Title

Teacher Labor Market Policy and the Theory of the Second Best

Authors

Bates, Michael
Dinerstein, Michael
Johnston, Andrew C
et al.

Publication Date

2025

DOI

10.1093/qje/qjae042

TEACHER LABOR MARKET POLICY AND THE THEORY OF THE SECOND BEST*

Michael Bates

Michael Dinerstein

Andrew C. Johnston

Isaac Sorkin

Abstract

We estimate a matching model of teachers and elementary schools with rich data on teacher applications and principal ratings from a large, urban district in North Carolina. Both teachers' and principals' preferences deviate from those that would maximize the achievement of economically disadvantaged students: teachers prefer schools with fewer disadvantaged students, and principal ratings are weakly related to teacher effectiveness. In equilibrium, these two deviations combine to produce a surprisingly equitable current allocation, where teacher quality is balanced across advantaged and disadvantaged students. To close achievement gaps, policies that address deviations on one side alone are ineffective or harmful, while policies that address both could substantially increase the achievement of disadvantaged students.

JEL: I24, J23, J45.

*We thank Nikhil Agarwal, Claudia Allende, Peter Arcidiacono, Natalie Bau, Christina Brown, Jediphi Cabal, Dennis Epple, Felipe Goncalves, Caroline Hoxby, Peter Hull, Adam Kapor, Lawrence Katz, Neale Mahoney, Richard Mansfield, Maxim Massenkoff, Derek Neal, Parag Pathak, Luigi Pistaferri, Jesse Shapiro, Camille Terrier, and numerous seminar and conference participants for helpful comments and conversations. Thanks to Ian Calaway for research assistance. We also thank representatives at an unnamed district, Kara Bonneau, and the North Carolina Education Research Data Center for data merging and access. Bates thanks the Blum Initiative and Sorkin thanks the Alfred P. Sloan Foundation for support. This paper subsumes drafts previously circulated under the titles "Teacher Labor Market Equilibrium and Student Achievement" and "Teacher Labor Market Equilibrium and the Distribution of Student Achievement." This research was approved by the Stanford IRB (#43865), the UC Merced IRB (UCM2019-88), and the UC Riverside IRB (HS-16-094). It was deemed non-human subjects research by the University of Chicago IRB (IRB17-1651). Mistakes are our own. Correspondence: Michael Dinerstein, Duke University, Department of Economics, 419 Chapel Dr, Durham, NC. E-mail: michael.dinerstein@duke.edu.

Achievement gaps between economically advantaged and disadvantaged children are large (Duncan and Magnuson, 2011). These achievement differences in childhood predict significant differences in adult wages (Neal and Johnson, 1996). Recent work argues that the relationship between achievement and earnings is causal—teachers that causally improve achievement also improve adult income (Chetty et al., 2014b). Because teachers vary widely in their effect on student achievement (Hanushek et al., 2004), adjusting the *allocation* of teachers to students may be one of the most potent policy instruments for helping the disadvantaged.

In this paper, we study the allocation of teachers to schools and assess the desirability of various policies. We combine rich data with the economics of two-sided markets to understand the current allocation and explore implications for policies.[1] We combine three empirical findings, each the subject of a large literature. First, within a district, teacher quality is balanced across disadvantaged and advantaged students. Second, teachers prefer to teach at schools with more advantaged students. Third, in hiring, principals mostly do not select high value-added teachers.

Conceptually, there is a tension between finding a balanced allocation and systematic teacher preferences for schools with advantaged students. We typically expect the entity that faces excess supply to have either lower prices or, if prices are restricted as in teacher labor markets, higher quality. In our setting, that intuition would imply that the schools with advantaged students would have better teachers. The third empirical finding offers a simple reconciliation: because principals do not hire (very much) on the basis of the value-added, they do not take advantage of the excess supply to hire better teachers.

This reconciliation is not just helpful in explaining the current allocation, it also provides useful guidance for policy. These results reflect "theory of the second best" (Lipsey and Lancaster (1956)) logic: when there are multiple deviations from policies that implement the first best, fixing any one can be ineffective or harmful. Suppose the planner wants to maximize the achievement of disadvantaged students. The district would approximately achieve the first best if teachers preferred teaching disadvantaged students and principals hired the best teachers. But when teachers' preferences and principal hiring both deviate from these benchmarks, policies that only target one side at a time may be ineffective or even harmful.

We empirically show that providing bonuses to teachers for teaching disadvantaged students as suggested by, e.g., Clotfelter et al. (2011) and Goldhaber et al. (2018), does not improve the quality of the teachers of disadvantaged students. The reason is simple: while bonuses increase the supply of teachers to schools with disadvantaged students, the principals do not take advantage of excess supply to hire better teachers. Similarly, providing bonuses or incentives to principals to hire based on value-added as suggested

[1] In the US, 92% of large districts now operate two-sided markets where both the teacher and principal must agree to the match (National Council on Teacher Quality, 2022).

by, e.g., Ballou (1996) and Jacob et al. (2018a), backfires in that now the schools with advantaged students use the excess supply to hire better teachers, generating the unequal outcomes we expect based on the teacher preferences. Implementing these two types of policies in tandem, however, achieves an allocation close to the first best.

We begin with a model of the teacher labor market that enables us to specify a first-best benchmark and how the market clears. Teachers apply to vacancies and principals hire among applicants. We assume that the equilibrium allocations are pairwise stable.[2] Given an objective of maximizing the achievement of disadvantaged students, the model provides a policy benchmark: the first-best allocation can be approximately achieved if teachers prefer positions with the most disadvantaged students and principals prefer to hire the most effective teachers.[3]

To estimate the model, we use detailed data from a large urban school district in North Carolina and we focus on the market for elementary school teachers. To evaluate the current allocation, we observe teachers linked to their students' yearly test scores. To understand teachers' preferences, we observe the full set and timing of job applications that teachers submit as well as the timing of job postings. To understand principals' behavior, we see the full set of applications the principal receives, notes the principal records about applications, interviews, and offers.

We first specify our empirical model of how teachers affect student math achievement. For our baseline, we allow teachers to have different value-added with disadvantaged and advantaged students (Condie et al., 2014; Delgado, 2023). But all of the results of the paper hold with a variety of more conventional homogeneous value-added models.

To identify teacher and principal preferences, we rely on relatively weak assumptions that allow for transparent identification. We focus on actions that are early in the process: the teacher decision to apply and the principal rating of applicants, rather than, say, the teacher decision to accept an offer or the principal decision to make an offer. This focus allows for strategy later on in the process. Based on institutional features and extensive analysis of applicant behavior, we argue that teachers apply non-strategically to vacancies when they are active, and principal rating behavior is non-strategic.

Our estimates of teacher and principal preferences broadly reproduce patterns in the literature that imply that teacher and principal preferences do not implement the first-best allocation. Relative to the literature, we have several advantages: we have actual choices, we observe choice sets, and we allow for preference heterogeneity. Consistent with the literature (e.g., Greenberg and McCall, 1974; Antos and Rosen, 1975;

[2] Our model fits in a recent literature considering allocation problems with non-choice outcomes (Agarwal et al., 2020; Ba et al., 2021; Cowgill et al., 2024; Dahlstrand, 2024).

[3] The statement is approximate because we allow for timing restrictions and comparative advantage.

Johnston, 2024), our key teacher finding is that teachers prefer schools with fewer disadvantaged students (though we find significant heterogeneity) so that teacher preferences differ from that would implement the first-best allocation in a market with uniform pay. Also consistent with the literature (e.g., Ballou, 1996; Boyd et al., 2011; Jacob et al., 2018b), our key principal finding is that the principal's preferred candidate is rarely the one that is most effective at raising student test scores. We can also reject the vertical preference model typically assumed in settings where choice sets are not observed (see Diamond and Agarwal, 2017). The lack of weight on test scores could reflect a lack of information or a lack of incentives, but regardless, principal behavior differs from that which implements the first-best allocation.

For our first main result, we show that advantaged and disadvantaged students have teachers of approximately equal strength. This pattern is present in raw test score gains, for a wide variety of value-added models, and for behavioral value-added. One notable feature of the current allocation is that disadvantaged students are more likely to have novice teachers for whom we cannot estimate value-added in our main models. We find a similar pattern when we rely either on a residual value-added estimator that just uses contemporaneous data and so also includes novice teachers, or when imputing value-added using observable characteristics.[4]

To understand why advantaged and disadvantaged students have teachers of approximately equal strength, we combine our estimates with the two-sided matching model. We show, consistent with the theory of the second best, that this result reflects multiple deviations combining to produce favorable allocations for disadvantaged students. If teachers only care about the number of disadvantaged students, then the allocation is little changed. The intuition is that because principals place little weight on value-added, they do not select more effective teachers from the larger applicant pool. Similarly, if principals only placed weight on value-added, then the outcome would be worse for disadvantaged students.

While these results explain the current allocation, they also inform the design of policies that the teacher preferences and principal hiring literatures recommend. Specifically, the basic message follows the theory of the second best: policies that only target one side of the market can be ineffective or even harmful.

While the current allocation is surprisingly balanced, it does not achieve the first-best allocation that maximizes the outcomes of disadvantaged students. We find that the first best would provide substantive gains: in a single year, reallocation could close one-fourteenth of the baseline achievement gap, while increasing average achievement. When extrapolating linearly, the achievement gap could largely be closed over twelve years of public education. Implementing extreme versions of teacher and principal policies jointly nearly achieves the first-best allocation.

[4]Papers that study differences in teacher value-added by level of student disadvantage tend to find little to no difference (e.g., Chetty et al. (2014b), Mansfield (2015), and Isenberg et al. (2022)). See Appendix I for a more detailed discussion of this literature.

We consider a number of extensions. We first show our estimates of the status quo are nearly identical when clearing the market at once or across multiple sub-periods. Most importantly, one argument for making principals more likely to hire high value-added teachers is that this affects the extensive margin of who works in the district. We therefore study an extension where we include teachers for whom we have to impute value-added. Making principals better at hiring does bring better teachers to the district, but given the structure of teacher preferences, advantaged (not disadvantaged) students largely benefit from the better teachers. Thus, the basic theory of the second best message persists even with an active extensive margin.

We also study more realistic teacher bonuses. We first consider one-sided bonuses that provide incentives for teachers to teach at schools with disadvantaged students. Such bonuses are only weakly effective because they do not affect how principals hire from the pool. If principals hire to maximize value-added (via some combination of information and incentives), then the teacher bonuses are effective.

To summarize, the unifying theme of this paper is the theory of the second best. Subsidizing one side of the market at a time can be ineffective or counterproductive, even when subsidizing both sides is beneficial, and the current allocation is balanced even though teachers' preferences suggest it would not be. Our results challenge the conclusions from the prior one-sided literatures by explaining an otherwise puzzling feature of the current allocation and reaching opposite conclusions about policy effectiveness. Reaching these conclusions requires rich data on the actions of both sides of the market.

This paper fits in a growing literature on equilibrium models of the teacher labor market.[5] These papers tend to fall into two camps. In the first, the hiring side of the market faces constraints imposed by the government (e.g., they must hire the most experienced applicant) such that the market is essentially one-sided (Bobba et al., 2024; Combe et al., 2022b; Elacqua et al., 2021; Tincani, 2021; Combe et al., 2022a).[6] We instead focus on two-sided labor markets, which characterize nearly all teacher labor markets in the US and the hiring of permanent teachers in many non-US settings. In the second camp, several papers study two-sided markets but infer preferences from data on equilibrium allocations (Boyd et al., 2013; Bates, 2020; Biasi et al., 2021). We instead observe the actions of each side of the market, which allows us to relax the strong assumptions necessary for identification in the absence of such data. We show that these assumptions deliver misleading conclusions about the relationship between teacher quality and student disadvantage in equilibrium as well as the desirability of commonly-suggested policies. Like us, Davis (2022), Ederer (2023), and Laverde et al. (2023) study two-sided markets with data on each side's actions. Unlike the

[5]This paper also relates to the industrial organization literature on information in matching markets which finds that incomplete information can limit student gains from being strategic (Kapor et al., 2020) and can lead to costly search (Chen and He, 2021; Arteaga et al., 2022).

[6]Bau (2022) studies an equilibrium model of school competition with school-student match effects.

first two papers, we estimate teacher quality based on student test scores instead of relying on observable teacher characteristics. We find that restricting teacher quality to vary only with observable characteristics changes the assessment of equilibrium and policy conclusions; for example, we find that the allocation is not balanced across advantaged and disadvantaged students in terms of teacher observables (like experience) despite parity on multiple direct measures of effectiveness. This finding comes from our detailed data linking teachers to students, and their test scores, at a finer level than many papers in the literature.[7]

Our study carries important lessons for the analysis of labor markets. Much of the labor literature, on topics such as wage inequality (e.g., Card et al., 2018) and amenities (e.g., Sorkin, 2018), relies on matched employer-employee data where researchers only observe equilibrium allocations. These markets are two-sided, which forces researchers to rely on the same identifying assumptions that led to misleading conclusions in the teachers literature. Our findings thus reinforce Oyer and Schaefer (2011)'s call for labor economists to study how firms hire workers and Card et al. (2018)'s suggestion that the labor literature on imperfect competition would benefit from "IO-style" case studies of particular markets.

I An equilibrium model of the teacher labor market

Here, we write down an equilibrium model of the within-district teacher labor market. The model clarifies the set of factors shaping the equilibrium, allows us to define the first-best allocation, and explains when the decentralized equilibrium attains the first-best allocation.

I.A Set-up

Teacher j derives utility u_{jk} from teaching at school k. School k's principal derives utility, v_{jk}, from hiring teacher j. Utility is non-transferable, as wages are set by the district and do not vary across assignments for a given teacher.

A teacher-school assignment produces value-added VA_{jk}. Because we are interested in the achievement of disadvantaged and advantaged students, we allow the value-added to depend on the student type. Specifically, let μ_{jm} be teacher j's value-added with students of type m, where $m \in \{0,1\}$ indicates whether a student is disadvantaged. Let n_{km} be the number of students in school k of type m. Then:

$$VA_{jk} = n_{k0}\mu_{j0} + n_{k1}\mu_{j1}. \tag{1}$$

[7] Laverde et al. (2023) link teachers to school-grades.

Finally, let $\mathcal{J}$ be the set of teachers, $\mathcal{K}$ be the set of schools, and assume for simplicity that the number of teachers and schools is the same. An assignment of teachers to classrooms is a one-to-one and onto function (bijection): $\phi : \mathcal{J} \rightarrow \mathcal{K}$ so that $\phi(j) = k$, the school k to which teacher j is assigned.[8] Denote by Φ the set of all possible assignments.

I.B First-best allocation

We are interested in policies that increase the achievement of disadvantaged students. We take as given the set of teachers and positions the district has and ask how to assign them. In Section VIII, we consider the set of teachers who apply in the transfer system and for whom we can estimate value-added: this set includes teachers who have previously taught anywhere in the state.

The district values the achievement of disadvantaged students:[9]

$$\max_{\phi \in \Phi} \left\{ \sum_{j \in \mathcal{J}} n_{k1} \mu_{j1} \right\}. \tag{2}$$

The structure of the first-best allocation is simple: rank teachers in descending order by value-added with disadvantaged students and rank classrooms in descending order by the number of disadvantaged students. Then assign the strongest teacher to the classroom with the largest number of disadvantaged students and so on.[10]

Because the paper's goal is to study the allocation of teachers, and not how best to use existing dollars, we do not include a budget constraint in the district's problem. As cost is still a relevant consideration in evaluating allocations, in Section IX we compare the effectiveness of policies that cost equal amounts.

I.C Decentralized equilibrium

Our equilibrium concept is (timing-constrained) pair-wise stability. Schools meet with all teachers who are in the market at the same time. Under a stable allocation, no teacher and school pair would prefer to jointly deviate and match (Roth and Sotomayor (1992), Definition 2.3). Stability is a natural assumption in decentralized markets as it says that pairwise gains from trade have been exhausted (i.e., the set of stable allocations are the core).

To model the empirical status quo, we assume (1) teachers and principals have the preferences we estimate for them and (2) the timing of the market follows that which we observed in the administrative

[8] For notational simplicity, each school has a single position. In the empirical model, schools may have multiple positions.

[9] In Appendix III, we include advantaged students' achievement.

[10] Table V (Part 1) shows that our results are very similar if we hold class sizes constant.

records, where not all matches are feasible. There is not necessarily a unique stable equilibrium. We model the status quo using the teacher-proposing deferred-acceptance algorithm (DA), which we use to find a stable equilibrium, not because DA is actually used.

When does the decentralized equilibrium correspond to the planner problem? Suppose that teachers rank schools according to the number of disadvantaged students ($u_{jk} \propto n_{k1} \forall j,k$) and principals rank teachers according to total output ($v_{jk} \propto VA_{jk} \forall j,k$). Then in the absence of comparative advantage or timing restrictions, the decentralized equilibrium—which is unique in this case—corresponds to the planner's solution. Notably, this combination of rankings is what the *joint* implementation of hard-to-staff school bonuses and guided principal hiring would achieve. Of course, the theory of the second best says that aligning only the teacher or the principal with the planner may not improve outcomes.

Empirically, we are then interested in the extent to which teacher and principal preferences align with those that decentralize the planner's solution. We are also interested in whether the other factors we have abstracted from—timing and comparative advantage—affect the gap between the decentralized equilibrium and the planner's solution.

II Data and institutional context

We use rich data on the labor market for elementary school teachers. Elementary schools are grades K to 5 (or sometimes 6). For the purpose of estimating the variance of classroom effects, we also use data from middle schools (grades 6 to 8) where teachers are more likely to teach multiple classrooms. The first type of data comes from the platform used to hire teachers in our focal district. We use this data to estimate teacher and principal preferences. The second type of data comes from staffing and achievement records from state accountability records. This data provides us with student-level test score data that we link to teachers and use to estimate value-added models. In addition, these records provide information about a variety of demographic characteristics of teachers and students as well as teachers' education and experience in the district. In this section, we briefly describe the data. See Appendix II for further details and Appendix Table A1 for summary statistics across samples.

II.A Job application and vacancy data

We obtained application records from our focal district's system, which spans 2010 through 2019 and records 346,663 job applications. In the system, schools post job vacancies, and applicants apply for jobs. The system also records various actions that principals take.

For every posted position, the vacancy files indicate the school, position title, and whether the position is full-time or part-time. We use the detail on the position title to isolate non-specialized elementary school teacher jobs (i.e., we omit elementary school jobs such as "literary facilitator elementary").

We use two features of the teacher file. First, the file records which vacancies the candidate applied to, and when she submitted the application. The timing information allows us to construct choice sets, which we detail in Section III. Second, the file records the city, zip code, and address where the teacher lives. This feature allows us to construct the commute time for each teacher-position combination.

We also have data in which principals record their assessments of teachers. Principals record their interest in different applicants, the equivalent of a "good" and a "bad" pile. Principals also record which candidates they invited to interview, which candidates were offered the position, and which candidates were hired.

II.B Administrative data

We link the platform data to state administrative records on teachers and students. For teachers, we have their experience, salary, licensing, certification status, test scores, class assignments, and the school where they work. For students, we have scores on standardized exams, grades, race, sex, and whether they qualify as disadvantaged based on Federal programs. Records on class assignments allow us to link teachers to students.

The North Carolina Education Research Data Center (NCERDC) matched the data from the job-market platform to the state's administrative data, using names, birth dates, and the last four digits of teachers' social security numbers. 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. Appendix Table A2 shows the share of newly hired teachers in the district that we find in our job market platform data. The lowest rate is 94% and in our focal year it is within rounding error of 100%.

The data show that student types vary considerably across teachers, which is driven by the sorting of students across schools. Appendix Figure A1 plots the fraction of a teacher's students that are economically disadvantaged. Almost a third of teachers have classrooms with almost entirely economically disadvantaged students. In Appendix Table A3, we show that this pattern reflects sorting of students across schools rather than across classrooms within schools. Specifically, the adjusted-R^2 of a regression predicting disadvantaged students is 0.4 using either school or classroom dummies. Similarly, the peer share of disadvantaged students that are disadvantaged is around 70% when using schools or classrooms. Given such student sorting across schools, different allocations of teachers to schools have the potential to yield very different learning gains for disadvantaged students.

II.C Market overview

Our district organizes a decentralized hiring and transfer process in which teachers choose where to apply and principals choose whom to hire. External and internal (transfer) applicants are pooled into one market. Here we describe the basic market structure.

Market organization: The school district runs a centralized online hiring platform, where each school posts openings. Teachers choose whether to apply to each posting.

Timing: We examine the "on-cycle" part of the market, which dictates hiring and transfers between school years. It begins in the winter, when the district notifies each school of known and expected attrition among the school's work force and of how many positions that school may hire. It ideally ends with filled positions by late August before the new school year. Similar to what Papay and Kraft (2016) find, some schools are unable to fill all positions by the start of the new school year.

Postings: The number of postings at a school reflects a combination of enrollment, budget, and the number of teachers who leave. All three pieces of information are not necessarily known before the main hiring season starts. This information delay generates variation within and across schools in the timing of postings. For example, late information about enrollment or budget fluctuations often necessitates late posting. Or if there is mid-year attrition, then the school would know long before hiring season started that there would be a vacancy, which allows for early posting.

Applications: An application consists of a variety of documents, including teacher certification and a brief diversity statement. The same set of documents applies to all positions. Thus, a prospective teacher faces a fixed cost of preparing materials but little marginal cost to apply to an additional posting.

Evaluation and hiring: When a teacher applies to a position, the hiring school receives her application materials through the platform. For teachers who previously taught in the district, principals may request or teachers may disclose a district-calculated measure of value-added.[11] The school's principal may then rate

[11] Raw growth scores have been available in North Carolina since 1997. In 2013, the district also began using Education Value-Added Assessment System (EVAAS) measures of value added. All teachers with such measures have an opportunity to reveal these evaluations with their initial application though some may not choose to do so.

the applications and choose to interview applicants on a rolling basis. For known positions at the beginning of the hiring period, there is a short window during which only transfers from within the district are able to apply. Schools can either hire from this pool or wait and consider more applicants.

If the principal wants to hire the candidate, she extends a job offer. The candidate has 24 hours to accept the offer, during which she might contact other schools that have shown interest. If the teacher accepts the offer, she commits to not accepting an alternate offer in the same cycle.

With a few small exceptions, teacher pay is determined by a mechanical formula that depends on degrees, certifications, and experience. These costs are borne by the district, so hiring a more experienced or credentialed teacher does not cost the principal more.

Eligibility: Teachers are eligible for positions if they have the necessary certification. We focus on the market for elementary-school classroom teachers because the common certification allows us to reliably classify which teachers are eligible. We can also infer elementary school teachers' quality from systematic gains in their students' test scores because teachers in these positions are typically responsible for instruction in the tested subjects.

III The vacancy posting and application process

In this section, we describe our model of teacher and principal actions in the labor market. We specify our model assumptions, consider how violations of the assumptions might manifest in the data, and show empirical evidence consistent with the assumptions. Our empirical analysis includes robustness checks around possible alternate assumptions. We defer a discussion of the pair-wise stability assumption until Section VIII.

Our models of teacher preferences and principal behavior assume that there is an action that reflects preference orderings, which leads to transparent arguments for how we identify preferences. There are several actions that teachers and principals take in order to form a match. A teacher decides to apply, the principal views the application and assigns a rating, the principal decides to select the teacher for an interview, the teacher accepts the interview, the principal decides to offer the teacher the job, and finally the teacher accepts the job. We use the earliest action we observe on both sides of the market to infer preference orderings: the teacher decision to apply and the principal rating decision. Conceptually, early stages are less susceptible to strategic considerations than later stages because it is free to drop out later. Moreover, by using the earliest action we allow later actions to be strategic.

III.A The teacher perspective

III.A.1 How we model applications

The district's labor market consists of potential teachers, indexed by j, and a set of positions, indexed by p. Each position is associated with a specific school, $k = k(p)$, and may be assigned to at most one teacher. The exception is the outside option ($p = 0$), which includes leaving the district or teaching and has unlimited capacity.

At the beginning of year t, each teacher has an assignment, denoted by c. For teachers new to the district, this assignment is the outside option ($c = 0$), while for incumbent teachers, the assignment is j's position in the prior year, $c = p(j, t-1)$. Teachers may always keep their initial assignment. On an exogenous date $r = r(j, t)$, teacher j enters the transfer system.[12] If she enters, then she is active in the transfer system until an exogenous end date, $r' = r'(j, t)$.

If the teacher enters the transfer system, then she may apply to any position p that is active at some point between r and r'. These positions comprise her choice set, $\mathcal{P}_{jt}$. There is no marginal cost to applying and there is no limit on the number of applications she can submit within the choice set. Let a_{jpt} be an indicator for whether teacher j applied to position p in year t. A teacher's application a_{jpt} is known only to position p and teacher j.

These assumptions lead teachers to treat the application process non-strategically by applying to any position with utility higher than her current position and the outside option. A teacher submits an application to position p if:

$$a_{jpt} = \mathbf{1}\{u_{jpt} > \max\{u_{jct}, u_{j0t}\}\}, \tag{3}$$

where u_{jpt} is teacher j's utility from working at position p in time t.

III.A.2 Model assumptions

There are three key assumptions that underlie this model of teacher application behavior. First, applications are non-strategic: if a position is more appealing than the outside option and current position, then the teacher applies. Second, the teacher considers all vacancies that overlap with her timing. Third, the set of positions the teacher sees is conditionally exogenous.

[12] We assume entry into the system is exogenous. We discuss selection into the system in Appendix IV.

Non-strategic applications: Assuming non-strategic applications is reasonable because of three institutional features. First, we focus on applications rather than interviews or the decision to accept the job. The application stage is less susceptible to strategic considerations than later stages because the teacher does not have to commit to an interview or accepting the job.[13] Second, the marginal cost of applying to a vacancy is effectively zero (it just requires clicking submit given already uploaded materials) so it is reasonable that a teacher just compares a given position to the outside option. Third, principals do not see the teacher's other applications, which limits complicated signaling stories.

Several empirical patterns are also consistent with non-strategic behavior. If teachers were instead strategic in submitting applications, then most models would imply a dynamic portfolio strategy where teachers might delay when they apply to a vacancy. We empirically investigate the frequency of delayed applications by constructing a measure of a teacher's wait time to apply to a vacancy. We calculate the time elapsed between the first day a teacher could have applied to a vacancy and the day the teacher actually applied to the vacancy, where we assume that the teacher only learns that a vacancy is available on days she logs into the system and applies. The top panel of Figure I shows that the median wait time to apply to vacancies that were already posted on the first day the teacher logged into the system (the "stock" of vacancies) is 0 days. The bottom panel shows that the median wait time to apply to vacancies that were posted after the first day the teacher applies (the "flow" of vacancies) is also 0 days. We thus find minimal waiting to apply to positions, such that teachers are unlikely to be engaging in dynamic portfolio strategies.

We similarly find no evidence of strategic delays resulting in non-applications.[14] Strategic non-applications imply asymmetric behavior according to market conditions. When a school posts two vacancies in a cycle, delaying an application is more useful for an early posting than a late posting. If applicants are trying to delay, then we might see higher application rates for the latter of the two vacancies. Appendix Table A6 shows that both the conditional (applied to the other position) and unconditional application rates are very similar for the earlier and later vacancies. This symmetry thus provides further evidence against the presence of strategic non-applications.[15]

[13]Even if teachers wanted to avoid the chance of having to take a future costly action (interview or offer), these actions are extremely rare. The mean number of interviews for a teacher is 0.2 (and 0.3 for an internal teacher), and a given position interviews on average only 2 teachers (Appendix Table A4). Thus, because a "successful" application is quite rare, it is hard for strategic considerations to enter.

[14]A strategic non-application requires that the vacancy closes while the applicant is waiting. But Appendix Table A5 shows that vacancies clear very slowly, especially early in the cycle.

[15]Even if non-applications were common, they would reduce our information about preferences but would not necessarily affect our results. In Table V (Part 2), we report a robustness check where we use the baseline estimates to simulate teacher utilities for each position. Among the positions each teacher actually applied to, we then convert the least preferred 20% of these to non-applications, provided there is at least one application remaining. We re-estimate the teacher preference model with the altered applications and find nearly identical results to the baseline.

Teachers consider all available vacancies: It is reasonable to assume teachers consider all available vacancies because teachers appear not to delay applications. If teachers were unaware of some open vacancies, then we would expect teachers to apply frequently after the first opportunity to do so. We see little evidence of such delayed applications. This pattern could reflect teachers missing a vacancy when it is posted and never searching for older vacancies. But we see the opposite – on the first day of applying, teachers apply to old and new vacancies, with a mean vacancy length of 23 days (Appendix Table A7, panel B).[16]

(Conditionally) Exogenous choice sets: Choice sets reflect a teacher's time in the market, which is an equilibrium outcome related to our behaviors of interest. For example, as we will find below, principals are slightly more likely to hire high value-added teachers, which would remove them from the market faster than teachers who do not receive offers. We therefore do not expect that choice sets will be identical on average across teacher types. Rather, we assume that conditional on observable characteristics, variation in teachers' choice sets, which our model links to variation in market entry and exit dates, is unrelated to teachers' idiosyncratic preferences for certain positions or position types.

We assess this assumption by examining teacher entry and exit patterns and vacancy posting patterns. On the teacher side, we find little evidence of strategic timing in entering the market. Comparing teachers with above and below median value-added, we find that they apply for positions at similar times in the cycle (see Appendix Table A9b). While a discernible relationship between value-added and entry timing could still be consistent with choice set exogeneity, the lack of a relationship suggests that entry may be close to random. Further, as we previously described, when entering the market teachers tend to apply to many jobs immediately, which suggests that teachers were not timing their entry for specific jobs. But in case teachers were targeting their entry for when an idiosyncratically desirable set of positions are posted, we conduct a robustness check (see Table V (Part 2)) where we estimate teacher preferences leaving out all vacancies that were posted within one week of when the teacher first started applying. Thus, these preference estimates reflect application behavior to positions posted well before or well after the day the applicant first applies.

The case for conditionally exogenous market exit is more complicated because one reason for exit—receiving and accepting a job offer—is possibly related to idiosyncratic preferences. But teachers exit the market for multiple reasons, and indeed we see that many teachers—including those who do not successfully transfer—stop applying long before the end of the hiring season (9% in April or before, 15% in May, 21% in June; see Appendix Table A7, panel C). This pattern suggests that much of exit is driven by shocks unrelated

[16] Appendix Table A8 shows statistics about the distribution of the time between when a vacancy is first posted and when a teacher applies for all applicants (median of 7 days), hired candidates (median of 5 days), interviewed candidates (median of 5 days) and positively assessed candidates (median of 7 days).

to accepting a job, or to the nature of the jobs being posted. Even for the teachers who leave the market by accepting a job, the job offer often comes well after the teacher applied.[17] This delay leaves a long period when the teacher may keep applying to more positions even while her preferred position is sitting on her application. To avoid further any potential relationship between when applicants leave the market and their idiosyncratic preferences for the positions available at that time, we conduct a robustness check (see Table V (Part 2)) where we estimate teacher preferences based only on vacancies that were available the day the teacher first applied for jobs that cycle.

On the school side, vacancy posting is spread throughout the hiring season. We split schools into Title I and non-Title I (Title I schools are high-poverty schools). Given results elsewhere in the paper, Title I schools on average are less sought after schools. Appendix Figure A2 looks at the distribution of first and last posting dates by type. The main feature of the graph is that postings are spread throughout the hiring season. Even within school, there is vast variation in the timing of postings across years: pooling across the years in our data, 85% of schools that post jobs in July also post jobs in April, and a similar pattern holds for schools with April postings (see Appendix Table A9c). The secondary feature of Appendix Figure A2 is that if anything more sought after vacancies (non-Title I) are active later in the hiring season.[18] Appendix Table A11 confirms this broad pattern for a variety of other student demographic characteristics. If we zoom in on multiple vacancies posted within school, then Appendix Table A12 shows that the earlier vacancy is more likely to hire a teacher with non-missing value-added but conditional on hiring a teacher with non-missing value-added, the later vacancy hires slightly better teachers. Combined, these features suggest that there is likely little correlation between teacher characteristics and the set of vacancies that they see.[19]

As a result of the preceding discussion, we construct a teacher's start (r) and end (r') (search) date as the dates of her first and last application, respectively. We thus estimate fairly large choice sets out of which teachers make a large number of choices, which helps us estimate preference heterogeneity. Specifically, the mean choice set size is 159 (median: 139), and the mean number of applications is 23 (median: 8).

[17] In Appendix Table A10, panel C, we show that 10% teachers are still applying to positions 23 days after the hired teacher did.

[18] Appendix Figure A3 shows that vacancy fill rates do not differ very much over the cycle or between Title I and non-Title I schools.

[19] Table V (Part 2) shows robustness to a seven-day buffer on both ends or to dropping teachers who only apply to one school. If choice sets are restricted, then fixing the deviations is further from first-best.

III.B The principal perspective

III.B.1 How we model principal behavior

Each position p is associated with a principal with the same index. Principal p derives value v_{jpt} from teacher j holding the position in year t. We model a principal as giving teacher j a positive rating ($b_{jpt} = 1$) if the value is positive: $v_{jpt} > 0$. A positive rating is at least one positive outcome: recording a positive note about the application, offering an interview, or extending a job offer. While we will often refer to values as reflecting utilities, principals may rank a teacher higher because of poor information rather than a utility comparison. Either interpretation is consistent with the paper's results and only affects the labeling of the hypothetical policy that would alter principals' choices.

III.B.2 Model assumptions

There are two assumptions underlying our model of principal behavior. First, principals value applicants who receive a positive outcome more than those who do not. Second, principals consider all applicants.

Principals value applicants with a positive outcome more than those without: The note-taking system is supportive of the first assumption. Principals may be strategic in deciding on interviews or offers if such actions are costly and a preferred teacher may have a low probability of accepting. Because the note-taking system allows principals to rate applicants with no direct consequences, principals can reveal their preferences while remaining strategic in consequential actions.[20]

Principals consider all applicants: The second assumption is reasonable because we see no relationship between when an applicant applied and the applicant's outcome. The applications that receive ratings are similar in timing to those that the principals do not rate (see Appendix Table A10).[21]

IV Production of student achievement

In this section, we first lay out the production model, which specifies teacher output at each school. Second, we describe our three-step estimation procedure and discuss parameter estimates. Third, we present a range of validation checks.

[20] While our assumptions allow for strategic interviews and offers, we do not find evidence that strategic behavior is common enough to affect our conclusions. Table V (Part 3) shows that results are robust to instead modeling principal behavior with a rank-order logit (where hires imply larger utilities than interviews, etc.), including where we restrict to only active choices (i.e., drop applications with no records in the note-taking system).

[21] Table V (Part 4) shows that results are robust to varying which applicants we assume principals consider.

IV.A Model

Given our interest in outcomes for disadvantaged students, we allow teacher value-added to differ between advantaged and disadvantaged students.[22] This choice follows the quickly expanding literature documenting match effects or allowing for comparative advantage (e.g., (Aucejo et al., 2022; Delgado, 2023; Graham et al., 2023; Biasi et al., 2021; Bau, 2022)). We show below (in Table V (Part 6)), however, that all of our conclusions are unchanged if we estimate the homogeneous model that is standard in the literature.

We use notation that follows Chetty et al. (2014a) and Delgado (2023). Let i index students and t index years, where t refers to the spring of the academic year, e.g., 2016 refers to 2015-2016. Each student i has an exogenous type $m(i,t) \in \{0,1\}$ in year t (whether the student is economically disadvantaged). Student i attends school $k = k(i,t)$ in year t and is assigned to classroom $c = c(i,t)$. Each classroom has a single teacher $j = j(c(i,t))$, though teachers may have multiple classrooms.

Student achievement depends on observed student characteristics, teacher value-added, school effects, time effects, classroom-student-type effects, and an error term. Formally, we model student achievement A^*_{it} as:

$$A^*_{it} = \beta_s X_{it} + \nu_{it}, \tag{4}$$

where X_{it} is a set of observed determinants of student achievement and

$$\nu_{it} = f(Z_{jt};\alpha) + \mu_{jmt} + \mu_k + \mu_t + \theta_{cmt} + \tilde{\varepsilon}_{it}. \tag{5}$$

Here, Z_{jt} is teacher experience (and f maps experience into output) and μ_{jmt} is teacher j's value-added in year t for student type m, excluding the return to experience. As in Chetty et al. (2014a), we allow a teacher's effectiveness to "drift" over time. μ_k captures school factors, such as an enthusiastic principal, while μ_t are time shocks. θ_{cmt} are classroom shocks specific to a student type, and $\tilde{\varepsilon}_{it}$ is idiosyncratic student-level variation. We make three standard assumptions to identify the model (see Appendix V).

Our object of interest is a forecast of teacher j's value-added from a hypothetical assignment to a new classroom (or set of classrooms) in school k. Define p_{kmt} as the proportion of type-m students in school k in year t. Given our model of match effects, a teacher's predicted mean value-added at school k in year t is:

$$VA^p_{jkt} = p_{k0t}\mu_{j0t} + p_{k1t}\mu_{j1t} + f(Z_{jt};\alpha), \tag{6}$$

[22] In robustness checks in Table V (Part 5), we consider two alternative splits of students: race and lagged student achievement. We find that our substantive conclusions are nearly identical.

such that a teacher's total value-added for n_{jkt} students is $VA_{jkt} = n_{jkt} VA^{p}_{jkt}$. We use data through $t-1$ from the whole state to forecast VA^{p}_{jkt} for assignments we see in the data and for counterfactual assignments.[23]

IV.B Estimation

We estimate our model in three steps using math scores and data from the whole state.[24] In the first step, we estimate the coefficients on student characteristics by regressing test scores (standardized at the state-level to have mean 0 and standard deviation 1 in each grade-year) on a set of student characteristics and classroom-student-type fixed effects. In the second step, we project the residuals (A_{it}) onto teacher fixed effects, school fixed effects, year fixed effects, and the teacher experience return function. In the final step, we form our estimate of teacher j's value-added in year t for type m (μ_{jmt}) as the best linear predictor based on the prior data in our sample (this prediction includes the experience function). Since in this final step we shrink the estimates, we understate the dispersion in match effects relative to the true dispersion. Using shrunken estimates and prior data implies that we use the information available to policy-makers. See Appendix V.B for estimation details and a discussion of what variation pins down parameters.

Alternative value-added models: We consider four alternative value-added models. The first is a homogeneous effects model, where we assume that teachers' effects on students are type-invariant, rather than allowing for comparative advantage. This model is restrictive relative to our baseline model, but increases our forecast precision. This model tests whether our results rely on comparative advantage or reduced forecast precision. The second model estimates the school effects differently: rather than including school fixed effects (as in, e.g., Jackson (2018) or Mansfield (2015)), we include school-level means of all of the covariates (as in, e.g., Chetty et al. (2014b)). This model tests whether our results depend on how we decompose effects into school and teacher components. Third, we use the Chetty et al. (2014b) estimator. Unlike our "homogeneous" value-added model, this model (a) forecasts using future test scores in addition to past test scores, (b) includes classroom controls like peer mean characteristics rather than school fixed effects, and (c) residualizes test scores using a teacher fixed effect rather than a teacher-year fixed effect. This model tests whether our results are robust to a more "standard" estimator. Finally, we also consider a simple residual estimator where just residualize contemporaneous test score gains for student characteristics. This estimator

[23] Our match effects model is unlikely captures all forms of match effects. But because the match effects vary at the same level as the social planner's objective – i.e., based on whether students are economically disadvantaged – any remaining orthogonal match effects do not affect the results.

[24] Focusing on a single subject allows us to rank all possible levels of output. We follow Biasi et al. (2021) in choosing math because it is typically more responsive to treatment (e.g., Rivkin et al. (2005), Kane and Staiger (2008), and Chetty et al. (2014a) for evidence). In Section VIII we show robustness to including a teacher's value-added on behavioral outcomes.

has the benefit that we can compute value-added for all teachers in the district in a given year and so allows us to directly address concerns about differential missigness of value-added between students of different types. See Appendix V.C for details.[25]

IV.C Validation of the match effects model

To validate our value-added model, we use a version of Chetty et al. (2014a)'s test for mean forecast unbiasedness. We predict a teacher j's value-added in school k in year t (μ_{jkt}) using data from all years prior to t. We then regress the realized mean student residuals in year t ($\bar{A}_{jt}$) on the prediction and test whether the coefficient on our prediction equals 1. Column (1) of Appendix Table A13 shows that the math value-added estimate is an unbiased predictor of residualized output, with a tight confidence interval around 1.06. Appendix Figure A4 shows that forecast unbiasedness holds throughout the distribution of teacher value-added. As our exercise will involve assigning teachers to new schools, forecast unbiasedness across "nearby" assignments may be weak validation for predicting output in "far away" assignments; for example, a teacher's ability with disadvantaged students estimated in a school with a small number of disadvantaged students might be a poor guide to their ability with disadvantaged students in a school with a large number of disadvantaged students. Therefore, we conduct additional tests that validate our estimates over moves similar to those in our counterfactuals. Column (4) of Appendix Table A13 shows mean forecast unbiasedness nearly holds for transferring teachers (with a coefficient of 0.98, not statistically different from 1) while the last two columns show mean forecast unbiasedness even for cases where teachers switch between classrooms with very different compositions or sizes.

We conduct a similar test for the comparative advantage component of value-added. In column (2) we compare our forecast of the difference in a teacher's value-added across (economically) disadvantaged and advantaged students with the realized test score difference. Again, we find that our estimates are nearly forecast unbiased. Appendix Figure A5 shows that forecast unbiasedness holds throughout the distribution. Appendix V.D further assesses the validity of the comparative advantage component of value-added, providing inference around relevant structural parameters (the estimated correlation between teacher value-added with advantaged and disadvantaged students is 0.86), likelihood tests, and additional validation around transferring teachers.

[25] Table V (Part 6) shows that our central conclusions do not depend on which value-added model we use.

V Teacher preferences

V.A Parameterization

We adopt a characteristics-based representation of teacher utilities over positions, which helps us to estimate preference heterogeneity. Teacher utilities over positions are:

$$u_{jpt} = -\gamma d_{jpt} + \pi_j \hat{VA}_{jpt} + \beta_j X_{pt} + \eta_{jt} + \varepsilon_{jpt}. \tag{7}$$

Teacher utility for the outside option is $u_{j0t} = \varepsilon_{j0t}$. d_{jpt} is the one-way commute time (in minutes) between the teacher and the position and will serve as a numeraire for exposition. VA_{jpt} is teacher j's total value added at position p in year t.

Predicted value-added, $\hat{VA}_{jpt}$, combines absolute and comparative advantage. We define a teacher's absolute advantage to be her predicted value-added at a representative school: $AA_{jt} = n_{0t}\hat{\mu}_{j0t} + n_{1t}\hat{\mu}_{j1t}$, where n_{mt} is the average number of type m students in a classroom in the district. Comparative advantage, CA_{jpt}, at a specific position is then the difference between predicted value-added at school $k(p)$ and absolute advantage: $CA_{jpt} = VA_{jpt} - AA_{jt}$. Because we control for absolute advantage in the person-time effects, the coefficient on $\hat{VA}_{jpt}$, π_j, captures the strength of teachers' preferences for schools where their comparative advantage is high. We allow for preference heterogeneity by including a random coefficient in π_j:

$$\pi_j = \bar{\pi} + \sigma^{VA}\nu_j^{VA}, \tag{8}$$

where $\nu_j^{VA} \sim^{iid} N(0,1)$.

X_{pt} is a vector of observed characteristics of positions: the fraction of a school's students that are (1) above the median in prior year math test scores (s), (2) Black (b), and (3) Hispanic (h), and (4) the average number of students in a class at the school that are economically disadvantaged (e). We allow for heterogeneous preferences:

$$\begin{aligned} \beta_j^e &= \beta_0^e + \beta_1^e AA_{jt} + \sigma^e \nu_{jt}^e \\ \beta_j^b &= \beta_0^b + \beta_1^b AA_{jt} + \beta_{j2}^b Black_j + \sigma^b \nu_{jt}^b \end{aligned} \tag{9}$$

where $Black_j$ is an indicator for teacher race category and ν is a vector of independent, standard normal random coefficients, which captures the standard deviation of idiosyncratic preferences. The equations for lagged achievement and Hispanic are parallel.[26]

[26] Table V (Part 7) shows that our results are robust to allowing for correlation in the random coefficients.

We follow Mundlak (1978) and Chamberlain (1982) and model η_{jt} using correlated random effects. We model teacher-year unobserved heterogeneity in preferences for teaching in the district as the sum of several components:

$$\eta_{jt} = \lambda Z_{jt} + \rho CM_{jt} + \sigma^{\eta} \nu^{\eta}_{jt}. \tag{10}$$

Z_{jt} are teacher-year characteristics – whether the teacher is in the district, whether the teacher is Black, whether the teacher is Hispanic, whether the teacher is female, the teacher's predicted value-added for economically disadvantaged students, the teacher's predicted value-added for non-economically disadvantaged students, and dummy variables for whether the teacher has 2-3 years of prior experience, 4-6 years of prior experience, or more than 6 years of prior experience. CM_{jt} is a set of teacher-year averages of the variables that vary across the job postings within teacher-year (value-added, commute time, interactions of teacher and school characteristics). Through CM_{jt}, we allow unobserved heterogeneity to be correlated with CA_{jpt} and X_{pt}. Finally, ν^{η}_{jt} is an independent standard normal random effect.[27]

ε_{jpt} is an iid Type I extreme value error. Let $V_{jpt} = u_{jpt} - \varepsilon_{jpt}$ be j's representative value for position p in year t. Then the distributional assumption on ε_{jpt} implies that:

$$Pr(a_{jpt} = 1) = \frac{exp(V_{jpt})}{1 + exp(V_{jct}) + exp(V_{jpt})} \text{ and } Pr(a_{jpt} = 1) = \frac{exp(V_{jpt})}{1 + exp(V_{jpt})}, \tag{11}$$

for teachers already in the district and teachers new to the district, respectively.

V.B Estimation and Identification

The data we use to estimate teacher preferences are applications to positions, and the method we use is maximum simulated likelihood, where we simulate from the normal distributions of the random coefficients. Let n index each simulation iteration and let $A_{jptn}(\theta)$ be the model-predicted probability that j applies to position p in year t in simulation iteration n at parameter vector θ. For each teacher j in year t, we construct the simulated likelihood as:

$$L_{jt} = \frac{1}{500} \sum_{n=1}^{500} \prod_{p \in \mathcal{P}_{jt}} (a_{jpt} A_{jptn}(\theta) + (1 - a_{jpt})(1 - A_{jptn}(\theta))), \tag{12}$$

[27] In Table V (Part 8), we consider binary logits, and show that our results are robust to either omitting random effects, or to including various combinations of teacher and school random and fixed effects.

where a_{jpt} is an indicator for whether j applied to p in the data. Our full simulated log likelihood function is:

$$l = \frac{1}{J} \sum_j log L_{jt}. \tag{13}$$

In Section III we argued that the institutions and data are consistent with teachers applying non-strategically. Under this assumption, the choices that teachers make identify preferences and preference heterogeneity. Heuristically, if within her choice set a teacher is more likely to apply to positions with a particular characteristic than a position without this characteristic, then we infer that the teacher has a preference for schools with this characteristic. Similar reasoning applies for mean coefficients, and observed and unobserved preference heterogeneity.

We seek to predict teachers' valuations over positions rather than causal effects of changes in characteristics on choices. In counterfactuals, we give utility bonuses as a function of school characteristics and so do not assume that teachers value money or these characteristics. As a convenient way to interpret magnitudes, we sometimes convert utility to minutes of commute time, which requires the stronger assumption that commute time is exogenous. We do not rely on having consistently estimated the causal effect of commute time, however, because we only make relative comparisons of the costs of various policies.

V.C Teacher Preference Estimates

Table I presents the teacher preference estimates. First, teachers prefer positions with more advantaged students. Second, teachers dislike positions with longer commutes. Finally, teachers have only slight preference toward positions where they have higher value-added.[28]

Responsiveness to school and match characteristics varies with observable and unobservable heterogeneity. For example, teachers with higher absolute advantage have relatively lower preferences for schools with more disdvantaged students. We also find a large positive same-race premium for Black teachers and schools with large fractions of Black students. In terms of unobserved heterogeneity, we typically find substantial dispersion in the random coefficients. For example, a standard deviation of the random coefficients on the number of disadvantaged is about the same as the mean valuation.

To help interpret the strength of—and heterogeneity in—some of these relationships, Panels (a) through (c) of Figure II show how the average rank of positions in teachers' preferences change as single characteristics change, as well as the 10th and 90th percentile of these positions in teachers' rankings. We do not

[28]We use the value-added forecast, $\hat{VA}_{jt}$, in our preference model. In Table V (Part 9), we show robustness to excluding value-added derived variables in our preference model.

hold other characteristics fixed so that, for example, when we study commute time, other characteristics of schools are potentially changing. The figure emphasizes that commute time is a powerful predictor of rankings: changing commute time from 5 minutes to 25 minutes decreases the average rank of a position (for the average teacher) from about the 70th percentile to the 50th percentile. Similarly, the fraction of students that are disadvantaged is a powerful predictor of ranking: across the support, the mean ranking moves by over 20 percentiles. In contrast, while teachers do pursue comparative advantage (see the coefficient in Table I), this relationship is quite weak: across the support of the data, varying teachers' comparative advantage only increases the rank of a position by a couple of percentiles. The figures also emphasize that there is substantial heterogeneity in teachers' rankings of positions: across the support of these characteristics, the range from the 10th percentile in the teacher distribution to the 90th is very large.

Hence, not only do teacher preferences deviate from those that would decentralize the planner's solution, they are negatively correlated. With minimal assumptions and data on real choices, we confirm the findings of the teacher preference literature regarding mean preferences but estimate considerable heterogeneity.

VI Principal behavior

VI.A Parameterization and identification

We adopt a characteristics-based model and parameterize v_{jpt} to be a linear function of position and teacher characteristics, a random effect, and an idiosyncratic teacher-position error:

$$v_{jpt} = \alpha_p W_{jpt} + \sigma_\kappa \kappa_{pt} + \upsilon_{jpt}. \tag{14}$$

To allow principal behavior to possibly align with output, W_{jpt} includes j's total predicted value-added at school $k(p)$.[29] We further include teacher characteristics: teacher prior experience (in bins of 2-3 years, 4-6 years, and 7+ years), whether the teacher has a Masters degree, whether the teacher is licensed, whether the teacher is certified, the teacher's Praxis score, whether the teacher is Black, whether the teacher is Hispanic, and whether the teacher is female.[30] Finally, we include a constant and interact whether the teacher is Black with the fraction of the school's students that are Black and whether the teacher is Hispanic with the fraction of the school's students that are Hispanic. We exclude salary because principals in our empirical context do

[29] We include predicted value-added, rather than realized value-added, in W_{jpt} so that principals only incorporate the information available at the time the application was received.

[30] We also include indicators for whether each covariate is missing. The Praxis test is a standardized teacher certification test administered by the Educational Testing Service.

not have to pay teacher salaries out of a school budget. We allow principals to have heterogeneous valuations over teachers based on W_{jpt} by letting α_p vary with whether the school has Title I status.

To capture heterogeneous outside options and variation in propensity to assign ratings, κ_{pt} is a normally distributed random effect. Finally, υ_{jpt} is i.i.d. Type I extreme value.

As with teachers, identification is straightforward given our characterization of the process in Section III. We observe the set of applications that a principal receives and we observe whether a principal gives an application a positive outcome. We interpret the decision to give an application a positive rating as a non-strategic and costless action. This interpretation allows us to infer principal valuations from their choices in a straightforward way: those that are rated positively are preferred to those that are not. Because we observe the ratings, even if *interviewing is costly* and so principals are strategic at this stage, then our identification assumption still holds. One might also worry that assigning a rating is costly, and so it is done strategically. To alleviate this concern, we show below that if we restrict attention to applications where a principal assigned a rating (either positive or negative), then our results are quantitatively identical (see Table V (Part 3)).

VI.B Estimates

Before presenting estimates from our baseline model, we consider what types of characteristics determine principal ratings. Appendix Table A14 presents the changes in pseudo-R^2s from including different sets of observable teacher characteristics. The main set of characteristics that explain ratings decisions are various observable characteristics of teachers: experience, licensing, certification, and Praxis scores. While one might think that these characteristics would predict value-added, in Appendix Table A15 we show that they have very limited predictive power. Indeed, value-added by itself or in addition to other characteristics generates very small changes in model fit.[31]

Despite the small explanatory power of value-added in principal decisions, Table II shows that principals do favor teachers with higher value-added in our baseline model.[32] We also observe significant heterogeneity, as Title I school principals rate Black and Hispanic teachers more positively than non-Title I principals do. To help interpret the strength of the value-added relationship, Panel (d) of Figure II shows that the mean percentile of teachers in principals' ratings goes from the 35th percentile to the 60th percentile across the support of projected value-added. Consistent with the idea that observed characteristics poorly predict

[31]EVAAS, the state of North Carolina's value-added measure, has even less explanatory power. As principals have access to this information, it is unlikely that the estimated weights principals place on value-added are due to measurement error in our estimates of value-added. Our results are quantitatively robust to significant amounts of attenuation. See Table V (Part 10).

[32]See Appendix VI for the likelihood, which closely parallels the one for teachers.

value-added, Appendix Figure A6 shows that if we omit value-added from the principal model then the relationship dramatically flattens. Nevertheless, the strength of this relationship is difficult to directly interpret. To assess the extent to which value-added explains principal decisions, in the model section below, we compare the allocations achieved using the estimated principal behavior to those with random principal behavior.

The relationship between value-added and principal rankings could reflect preferences or information. Distinguishing between these does not affect analysis of the current allocation (or the counterfactuals) because we will compare how principals currently act with a proposed alternative ranking.[33] But if incomplete information explains principal rankings, our empirical strategy might use the data differently. For example, we use principals' notes for identification because we can then allow for strategic interviewing or offering. But if interview or offer decisions deviate from notes because information resolves (rather than strategy), then we would want to use interviews and offers and not the notes. We show, however, in Table V (Part 3) that principal models estimated using only offers delivers nearly identical results.

Hence, consistent with the previous literature, principal valuations deviate from those that would implement the planner's solution Whether the positive relationship between rankings and value-added is strong enough to generate allocations close to the planner's solution depends on how both sides combine in equilibrium.

VII The current allocation

With our model estimates, specifically of teacher value-added, we now discuss the current allocation of teachers across schools.

Student and teacher characteristics: Table III presents properties of the current allocation where we report student-weighted results when we split students by our measure of economic disadvantage. We report results in our focal district, as well as in all other districts in North Carolina. Disadvantaged students are more likely to be minorities. Disadvantaged students also have teachers with worse observed characteristics. Specifically, they are less experienced, less likely to have a graduate degree, a regular license, be certified, and have lower Praxis test scores (a standardized test).

[33] Appendix Table A14 shows that the model's explanatory power decreases when using the readily-available EVAAS measure.

Test scores and teacher value-added: Between advantaged and disadvantaged students, there are large achievement gaps in levels. But in gains, we see no gaps. This "raw" data fact hints that there are not large differences in learning across schools, which suggests that the average quality of teachers is likely similar.

Looking across a variety of measures of teacher value-added, the broad pattern is that disadvantaged students have teachers of similar strength to advantaged students. This pattern is true both in our focal district, as well as in the rest of North Carolina. As we mentioned in the introduction, this finding is not new to us and has been found in many districts across the United States (see footnote 4). Specifically, with our baseline value-added model, we find equivalent value-added with advantaged and disadvantaged students among teachers of advantaged and disadvantaged students. Our alternative value-added models find similar patterns: with homogeneous value-added, advantaged students have a slight advantage and this grows slightly with the Chetty et al. (2014b) estimator and an estimator that uses school mean characteristics rather than school fixed effects. The estimators that find a slight advantage for advantaged students in our focal district tend to find smaller differences in the rest of North Carolina. The table also reports measures of behavioral value-added (see Appendix V.E for details on how we construct behavioral value-added) and shows that they are approximately balanced across advantaged and disadvantaged students.

Other student classifications: In Appendix Tables we present similar sets of summary statistics for a wide variety of alternative "splits" of students: splitting students by race (Appendix Table A16), by lagged achievement (Appendix Table A17), by a measure of persistent disadvantage (Appendix Table A18, and see Appendix Table A19 for the relationship between disadvantage and persistent disadvantage in our sample), and splitting by school characteristics (high-poverty vs. not) rather than by student characteristics (Appendix Table A20). The basic patterns persist across all these variants. We emphasized in Section II.B that there is minimal within-school sorting of students across classrooms. Validating the lack of within-school sorting, Appendix Table A21 shows similar patterns when we measure the advantaged-disadvantaged gap using school averages of teacher characteristics.[34]

Missing value-added: One critique of this finding is that it refers to teachers for whom we can estimate value-added, and disadvantaged students are especially likely to have inexperienced teachers for whom we cannot estimate value-added. Table III shows that disadvantaged students are more than twice as likely to have a teacher for whom we cannot estimate value-added.

[34] Appendix Figure A7 shows the result visually. If we classify schools by their mean teacher value-added, we find that the share of disadvantaged students is weakly increasing in the school's mean teacher value-added. This pattern holds for three different value-added measures.

To address the concern about differential missingness, we report results of the residual value-added estimator (the teacher's mean of A_{it}, in the notation of Appendix V), which only uses data from the current year and so can be estimated for all teachers. This value-added estimator finds similar patterns. As an alternate measure, we impute value-added for the teachers for whom we cannot estimate value-added. At a high level, we use the set of observed characteristics in the top portion of the table (Appendix VII details the exact imputation process). Naturally, since one of the themes of this paper is that observed characteristics poorly predict value-added, there is a limit to how good the imputation model can be, though this may simply reflect that principals also have limited information. The main finding is that including inexperienced teachers does not alter the central message of the table that disadvantaged and advantaged students have teachers of similar strength.

VIII Understanding the current allocation

To understand how the current allocation is generated, we simulate the market equilibrium by combining the estimated market timing from Section III, the estimated match-specific output from Section IV, the estimated teacher preferences from Section V, and the estimated principal valuations from Section VI.

VIII.A Simulation details

We consider allocating the set of *teachers who apply for positions* in the district in the 2015-2016 cycle, including teachers who are not currently in the district. We restrict attention to the teachers for whom we can compute value-added, which includes teachers who have previously taught anywhere in the state. This restriction drops a large number of teachers: we end up with 178 elementary school teachers and 296 positions. To avoid the possibility of artificial imbalance playing a role in our estimates (see Ashlagi et al. (2017)), in each of 400 simulation runs we randomly drop positions so that there are the same number of teachers and positions. In Appendix VIII.A, we allow teachers to outnumber positions.

While we estimate a distribution of random coefficients, in simulations we use the single draw of the random coefficients per teacher and principal that maximizes the likelihood for the teacher or principal. We draw i.i.d. type I errors for ε_{jpt} and υ_{jpt}.

In using DA to find stable allocations, we have teachers and principals submit rankings according to their true preferences. If there are multiple equilibria, then for one side of the market it is not a dominant strategy to report truthfully. Below we show, however, that the equilibrium is essentially always unique and so truthful reporting is a dominant strategy.

For teachers and vacancies that are not in each other's choice sets, we assign a large negative number to the valuations. We do not include an outside option when we run DA. Given that we impose balanced markets, all teachers are hired and all positions are filled (in Appendix VIII.A some teachers are not hired).

VIII.B Model fit

We now turn to the fit of the model under the status quo. Because we estimate several model components fairly directly from data, fit largely highlights how well our market equilibrium assumption (pairwise stability) performs. Figure III shows that the model matches the basic qualitative patterns in the data: schools with a larger share of disadvantaged students have teachers (a) with stronger absolute advantage, (b) with comparative advantage in teaching economically disadvantaged students, (c) less likely to be experienced, and (d) more likely to be Black. Quantitatively, the model almost exactly matches the slope for teacher experience and whether teachers are Black. The model underpredicts the slope in absolute advantage.

Figure IV (and Table IV) shows that in the estimated status quo, disadvantaged students are assigned slightly *better* teachers than advantaged students. This feature matches the data.

VIII.C The importance of second-best reasoning

In the last section, we documented that advantaged students have no more effective teachers than disadvantaged students. Relative to the structure of teacher preferences, this balance is surprising in that teachers' revealed preference is strongly averse to teaching at schools with disadvantaged students. In this section, we explain this result through the economics of two-sided markets and the theory of the second best.

A couple of subtle explanations play no role in explaining the current allocation. First, there is no room for equilibrium selection. Changing the equilibrium from the teacher-proposing equilibrium to the school-proposing equilibrium has no effect on the allocation. Second, timing has little role. Making all vacancies and teachers active at the same time increases output slightly for advantaged students and barely decreases it for disadvantaged students. We show these and other allocations in Figure IV and Table IV.

Aligning teacher and principal preferences with the planner's solution shows that there are important interactions between both sides of the market, such that thinking about one side at a time leads to ineffective or harmful policy ideas. First, if teachers had preferences that would decentralize the planner's solution—they only care about the number of disadvantaged students in a school—then the allocation is little changed. Thus, a natural teacher-side policy is ineffective. Second, if principals had preferences that would decentralize the planner's solution—they only care about the output in the match—then the allocation is worse for equity and resembles what we might expect based on the structure of teacher preferences.[35] Thus, a natural policy based on one-sided reasoning is harmful.

One-sided reasoning is misleading here because of the theory of the second best: preferences on both sides of the market deviate from the preferences that decentralize the planner's solution, but these deviations interact to generate surprisingly favorable allocations. Were we to eliminate the deviation on the principal side of the market and have principals order teachers by value-added, then the strongest teachers would reach their most preferred schools. Given the structure of teacher preferences, this change would lead advantaged students to have much more effective teachers. Hence, by placing weight on factors other than value-added, principals "push back" on teacher preferences and overcome differences in applicant pools across positions.

Reaching these conclusions required an equilibrium model and data to identify preferences from actions rather than equilibrium assignments. With data only on equilibrium assignments, typically one assumes that one side of the market has vertical preferences, which fills in the choice sets for the other side of the market (see Diamond and Agarwal, 2017). If we had (incorrectly) assumed principals have vertical preferences in value-added, then we would have concluded that the status quo was very unfavorable to disadvantaged students, and teacher bonus policies by themselves were effective.

Figure IV (and Table IV) shows that there are substantial gains in the first-best. Disadvantaged students gain about 0.06σ, or about one-fourteenth of the unconditional achievement gap that we document in Table III. While these numbers refer only to teachers in the transfer system, in Appendix Table A22 we show that these gains are similar if we look at all teachers in the district. Naturally, these gains are not costless—they come somewhat at the expense of advantaged students, whose teacher quality suffers, but total output still increases (by about 0.021σ in the transfer sample, which is similar to the 0.016σ in the whole sample).

Finally, Figure IV (and Table IV) shows that the combination of the two policies mentioned above—teachers rank schools based on the number of disadvantaged students and principals rank teachers based on projected output—achieves 94% of the first best (the remaining gap is due to comparative advantage and timing). Thus, in Section IX we study policies that move us closer to this point.

VIII.D Parameterizing teacher preferences and principal behavior using the model

In the previous subsection we emphasized stylized features of teacher preferences and principal behavior to explain our results. First, teachers prefer not to teach at schools with disadvantaged students. Second, principals do not place very much weight on value-added.

[35] The allocation is also worse for efficiency: per student output declines by about 0.009σ.

To more directly parameterize the inequity for disadvantaged students implied by the structure of teacher preferences, in Appendix Table A23 we display the results of two exercises. First, we ignore market clearing and assign each teacher her preferred position so some positions have multiple teachers and some have none. Conditional on receiving a teacher, disadvantaged students do as well as, or better than, advantaged students. The assignment rate, however, is dramatically different across advantaged and disadvantaged students. Thus, if we follow teachers' preferences and ignore market clearing, then few disadvantaged students would receive teachers. Second, we impose a market clearing mechanism that lets teachers' preferences matter the most. Specifically, we clear the market using a serial dictatorship ordered by teacher's value-added with disadvantaged students. Here, disadvantaged students do dramatically worse.

To more directly parameterize principal behavior, we simulate equilibrium allocations where we give principals random preferences over teachers. Table IV shows that this allocation is very similar for disadvantaged (and advantaged students) to the status quo. Thus, the loose heuristic that principals hire essentially randomly is a decent approximation to the data. Relatedly, Appendix Table A24 shows that principals make mistakes in the sense that they have much better teachers in their choice set than the ones they either rate positively, interview, or hire.

VIII.E Different objectives

Our social planner maximizes disadvantaged students' output. Here, we consider how our results might change with alternate objectives.

First, the social planner may place weight on other forms of output, not just math test scores. We estimate teachers' value-added on an index of behavioral outcomes and find that behavioral value-added is still balanced across advantaged and disadvantaged students (Table III, see Appendix V.E for details on how we construct behavioral value-added).

Second, the social planner may place weight on other agents, not just disadvantaged students. First, the social planner may place equal weight on all students. We formalize this objective in Appendix III. Row 11 of Table V shows that aligning principals' preferences with the social planner's objective function still lowers total academic achievement. Aligning teachers' preferences with the social planner's, though, can lead to some total output gains. When we implement both fixes the output is only 70% of the way from the status quo to the first best, whereas for equity it is 94% of the way there. The reason is the lack of personalized pricing. Specifically, making preferences be based on output is equivalent to assigning a price per unit of output. To maximize output, however, the planner wants to "personalize" prices by allowing them to depend on the specific combination of value-added with advantaged and disadvantaged students. Thus, from the perspective of efficiency, the lack of personalized pricing limits sorting on comparative advantage and is an important barrier to achieving desirable allocations. Second, the social planner may place weight on teacher utility. In Section IX, we constrain the policies we consider to make each teacher weakly better off than in the status quo.

VIII.F Robustness and extensions

In Table V (panels 1 through 11), we report the robustness checks we have mentioned throughout the text. When we study equity, the following three findings are robust across all of these alternatives: first, there are large gains from moving to the first-best; second, fixing one of the deviations from what decentralizes first best (making teachers value the number of disadvantaged students or principals maximize value-added) is either ineffective or harmful; and third, that fixing both comes close to implementing the first-best (the exception is when we restrict teacher choice sets to the first day because the timing constraints bite more).

In Appendix VIII, we consider two extensions. First, we consider an active extensive margin. We show that even with an active extensive margin the basic theory of the second message persists. The intuition is that improving principal hiring brings better teachers to the district, but this improvement only benefits advantaged students (Panel 12 of Table V). Second, we clear the market in sub-periods rather than all at once. We find that our estimates of the status quo allocation are nearly identical (Panel 13 of Table V).

IX Teacher bonus counterfactuals

In this section, we consider policies that may move the allocation closer to the first-best. We compare teacher bonus policies that cost the district equivalent amounts while holding all teachers harmless. We then interact these bonuses with principal-side policies.

IX.A Implementation details

The district offers a two-part bonus on the basis of a teacher-position characteristic, z_{jpt}, where each teacher receives a lump-sum amount, b_0, and a bonus b_1 per unit of characteristic, z_{jpt}. Teacher j's utility for teaching at position p in year t is

$$\tilde{u}_{jpt} = u_{jpt} + \gamma(b_0 + b_1 z_{jpt}), \tag{15}$$

where we multiply by the commute time coefficient (γ) to express bonus spending in minutes of commute time. For each b_1, we solve for the teacher-optimal stable equilibrium assignments, where $p^*(j)$ is j's assigned position, given the bonus size and the object that generates the bonus. Thus, because we give teachers utility directly for the characteristic, we do not use our estimated coefficients on the characteristics.[36]

To focus on policies that are likely to receive teachers' support, we make each teacher weakly better off than in the status quo equilibrium. We set the transfer such that the teacher with the worst change is indifferent. This lump-sum transfer can be either positive or negative. Thus, the district's total bonus to a teacher depends both on the choice of how much to compensate for the characteristic and how it changes the allocation.

We examine bonus schemes over two objects. First, we study bonuses based on the number of disadvantaged students the teacher has ($n_{k(p)1t}$). These bonuses mimic the hard-to-staff school bonuses that some districts have piloted. Second, we interact school and teacher characteristics by considering bonuses based on a teacher's absolute advantage times the number of disadvantaged students: $((p_{0t}\hat{\mu}_{j0t} + (1-p_{0t})\hat{\mu}_{j1t})n_{k(p)1t})$.[37]

IX.B Results

Panel (a) of Figure V shows the effect of these two bonus schemes on disadvantaged students' test scores when principals hire according to their estimated preferences. The top line shows achievement in the the first-best allocation. To allow for comparisons across bonus schemes, the horizontal axis is the total realized spending per teacher (normalized to be in minutes of commute time per teacher).

We have three results, all of which reflect the theory of the second best. First, untargeted bonuses for teaching disadvantaged students are relatively ineffective in raising disadvantaged students' test scores. Second, targeted bonuses that pay the best teachers more for teaching disadvantaged students are more effective than untargeted bonuses because they jointly address deviations on both sides of the market. Specifically, the applicant pool only expands among the best teachers, so then the principals' difficulties in identifying good teachers matters less (a random teacher in the new applicant pool is better). Third, the bonuses eventually become less effective as they grow larger. Here, larger bonuses expand the applicant pool for disadvantaged schools, but the larger pool causes the deviation in principal preferences to matter more.

Effective policy needs to address the deviations jointly. In Panel (b) of Figure V we consider the effect of the teacher bonus schemes when principals hire according to value-added. Such hiring rules may be induced by a combination of an information intervention and principal bonuses for hiring effective teachers.

[36] We compare the effectiveness of bonuses with equivalent utility costs. Because we use the same conversion factor for all schemes, the conversion factor does not affect the comparisons.

[37] We implement the bonus schemes starting from the status quo with relaxed timing constraints.

We again have three results. First, as in the prior section, if teacher bonuses are small such that estimated teacher preferences largely guide applications, then principals hiring according to value-added leads to large decreases in disadvantaged students' test scores. Fixing the deviation on the principal side, but hardly closing the teacher deviation, has a large negative effect relative to the status quo. Second, as teacher bonuses get larger, principals hiring according to value-added make the teacher bonuses particularly effective. At the equivalent of about 50 minutes of commute time per teacher, the bonuses have nearly reached the first-best. That teacher bonus effectiveness is increasing in principal bonuses (or information interventions) reflects the interaction of the two sides of the market. Third, for some levels of spending untargeted bonuses now outperform targeted bonuses. Because the principal deviation has been closed, the targeting of bonuses is no longer needed. In fact, such targeting is now counter-productive.

X Discussion

We have studied the equity consequences of the within-district allocation of teachers to schools. We consider both the current allocation and alternative policies. To approximately decentralize the first-best that maximizes disadvantaged students' achievement, teachers would need to prefer schools with more disadvantaged students and principals would need to prefer higher value-added teachers. Using rich data from the teacher transfer system that allows us to observe actions, we show that both sides' preferences deviate from these. Nonetheless, and consistent with the theory of the second best, these two deviations interact to generate a surprisingly equitable allocation, where disadvantaged students do not have worse teachers than advantaged students. In terms of policy, and again consistent with the theory of the second best, fixing one deviation at a time is either ineffective or harmful. Fixing both deviations could close about a fourteenth of the achievement gap per year.

More broadly, this paper has demonstrated the value of using rich data to study the functioning of particular labor markets. Our data allows us to estimate the behavior of the main agents in the market, rather than relying on strong assumptions to infer these from the observed equilibrium. In so doing, we have arrived at surprising conclusions about the determinants of the equilibrium and the design of policies. Presumably, other labor markets would also benefit from such analysis.

UNIVERSITY OF CALIFORNIA, RIVERSIDE, UNITED STATES

DUKE UNIVERSITY AND NATIONAL BUREAU OF ECONOMIC RESEARCH, UNITED STATES

UNIVERSITY OF CALIFORNIA, MERCED AND NATIONAL BUREAU OF ECONOMIC RESEARCH, UNITED STATES

STANFORD UNIVERSITY AND NATIONAL BUREAU OF ECONOMIC RESEARCH, UNITED STATES

References

Agarwal, Nikhil, Charles Hodgson, and Paulo Somaini, "Choices and Outcomes in Assignment Mechanisms: The Allocation of Deceased Donor Kidneys," Working Paper 28064, National Bureau of Economic Research 2020.

Antos, Joseph R. and Sherwin Rosen, "Discrimination in the market for public school teachers," *Journal of Econometrics*, 3 (1975), 123–150.

Arteaga, Felipe, Adam J Kapor, Christopher A Neilson, and Seth D Zimmerman, "Smart matching platforms and heterogeneous beliefs in centralized school choice," *The Quarterly Journal of Economics*, 137 (2022), 1791–1848.

Ashlagi, Itai, Yash Kanoria, and Jacob D Leshno, "Unbalanced random matching markets: The stark effect of competition," *Journal of Political Economy*, 125 (2017), 69–98.

Aucejo, Esteban M, Patrick Coate, Jane Cooley Fruehwirth, Sean Kelly, and Zachary Mozenter, "Teacher Effectiveness and Classroom Composition: Understanding Match Effects in the Classroom," *Economic Journal*, 132 (2022), 3047–3064.

Ba, Bocar, Patrick Bayer, Nayoung Rim, Roman Rivera, and Modibo Sidibé, "Police Officer Assignment and Neighborhood Crime," Working Paper 29243, National Bureau of Economic Research 2021.

Ballou, Dale, "Do public schools hire the best applicants?," *The Quarterly Journal of Economics*, 111 (1996), 97–133.

Bates, Michael, "Public and private employer learning: Evidence from the adoption of teacher value added," *Journal of Labor Economics*, 38 (2020), 375–420.

Bau, Natalie, "Estimating an equilibrium model of horizontal competition in education," *Journal of Political Economy*, 130 (2022), 1717–1764.

Biasi, Barbara, Chao Fu, and John Stromme, "Equilibrium in Teachers' Labor Markets with Flexible Pay," Working Paper 28530, National Bureau of Economic Research 2021.

Bobba, Matteo, Tim Ederer, Gianmarco Leon-Ciliotta, Christopher Neilson, and Marco Nieddu, "Teacher compensation and structural inequality: Evidence from centralized teacher school choice in Perú," Working Paper 29068, National Bureau of Economic Research 2024.

Boyd, Don, Hamp Lankford, Susanna Loeb, Matthew Ronfeldt, and Jim Wyckoff, "The role of teacher quality in retention and hiring: using applications to transfer to uncover preferences for teachers and schools," *Journal of Policy Analysis and Management*, 30 (2011), 88–110.

Boyd, Donald, Hamilton Lankford, Susanna Loeb, and Jim Wyckoff, "Analyzing the determinants of the matching of public school teachers to jobs: disentangling the preferences of teachers and employers," *Journal of Labor Economics*, 31 (2013), 83–117.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 36 (2018), S13–S70.

Chamberlain, Gary, "Multivariate regression models for panel data," *Journal of Econometrics*, 18 (1982), 5–46.

Chen, Yan and Yinghua He, "Information acquisition and provision in school choice: an experimental study," *Journal of Economic Theory*, 197 (2021), 105345.

Chetty, Raj, John N. Friedman, and Jonah E. Rockoff, "Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates," *American Economic Review*, 104 (2014), 2593–2632.

—, John N Friedman, and Jonah E Rockoff, "Measuring the impacts of teachers II: Teacher value-added and student outcomes in adulthood," *American Economic Review*, 104 (2014), 2633–79.

Clotfelter, Charles T, Helen F Ladd, and Jacob L Vigdor, "Teacher mobility, school segregation, and pay-based policies to level the playing field," *Education Finance and Policy*, 6 (2011), 399–438.

Combe, Julien, Olivier Tercieux, and Camille Terrier, "The design of teacher assignment: Theory and evidence," *The Review of Economic Studies*, 89 (2022), 3154–3222.

—, Umut Dur, Olivier Tercieux, Camille Terrier, and M Utku Unver, "Market Design for Distributional Objectives in (Re) assignment: An Application to Improve the Distribution of Teachers in Schools," Working Paper 2022.

Condie, Scott, Lars Lefgren, and David Sims, "Teacher heterogeneity, value-added, and education policy," *Economics of Education Review*, 40 (2014), 76–92.

Cowgill, Bo, Jonathan Davis, B Pablo Montagnes, and Patryk Perkowski, "Stable Matching on the Job? Theory and Evidence on Internal Talent Markets," *Management Science*, forthcoming (2024).

Dahlstrand, Amanda, "Defying Distance? The Provision of Services in the Digital Age," Working Paper 2024.

Davis, Jonathan, "The Short and Long Run Impacts of the Deferred Acceptance Algorithm: Evidence from Matching Teach For America Teachers to Schools," Working paper 2022.

Delgado, William, "Disparate Teacher Effects, Comparative Advantage, and Match Quality," Working Paper 2023.

Diamond, William and Nikhil Agarwal, "Latent indices in assortative matching models," *Quantitative Economics*, 8 (2017), 685–728.

Duncan, Greg J and Katherine Magnuson, "The nature and impact of early achievement skills, attention skills, and behavior problems," *Whither opportunity*, (2011), pp. 47–70.

Ederer, Tim, "Labor market dynamics and teacher spatial sorting," Technical Report 2023.

Elacqua, Gregory, Anne Sofie Westh Olsen, Santiago Velez-Ferro et al., "The Market Design Approach to Teacher Assignment: Evidence from Ecuador," Technical Report, Inter-American Development Bank 2021.

Goldhaber, Dan, Vanessa Quince, and Roddy Theobald, "Has it always been this way? Tracing the evolution of teacher quality gaps in US public schools," *American Educational Research Journal*, 55 (2018), 171–201.

Graham, Bryan S., Geert Ridder, Petra M. Thiemann, and Gema Zamarro, "Teacher-to-classroom assignment and student achievement," *Journal of Business and Economic Statistics*, 41 (2023), 1328–1340.

Greenberg, David and John McCall, "Teacher mobility and allocation," *Journal of Human Resources*, 9 (1974), 480–502.

Hanushek, Eric A., John F. Kain, and Steven G. Rivkin, "Why public schools lose teachers," *Journal of Human Resources*, 39 (2004), 326–354.

Isenberg, Eric, Jeffrey Max, Philip Gleason, and Jonah Deutsch, "Do Low-Income Students Have Equal Access to Effective Teachers?," *Educational Evaluation and Policy Analysis*, 44 (2022), 234–256.

Jackson, C Kirabo, "What do test scores miss? The importance of teacher effects on non–test score outcomes," *Journal of Political Economy*, 126 (2018), 2072–2107.

Jacob, Brian A, Jonah E Rockoff, Eric S Taylor, Benjamin Lindy, and Rachel Rosen, "Teacher applicant hiring and teacher performance: Evidence from DC public schools," *Journal of Public Economics*, 166 (2018), 81–97.

— , — , — , — , and — , "Teacher applicant hiring and teacher performance: Evidence from DC public schools," *Journal of Public Economics*, 166 (2018), 81–97.

Johnston, Andrew C., "Preferences, Selection, and the Structure of Teacher Pay," *American Economic Journal: Applied Economics*, (2024).

Kane, Thomas J and Douglas O Staiger, "Estimating teacher impacts on student achievement: An experimental evaluation," Working Paper 14607, National Bureau of Economic Research 2008.

Kapor, Adam J, Christopher A Neilson, and Seth D Zimmerman, "Heterogeneous beliefs and school choice mechanisms," *American Economic Review*, 110 (2020), 1274–1315.

Laverde, Mariana, Elton Mykerezi, Aaron Soujourner, and Aradhya Sood, "Gains from Reassignment: Evidence from a Two-Sided Teacher Market," Working Paper 2023.

Lipsey, Richard G. and Kelvin Lancaster, "The general theory of second best," *The Review of Economic Studies*, 24 (1956), 11–32.

Mansfield, Richard K, "Teacher quality and student inequality," *Journal of Labor Economics*, 33 (2015), 751–788.

Mundlak, Yair, "On the pooling of time series and cross section data," *Econometrica*, 46 (1978), 69–85.

National Council on Teacher Quality, "Teacher contract database," 2022.

Neal, Derek A and William R Johnson, "The role of premarket factors in black-white wage differences," *Journal of Political Economy*, 104 (1996), 869–895.

Oyer, Paul and Scott J Schaefer, "Personnel economics: Hiring and incentives," *Handbook of Labor Economics*, 4 (2011), 1769–1823.

Papay, John P and Matthew A Kraft, "The productivity costs of inefficient hiring practices: Evidence from late teacher hiring," *Journal of Policy Analysis and Management*, 35 (2016), 791–817.

Rivkin, Steven G, Eric A Hanushek, and John F Kain, "Teachers, schools, and academic achievement," *Econometrica*, 73 (2005), 417–458.

Roth, Alvin E. and Marilda A. Oliveira Sotomayor, *Two Sided-Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge University Press, 1992.

Sorkin, Isaac, "Ranking firms using revealed preference," *Quarterly Journal of Economics*, 133 (2018), 1331–1393.

Tincani, Michela M, "Teacher Labor Markets, School Vouchers, and Student Cognitive Achievement: Evidence from Chile," *Quantitative Economics*, 12 (2021), 173–216.

Table I: Teacher preference estimates

	Estimate	Standard Error
Constant	4.370	3.803
Commute Time	-0.058	0.001
Commute Time Missing	-1.415	0.033
Value Added	0.075	0.021
St Dev Value Added RC	0.017	0.009
School Characteristics and Interactions		
N Disadv.	-0.034	0.003
Frac. Black	-0.656	0.123
Frac. Hispanic	0.141	0.122
Frac. Above Med. Achiev.	0.350	0.137
Abs Adv x N Disadv.	-0.060	0.033
Abs Adv x Frac. Black	-0.206	1.062
Abs Adv x Frac. Hispanic	1.251	1.128
Abs Adv x Frac. Above Med. Achiev.	-1.181	1.259
Black x Frac. Black	1.431	0.191
Hispanic x Frac. Hispanic	0.697	0.829
St Dev N. Disadv. RC	0.032	0.002
St Dev Frac. Black RC	1.478	0.067
St Dev Frac. Hispanic RC	1.513	0.090
St Dev Frac. Above Med. Achiev. RC	1.749	0.053
Teacher Characteristics		
VA Non-Disadv. Students	-0.583	0.415
VA Disadv. Students	0.373	0.487
In District	-0.039	0.078

	Estimate	Standard Error
Black	0.924	1.423
Hispanic	5.430	4.672
Female	-0.091	0.093
Experience 2-3	0.077	0.137
Experience 4-6	-1.166	0.125
Experience 7+	-1.240	0.112
St Dev Random Effect	1.558	0.038
Chamberlain-Mundlak Device		
N. Disadv. Mean	0.397	0.153
Commute Time Mean	0.011	0.005
Commute Time Missing Mean	0.741	0.190
Value Added Mean	0.107	0.768
Frac. Black Mean	-4.218	3.695
Frac. Hispanic Mean	-13.589	3.879
Frac. Above Med. Ach. Mean	10.079	5.127
Abs Adv x N Disadv. Mean	-0.434	0.805
Abs Adv x Frac. Black Mean	-0.162	17.202
Abs Adv x Frac. Hispanic Mean	3.557	18.942
Abs Adv x Frac. Above Med. Achiev. Mean	1.428	20.823
Black x Frac. Black Mean	-4.718	3.274
Hispanic x Frac. Hispanic Mean	-20.910	19.121
Number of Students Mean	-0.440	0.123
Sample Size: N Applicants	866	
Sample Size: N Obs	128,264	

The two columns of the table report coefficients from the same model. The table shows teacher preference coefficients, estimated using maximum simulated likelihood. We model the probability that a teacher applies to a position where the alternate options are not teaching in the district or keeping the current position. Random coefficients (“RC”) are independent and simulated from the standard normal distribution. We model unobserved teacher-year heterogeneity using a Mundlak (1978) and Chamberlain (1982) device, taking the mean of each covariate across an applicant’s choices. Commute time is measured in minutes, value added is total predicted output. Experience below 2 years is the omitted category.

Table II: Principal valuation estimates

	Estimate	Standard Error
Characteristics		
Constant	-5.553	0.529
St Dev Random Effect	1.398	0.021
Title I	0.611	0.682
Value Added	0.095	0.027
Value Added x Title I	0.036	0.036
Experience 2-3	0.360	0.131
Experience 2-3 x Title I	-0.043	0.167
Experience 4-6	0.182	0.119
Experience 4-6 x Title I	0.080	0.162
Experience 7+	0.037	0.095
Experience 7+ x Title I	-0.315	0.127
Experience Missing	-0.356	0.068
Experience Missing x Title I	0.437	0.092
Masters	0.055	0.112
Masters x Title I	0.258	0.142
Black	-0.972	0.235
Black x Title I	1.773	0.475
Black x Frac. Black	0.646	0.280
Black x Frac. Black x Title I	-0.512	0.532
Hispanic	-0.651	0.456
Hispanic x Title I	0.502	0.566
Hispanic x Frac. Hispanic	2.277	2.230
Hispanic x Frac. Hispanic x Title I	-1.773	2.364

	Estimate	Standard Error
Female	-0.009	0.107
Female x Title I	0.071	0.134
Gender Missing	0.854	0.480
Gender Missing x Title I	-0.576	0.647
Race Missing	-0.454	0.227
Race Missing x Title I	0.330	0.270
VA Missing	0.488	0.090
VA Missing x Title I	-0.201	0.126
Praxis	0.169	0.054
Praxis x Title I	0.007	0.068
Praxis Missing	-0.139	0.066
Praxis Missing x Title I	0.121	0.083
Grad Deg	0.157	0.069
Grad Deg x Title I	-0.234	0.088
Grad Deg Missing	-0.138	0.731
Grad Deg Missing x Title I	-0.415	0.834
Certified	0.998	0.678
Certified x Title I	-1.015	0.811
Certified Missing	0.244	0.671
Certified Missing x Title I	-0.792	0.801
Licensed	0.955	0.429
Licensed x Title I	0.462	0.457
Sample Size: N Positions	1,824	
Sample Size: N Obs	343,161	

The two columns of the table report coefficients from the same model. The table shows principal valuation coefficients, estimated using maximum simulated likelihood. We model the probability that a principal submits a positive outcome (hire, interview, positive rating) for an application. Random effects are simulated from the normal distribution. Experience below 2 years is the omitted category. Value-added is total predicted output.

Table III: Summary statistics for 2015-16, by economic disadvantage

	Focal, Adv	Focal, Disadv	Other, Adv	Other, Disadv
Students				
White (%)	64.61	9.11	75.58	35.09
Black (%)	17.04	51.78	9.54	32.63
Hispanic (%)	6.77	32.58	6.00	23.90
Student performance (level scores)				
Math	0.70	-0.16	0.43	-0.30
Student performance (gain scores)				
Math	0.07	0.07	-0.01	0.00
Teachers				
Experience (% of teachers)				
0 years	4.32	10.99	3.35	4.85
1-2 years	10.45	17.20	6.90	9.80
3-5 years	17.32	19.30	11.21	12.84
6-12 years	29.48	23.01	26.72	26.19
13 or more years	38.43	29.49	51.82	46.32
Graduate degree (%)	45.20	43.34	39.65	37.43
Regular license (%)	94.69	85.60	95.69	93.44
NBPTS certified (%)	16.08	6.82	14.27	9.95
Praxis score	0.37	0.03	0.29	0.13
Age	39	37	41	40
Mean math value-added				
Baseline, econ disadv	0.02	0.02	-0.01	-0.01
Baseline, econ adv	0.01	0.01	-0.01	-0.01
Homogeneous	0.02	0.01	-0.00	-0.01
CFR	0.09	0.07	0.01	0.01
Using school means	0.15	0.13	0.07	0.08
Imputed, econ disadv	-0.03	-0.01	-0.01	-0.01
Imputed, econ adv	0.01	0.01	-0.01	-0.01
Fraction imputed	0.16	0.41	0.14	0.22
Residual	0.02	0.02	-0.03	0.00
Mean behavioral value-added				
Baseline	-0.01	0.01		
Sample size				
Number of students	12,329	22,628	122,903	197,028

The table shows mean student and teacher characteristics in our sample for the 2015-16 school year. We split the sample into whether the student is in our focal district ("Focal") or in the rest of North Carolina ("Other") and whether he or she is economically advantaged ("Adv") or disadvantaged ("Disadv"). Math scores are standardized to have mean 0 and standard deviation 1 at the state-grade-year level. The alternate VA estimators are (a) a homogeneous value-added model with constant effects across student types, (b) using the estimator from Chetty et al. (2014a), and (c) a model that uses school mean characteristics rather than school fixed effects. Imputed value-added predicts value-added based on how observable teacher characteristics predict value-added. "Residual" is the unshrunken value-added for 2015-16, which has no missing values.

Table IV: Current allocation, alternative policies, and first-best

Description	VA disadv.	VA adv.	mean VA
Status quo	-0.026	-0.040	-0.031
Noisy hiring	-0.025	-0.037	-0.029
School propose	-0.026	-0.040	-0.031
All options (timing)	-0.030	-0.033	-0.031
Principals rank by VA	-0.058	-0.003	-0.040
Teachers rank by N disadv.	-0.026	-0.041	-0.031
Previous two changes	0.032	-0.101	-0.013
First best	0.036	-0.101	-0.010

This table displays numbers corresponding to the allocations plotted in Figure IV, as well as the overall achievement per student. The status quo uses teacher and principal estimated preferences and restricted choice sets, and solves for the teacher proposing stable allocation. Noisy hiring maintains the status quo but replaces the estimated principal behavior with a random valuation of applicants. School propose takes the status quo and instead uses the school proposing stable allocation. All options relaxes the timing constraint in the status quo. Teachers rank by N disadvantage changes the teacher preferences in the status quo. Principals rank by VA changes the principal preferences in the status quo. Previous two changes takes the status quo and replaces the teacher preferences with teachers ranking on the number of disadvantaged students and principals ranking by value added. The first-best is the allocation where the planner maximizes the achievement of disadvantaged students. We report averages over 400 simulations.

Table V: Robustness: disadvantaged achievement

	Status quo	All options	Principal Max VA	Teach Max N Dis	Previous Two	First best
Baseline	-0.026	-0.030	-0.058	-0.026	0.032	0.036
1. Hold class sizes constant: baseline uses class size						
Constant class size	-0.027	-0.030	-0.049	-0.027	0.017	0.020
Constant class size (CFR)	-0.026	-0.029	-0.051	-0.034	0.019	0.022
2. Vary choice set construction for teachers						
7 day buffer	-0.028	-0.031	-0.062	-0.024	0.032	0.036
First day choice sets only	-0.030	-0.031	-0.046	-0.026	0.015	0.036
Drop single app. teachers	-0.025	-0.029	-0.055	-0.025	0.030	0.035
Donut	-0.034	-0.036	-0.062	-0.032	0.024	0.031
Drop 20 percent of apps.	-0.027	-0.031	-0.060	-0.026	0.032	0.036
3. Estimate principal preferences using rank order logit: baseline is binary logit						
All data	-0.027	-0.030	-0.058	-0.031	0.032	0.036
Active choices	-0.024	-0.026	-0.059	-0.036	0.032	0.036
Hire outcome only	-0.025	-0.028	-0.058	-0.027	0.032	0.036
4. Vary window in which we estimate principal preferences: baseline is all applications						
W/in 2 weeks of hire	-0.026	-0.030	-0.059	-0.026	0.032	0.036
First half	-0.027	-0.029	-0.058	-0.025	0.032	0.036
Second half	-0.025	-0.027	-0.058	-0.031	0.032	0.036
5. Vary student type split: baseline is economic disadvantage						
Achievement	-0.026	-0.031	-0.055	-0.025	0.028	0.034
Race	-0.026	-0.029	-0.054	-0.027	0.024	0.029
6. Alternative value-added models						
Homogenous	-0.026	-0.032	-0.052	-0.024	0.032	0.044
Using school means	-0.026	-0.031	-0.055	-0.033	0.037	0.050
CFR	-0.026	-0.028	-0.061	-0.034	0.035	0.044
7. Allow for correlated random coefficients in teacher preferences						
Corr. R.C.	-0.026	-0.029	-0.061	-0.026	0.032	0.036
8. Vary teacher preference specification to use binary logit						
No REs or FEs	-0.026	-0.028	-0.049	-0.026	0.032	0.036
Teacher REs, School FEs	-0.026	-0.028	-0.043	-0.026	0.032	0.036
Teacher FEs, School FEs	-0.027	-0.028	-0.046	-0.026	0.032	0.036
9. Omit value-added from teacher preferences						
No VA	-0.027	-0.030	-0.061	-0.026	0.032	0.036
10. Multiply value-added coefficients by 10 in principal model						
Multiply by 10	-0.029	-0.034	-0.059	-0.021	0.032	0.036
11. Efficiency objective: outcome is mean achievement						
	-0.026	-0.027	-0.035	-0.006	-0.000	0.011
12. Impute value-added for teachers without value added						
Balanced	-0.026	-0.029	-0.039	-0.041	-0.015	0.019
Teachers long	-0.026	-0.022	-0.019	-0.048	-0.019	0.031
13. The role of timing						
One period, period 1	-0.008	N/A	N/A	N/A	N/A	N/A
One period, period 2	-0.018	N/A	N/A	N/A	N/A	N/A
One period, period 3	-0.066	N/A	N/A	N/A	N/A	N/A
Three periods, overall	-0.027	N/A	N/A	N/A	N/A	N/A

The table shows robustness checks for our main results. Each of the robustness checks have been explained and motivated throughout the text and prior footnotes. The columns show the value-added of teachers assigned to disadvantaged students. The status quo uses teacher and principal estimated preferences and restricted choice sets, and solves for the teacher proposing stable allocation. All options relaxes the timing constraint in the status quo. Teachers rank by N disadvantage changes the teacher preferences in the status quo. Principals rank by VA changes the principal preferences in the status quo. Previous two changes takes the status quo and replaces the teacher preferences with teachers ranking on the number of disadvantaged students and principals ranking by value added. The first-best is the allocation where the planner maximizes the achievement of disadvantaged students. For comparability, we normalize the status quo outcome to be the same as in the baseline for rows where the value-added model, population or outcome is different: Panel 1 (row 2), Panels 5, 6, 11, and 12. We report averages over 400 simulations.

Figure I: Wait time to apply to vacancies

(a) Stock of vacancies

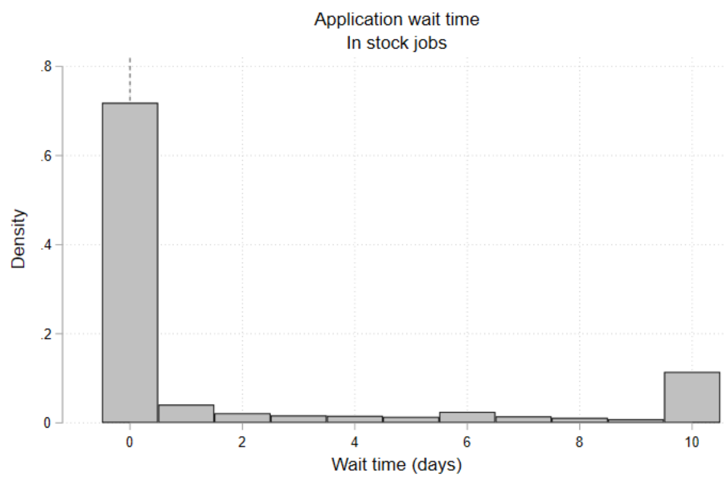

(b) Flow of vacancies

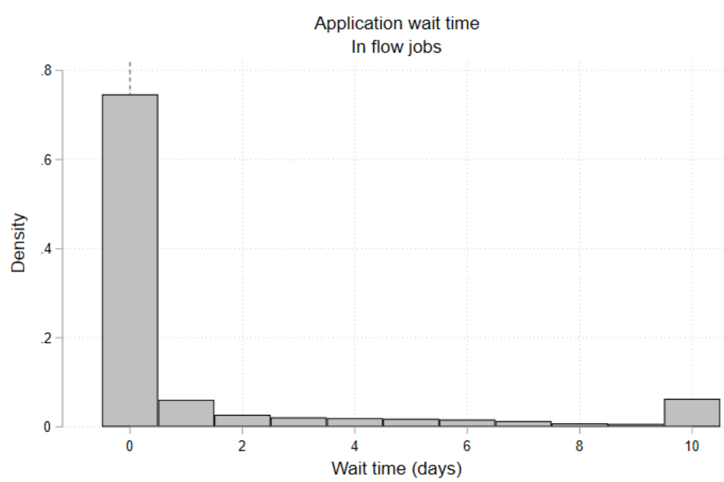

The figures show the wait time for applicants to apply to vacancies. In Panel A, we look at vacancies that were "in stock" (already posted) on the day the teacher first applied on the platform. We plot the "leave one out" wait time, where we omit one job the teacher applied to on the first day. In Panel B we look at the wait time to apply to vacancies that were posted after the teacher first applied on the platform. We measure wait time as the time from when the teacher first applied to another job (once the focal position is posted) until they apply to the posted job. The final category corresponds to waiting at least 10 days. The median wait time is zero in both figures.

Figure II: Bivariate preference relationships

(a) Teachers: commute time

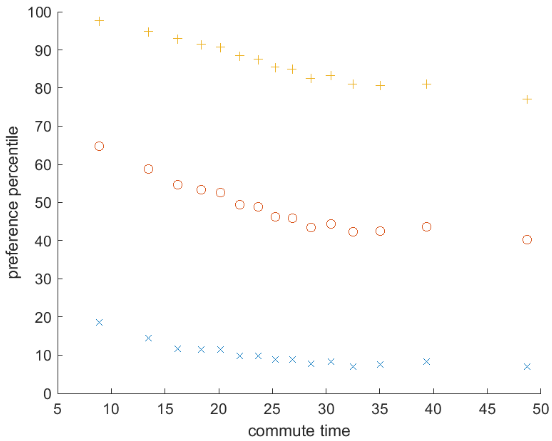

(b) Teachers: number disadvantaged

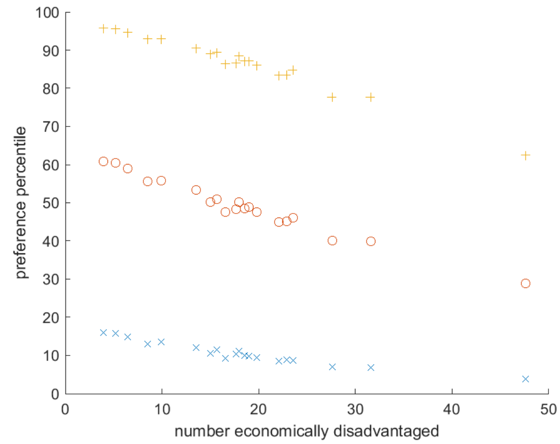

(c) Teachers: output

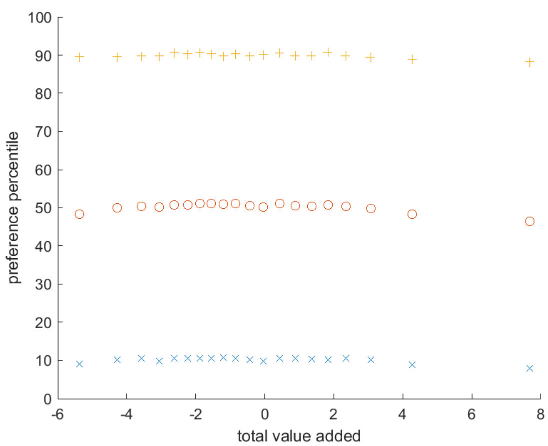

(d) Principals: output

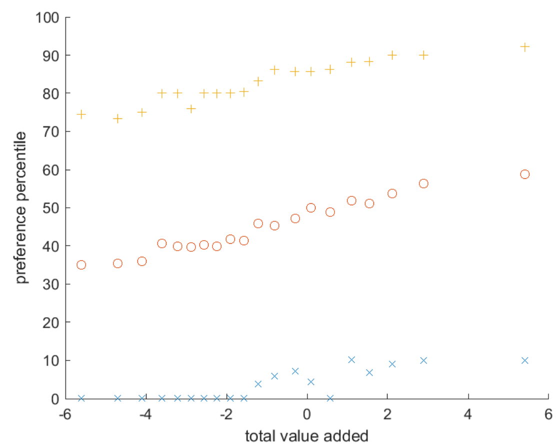

This figure shows binscatters of bivariate relationships between characteristics and preferences. The middle set of points (red circle) is the mean percentile, while the top (orange cross) and bottom (blue x) sets of points are the pointwise 10th and 90th percentiles, respectively. In Panels (a)-(c), we show the bivariate relationship between characteristics in the teacher preference model and how teachers rank positions by estimating each teacher's ranking over positions and ordering positions from a teacher's most preferred (100) to least preferred (0). In Panel (d), we estimate show the bivariate relationship between characteristics in the principal model and principal rankings. We estimate each principal's ranking over teachers and order teachers from a principal's most preferred (100) to least preferred (0).

Figure III: Model fit

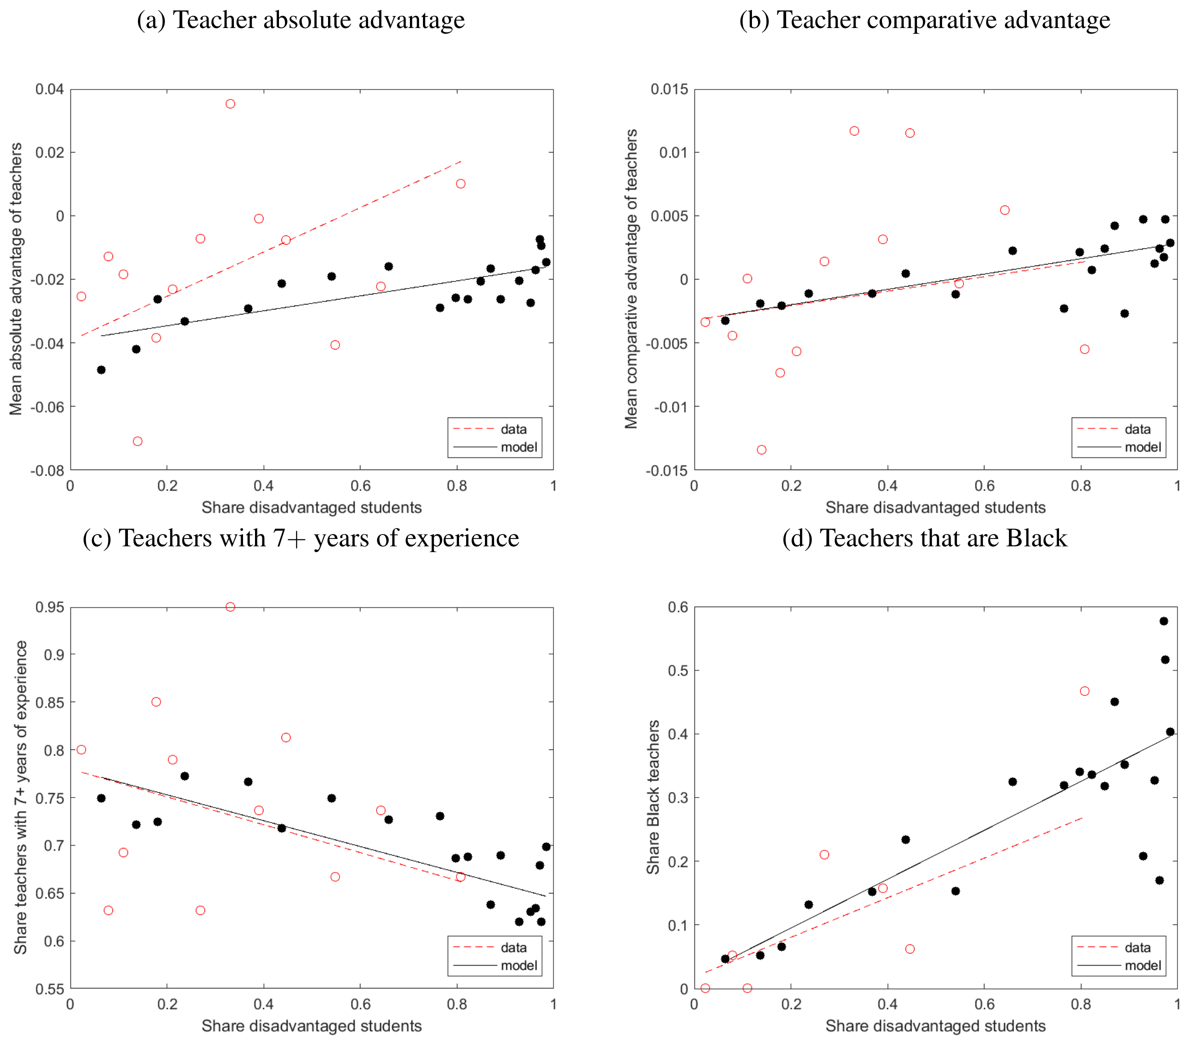

(a) Teacher absolute advantage

(b) Teacher comparative advantage

(c) Teachers with 7+ years of experience

(d) Teachers that are Black

This figure compares the allocations implied by the model to the allocations we observe in the data. The solid line presents the fitted values and the dots represent the binscatter. 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. Positions are sorted on the x-axis by the share of disadvantaged students in the school. The intercepts are normalized to be equal. In panel (a), absolute advantage is the average of the teacher's ability with advantaged and disadvantaged students, weighted by the share of these students in the district as a whole. In panel (b), comparative advantage is the difference between value-added with disadvantaged and advantaged students. In panel (c), the outcome is the share of teachers with 7 or more years of experience in the state of the North Carolina. In panel (d), the outcome is the share of teachers at the school that are Black.

Figure IV: Current allocation, alternative policies, and first-best

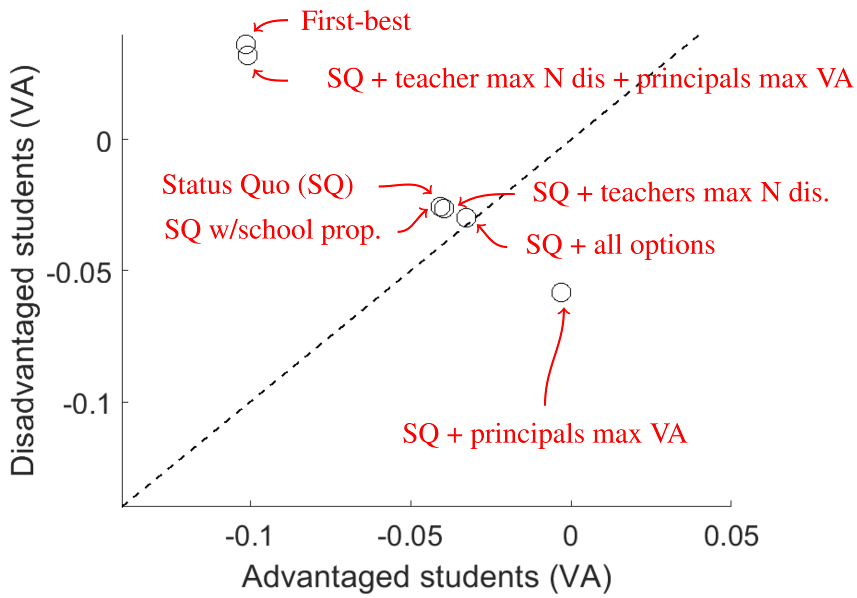

This figure simulates the trade-off between student achievement for economically advantaged and disadvantaged students. The axes refer to per student achievement. The status quo uses teacher and principal estimated preferences and restricted choice sets, and solves for the teacher proposing stable allocation. The status quo with school proposing replaces the teacher proposing deferred acceptance algorithm with school proposing; this point is the only one in the figure that uses school proposing deferred acceptance. The status quo plus all options relaxes the timing restrictions and allows teachers to match with any position; this point is the only one in the figure that relaxes timing constraints. The status quo plus teachers max N disadvantaged replaces the estimated teacher preferences with the assumption that teachers seek to maximize the number of disadvantaged students they teach. The status quo plus principals maximize value-added replaces the estimated principal behavior with the assumption that principals seek to hire teachers to maximize the achievement of their students. The status quo plus teachers maximize the number of disadvantaged and principals maximize value-added replaces estimated with teacher and principal preferences with these assumptions. The first-best is the allocation where the planner maximizes the achievement of disadvantaged students. The Figure plots averages over 400 simulations.

Figure V: Teacher bonus schemes

(a) Principals' estimated preferences   (b) Principals value output

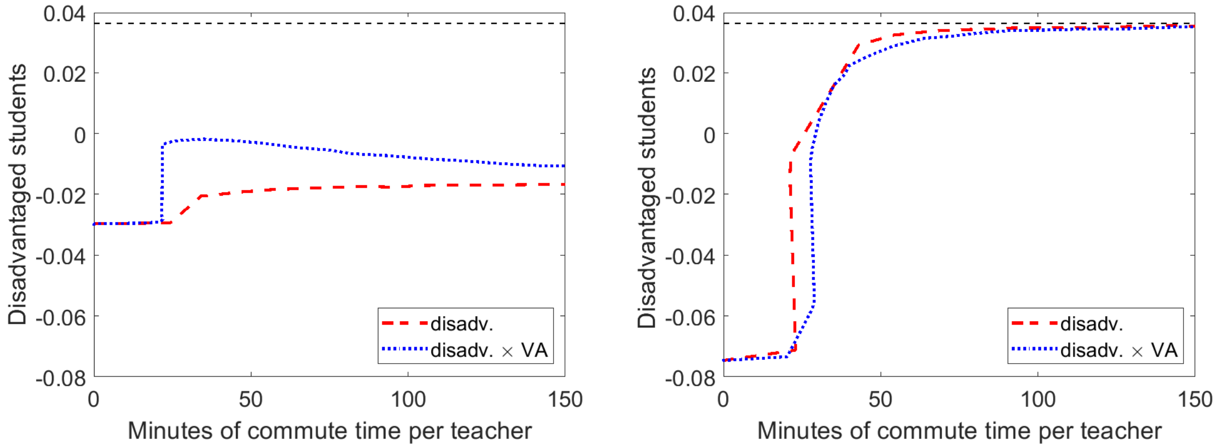

This figure shows the effect of teacher bonus schemes on the achievement of disadvantaged students. The y-axis is per student achievement. The x-axis shows the cost of the policy per teacher, which we express in minutes of commute time per teacher. The y-axis shows the benefits in terms of achievement per disadvantaged student. We consider two policies: subsidizing the position based on the number of disadvantaged students in the position, and subsidizing the position based on number disadvantaged interacted with the teacher's absolute advantage. In the left panel, we take as the baseline allocation the status quo without timing constraints. In the right panel, we take as the baseline allocation one where principals maximize output without timing constraints. The horizontal dashed lines show the output in the first-best.