A Data Appendix

A.1 Student-level data

We use student records from the NCERDC over the years of 2006-2007 through 2017-2018 to measure multi-dimensional teacher productivity in raising math test scores. This provides 8,177,312 student-year observations. We focus on math teachers in grades 4 through 8 to capture the majority of teachers with prior performance data who enter the applicant pool. We use third to seventh grade math and reading scores as lagged achievement. Test score data as well as student demographics such as ethnicity, gender, gifted designation, disability designation, whether the student is a migrant, whether the student is learning English, whether the student is economically disadvantaged, test accommodations, age, and grade come from the NCERDC master-build files. We use only data from standard end-of-grade exams. This leaves us with 5,322,896 student-year observations.

Beginning in the 2006-2007 school year, the state began recording course membership files linking students directly to courses and instructors. Prior to this change, teachers were linked to students through data on the proctors of the end-of-course exams. The new course membership files provide stronger teacher–subject-student links than the previous system, in which teachers were more frequently linked to the wrong subject (Harris and Sass, 2011).

With the course membership files, we still must determine which teacher is most responsible for teaching math. We use a tiered system. We use course codes (starting with "20") and course names (including text "math," "alg," "geom," and "calc") to do so. We also want to prioritize standard classes as opposed to temporary or supplemental instruction (course names including text such as "study," "special," "resource," "pullout," "remed," "enrich," "indiv," and "except"). We assign students to the teacher most likely to be the math teacher according to the following rules: (1) Students are assigned first to a high-certainty math teacher (the course code and title indicate a standard math class without mention of supplemental instruction). (2) Students with self-contained teachers are assigned to that teacher if there is no high-certainty math teacher present. (3) Students with course codes and course titles indicating math teachers but no self-contained teachers or high-certainty math teachers are assigned to those middle-certainty math teachers. (4) Students with a teacher of a course that either has a math code or a math course title but no other math course or self-contained teacher are assigned to those low-certainty math teachers. (5) Students with a science course code but no math course or self-contained courses are assigned to their science teachers to accommodate recent trends of math and science block scheduling. We exclude classes in which more than half the class requires special accommodations. Ultimately, our sample for constructing teacher value-added measures is composed of 5,159,337 student-year observations providing measures for 38,566 teachers.

A.2 Application and vacancy data

Our application and vacancy data cover the 2010-2019 cycles. We restrict our sample to applications and vacancies for on-cycle, standard elementary school positions. We show how these restrictions change the sample in Appendix Table A1.

We define on-cycle as positions that receive their first applications of a cycle between April 1 and August 15.

We select standard elementary school positions by filtering on the vacancy type (”instructional”) and the vacancy title. Seventy percent of posted vacancies are for instructional positions. We require that the position indicate elementary school grades by having at least one of the following text strings in the title: “k-”, “3rd”, “4th”, “5th”, “-5”, “-6”, “4-6”, or ”elem”. 39% of vacancies include at least one of these strings in the title.

We then exclude positions with specific subjects mentioned in the title or indications that the position is non-standard (“specialized”, “end of year”, “interim”, “assistant”, “virtual”, “resource”, “itinerant”, “exchange”, “extensions”, “immersion”, “academic support”, “temporary”, “continuous”, “early end”, “interventionist”, or “substitute”). With all of the restrictions above, our final sample consists of 20% of the full set of applications, 25% of the full set of applicants, and 7% of the full set of vacancies.

We code the application’s outcome into whether the candidate is hired (“Accepted-Pending Licensure”, “Hired”, “Hiring Request in Process”, “Offer Accepted”), declines an offer (“Offer Declined”), offered an interview (“Completed BEI Interview”, “Contact for Interview”, “Interview Scheduled”, “Invited to Complete Virtual Interview”, “Invited to Interview”, “Recommended for Interview (By Request)”), or given a positive rating (“1st Choice”, “2nd Choice”, “Highly Recommend for Interview”, “Recommend”, “Recommend for Interview”, “Recommendation Accepted”, “Strong Candidate”). These categories are encodings of a single variable, so they are mutually exclusive (i.e., if a candidate is hired, the prior outcome may be overwritten). For robustness analysis, we also split up the remaining applications into middle ratings (“Attended Info Session/Class”, “Hold for Later Consideration”, “Invited to Info Session/Class”, “Possible recommend for interview”, “Recommend with Hesitation”), negative ratings (“Failed Job Questionnaire”, “Incomplete Application”, “Ineligible Selection”, “Not Good Fit”, “Not Qualified”, “Pool - Ineligible”, “SS - INELIGIBLE”, “Screened - Not Selected”), withdrawals (“Candidate Withdrew Interest”), or no evaluation (“Eligible Selection”, “New”, “Pool - Eligible”, “Pool Candidate”).

A.3 Matching across datasets

For this project the North Carolina Education Research Data Center (NCERDC) combined records held there on teacher work histories, school characteristics, and student achievement with data provided by a large urban school district containing further personnel files, open positions within the school district, and applications for those positions. They performed an interactive fuzzy match using names and birth year. For teachers who had a sufficiently good match (that is, a unique name-birth-year combination), we have a de-identified ID that allows us to connect their platform data to their staffing records and students' achievement.

The NCERDC reports that of the 74,395 applicants to positions, 29,008 are matched to NCERDC records. Many of these applicants never teach in the state and thus would not be expected to match. Of the 26,983 employees listed within the district, 20,966 are matched to NCERDC records. However, the match rate is much better among personnel who teach tested subjects. Of the 13,982 teachers with EVAAS scores in the district, 13,865 are matched to the NCERDC data.

A.4 Sample characteristics

Returning to Appendix Table A1, we see how the sample's characteristics varies with sample restrictions. The "Elementary Sample" restricts to on-cycle elementary school instructional positions without specialization, the "Value-Added Sample" further restricts to teachers with value-added forecasts based on prior years, and the "2015 Sample" further restricts to the 2015 application cycle. We use the "Elementary Sample" for estimating principal preferences, the "Value-Added Sample" for estimating teacher preferences, and the "2015 Sample" for estimating counterfactual allocations.

We see a few expected patterns based on the sample restrictions. For the last two columns, we require teachers to have value-added forecasts based on data from prior years. This restrictions leads us to a more experienced sample of teachers. These teachers are more likely both to already be in the district and to transfer to a new school (from a prior school or from out of district). We also see these teachers have lower application rates, perhaps because many already have in-district placements. We see little change in the teacher sample's mean value-added (by student type or at a representative school) or choice set size. The mean characteristics in the positions sample also change minimally with the sample restrictions.

B Omitted details on value-added model: assumptions, results, and validation

B.1 Formal statement of assumptions for value-added model

Here we formally state the assumptions that were informally discussed in Section 3.

Assumption 1 (Exogeneity and stationarity of classroom and student-level shocks)**.** *Classroom-student-type shocks (θ_{cmt}) are independent across classrooms and independent from teachers and schools. Classroom-student-type shocks follow a stationary process:*

$$\mathbb{E}[\theta_{c0t}|t] = \mathbb{E}[\theta_{c1t}|t] = 0 \tag{A1}$$

$$Var(\theta_{c0t}) = \sigma^2_{\theta_0},\ Var(\theta_{c1t}) = \sigma^2_{\theta_1},\ Cov(\theta_{c0t}, \theta_{c1t}) = \sigma_{\theta_0\theta_1} \tag{A2}$$

for all t.

Student-level idiosyncratic variation is independent across students and independent from teachers and schools. Student-level shocks follow a stationary process depending on the student's type:

$$\mathbb{E}[\tilde{\varepsilon}_{it}|t] = 0 \tag{A3}$$

$$Var(\tilde{\varepsilon}_{it}) = \sigma^2_{\varepsilon m} \ \textit{for } m = 0, 1 \tag{A4}$$

for all t.

Assumption 2 (Joint stationarity of teacher effects)**.** *The non-experience part of teacher value-added for each student type follows a stationary process that does not depend on the teacher's school. The covariances between the teacher's value-added across student types depend only on the number of years elapsed:*

$$\mathbb{E}[\mu_{j0t}|t] = \mathbb{E}[\mu_{j1s}|t] = 0 \tag{A5}$$

$$Var(\mu_{j0t}) = \sigma^2_{\mu_0},\ Var(\mu_{j1t}) = \sigma^2_{\mu_1},\ Cov(\mu_{j0t}, \mu_{j1t}) = \sigma_{\mu_0\mu_1} \tag{A6}$$

$$Cov(\mu_{j0t}, \mu_{j0,t+s}) = \sigma_{\mu_0 s},\ Cov(\mu_{j1t}, \mu_{j1,t+s}) = \sigma_{\mu_1 s} \tag{A7}$$

$$Cov(\mu_{j0t}, \mu_{j1,t+s}) = \sigma_{\mu_0\mu_1 s} \tag{A8}$$

for all t.

Assumption 3 (Independence of drift and school effects)**.** *Let $\bar{\mu}_{jm}$ be teacher j's mean value-added for student type m. Let k be j's assigned school in year t. Then:*

$$(\mu_{jmt} - \bar{\mu}_{jm}) \perp \mu_k \ \textit{for } m = 0, 1. \tag{A9}$$

B.2 Additional details on estimation

In the first step, we estimate β_l by regressing test scores (standardized to have mean 0 and standard deviation 1 in each grade-year) on a set of student characteristics (X_{it}) and classroom-student-type fixed effects:

$$A^*_{it} = \beta_s X_{it} + \lambda_{cmt} + \upsilon_{it}. \tag{A10}$$

For characteristics, we include ethnicity, gender, gifted designation, disability designation, whether the student is a migrant, whether the student is learning English, whether the student is economically disadvantaged, test accommodations, age, and grade-specific cubic polynomials in lagged math and lagged reading scores. We subtract the estimated effects of the student characteristics to form the first set of residuals, $\hat{\nu}_{it}$:[19]

$$\hat{\nu}_{it} = A^*_{it} - \hat{\beta}_s X_{it}. \tag{A11}$$

These student-level residuals include teacher, school, and classroom components, as well as idiosyncratic student-level variation.

In the second step, we project the residuals onto teacher fixed effects, school fixed effects, and the teacher experience return function. Following the literature, we specify the experience return function as separate returns for every level of experience up to 6 years, and then a single category of experience of at least 7 years:

$$\hat{\nu}_{it} = \sum_{e=1}^{6} \alpha^e \mathbb{1}\{Z_{jt} = e\} + \alpha^7 \mathbb{1}\{Z_{jt} \geq 7\} + \mu_{jm} + \mu_k + \mu_t + \varepsilon_{it}, \tag{A12}$$

where $\varepsilon_{it} = (\mu_{jmt} - \mu_{jm}) + \theta_{cmt} + \tilde{\varepsilon}_{it}$. We then form a second set of student-level residuals by subtracting off the estimated school and experience effects:

$$A_{it} = \hat{\nu}_{it} - \left(\sum_{e=1}^{6} \hat{\alpha}^e \mathbb{1}\{Z_{jt} = e\} + \hat{\alpha}^7 \mathbb{1}\{Z_{jt} \geq 7\} + \hat{\mu}_k + \hat{\mu}_t \right). \tag{A13}$$

We aggregate these student-level residuals into teacher-year mean residuals for each student type: $\bar{A}_{jmt}$. Let $\mathbf{A}_j^{-t}$ be a vector of mean residuals for each student type-year that j teaches in the data, prior to year t.

In the final step, we form our estimate of teacher j's value-added (net of experience effects) in year t for type m as the best linear predictor based on the prior data in our sample:

$$\hat{\mu}_{jmt} \equiv \mathbb{E}^* \left[\mu_{jmt} | \mathbf{A}_j^{-t} \right] = \psi_m' \mathbf{A}_j^{-t}, \tag{A14}$$

where ψ_m is a vector of reliability weights. We estimate ψ_m following Delgado (2021). Our estimate of teacher j's composite value-added at school k in year t is:

$$\widehat{VA}_{jkt} = p_{k0t} \hat{\mu}_{j0t} + p_{k1t} \hat{\mu}_{j1t} + f(Z_{jt}; \hat{\alpha}). \tag{A15}$$

Variation in the data: We now discuss the variation in the data that pins down key parameters. The coefficient on student characteristics uses how test scores vary with within-classroom-student type variation in student characteristics.[20] The school effects use the change in (student) output

[19] Here we deviate from the standard notation, by introducing $\hat{\nu}_{it}$. Our procedure has two residualization steps because we include classroom-student type fixed effects in the first step, which would subsume the teacher and school fixed effects. We thus decompose student residuals into teacher and school components in a second step.

[20] Because we include classroom-student-type fixed effects, our model allows for an arbitrary correlation between students' characteristics and the quality of their assigned teachers. Allowing such correlation is important in a context where teachers have some control over where they work.

when teachers switch schools, beyond what would be predicted by drift and by the change in student type composition. Heuristically, if teachers' output regularly increases when teachers transfer to a certain school, then we would estimate a high school effect. The teacher mean effects for each student type are pinned down by relative increases in students' (residualized) test scores across different teachers. We are able to rank teachers both within and across schools, provided teachers and schools are in a set connected by transfers so that we can identify the school effects.

Finally, we identify the parameters of the teacher value-added distribution and the drift process based on the stationarity assumptions and the observations of teachers across years, classrooms, and student types. As an example, the variance of the teacher effects for student type m is identified by the covariance between a teacher's mean student residuals for student type m in two different classrooms in the same year.[21] With our assumptions that classroom and student shocks are uncorrelated across classrooms, the only reason a teacher's students would have similar (residualized) outcomes is the teacher's value-added.

B.3 Testing for comparative advantage

Our measures forecast teachers' future value-added without bias. Our high estimated correlation between a teacher's effectiveness with the two student types raises the question of whether our estimates of comparative advantage simply reflect statistical noise. We perform three exercises to test our multi-dimensional value-added model versus a single-dimensional model.

First, we estimate standard errors and confidence intervals for the structural parameters in our production model. The estimated correlation in teacher value-added across student types is 0.86. We can, however, decisively reject a correlation of 1 as the bootstrap standard error is 0.035, with a 95% confidence interval of $(0.73, 0.87)$ (Appendix Table A2).

Second, we perform a likelihood-ratio test comparing our model with a model with one-dimensional teacher value-added. We take the mean residuals at the level of the teacher-classroom-student type, $\bar{A}_{jcmt}$, and collect a teacher's mean residuals across classrooms and student types, which come from a normal distribution:

$$\begin{pmatrix} \bar{A}_{jc1t} \\ \bar{A}_{jc'2t} \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2_{\mu_1} + \sigma^2_{\theta_1} + \frac{\sigma^2_{\varepsilon_1}}{N_{jc0t}} & \sigma_{\mu_1\mu_2} \\ \sigma_{\mu_1\mu_2} & \sigma^2_{\mu_2} + \sigma^2_{\theta_2} + \frac{\sigma^2_{\varepsilon_2}}{N_{jc2t}} \end{pmatrix} \right). \tag{A16}$$

We compare the likelihoods across our baseline model and an alternate model of homogeneous value-added where $\sigma^2_{\mu_1} = \sigma^2_{\mu_2}$, $\sigma^2_{\theta_1} = \sigma^2_{\theta_2}$, $\sigma^2_{\varepsilon_1} = \sigma^2_{\varepsilon_2}$, and $\sigma_{\mu_1\mu_2} = 0$. Our likelihood-ratio test has 4 degrees of freedom, and we reject the homogeneous value-added model in favor of the heterogeneous model, with a test statistic of 610, so the p-value is arbitrarily small ($p < 0.0001$).[22]

[21] In our setting many elementary school teachers have students from multiple classes. The prevalence of multiple classrooms is increasing over time (Appendix Table A9).

[22] We restrict the sample to one randomly-chosen vector of mean residuals per teacher so that the observations in our

Third, we fix a teacher's type according to whether she is above or below the median in comparative advantage in teaching economically disadvantaged students in pre-transfer schools. We then test whether changes in the share of economically disadvantaged students differentially predict changes in student test score residuals ($\hat{v}_{it}$ from equation A13) in post-transfer schools by teacher-type. The logic of the test is as follows. Under a homogeneous value-added model, changes in the share of economically disadvantaged students should have no bearing on changes in teacher productivity across schools. If our estimated comparative advantage is meaningful, however, then as the share of disadvantaged students rises, teachers with a comparative advantage in teaching disadvantaged students should see gains in average productivity relative to teachers with a comparative advantage in teaching economically advantaged students. Accordingly, we regress across-transfer changes in teacher-by-school average student residuals on across-transfer changes in the share of disadvantaged students interacted with teachers' type. The results appear in Appendix Table A10. For teachers with a comparative advantage in teaching advantaged students in pre-transfer schools, productivity falls as the share of disadvantaged students rises (p-value=0.043). In contrast, for teachers with a comparative advantage in teaching disadvantaged students, productivity rises as the share of disadvantaged students rises (p-value=0.014). These findings indicates that comparative advantage is persistent across settings and predictive of match-specific productivity.

C Within-school assignments

Our analysis focuses on the allocation of teachers across schools in a district. Another margin of allocation could be within-school assignment of teachers based on class size or composition. Ignoring this margin could affect our results in two ways. First, we could understate the potential allocation gains (or even focus on the less important margin). In Table 2 we show that the gains to within-school reallocation are much smaller than the gains from reallocation across schools.

Second, if within-school position characteristics are endogenous, our preference model might be misspecified. For example, suppose that an experienced teacher can negotiate for the Honors class at a school but the inexperienced teacher cannot. We assess this possibility in two ways.

C.1 Persistence of classroom characteristics

If within-school assignment characteristics were endogenous and a function of the teacher's type, we would expect persistence in these characteristics over time. In Appendix Tables A11 and A12, we show the autocorrelations in number of students taught by a teacher and the fraction of students that are economically disadvantaged. In each table's top panel, we show the school-level autocorrelation. We find that differences across schools – which we leverage in our analysis – are fairly persistent.

likelihood are independent. We also find a similar test statistic when we use mean residuals, $\bar{A}_{jcmt}$, from a model where the fixed effects in the residualizing steps are not separated by student type.

In each table's bottom panel, we show the teacher-level autocorrelation where we residualize by school-year fixed effects to isolate the within-school deviation. These within-school differences across teachers – which we do not leverage in our analysis – are not persistent at all.

C.2 Do teachers bargain over student assignment on the job market?

Second, we examine how students are assigned to teachers within and across schools. This question is of particular interest since we would like to know whether teachers bargain with principals over their student assignments. Are sought-after teachers assigned "preferable" class compositions? The primary teacher characteristic we use is *experience*, which principals value and is reliably measured in our data. We first explore the relationship visually. Student attributes have a linear relationship with log(experience), so we estimate models in which the outcome variables are student attributes and the primary explanatory variable is a teacher's log(experience). In regressions, standard errors are clustered by teacher and by year.

Without controlling for school setting, there is a strong relationship between experience and student attributes (see Appendix Table A13). More experienced teachers are assigned higher-scoring students, fewer economically disdavantaged students, more students designated as gifted, and fewer Black students.

Much sorting takes place across schools—as teachers gain experience, they sort to more suburban schools where students are less economically disadvantaged and higher achieving. In the basic cross section, we find that a 100 percent increase in experience reduces poverty shares by 0.037 percentage points (significant at the 0.001 level). When we control for year and school fixed effects, the coefficient on (log) experience falls by over 80 percent to 0.006 (significant at the 0.001 level). We examine the gradient among newly hired teachers. This group is of particular relevance because applicants (as opposed to teachers not applying to new jobs) are the teachers we consider in our counterfactual exercises. When looking at this group, we find no significant relationship between experience and disadvantaged-student assignment, conditional on school-year fixed effects. This suggests that principals do not sort students to teachers based on experience within a school, and indicates that bargaining over student characteristics is unlikely.

The pattern of sorting Black students to teachers is quite similar. We find that doubling teacher experience reduces the Black share of a teacher's class by 0.033 percentage points (significant at the 0.001 level). When looking within a school, the experience gradient falls by 97 percent—the sorting of Black students to teachers is almost exclusively across schools. When we examine the relationship among new hires, the relationship is even smaller and statistically insignificant. The gradient between student test scores and teacher experience attenuates by 90 percent when accounting for school-year fixed effects. There still is a small, systematic difference in test scores which appears to arise from hiring more experienced teachers to serve in gifted-and-talented classrooms. We see very experienced teachers assigned somewhat less desirable class assignments than would be predicted by the rest of the support. It may be that schools encourage older teachers to leave by giving them more difficult class compositions.

In summary, among new teachers, experience does not significantly predict assignment to disadvantaged students or Black students within schools. There is a small experience gradient for higher achieving students among new teachers. It seems teachers of gifted-and-talented classrooms tend to be senior.

D Heterogeneity in application rate gap between Title I and non-Title schools

To showcase unobservable preference heterogeneity, we focus on teacher preferences over a binary characteristic: whether the school has Title I designation. Appendix Table A16 shows that on average teachers are less likely to apply to Title I schools. The application rate to non-Title I schools is almost 16% and to Title I schools is about 14%, and leaving an application gap of close to 2 percentage points (or 10%).

The second and third columns of Appendix Table A16 show why we are able to estimate heterogeneity precisely: the median number of applications choices that each teacher makes is over 65 for both Title I positions and non-Title I positions. Thus, teachers' application sets have the potential to include many or few Title I positions.

Appendix Figure A20 shows that the mean gap in application rates across school types masks substantial heterogeneity. For each teacher, we calculate the gap in application rates (for positions in the teacher's choice set) between Title I and non-Title I schools, and then we plot the distribution of the gap. Visually, the distribution almost appears centered on zero (the median is 0.003). But there is substantial dispersion: the standard deviation of the raw gaps is 0.134.

Naturally, the standard deviation of the raw gap overstates the extent of dispersion because it incorporates noise. We now describe and implement a simple minimum distance estimator for the true standard deviation of the applicant gap. For each teacher j we observe a_{j1} applications sent to a Title I school and c_{j1} is the number of Title I vacancies in the teacher's choice set. We can then estimate

$$\hat{p}_{j1} = \frac{a_{j1}}{c_{j1}}$$

or teacher j's application probability to a Title I school.

Using the natural notation for a "not-Title I" school, we also have:

$$\hat{p}_{j0} = \frac{a_{j0}}{c_{j0}}$$

We can then compute the "gap", or Title I penalty, as

$$\hat{g}_j = \hat{p}_{j1} - \hat{p}_{j0}.$$

We are then interested in the distribution of these gaps – e.g., the standard deviation (sd) of g_j. Naturally, taking $sd(\hat{g}_j)$ will result in an over-estimate of the amount of dispersion.

We fit the following model.

$$p_{j0} = \hat{\bar{p}}_0$$

$$p_{j1} = N(\hat{\bar{p}}_1, \sigma)$$

where $\hat{\bar{p}}_.$ are the population average application probabilities and σ is a parameter to estimate. We estimate σ by simulated method of moments where the moment to match is $sd(\hat{g}_j)$ and we simulate data from the model embedded in the previous two equation using the observed $\{c_{j0}, c_{j1}\}$.

We find that the estimated standard deviation is 0.114, so the visual depiction of noise is in line with the underlying truth. If we take the minimum distance estimate at face value, while on average teachers have higher application probabilities to non-Title I schools, about 44% of teachers have higher application probabilities to Title I instead. Hence, this suggests that even though on average teachers prefer non-Title I schools, there is a substantial amount of heterogeneity in the applicant pool. Depending on how such preference heterogeneity maps into the existing allocation of teachers to schools, policies that make Title I schools more attractive could have small or large effects on teachers' application rates.

E Selection into the transfer market

What explains the differences in student gains between Table 2 and the results depicted in Figure 7b? Here, we compare the teacher transfer market to the broader sample of teachers and positions. We first consider the representation of schools in the transfer market. Unsurprisingly, we see significant over-representation for positions in schools with high proportions of economically disadvantaged students. Appendix Table A14 shows that a 10 percentage point increase in the share of economically disadvantaged students is associated with 0.15 more positions posted. Because the overrepresented type of school is already the more common one (more than half of the students in our district are economically disadvantaged) this means that gains from sorting on comparative advantage are going to be understated in our transfer sample.

The pattern is less pronounced for the number of students a teacher instructs. Though the point estimate implies that an additional 10 students per teacher lowers the number of positions a school posts by 0.2, this relationship is largely driven by outliers, as shown in Appendix Figure A19.

To examine the selection of teachers into the transfer market, we first look at four cohorts, 2010-2013, such that we can follow them for five years. We further restrict attention to those for whom we can measure productivity, leaving us with 553 teachers who entered the state's data during those years. Of those, 207 applied to transfer at some point during the first five years. Only 124 remain in their original school and have not applied to transfer within five years of entering the district. The remaining 287 leave the district. Appendix Table A15 shows that there is very little difference in comparative advantage between teachers who applied to transfer and the teachers who did not. Teachers who apply for transfer have lower and less variable absolute advantage.

Accordingly, it is unlikely that the difference in per-student potential gains is due to teacher selection into transferring, particularly with regard to comparative disadvantage (with disadvantaged students). It is possible that the small differences in absolute disadvantage interacted with the under-representation of large classes accounts for some of the gap. The clearest selection into the transferring market, however, comes from the over-representation of schools with a high concentration of disadvantaged students. With a limited distribution of schools, there is less room to realize the gains from teachers' comparative advantages.

F Principal preferences estimation

We estimate principal preferences via maximum simulated likelihood, where we simulate from the normal distributions of the random effect at the level of the position-year. Let n index each simulation iteration and let $B_{jptn}(\theta)$ be the model-predicted probability that p rates j positively in year t in simulation iteration n at parameter vector θ. For each position p in year t, we construct the simulated likelihood as:

$$L_{pt} = \frac{1}{100}\sum_{n=1}^{100}\prod_{j\in\mathcal{J}_{pt}}(b_{jpt}B_{jptn}(\theta)+(1-b_{jpt})(1-B_{jptn}(\theta))), \tag{A17}$$

where $\mathcal{J}_{pt}$ is the set of teachers who applied to a position p in year t and b_{jpt} is an indicator for whether p rated j positively in the data. Our full simulated log likelihood function is:

$$l = \frac{1}{P}\sum_{p} log L_{pt}. \tag{A18}$$

Table A1: Sample and summary statistics

	Full Sample	Elementary Sample	Value-Added Sample	2015 Sample
Applications				
N	2,163,711	337,754	13,819	2,702
On-Cycle	0.68	1.00	1.00	1.00
Instructional	0.70	1.00	1.00	1.00
Elementary	0.39	1.00	1.00	1.00
Applicants				
N	104,795	14,864	867	178
Female		0.92	0.87	0.89
Black		0.24	0.30	0.25
Hispanic		0.03	0.01	0.03
In-District		0.12	0.43	0.44
Choice Set Size		159.10	151.14	151.35
Application Rate		0.18	0.11	0.10
Transferred		0.23	0.43	0.51
Mean Commute Time		17.78	22.57	22.50
Experience		5.81	9.22	9.89
VA Econ Adv		-0.03	-0.03	-0.04
VA Econ Disadv		-0.02	-0.02	-0.03
Abs Adv		-0.03	-0.03	-0.03
Comp Adv in Econ Disadv		0.01	0.01	0.01
Positions				
N	38,921	1,824	1,784	296
Choice Set Size		1,293.54	71.89	88.63
Application Rate		0.14	0.11	0.10
Mean Class Size		26.40	26.40	25.69
Frac Econ Disadv		0.65	0.65	0.68
Frac Black		0.43	0.43	0.45
Frac Hispanic		0.24	0.24	0.25

The table shows count or mean statistics across different samples. The "Full Sample" includes all of the raw data, the "Elementary Sample" restricts to on-cycle elementary school instructional positions without specialization, the "Value-Added Sample" further restricts to teachers with value-added forecasts based on prior years, and the "2015 Sample" further restricts to the 2015 application cycle (for positions in the 2016 school year). We use the "Elementary Sample" for estimating principal preferences, the "Value-Added Sample" for estimating teacher preferences, and the "2015 Sample" for estimating counterfactual allocations. We do not include mean statistics for applicants and positions for the complete sample because we built the data on the subsample. Commute time is measured in minutes, absolute advantage is value-added at the representative school in the district, and choice set size is the number of positions in a teacher's choice set (Applicants panel) or the number of teachers with the position in their choice set (Positions panel).

Table A2: Teacher Value-Added Structural Parameters

	Estimates	Standard Errors	95% CI Lower Bound	95% CI Upper Bound
$\sigma_{\varepsilon 1}$	0.450	0.000	0.456	0.457
$\sigma_{\varepsilon 2}$	0.470	0.000	0.477	0.479
$\sigma_{\theta 1}$	0.110	0.007	0.108	0.137
$\sigma_{\theta 2}$	0.088	0.015	0.089	0.143
correlation($\theta_{c0t}, \theta_{c1t}$)	0.657	0.162	0.126	0.844
$\sigma_{\mu 1}$	0.249	0.007	0.262	0.284
$\sigma_{\mu 2}$	0.243	0.015	0.254	0.316
correlation(μ_{j0t}, μ_{j1t})	0.859	0.035	0.729	0.872

The table shows the estimates of a subset of the structural parameters of the production model – specifically the parameters corresponding to contemporaneous output. Non-disadvantaged students have index 1 while disadvantaged students have index 2. ε is the student-year idiosyncratic component, θ captures classroom effects, and μ describes a teacher's value-added. The remaining structural parameters describe the drift process of teacher value-added over time. Standard errors and confidence intervals are estimated with 100 bootstrap iterations.

Table A3: Estimated Experience Returns to Teacher Value-Added

	1	2	3	4	5	6	7+
Estimate	0.056	0.077	0.083	0.088	0.088	0.091	0.070
Standard Error	0.004	0.004	0.005	0.005	0.005	0.005	0.005

The table shows the estimated experience returns for math test scores, where the scores have been normalized to have mean 0 and standard deviation 1 for students in a given grade-year. Columns designate the number of prior years of experience. The omitted category is teachers with no prior experience.

Table A4: Potential Gains from Reassignment – Test Score Percentiles

	Per-Student Gains (σ)	As a Fraction of (Best-Actual)	Non-Disadvantaged	Disadvantaged
Alternate Allocations				
Best	0.050		0.089	0.016
Random	-0.003	-0.19	0.017	-0.021
Worst	-0.053	-3.63	-0.052	-0.054
Alternate Policies				
Best w/i School	0.012	0.98	0.015	0.010
Replace Bottom 5% of Teachers	0.011	0.74	0.012	0.010
Targeting Student Types				
Max Non-Disadvantaged VA	0.024	1.66	0.130	-0.072
Max Disadvantaged VA	0.016	1.12	-0.047	0.074

The table shows the potential gains from reassignments of teachers to different schools. Test scores are constructed as the raw score percentile (from 0 to 1), where percentiles are calculated for each grade-year in the state. We then normalize the test scores to be in standard deviation units based on the standard deviation of the uniform distribution. The sample is all teachers with non-missing value-added measures in 2016, along with their corresponding 2016 assignments. Gains come from better matching of teachers to students, as teachers' effectiveness may differ across student types. The first column shows the per-student gains from various allocations relative to the actual allocation. The second column shows the gain as a fraction of the full difference between the best (output-maximizing) and actual allocations. The third and fourth columns show the per-student gains, relative to the actual allocation, for non-disadvantaged and disadvantaged students. The best within school allocation only changes the teacher-classroom assignments within a school. "Replacing Bottom 5% of Teachers" refers to replacing the bottom 5% of teachers according to realized per-student output with teachers with median value-added for each student type. The targeting student types allocations are the ones that maximize per-student output for students of one type only. We assign classrooms the mean student composition and class sizes in that school in 2016 in all allocations except the "Best w/i School" and "Constant Class Size" allocations.

Table A5: Potential Gains from Reassignment – Constant Class Size

	Per-Student Gains (σ)	As a Fraction of (Best-Actual)	non-Disadvantaged	Disadvantaged
Alternate Allocations				
Best	0.021		0.018	0.023
Random	0.000	0.02	0.002	-0.001
Worst	-0.020	-0.94	-0.015	-0.023
Alternate Policies				
Best w/i School	0.003	0.16	0.007	0.000
Replace Bottom 5% of Teachers	0.017	0.82	0.023	0.012
Targeting Student Types				
Max Non-Disadvantaged VA	0.003	0.17	0.123	-0.090
Max Disadvantaged VA	0.005	0.22	-0.111	0.096

The table shows the potential gains from reassignments of teachers to different schools where each school has the same number (but possibly different composition) of students per class. The sample is all teachers with non-missing value-added measures in 2016, along with their corresponding 2016 assignments. Gains come from better matching of teachers to students, as teachers' effectiveness may differ across student types. The first column shows the per-student gains from various allocations relative to the actual allocation. Gains are measured in student standard deviations (σ). The second column shows the gain as a fraction of the full difference between the best (output-maximizing) and actual allocations. The third and fourth columns show the per-student gains, relative to the actual allocation, for non-disadvantaged and disadvantaged students. The best within school allocation only changes the teacher-classroom assignments within a school. "Replacing Bottom 5% of Teachers" refers to replacing the bottom 5% of teachers according to realized per-student output with teachers with median value-added for each student type. The targeting student types allocations are the ones that maximize per-student output for students of one type only. We assign classrooms the mean student composition and class sizes in that school in 2016 in all allocations except the "Best w/i School" and "Constant Class Size" allocations.

Table A6: Same-Race and Same-Gender Effects on Test Scores

	Student Res
Black Teacher - Black Student	0.00225
	(0.00164)
Hispanic Teacher - Hispanic Student	-0.00556
	(0.00549)
Female Teacher - Female Student	0.00478
	(0.000550)
Fixed Effects	Teacher, School
Mean DV	0.0000115
Clusters	37940
N	5158740

An observation is a student-year and the outcome is the student's math score residualized by student demographics including lagged scores, school fixed effects, and teacher experience measures. The regressors include measures of demographic match between student and teacher. The regression includes school fixed effects and teacher fixed effects. Standard errors are clustered at the teacher level.

Table A7: Teacher Value-Added Structural Parameters with Alternate Forms of Heterogeneity

	Race	Achievement
$\sigma_{\varepsilon 1}$	0.465	0.481
$\sigma_{\varepsilon 2}$	0.457	0.439
$\sigma_{\theta 1}$	0.091	0.099
$\sigma_{\theta 2}$	0.110	0.102
correlation($\theta_{c0t}, \theta_{c1t}$)	0.637	0.628
$\sigma_{\mu 1}$	0.233	0.240
$\sigma_{\mu 2}$	0.261	0.282
correlation(μ_{j0t}, μ_{j1t})	0.900	0.844

The table shows the estimates of a subset of the structural parameters of production models with alternate forms of heterogeneous teacher effects – specifically by race and prior achievement. In the first column, non-white students have index 1 while White students have index 2. In the second column, students with below median prior math achievement have index 1 while students with above median prior math achievement have index 2. ε is the student-year idiosyncratic component, θ captures classroom effects, and μ describes a teacher's value-added. The remaining structural parameters describe the drift process of teacher value-added over time.

Table A8: Application timing

	Obs	Mean days	Median days	Share 0 days
Stock	196,779	3.6	0	0.72
Flow	146,382	2.1	0	0.75

(a) Wait times until applying

	Obs	Mean fraction of days	Mean fraction of applications	Mean days since posting
First day	14,864	0.61	0.65	23.47
Subsequent days	40,850	0.14	0.13	11.55

(b) First day versus subsequent days

	Obs	April or before	May	June	July	August
First day (all teachers)	14,864	0.20	0.25	0.22	0.18	0.15
Last day (all teachers)	14,864	0.09	0.15	0.21	0.26	0.29
First day (transfers)	2,547	0.27	0.30	0.24	0.14	0.05
Last day (transfers)	2,547	0.10	0.17	0.25	0.29	0.19

(c) Timing of first and last days

The tables show statistics related to application timing. Panel (a) shows how long it took an applicant to apply to positions that were in "stock" (already posted) on the day the teacher first applied on the platform or in "flow" (posted after the day the teacher first applied on the platform). Panel (b) shows application statistics for the first day a teacher applied on the platform in a cycle versus subsequent days. "Mean days since posting" is the mean number of days a vacancy had been posted at the time the teacher applied. Panel (c) shows the (monthly) timing of when an applicant's first and last application days of the cycle occurred. "All teachers" includes all applicants while "transfers" includes just teachers who ended up in new schools.

Table A9: Multi-classroom teacher prevalence

Year	All	Grade 4	Grade 5	Grade 6	Grade 7	Grade 8
2012	0.264	0.109	0.187	0.618	0.621	0.631
2013	0.287	0.124	0.210	0.636	0.631	0.649
2014	0.300	0.152	0.227	0.633	0.625	0.644
2015	0.363	0.256	0.345	0.615	0.598	0.602
2016	0.391	0.305	0.392	0.595	0.591	0.595
2017	0.385	0.291	0.399	0.612	0.569	0.596
2018	0.393	0.307	0.425	0.596	0.586	0.578
Estimation sample	0.417					

The table shows the prevalence of teachers having multiple classrooms, separately by teacher's grade and year. The sample includes teachers for whom we can calculate math value-added. Our estimation sample consists of teachers, with value-added forecasts, who applied to elementary school positions.

Table A10: Predicting Student Residuals by Teacher Type

	Student res	Student res
Share disadvantaged	-0.0549	-0.0409
	(0.0251)	(0.0202)
Share disadvantaged x CA in disadvantaged	0.0820	0.0697
	(0.0356)	(0.0283)
Num teachers	3214	3214
Num students	157671	157671
Mean CA	-0.00805	-0.00805
SD CA	0.0624	0.0624
Controls	No	Yes

The table assesses whether changes in the share of economically disadvantaged students predict changes in student test score residuals differently by teacher type across transfers. Teacher type is defined by comparative advantage in pre-transfer schools, with “CA in disadvantaged” an indicator for whether the teacher is above median in comparative advantage in teaching disadvantaged students. The outcome is changes in average teacher-by-school student residuals across transfers. “Share disadvantaged” is the change in the average share of economically disadvantaged students teacher j taught when moving from one school to another. Controls include a cubic in average experience in the school, an indicator for experience missingness, and transfer year indicators. Standard errors are clustered at the teacher level.

Table A11: Autocorrelations in class size

Class size, school level

Variables	Class Size t	Class Size t-1	Class Size t-2	Class Size t-3	Class Size t-4
Class Size t	1.0000				
Class Size t-1	0.7329	1.0000			
	(0.0000)				
Class Size t-2	0.6248	0.6966	1.0000		
	(0.0000)	(0.0000)			
Class Size t-3	0.4093	0.5261	0.6598	1.0000	
	(0.0000)	(0.0000)	(0.0000)		
Class Size t-4	0.3722	0.3746	0.4365	0.5796	1.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	

Nb. obs. : 247

Class size, teacher level

Variables	(Res.) Size t	(Res.) Size t-1	(Res.) Size t-2	(Res.) Size t-3	(Res.) Size t-4
(Res.) Size t	1.0000				
(Res.) Size t-1	0.3668	1.0000			
	(0.0000)				
(Res.) Size t-2	0.2688	0.3717	1.0000		
	(0.0000)	(0.0000)			
(Res.) Size t-3	0.2900	0.1272	0.2699	1.0000	
	(0.0000)	(0.0186)	(0.0000)		
(Res.) Size t-4	0.1173	0.1438	0.0698	0.3098	1.0000
	(0.0301)	(0.0077)	(0.1978)	(0.0000)	

Nb. obs. : 342

The table shows correlations (within unit) between class size in one year and class size in a prior year. In the top panel, a unit of analysis is a school and class size is the mean across all of the school's classrooms (that generate math test scores). In the bottom panel, a unit of analysis is a teacher and class size is residualized by school-year fixed effects such that residual class size compares how a teacher's class size deviates from the school-year mean.

Table A12: Autocorrelations in class composition

Class composition, school level

Variables	Frac Disadv t	Frac Disadv t-1	Frac Disadv t-2	Frac Disadv t-3	Frac Disadv t-4
Frac Disadv t	1.0000				
Frac Disadv t-1	0.9602	1.0000			
	(0.0000)				
Frac Disadv t-2	0.9430	0.9555	1.0000		
	(0.0000)	(0.0000)			
Frac Disadv t-3	0.9363	0.9370	0.9496	1.0000	
	(0.0000)	(0.0000)	(0.0000)		
Frac Disadv t-4	0.9435	0.9467	0.9554	0.9775	1.0000
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	

Nb. obs. : 247

Class composition, teacher level

Variables	(Res.) Dis t	(Res.) Dis t-1	(Res.) Dis t-2	(Res.) Dis t-3	(Res.) Dis t-4
(Res.) Dis t	1.0000				
(Res.) Dis t-1	0.3170	1.0000			
	(0.0000)				
(Res.) Dis t-2	0.2898	0.3200	1.0000		
	(0.0000)	(0.0000)			
(Res.) Dis t-3	0.1524	0.2076	0.3723	1.0000	
	(0.0047)	(0.0001)	(0.0000)		
(Res.) Dis t-4	0.0921	0.0512	0.2203	0.3925	1.0000
	(0.0889)	(0.3450)	(0.0000)	(0.0000)	

Nb. obs. : 342

The table shows correlations (within unit) between class composition (fraction of students that are economically disadvantaged) in one year and class composition in a prior year. In the top panel, a unit of analysis is a school and class composition is the (weighted) mean across all of the school's classrooms (that generate math test scores). In the bottom panel, a unit of analysis is a teacher and class composition is residualized by school-year fixed effects such that residual class composition compares how a teacher's class composition deviates from the school-year mean.

Table A13: Teacher experience and student assignment

	(1) Outcome	(2) Outcome	(3) Outcome	(4) Outcome	(5) Outcome
Outcome: Share economically disadvantaged students assigned					
log(experience)	-0.0369	-0.0311	-0.0063	-0.0029	-0.0021
	(0.0013)	(0.0028)	(0.0005)	(0.0011)	(0.0011)
Outcome: Share Black students assigned					
log(experience)	-0.0331	-0.0195	-0.0010	-0.0008	-0.0005
	(0.0010)	(0.0023)	(0.0004)	(0.0008)	(0.0010)
Outcome: Average student lagged math score					
log(experience)	0.0887	0.0474	0.0461	0.0173	0.0115
	(0.0023)	(0.0049)	(0.0016)	(0.0033)	(0.0041)
Outcome: Share gifted status					
log(experience)	0.0231	0.0106	0.0161	0.0053	0.0074
	(0.0007)	(0.0014)	(0.0006)	(0.0012)	(0.0016)
New only		X		X	X
Year FE			X	X	
School FE			X	X	
School-year FE					X
N	1,879,666	258,723	1,879,666	258,723	258,723

Standard errors in parentheses.

The table shows separate regression results for different outcomes on the log of a teacher's prior experience. Outcomes are mean characteristics of the students in a teacher's classroom. "New only" indicates that the sample only includes teachers new to the school; thus, the regression compares outcomes across teachers new to the school depending on the teacher's experience.

Table A14: Predicting posted positions

	(1) Positions	(2) Positions
Class size	-0.0199	
	(0.0125)	
Fraction disadvantaged		1.503
		(0.544)
N	116	116

An observation is a school-year. The outcome is the number of positions posted in an application cycle and the regressors are characteristics of the school's mean class. Robust standard errors are in parentheses.

Table A15: Transferring and non-transferring teachers' value added

	(1)			(2)		
	Did not apply			Applied to transfer		
	mean	sd	count	mean	sd	count
Comparative advantage	0.0001	0.0351	528	-0.0002	0.0367	506
Absolute advantage	0.0034	0.1210	528	0.0219	0.1508	506

The table shows the means and standard deviations of absolute and comparative advantage for teaching economically advantaged students by whether the teacher ever submits an application to transfer. An observation is a teacher with a value-added forecast. These are pooled over years 2010 through 2018.

Table A16: Applications to Title I and non-Title schools

	Obs	Mean choice set	Median	Mean prob.	25th	50th	75th	Std. dev.	Overall mean prob.
Title I	14,747	85.3	68	0.176	0.010	0.056	0.264	0.237	0.137
non-Title I	14,747	74.0	66	0.176	0.013	0.084	0.270	0.217	0.155
Gap	14,747			-0.001	-0.049	0.003	0.041	0.134	-0.018

The table shows application statistics to positions at Title I and non-Title I schools. Columns (2) and (3) show the mean and median choice set sizes for an applicant. "Gap" shows the difference in statistics across the two school types.

Figure A1: Value-Added Drift Parameters

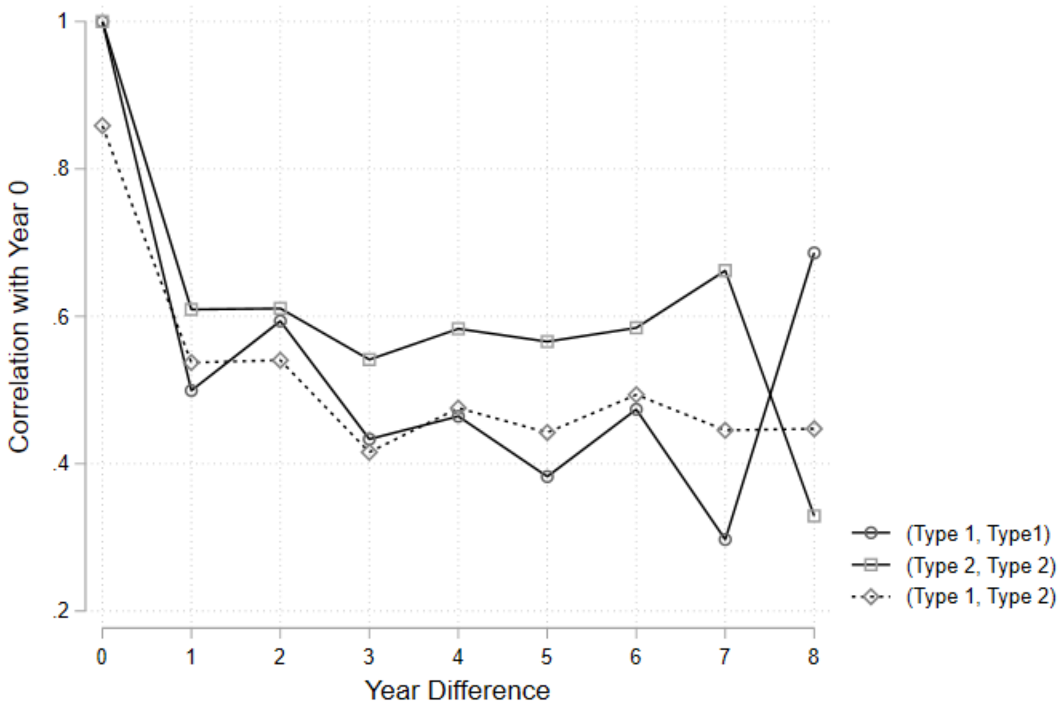

The figure shows the estimated correlations between teacher value-added in different years. The x-axis captures the year difference between the teacher's value-added measures. The three lines reflect correlations in teacher value-added within student type (1 for non-disadvantaged students, 2 for disadvantaged students) or across student type.

Figure A2: Math Comparative Advantage Forecast Unbiasedness

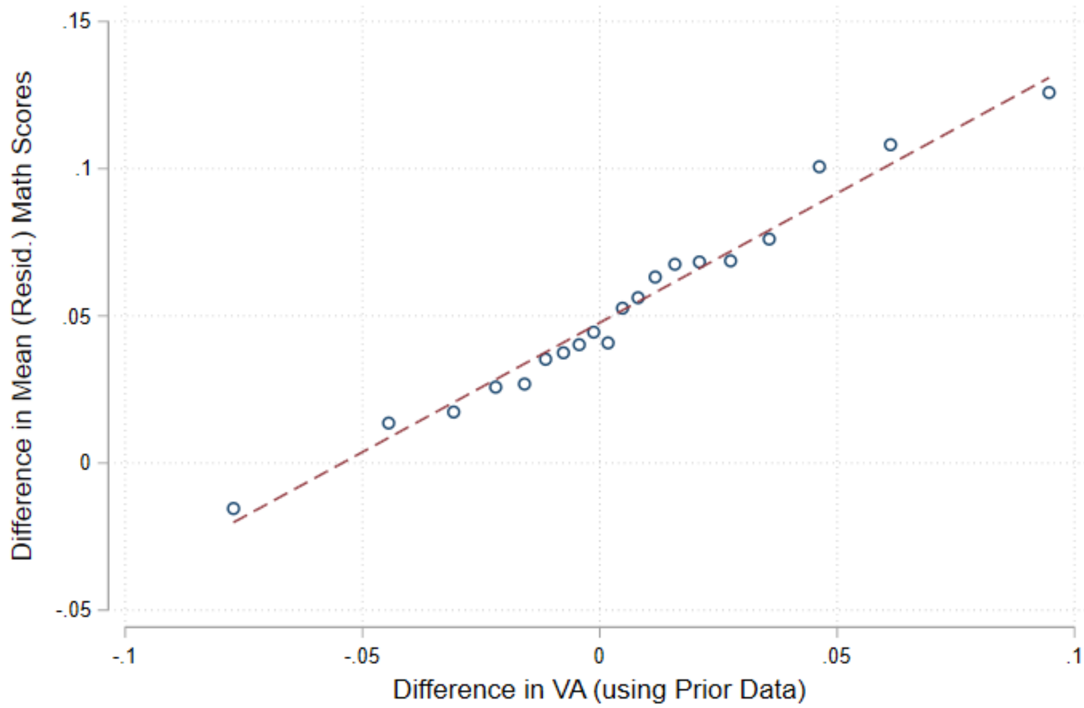

The figure is a binscatter, where an observation is a teacher-year and "Difference in VA" is the difference in a teacher's math value-added between economically disadvantaged and advantaged students. Value-added estimates are predictions using data from prior years. Units are student standard deviations. The y-axis is the difference in mean student math test score, residualized by student demographics including lagged scores, school fixed effects, and teacher experience measures. The mean is taken over all students (of a given type) for a given teacher-year and the difference is between a teacher's economically disadvantaged and advantaged students.

Figure A3: Transfer event study

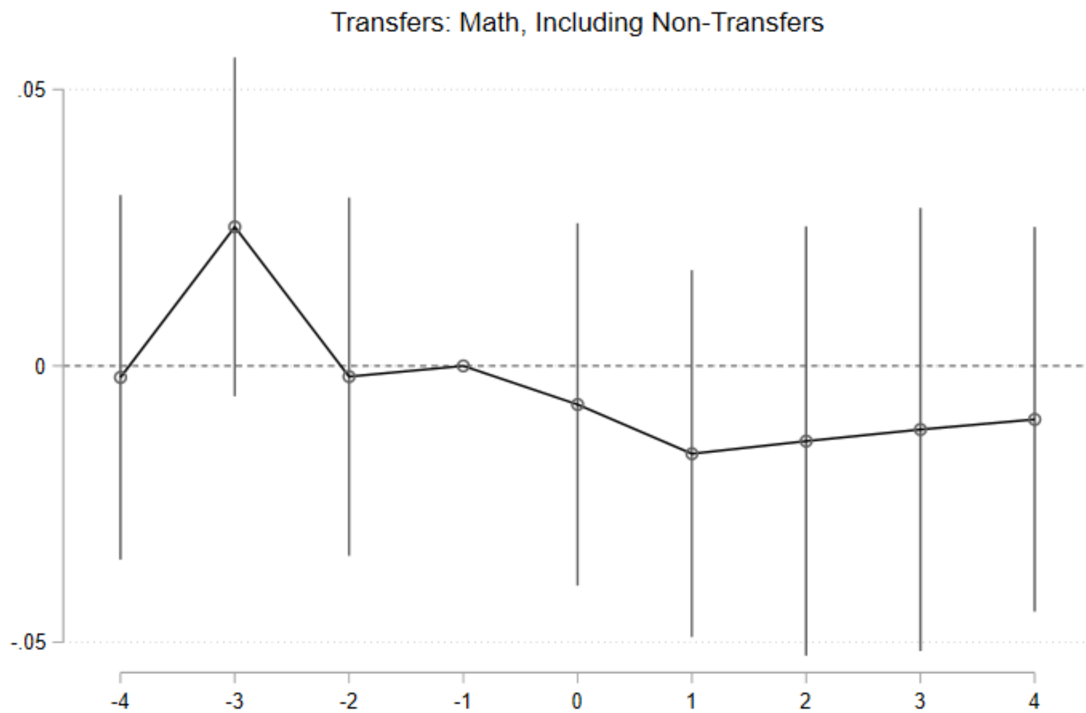

The figure shows event study coefficient estimates and 95% confidence intervals. The outcome is residualized math test score (residualized by student observables including lagged scores, school fixed effects, and an experience function), in student standard deviation units. The event is the teacher's first transfer from one school to another school in the state, where non-transfers do not have an event. We include teacher and year fixed effects and follow Sun and Abraham (2021) in constructing the estimates.

Figure A4: Forecast Unbiasedness for Large Changes in Class Size

(a) Large decreases in class size

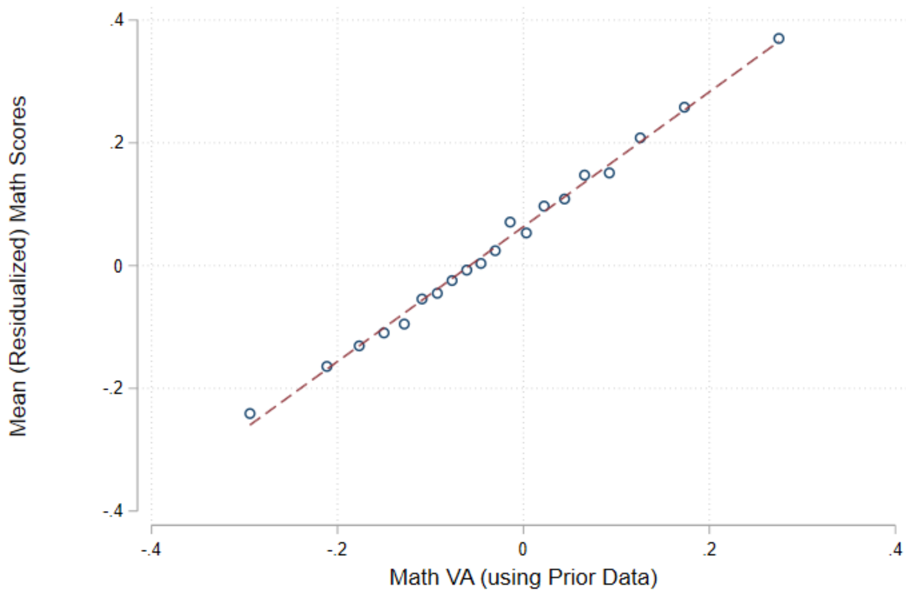

(b) Large increases in class size

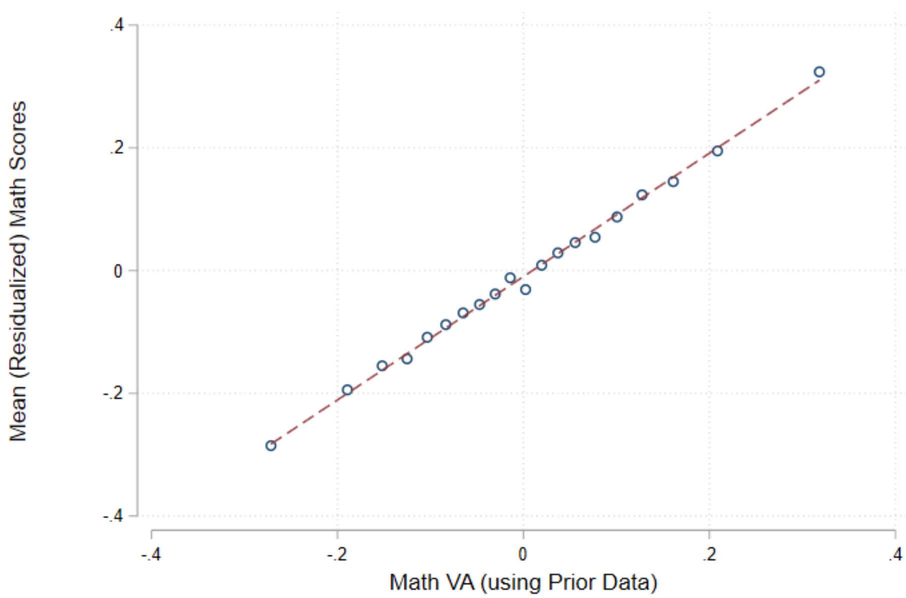

The figure shows a binscatter of student residual test scores by value-added prediction where an observation is a teacher-year. For decreases, the sample consists of all teachers where the class size used for prediction exceeds the class size in the target by more than 10 students. For increases, the sample consists of all teachers where the class size used for prediction is less than the class size in the target by more than 10 students.

Figure A5: Value-Added Distribution

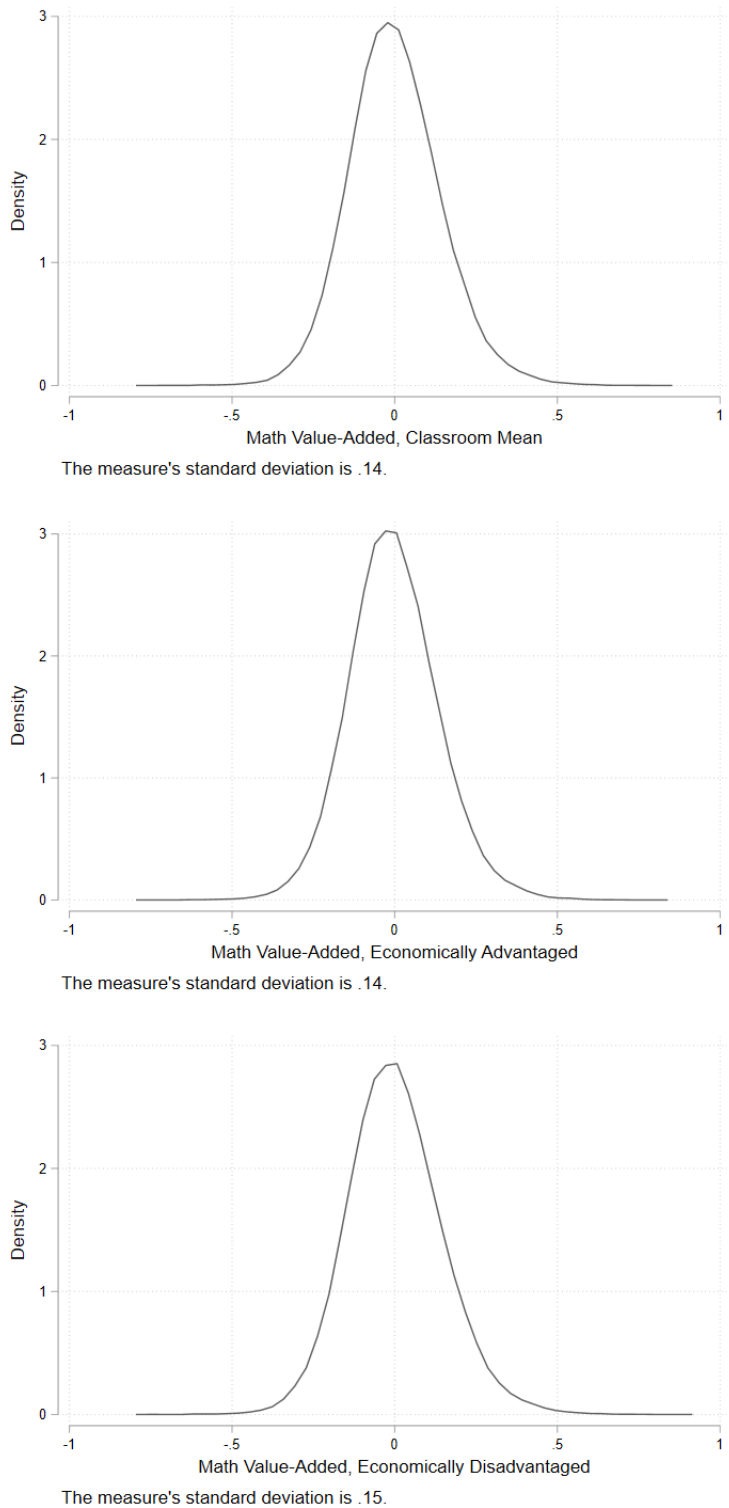

The figures show kernel density plots of our forecast of a teacher's value-added in a given year at the school they actually teach at (panel A), for economically advantaged students (panel B), and for economically disadvantaged students (panel C). The forecast uses only data from prior years. The units are student standard deviations.

Figure A6: Gains from teacher replacement

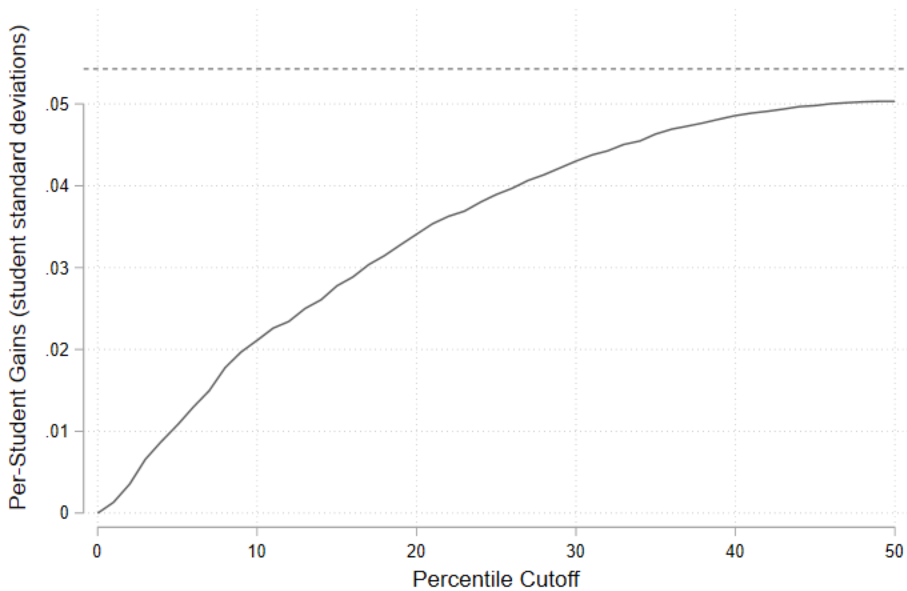

With class size variation

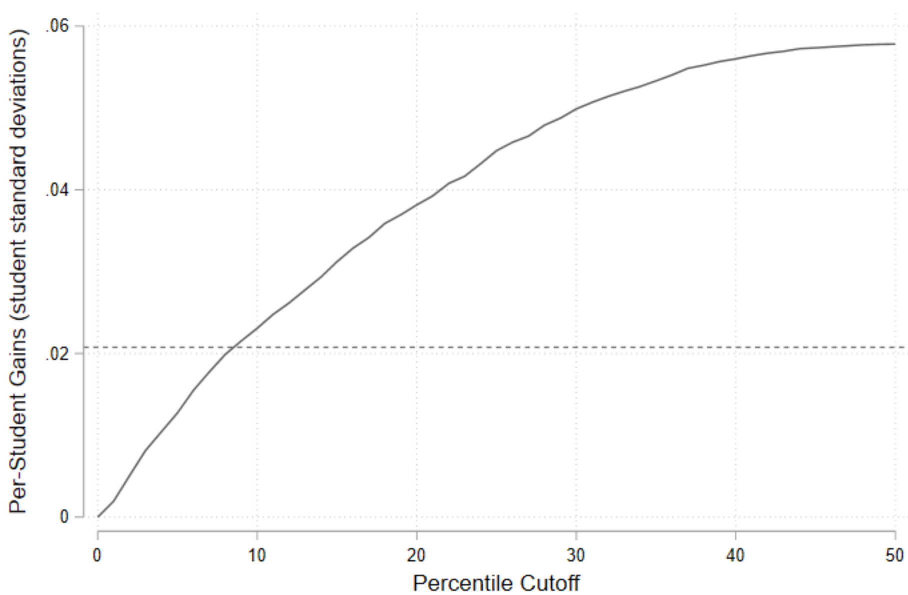

Constant class size

This Figure shows the results from policies that replace the X% of low-performing teachers with median value-added teachers, where the x-axis shows different values of X. The sample is the 2016 teachers with value-added forecasts. We assess performance based on realized value-added in the data (i.e., at the schools and classrooms a teacher is actually at in the data), and the median value-added teacher has median values for both dimensions of value-added. The y-axis is per-student gains in achievement. The top panel uses class size variation while the bottom panel imposes constant class sizes (at the district mean). The horizontal dashed lines are the gains from the output-maximizing allocation of existing teachers across schools in the district.

Figure A7: Student Gains by Fraction of Teachers Reassigned

The figure plots the potential per-student math test score gains (in student standard deviation units) as a function of the fraction of teachers that are assigned to a school different than their actual school. The sample consists of the 2016 teachers with math value-added scores.

Figure A8: Changes in a teacher's classroom composition and size between the output-maximizing and actual allocations

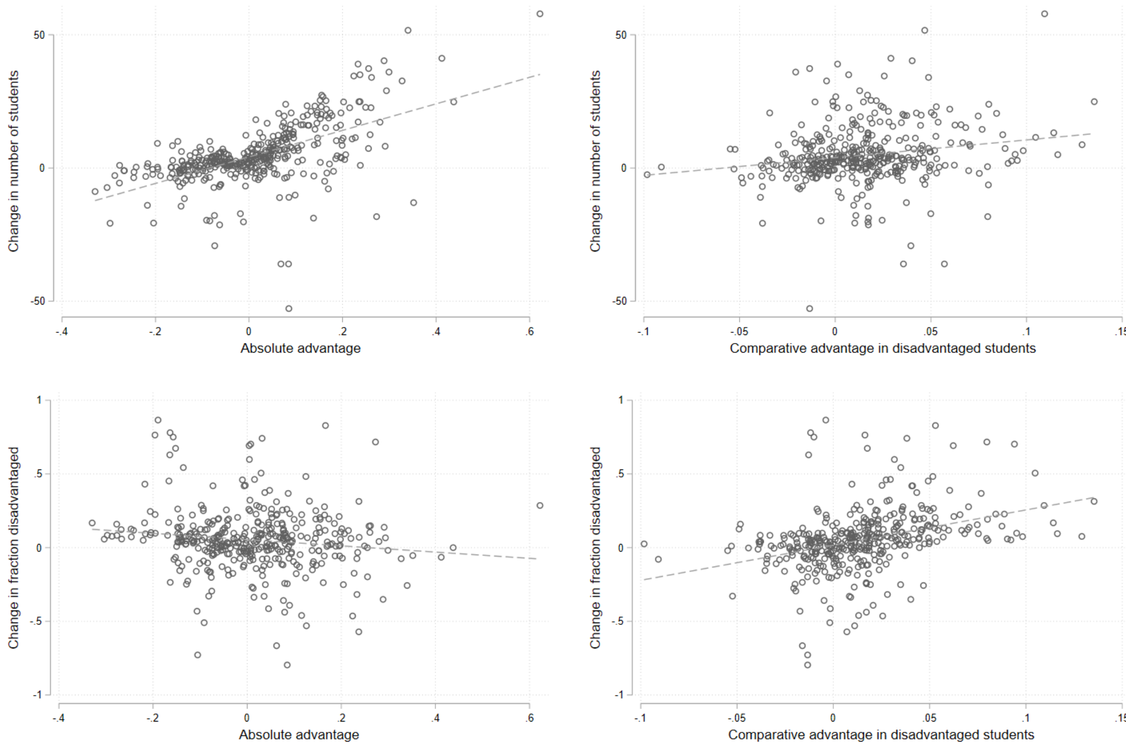

The figures show scatterplots and lines of best fit for the 2016 sample of teachers with value-added scores. In the top row, the variable of interest is the difference in the number of students a teacher teaches between the output-maximizing and actual allocations. Positive numbers are teachers who have more students in the output-maximizing allocation than in the actual. In the bottom row, the variable of interest is the difference in the fraction of disadvantaged students a teacher teaches between the output-maximizing and actual allocations. In the left column, teachers are ordered on the x-axis by absolute advantage (value-added at a representative school). In the right column, the teachers are sorted by comparative advantage in teaching economically disadvantaged students.

Figure A9: Optimal teacher placement relative to placement that generated value-added

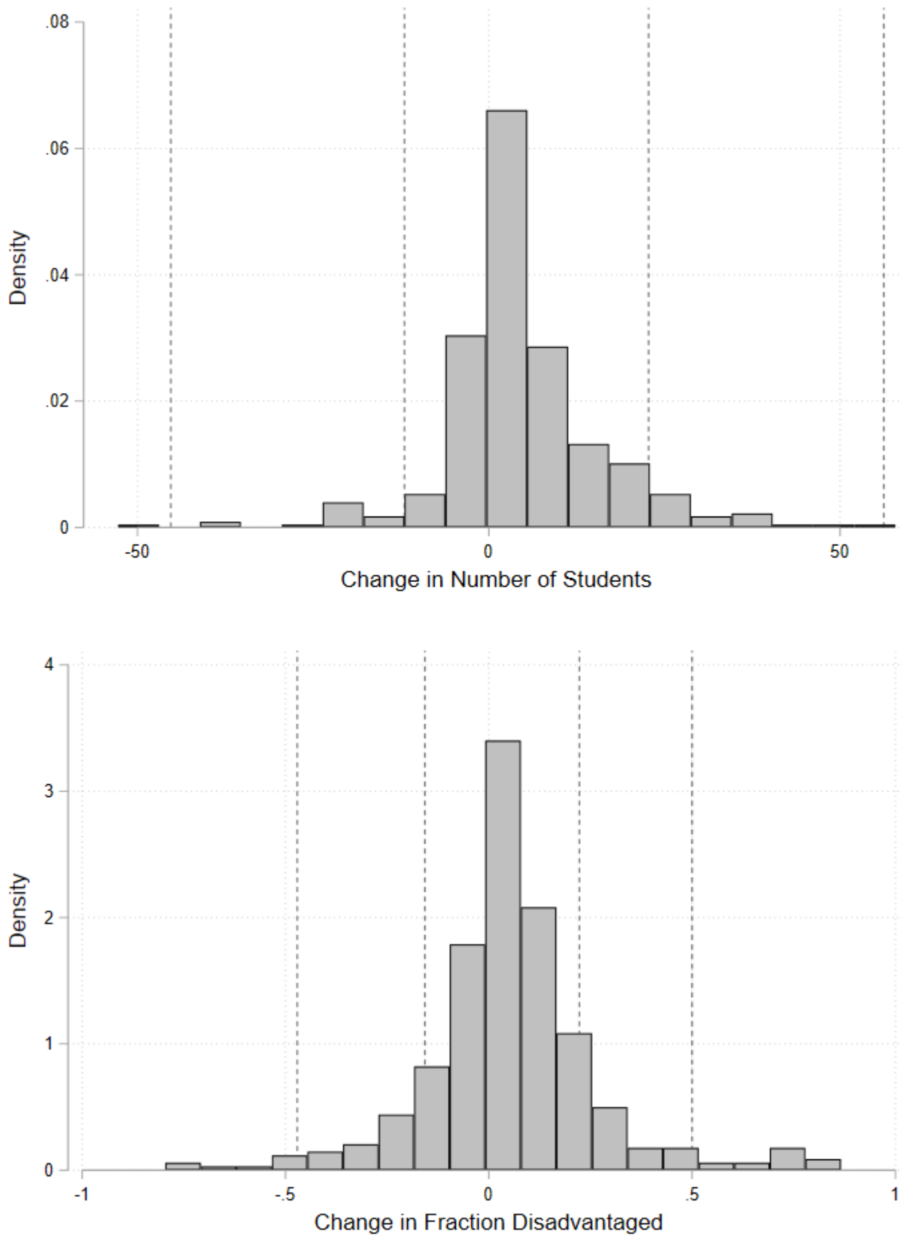

The figures show histograms for the 2016 sample of teachers with value-added scores. In the top panel, the variable of interest is the difference in the number of students a teacher teaches between the output-maximizing allocation and the classrooms that generated the teacher's value-added forecast. Positive numbers are teachers who have more students in the output-maximizing allocation than in the estimation data. In the bottom row, the variable of interest is the difference in the fraction of disadvantaged students a teacher teaches between the output-maximizing allocation and the classrooms that generated the teacher's value-added forecast. The vertical dashed lines represent the 1st, 10th, 90th, and 99th percentiles of the distribution we use for validation of our value-added measures in Table 1.

Figure A10: Simulations of two forms of misspecification

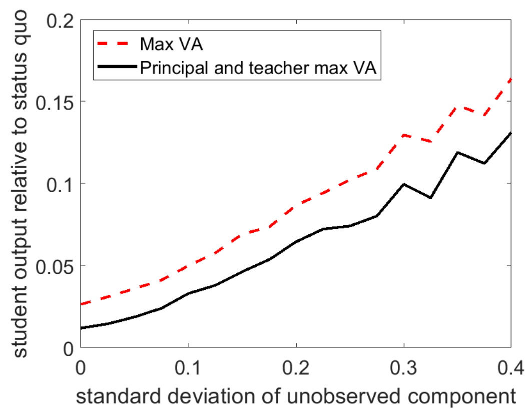

Unmodelled match effects

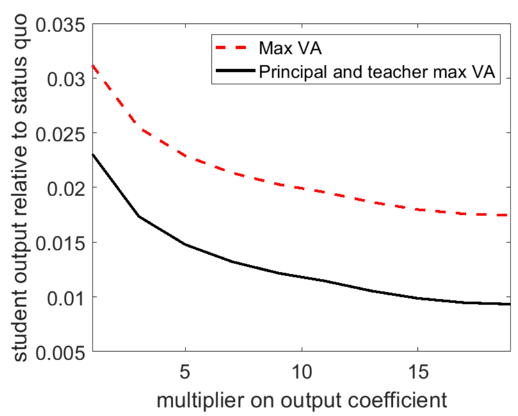

Attenuation in coefficient on output in teacher preferences

The figures show results from simulation exercises where we vary parameters related to match effects. In each figure the y-axis is the mean student achievement relative to the status quo. The dashed red line is the achievement in the output-maximizing allocation while the solid black line is the achievement in the equilibrium where principals and teachers each have preferences in order of value-added produced. The top panel adds an iid unobserved component to match effects, where the x-axis is the standard deviation of this component. The bottom panel varies the coefficient in teacher preferences on value-added. If our model misses match effects that teachers are aware of, then the preference coefficient might be attenuated. The x-axis in the bottom panel shows by how much we multiply our estimated coefficient on value-added.

Figure A11: Market timing

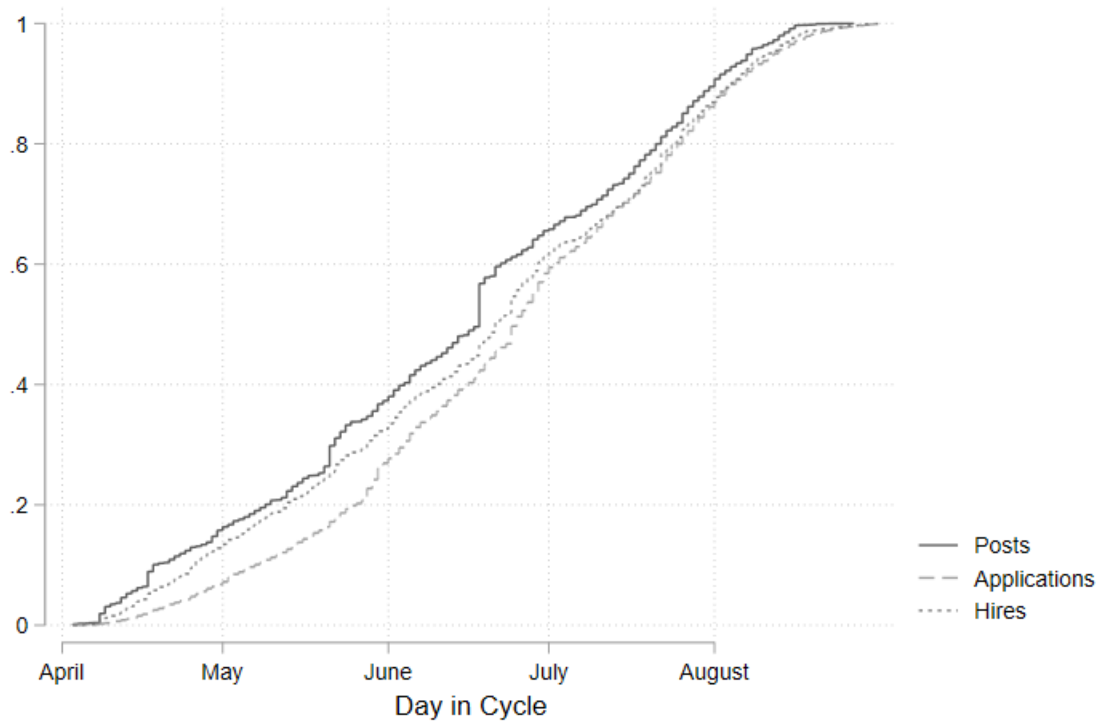

The figure shows CDFs for postings, applications, and hires (the application date of the application that led to a hire).

Figure A12: Application probabilities against commute time

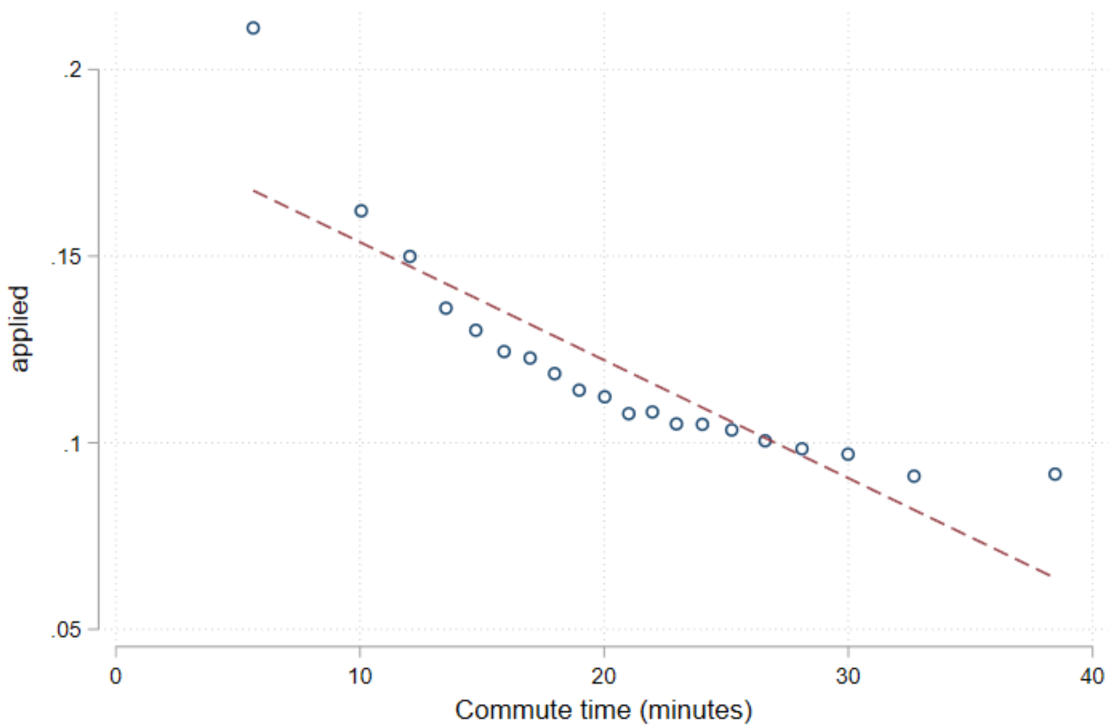

This Figure plots the probability of applying against commute time (measured in one-way minutes). The Figure residualizes for applicant fixed effects.

Figure A13: Number of teachers by fraction economically disadvantaged

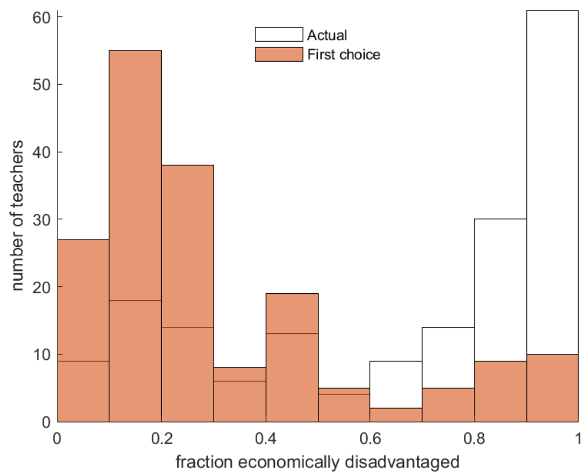

This Figure plots histograms of the number of teachers, by the fraction of students who are economically disadvantaged. The histograms are for the actual positions in the data (in white) and the positions teachers would have if they could all have their top choice (in red).

Figure A14: Model fit: teacher serial dictatorship based on absolute advantage (descending)

(a) Teacher absolute advantage

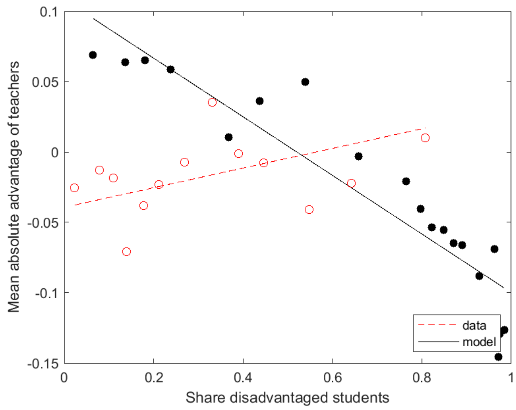

(b) Teacher comparative advantage

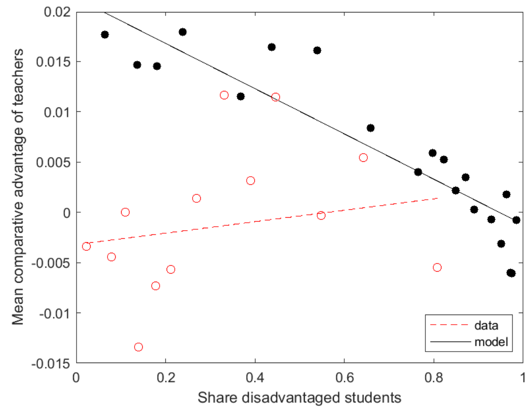

(c) Teachers with 7+ years of experience

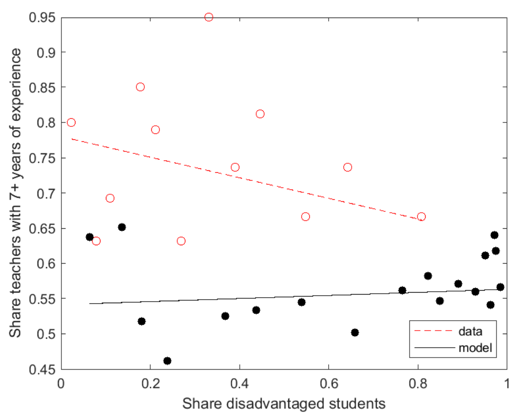

(d) Teachers that are Black

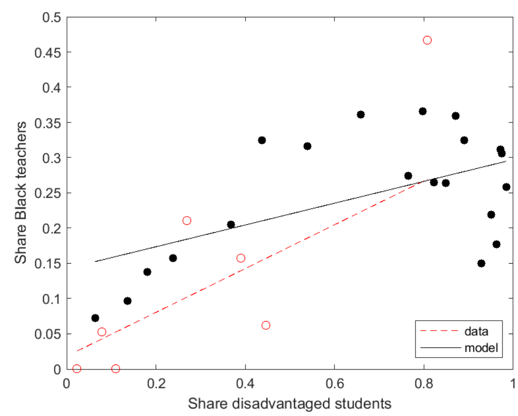

This Figure compares the allocations implied by a model in which the allocation is determined by a serial dictatorship where teachers go in descending order of their absolute advantage to the allocations we observe in the data. The data refers to all teachers in the district. The model refers to the teachers who apply in the transfer system for whom we have value-added scores.

Figure A15: Model fit: teacher serial dictatorship based on experience (descending)

(a) Teacher absolute advantage

(b) Teacher comparative advantage

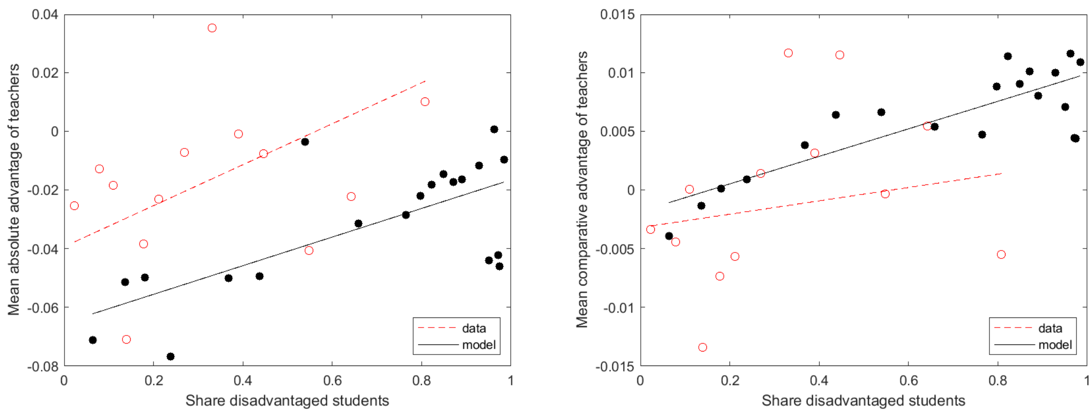

(c) Teachers with 7+ years of experience

(d) Teachers that are Black

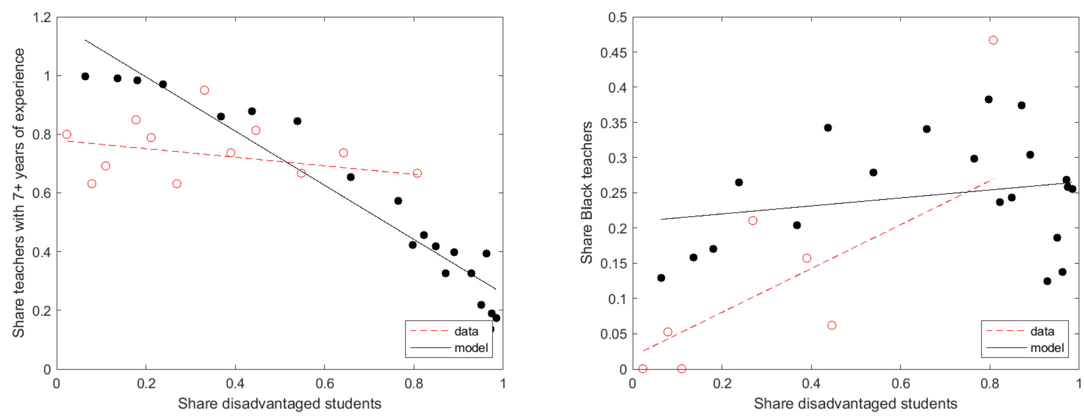

This Figure compares the allocations implied by a model in which the allocation is determined by a serial dictatorship where teachers go in descending order of their experience to the allocations we observe in the data. The data refers to all teachers in the district. The model refers to the teachers who apply in the transfer system for whom we have value-added scores.

Figure A16: Model fit: school serial dictatorship based on fraction disadvantaged (descending)

(a) Teacher absolute advantage

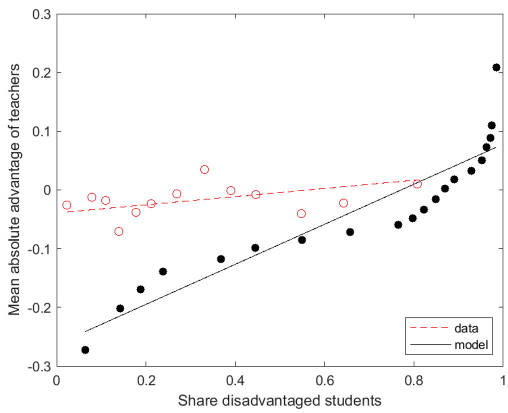

(b) Teacher comparative advantage

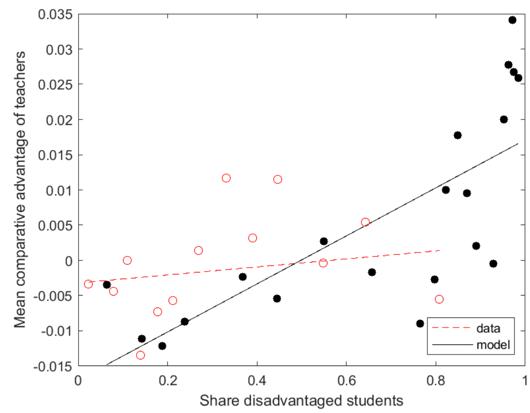

(c) Teachers with 7+ years of experience

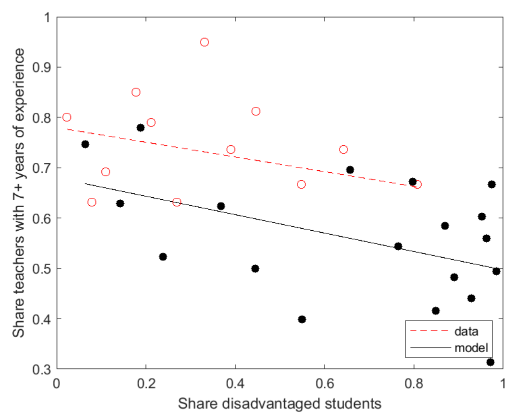

(d) Teachers that are Black

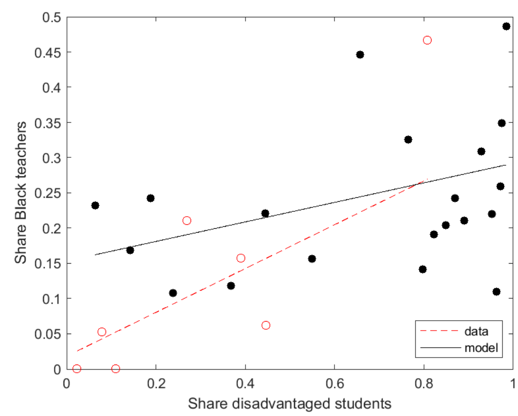

This Figure compares the allocations implied by a model in which the allocation is determined by a serial dictatorship where schools go in descending order of their fraction of disadvantaged students to the allocations we observe in the data. The data refers to all teachers in the district. The model refers to the teachers who apply in the transfer system for whom we have value-added scores.

Figure A17: Features of classes and teachers – transfer sample

(a) Class size and fraction disadvantaged

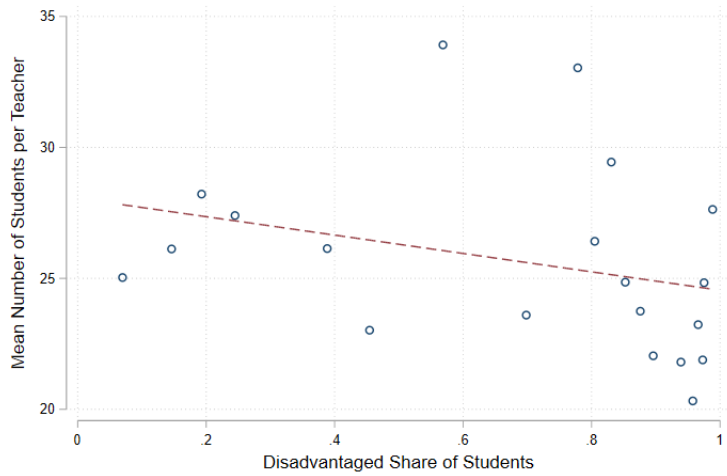

(b) Comparative advantage for disadvantaged students, and absolute advantage

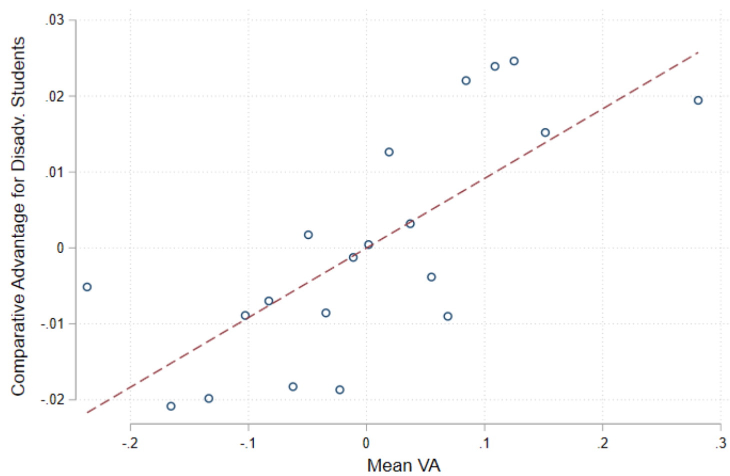

The figures show binscatters related to classroom characteristics and teacher characteristics in the transfer sample used for counterfactual analysis. The top panel shows the relationship between a school's (mean) disadvantaged share of students and a school's (mean) number of students per teacher. The bottom panel shows the relationship between a teacher's absolute advantage (x-axis) and comparative advantage in teaching economically disadvantaged students (y-axis). For this figure, absolute advantage is the average value-added across students types (rather than the value-added at a representative school) to avoid mechanical correlations between absolute and comparative advantage.

Figure A18: First-best allocation's placement of teachers, by absolute advantage

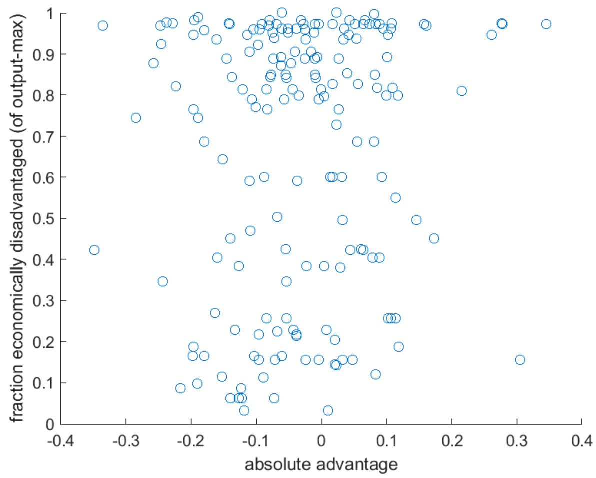

This Figure plots the first-best allocation in our transfer sample, where we divide teachers by absolute advantage and positions by fraction of students that are economically advantaged. Each point is an assignment of a teacher to a position.

Figure A19: Postings selection in the transfer market

(a) Positions and fraction of disadvantaged students

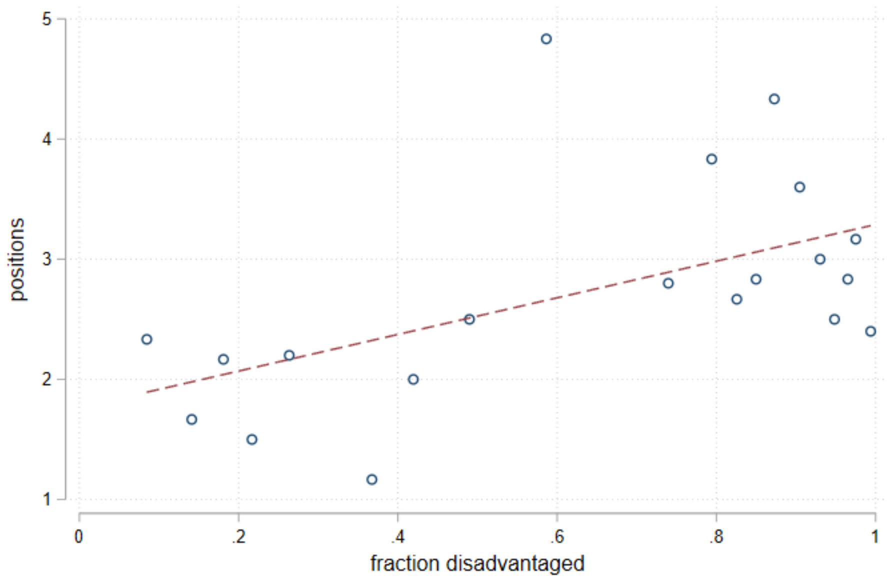

(b) Positions and class size

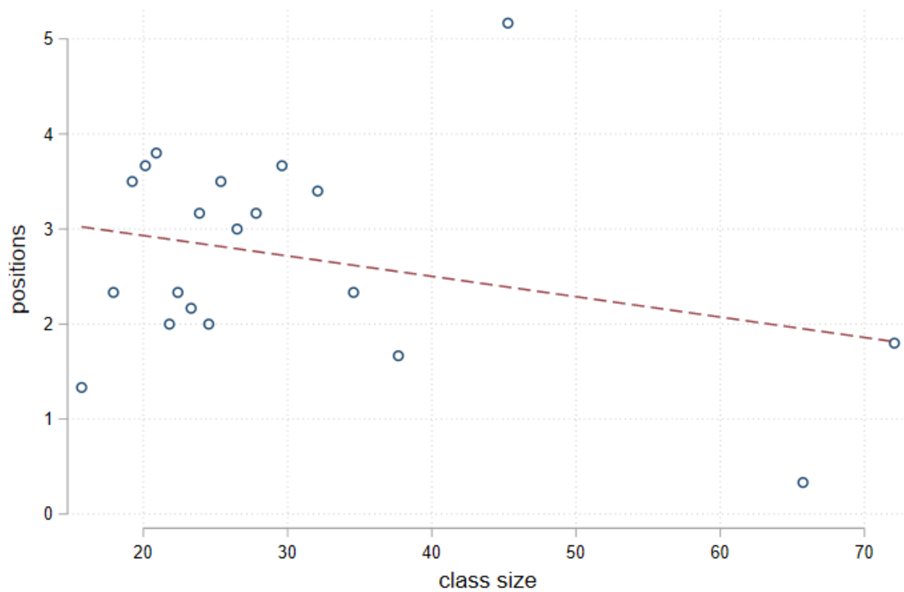

This figure shows the relationship between number of positions posted and (a) a school's fraction of students that are economically disadvantaged and (b) a school's class size. An observation is a school.

Figure A20: Title I Application Gap

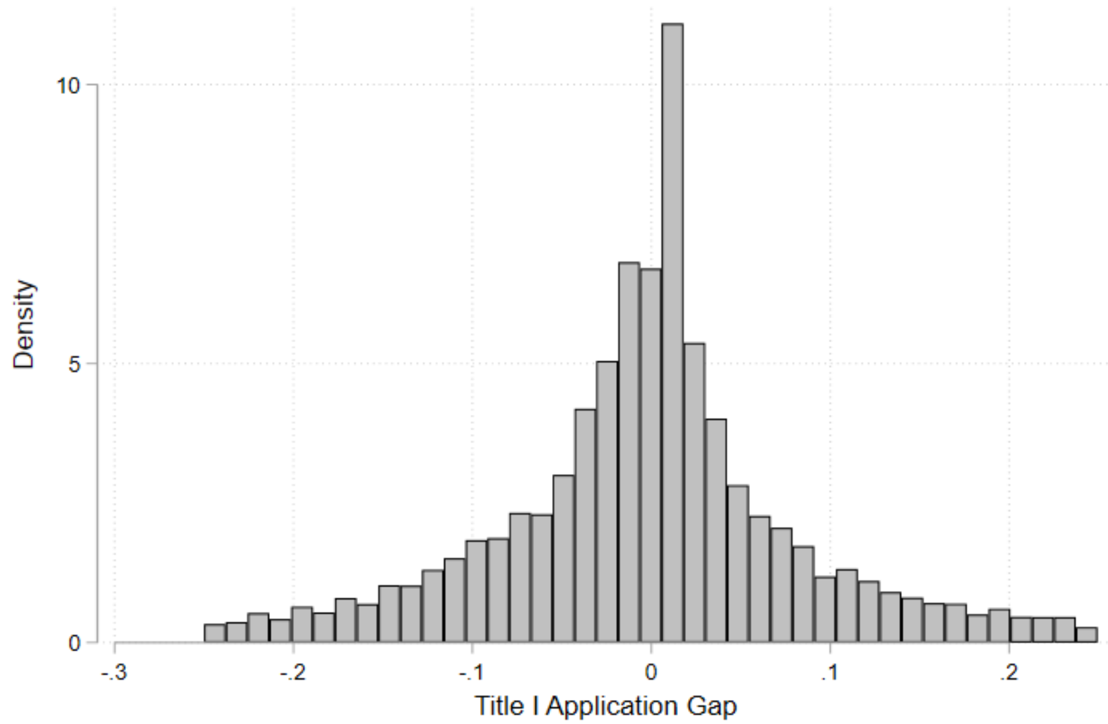

This Figure plots the distribution of the individual-level Title I application rate minus the individual-level non-Title application rate. Thus, the positive entries indicate that a teacher applies to a greater share of the Title I schools in their choice set than to the non-Title I schools.