristics of the assigned and realized teacher in any of the regression fits (see Appendix Table A.3).

Assumption 2 states that differences between realized and assigned unobserved peer quality cannot be predicted by assigned teacher observables, conditional on own baseline achievement and the baseline achievement of assigned peers. Again, we can only perform an indirect test of this assumption. To do so we regress differences between the assigned and realized characteristics of classroom peers onto the FFT of the assigned teacher, controlling for own baseline test scores and assigned peers’ average baseline test scores. We do not find, consistent with Assumption 2, that teacher FFT predicts differences between the characteristics of the assigned and realized peers (see Appendix Tables A.4 and A.5).

Finally, we directly assess restriction (25), which implies that the baseline test scores of realized peers cannot be predicted by assigned teacher FFT. We test this restriction by

regressing the baseline test scores of realized peers onto a student’s baseline test score, the baseline test scores of her assigned peers, as well as the FFT of her assigned teacher. We find that this restriction also holds in our data (see Appendix Table A.6).

4 Computing reallocation effects

To construct estimates of the average effect of reassigning teachers across classrooms we proceed in three steps. First, as outlined in Section 2, we estimate (a subset of) the parameters of the educational production function. Second, we feed these parameters into a linear programming problem to compute an optimal assignment (i.e., an assignment that maximizes aggregate achievement). Third, we compare the aggregate outcome under the optimal assignment to the aggregate outcome under the status quo assignment, as well as to that under a worst assignment (i.e., one which minimizes aggregate achievement). These comparisons provide a sense of the magnitude of achievement gains available from improved teacher assignment policies.

Because the status quo assignment in the MET experiment is generated by randomization, one might worry that this not a realistic point of comparison. To address this concern, we show that the joint distribution of FFT and student baseline achievement in the year prior to the experiment – a year without randomization – looks very similar to the joint distribution of FFT and student baseline achievement under random assignment (see Appendix Figure A.2).

Estimation of the education production function

The specification that we use to estimate the parameters of the educational production function coincides with equation (23). Since randomization was carried out within randomization blocks, we additionally include randomization block fixed effects in this regression model. We then estimate the model’s parameters by the method of instrumental variables (IV) as described in Section 2. Tests for instrument relevance following Sanderson and Windmeijer (2016) confirm that our IV estimates are unlikely to suffer from weak-instrument bias (see Appendix Table A.7).

Computing and characterizing an optimal assignment

The aim of the analysis is to find an assignment of teachers to classrooms that improves student outcomes. What is meant by an “optimal” assignment depends, of course, on the objective function.

In this paper we choose to maximize aggregate outcomes (i.e., the sum of all students’ test scores). This is the “simplest” objective we can consider, it is straightforward to compute,

justifiable from a utilitarian perspective, and easy to interpret. We would like to emphasize, however, that our analysis can be modified to accommodate other objective functions – in some cases easily, in others with more difficulty. In practice, for example, the social planner may care about both the aggregate outcome as well as inequality, especially across identifiable subgroups.

Our intention is not to advocate for maximization of the aggregate outcome in practice, although doing so is appropriate in some circumstances; rather this objective provides a convenient starting point and makes our analysis comparable to that of other educational policy evaluations (e.g., the typical regression-based class-size analysis provides an estimate of the effect of class-size on *average* achievement).

To determine the optimal allocation, we feed the *estimated* parameters from equation (23), specifically $\hat{\delta}$, $\hat{\eta}$, and $\hat{\lambda}$, into a linear program. Given these parameters, we can compute, for each student, her predicted outcome when taught by a low-, middle-, or high-FFT teacher, leaving the original classroom composition unchanged.

Formally, for each student i , we compute three counterfactual outcomes $\hat{Y}_i(w)$, $w \in \{w_L, w_M, w_H\}$, corresponding to an assignment to a low-, middle- or high-FFT teacher. Aggregated to the classroom level this yields, in an abuse of notation, three counterfactual classroom-level test score aggregates, $\hat{Y}_c(w) = \sum_{i \in s(c)} \hat{Y}_i(w)$, with $w \in \{w_L, w_M, w_H\}$ (here $s(c)$ denotes the set of indices for students in classroom c).

By aggregating the counterfactual outcomes to the *classroom* level, we transform the assignment problem from a many-to-one matching problem to a one-to-one matching problem. This approach is only suitable because the configuration of students across classrooms remains fixed; we only consider the effects of reassigning teachers across existing classrooms of students.²⁶ The one-to-one matching problem is a special linear program, a transportation problem, which is easily solvable.²⁷

There are a few additional constraints we impose to make the reallocation exercise realistic. First, we do not allow teachers to be reassigned across districts or across school types (i.e., elementary or middle school). We also present, as a sensitivity check, a version of the allocation where we only allow teachers to be reassigned within their randomization block. Second, there are a few teachers that teach several sections of a class. In this case, we treat these sections as clusters, and allocate one teacher to all sections in each such cluster.

When computing counterfactual outcomes for each student, we omit the nonparametric component, $h(x, \bar{x})$, as well as the randomization block fixed effects and the main effects of own baseline achievement. This is appropriate because we report the difference between the sum of aggregate outcomes for two allocations where the nonparametric component, the fixed

²⁶As noted earlier the effects of classroom composition changes on student achievement are unidentifiable here in any case.

²⁷Appendix C contains the details on how we specify the transportation problem. To compute the optimal assignment, we use the function `lpSolve` in *R*.

effects, and the main effects of own baseline achievement cancel out.

In addition to the optimal assignment, we also compute a “worst” possible assignment, i.e., an assignment that minimizes aggregate test scores. The difference between the aggregate outcome for the best and worst assignment is the maximal reallocation gain. This provides an upper bound on the magnitude of student achievement gains that teacher reassignments based on FFT and prior student achievement might yield in practice.

Average Reallocation Effects

An individual reallocation effect, or reallocation gain, is defined as the difference between an individual student’s outcome under two assignments. For example, one can compute, for each student, $\widehat{Y}_i(\widetilde{W}_i^{opt})$, i.e., the predicted outcome under the optimal allocation, where \widetilde{W}_i^{opt} takes the values w_L , w_M , or w_H , depending on whether the student would be assigned to a low-, middle-, or high-FFT teacher in an optimal allocation. Similarly, one can compute the same parameter for each student under the status quo, which we denote as $\widehat{Y}_i(W_i)$. In MET schools the status quo assignment was induced by random assignment of teachers to classrooms within randomization blocks. The allocation of teachers across schools within a given district was, of course, non-random, and possibly non-optimal from the standpoint of maximizing the aggregate outcome. An individual reallocation gain can then be computed as the difference between these two outcomes, $\widehat{Y}_i(\widetilde{W}_i^{opt}) - \widehat{Y}_i(W_i)$. The optimal assignment is not the assignment that maximizes the predicted outcome for student i , but the assignment that maximizes the aggregate outcome across all the classrooms in the sample (subject to the various constraints we impose on the problem – like ruling out reassignments across school districts).

These individual gains can be aggregated in many ways in order to create policy-relevant parameters. We start by defining our key parameter of interest, the *average reallocation effect*, as

$$\widehat{ARE} = \frac{1}{N} \sum_{i=1}^N \left[\widehat{Y}_i(\widetilde{W}_i^{opt}) - \widehat{Y}_i(W_i) \right]. \quad (28)$$

We also consider conditional average reallocation effects, that is, average reallocation effects for students with varying baseline characteristic x (e.g., students with low, middle, and high baseline test scores):

$$\widehat{ARE}(x) = \frac{1}{\sum_{i=1}^N \mathbf{1}(X_i = x)} \sum_{i=1}^N \mathbf{1}(X_i = x) \left[\widehat{Y}_i(\widetilde{W}_i^{opt}) - \widehat{Y}_i(W_i) \right]. \quad (29)$$

Below we show that only about one-half of MET students would experience a teacher change when switching from the status quo teacher assignment to an optimal one. The

reallocation effect for students in classrooms not assigned a different teacher is, of course, zero. This motivates a focus on the average achievement gain among those students who *are* assigned a different teacher. We define the reallocation effect *conditional on being reassigned* as:

$$\widehat{AREC}(x) = \frac{1}{\sum_{i=1}^N \mathbf{1}(X_i = x) \mathbf{1}(W_i \neq \tilde{W}_i^{opt})} \times \sum_{i=1}^N \mathbf{1}(X_i = x) \mathbf{1}(W_i \neq \tilde{W}_i^{opt}) \left[\widehat{Y}_i(\tilde{W}_i^{opt}) - \widehat{Y}_i(W_i) \right]. \quad (30)$$

We can, of course, not condition on x as well.

Finally, the maximal reallocation gain is obtained if we compare the average outcome under an optimal assignment with that under a worst assignment

$$\widehat{ARE}^{max}(x) = \frac{1}{\sum_{i=1}^N \mathbf{1}(X_i = x)} \sum_{i=1}^N \mathbf{1}(X_i = x) \left[\widehat{Y}_i(\tilde{W}_i^{opt}) - \widehat{Y}_i(\tilde{W}_i^{worst}) \right], \quad (31)$$

where $\widehat{Y}_i(\tilde{W}_i^{worst})$ is the predicted outcome for student i under an allocation that minimizes aggregate test scores.

Inference on reallocation effects

We use the Bayesian bootstrap to quantify our (posterior) uncertainty about AREs. We treat each teacher-classroom pair as an i.i.d. draw from some unknown (population) distribution. Following [Chamberlain and Imbens \(2003\)](#), we approximate this unknown population by a multinomial, to which we assign an improper Dirichlet prior. This leads to a posterior distribution which (i) is also Dirichlet and (ii) conveniently only places probability mass on data points observed in our sample.

Mechanically, we draw C standard exponential random variables and weigh each student in section/classroom $c = 1, \dots, C$ with the c^{th} weight (i.e., all students in the same section have the same weight). Using this weighted sample, we compute the IV regression fit, solve for the optimal (worst) assignment, and compute the various reallocation effects. We repeat this procedure 1,000 times. This generates 1,000 independent draws from the posterior distribution of the ARE.

Formally this approach to inference is Bayesian. Consequently the “standard errors” we present for our ARE estimates summarize dispersion in the relevant posterior distribution (not variability across repeated samples).²⁸ An alternative, frequentist, approach to inference

²⁸The “standard errors” for the AREs are constructed as follows. Let $\hat{\theta}^{(b)}$ be the estimate of the reallocation effect in b^{th} the bootstrap sample (or equivalently the b^{th} draw from the posterior distribution for θ) and

is provided by [Hsieh et al. \(2018\)](#). They transform the problem of inference on the solution to a linear program into inference on a set of linear moment inequalities. If the binding constraints are the same over the bootstrap distribution, then inference based on the Bayesian bootstrap will be similar to that based on moment inequalities (see also [Graham et al., 2007](#) and [Bhattacharya, 2009](#)).

We wish to emphasize that we make no claims about the frequentist properties of our procedure. Given that our primary exercise involves solving a social planning problem, the Bayesian approach – which involves summarizing the posterior distribution of AREs – is both principled as well as convenient.

5 Results

Regression results

Before reporting parameter estimates for our preferred model – equation (23) above with three categories of FFT, own and peer baseline achievement each – we present those for a more conventional model with teacher FFT, student and peer average baseline test scores all entering linearly. For this specification we leave the FFT and baseline test scores undiscretized. These initial results replicate and expand upon the prior work of [Garrett and Steinberg \(2015\)](#), who study the impact of FFT on student achievement in (approximately) the same sample.²⁹

These initial results indicate that teacher FFT does not affect students’ test scores on average (see Table 2, Panel A). This is consistent with the findings of [Garrett and Steinberg \(2015\)](#). Table 2 presents estimates, using the instrumental variables described in Section 2. Here this involves using the FFT score of a student’s randomly *assigned* teacher as an instrument for that of her *realized* teacher (Panels A–C), and the average baseline test score of her assigned peers as an instrument for the average baseline test score of her realized peers (Panel B).

OLS estimates of the same model are reported in Appendix Table A.8. The coefficient on FFT is statistically significant and positive in the OLS results. The discrepancy between the OLS and IV results may reflect the impact of correcting for non-compliance, as described above, or simply reflect the greater sampling variability of the IV estimates.

Next we add the interactions of teacher FFT with both the baseline student, and peer

$\hat{\theta}$ the reallocation effect in the original sample. Consider the centered statistic $\hat{\theta}^{(b)} - \hat{\theta}$; let $q(0.025)$ and $q(0.975)$ be the 0.025 and 0.975 quantiles of its bootstrap/posterior distribution. We then construct the interval $[\hat{\theta} - q(0.975), \hat{\theta} - q(0.025)]$ (see [Hansen, 2020](#)). We compute the corresponding standard error by dividing this interval by $2\Phi^{-1}(0.975) \approx 3.92$, where $\Phi(\cdot)$ is the standard normal CDF.

²⁹[Aucejo et al. \(2019\)](#) study the impact of FFT in the MET data, but focus on math classrooms in elementary schools.

Table 2: IV regression results of the linear model. Dependent variables: student test score outcomes

		(1)	(2)	(3)	(4)	(5)	(6)
		A. Only teacher effects		B. Full model		C. Without teacher \times peer interactions	
		Math	ELA	Math	ELA	Math	ELA
δ	FFT	0.013 (0.093)	0.079 (0.076)	-0.004 (0.091)	0.104 (0.075)	-0.004 (0.091)	0.076 (0.075)
η	FFT \times baseline			0.096** (0.042)	0.073* (0.039)	0.098** (0.046)	0.049 (0.040)
λ	FFT \times avg. peer baseline			0.016 (0.103)	-0.172 (0.197)		
β	Baseline	0.749*** (0.011)	0.693*** (0.011)	0.505*** (0.109)	0.512*** (0.099)	0.498*** (0.119)	0.565*** (0.103)
R^2		0.700	0.632	0.700	0.621	0.700	0.632
N		8,534	9,641	8,534	9,641	8,534	9,641

Note: The dependent variables are subject-specific test score outcomes in math and ELA. The specifications include linear terms for FFT, individual and peer baseline test scores. The instrumental variables are based on assigned teacher FFT (Panels A–C) and assigned peer baseline test scores (Panel B). All regressions control for the $h(x, \bar{x})$ function (see Section 2) and for randomization block fixed effects. Analytic standard errors, clustered by randomization block, are in parentheses.

** significant at the 1%-level ** significant at the 5%-level * significant at the 10%-level.

average, test scores. This provides an initial indication of whether any complementarity between teacher FFT and student baseline achievement is present. These results suggest that high-FFT teachers are more effective at teaching students with higher baseline scores (see Table 2, Panels B and C). This result is significant at the 5-percent level in the math sample, but only weakly significant in the ELA sample. The magnitudes and precision of the teacher-student match effects are similar across the IV and OLS estimates (the latter reported in Appendix Table A.8). The coefficient on the interaction between teacher FFT and peer average baseline achievement (Panel B) is poorly determined (whether estimated by IV or OLS).

In Table 3 we report IV estimates of our preferred 3×3 model. This corresponds to equation (23) with both W_i and X_i consisting of dummy variables for middle and high FFT and baseline test scores respectively. In this specification, the teacher FFT main effects are insignificant. We do, however, find positive match effects between teacher FFT and student baseline scores. These are significant for the high-FFT teachers (i.e., teachers with a “proficient” score on average). Students with middle or high baseline test score levels score significantly higher on end-of-year achievement tests when matched with a high-FFT teacher, compared to students with low baseline test scores (see Table 3, Panels B and C).³⁰

³⁰In contrast, the coefficients on interactions of the teacher FFT dummy variables with the peer composition

Table 3: IV regression results of the 3×3 model. Dependent variables: student test score outcomes

		(1)	(2)	(3)	(4)	(5)	(6)
		A. Only teacher effects		B. Full model		C. Without teacher \times peer interactions	
		Math	ELA	Math	ELA	Math	ELA
δ	FFT middle	0.069 (0.053)	-0.029 (0.048)	-0.145 (0.138)	-0.219 (0.173)	0.027 (0.060)	-0.082 (0.059)
	FFT high	0.038 (0.067)	-0.058 (0.065)	-0.547 (0.380)	-0.137 (0.203)	-0.155 (0.103)	-0.154* (0.083)
η	FFT middle						
	\times baseline middle			0.040 (0.060)	0.057 (0.067)	0.052 (0.060)	0.067 (0.063)
	\times baseline high			0.015 (0.084)	0.076 (0.081)	0.050 (0.076)	0.097 (0.082)
	FFT high						
	\times baseline middle			0.184** (0.082)	0.150** (0.074)	0.196** (0.084)	0.127* (0.072)
	\times baseline high			0.226** (0.099)	0.187** (0.092)	0.265** (0.100)	0.149 (0.095)
λ	FFT middle						
	\times fraction peers middle			0.318 (0.299)	0.288 (0.401)		
	\times fraction peers high			0.239 (0.205)	0.161 (0.288)		
	FFT high						
	\times fraction peers middle			0.641 (0.513)	0.227 (0.389)		
	\times fraction peers high			0.460 (0.409)	-0.355 (0.343)		
β	Baseline middle	0.888*** (0.077)	0.739*** (0.084)	0.843*** (0.092)	0.699*** (0.096)	0.840*** (0.092)	0.682*** (0.094)
	Baseline high	1.622*** (0.103)	1.555*** (0.105)	1.578*** (0.124)	1.503*** (0.118)	1.550*** (0.119)	1.470*** (0.115)
R^2		0.617	0.554	0.616	0.552	0.617	0.555
N		8,534	9,641	8,534	9,641	8,534	9,641

Note: The dependent variables are subject-specific test score outcomes in math and ELA. The instrumental variables are based on assigned teacher FFT (Panels A–C) and assigned peer baseline test scores (Panel B). All regressions control for the $h(x, \bar{x})$ function (see Section 2) and for randomization block fixed effects. Analytic standard errors, clustered by randomization block, are in parentheses.

*** significant at the 1%-level ** significant at the 5%-level * significant at the 10%-level.

The OLS estimation results, reported in the appendix, are qualitatively similar to the IV ones; although the interactions of the middle-FFT dummy variable with the middle- and high-baseline student test score dummies are also significant when estimated by OLS (see Appendix Table A.9).

To compute optimal assignments and average reallocation effects, we use the IV estimates presented in columns 5–6 of Table 3. These specifications omit the FFT-by-peer composition interaction terms, whose coefficients are poorly determined in all specifications (whether fitted by IV or OLS). Omitting these terms has little effect on either the location or the precision of the coefficients on the FFT-by-baseline interactions. In Appendix Tables A.15 and A.16 we also present average reallocation effects based upon the specifications in columns 3–4 of Table 3. These ARE estimates are larger, albeit less precisely determined.

Characterization of an optimal reassignment

If an optimal assignment is similar to the status quo one, then large reallocation effects are unlikely. Therefore, before quantifying any reallocation effects, we first discuss how an optimal assignment and a worst assignment differ from the status quo.

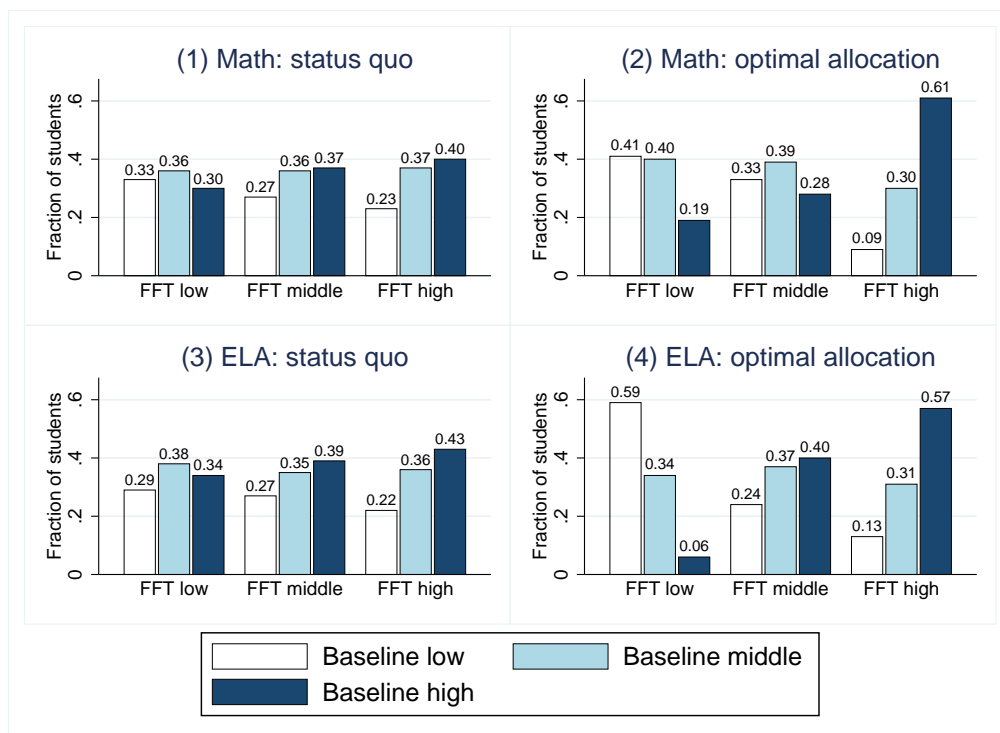
Our primary reassignment policy considers reassignments of teachers across classrooms district-wide (that is, we allow teachers to move to a different school within their district). We do restrict teachers to teach at the same level (e.g., elementary school teachers do not move to middle school classrooms). Under this scenario we find that, when moving from the status quo to an optimal assignment, 49 percent of the students in the math sample, and 47 percent of the students in the ELA sample, are assigned to a new teacher. The balance of students remain with their status quo teacher.

It is interesting to examine how the reallocation changes the “assortativeness” of the assignment. An assignment is characterized as positive assortative if students with higher baseline test scores are more frequently matched with higher-FFT teachers, compared to a random assignment. Indeed, we observe this in our data for the optimal assignment. Figure 1 displays, for each level of FFT, the distribution of student baseline test scores in the average class a teacher of that FFT level is assigned to. In both the math and in the ELA sample, the optimal allocation is more assortative than the status quo. In the status quo in math, for instance, a high-FFT teacher has on average 23 percent of students with low baseline test scores and 40 percent of students with high baseline test scores. In the optimal allocation, the fraction of students with low baseline test scores drops to 9 percent, and the fraction of students with high baseline test scores increases to 61 percent on average. The optimal allocation in the ELA sample displays a similar pattern. The positive assortativeness arises in the optimal allocation because students with high baseline test scores benefit more from

variables are imprecisely determined.

a high-FFT teacher, compared to students with low baseline test scores (see Table 3). The worst allocation displays negative assortativeness (see Appendix Figure A.3).

Figure 1: Assortativeness of the optimal allocation in comparison with the status quo



Note: The figure compares the status quo allocation (Panels 1 and 3) with the optimal allocation (Panels 2 and 4). The bars represent the fractions of students with low, middle, and high baseline test scores that a teacher with low, middle, and high FFT is assigned to on average under each of the allocations. For each teacher type, the fractions add up to 1. The optimization is carried out within school types (elementary or middle school) and districts.

Average reallocation effects

Tables 4 and 5 present ARE estimates. In Panel A an optimal assignment is compared with the status quo; while in Panel B optimal assignments are compared with “worst” assignments. We presents results for all students (Panels A.I and B.I), as well as for just those students who experience a teacher change as part of the reallocation (Panels A.II and B.II).

In the math sample (Table 4), the optimal allocation improves average test scores by 1.7 percent of a test score standard deviation compared to the status quo. This effect is precisely determined with a Bayesian bootstrap standard error of 0.6 percent. The reported effects are largely driven by students with high baseline test scores. These students gain 2.8 percent of a test score standard deviation on average; students with middle and low baseline test scores, in contrast, gain only 1.2 and 1.4 percent of a test score standard deviation, respectively (on average).

Since only half of the students experience a change in their teacher, the average effect represents an equal-weighted mixture of a zero effect and a positive effect on those students who do experience a change in teachers. The average effect for the latter group is 3.6 percent of a test score standard deviation (Panel A.II, SE = 1.2 percent).

Table 4: Average reallocation gains in math. Gains expressed in test score standard deviations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Optimal versus status quo								
	A.I Full sample				A.II Conditional on being reallocated			
	all students	high	middle	low	all students	high	middle	low
Gain	0.017	0.028	0.012	0.014	0.036	0.059	0.028	0.026
SE	(0.006)	(0.014)	(0.008)	(0.011)	(0.012)	(0.029)	(0.019)	(0.019)
N	8,534	2,332	3,108	3,094	4,107	1,121	1,380	1,606
Panel B. Optimal versus worst allocation								
	B.I Full sample				B.II Conditional on being reallocated			
	all students	high	middle	low	all students	high	middle	low
Gain	0.040	0.072	0.019	0.038	0.060	0.103	0.032	0.053
SE	(0.012)	(0.034)	(0.015)	(0.026)	(0.019)	(0.048)	(0.022)	(0.037)
N	8,534	2,332	3,108	3,094	5,746	1,626	1,912	2,208

Note: The table shows the average reallocation gains from implementing the optimal assignment instead of the random assignment (status quo) in Panel A, and the average reallocation gains from implementing the optimal assignment instead of the worst assignment in Panel B. The gains are expressed in test score standard deviations. The computations are based on a 3×3 model without teacher-by-peer interactions (see Table 3, column 5). The results are presented separately for the full sample of students (columns 1–4), and for the sample of classrooms that would get a new teacher as a result of the reallocation (columns 5–8). The reassignments are carried out within school types and districts. Standard errors are in parentheses and computed using the Bayesian bootstrap with 1,000 replications (see Section 4). High/middle/low: students in the top/middle/bottom tercile of the baseline test score distribution.

The comparison of an optimal allocation with a worst allocation yields improvements that are about twice as large. Relative to a worst allocation, an optimal allocation improves test scores by 4.0 percent of a standard deviation on average (SE = 1.2 percent). The gains are 7.2 percent of a standard deviation for students with high baseline test scores, and 1.9 and 3.8 percent of a standard deviation for students with middle and low baseline test scores, respectively. If one considers only those students who are reassigned to a new teacher, the reallocation effect amounts to 6.0 percent of a standard deviation on average (Panel B.II, SE = 1.9 percent).

One way to benchmark the magnitude of these AREs is to compare them with the effects of hypothetical policies aimed at improving teacher value-added measures (VAMs). As noted in the introduction, such policies are controversial, as is the evidence marshalled to support them. Here we offer no commentary on the advisability of actually adopting VAM-guided

teacher personnel policies; nor do we offer an assessment of VAM studies. Rather we simply use these studies, and the policy thought experiments they motivate, to benchmark our ARE findings.

Teacher value-added is typically conceptualized as an invariant intercept-shifter, which uniformly raises or lowers the achievement of all students in a classroom. In this framework replacing a low value-added teacher with a high one will raise achievement for all students in a classroom. [Rockoff \(2004\)](#) estimates that the standard deviation of the population distribution of teacher value-added (in a New Jersey school district) is around 0.10 test score standard deviations in both math and reading. Recent studies find somewhat higher estimates: [Chetty et al. \(2014a\)](#) estimates that the standard deviation of teacher value-added is 0.16 in math and 0.12 in reading; similarly, [Rothstein \(2017\)](#) finds values of 0.19 in math and 0.12 in reading.

Using a standard deviation of 0.15 we can consider the effect of a policy which removes the bottom $\tau \times 100$ percent of teachers, sorted by VAM, and replaces them with teachers at the $\tilde{\tau}^{th}$ quantile of the VAM distribution. Under normality the effect of such a policy on average student achievement is to increase test scores by

$$(1 - \tau) \sigma \frac{\phi\left(\frac{q_\tau}{\sigma}\right)}{1 - \Phi\left(\frac{q_\tau}{\sigma}\right)} + \sigma \Phi^{-1}(\tilde{\tau})$$

standard deviations. Setting $\tau = 0.05$ and $\tilde{\tau} = 0.75$ this expression gives an estimate of the policy effect of 0.021 (i.e., 2.1 percent of a test score standard deviation). This is comparable to the average effect on math achievement associated with moving from the status quo MET assignment to an optimal one. In practice correctly identifying, and removing from classrooms, the bottom five percent of teachers would be difficult to do. Replacing them with teachers in the top quartile of the VAM distribution perhaps even more so. Contextualized in this way the AREs we find are large.

An attractive feature of the policies we consider is that they are based on measurable student and teacher attributes, not noisily measured latent ones. At the same time we are mindful that most school districts would not find it costless to reallocate teachers freely across classrooms and schools.

Another way to calibrate the size of the effects we find is as follows. We find that when implementing an optimal teacher-to-classroom assignment, only about one half of students experience a change in teachers. We find that test scores increase by about 3.6 percent of a standard deviation for *these* students. This is comparable to increasing the average VAM of these students' teachers from zero (i.e., the median) to the 0.6 quantile of the teacher VAM distribution. Again, increasing the VAM of half of all teachers by such an amount may be

difficult to do in practice.³¹

For English language arts (ELA) achievement we find smaller reallocation effects. Moving from the status quo to an optimal allocation is estimated to raise achievement by 0.8 percent of a test score standard deviation (SE = 0.6 percent). As with math, these gains are concentrated among students with high baseline scores who are assigned a new teacher. These students experience an average gain of 5 percent of a test score standard deviation (SE = 2.5 percent).

Table 5: Average reallocation gains in ELA. Gains expressed in test score standard deviations.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Optimal versus status quo								
	A.I Full sample				A.II Conditional on being reallocated			
	all students	high	middle	low	all students	high	middle	low
Gain	0.008	0.024	0.002	0.004	0.017	0.050	0.004	0.008
SE	(0.006)	(0.011)	(0.007)	(0.010)	(0.012)	(0.025)	(0.013)	(0.020)
N	9,641	2,480	3,402	3,759	4,529	1,167	1,627	1,735
Panel B. Optimal versus worst allocation								
	B.I Full sample				B.II Conditional on being reallocated			
	all students	high	middle	low	all students	high	middle	low
Gain	0.018	0.057	0.003	0.005	0.025	0.080	0.004	0.007
SE	(0.010)	(0.028)	(0.011)	(0.019)	(0.015)	(0.038)	(0.017)	(0.024)
N	9,641	2,480	3,402	3,759	6,675	1,770	2,229	2,676

Note: The table shows the average reallocation gains from implementing the optimal assignment instead of the random assignment (status quo) in Panel A, and the average reallocation gains from implementing the optimal assignment instead of the worst assignment in Panel B. The gains are expressed in test score standard deviations. The computations are based on a 3×3 model without teacher-by-peer interactions (see Table 3, column 6). The results are presented separately for the full sample of students (columns 1–4), and for the sample of classrooms that would get a new teacher as a result of the reallocation (columns 5–8). The reassignments are carried out within school types and districts. Standard errors are in parentheses and computed using the Bayesian bootstrap with 1,000 replications (see Section 4). High/middle/low: students in the top/middle/bottom tercile of the baseline test score distribution.

In sum implementing an optimal assignment generates improvements in test score outcomes across the distribution of baseline achievement. Yet, while the gains are large for students with high baseline test scores – up to 10 percent of a standard deviation for students with high baseline test scores in math – for students with middle and low baseline test scores, the gains are smaller. Our results suggest that reallocations matter more in math,

³¹One may also use effect sizes of educational interventions in general as point of comparison. Based on a meta-study by Kraft (forthcoming), an effect size of 4 percent of a standard deviation in math qualifies as a “medium effect,” corresponding to the 40th percentile of effect sizes found across 750 educational interventions in pre-K–12 settings in the US (see Table 1 in Kraft, forthcoming).

compared to ELA. This is in line with the value-added literature, which consistently reports higher value-added in math compared to ELA.

Our results provide strong evidence of complementarity between student baseline test scores and teacher practices. This complementarity, if taken into account when assigning teachers to classrooms, appears to make a difference in students' performance across the distribution of baseline achievement. Teacher reassignments that maximize the average test score, however, will not close the achievement gap between students with low versus high baseline achievement levels. They actually widen this achievement gap. A different objective function, one that values equity as well as "efficiency," may be preferred in practice.

It is also possible that assignments based on other student/teacher attributes might both raise average achievement and narrow achievement gaps. For example, if literacy in a student's native language helps develop literacy in English, then perhaps bilingual teachers could help to raise the achievement levels of English language learners. Many other examples might merit exploration.

6 Sensitivity checks

Teacher and student categories

Working with a discrete categorization of baseline achievement and teacher FFT provides a simple, but also highly flexible, way of capturing non-linearities in educational production. We tested the sensitivity of our results to alternative categorizations of teachers and students. Specifically, we test (a) a coarser specification with two levels of teacher FFT (cutoff point at 2.5) and two levels of student baseline test scores and (b) a finer specification with four levels of teacher FFT (cutoff points at 2.25, 2.5, and 2.75) and four levels of student baseline test scores.

In the coarser specification, we do not detect any significant match effects between teacher FFT and student baseline test scores (see Appendix Table A.10): the model with only two levels of teacher FFT does not capture the positive effect of the high-FFT teachers on students with middle or high baseline test scores. By contrast, the results of the finer specification, with four levels of FFT, are very similar to the results of our preferred specification (see Appendix Table A.11).

Measure of teaching practices

We also assess the sensitivity of our results to the measure of teaching practices that we use in the analysis. The MET data also contains an alternative measure, the CLASS (Classroom Assessment Scoring System), for a subset of the sample (6,320 observations in the math sample and 6,999 observations in the ELA sample). The CLASS is a teacher observation

protocol that uses different domains and evaluation criteria than the FFT (see Section B.5 in the Data Appendix for details on the observation protocol and on how we process the data). The CLASS is also widely used in research on teacher quality (e.g. Araujo et al., 2016). We find that our results are similar when using the CLASS instead of the FFT (see Appendix Table A.12).

Restrictions on the reassignment process

In our preferred optimization procedure, we optimize the assignment of teachers to classrooms within types of schools (elementary or middle school) and school districts. We choose these restrictions because teachers might not be willing or able to teach in a different school type or a different school district. As a sensitivity check, we also calculated the results of both a more restrictive and a less restrictive optimization procedure.

Our least restrictive allocation optimizes the assignment within school types, but allows for reassignments across districts. The magnitudes of the reallocation effects in this case are only slightly larger than those where reallocations are within districts (see Appendix Tables A.13 and A.14, columns 1–4). This finding may reflect the similarity of the distributions of baseline student achievement and teacher FFT in MET school districts. It would be interesting to repeat our analysis in a metro area consisting of an urban core district and multiple suburban ones. In such a setting it seems plausible that moving teachers across school districts might raise average achievement.

Our most restrictive allocation reassigns teachers only within school-grade-subject cells (i.e., randomization blocks). This restriction is strong because each randomization block typically contains only two or three sections. Under this restriction, the reallocation gains are overall negligible (see Appendix Tables A.13 and A.14, columns 5–8).

Accounting for teacher-peer match effects in the assignment

Our main results for the reallocation effects are based on a model without teacher-to-peer match effects (i.e., in a model which imposes the restriction that $\lambda = 0$), since these effects are very noisily estimated using our data. Using a model which does not impose this restriction generates reallocation effects about twice as large. These effects are also more noisily measured (see Appendix Tables A.15 and A.16).

7 Conclusion

We provide an econometric framework that allows us to semiparametrically characterize complementarity between teaching practices and student baseline test scores. Our framework

exploits the random assignment of teachers to classrooms available in the MET dataset, while formally dealing with non-compliance by both teachers and students.

Our results show that the potential gains associated with an outcome-maximizing assignment of teachers to classrooms are large. They are comparable to fairly large (hypothetical) interventions to raise teacher VAM. An attractive feature is that they are, at least in theory, resource neutral. No new teachers are required to implement the policies we consider.

Our focus on the objective of maximizing the population average test score leads to an optimal assignment that increases the gap between less and more prepared students. We could also consider objective functions that do not focus on average test scores, but instead on test score gaps or proficiency levels. We could, for instance, choose an assignment which maximizes the number of students who reach a “proficient” level on their end-of-year assessment. It is possible that this objective function would suggest a less assortative assignment: students with high baseline scores would likely reach the proficient level regardless of their teacher’s FFT, while the incremental effect of having a high-FFT teacher on the probability of attaining proficiency may be larger for students with lower baseline test scores. Outcomes other than math and ELA achievement (e.g., socio-emotional skills) may also be of interest.

We consider this paper as a first pass that establishes the feasibility of recovering match effects from imperfect experimental data and shows that the resulting reallocation effects that depend upon these match effects and the supply of teachers are substantial.

Acknowledgements

We thank seminar audiences at the 2017 All-California Econometrics Conference at Stanford University, UC Riverside, NESG, University of Duisburg-Essen, Tinbergen Institute Amsterdam, Lund University, IFN Stockholm, the University of Bristol, and the University of Southern California for helpful feedback. We thank Tommy Andersson, Kirabo Jackson, Magne Mogstad, Hessel Oosterbeek, Daniele Paserman, and Hashem Pesaran for useful comments and discussions. All the usual disclaimers apply. Financial support for Graham was provided by the National Science Foundation (SES #1357499, SES #1851647). Thiemann was affiliated with USC Dornsife INET while working on this paper.

References

- AI, C. AND X. CHEN (2003): “Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions,” *Econometrica*, 71, 1795 – 1843.
- ARAUJO, M. C., P. CARNEIRO, Y. CRUZ-AGUAYO, AND N. SCHADY (2016): “Teacher quality and learning outcomes in kindergarten,” *The Quarterly Journal of Economics*, 131, 1415 – 1453.
- AUCEJO, E., P. COATE, J. C. FRUEHWIRTH, S. KELLY, AND Z. MOZENTER (2019): “Teacher Effectiveness and Classroom Composition,” *Manuscript*.
- BHATTACHARYA, D. (2009): “Inferring optimal peer assignment from experimental data,” *Journal of the American Statistical Association*, 104, 486 – 500.
- BHATTACHARYA, D. AND P. DUPAS (2012): “Inferring welfare maximizing treatment assignment under budget constraints,” *Journal of Econometrics*, 167, 168–196.
- BONHOMME, S. (2020): “Teams: heterogeneity, sorting and complementarity,” Tech. rep., University of Chicago, Department of Economics.
- CHAMBERLAIN, G. AND G. W. IMBENS (2003): “Nonparametric applications of Bayesian inference,” *Journal of Business and Economic Statistics*, 21, 12 – 18.
- CHEN, X., O. LINTON, AND I. VAN KEILEGOM (2003): “Estimation of semiparametric models when the criterion function is not smooth,” *Econometrica*, 71, 1591 – 1608.
- CHETTY, R., J. N. FRIEDMAN, AND J. E. ROCKOFF (2012): “Great Teaching,” *Education Next*, 12, 59 – 64.
- (2014a): “Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates,” *American Economic Review*, 104, 2593 – 2632.
- (2014b): “Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood,” *American Economic Review*, 104, 2633 – 2679.
- CLOTFELTER, C. T., H. F. LADD, AND J. L. VIGDOR (2006): “Teacher-student matching and the assessment of teacher effectiveness,” *Journal of human Resources*, 41, 778–820.
- COHEN-VOGEL, L., L. FENG, AND L. OSBORNE-LAMPKIN (2013): “Seniority provisions in collective bargaining agreements and the ‘teacher quality gap’,” *Educational Evaluation and Policy Analysis*, 35, 324 – 343.
- CONDIE, S., L. LEFGREN, AND D. SIMS (2014): “Teacher heterogeneity, value-added and education policy,” *Economics of Education Review*, 40, 76–92.

- DANIELSON, C. (2011): “The Framework for Teaching Evaluation Instrument, 2011 Edition,” Retrieved on May 31, 2020, from <https://danielsongroup.org/downloads/2011-framework-teaching-evaluation-instrument>.
- DARLING-HAMMOND, L. (2015): “Can Value Added Add Value to Teacher Evaluation?” *Educational Researcher*, 44, 132 – 137.
- DEE, T. S. (2004): “Teachers, race, and student achievement in a randomized experiment,” *Review of Economics and Statistics*, 86, 195 – 210.
- (2005): “A Teacher like Me: Does Race, Ethnicity, or Gender Matter?” *The American Economic Review*, 95, 158 – 165.
- (2007): “Teachers and the gender gaps in student achievement,” *Journal of Human resources*, 42, 528–554.
- EGALITE, A. J. AND B. KISIDA (2018): “The effects of teacher match on students’ academic perceptions and attitudes,” *Educational Evaluation and Policy Analysis*, 40, 59–81.
- GARRETT, R. AND M. P. STEINBERG (2015): “Examining Teacher Effectiveness Using Classroom Observation Scores: Evidence From the Randomization of Teachers to Students,” *Educational Evaluation and Policy Analysis*, 37, 224 – 242.
- GONG, J., Y. LU, AND H. SONG (2018): “The effect of teacher gender on students’ academic and noncognitive outcomes,” *Journal of Labor Economics*, 36, 743–778.
- GRAHAM, B. (2008): “Identifying social interactions through conditional variance restrictions,” *Econometrica*, 76, 643 – 660.
- GRAHAM, B. S. (2011): “Econometric methods for the analysis of assignment problems in the presence of complementarity and social spillovers,” in *Handbook of social economics*, ed. by J. Benhabib, A. Bisin, and M. O. Jackson, Amsterdam: North-Holland, vol. 1B, 965 – 1052.
- GRAHAM, B. S., G. W. IMBENS, AND G. RIDDER (2007): “Redistributive effects for discretely-valued inputs,” IEPR Working Paper 07.7, University of Southern California.
- (2010): “Measuring the effects of segregation in the presence of social spillovers: a nonparametric approach,” Working Paper 16499, NBER.
- (2014): “Complementarity and aggregate implications of assortative matching: a nonparametric analysis,” *Quantitative Economics*, 5, 29 – 66.

- (2020): “Identification and Efficiency Bounds for the Average Match Function Under Conditionally Exogenous Matching,” *Journal of Business & Economic Statistics*, 38, 303 – 316.
- GRISSOM, J. A., D. KALOGRIDES, AND S. LOEB (2015): “The micropolitics of educational inequality: The case of teacher–student assignments,” *Peabody Journal of Education*, 90, 601 – 614.
- HANSEN, B. E. (2020): “Econometrics,” Retrieved on July 4, 2020, from <https://www.ssc.wisc.edu/~bhansen/econometrics/>.
- HANUSHEK, E. (1971): “Teacher Characteristics and Gains in Student Achievement: Estimation Using Micro Data.” *American Economic Review*, 61, 280 – 288.
- HANUSHEK, E. A., J. F. KAIN, AND S. G. RIVKIN (2004): “Disruption versus Tiebout improvement: The costs and benefits of switching schools,” *Journal of Public Economics*, 88, 1721 – 1746.
- HSIEH, Y.-W., X. SHI, AND M. SHUM (2018): “Inference on estimators defined by mathematical programming,” *Manuscript*.
- JACKSON, C. K. (2013): “Match quality, worker productivity, and worker mobility: Direct evidence from teachers,” *Review of Economics and Statistics*, 95, 1096–1116.
- JOSHI, E., S. DOAN, AND M. G. SPRINGER (2018): “Student-teacher race congruence: New evidence and insight from Tennessee,” *AERA Open*, 4, 2332858418817528.
- KALOGRIDES, D., S. LOEB, AND T. BETEILLE (2011): “Power Play? Teacher Characteristics and Class Assignments,” Working Paper 59, National Center for Analysis of Longitudinal Data in Education Research.
- KANE, THOMAS, J., D. F. MCCAFFREY, T. MILLER, AND D. O. STAIGER (2013): “Have we identified effective teachers? Validating measures of effect teaching using random assignment,” Met project research paper, Bill & Melinda Gates Foundation.
- KRAFT, M. A. (forthcoming): “Interpreting Effect Sizes of Education Interventions,” *Educational Researcher*.
- LIM, J. AND J. MEER (2017): “The impact of teacher–student gender matches random assignment evidence from South Korea,” *Journal of Human Resources*, 52, 979–997.
- LOEB, S., J. SOLAND, AND L. FOX (2014a): “Is a Good Teacher a Good Teacher for All? Comparing Value-Added of Teachers with Their English Learners and Non-English Learners,” *Educational Evaluation and Policy Analysis*, 36, 457 – 475.


- (2014b): “Is a good teacher a good teacher for all? Comparing value-added of teachers with their English learners and non-English learners,” *Educational Evaluation and Policy Analysis*, 36, 457–475.
- LUSHER, L., D. CAMPBELL, AND S. CARRELL (2018): “TAs like me: Racial interactions between graduate teaching assistants and undergraduates,” *Journal of Public Economics*, 159, 203–224.
- MANSKI, C. F. (1993): “Identification of endogenous social effects: the reflection problem,” *Review of Economic Studies*, 60, 531 – 542.
- McFARLAND, J. ET AL. (2019): “The Condition of Education 2019,” NCES 2019-144, National Center for Education Statistics.
- MORGANSTEIN, D. AND R. WASSERSTEIN (2014): “ASA Statement on Value-Added Models,” *Statistics and Public Policy*, 1, 108 – 110.
- ROBINSON, P. M. (1988): “Root-N-Consistent Semiparametric Regression,” *Econometrica*, 56, 931 – 954.
- ROCKOFF, J. E. (2004): “The Impact of Individual Teachers on Student Achievement: Evidence from Panel Data,” *American Economic Review*, 94, 247 – 252.
- ROTHSTEIN, J. (2010): “Teacher Quality in Educational Production: Tracking, Decay, and Student Achievement 2010,” *Quarterly Journal of Economics*, 125, 175 – 214.
- (2017): “Measuring the impacts of teachers: Comment,” *American Economic Review*, 107, 1656–84.
- SANDERSON, E. AND F. WINDMEIJER (2016): “A weak instrument F-test in linear IV models with multiple endogenous variables,” *Journal of Econometrics*, 190, 212 – 221.
- SNYDER, T. D. ET AL. (2017): “Digest of Education Statistics,” NCES 2018-070, National Center for Education Statistics.
- STOCK, J. H. AND M. YOGO (2005): “Testing for weak instruments in Linear IV regression,” in *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Cambridge University Press, 80 – 108.
- TRAPP, A. C., A. TEYTELBOYM, A. MARTINELLO, T. ANDERSSON, AND N. AHANI (2018): “Placement Optimization in Refugee Resettlement,” Tech. rep., Lund University, Department of Economics.

U.S. DEPARTMENT OF EDUCATION (2020): “Teacher Performance Evaluations in U.S. Public Schools,” NCES 2020-133, <https://nces.ed.gov/pubs2020/2020133.pdf>, online; accessed 21 January 2021.

WHITE, M., B. ROWEN, G. ALTER, L. BLANKENSHIP, C. GREENE, AND S. WINDISH (2019): *User Guide to Measures of Effective Teaching Longitudinal Database (MET LDB)*, Ann Arbor: Inter-University Consortium for Political and Social Research, The University of Michigan.

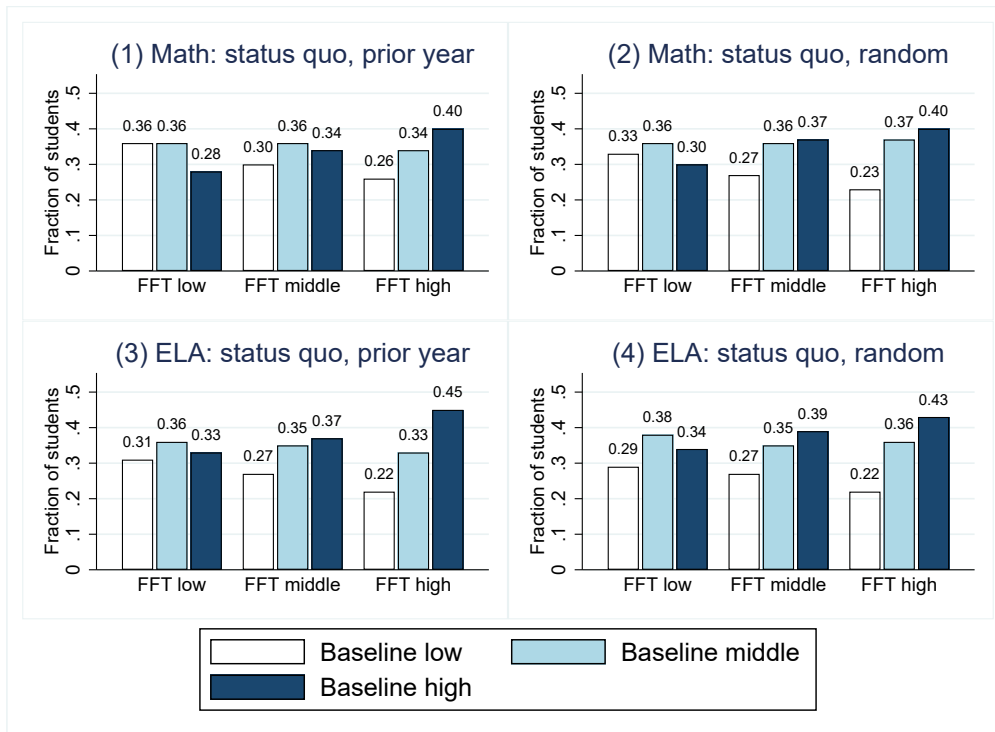
A Figures and Tables

Figure A.1: Distribution of teacher FFT and student baseline test scores



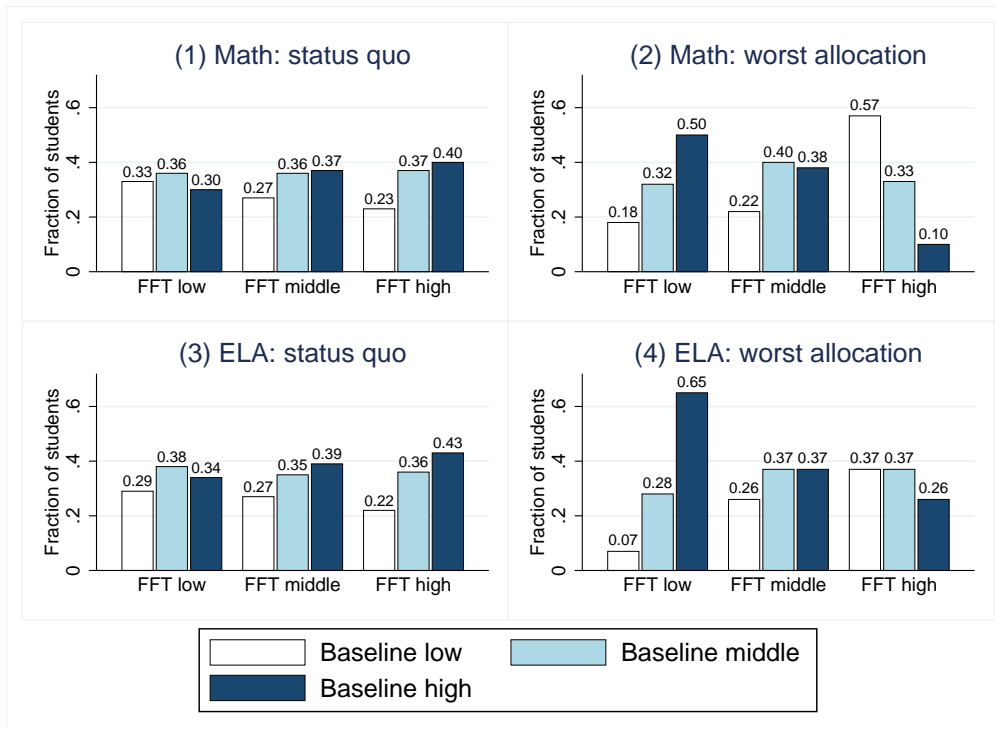
Note: Distribution of teacher FFT (Panels 1 and 3) and student baseline test scores (Panels 2 and 4) in the estimation sample. Student baseline test scores are section averages.

Figure A.2: Assortativeness of the status quo in the year prior to randomization (AY 2009/10) and of the status quo in the year where teachers were randomly assigned (AY 2010/11)



Note: The figure compares the status quo allocation in the year prior to the randomized field experiment (Panels 1 and 3) with the status quo when the random allocation was implemented (Panels 2 and 4). In the year prior to the field experiment, the status quo was determined by schools and school principals. In the year of the field experiment, MET researchers conducted the randomization of teachers to classrooms within schools. The bars represent the fractions of students with low, middle, and high baseline test scores that a teacher with low, middle, and high FFT is assigned to on average under each of the allocations. For each teacher type, the fractions add up to 1. Sample sizes math: 33,913 students and 955 teachers in (Panel 1); 8,534 students and 614 teachers (Panel 2). Sample sizes ELA: 35,068 students and 1,015 teachers (Panel 3); 9,641 students and 649 teachers (Panel 3). The sample sizes are smaller in Panels 2 and 4 because some schools and teachers opted out of the field experiment.

Figure A.3: Assortativeness of the worst allocation in comparison with the status quo



Note: The figure compares the status quo allocation (Panels 1 and 3) with the worst allocation (Panels 2 and 4). The bars represent the fractions of students with low, middle, and high baseline test scores that a teacher with low, middle, and high FFT is assigned to on average under each of the allocations. For each teacher type, the fractions add up to 1. The optimization is carried out within school types (elementary or middle school) and districts.

Table A.1: Summary statistics of the estimation samples

	(1)	(2)	(3)	(4)
	Math sample		ELA sample	
	Mean	SD	Mean	SD
Student characteristics				
Age	10.32	(1.56)	10.28	(1.48)
Fourth grade	26%	-	26%	-
Fifth grade	35%	-	35%	-
Sixth grade	16%	-	15%	-
Seventh grade	10%	-	12%	-
Eighth grade	12%	-	12%	-
Male	49%	-	49%	-
Gifted	6%	-	11%	-
Special education	7%	-	7%	-
ELL	15%	-	13%	-
Free/reduced-price lunch	58%	-	57%	-
White	27%	-	28%	-
Black	31%	-	30%	-
Hispanic	31%	-	33%	-
Asian	8%	-	7%	-
Other race	3%	-	3%	-
Subject test score 2010 (baseline)	0.11	(0.89)	0.17	(0.93)
Subject test score 2011 (outcome)	0.15	(0.90)	0.17	(0.91)
Teacher characteristics				
Male	14%	-	11%	-
White	65%	-	64%	-
Black	28%	-	29%	-
Hispanic	6%	-	6%	-
Other race	2%	-	1%	-
Years in district	7.39	(6.73)	7.36	(6.33)
Master's or higher	40%	-	37%	-
FFT	2.53	(0.30)	2.59	(0.30)
Classroom characteristics				
Class size	25.54	(5.48)	25.60	(5.44)
Sample size				
Students	8,534		9,641	
Teachers	614		649	
Schools	153		160	

Note: Summary statistics of the estimation samples. Standard deviations are in parentheses. ELL: English language learner. FFT: Framework for teaching. For details on the sample construction, see Appendix B.

Table A.2: Balancing tests. Dependent variable: FFT of the assigned teacher

	(1)	(2)	(3)	(4)
	Math		ELA	
	coeff	SE	coeff	SE
Baseline test score	0.004	(0.005)	0.000	(0.004)
Age	-0.003	(0.006)	-0.006	(0.005)
Male	0.001	(0.005)	-0.002	(0.003)
Gifted	-0.015	(0.015)	0.010	(0.016)
Special education	0.013	(0.012)	-0.002	(0.009)
ELL	-0.004	(0.011)	-0.014	(0.011)
Black	-0.004	(0.008)	-0.010	(0.007)
Hispanic	-0.013	(0.009)	0.002	(0.006)
Asian	0.009	(0.014)	0.016	(0.011)
Other race	0.014	(0.012)	0.010	(0.010)
Free/reduced-price lunch	0.005	(0.010)	-0.002	(0.006)
F-test for joint significance (p-value)	0.465		0.241	
R^2	0.598		0.635	
N	8,518		9,630	

Note: The table presents results of OLS regressions of the FFT of the assigned teacher on individual student characteristics. Free/reduced-price lunch eligibility is coded as 0 for students who do not have information on this variable in the dataset. All regressions control for randomization block fixed effects. Analytic standard errors, clustered by randomization block, are in parentheses. ELL: English language learner. FFT: Framework for teaching.

Table A.3: Tests of Assumption 1. Dependent variables: Differences between assigned and realized teacher characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
Panel I. Math sample						
	Male	Teacher's race			Experience	Master's degree
		white	black	hispanic	(years)	
FFT, assigned teacher	0.009 (0.060)	-0.039 (0.107)	-0.061 (0.068)	0.059 (0.059)	-2.727 (2.236)	-0.018 (0.099)
Baseline test score	0.003 (0.005)	-0.003 (0.005)	0.005 (0.004)	-0.002 (0.002)	0.071 (0.082)	-0.006 (0.006)
Avg. peer baseline test score, assigned peers	0.027 (0.076)	0.037 (0.058)	0.007 (0.034)	-0.037 (0.041)	0.738 (0.931)	-0.052 (0.089)
F-test for joint significance (p-value)	0.956	0.799	0.537	0.731	0.635	0.849
R^2	0.149	0.176	0.108	0.039	0.163	0.154
N	8,163	8,098	7,689	7,689	5,585	6,083
Panel II. ELA sample						
	Male	Teacher's race			Experience	Master's degree
		white	black	hispanic	(years)	
FFT, assigned teacher	0.008 (0.067)	0.036 (0.085)	-0.037 (0.071)	0.031 (0.028)	0.500 (0.799)	-0.003 (0.060)
Baseline test score	0.003 (0.002)	-0.000 (0.004)	0.002 (0.003)	-0.002 (0.002)	-0.008 (0.042)	0.001 (0.004)
Avg. peer baseline test score, assigned peers	-0.008 (0.027)	0.102** (0.046)	-0.041 (0.031)	-0.056* (0.029)	0.220 (0.436)	0.049 (0.049)
F-test for joint significance (p-value)	0.300	0.116	0.414	0.063	0.837	0.779
R^2	0.090	0.130	0.053	0.050	0.142	0.111
N	9,183	9,183	9,183	9,183	6,791	6,963

Note: The table presents results of OLS regressions of differences between assigned and realized teacher characteristics on individual baseline test scores, the FFT of the assigned teacher, and the average baseline test scores of the assigned peers. Each column represents a regression with a different dependent variable. All regressions control for randomization block fixed effects. Analytic standard errors, clustered by randomization block, are in parentheses. Teachers with ethnicity "white" and ethnicity "other" are pooled into one category. FFT: Framework for teaching.

**significant at the 5%-level, *significant at the 10%-level.

Table A.4: Tests of Assumption 2 in the math sample. Dependent variables: differences between assigned and realized peer characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Difference between assigned and realized peer characteristics	Age	Male	Gifted	Special education	ELL	FRL	Race/ethnicity			
							White	Black	Hispanic	Asian
FFT, assigned teacher	0.008 (0.019)	0.015 (0.014)	0.006 (0.009)	-0.017 (0.011)	0.006 (0.012)	-0.005 (0.016)	0.004 (0.011)	-0.010 (0.011)	0.006 (0.013)	-0.001 (0.007)
Baseline test score	-0.006 (0.003)	0.000 (0.001)	0.004** (0.002)	-0.001 (0.001)	0.000 (0.001)	0.000 (0.002)	0.001 (0.001)	-0.002 (0.002)	0.000 (0.001)	0.001* (0.001)
Avg. peer baseline test score, assigned peers	-0.030* (0.017)	-0.006 (0.013)	-0.013 (0.009)	0.009 (0.008)	0.033** (0.011)	-0.001 (0.014)	-0.012 (0.010)	0.012 (0.011)	-0.007 (0.011)	0.008 (0.005)
R^2	0.124	0.264	0.222	0.240	0.200	0.241	0.207	0.232	0.247	0.185
N	8,525	8,534	8,534	8,527	8,534	5,976	8,534	8,534	8,534	8,534

Note: The table presents results of OLS regressions of differences between assigned and realized peer characteristics on the FFT of the assigned teacher, individual baseline test scores, and the average baseline test scores of the assigned peers. Each column represents a regression with a different dependent variable. All regressions control for randomization block fixed effects. Analytic standard errors, clustered by randomization block, are in parentheses. ELL: English language learner, FRL: Free/reduced-price lunch eligible, FFT: Framework for teaching.

** significant at the 5%-level * significant at the 10%-level.

Table A.5: Tests of Assumption 2 in the ELA sample. Dependent variables: differences between assigned and realized peer characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Difference between assigned and realized peer characteristics	Age	Male	Gifted	Special education	ELL	FRL	Race/ethnicity			
							White	Black	Hispanic	Asian
FFT, assigned teacher	0.000 (0.020)	0.008 (0.014)	0.008 (0.009)	-0.012 (0.011)	0.022 (0.017)	-0.003 (0.014)	-0.008 (0.012)	-0.001 (0.013)	0.012 (0.014)	-0.002 (0.009)
Baseline test score	-0.005** (0.003)	-0.001 (0.001)	0.002 (0.001)	0.000 (0.001)	-0.001 (0.001)	-0.002 (0.001)	0.001 (0.001)	0.000 (0.001)	-0.001 (0.001)	0.001 (0.000)
Avg. peer baseline test score, assigned peers	-0.009 (0.013)	0.000 (0.011)	-0.029*** (0.007)	0.013 (0.009)	0.028** (0.009)	0.007 (0.011)	-0.002 (0.008)	0.001 (0.008)	0.001 (0.010)	0.000 (0.006)
R^2	0.098	0.241	0.147	0.208	0.178	0.238	0.203	0.254	0.218	0.215
N	9,635	9,641	9,641	9,636	9,641	7,270	9,641	9,641	9,641	9,641

Note: The table presents results of OLS regressions of differences between assigned and realized peer characteristics on the FFT of the assigned teacher, individual baseline test scores, and the average baseline test scores of the assigned peers. Each column represents a regression with a different dependent variable. All regressions control for randomization block fixed effects. Analytic standard errors, clustered by randomization block, are in parentheses. ELL: English language learner, FRL: Free/reduced-price lunch eligible, FFT: Framework for teaching.

** significant at the 5%-level * significant at the 10%-level.

Table A.6: Tests of restriction (25). Dependent variables: average peer baseline test scores of realized peers

	(1)	(2)
	Average peer baseline test score, realized peers	
	Math sample	ELA sample
FFT, assigned teacher	-0.015 (0.029)	0.036 (0.029)
Baseline test score	0.031*** (0.008)	0.023*** (0.005)
Avg. peer baseline test score, assigned peers	0.744 (0.034)	0.769 (0.030)
R^2	0.881	0.904
N	8,534	9,641

Note: The table presents results of OLS regressions of the average peer baseline test scores of the realized peers on individual baseline test scores, FFT of the assigned teacher, and the average peer baseline test scores of the assigned peers in the math and ELA samples. All regressions control for randomization block fixed effects. Analytic standard errors, clustered by randomization block, are in parentheses. FFT: Framework for teaching.

*** significant at the 1%-level.

Table A.7: Tests for instrument relevance

	(1)	(2)	(3)
Panel A. Math			
	First-stage F-statistics		
	unconditional		SW
Variable	F(10,238)	p-value	F(1, 238)
FFT middle	17.57	(0.000)	60.17
FFT high	11.62	(0.000)	57.95
FFT middle \times baseline middle	82.49	(0.000)	498.96
FFT high \times baseline middle	33.30	(0.000)	387.27
FFT middle \times baseline high	39.59	(0.000)	322.43
FFT high \times baseline high	36.92	(0.000)	378.59
FFT middle \times fraction peers middle	17.47	(0.000)	113.27
FFT high \times fraction peers middle	12.85	(0.000)	102.70
FFT middle \times fraction peers high	8.07	(0.000)	182.12
FFT middle \times fraction peers high	12.01	(0.000)	137.55
Panel B. ELA			
	First-stage F-statistics		
	unconditional		SW
Variable	F(10,238)	p-value	F(1, 238)
FFT middle	26.21	(0.000)	75.61
FFT high	22.34	(0.000)	102.87
FFT middle \times baseline middle	146.86	(0.000)	836.59
FFT high \times baseline middle	121.72	(0.000)	1500.06
FFT middle \times baseline high	128.34	(0.000)	1226.48
FFT high \times baseline high	88.44	(0.000)	1724.38
FFT middle \times fraction peers middle	25.06	(0.000)	164.19
FFT high \times fraction peers middle	21.57	(0.000)	213.78
FFT middle \times fraction peers high	13.9	(0.000)	362.49
FFT middle \times fraction peers high	13.21	(0.000)	477.11

Note: The table presents weak instrument F-tests following [Sanderson and Windmeijer \(2016\)](#). Columns 1 and 2 present the unconditional F-statistics from first-stage regressions of equation (23). Column 3 presents the Sanderson-Windmeijer conditional F-statistics for each of the endogenous variables. The F-statistics in column 3 need to be compared to the critical values of the Stock-Yogo weak identification F-test ([Stock and Yogo, 2005](#)). All of the F-statistics exceed the Stock-Yogo critical values for a 5% maximal IV relative bias (critical value is 20.74).

Table A.8: OLS regression results of the linear model. Dependent variables: student test score outcomes

		(1)	(2)	(3)	(4)	(5)	(6)
		A. Only teacher effects		B. Full model		C. Without teacher \times peer interactions	
		Math	ELA	Math	ELA	Math	ELA
δ	FFT	0.088** (0.040)	0.118** (0.045)	0.081** (0.040)	0.116** (0.045)	0.080** (0.040)	0.116** (0.044)
η	FFT \times baseline			0.060* (0.031)	0.088** (0.032)	0.059** (0.028)	0.078** (0.030)
λ	FFT \times avg. peer baseline			-0.006 (0.066)	-0.066 (0.081)		
β	Baseline	0.749*** (0.011)	0.690*** (0.011)	0.595*** (0.079)	0.460*** (0.081)	0.598*** (0.073)	0.487*** (0.078)
R^2		0.700	0.633	0.701	0.633	0.701	0.633
N		8,534	9,641	8,534	9,641	8,534	9,641

Note: The dependent variables are subject-specific test score outcomes in math and ELA. The specifications include linear terms for FFT, individual and peer baseline test scores. All regressions control for the $h(x, \bar{x})$ function (see Section 2) and for randomization block fixed effects. Analytic standard errors, clustered by randomization block, are in parentheses. FFT: Framework for teaching.

*** significant at the 1%-level ** significant at the 5%-level * significant at the 10%-level.

Table A.9: OLS regression results of the 3×3 model. Dependent variables: student test score outcomes

		(1)	(2)	(3)	(4)	(5)	(6)
		A. Only teacher effects		B. Full model		C. Without teacher \times peer interactions	
		Math	ELA	Math	ELA	Math	ELA
δ	FFT middle	0.062*	0.028	0.032	-0.058	0.032	-0.127**
		(0.032)	(0.034)	(0.097)	(0.091)	(0.038)	(0.045)
	FFT high	0.066	0.048	-0.060	-0.043	-0.016	-0.141**
		(0.041)	(0.043)	(0.113)	(0.125)	(0.052)	(0.054)
η	FFT middle						
	\times baseline middle			0.030	0.108**	0.033	0.102**
				(0.047)	(0.054)	(0.045)	(0.051)
	\times baseline high			0.039	0.136**	0.053	0.137**
				(0.060)	(0.067)	(0.055)	(0.062)
	FFT high						
	\times baseline middle			0.100*	0.113*	0.098*	0.094
				(0.059)	(0.060)	(0.059)	(0.059)
	\times baseline high			0.120*	0.225**	0.104	0.193**
				(0.068)	(0.075)	(0.067)	(0.072)
λ	FFT middle						
	\times fraction peers middle			-0.024	-0.204		
				(0.198)	(0.187)		
	\times fraction peers high			-0.045	-0.126		
				(0.187)	(0.154)		
	FFT high						
	\times fraction peers middle			0.155	-0.187		
				(0.224)	(0.232)		
	\times fraction peers high			0.460	-0.355		
				(0.409)	(0.343)		
β	Baseline middle	0.911***	0.823***	0.782***	0.713***	0.783***	0.723***
		(0.072)	(0.075)	(0.076)	(0.093)	(0.075)	(0.095)
	Baseline high	1.678***	1.636***	1.551***	1.443***	1.544***	1.454***
		(0.105)	(0.101)	(0.111)	(0.113)	(0.106)	(0.116)
R^2		0.618	0.557	0.570	0.523	0.570	0.523
N		8,534	9,641	8,534	9,641	8,534	9,641

Note: The dependent variables are subject-specific test score outcomes in math and ELA. All regressions control for the $h(x, \bar{x})$ function (see Section 2) and for randomization block fixed effects. Analytic standard errors, clustered by randomization block, are in parentheses. FFT: Framework for teaching.

*** significant at the 1%-level ** significant at the 5%-level * significant at the 10%-level.

Table A.10: IV regression results of the 2×2 model. Dependent variables: student test score outcomes

		(1)	(2)	(3)	(4)	(5)	(6)
		A. Only teacher effects		B. Full model		C. Without teacher \times peer interactions	
		Math	ELA	Math	ELA	Math	ELA
δ	FFT high	0.072*	-0.060	0.088	0.170	0.056	-0.027
		(0.038)	(0.043)	(0.087)	(0.107)	(0.047)	(0.046)
η	\times baseline high			0.039	-0.014	0.030	-0.062
				(0.053)	(0.055)	(0.052)	(0.055)
λ	\times fraction peers high			-0.068	-0.422*		
				(0.149)	(0.229)		
β	Baseline high	0.935***	0.927***	0.920***	0.938***	0.924***	0.956***
		(0.058)	(0.054)	(0.058)	(0.060)	(0.057)	(0.060)
R^2		0.545	0.485	0.545	0.481	0.545	0.485
N		8,534	9,641	8,534	9,641	8,534	9,641

Note: The dependent variables are subject-specific test score outcomes in math and ELA. The instrumental variables are based on assigned teacher FFT (Panels A–C) and assigned peer baseline test scores (Panel B). All regressions control for the $h(x, \bar{x})$ function (see Section 2) and for randomization block fixed effects. Analytic standard errors, clustered by randomization block, are in parentheses. FFT: Framework for teaching.

*** significant at the 1%-level ** significant at the 5%-level * significant at the 10%-level.

Table A.11: IV regression results of the 4×4 model. Dependent variables: student test score outcomes

		(1)	(2)	(3)	(4)	(5)	(6)
		A. Only teacher effects		B. Full model		C. Without teacher \times peer interactions	
		Math	ELA	Math	ELA	Math	ELA
δ	FFT lower middle	0.146** (0.063)	-0.037 (0.048)	0.029 (0.307)	0.097 (0.188)	0.114 (0.092)	-0.056 (0.059)
	FFT upper middle	0.029 (0.059)	0.013 (0.055)	-0.656* (0.386)	-0.222 (0.268)	-0.035 (0.078)	-0.114 (0.072)
	FFT high	0.017 (0.069)	-0.006 (0.060)	-1.622** (0.815)	-0.044 (0.213)	-0.162 (0.122)	-0.132 (0.086)
η	FFT lower middle						
	\times baseline lower middle			0.111 (0.096)	0.084 (0.075)	0.083 (0.095)	0.074 (0.072)
	\times baseline upper middle			0.028 (0.104)	0.023 (0.087)	-0.001 (0.107)	0.007 (0.079)
	\times baseline high			0.053 (0.125)	0.028 (0.107)	0.016 (0.124)	0.002 (0.104)
	FFT upper middle						
	\times baseline lower middle			0.101 (0.087)	0.103 (0.070)	0.096 (0.078)	0.103 (0.069)
	\times baseline upper middle			0.001 (0.087)	0.141* (0.084)	0.015 (0.083)	0.142* (0.080)
	\times baseline high			0.110 (0.115)	0.214** (0.092)	0.111 (0.100)	0.222** (0.093)
	FFT high						
	\times baseline lower middle			0.269** (0.120)	0.144* (0.075)	0.200* (0.104)	0.118 (0.074)
	\times baseline upper middle			0.136 (0.111)	0.232** (0.089)	0.140 (0.113)	0.181** (0.085)
	\times baseline high			0.266** (0.130)	0.251** (0.097)	0.253** (0.124)	0.182* (0.099)
β	Baseline lower middle	0.635*** (0.131)	0.780*** (0.098)	0.517** (0.195)	0.747*** (0.105)	0.563*** (0.156)	0.721*** (0.107)
	Baseline upper middle	1.228*** (0.160)	1.273*** (0.116)	1.143*** (0.218)	1.246*** (0.124)	1.201*** (0.177)	1.203*** (0.127)
	Baseline high	1.913*** (0.172)	2.170*** (0.141)	1.750*** (0.246)	2.121*** (0.148)	1.835*** (0.201)	2.072*** (0.151)
λ			✓	✓			
R^2		0.649	0.592	0.638	0.589	0.648	0.593
N		8,534	9,641	8,534	9,641	8,534	9,641

Note: The dependent variables are subject-specific test score outcomes in math and ELA. The instrumental variables are based on assigned teacher FFT (Panels A–C) and assigned peer baseline test scores (Panel B). All regressions control for the $h(x, \bar{x})$ function (see Section 2) and for randomization block fixed effects. The results for the coefficients on teacher-peer match effects (Panel B) are omitted. Analytic standard errors, clustered by randomization block, are in parentheses. FFT: framework for teaching.

*** significant at the 1%-level ** significant at the 5%-level * significant at the 10%-level.

Table A.12: IV regression results of the 3×3 model with CLASS as measure of teaching practices.
 Dependent variables: student test score outcomes

		(1)	(2)	(3)	(4)	(5)	(6)	
		A. Only teacher effects		B. Full model		C. Without teacher \times peer interactions		
		Math	ELA	Math	ELA	Math	ELA	
δ	CLASS middle	0.039 (0.061)	0.071 (0.067)	-0.214 (0.489)	0.115 (0.235)	0.008 (0.103)	0.038 (0.072)	
	CLASS high	-0.050 (0.079)	0.010 (0.088)	-0.521 (0.520)	0.085 (0.262)	-0.211* (0.117)	-0.099 (0.107)	
η	CLASS middle \times baseline middle			0.075 (0.105)	0.054 (0.070)	0.074 (0.098)	0.058 (0.067)	
	\times baseline high			0.008 (0.115)	0.040 (0.092)	0.019 (0.109)	0.048 (0.091)	
	CLASS high \times baseline middle			0.201* (0.103)	0.177** (0.076)	0.198** (0.100)	0.168** (0.077)	
	\times baseline high			0.253** (0.114)	0.157* (0.085)	0.240** (0.106)	0.137 (0.088)	
	λ	CLASS middle \times fraction peers middle			0.451 (0.934)	-0.165 (0.473)		
		\times fraction peers high			0.153 (0.373)	-0.040 (0.306)		
CLASS high \times fraction peers middle				0.741 (0.916)	-0.168 (0.375)			
\times fraction peers high			0.036 (0.442)	-0.324 (0.382)				
β	Baseline middle	0.929*** (0.095)	0.870*** (0.087)	0.819*** (0.148)	0.798*** (0.100)	0.820*** (0.143)	0.798*** (0.099)	
	Baseline high	1.686*** (0.118)	1.645*** (0.115)	1.572*** (0.158)	1.584*** (0.122)	1.562*** (0.152)	1.586*** (0.124)	
R^2		0.607	0.557	0.603	0.557	0.607	0.558	
N		6,320	6,999	6,320	6,999	6,320	6,999	

Note: The dependent variables are subject-specific test score outcomes in math and ELA. The instrumental variables are based on assigned teacher CLASS (Panels A–C) and assigned peer baseline test scores (Panel B). For details on the CLASS measure of teaching practices, see Appendix B. All regressions control for the $h(x, \bar{x})$ function (see Section 2) and for randomization block fixed effects. Analytic standard errors, clustered by randomization block, are in parentheses.

***significant at the 1%-level, **significant at the 5%-level, *significant at the 10%-level.

Table A.13: Average reallocation gains in math: sensitivity to restrictions on possible assignments. Gains expressed in test score standard deviations.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Optimal versus status quo								
	A.I Within school type				A.II Within randomization block			
	all students	high	middle	low	all students	high	middle	low
Gain	0.019	0.031	0.015	0.015	0.005	0.009	0.004	0.003
SE	(0.007)	(0.015)	(0.010)	(0.012)	(0.002)	(0.004)	(0.003)	(0.003)
N	8,534	2,332	3,108	3,094	8,534	2,332	3,108	3,094
Panel B. Optimal versus worst allocation								
	B.I Within school type				B.II Within randomization block			
	all students	high	middle	low	all students	high	middle	low
Gain	0.048	0.089	0.025	0.041	0.011	0.020	0.008	0.008
SE	(0.015)	(0.039)	(0.019)	(0.028)	(0.003)	(0.009)	(0.006)	(0.007)
N	8,534	2,332	3,108	3,094	8,534	2,332	3,108	3,094

Note: The table shows the average reallocation gains from implementing the optimal assignment instead of the random assignment (status quo) in Panel A, and the average reallocation gains from implementing the optimal assignment instead of the worst assignment in Panel B. The gains are expressed in test score standard deviations. The computations are based on the 3×3 model without teacher-by-peer interactions (see Table 3, column 5). Panels A.I and B.I display the reallocation effects when the reallocation is carried out within school types (elementary or middle school) and across school districts. Panels A.II and B.II display the reallocation effects when the reallocation is carried out within randomization blocks. Standard errors are in parentheses and computed using the Bayesian bootstrap with 1,000 replications (see Section 4). High/middle/low: students in the top/middle/bottom tercile of the baseline test score distribution.

Table A.14: Average reallocation gains in ELA: sensitivity to restrictions on possible assignments.
Gains expressed in test score standard deviations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Optimal versus status quo								
	A.I Within school type				A.II Within randomization block			
	all students	high	middle	low	all students	high	middle	low
Gain	0.009	0.027	0.002	0.004	0.002	0.007	0.001	0.001
SE	(0.007)	(0.013)	(0.007)	(0.012)	(0.002)	(0.004)	(0.002)	(0.003)
N	9,641	2,480	3,402	3,759	9,641	2,480	3,402	3,759
Panel B. Optimal versus worst allocation								
	B.I Within school type				B.II Within randomization block			
	all students	high	middle	low	all students	high	middle	low
Gain	0.021	0.066	0.004	0.006	0.005	0.017	0.001	0.002
SE	(0.013)	(0.034)	(0.013)	(0.022)	(0.004)	(0.009)	(0.004)	(0.007)
N	9,641	2,480	3,402	3,759	9,641	2,480	3,402	3,759

Note: The table shows the average reallocation gains from implementing the optimal assignment instead of the random assignment (status quo) in Panel A, and the average reallocation gains from implementing the optimal assignment instead of the worst assignment in Panel B. The gains are expressed in test score standard deviations. The computations are based on the 3×3 model without teacher-by-peer interactions (see Table 3, column 6). Panels A.I and B.I display the reallocation gains when the reallocation is carried out within school types (elementary or middle school) and across school districts. Panels A.II and B.II display the reallocation effects when the reallocation is carried out within randomization blocks. Standard errors are in parentheses and computed using the Bayesian bootstrap with 1,000 replications. High/middle/low: students in the top/middle/bottom tercile of the baseline test score distribution.

Table A.15: Average reallocation gains in math: sensitivity to inclusion of teacher-by-peer interactions. Gains expressed in test score standard deviations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A. Optimal versus status quo												
	A.I Within school type				A.II Within school type and district				A.III Within randomization block			
	all students	high	middle	low	all students	high	middle	low	all students	high	middle	low
Gain	0.035	0.059	0.026	0.027	0.034	0.057	0.025	0.025	0.009	0.016	0.008	0.006
SE	(0.020)	(0.028)	(0.017)	(0.026)	(0.022)	(0.030)	(0.018)	(0.031)	(0.006)	(0.008)	(0.006)	(0.006)
N	8,534	2,332	3,108	3,094	8,534	2,332	3,108	3,094	8,534	2,332	3,108	3,094
Panel B. Optimal versus worst allocation												
	B.I Within school type				B.II Within school type and district				B.III Within randomization block			
	all students	high	middle	low	all students	high	middle	low	all students	high	middle	low
Gain	0.098	0.170	0.070	0.073	0.085	0.146	0.061	0.064	0.022	0.035	0.019	0.015
SE	(0.048)	(0.093)	(0.043)	(0.053)	(0.047)	(0.089)	(0.046)	(0.054)	(0.012)	(0.018)	(0.013)	(0.012)
N	8,534	2,332	3,108	3,094	8,534	2,332	3,108	3,094	8,534	2,332	3,108	3,094

Note: The table shows the average reallocation gains from implementing the optimal assignment instead of the random assignment (status quo) in Panel A, and the average reallocation gains from implementing the optimal assignment instead of the worst assignment in Panel B. The gains are expressed in test score standard deviations. The computations are based on the 3×3 model with teacher-by-peer interactions (see Table 3, column 3). Columns 1–4 display the reallocation effects when the reallocation is carried out within school types (elementary or middle school) and across districts, columns 5–8 display the reallocation effects when the reallocation is carried out within school types and districts, and columns 9–12 display the reallocation effects when the reallocation is carried out within randomization blocks. Standard errors are in parentheses and computed using the Bayesian bootstrap with 1,000 replications. High/middle/low: students in the top/middle/bottom tercile of the baseline test score distribution.

Table A.16: Average reallocation gains in ELA: sensitivity to inclusion of teacher-by-peer interactions. Gains expressed in test score standard deviations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A. Optimal versus status quo												
	A.I Within school type				A.II Within school type and district				A.III Within randomization block			
	all students	high	middle	low	all students	high	middle	low	all students	high	middle	low
Gain	0.044	0.042	0.043	0.045	0.039	0.034	0.040	0.042	0.012	0.011	0.011	0.014
SE	(0.027)	(0.023)	(0.024)	(0.044)	(0.023)	(0.019)	(0.020)	(0.037)	(0.008)	(0.005)	(0.007)	(0.011)
N	9,641	2,480	3,402	3,759	9,641	2,480	3,402	3,759	9,641	2,480	3,402	3,759
Panel B. Optimal versus worst allocation												
	B.I Within school type				B.II Within school type and district				B.III Within randomization block			
	all students	high	middle	low	all students	high	middle	low	all students	high	middle	low
Gain	0.091	0.070	0.080	0.116	0.080	0.061	0.070	0.101	0.027	0.023	0.021	0.035
SE	(0.052)	(0.037)	(0.043)	(0.085)	(0.044)	(0.031)	(0.035)	(0.069)	(0.016)	(0.011)	(0.012)	(0.026)
N	9,641	2,480	3,402	3,759	9,641	2,480	3,402	3,759	9,641	2,480	3,402	3,759

Note: The table shows the average reallocation gains from implementing the optimal assignment instead of the random assignment (status quo) in Panel A, and the average reallocation gains from implementing the optimal assignment instead of the worst assignment in Panel B. The gains are expressed in test score standard deviations. The computations are based on the 3×3 model with teacher-by-peer interactions (see Table 3, column 4). Columns 1–4 display the reallocation effects when the reallocation is carried out within school types (elementary or middle school) and across districts, columns 5–8 display the reallocation effects when the reallocation is carried out within school types and districts, and columns 9–12 display the reallocation effects when the reallocation is carried out within randomization blocks. Standard errors are in parentheses and computed using the Bayesian bootstrap with 1,000 replications. High/middle/low: students in the top/middle/bottom tercile of the baseline test score distribution.

B Data Appendix

B.1 Construction of the dataset from the MET files

Our dataset combines eight different data files from the MET study (2018 release): the randomization file, the teacher file, the class section file, the student file, two classroom observation score files (the CLASS and the FFT file), as well as the district-wide files for the school years 2009/10 and 2010/11.

The basis of our data construction is the randomization file. It contains identifiers for all students who were randomly assigned to a teacher in the second year of the MET study, and an identifier (identifiers) for their assigned teacher(s). Through the teacher identifiers, we merge this dataset to the FFT and CLASS files, and thus obtain the FFT and CLASS of the assigned teachers. Moreover, we use the teacher identifiers to merge the data to the teacher file, in order to obtain teacher background characteristics of the assigned teachers.

Student characteristics and test score outcomes come from the student file and the district-wide files. Through a student identifier, we first merge the randomization file with the student file, which contains individual information for all students who were part of the MET study and still had a MET teacher at the end of the school year; i.e., the student file does not contain any information on students who switch to a non-MET teacher within the same school, a non-MET school, or a non-MET district. We use the student file to extract students' demographic characteristics, baseline test scores (test scores in school year 2009/10), and test score outcomes (test scores in school year 2010/11) in math and ELA. While there are no missings for student background characteristics, some of the students have missing baseline test scores or missing test score outcomes. Therefore, we obtain the missing test score information from the district-wide files.

To construct information on the peer group composition in the assigned classroom, we average the student background characteristics and baseline test scores at the level of the assigned teacher and randomization block, since each teacher can only be assigned to one classroom within a randomization block. To be precise, we construct the leave-own-out mean, i.e., the classroom mean excluding the student herself.

Information on the realized teacher and the realized peers are constructed based on the student file. This file contains an identifier for the section that the student attended in school year 2010/11, as well as an identifier for the teacher who taught the section. Based on this teacher identifier, we merge information on the FFT, CLASS, and background characteristics of the realized teacher. Furthermore, we construct leave-own-out means of peer characteristics based on the section identifiers. We add information on the size of the realized classroom from the class section file.

B.2 Construction of the estimation sample

In our analysis, we focus on students in grades 4-8 who were randomized to a teacher before the start of the school year. We create separate samples for math and ELA. In total, the randomization sample contains information on about 16,000 students who were randomized to a teacher in math and ELA in grades 4-8.

Table B.1: Sample construction from the randomization file

Restrictions	(1)	(2)	(3)	(4)
	Math		ELA	
	N	percent of base sample	N	percent of base sample
1. Students in randomization file	15,749	-	16,252	-
2. Record in the student file	10,268	-	11,271	-
3. At least two classrooms per randomization block (base sample)	9,824	100%	10,856	100%
4. Baseline test scores and test score outcomes available	9,245	94%	10,136	93%
5. FFT of the assigned teacher available	9,066	92%	10,057	93%
6. FFT of the realized teacher available	8,724	89%	9,767	90%
7. Information on the assigned peers available	8,718	89%	9,762	90%
8. Information on the realized peers available	8,717	89%	9,761	90%
9. At least two classrooms per randomization block after applying all restrictions	8,534	87%	9,641	89%

Note: Sample restrictions used to create the estimation sample from MET data files, school year 2010/11.

Table B.1 details further restrictions that we apply to construct our estimation dataset. We require each student to have a record in the student file, since we use this file to identify the realized teacher, the realized peer group, and the student background variables. About two-thirds of the students in the randomization sample can be identified in the student file. Moreover, we restrict our sample to only randomization blocks with at least two classrooms. The resulting dataset forms our base sample.

After constructing our base sample, we remove observations with missing information on baseline test scores or test score outcomes, with missing information on the FFT of the assigned or realized teacher, and with missing information on the baseline test scores of the assigned or realized peers. We further remove all randomization blocks with only one classroom after applying these restrictions. Our resulting estimation sample contains information on about 8,500 students in math and 9,600 students in ELA. Thus, we retain about 87 percent of observations from the base sample in math, and 89 percent of observations from the base sample in ELA.

Our sample size is close to the sample size reported by [Garrett and Steinberg \(2015\)](#). We obtain a slightly larger sample size because we complete missing test score information in the student file with information from the district-wide files. These files are part of the 2018 MET release and were not available when [Garrett and Steinberg \(2015\)](#) published their study.

Our estimation dataset does not contain any missing information for the main variables used in the analysis; however, it contains some missings in teacher and student demographics. With the exception of free/reduced-price lunch eligibility (30 percent missings in math and 25 percent missings in ELA), the student background variables contain (virtually) no missings. Further missings occur in the teacher demographics: teacher gender and race/ethnicity are missing for 3–4 percent of the estimation sample. Teacher experience and teachers' education

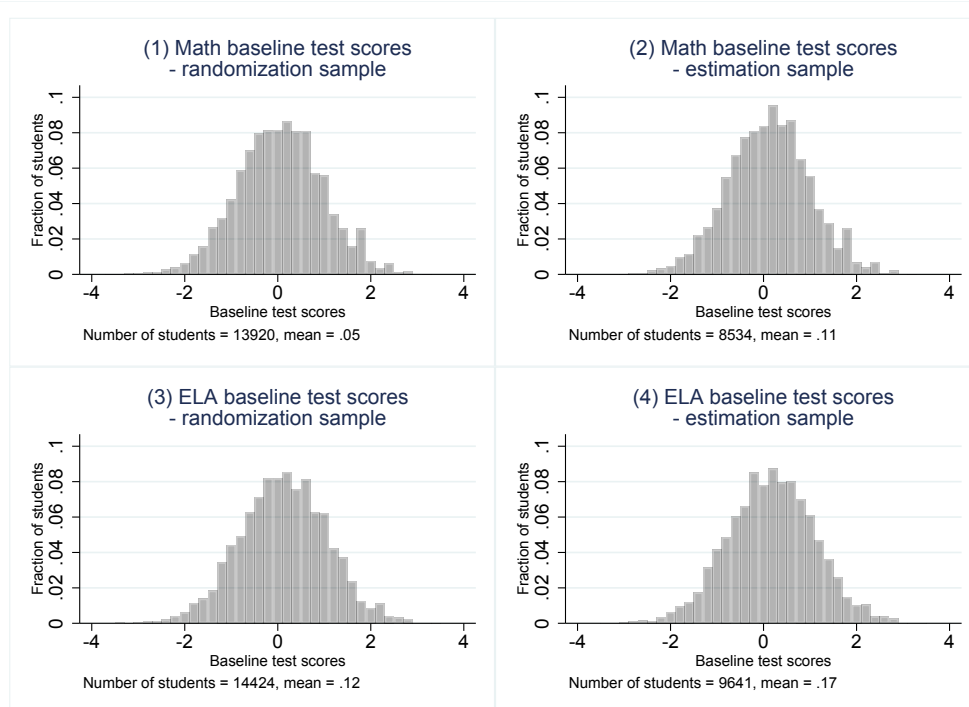
are missing for about 30 percent of the estimation sample each, because one school district did not provide information on teachers' education, and another district did not provide information on teachers' experience. Full information on teacher demographics is only available for 42 percent of the math sample, and for 48 percent of the ELA sample.

B.3 Sample comparisons

This section investigates in how far the randomization sample differs from the estimation sample. It compares the distribution of student baseline test scores and teacher FFT in the original randomization sample with the distribution in the estimation sample.

The randomization sample consists of about 15,700 students in math and 16,300 students in ELA; baseline test scores are available for about 13,900 students in math and for about 14,400 students in ELA (see Figure B.1). The remaining student observations can neither be matched to the district-wide files nor be matched to the student file. Our estimation sample contains only students in the student file – because we use this file to identify realized teachers and peers – and is thus considerably smaller (about 8,500 students in math and about 9,600 students in ELA). Students in the estimation dataset have on average higher baseline test scores than the students in the randomization dataset. The difference amounts to 0.06 standard deviations in math and to 0.05 standard deviations in ELA.

Figure B.1: Distribution of baseline test scores in the randomization sample and the estimation sample

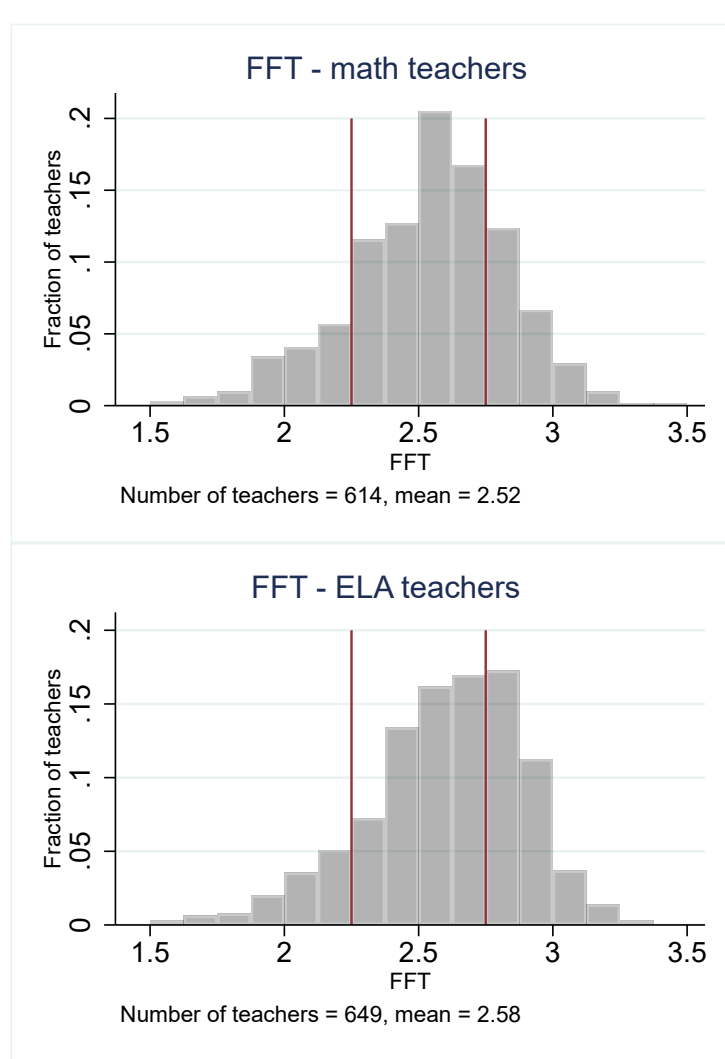


Note: The figure compares the distributions of student baseline test scores across the randomization sample and the estimation sample for math (Panels 1 and 2) and ELA (Panels 3 and 4).

Teacher FFT does not differ appreciably between the randomization sample and the estimation sample (see Figure B.2). In the math randomization sample for grades 4–8, 666

teachers have non-missing information on FFT; the estimation sample includes 614 teachers. In both samples, the average FFT is 2.52. In the ELA randomization sample for these grades, 705 teacher have non-missing FFT, and the estimation sample contains 649 teachers. The average FFT in both samples are nearly identical with 2.57 and 2.58.

Figure B.2: Distribution of FFT in the randomization sample and the estimation sample



Note: The figure compares the distributions of teacher FFT across the randomization sample and the estimation sample for math teachers (Panels 1 and 2) and ELA teachers (Panels 3 and 4). The red lines denote the FFT cutoffs we use to classify teachers in our preferred specification.

B.4 Comparison of the FFT with measures of teacher quality

In the estimation sample, teacher FFT is positively correlated with a number of other teacher quality measures in the MET data: alternative observational measures (CLASS, MQI, PLATO), value-added, content knowledge in ELA, and principal ratings.

Table B.2: Correlation of the FFT with measures of teacher quality

	(1)	(2)	(3)	(4)	(5)	(6)
	N	Correlation with FFT Corr.	p-value	Mean by FFT category low	middle	high
Panel A. Math sample						
Classroom observations						
CLASS	483	0.672	(0.000)	3.84	4.34	1.42
MQI	608	0.396	(0.000)	1.23	1.32	8.42
Experience	443	0.062	(0.190)	6.40	7.56	8.24
Education (Master's degree)	468	0.066	(0.151)	0.31	0.42	0.45
Value added	614	0.162	(0.000)	-0.06	0.01	0.04
Content knowledge	476	-0.088	(0.054)	59.46	55.61	54.46
Principal ratings						
Top 5	445	0.097	(0.042)	0.16	0.19	0.23
Top 10	445	0.111	(0.019)	0.46	0.45	0.57
Top 50	445	0.133	(0.005)	0.70	0.80	0.88
Panel B. ELA sample						
Classroom observations						
CLASS	532	0.691	(0.000)	3.79	4.37	4.70
PLATO	637	0.601	(0.000)	2.17	2.58	2.70
Experience	480	-0.016	(0.734)	7.01	7.47	7.04
Education (Master's degree)	495	0.004	(0.926)	0.31	0.44	0.31
Value added	645	0.116	(0.003)	0.00	0.00	0.03
Content knowledge	499	0.142	(0.001)	62.36	64.71	67.69
Principal ratings						
Top 5	467	0.210	(0.000)	0.07	0.18	0.24
Top 10	467	0.227	(0.000)	0.36	0.48	0.60
Top 50	467	0.22	(0.000)	0.64	0.82	0.85

Note: The table displays correlations between the FFT and different measures of teacher quality, based on the estimation sample (Columns 2–3). The full sample consists of 614 teachers in math (Panel A) and 649 teachers in ELA (Panel B). The table also displays the means of the measures by FFT category (Columns 4–6). CLASS: Classroom observation scoring system, MQI: Mathematical Quality of Instruction, PLATO: Protocol for Language Arts Teaching Observations. For details on these measures, see [White et al. \(2019\)](#).

B.5 Information on the CLASS measure

Information from the MET User Guide. The following information about the CLASS (Classroom Assessment Scoring System) protocol is taken from the MET User Guide (White et al., 2019, p. 32–33):

“CLASS is an observational protocol designed to measure the extent to which teachers effectively support children’s social and academic development. Two different versions of CLASS were used in the MET Study: the Upper Elementary (Grades 4-5) and the Secondary (Grades 6-9).

“The CLASS instrument is divided into three broad domains of measurement: Emotional Support, Classroom Organization, and Instructional Support. Each domain, in turn, is measured by a number of dimensions. The domain “Emotional Support,” for example, refers to the emotional tone in a classroom, which can be measured along four dimensions: positive climate, negative climate, teacher sensitivity, and regard for student perspectives. The domain “Classroom Organization” refers to the ways a classroom is structured to manage students’ behavior, time, and attention, which can be measured along three dimensions: behavior management, productivity, and instructional learning formats. The domain “Instructional Supports” refers to the ways a teacher provides supports to encourage student conceptual understanding and student problem solving and can be measured along four dimensions: content understanding, analysis and problem solving, instructional dialogue, and quality of feedback. [...]

“CLASS scoring is done using a detailed scoring rubric. In this rubric, a classroom is scored on each instructional dimension at 15-minute intervals using a 7-point scale. For the MET Study, only the first 30 minutes of each video was scored. Scores are assigned based on anchor descriptions of what is to be observed in order for a classroom to be scored at “high,” “mid,” and “low” points on the 7-point scale. In the MET Study, dimension scores are often aggregated to higher levels of analysis simply by averaging raters’ scores to get a single segment score and then calculating the harmonic mean of segment scores across all segments for a particular target of measurement (e.g., a day, a class section, a teacher). Standard errors of measurement for these derived scores are not generally reported.”

Use of the CLASS measure in our study. We use the CLASS measure to test the sensitivity of our results to the classroom observation protocol used. To construct a unique measure for each teacher, we take the average across the three CLASS domains.

Our estimation sample contains 466 teachers in the math sample and 495 teachers in the ELA sample. In the math sample, CLASS ranges from 2.54 to 5.58, with a mean of 4.34. In the ELA sample, CLASS ranges from 2.95 to 5.58, with a mean of 4.39. We split teachers in three categories according to this measure. We choose the cutoff values of 4 and 4.5 to carry out the split. Both in the math and in the ELA sample, 20 percent of the teachers are classified as low, 42 percent as middle, and 38 percent as high.

C Optimal allocation: linear program

The linear program aims at maximizing the aggregate test score outcomes in the data. We use $\hat{Y}_c(w)$ to denote the predicted aggregate outcome of classroom c when assigned a teacher of level w , where $w = w_L$ when assigned a low-, $w = w_M$ when assigned a middle-, and $w = w_H$ when assigned a high-FFT teacher (see Section 4 for details). We use C to denote the total number of teachers in the dataset, and C_w to denote the number of teachers of level w in the dataset. We define an indicator variable α_{cw} , which takes the value 1 if classroom c is taught by a teacher of level w , and 0 otherwise. We also define an assignment matrix A , which contains all α_{cw} . The linear program can then be written as:

$$\max_A \sum_{c=1}^C \sum_{w \in \{w_L, w_M, w_H\}} \alpha_{cw} \hat{Y}_c(w)$$

subject to

$$\sum_{w \in \{w_L, w_M, w_H\}} \alpha_{cw} = 1 \quad \forall c \in C$$

$$\sum_{c=1}^C \alpha_{cw} = C_w \quad \text{for } w \in \{w_L, w_M, w_H\}$$

$$\alpha_{cw} \in \{0, 1\}$$

This is a transportation problem with $C + 3 + (3 \times C)$ constraints. We solve the transportation problem in R using `lpSolve`.