

Information Frictions and the Labor Market for Public School Teachers*

Mark Colas

Chao Fu

University of Oregon

University of Wisconsin and NBER

March 24, 2026

Abstract

Information frictions between current and prospective employers can distort worker mobility, market entry, and the allocation of talent across firms. Reducing these frictions may improve matching efficiency but can also intensify sorting and lead to equity concerns in sectors providing essential services. We study these trade-offs in the decentralized labor market for public school teachers. We develop and estimate an equilibrium model in which schools make hiring offers under asymmetric information about teacher quality and teachers choose among offers and an outside option. We take the model to administrative data from the Houston Independent School District, which introduced a reform that made multidimensional measures of teacher quality observable to all schools. Counterfactual simulations reveal that (1) making district teachers' quality observable to all district schools improves average teacher quality at the district level and in the top and bottom quartiles of schools ranked by student performance, while decreasing it in other schools; (2) budget-neutral bonus programs that incentivize high-quality teachers to teach in low-performing schools can increase overall teacher quality and reduce cross-school inequalities; and (3) these programs are more effective in markets with cross-employer information symmetry.

*We thank John Stromme for excellent research assistance. We thank the staff at the Houston Independent School District and the Houston Education Research Consortium for their help with data access. We thank Chris Taber, Ken Wolpin, participants at the Econometric Society Conference in Dynamic Structural Econometrics, Stanford Institute of Theoretical Economics, Chinese Economic Association Conference, China Labor Economists Forum, Economics of Education Workshop at TSE, and seminar participants at ASU, Boston College, Boston U., Carnegie Mellon, Chicago Booth, Notre Dame, Tulane, U Minnesota, UNC, UPenn, U Nebraska, U Rochester, U Wisconsin, and Vanderbilt for helpful comments. This work benefited from access to the University of Oregon's high-performance computing cluster, Talapas. This research is supported by NSF Grant SES-1947674. The views expressed in this paper are those of the authors and not necessarily those of the funders or persons named here. Contacts: Mark Colas, University of Oregon, mcolas@uoregon.edu; Chao Fu, University of Wisconsin, cfu@ssc.wisc.edu.

1 Introduction

In many labor markets, employers learn about worker productivity by observing workers on the job, while outside firms observe only noisy signals. This informational asymmetry can generate a winner’s curse problem: firms may hesitate to hire workers from other employers because mobility itself can signal low quality, thereby obstructing efficient labor allocation (e.g., Waldman, 1984; Greenwald, 1986; Chang and Wang, 1996; Waldman and Zax, 2016). Reducing information frictions can therefore improve the sorting of workers across firms and increase total output.

In essential-services sectors, such as healthcare, public safety, and education, the implications of information frictions are more complex. In a typical labor market, the benchmark for optimal worker-firm allocation is total output maximization; how that output is distributed across firms is second-order. By contrast, in labor markets for essential-service workers, both aggregate efficiency and the distribution of quality services across consumers from different backgrounds are central concerns.¹ Reducing information frictions therefore involve important equity–efficiency tradeoffs. On the one hand, reducing information frictions makes it easier for productive workers to move to more desirable jobs, potentially exacerbating inequality if such jobs are concentrated in advantaged communities. On the other hand, greater transparency improves upward job mobility and can attract higher-quality workers into the market, raising overall quality.

We study these issues in the labor market for public school teachers. This setting is not only important—effective teachers play an instrumental role in students’ long-run outcomes and intergenerational skill transmission (e.g. Chetty et al., 2014; Hanushek and Rivkin, 2010; Hanushek, 2011; Jackson, 2018; Abbott et al., 2019)—but also particularly well suited for studying information frictions for four reasons. First, teacher effectiveness is only weakly correlated with observable characteristics such as education and experience, making it difficult to identify effective teachers at the time of hire. Second, schools learn about teacher quality over time through classroom observation and student performance, generating scope for asymmetric information between current and prospective employers.² Third, as is common in the public sector, teacher compensation is often rigid and varies little across schools, making the relative attractiveness of a school largely dependent on non-pecuniary factors, such as the students it serves. Fourth, distributional concerns are first-order in this setting. If

¹Worker welfare considerations arise in all labor markets but are not the focus of this paper. Throughout, we focus on the quality of workers in the market overall (efficiency) and its distribution across localities (equity).

²For studies on the difficulty of predicting teacher effectiveness, see, e.g., Staiger and Rockoff (2010) and Rivkin et al. (2005); for evidence that schools learn about the quality of their teachers, see, e.g., Jacob and Lefgren (2008), Chingos and West (2011), Rockoff et al. (2012) and Bates (2020)

teachers prefer schools with fewer disadvantaged students, reducing information asymmetries can shift highly effective teachers toward more advantaged schools and widen disparities.

Our analysis exploits an information intervention in the Houston Independent School District (HISD), the eighth-largest school district in the United States. In the 2010–11 school year, HISD introduced the Effective Teachers Initiative (ETI), centered on a rigorous teacher evaluation system. By combining in-person observations (e.g., how a teacher interacts with students, colleagues, and parents) with estimates of teachers’ value-added (VA) to test scores, ETI provided a multidimensional assessment of each teacher’s quality and made this information accessible to all schools in the district. Notably, some dimensions of a teacher’s quality as measured by ETI (e.g., quality revealed via in-person observations) could be learned by her employer through daily observations but were likely unobservable to other schools before ETI. Therefore, ETI largely broke the information asymmetry between a teacher’s employer and other schools in the district.

Using HISD as a platform, we develop an equilibrium model of the labor market for public school teachers, accounting for the role of information frictions. In the model, the teacher distribution, consisting of both market incumbents and potential entrants, is characterized by publicly observable attributes and by quality measures that are fully observed by the teacher and her current employer but only partly known to prospective employers. A teacher cares about her salary and a school’s characteristics, and chooses her most preferred option among schools offering her a job or the outside option, subject to moving costs. Each school aims to fill its capacity with its most preferred teachers by extending job offers to teachers it wishes to hire, given its beliefs about (1) each teacher’s probability of accepting their job offer and (2) the expected quality of a teacher *conditional* on her acceptance — two equilibrium objects resulting from decisions by all schools and teachers.

The model highlights how a transparent teacher evaluation program can affect the equilibrium teacher-school sorting. When the incumbent employer is the only school fully informed about its teachers’ quality, schools may be apprehensive about hiring a teacher from another school in fear of the winner’s curse: Given the cost involved in changing jobs, if a teacher originally working in school s accepts an offer from another school, one possible explanation is that she has been found to be ineffective and dismissed by s .³ As a result, teachers may get few offers from other schools and therefore rarely switch jobs. Programs such as ETI evaluate the effectiveness of market incumbent teachers and provide schools with credible information about teachers working in other schools. This breaks the information asymmetry

³In our model, the market cannot observe whether a teacher is laid off or quits voluntarily. This is in contrast to the setup in, e.g., Gibbons and Katz (1991), where layoff is public information and serves as a signal of the worker’s (low) ability.

and affects both teacher-school sorting within the market and teachers' entry and exit.

We estimate our model using the post-ETI data and use the pre-ETI data to validate the estimated model. The model replicates the data patterns well in both periods. Our estimates suggest that all else equal, teachers prefer schools with higher fractions of high-performing students and that this preference is stronger among higher-quality teachers. Meanwhile, relative to experience and education, schools significantly value teachers' quality, especially teachers' VA. Moreover, preferences for higher-VA teachers are stronger among schools serving higher fractions of high-performing students. That is, preferences on *both* sides of the market exhibit complementarity. Without frictions, these preferences would directly lead to assortative matching between teachers and schools, leaving lower-performing schools behind.⁴ The sorting can be less assortative with information frictions, as schools are reluctant to hire teachers from other schools in fear of the winner's curse.

We use the estimated model to quantify the role of information frictions. Holding *all other initial conditions fixed* at their post-ETI levels, we simulate the equilibrium under a counterfactual asymmetric information environment, where prospective employers do not observe teachers' quality. Contrasting this equilibrium with the (baseline) post-ETI equilibrium, we find that removing information asymmetry between current and prospective employers improves district-level teacher quality by incentivizing higher-quality teachers to enter the district. Moreover, dividing schools into four groups by a school's fraction of high-achieving students, we find that average teacher quality improves in the top and the bottom 25% of schools, while decreasing that in other schools. These changes result from the interplay of information's intensive- and extensive-margin impacts on the market. At the intensive margin, symmetric information benefits top schools because it makes it easier for these schools, which are preferred by most teachers, to recruit effective teachers from other schools. At the extensive margin, symmetric information induces higher-quality teachers to enter the market. These entrants, although less effective than incumbent teachers on average, are more effective than many teachers working in low-performing schools; therefore, these schools benefit most from this extensive-margin effect as they can replace ineffective incumbents with these entrants.

With the goal of providing students with equal access to high-quality teachers, many school districts have implemented programs, such as teacher bonuses, aimed at improving teacher quality in low-performing schools. Given the role of information frictions in shaping equilibrium allocations, it is natural to hypothesize that the impacts of such programs may vary with the market's information environment, although it is not clear how. To shed

⁴Throughout the paper, we use "low-performing schools" and "schools with more low-performing students" interchangeably.

light on this, we introduce counterfactual teacher bonus programs to incentivize high-quality teachers to teach in low-performing schools. We evaluate these programs in a market with information asymmetry between employers and an otherwise-identical market without such information asymmetry. We find that these bonus programs can both increase overall teacher quality on the market and reduce the teacher-quality gap between high- and low-performing schools, even though we hold the total per-teacher compensation fixed for the district by reducing teachers' base salaries. More importantly, these programs are more effective when implemented in a market with symmetric information between employers.

Beyond the teacher labor market, our findings illustrate a broader lesson: Information frictions shape not only the sorting of workers across jobs but also workers' decisions to participate in the market at all. Both margins affect efficiency and the distribution of talent. In sectors where equitable access to high-quality services is a key concern, such as healthcare, public safety, and education, policies that affect information should account for both margins. Moreover, our findings suggest a fruitful direction for policy design: exploiting the interaction between information and compensation policies to improve both efficiency and equity.

Related Literature Our paper contributes to the broad literature on the labor market of teachers, and in particular, the strand that studies this market via the lens of structural models. Among these studies, a large subset focuses on the supply side. For example, Dolton and Klaauw (1999) study teachers' decision to leave the profession; Stinebrickner (2001a), Stinebrickner (2001b), Wiswall (2007), and Lang and Palacios (2018) model individuals' dynamic choices between teaching and non-teaching options; Behrman et al. (2016) further breaks down the teaching option into teaching in one of three types of schools; Boyd et al. (2005) and Scafidi et al. (2007) find that teachers prefer schools with fewer low-achieving and minority students.

A smaller subset of work considers both sides of the market. Tincani (2021) estimates an equilibrium model where a representative private school sets teacher wages and tuition, workers choose among teaching in public schools (which are passive in the model), teaching in the private school, and non-teaching options, and households choose between public and private schools. Focusing on the market of public school teachers as we do, Boyd et al. (2013) estimate a two-sided matching model, assuming that the observed teacher-school matches are stable; using data from Peru, Bobba et al. (2024) and Ederer (2023) study teacher-school sorting in a centralized application-assignment environment. In the U.S. setting, Bates et al. (2025) use information on teachers' application decisions and schools' ranking of applicants to study teacher-school allocation within a district, allowing teacher value-added to differ between advantaged and disadvantaged students; Laverde et al. (2025) model student achievement and teacher and school decisions, accounting for the potential selection

of teachers with different comparative advantages in test score gains. Both these papers are interested in the potential gains in students test scores by reassigning teachers across schools. Using a market equilibrium model, Biasi et al. (2024) study how public school districts use wage and hiring strategies to compete for better teachers in a setting where districts have control over teacher pay.

Our paper complements these studies well. First, we explicitly account for information frictions and examine how policies that reduce information asymmetry may affect the market equilibrium. Second, rather than focusing on a closed market, we examine teacher-school matching both within the market and at the entry-exit margin; this extensive margin is found to have qualitatively and quantitatively important equity-efficiency implications. Third, we study constrained optimal policy designs, taking into account the interactions between schools and teachers in a decentralized market setting. This approach differs from previous studies on centralized markets (e.g., Bobba et al., 2024; Ederer, 2023), as well as studies using centralized allocation algorithms as an elegant way to approximate decentralized market equilibrium (e.g., Bates et al., 2025). Our approach, based on the decentralized nature of the market we study, allows us to design easy-to-implement policies that *induce* more equitable and efficient market equilibrium sorting. Finally, our data contain teacher quality measures that are plausibly valued by schools, including teachers' VA to test scores and, more uniquely, quality revealed via in-person observations. Post ETI, these measures are observed by all schools. This setting allows us to better infer how much schools value these qualities, while in other settings, the estimated school preferences often confound real preferences and information friction, as acknowledged by Bates et al. (2025): An estimated weak preference for teacher quality can arise because schools do not care much about it and/or because they are ill-informed about it.

There is an extensive literature, both theoretical and empirical, on the impact of incomplete information on labor market allocation, with one strand focusing on the impact of asymmetric information between current and prospective employers.⁵ In the context of the teacher labor market, Bates (2020) develops novel tests of cross-school information asymmetry. Using a difference-in-differences approach, he examines how teachers' within-district movement and exits were affected by the adoption of value-added (VA) measures by two North Carolina school districts. He finds direct evidence of asymmetric information between employers. His findings, based on quasi-experimental data variation from a different state,

⁵Examples of theoretical papers include Waldman (1984); Greenwald (1986); Lazear (1986); Milgrom and Oster (1987); Riordan and Staiger (1993); and Laing (1994). Examples of empirical papers include Gibbons and Katz (1991); Chang and Wang (1996); Acemoglu and Pischke (1998); Schönberg (2007); Zhang (2007); Pinkston (2009); Hu and Taber (2011); DeVaro and Waldman (2012); Kahn (2013); Bognanno and Melero (2016); Cassidy et al. (2016); Waldman and Zax (2016); and Wu (2025).

lend external support to our finding that asymmetric information plays a key role in the teacher labor market. Complementing his work, we use a structural estimation approach to recover teacher and school preferences underlying the observed equilibrium outcomes, which we then use to conduct counterfactual policy evaluations. We also study how information structure can affect teacher entry.

Our paper also contributes to the literature on teachers' mobility and educational inequality (see Jacob, 2007; Neal, 2011; Jackson et al., 2014, for reviews). Previous studies have found that better teachers tend to teach in schools with more advantaged students (e.g., Lankford et al., 2002; Ingersoll et al., 2004; Clotfelter et al., 2005; Mansfield, 2015), and that teacher mobility is closely related to schools' student composition (e.g., Hanushek et al., 2004). Kraft et al. (2020) find that teacher evaluation reforms increased the number of new teachers from more competitive colleges but decreased the overall supply of teachers. Morgan et al. (2026) find that performance-based compensation differentials can effectively attract and retain effective teachers in disadvantaged schools. Specifically in the context of HISD, Brehm et al. (2017) and Imberman and Lovenheim (2015) study the effect of ASPIRE; Cullen et al. (2021) analyze the impact of ETI on teacher exit and student outcomes and find that ineffective teachers were more likely to exit the district after ETI.⁶ Building on these papers, we study the interaction of ETI and ASPIRE via the lens of an equilibrium model.

2 Background

The Houston Independent School District (HISD) serves as a community school district for most of the city of Houston and several nearby and insular municipalities in addition to some unincorporated areas. Like most districts in Texas, HISD has two features that distinguish it from many districts in other states: (1) HISD is independent of the city and all other municipal and county jurisdictions;⁷ (2) teachers do not have tenure and most of them are on one-year contracts, leaving schools with significant latitude in hiring and dismissal decisions.

Over the past 20 years, HISD has implemented several policies aimed at improving teacher effectiveness and educational equality, including two major policies that affected the market in our sample period (2006-07 to 2014-15). The first policy is a performance pay scheme (ASPIRE), introduced in 2006-07, which rewards teachers and administrators for raising students' test scores. The second policy is the Effective Teachers Initiative (ETI), introduced

⁶Studies using data from other places in the U.S. have also found that more rigorous teacher evaluation systems increase turnover among lower-performing teachers (Loeb et al., 2015; Steinberg and Sartain, 2015; Rodriguez et al., 2020).

⁷In 2023, after our sample period, the Texas Education Agency assumed direct control of the HISD.

in 2010-11. We describe these two policies briefly here and in detail in Appendix D.

ASPIRE Like most other U.S. school districts, HISD uses a rigid experience-education schedule to pay its teachers. ASPIRE relaxes this pay rigidity by providing teachers with performance-based bonuses. Over the period we study, the generosity and detailed eligibility criteria of ASPIRE rewards varied across years, but its structure is as follows. A teacher-level value-added measure known as the EVAAS score is calculated by the SAS Institute based on student growth over the year on state-level standardized tests. At the end of an academic year, a teacher can receive up to two types of salary bonuses based on EVAAS scores: an individual reward and a campus reward. Individual rewards are given to teachers whose EVAAS scores exceed certain thresholds. Campus rewards are granted to all eligible teachers employed at schools that achieve high school-level EVAAS scores.⁸ When a campus is eligible, a teacher’s eligibility for campus rewards primarily depends on being a salaried (not substitute or associate) teacher and not exceeding a specified limit of instructional days absent. Starting in 2010-11, teachers with very low EVAAS scores (approximately 15-20% of teachers) were disqualified from receiving any ASPIRE awards.

In terms of teacher pay, ASPIRE introduces pay variation both across teachers within the same experience-education group and across schools for individual teachers, with the latter being key to identifying teachers’ pay preferences. In terms of information, a teacher’s EVAAS score is confidential but the list of ASPIRE winners and their reward amounts are publicly announced. This information, combined with the publicly-known ASPIRE formula, could provide prospective employers with some information about a teacher’s VA.

ETI ETI was designed to improve teacher quality through more effective recruitment at the front end, individualized professional development in the middle, and targeted retention and exit on the back end. As detailed in Appendix D.1, the cornerstone of ETI is a teacher evaluation system, which was designed by the district during the 2010-11 school year with inputs from stakeholders and formally approved by the school board in spring of 2011. By the 2012-13 school year, the ETI evaluation had fully developed to include three component scores: 1) instructional practice, concerning how well a teacher plans lessons and instructs students, 2) professional expectations, concerning how well a teacher interacts with colleagues and parents, follows school policies, and participates in professional development, and 3) student performance, measured primarily by EVAAS scores for teachers in our sample. The first two components are based on in-person observations of the teacher in class (e.g., communicates concepts with students) and out of class (e.g., collaborates with

⁸For example, during the 2013-14 school year, at the individual level, a teacher received \$5,000 or \$10,000, depending on which range their EVAAS scores fall in; at the campus level, teachers at schools in the top quintile of school-level EVAAS scores were eligible to receive the campus reward.

colleagues) by at least one certified appraiser. The appraiser, who also provides teachers with individualized feedback, can be the teacher’s supervisor or a person approved by the school board.

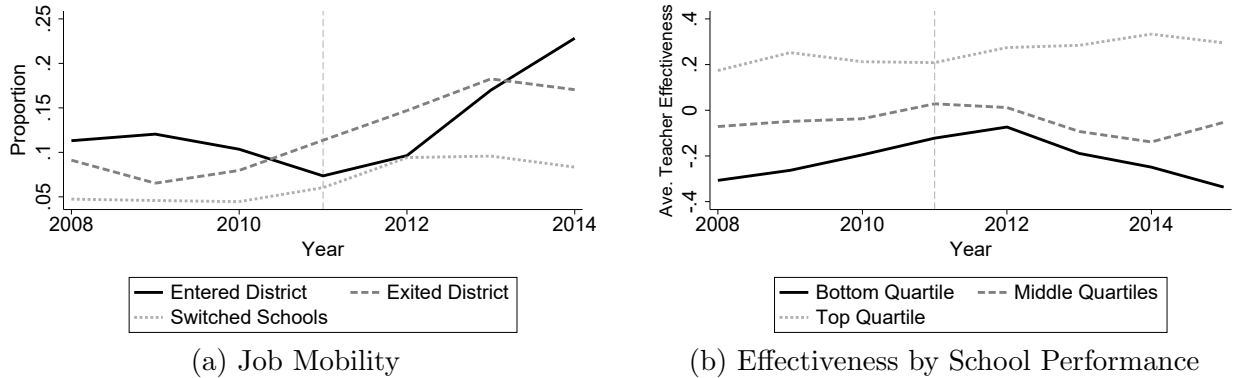
The ETI component scores are stored by the district in each teacher’s personnel file and made accessible to all district schools. In addition, each teacher is entitled to receive a copy of the evaluation. Because ETI scores are recorded at the end of a school year, hiring decisions for 2011-12 and later years are made under the ETI information environment. In earlier years (including 2010-11), a prospective employer’s information was much more limited: For example, they had no information about teacher’s quality reflected by in-person evaluations. That is, ETI significantly changed the information structure in the HISD teachers’ market. In this paper, we focus on this most salient aspect of ETI, while abstracting from its other components (namely, providing teachers with regular feedback on their progress and opportunities for development).

Discussion The institutional features in HISD make it a useful setting for studying how information about teacher quality shapes hiring decisions in decentralized teacher labor markets. On the one hand, the fact that schools’ hiring choices are less constrained by tenure and long-term contracts in Texas allows us to focus more directly on the role of information. On the other hand, while tenure protections are stronger in many other settings, the qualitative mechanisms we study, including how information shapes schools’ hiring behavior, teacher mobility, and policy responses, are not unique to Texas or HISD. Large urban districts commonly rely on school-level hiring and employ probationary or short-term contracts for a substantial share of teachers.⁹ Moreover, the policy interventions in HISD—rewards for high-performing teachers and formal evaluation systems to generate and disseminate information about teacher quality—are either common or gaining popularity.¹⁰ As a result, while the magnitude of the effects we estimate may depend on local institutional details, the qualitative equity-efficiency tradeoffs we highlight and the policy mechanisms we study are likely to generalize to other teacher labor markets, and more generally, labor markets for essential-service workers.

⁹In the U.S., several states have significantly curtailed traditional teacher tenure or replaced it with renewable annual or multi-year contracts (National Education Association, 2023). Fixed-term contracts are also common in the EU. In Spain, Italy, Austria, and Portugal, more than two-thirds of teachers under age 35 are on fixed-term contracts (EACEA, 2021).

¹⁰This applies not only to the U.S., but also Europe, where formal teacher appraisal and support systems are increasingly used to provide feedback, link professional development to evaluation outcomes, and link evaluation to career progression and incentives (European Commission, 2018; EACEA, 2021).

Figure 1: Job Mobility and Average Effectiveness by School Performance



Note: Left: “Exited District” is defined as the fraction of district teachers in the corresponding year who did not work in the district in the following year. “Switched Schools” is defined as the fraction of district teachers in the corresponding year who worked at a different school in the district in the following year. “Entered District” is the number of teachers who entered the district divided by the number of all teachers working in the district in the corresponding year. Right: Teacher effectiveness is measured as average EVAAS score over all years for a given teacher. The dotted line corresponds to the elementary schools with the most students who met standards before 2011, the solid line corresponds to the elementary schools with the least students who met standards before 2011, and the dashed line with the middle quartiles.

3 A First Glance at the Data

As detailed later in the data section, our data consist of linked administrative data sets from HISD at the school and teacher level between 2007 and 2015; we focus on the market for Grades 3-5 math and reading teachers. In this section, we provide a first glance at this market before and after the introduction of ETI in 2010-11. For the rest of the paper, we refer the academic year using the calendar year of the spring semester (e.g. 2011 for 2010-11).

Teacher Job Mobility Figure 1a shows, for each year $t = 2008, \dots, 2014$, three measures of job mobility: (1) the exit rate, i.e., the fraction of district teachers who exited the district in $t + 1$ (the dashed line), (2) the job-switching rate, i.e., the fraction of district teachers who transferred to a different school in the district in $t + 1$ (the dotted line), and (3) the number of period- t entrants as a fraction of all teachers working in this market in t (the solid line). All three types of job mobility became much more frequent after ETI. For example, the fraction of teachers who exited the district dipped in 2009 (concurrent with the great recession), but by 2014 this fraction was nearly twice as high as that in 2008.

Teacher-School Sorting Figure 1b shows a school’s average teacher effectiveness over time for three groups of schools: the top 25% (the dotted line), the middle 50% (the dashed line), and the bottom 25% (the solid line), ranked by the percentage of each school’s Grades 3-5 students who met the state testing standards between 2008 and 2011. Figure 1b uses a time-invariant measure of teacher effectiveness (one’s average EVAAS score across all years observed during our sample period) and a time-invariant measure of school-level achievement

that is not affected by ETI (the percentage of each school’s Grades 3-5 students who met the state testing standards between 2008 and 2011). As such, changes in a school’s average teacher effectiveness reflect changes in its teacher composition.¹¹ Across all years, the average teacher effectiveness was higher in schools with higher (pre-ETI) achievement, suggesting positive assortative matching. Post ETI, average teacher effectiveness increased in the top school group but decreased in the other school groups. As a result, the teacher quality gap between the top and the lower school groups widened after ETI. We find similar patterns if we divide schools by the fraction of minority students or by the fraction of economically disadvantaged students (Appendix Figures B5 and B6): Post ETI, average teacher quality increased (decreased) in schools with fewer (more) minority or disadvantaged students.

Figures 1a and 1b provide some suggestive evidence that, under ETI, teachers’ job mobility increased, and teacher-school sorting became more assortative, leading to potential concerns over educational inequality. To formally assess the role of information in shaping these equilibrium outcomes and, *more importantly*, to examine how we can exploit the interaction of information and other counterfactual educational policies to improve educational equity and efficiency, we build the following equilibrium model.

4 Model

We model the equilibrium in a labor market for public school teachers, where schools make job offers to teachers they want to hire, and both incumbent teachers and potential entrants choose their most preferred feasible option.¹² Without a transparent teacher evaluation system (e.g., ETI), the market is subject to information asymmetry between a teacher’s current employer and prospective employers.

4.1 Primitives

There are S schools in a district/market ($s = 1, \dots, S$). A school is endowed with κ_s teaching positions (capacity), and a vector of characteristics $z_s = \{z_{sk}\}_{k=1}^K$, including the composition of its student body (z_{s1}). There is a distribution of teachers $F(x, q, v, s_0)$, each characterized by a vector x (experience, education), the school she works in at the beginning of the model s_0 ($s_0 \in \{1, \dots, S\}$ for district incumbents and $s_0 = 0$ for potential entrants), and two quality measures: q and v . In particular, q is quantifiable by outputs, as measured by one’s value-added (EVAAS) to students’ achievement; v is reflected in one’s daily teaching and

¹¹In Appendix Figure B4, we use each teacher’s EVAAS score in the year in question and repeat the exercise. That figure shows patterns very similar to, but noisier than, those in Figure 1b.

¹²As is true in most public school districts, HISD follows a rigid schedule of teacher salaries based on experience and education. Accordingly, schools in our model can only choose whom to hire but not how they are paid.

Table 1: Schools' Information Set $\Gamma_s(x, q, v, s_0)$

	District Incumbent Teacher		Potential Entrant
	Current school ($s = s_0$)	Other schools ($s \neq s_0$)	All schools
Pre-ETI	x, q, v, s_0	$x, A(q), s_0$	x, s_0
ETI	x, q, v, s_0	x, q, v, s_0	x, s_0

Column 1 gives the current employer school's information set about a teacher. Column 2 gives other schools' information set about a district incumbent teacher. Column 3 gives all schools' (common) information set about a potential entrant.

professional activities, as measured by ETI in-person evaluations.

The timing of the model is as follows:

1. Schools simultaneously make job offers.
2. Each teacher observes her contemporaneous preference shocks for the outside option and for each school $\epsilon = \{\epsilon_s\}_0^S$ and chooses her most preferred option within her offer set, which always includes the outside option.

4.1.1 Information Structure

School characteristics (z, κ) and teacher characteristics (x, s_0) are public information. A teacher has full information about her own characteristics (x, q, v, s_0) and observes her contemporaneous preference shocks ϵ right before making her decisions. The researcher has information on school characteristics (z, κ) , teacher characteristics (x, q, v, s_0) for all incumbents and entrants, and the distribution of (x, q, v) for potential entrants.

Letting $\Gamma_s(x, q, v, s_0)$ be the set of information available to school s about teacher (x, q, v, s_0) , the first two columns of Table 1 summarize a school's set of information about a district incumbent teacher for her current employer (Column 1) and other schools (Column 2); Column 3 summarizes all schools' (common) information set about a potential entrant. Before ETI, a teacher and her employer s_0 have full information about (q, v) ; other schools have no information about v and some information about q , the latter being governed by a function $A(q)$. Rather than taking a firm stand on the informational contents of $A(q)$, we hypothesize two alternative forms of $A(q)$ based on the policy background (Section 2) and our conversations with HISD; we will use the pre-ETI data to infer which information environment is more plausible in Section 5. After ETI, q and v are observable to all schools. The structure of information about a potential entrant is the same before and after ETI and common across schools: A potential entrant's (q, v) are unobservable to all schools. After a teacher enters the market and teaches for one period, she becomes an incumbent and therefore the structure of information about an incumbent teacher applies.

4.1.2 Preferences

Schools Schools' preferences over teacher characteristics are allowed to differ by school characteristics z_s , given by

$$B(x, q, v, z_s), \quad (1)$$

where $B(\cdot)$ is weakly increasing in x , q , and v .¹³

Teachers For a teacher with (x, v, q, s_0) , the utility of teaching in school s , net of moving costs, is given by

$$U_s(x, q, v, s_0) = W(x, q, v, z_s) + z_s \alpha_0 + z_{s1}(\alpha_1 q + \alpha_2 v) - I(s \neq s_0)(d_1 + d_2 x_1), \quad (2)$$

where $W(\cdot)$ is the pay function set by the district.¹⁴ The vector α_0 measures teachers' preferences for school characteristics; α_1 and α_2 capture potentially different preferences for student composition (z_{s1}) among teachers with different q and v , respectively; d_1 and d_2 capture the cost of moving and how it may vary by experience.

The value of the outside option, net of moving costs, is given by

$$U_0(x, q, v, s_0) = u_0(x, q, v, s_0) - I(0 \neq s_0)(d_1 + d_2 x_1). \quad (3)$$

The value of the outside option, $u_0(\cdot)$, varies by teacher characteristics (x, q, v) and by whether the teacher is a district incumbent or potential entrant.

Remark 1 *The ASPIRE component of $W(\cdot)$ includes both individual and campus awards. The campus award structure provides pay variation across schools for the same teacher, which is critical for identifying teachers' preferences with respect to pay. However, we need to make an assumption as to how teachers predict which campuses are eligible for campus awards. If we assumed that teachers are fully rational, we would have to deal with a potential multiple equilibrium problem because a teacher's pay in school s depends on which other teachers also accept offers from s . To avoid this complication, we assume that a teacher uses schools' previous years' campus award eligibility (a component of z_s) in forming her belief about how much she will be paid in each school. Empirically, past eligibility predicts future eligibility very well because the set of eligible campuses was quite stable in our sample period.*

¹³We assume that conditional on (x, q, v) , schools do not care about a teacher's origin s_0 . Given that we only observe accepted offers, it is hard to distinguish between teachers' moving costs and schools' direct preference for its current employees. We therefore assume the latter away.

¹⁴In the status quo, $W(\cdot)$ includes the teacher salary (a rigid function of x) and the ASPIRE awards, the latter vary by q and z_s but not by v . In our counterfactual policy analysis, we will explore alternative bonus formulae.

Similarly, when calculating a teacher's individual-level ASPIRE award, we assume that the EVAAS score cutoffs for awards are fixed.

4.2 Teachers' Problem

4.2.1 District Incumbents

Observing her preference shocks $\epsilon = \{\epsilon_s\}_{s=0}^S$ and job offers, a district incumbent teacher (x, q, v, s_0) solves the following problem

$$V_1(x, q, v, s_0, \epsilon) = \max_{s: o_s(x, q, v, s_0)=1} \{U_s(x, q, v, s_0) + \epsilon_s\}, \quad (4)$$

where $o_s(x, q, v, s_0)$ is the indicator of whether or not teacher (x, q, v, s_0) has an offer from school $s = 1, \dots, S$; $o_0(\cdot) = 1$, i.e., the outside option is always available. We assume that preference shock ϵ 's are i.i.d., drawn from a Type-1 extreme value distribution with mean 0 and scale parameter σ_ϵ .

4.2.2 Potential Entrants

A potential entrant's quality (q, v) is unobservable to district schools at the time of hiring. However, once a teacher has worked in the district, her employer fully observes her quality. Depending on the market's informational structure, this information may also become available to other schools. Consequently, while an entrant's job offers do not depend on (q, v) (initially unknown to schools), once becoming an incumbent, her opportunities depend on both her employer's identity and her revealed quality (q, v) (Equation (4)).

To capture these dynamics, we model the payoff of working in school s for an entrant $(x, q, v, s_0 = 0)$ as the sum of her current payoff and her expected future payoff:

$$U_s(x, q, v, 0) + \beta E[V_1(x, q, v, s_0 = s, \epsilon')]. \quad (5)$$

In (5), $U_s(\cdot)$ is the utility of working in school s defined in (2); $\beta E[V_1(x, q, v, s_0 = s, \epsilon')]$ is the discounted (at the rate β) expected value of being an incumbent in s , where $V_1(\cdot)$ is defined in (4) and the expectation is taken over next-period's preference shocks.

With this payoff structure, we define a potential entrant's problem as

$$V_0(x, q, v, 0, \epsilon) = \max \left\{ \begin{array}{l} \max_{s: o_s(x, q, v, 0)=1} \{U_s(x, q, v, 0) + \beta E[V_1(x, q, v, s_0 = s, \epsilon')]\} + \epsilon_s, \\ U_0(x, v, q_0, 0) + \epsilon_0 \end{array} \right\}, \quad (6)$$

where ϵ is one's current preference shocks.

Discussion: We model teachers’ problems differently for incumbents and potential entrants due to the following consideration.¹⁵ To assess the role of information structure for the teacher labor market, we need to allow teachers’ self-selection into the market to depend on the information structure; our specification of a potential entrant’s problem allows us to capture this dependence in a parsimonious fashion: As we discuss further when we present the schools’ problem, the effect of (q, v) on an incumbent’s job offer set, and thus $V_1(\cdot)$, varies with the market’s informational environment. As such, (6) implies that the pool of teachers who self-select to enter the market will depend on the information structure.¹⁶

4.3 Schools’ Problem

A school seeks to fill its capacity with its most preferred teachers it can attract. Letting $O_s(\Gamma_s(\cdot)) \in [0, 1]$ be the offer decision rule of school s given its information $\{\Gamma_s(x, q, v, s_0)\}$ specified in Table 1, $O_s(\Gamma_s(\cdot))$ solves the following problem (detailed in Appendix B)

$$\begin{aligned} \max_{O_s(\Gamma_s(\cdot)) \in [0, 1]} & \left\{ \int O_s(\Gamma_s(x, q, v, s_0)) h_s(\Gamma_s(x, q, v, s_0)) \bar{B}_s(\Gamma_s(x, q, v, s_0)) dF(x, q, v, s_0) \right\} \quad (7) \\ \text{s.t.} & \int O_s(\Gamma_s(x, q, v, s_0)) h_s(\Gamma_s(x, q, v, s_0)) dF(x, q, v, s_0) \leq \kappa_s. \end{aligned}$$

From the perspective of school s , given a teacher about whom s observes $\Gamma_s(x, q, v, s_0)$, $h_s(\Gamma_s(x, q, v, s_0))$ is the probability that she will accept an offer from s , while $\bar{B}_s(\Gamma_s(x, q, v, s_0))$ is the expected value of $B(x, q, v, z_s)$ conditional on the teacher accepting the offer. Subject to its capacity constraint, a school’s offer decision rule $O_s(\Gamma_s(\cdot))$ maximizes the total expected value of teachers it hires, where the expectation is based on its information set $\Gamma_s(\cdot)$. This decision rule governs $\{o_s(x, q, v, s_0)\}_{(x, q, v, s_0)}$, i.e., whether or not each teacher receives an offer from s . All teachers that look identical to school s are treated equally.

As detailed in Appendix A.1, both $h_s(\cdot)$ and $\bar{B}_s(\cdot)$ are derived from teachers’ acceptance probabilities; we denote by $P_s(x, q, v, s_0)$ the probability that a teacher characterized by (x, q, v, s_0) will accept an offer from s . From equations (4) and (6), it is clear that $P_s(\cdot)$ depends not only on teachers’ preferences but also on other schools’ offer decisions $\{O_{s'}(\cdot)\}_{s' \neq s}$: All else being equal, a teacher is more likely to accept an offer from s if she has fewer options, and in particular, if she is fired by her employer. When a teacher’s employer is the only school fully informed of her quality, schools may be apprehensive about hiring teachers from other schools in fear of the “winner’s curse,” potentially limiting her job prospects and

¹⁵As is the case in any discrete choice model, welfare *levels* are not comparable across individuals, and in particular, between incumbents and potential entrants.

¹⁶Ideally, one would capture this dependence with a fully dynamic (life-cycle) individual decision model. Unfortunately, embedding such individual dynamics within our *equilibrium* framework would make the exercise intractable; our specification captures the essence in a parsimonious way.

restricting overall mobility.

Relaxing information frictions can ease job transitions for high-quality teachers, allowing them to receive more offers and move to more desirable schools. On the one hand, this can exacerbate the inequality between advantaged schools and disadvantaged schools, if most teachers prefer the former. On the other hand, given the easier upward job mobility under information transparency, high-quality teachers may be induced to enter the market and accept positions in less desirable schools upon entry. Schools can replace their ineffective incumbent teachers with these entrants. As a result, teacher effectiveness may improve both market-wide and in disadvantaged schools. Given these two offsetting forces, the ultimate equity-efficiency implications of information interventions become an empirical question.

4.4 Equilibrium

Definition 2 *A market equilibrium is given by $\{s^*(\cdot), O_s^*(\Gamma_s(\cdot)), P_s(\cdot)\}$ such that*

- 1) *Given $\{o_s(x, q, v, s_0)\}_s$ implied by $\{O_s^*(\Gamma_s(x, q, v, s_0))\}_s$, $s^*(x, q, v, s_0, \epsilon)$ solves the teacher's problem for all (x, q, v, s_0, ϵ) .*
- 2) *Given its belief about $\{P_s(x, q, v, s_0)\}_{(x, q, v, s_0)}$, $\{O_s^*(\Gamma_s(x, q, v, s_0))\}_{(x, q, v, s_0)}$ solves the school's problem for all s .*
- 3) *$\{P_s(\cdot)\}_s$ is consistent with teacher and school decisions $\{s^*(\cdot), O_s^*(\Gamma_s(\cdot))\}$.*

4.4.1 Simplification

For school s to solve its problem, it is sufficient to know teachers' acceptance probabilities $P_s(\cdot)$. Therefore, an equilibrium boils down to a fixed point of $\{P_s(\cdot)\}_s$. However, forming exact beliefs about the high dimensional (non-parametric) objects $\{P_s(\cdot)\}$ is a daunting task for any decision maker.¹⁷ As a feasible alternative, we assume that schools make decisions based on a simplified parametric belief function $\tilde{P}_s(\cdot; \zeta)$:

$$\tilde{P}_s(x, q, v, s_0; \zeta) = \frac{\exp(g(x, q, v, s_0, s; \zeta))}{1 + \exp(g(x, q, v, s_0, s; \zeta))}, \quad (8)$$

where

$$\begin{aligned} g(\cdot) = & \zeta_0 \frac{W(x, q, v, z_s) - \bar{w}(x, q, v)}{\sigma_{w(x, q, v)}} + x\zeta_1 + q\zeta_2 + v\zeta_3 \\ & + z_s\zeta_4 + z_{s1}(\zeta_5q + \zeta_6v) + I(s = s_0)(\zeta_7 + \zeta_8x_1) + I(s_0 = 0)\zeta_9 \\ & + I(s_0 > 0, s \neq s_0)(\zeta_{10}q + \zeta_{11}v) + I(s_0 = 0)(\zeta_{12}q + \zeta_{13}v). \end{aligned} \quad (9)$$

¹⁷Alternatively, a school can derive $\{P_s(\cdot)\}$ based on its belief about other schools' strategies, which is also a rather high-dimensional object.

This simplified belief function captures all the factors governing its counterpart $\{P_s(\cdot)\}$. The first term in (9) relates to how acceptance probabilities vary with the relative pay attractiveness in school s . In particular, $\bar{w}(x, q, v)$ and $\sigma_{w(x, q, v)}$ are the cross-school average and standard deviation (std) of pay for a teacher with attributes (x, q, v) according to the district’s pay function $W(\cdot)$; we measure relative pay attractiveness by the standard deviation of pay at school s relative to the teacher’s average pay in the district. The other terms in the first row relate to the overall quality of the teacher: A school should expect more competitors for a better teacher. The second row of (9) mirrors teachers’ preferences for schools as in (2). The third row captures the effect of information. For district incumbent teachers ($s_0 > 0$), ζ_{10} and ζ_{11} capture a prospective employer’s belief about the quality of teachers coming from other schools. With information asymmetry, a school s may expect it to be harder to hire a teacher with high q and/or high v from another school, because the informed school (s_0) is very likely to have made an offer to the teacher. In that case, ζ_{10} and ζ_{11} are expected to be negative. Finally, ζ_{12} and ζ_{13} capture schools’ belief about how potential entrants self-select into the district in terms of q and v .

For the rest of this paper, we will study the equilibrium with the simplified belief $\tilde{P}_s(\cdot; \zeta)$, which approximates the true acceptance probabilities very well.¹⁸ Solving for the equilibrium with the simplified belief boils down to finding the vector ζ that makes $\tilde{P}_s(\cdot; \zeta)$ defined in (8) consistent with acceptance probabilities $P_s(\cdot)$ derived from teachers’ optimal decisions. Notice that ζ is *equilibrium-specific and policy variant*. In our counterfactual policy simulations, we will search for the equilibrium-consistent vector ζ for each given policy.

4.5 Model Discussion

For tractability, we abstract from several important aspects. First, we take a school’s characteristics z_s as given. In particular, we assume away systematic household re-sorting in response to our policy interventions that would change school-level student characteristics. In our sample period, school-level student demographics and achievement were quite stable and seemingly unaffected by ETI.¹⁹ Although we do not expect our counterfactual policies to significantly affect schools’ student bodies, readers should be aware of this limitation when interpreting our results.

Second, we abstract from the effect of financial incentives and information on individual teachers’ effort and effectiveness. A large literature has focused on the effect of financial

¹⁸The R^2 is 0.988 when we regress $P_s(\cdot)$ on the simplified beliefs $\tilde{P}_s(\cdot; \zeta)$.

¹⁹Between 2007 and 2015, in terms of student achievement (z_1), 82% (93%) of schools in the bottom (top) quartile before 2011 remain in the bottom (top) quartile post 2011; in terms of the fraction of minority students, 95% (85%) of schools in the bottom (top) quartile before 2011 remain in the bottom (top) quartile post 2011.

incentives on students’ test scores; the findings are mixed.²⁰ We complement this literature by focusing on a different channel through which information and financial incentives may affect education, i.e., information and incentives may affect teacher-school sorting and teacher exit and entry. To the extent that teachers may improve their effectiveness in response to financial incentives and information, our counterfactual policy results could understate the total policy effects. In Appendix E.5, we examine pre- vs post-ETI within-teacher differences in a number of proxies of teacher effort and effectiveness, including teacher attendance and EVAAS scores. Although we do not find significant differences in any of these measures, this exercise does not rule out the effect of ETI on teacher effort.

Third, we model the outside option value as a function of teachers’ characteristics. A more comprehensive approach would explicitly model outside employers’ reactions.²¹ This exercise is very challenging and data-demanding, which we leave for future work.

5 Estimation

As we describe in detail in the next section, we have obtained data from both the pre- and the post-ETI periods. To the extent that ETI introduced non-trivial changes to the information structure, if our model, estimated under one regime, can fit the data under the other regime, it will increase our confidence in the model’s ability to make credible counterfactual policy predictions. Both model validation and counterfactual policy simulations are based on the premise that model parameters are fixed across different scenarios. However, our pre-ETI data are from 2010-11 (a great recession year), while our post-ETI data are from 2013-14. It is unlikely that the value of the outside option (e.g., non-teaching jobs) would be the same between these two years. As such, we allow (only) the outside option value parameters to differ between the pre- and post-ETI sample periods. We estimate our model in two stages, leaving enough information out of the estimation to conduct model validation.

Stage 1 We estimate all model parameters Θ using only the post-ETI data via indirect inference, which involves 1) computing from the data a set of “auxiliary models” that summarize the patterns in the data, and 2) repeatedly simulating data with the structural model, computing corresponding auxiliary models with the simulated data, and searching for model parameters such that the auxiliary models from the simulated data match those from 1). In particular, letting $\bar{\beta}$ denote our chosen set of auxiliary model parameters computed from

²⁰Studies using data from outside of the US have found evidence that financial incentives for teachers affect student achievement (Muralidharan and Sundararaman, 2011; Duflo et al., 2012; Lavy, 2002; Atkinson et al., 2009; Glewwe et al., 2010). However, incentive programs implemented in the US have yielded mixed results, e.g., Fryer (2013); Imberman and Lovenheim (2015); Dee and Wyckoff (2015); Brehm et al. (2017).

²¹We found no well-documented evidence that other districts implemented pay schemes in response to HISD’s introduction of the ASPIRE program or ETI.

data and $\widehat{\beta}(\Theta)$ be the corresponding auxiliary model parameters obtained from simulating a large dataset from the model (parameterized by Θ) and computing the same estimators, the estimated vector of structural parameters is the solution

$$\widehat{\Theta} = \operatorname{argmin}_{\Theta} \left\{ [\widehat{\beta}(\Theta) - \bar{\beta}]' W [\widehat{\beta}(\Theta) - \bar{\beta}] \right\},$$

where W is a weighting matrix. In this stage, we target auxiliary models designed to extract all identifying information from the post-ETI data. The vector of parameter estimates $\widehat{\Theta}$ is used to conduct our counterfactual policy analysis in Section 8.

Stage 2 We allow parameters related to teachers' outside options to differ in the pre-ETI year. Denoting Θ^{sub} as the subvector of model parameters excluding outside value parameters, we fix $\widehat{\Theta}^{sub}$ estimated from Stage 1 and estimate the outside-value parameters ($\widehat{\Theta}^{out}$) in the pre-ETI period, for a given information structure $A(q)$ (Table 1), by matching exit-and-entry related moments *only* from the pre-ETI data. Specifically, we hypothesize two cases about $A(q)$, i.e., how much prospective employers know about the q of a teacher working in another district school.

Case 1) $A(q) = q$, i.e., prospective employers know q , since a teacher can present her EVAAS scores to schools;

Case 2) $A(q)$ is such that prospective employers know about q up to the information revealed by the ASPIRE award outcomes (whether or not the teacher's q is above certain thresholds). In both cases, a prospective employer will use $(A(q), x, s_0)$ to infer a teacher's expected value conditional on acceptance as stated in the school's problem (7).

For a given $A(q)$ case, we estimate the $\widehat{\Theta}^{out}$ associated with it, together with Stage-1 estimates $\widehat{\Theta}^{sub}$, we simulate the model and use all the other auxiliary models computed from the pre-ETI data to gauge the fit. Out of the two $A(q)$ cases, the one that fits the pre-ETI data better is the information environment that is more plausible. Moreover, how well this preferred case fits the pre-ETI data serves to validate our model.

The value of re-estimating the pre-ETI outside-value parameters goes beyond model validation; it allows us to infer which informational environment is more plausible pre-ETI. Nevertheless, we also conduct a validation exercise by simulating the model given the pre-ETI initial conditions and *all* Stage-1 estimates, including the outside-value parameters.

5.1 Identification

A challenge involved in identification is the fact that the researcher observes only the accepted offers, rather than all offers made. This makes it hard to separate teachers' preferences

from schools’ preferences (Menzel, 2015).²² Some studies tackle this problem by using data on applications and schools’ rankings of applicants (e.g., Ederer, 2023; Fabre, 2023; Bates et al., 2025); others exploit instrumental variables that influence one side’s payoffs while being excluded from the other side’s preferences (Agarwal, 2015; Laverde et al., 2025). We take a different approach: Assuming information symmetry between the researcher and the school under ETI, we exploit the school’s optimization problem and extract information from observed matches that allows us to overcome this identification challenge. Below, we first provide the main argument and then provide further discussion.

The expected marginal benefit of hiring a teacher for school s given its information set $\Gamma_s(\cdot)$ is given by $\bar{B}_s(\Gamma_s(x, q, v, s_0))$. Under ETI, $\Gamma_s(x, q, v, s_0) = (x, q, v, z_s)$ for district incumbent teachers and therefore the expected marginal benefit of hiring a district incumbent is simply $B(x, q, v, z_s)$, which is weakly increasing in each of the three observed traits: x , q , and v . The marginal cost is the shadow price of a slot, which is common for all teachers. If school s hired district incumbent teacher i , then district incumbent teacher j , who has weakly higher x , q and v than i , must also have had an offer from s under ETI. From the observed “seed pair” ((s, i) in this example), we can infer offers for other incumbent teachers (j in this example). Together with the observed offers, these inferred offers allow us to construct, for each district incumbent teacher, a subset \tilde{O}_i of all offers O_i she received, which always includes the outside option (for 1605 out of 1620 incumbent teachers, \tilde{O}_i contains at least 2 schools). Teachers’ choices within \tilde{O}_i inform us of their preferences.²³

We can also leverage observed offers in a different way to learn about schools’ preferences. First, given that teachers are subject to preference shocks with an unbounded support, as the number of teachers per school grows large, the lowest (x, q, v) among district incumbents working in s is the lowest (x, q, v) that school s is willing to hire (since an offer would be accepted by some teacher). This identifies schools’ preferences over teachers. However, in practice, many schools are small, therefore, our identification also relies on the next observation.

Second, given that the distribution of teachers’ preferences is revealed from their choices within \tilde{O}_i , we can predict the probability that a teacher would choose to work in each school

²²In the context of two-sided matching markets with non-transferable utility, Menzel (2015) show that under pairwise stable matching, the joint surplus is nonparametrically identified; yet, without additional assumptions, it is impossible to separately identify preferences on each side of the market. Ederer (2023) extend this framework to a dynamic setting.

²³Multinomial discrete-choice models can be point-identified using a subset of choices, parametrically (e.g., McFadden, 1978) and semiparametrically (e.g., Fox, 2007). In a framework much more flexible than ours, Barseghyan et al. (2021) allow for *unrestricted correlation* between choice sets and preferences and characterize the sharp identification region of model parameters. We build on insights from these studies to design our auxiliary model Aux 1a (Section 5.2), which is used to extract information useful for identification.

if they had offers from all schools. As long as (some) schools are selective (i.e., they do not make offers to all teachers), this predicted distribution of teacher-school matches will be systematically different from the observed matches, because a teacher can choose a school s only if they had an offer from s . That is, this discrepancy is driven by schools’ offer decisions, which are governed by schools’ preferences.

This argument can be illustrated with a simple example (see Appendix E.1 for a graphical illustration). Suppose all teachers prefer School 1 to School 2. If teacher i works in School 1 while teacher j works in School 2, this implies that School 1 prefers i to j . The same argument applies when teachers have preference shocks: Accounting for teachers’ preference shocks, if teachers systematically prefer School 1 over School 2, then, School 1 must prefer their hires over (most) teachers working in School 2. As long as the distribution of (x, q, v) in School 1 does not systematically dominate that of (x, q, v) in School 2 in all dimensions, we can infer how much School 1 cares about x and v relative to q . That is, given teachers’ preferences, the realized distribution of teacher-school matches identifies schools’ preferences.

Discussion To clarify our identification strategy, we discuss its three steps in turn and highlight which assumptions are necessary and which are made for convenience.

Step 1: We construct a subset of offers for each teacher by exploiting schools’ optimal decisions given their information sets. Under ETI, a school has full information about teachers’ attributes (x, q, v) , which allows us to rank teachers and construct \tilde{O}_i based on (x, q, v) , given the following two assumptions.²⁴

A1: Schools cannot discriminate teachers by factors other than (x, q, v) . If some job offers were made based on factors we cannot observe, the inferred \tilde{O}_i might include infeasible options for some teachers, thus introducing biases in the inferred teacher preferences based on \tilde{O}_i . It is likely that such biases in our setting are smaller than they would be in a usual setting: Under ETI, schools and the researcher observe rich measures of teacher quality, including teachers’ VA to test scores (q), and, more uniquely, quality revealed via in-person observations (v). This setting makes the information about teachers more symmetric between the school and the researcher than in a usual setting. Nevertheless, this assumption might be violated in some teacher-school matches. As a robustness check, we re-construct inferred offer sets based on a stricter rule: From an observed “seed” school-teacher pair (s, i) , we include s in teacher j ’s offer set only if j has q and v at least 0.3 standard deviations higher than those of i , rather than just weakly higher. Doing so significantly reduces the size of \tilde{O}_i ,

²⁴We cannot use the same argument to construct \tilde{O}_i pre ETI because the information about (q, v) is asymmetric between a teacher’s employer and other schools. Therefore, we use the post-ETI sample rather than the pre-ETI sample for Stage-1 estimation. Similarly, because schools only observe potential entrants’ x but not their (q, v) , we use only district incumbents to find seed pairs and to construct \tilde{O}_i .

however, our auxiliary models remain robust (Appendix Table B11).

A2: There are no job posting costs. This assumption is plausible because in reality, vacancies in HISD are posted publicly online.²⁵ We also assume that teachers get offers without having to apply. This assumption does not affect our inference of teacher preferences because the following two cases are observationally equivalent in our framework: (i) school s made an offer that teacher i rejected; (ii) teacher i was eligible for a job in s but did not apply. Both cases imply that s was not sufficiently attractive to teacher i .²⁶

Step 2: Given the inferred subset of offers, we identify teachers’ preferences from their choices within \tilde{O}_i , based on the insight from Fox (2007), who proves semiparametric identification of multinomial discrete-choice models using a subset of options. Besides standard regularity conditions, the identification relies on two conditions. The first condition is rank-ordering: Choices with higher deterministic payoffs are more likely to be chosen by a given agent. The second condition is exchangeability, also known as label independence. These two conditions hold in a large set of multinomial discrete-choice models (with additively-separable error terms), including ours.²⁷ Provided that preference shocks are additively separable, specific distributional assumptions are not necessary; we assume Type-1 extreme-value distribution for convenience. District incumbents’ choices within \tilde{O}_i reveal teachers’ evaluation of school characteristics relative to wages, governed by preference parameters that are common among all teachers. In addition, these choices also inform us of how much district incumbents value the outside option. The entry rates of potential entrants with different characteristics in turn inform us of how much they value the outside option.

Step 3: Given teachers’ preferences, we identify schools’ preferences from the realized set of teacher–school matches. Unlike the teacher-side of the problem, we know each school’s choice set—all incumbent teachers and potential entrants. Given A1 and A2, as long as a school does not make offers to all teachers, its preference is revealed by its offer decisions. Moreover, provided that schools agree that teacher i is weakly better than j if i has higher (x, q, v) , we can allow schools to differ in how they rank teachers in ways that are more

²⁵In Appendix D.3, we show a screenshot of the HISD job posting website (<https://www.applitrack.com/houstonisd/onlineapp/>) from April 2021. Our conversation with HISD staff confirms that posting a job on this platform is rather simple.

²⁶If it is costly for teachers to apply for jobs (more so for jobs in schools other than one’s current school), these costs would be absorbed in teachers’ moving costs in our model.

²⁷The first condition requires that for a given agent, choice probabilities are rank ordered by deterministic payoffs (random coefficient models violate this condition), but two agents with only slightly different deterministic payoffs can make choices with different probabilities. The second condition holds for any collection of i.i.d. random variables.

flexible than our current (parsimonious) specification.²⁸

5.2 Auxiliary Models

The identification argument in Section 5.1 suggests targeting statistics summarizing 1) teachers’ choices within the inferred offer subset \tilde{O}_i and 2) the realized distribution of teacher-school matches. Notice that, although certain auxiliary models are intuitively more informative about certain structural parameters than others (as we explained above), the identification of the model relies on using information extracted from all auxiliary models and we target the following auxiliary models *jointly*. To provide more evidence on the mapping between data and parameters, in Appendix E.6, we perturb structural parameters one by one and measure the responses of subsets of auxiliary models.

Aux 1: Coefficients from two regressions of the following form

$$y_{is} = \beta_1 W(x_i, q_i, v_i, z_s) + z'_s \beta_2 + z_{s1} (\beta_3 q_i + \beta_4 v_i) + I(s = s_{0i}) (\beta_5 + \beta_6 x_{i1}) \quad (10) \\ + I(s = 0) [\beta_7 + x'_i \beta_8 + \beta_9 q_i + \beta_{10} v_i] + \psi_i + \varepsilon_{is},$$

where $y_{is} = 1$ if teacher i is matched with school s , 0 otherwise. The right-hand-side variables are the same as those enter teacher’s preferences, including wage $W(\cdot)$, the vector of school characteristics z_s (including z_{s1}), z_{s1} interacted with q_i and v_i , whether or not s is one’s current school $I(s = s_{0i})$ (in itself and interacted with experience), whether or not s is the outside option $I(s = 0)$ (in itself and interacted with teacher characteristics). In addition, we include a teacher dummy ψ_i to relate all (i, s) observations associated with teacher i .²⁹ The two regressions differ in the number of observations, reflecting the identification argument in Section 5.1.

- (a) In the first regression, i ’s are all district incumbent teachers whose inferred subsets of offers \tilde{O}_i contain more than one option, and an observation (i, s) is a teacher-district pair in these inferred subsets.
- (b) In the second regression, an observation is any incumbent teacher-destination pair, with $\#incumbents \times (S + 1)$ total observations.

²⁸For example, we can allow for any interaction of teachers’ and schools’ traits, provided that $B(\cdot)$ is weakly increasing in teacher traits. Empirically, given the small sample size of schools, we keep the specification parsimonious and it explains the data well.

²⁹While conditional logit regressions may seem a more intuitive choice for summarizing discrete choices, they are computationally too expensive to run within each iteration of the estimation process. We use a linear regression specification with teacher dummies, which capture the idea that the same teacher is choosing one school out of a given choice set.

Aux 2: Average school-level teacher characteristics within each group of schools by quartiles of z_{s1} , which supplement Aux 1b.

Aux 3: Moments capturing job mobility:

- (a) Average entrants' characteristics and the characteristics z_{s1} of schools where an average entrant works: $E(z_1|entrants)$;
- (b) Cross moments of entrants' characteristics and their employers' characteristics z_{s1} : $E(xz_1|entrants)$, $E(qz_1|entrants)$, $E(vz_1|entrants)$;
- (c) Incumbent teachers' exit rate and within-district job switching rate, which supplement Aux 1a;
- (d) The ratio between the number of entrants and that of incumbents who stayed.

6 Data

Our data consist of two linked data sets from HISD at the teacher and school levels for years 2007 to 2015. Data are reported by academic year, and referenced using the calendar year of the spring semester (e.g. 2013 for 2012-13). Here, we provide a brief summary of our data and how we construct data counterparts of our model components; in Appendix C, we provide further details and additional summary statistics.

Teacher Data: For information on district incumbent teachers, we use administrative data from HISD, covering all individuals employed by HISD between 2008 and 2015. This panel contains information about teachers' education, years of teaching experience, EVAAS scores, school and class identifiers, and grades and subjects taught. For post-ETI years, we also observe each teacher's yearly ETI component scores. For information on potential entrants, we use data from the Texas Academic Performance Reports, which provide campus-level data on public schools in Texas. From these reports, we use information on the distribution of teachers' experience and education (x) in all the 64 non-HISD districts located in the Houston metropolitan statistical area (MSA).³⁰

School Data: We use information provided by the Common Core and the Texas Education Agency on school characteristics, including funds received and proportions of students in each grade meeting state-set satisfactory levels in standardized tests.

6.1 Empirical Definitions

To map our equilibrium model to the data, we need to introduce some empirical definitions.

³⁰Ideally, we would also use the information on private-school teachers to construct the distribution of potential entrants. Unfortunately, we do not have information on these teachers, who account for a small fraction of teachers in Texas (e.g. lower than 10% in 2020).

The Market Our model is in a static equilibrium setting. For estimation and counterfactual policy analysis, we use data from 2014, allowing three years for the market to adjust following the introduction of ETI in 2011. Additionally, since teachers began receiving all three ETI component scores in 2013, prospective employers had information on all three ETI component scores when making hiring decisions in 2014. To validate the estimated model, we simulate the market equilibrium under initial conditions in 2011. ETI scores were assigned to teachers starting at the *end* of the 2011 academic year, allowing us to observe ETI scores of market participants in 2011; however, hiring in 2011 was made in the pre-ETI information environment.

In both 2011 and 2014, we focus on the market for public-school non-substitute full-time math and reading teachers in Grades 3-5 for the following reasons. First, we exclude substitute and part-time teachers, as these workers typically face different types of contracts than regular full-time teachers.³¹ Second, we follow the literature and exclude secondary school teachers because, unlike elementary school teachers, secondary school teachers often teach multiple grades, making it hard to identify individual teacher contributions (Kane and Staiger, 2008; Chetty et al., 2014). Among elementary school teachers, value-added measures (EVAAS scores) are obtainable for those who teach math or language in Grades 3-5. The cost of this focus is the implicit assumption that a school’s capacity constraint for these teachers does not interact with that for other teachers. In addition, we dropped 295 teachers who had missing values for key variables. Within HISD, the estimation sample contains 169 schools and 2,033 teachers (incumbents and entrants); and the validation sample contains 171 schools and 1,970 teachers.

Although all qualified individuals can be counted as potential entrants, it is reasonable to focus on the group who have non-trivial probabilities of entering the market as defined above. Therefore, for potential entrants, we focus on teachers employed in all 64 non-HISD districts within the Houston MSA. Among these, we use information on non-substitute full-time teachers with lower than 30 years of experience who taught Grades 3-5 in 2013, including novice teachers.³²

Schools *School Characteristics* z_s include three variables: the fraction of the school’s students who passed the state’s standards on standardized math tests (z_{s1}), funding per teacher net of teacher salaries (z_{s2}) (henceforth funding per teacher), and the school’s previous-year

³¹We require that teachers teach at least 75% of total classroom hours for a class of a year to be included in our sample. Substitute teachers make up fewer than 1% of public school teachers in Texas.

³²The dataset we use does not contain information on the subjects taught by the teachers. The exclusion of teachers with over 30 years of experience from potential entrants is largely consistent with the data: Of the 498 teachers who entered HISD as a full-time Grade 3-5 math/language teacher in 2014, only 3 had over 30 years of experience.

campus award eligibility (z_{s3}) as discussed in Remark 1.

School Capacity κ_s Assuming that the data are generated from an equilibrium and therefore schools' capacity constraints bind, κ_s is the number of teachers in our estimation (validation) sample who worked in s in 2014 (2011).

Teachers *Teacher Characteristics* x include experience (x_1) and an indicator of whether one has a graduate degree (x_2). For data confidentiality reasons, we model x_1 in coarse experience groups: 0 (new teachers), 1-2, 3-4, 5-8, 9-13, 14-29, and 30 or more years (retirement-age teachers).³³ This grouping respects the experience-education wage schedule in HISD and is comparable with previous literature. It should be noted that the potential entrant pool includes both new and experienced teachers.

Teacher Quality (q, v) : For q , we use the average of a teacher's reading and math EVAAS scores prior to the hiring year in question. We standardize q to be of mean zero and standard deviation 1 within each year among all Grade 3-5 math and reading teachers in the district. To construct v , we combine ETI's two in-person evaluation measures: instructional practice and professional expectations. We describe this procedure, and how we calculate v for teachers in pre-ETI years, in Appendix C.1. This procedure gives us (q, v) for all HISD incumbents and those who entered, and hence the joint distribution of (x, q, v) among HISD teachers.

For potential entrants who did not enter, we do not observe their (q, v) . To obtain the joint distribution of $F(x, q, v)$ for all potential entrants, we use the distribution of their x and the conditional distribution of $F(q, v|x)$ among HISD teachers, the latter assumed to be the same as that among all *potential* entrants. That is, we assume $F(q, v|x)$ is common for all districts within the Houston MSA. Notice that the distribution $F(q, v|x, entrant)$ is endogenous and therefore would be different from the distribution $F(q, v|x)$ among all *potential* entrants. We will use our model to explain the endogenous distribution $F(x, q, v|entrant)$.

Teacher's Origin School (s_0) : We define district incumbent teachers in year t as those who had worked in HISD before t since 2008. An incumbent teacher's s_0 is their most recent employer in the district prior to t . We define district entrants in t as those who entered HISD but had never worked in HISD before t (since 2008). Entrants and other potential entrants have $s_0 = 0$.

³³To save on the number of structural and auxiliary model parameters, we use the median within each experience group as x_1 in teachers' moving cost function and in deriving summary statistics and auxiliary models. Appendix C.2 explains how we construct the initial joint distribution of teachers across schools without being able to export teacher-level data.

Table 2: Summary Statistics (Post-ETI)

A. Teacher Characteristics	All	Incumbents	Entrants
Experience (x_1)	9.74 (8.90)	11.07 (8.89)	4.51 (6.72)
Graduate Degree (x_2)	0.28	0.29	0.23
q	-0.06 (1.02)	0.01 (1.00)	-0.34 (1.06)
v	-0.09 (0.85)	0.02 (0.81)	-0.53 (0.86)
$corr(q, v)$	0.27	0.25	0.23
# Teachers	2,033	1,620	413
B. School Characteristics	All	1st Quart. z_1	4th Quart. z_1
z_1 (Frac. students meeting std.)	0.73 (0.15)	0.52 (0.07)	0.90 (0.04)
Funding per teacher (\$1,000)	21.35 (6.51)	25.48 (9.76)	20.74 (4.82)
Capacity (#teaching slots)	10.23 (3.86)	7.67 (3.15)	12.10 (4.29)
# Schools	169	42	42

Panel A shows teacher-level statistics. Cross-teacher standard deviations in parentheses. Panel B shows school-level statistics for all schools, and schools in the top and bottom quartiles of percent of students meeting testing standards. Cross-school standard deviations in parenthesis.

6.2 Summary Statistics

Table 2 summarizes our post-ETI sample. Panel A shows average characteristics among all HISD teachers, HISD incumbents, and entrants; cross-teacher standard deviations are in parentheses.³⁴ Compared to incumbent teachers, entrants have lower experience, education, and quality as measured by q and v . The two quality measures, q and v , are weakly correlated.

Panel B of Table 2 summarizes school-level statistics, with cross-school standard deviations shown in parentheses. We present statistics for all schools in our estimation sample and separately for schools ranked in the bottom and the top quartiles in the distribution of z_1 , i.e., the fraction of a school’s students who met the state standard (high-performing henceforth). Schools with fewer high-performing students tend to have fewer teaching slots but larger funding per teacher. The latter arises partly because of the district’s “targeted interventions” favoring low-performing schools.³⁵

Table 3 shows the equilibrium teacher-school sorting, entry and exit in 2014. Treating each school as an observation, Panel A shows school-level teacher characteristics among schools in different quartiles of the z_1 distribution (cross-school std dev shown in parentheses). Teacher-school matching exhibits a strong assortative pattern, in terms of experience, q , and v . Panel B shows that in 2014, 18.8% of district incumbent teachers exited HISD and 6.6% of them switched schools within HISD. Among all teachers working in our 2014 market, the ratio between entrants and incumbents was 23.9%.

³⁴We standardize q and v using all math and reading teachers in Grades 3-5, some of whom are excluded from our sample as explained earlier. Therefore, the q and v in our estimation sample do not have exactly mean 0 and standard deviation of 1.

³⁵School funding from the district is determined by a formula which accounts for number of students, demographic composition of students, and “targeted interventions.”

Table 3: Outcome Sorting, Entry, and Exit (post-ETI)

A. Average Employee Characteristics post-ETI (Aux 2)						
School Group by z_1	Experience		q		v	
Quartile 1	8.00	(3.67)	-0.46	(0.58)	-0.24	(0.36)
Quartile 2	8.85	(2.94)	-0.18	(0.43)	-0.15	(0.30)
Quartile 3	9.99	(4.07)	0.01	(0.37)	-0.09	(0.29)
Quartile 4	10.49	(3.63)	0.24	(0.26)	0.03	(0.21)

B. Job Mobility						
Incumbents			Entrants			
Exit HISD	18.8%		$\frac{\#Entrants}{\#Stayers}$		23.9%	
Within-HISD Job Switch	6.6%					

Panel A shows teacher-level statistics with cross-teacher standard deviations in parentheses. Panel B shows school-level statistics for all schools, and schools in the top and bottom quartiles of percent of students meeting testing standards. Cross-school standard deviations are shown in parentheses.

7 Estimation Results

7.1 Parameter Estimates

Table 4 shows estimates of selected model parameters and their standard errors (in parentheses) derived numerically via the Delta Method; other parameter estimates are reported in Table A1. Panel A of Table 4 shows teacher preferences for schools. For an average teacher, schools with higher fractions of students meeting the state standard (z_1) are more desirable. Moreover, coefficients of the interactions between z_1 and both measures of teacher quality— q (VA) and v (in-person evaluation of one’s teaching and professional practice)—are estimated to be positive. That is, higher-quality teachers have stronger preferences for schools with academically stronger student bodies. For example, a teacher with average values of q and v would put a premium of \$4,636 on a school with $z_1 = 0.9$ over an otherwise identical school with $z_1 = 0.5$; this premium is \$12,177 (-\$2,905) for a teacher whose q and v are both ranked at the 80th (20th) percentile.

Although statistically insignificant, the point estimate of the parameter associated with z_2 is positive. That is, teachers prefer schools with larger (non-teacher salary) funding per teacher, *all else being equal*. Recall that Table 2 shows a negative correlation between a school’s student body (z_1) and its funding (z_2). Since teachers value both of these factors, Texas’s funding formula, which favors low- z_1 schools, has increased the attractiveness of such schools, but only to a limited extent. For example, at the status quo, an average teacher places a \$3,800 premium on an average school in the top quartile of z_1 compared to an average school in the bottom quartile of z_1 ; if we removed the cross-school funding (z_2) difference, this premium would increase to \$4,300.

Finally, consistent with findings from previous studies on worker mobility (e.g., Kennan

and Walker, 2011), we find that teachers have high moving costs, which grow with teacher experience, and that they face highly-dispersed preference shocks (σ_ϵ): High average moving costs help explain the lack of teacher mobility in general, while preference shocks absorb idiosyncratic reasons for mobility.

Panel B presents the functional form of $B(\cdot)$, which governs school preferences, along with the corresponding parameter estimates. As detailed in Appendix A.2.2, this functional form ensures that schools’ valuations of teachers are increasing in x , q and v . Additionally, it allows school preferences to vary with the school’s proportion of high-achieving students (z_1). We find that schools highly value teachers’ VA, q ; the strength of this preference increases with a school’s fraction of high-performing students (z_1). Relative to q , a teacher’s quality as measured by v (in-person evaluation scores) is less valuable to schools in general. Our point estimates also indicate that v is valued relatively more by schools with more low-achieving students (lower z_1), although the estimate is statistically insignificant. *Conditional* on teacher quality q and v , we find that schools directly care about whether or not a teacher has ever taught, but beyond that, they care little about a teacher’s additional years of experience and education.

Although our model is silent about why schools put different weights on various traits of a teacher, the resulting estimates are reasonable. For example, given the general emphasis on test scores as performance measures for both students and schools, it is unsurprising that schools care a lot about teachers’ VA (q) relative to any other teacher trait.³⁶

Our estimates have major implications for educational equity and how it may be affected by the information structure. We find that preferences on *both* sides of the market exhibit complementarity: Teachers, especially high-quality teachers, prefer high-performing schools, while high- q teachers are strongly preferred by all schools but more so by high-performing schools. Without frictions, these preferences would directly lead to assortative matching between teachers and schools, leaving lower-performing schools behind. Information frictions can mitigate this assortative matching, as schools may be reluctant to hire teachers from other schools in fear of the winner’s curse.

7.2 Model Fit and Model Validation

Model Fit Table A2 shows within-sample model fits of auxiliary regressions (Aux 1 in Section 5.2) that characterize the allocation of incumbent teachers across district schools and the outside option. With a few exceptions, the model closely matches these targeted

³⁶It is also reasonable to find that a teacher’s quality v is valued more by schools with more low-achieving students (although this estimate is statistically insignificant): These schools are more likely have students with behavioral problems (e.g., Tippett and Wolke, 2014; Mahvar et al., 2018; Badger et al., 2023), making a teacher’s classroom management skills (a key component of v) very important.

Table 4: Parameter Estimates

A. Teacher Preferences				
wage (\$1,000)	1	normalized	funding (\$1,000 per teacher)	0.10
z_1 (Frac. students meeting standard)	11.59		$I(s \neq s_0)$ (moving)	-70.24
$z_1 \times q$	15.29		$I(s \neq s_0) \times$ experience	-1.11
$z_1 \times v$	7.11		σ_ϵ	10.08
B. School Prefs $B(x, q, v, z_s) = xb_0 + \exp(b_1 z_{s1}) \ln(q - \underline{q}) + \exp(b_2 + b_3 z_{s1}) \ln(v - \underline{v})$				
q :			Yrs of Experience:	
b_1	1.01		1-2	1.16
v :			3-4	1.48
b_2	-0.63		5-8	1.48
b_3	-0.43		9-13	1.68
Graduate degree	0.09		≥ 14	1.68

Std errors (in parentheses) are derived numerically via the Delta Method.

coefficients. Table A3 shows model fits for the moments of teacher characteristics by school groups (Aux 2) and job mobility (Aux 3). The overall fit is good, although the model slightly overpredicts the exit rate and the assertiveness of teacher-school matches.

Model Validation with Re-estimated Outside-Value Parameters As described in Table 1, in pre-ETI periods, a teacher’s current employer was the only school fully informed of the teachers’ quality; other schools had no information on v and might have some information on q , captured by a function $A(q)$. We have estimated the pre-ETI outside value parameters by matching only entry and exit moments for each of the two $A(q)$ specifications described in Section 5: Case 1) $A(q) = q$, i.e., prospective employers know q , and Case 2) $A(q)$ is such that prospective employers know about q up to the information revealed by the ASPIRE award outcomes. Using these parameter estimates (Appendix Table A1), together with the non-outside value parameters from Stage-1 estimation, we simulate the model under each given $A(q)$ and the initial conditions from 2011 data.

Tables A4, A5, and B10 compare model-predicted 2011 equilibrium outcomes, for each of the $A(q)$ cases, with the data. Although both these models fit well the targeted moments (entry and exit), Case 2) is closer to the pre-ETI environment, as measured by the data-model distance of non-targeted auxiliary models.³⁷ Under Case 2), our model match reasonably well the pre-ETI distribution of teacher-school matches, as captured by the pre-ETI non-targeted Aux 1 and Aux 2. For example, although it under-predicts the average teacher quality in low-performing schools pre-ETI, overall, it replicates the fact that the assortativeness of teacher-school allocation was much weaker in the pre-ETI period than the post-ETI period.

³⁷Using the inverse of variance of each target as the weights, the overall data-model distance of non-targeted auxiliary models is 51% higher under Case 1) compared to Case 2).

Table 5: Schools' Information Set $\Gamma_s(x, q, v, s_0)$ for Counterfactual Simulations

	District Incumbent Teacher		Potential Entrant
	Current school	Other schools	All schools
Symmetric (baseline)	x, q, v, s_0	x, q, v, s_0	x, s_0
Asymmetric (counterfactual)	x, q, v, s_0	x, s_0	x, s_0

Column 1 gives the current employer school's information set about a teacher. Column 2 gives other schools' information set about a district incumbent teacher. Column 3 gives all schools' (common) information set about a potential entrant.

Model Validation with All Post-ETI Parameter Estimates Given the impact of the economic downturn in the pre-ETI year, without re-estimating the outside-value parameters, the model is not expected to fit the market-level teacher quality. This is reflected in the last row of Table A6, which shows the teacher quality in an average school, data versus model. However, given that teachers' valuation of one school versus another is (assumed to be) the same before and after ETI, the model can be validated by its performance in predicting the relative quality of teachers across schools. Rows 1-4 of Table A6 present the average teacher quality (q, v) in each school group relative to that in an average school, data versus model. The model is able to replicate these relative values well.

8 Counterfactual Policy Simulations

Using our estimated model, we evaluate the role of information asymmetry in the market and assess how it affects the effectiveness of policies designed to incentivize teachers to work in lower-performing schools. It should be noted that, given the data at our disposal, we treat HISD as a market; the quantitative equity-efficiency implications from our policy simulations should be interpreted accordingly. However, the lesson learned from these exercises, namely the role of information and its interaction with compensation policies, is useful for teacher markets defined more broadly, e.g., the public school sector (vs. private schools) and the teaching profession (vs. other occupations).

8.1 Equity-Efficiency Implications of Information

Holding all the other initial conditions fixed at their post-ETI levels, we compare two market equilibria, each arising from an information environment displayed in Table 5: the baseline symmetric-information environment versus a counterfactual environment where prospective employers have no information on a teacher's (q, v) .³⁸

Table 6 presents the results. Panel A shows that information increases the within-district job switching rate: 6.4% in the symmetric-info case versus 5.1% in the asymmetric-info case.

³⁸We have also conducted a counterfactual where prospective employers know q up to the ASPIRE award results. As expected, the equilibrium outcomes in that case (Appendix E.2) lie in between those in the two environments shown in Table 6.

Table 6: Symmetric vs Asymmetric Information Equilibrium

	Symmetric		Asymmetric	
A. Job Mobility (%)				
Exit Rate	20.7		20.7	
Job Switch Rate	6.4		5.1	
$\frac{\#Entrants}{\#Stayers}$	22.8		22.6	
B. Average Entrant Characteristics				
q	-0.36		-0.39	
v	-0.55		-0.56	
C. Teacher Quality				
	q		v	
	Symmetric	Asymmetric	Symmetric	Asymmetric
Overall	-0.057	-0.062	-0.080	-0.084
By school z_1 : Quart 1	-0.50	-0.52	-0.25	-0.28
Quart 2	-0.19	-0.17	-0.16	-0.14
Quart 3	0.06	0.08	-0.03	-0.01
Quart 4	0.25	0.20	0.06	0.02
Gap (Q4-Q1)	0.75	0.72	0.31	0.30

Panel A shows simulated job mobility statistics under symmetric information and asymmetric information.

Panel B shows average entrant characteristics under symmetric information and asymmetric information.

Panel C shows average q and v of teachers overall and within school quartiles under symmetric information and asymmetric information.

In contrast, entry and exit rates are very similar between these two cases. Despite its null effect on entry and exit rates, information does affect the composition of teachers: Panel B shows that relative to a market with asymmetric information, a market with symmetric information attracts higher-quality teachers to enter. As a result, Panel C shows that the overall teacher quality in the district, as measured by both q and v , is slightly higher when information is symmetric.

Panel C also zooms into quartiles of schools ranked by a school's fraction of high-achieving students (z_1). Making information symmetric across employers improves the average teacher quality in the top and the bottom 25% of schools, while decreasing it in the middle two quartiles. These non-monotonic effects arise for the following reasons. First, with symmetric information, all schools are better at identifying high-quality teachers, but it is the top schools that are best able to poach high-quality teachers from other schools because they are preferred by most teachers, especially high-quality teachers (Table 4). Second, schools in the middle lose more to top schools than bottom schools do because fewer high-quality teachers initially work in bottom schools.³⁹ Third, as shown in Panel B, symmetric information induces higher-quality teachers to enter the market. These entrants, although less effective than incumbent teachers on average, are more effective than many teachers working in low-performing schools; therefore, these schools benefit most from this extensive-margin effect.

³⁹Only 17% of teachers initially working in the bottom-quartile schools are in the top quartile of q , this fraction is 21% in the second-quartile schools and 27% in the third-quartile schools.

Table 7: Teacher Bonus (\$/year)

Teacher χ	School z_1		
	1st 12.5%	2nd 12.5%	other schools
4th quartile	10,000	5,000	0
3rd quartile	8,000	4,000	0
below median	0	0	0

Size of teacher bonus as function of teacher quality measure, χ , and school's fraction of high achieving students, z_1 .

8.2 Teachers' Bonus Programs and Information

To provide students with equal access to high-quality teachers, many school districts have implemented various policies that favor low-performing schools, including giving bonuses to high-quality teachers working in low-performing schools.⁴⁰ We design counterfactual teachers' bonus programs in the same spirit.

Specifically, letting $\chi(q, v)$ be a measure of teacher quality, our program incentivizes higher- χ teachers to teach in schools ranked in the bottom 25% of the z_1 distribution (i.e., schools with the lowest fraction of high-achieving students). We consider three quality measures: 1) $\chi(q, v) = q$, 2) $\chi(q, v) = q + v$, and 3) $\chi(q, v) = v$. Table 7 specifies the amount of bonus a teacher (incumbent or entrant) receives for working in a given school as a function of teacher quality $\chi(q, v)$ and school characteristics z_1 .

Remark 3 *Bonuses are typically rewarded to qualifying teachers at the end of the academic year, as is the case for ASPIRE awards. Following this practice, our bonus programs do not require that the district observe a teacher's quality upfront (e.g., the district may not observe a teacher's quality upon entry).*

Given the role of information frictions in shaping equilibrium teacher-school matches, it is natural to hypothesize that the impacts of bonus programs could also vary with the market's information environment, although it is not clear how. To shed light on this, we examine the effects of our counterfactual bonus programs separately on a market with symmetric information and an otherwise-identical market with asymmetric information, as described in Table 5. Furthermore, to explore the possibility of improving educational efficiency and equity at no additional monetary costs, we impose a budget-neutral condition for our bonus programs. We do so by reducing teachers' base wages $W(\cdot)$ proportionally by a fraction τ when a teacher bonus program is in place; τ is calibrated for each counterfactual simulation such that the district's total expenditure on teacher compensation is the same with and

⁴⁰For example, in 2016, the Dallas Independent School District introduced the Accelerating Campus Excellence program, which provided up to a \$10,000 stipend for high-performing teachers who worked in disadvantaged schools.

Table 8: Effects of Bonus Programs

	A. $\chi(q, v) = q$				B. $\chi(q, v) = q + v$				C. $\chi(q, v) = v$			
	Symmetric		Asymmetric		Symmetric		Asymmetric		Symmetric		Asymmetric	
	Δq	Δv	Δq	Δv	Δq	Δv	Δq	Δv	Δq	Δv	Δq	Δv
Overall	0.05	0.003	0.05	0.01	0.04	0.02	0.04	0.02	0.01	0.02	0.01	0.03
By school z_1 :												
Quartile 1	0.22	0.05	0.19	0.05	0.18	0.10	0.14	0.09	0.09	0.13	0.04	0.12
Quartile 2	0.02	0.01	0.03	0.01	0.02	0.01	0.02	0.02	0.01	0.004	0.004	0.01
Quartile 3	0.003	-0.01	0.01	-0.002	0.003	-0.01	0.01	0.001	-0.01	-0.01	-0.001	0.002
Quartile 4	-0.02	-0.02	0.01	-0.01	-0.02	-0.02	0.01	-0.002	0.01	-0.01	0.001	0.002

Changes in average teacher q and v overall and in schools within each z_1 quartile when adding teacher bonus programs.

“Symmetric” (“Asymmetric”): a market with the information structure shown in Row 1 (Row 2) of Table 5.

without the bonus program; the budget-balancing τ ranges between 1.2% and 1.6% (detailed in Appendix Table B3). Although a 1.2-1.6% reduction in base wages is arguably minor, to the extent that the district may face resistance for any reduction in teachers’ base wages, in Appendix E.4, we also conduct counterfactual bonus experiments without changing base wages, which requires the district to spend more. The policy effects in those cases are qualitatively similar, and unsurprisingly, quantitatively larger.

The effects of the budget-neutral bonus programs are shown in Table 8. Panels A, B, and C correspond to bonus programs using each of the three quality measures, respectively. Within each panel, the columns labeled “Symmetric” show the program effect (i.e., the difference in each outcome with versus without the bonus) on a market with symmetric information; the columns labeled “Asymmetric” show the program effect on a market with asymmetric information. Our main findings are as follows:

Result 1: Under both types of market informational environments, all of our bonus programs improve both equity and efficiency. By making the targeted (low-performing) schools more attractive than they used to be, these programs increase teacher quality at targeted schools. Moreover, by attracting higher-quality teachers to enter the market, all of these programs also improve the overall teacher quality in the district. As a result, the improvement at targeted schools is accompanied with very small, and mostly positive, changes in other schools, except for the highest-performing school group, which tends to experience slight decreases in their teacher quality.

For example, in Panel A, we set $\chi(q, v) = q$, i.e., the bonus is based only on EVAAS scores. On a market with symmetric information, this bonus program leads to an increase in overall q (v) by 0.05 (0.003), thus improving efficiency or total teacher quality. As intended, the program is especially successful at attracting high- q teachers to low-performing schools.

Result 2: Quantitatively, all of our programs are more effective at improving teacher quality

in targeted schools when implemented in a market with symmetric information, suggesting a fruitful interaction between targeted teachers’ bonus policies and information intervention policies.

Comparing the columns with symmetric information (“Symmetric”) to those with asymmetric information (“Asymmetric”), the improvement in teacher quality in targeted (the bottom quartile) schools is larger in the former case across all three panels. For example, in Panel B, the improvement in the average q among teachers in the targeted schools is 0.14 in the asymmetric-info case and 0.18 in the symmetric-info case (a 29% difference). This occurs because with symmetric information, targeted schools can identify and use their bonus advantages to poach high-quality teachers working in other schools more effectively.

Result 3: Due to the positive correlation between q and v , all three bonus programs improve average teacher quality in both dimensions for targeted schools and at the district level. This holds even when only one of the two measures is used in the bonus formula (Panels A and C). However, the improvement in q tends to be larger than that in v , except when bonuses depend only on v ; this is largely driven by schools’ preferences, who value q more than v .

Specifically, as expected, when bonuses are tied to q only, average teacher v improves much less than q in targeted schools (Panel A). In Panel B, we reward both q and v equally by setting $\chi(q, v) = q + v$. Although this program is more successful than the first bonus program (Panel A) at increasing the average v in targeted schools, its effect on q is still larger than its effect on v . This is because, whenever possible, schools—including the low-performing schools—prioritize hiring higher- q teachers: Even though low-performing schools care more about v than high-performing schools do, all schools care about q more than v . Moreover, given that q is more valued by schools, this bonus program is also more successful at attracting high- q teachers to enter the market, resulting a larger overall district-level improvement in q than in v . Given these findings, we further examine the case when bonuses are based only on v , i.e., $\chi(q, v) = v$ in Panel C. This bonus program is able to “counteract” schools’ preferences and induce a larger improvement in v than in q both in targeted schools and at the district level.

9 Conclusion

Information frictions between current and prospective employers shape not only worker mobility but also entry and the allocation of talent across firms. In this paper, we study these forces in the decentralized labor market for public school teachers, a setting in which both efficiency and equity are central concerns.

We develop and estimate an equilibrium model of teacher mobility under asymmetric

information and apply it to administrative data from the Houston Independent School District, where a reform made multidimensional measures of teacher quality observable across schools. The model captures how information affects both sorting across schools and teachers' entry and exit decisions, and it allows us to quantify the equilibrium and distributional effects of increased transparency.

Counterfactual simulations using the estimated model reveal three main findings. First, making information symmetric between employers improves average teacher quality overall and in bottom- and top-quartile schools ranked by student performance, but hurts schools in the middle. This result highlights the role of information frictions in shaping the market equilibrium at the intensive margin (teacher-school sorting) and the extensive margin (teacher entry and exit). Second, budget-neutral policies that incentivize high-quality teachers to work in low-performing schools can increase overall teacher quality while reducing inequality across schools. Third, these policies are more effective when combined with transparent teacher evaluation.

These results point to a more general implication: Information shapes not only the allocation of workers within a market, but also workers' decision to enter or exit it. In markets where equitable access to high-quality services is a key concern, such as healthcare, public safety, and education, information policies should account for both margins, as they jointly determine the overall quality of the workforce and its distribution, and hence the efficiency and equity of service provision. Further, our findings suggest a fruitful interaction between information and compensation policies to improve efficiency and equity.

Our framework leaves several extensions for future work. First, with additional data, one could incorporate decisions by private schools and consider competition not only among public schools, but also between public and private sectors. Second, incorporating household sorting (e.g., Epple and Sieg, 1999; Epple and Romano, 2003; Ferreyra, 2007; Epple and Ferreyra, 2008) would allow for a more comprehensive evaluation of targeted programs in the long run. Third, adding teachers' effort choices to our framework would capture an additional dimension along which incentive programs may improve efficiency and equity.

References

- Abbott, B., G. Gallipoli, C. Meghir, and G. L. Violante (2019). Education policy and intergenerational transfers in equilibrium. *Journal of Political Economy* 127(6), 2569–2624.
- Acemoglu, D. and J.-S. Pischke (1998). Why do firms train? theory and evidence. *The Quarterly Journal of Economics* 113(1), 79–119.
- Agarwal, N. (2015). An empirical model of the medical match. *American Economic Review* 105(7), 1939–1978.
- Atkinson, A., S. Burgess, B. Croxson, P. Gregg, C. Propper, H. Slater, and D. Wilson (2009). Eval-

- uating the impact of performance-related pay for teachers in England. *Labour Economics* 16(3), 251–261.
- Badger, J. R., M. Zaneva, R. P. Hastings, M. R. Broome, R. Hayes, P. Patterson, N. Rose, S. Clarkson, J. Hutchings, and L. Bowes (2023). Associations between school-level disadvantage, bullying involvement and children’s mental health. *Children* 10(12), 1852.
- Barseghyan, L., M. Coughlin, F. Molinari, and J. C. Teitelbaum (2021). Heterogeneous choice sets and preferences. *Econometrica* 89(5), 2015–2048.
- Bates, M. (2020). Public and private employer learning: Evidence from the adoption of teacher value added. *Journal of Labor Economics* 38(2), 375–420.
- Bates, M., M. Dinerstein, A. C. Johnston, and I. Sorokin (2025). Teacher labor market policy and the theory of the second best. *The Quarterly Journal of Economics* 140(2), 1417–1469.
- Behrman, J. R., M. M. Tincani, P. E. Todd, and K. I. Wolpin (2016). Teacher quality in public and private schools under a voucher system: The case of Chile. *Journal of Labor Economics* 34(2), 319–362.
- Biasi, B., C. Fu, and J. Stromme (2024). Equilibrium in the market for public school teachers: District wage strategies and teacher comparative advantage. Technical report, National Bureau of Economic Research.
- Bobba, M., T. Ederer, G. Leon-Ciliotta, C. Neilson, and M. G. Nieddu (2024). Teacher compensation and structural inequality: Evidence from centralized teacher school choice in Perú. Technical report, National Bureau of Economic Research.
- Bognanno, M. and E. Melero (2016). Promotion signals, experience, and education. *Journal of Economics & Management Strategy* 25(1), 111–132.
- Boyd, D., H. Lankford, S. Loeb, and J. Wyckoff (2005). Explaining the short careers of high-achieving teachers in schools with low-performing students. *American Economic Review* 95(2), 166–171.
- Boyd, D., H. Lankford, S. Loeb, and J. Wyckoff (2013). Analyzing the determinants of the matching of public school teachers to jobs: Disentangling the preferences of teachers and employers. *Journal of Labor Economics* 31(1), 83–117.
- Brehm, M., S. A. Imberman, and M. F. Lovenheim (2017). Achievement effects of individual performance incentives in a teacher merit pay tournament. *Labour Economics* 44, 133–150.
- Cassidy, H., J. DeVaro, and A. Kauhanen (2016). Promotion signaling, gender, and turnover: New theory and evidence. *Journal of Economic Behavior & Organization* 126, 140–166.
- Chang, C. and Y. Wang (1996). Human capital investment under asymmetric information: The pigovian conjecture revisited. *Journal of Labor Economics* 14(3), 505–519.
- Chetty, R., J. N. Friedman, and J. E. Rockoff (2014). Measuring the impacts of teachers i: Evaluating bias in teacher value-added estimates. *American Economic Review* 104(9), 2593–2632.
- Chingos, M. M. and M. R. West (2011). Promotion and reassignment in public school districts: How do schools respond to differences in teacher effectiveness? *Economics of Education Review* 30(3), 419–433.
- Clotfelter, C. T., H. F. Ladd, and J. Vigdor (2005). Who teaches whom? Race and the distribution

- of novice teachers. *Economics of Education Review* 24(4), 377–392.
- Cullen, J. B., C. Koedel, and E. Parsons (2021). The compositional effect of rigorous teacher evaluation on workforce quality. *Education Finance and Policy* 16(1), 7–41.
- Dee, T. S. and J. Wyckoff (2015). Incentives, selection, and teacher performance: Evidence from impact. *Journal of Policy Analysis and Management* 34(2), 267–297.
- DeVaro, J. and M. Waldman (2012). The signaling role of promotions: Further theory and empirical evidence. *Journal of Labor Economics* 30(1), 91–147.
- Dolton, P. and W. v. d. Klaauw (1999). The turnover of teachers: A competing risks explanation. *Review of Economics and Statistics* 81(3), 543–550.
- Duflo, E., R. Hanna, and S. P. Rya (2012). Incentives work: Getting teachers to come to school. *American Economic Review* 102(4), 1241–1278.
- EACEA (2021). Teachers in europe: Careers, development and well-being. Technical report, Education, Audiovisual and Culture Executive Agency (EACEA), European Commission, Luxembourg. Eurydice Report.
- Ederer, T. (2023). Labor market dynamics and teacher spatial sorting. Technical report, Working Paper.
- Epple, D. and M. M. Ferreyra (2008). School finance reform: Assessing general equilibrium effects. *Journal of Public Economics* 92(5-6), 1326–1351.
- Epple, D. and H. Sieg (1999). Estimating equilibrium models of local jurisdictions. *Journal of Political Economy* 107(4), 645–681.
- Epple, D. N. and R. Romano (2003). Neighborhood schools, choice, and the distribution of educational benefits. In *The economics of school choice*, pp. 227–286. University of Chicago Press.
- European Commission (2018). Boosting teacher quality. Technical report, European Commission. EU study identifying teacher appraisal and incentive systems as key policy instruments in teacher quality policy.
- Fabre, A. (2023). The geography of higher education and spatial inequalities. Technical report, Institute for Fiscal Studies.
- Ferreyra, M. M. (2007). Estimating the effects of private school vouchers in multidistrict economies. *American Economic Review* 97(3), 789–817.
- Fox, J. T. (2007). Semiparametric estimation of multinomial discrete-choice models using a subset of choices. *The RAND Journal of Economics* 38(4), 1002–1019.
- Fryer, R. G. (2013). Teacher incentives and student achievement: Evidence from New York City Public Schools. *Journal of Labor Economics* 31(2), 373–407.
- Gibbons, R. and L. F. Katz (1991). Layoffs and lemons. *Journal of Labor Economics* 9(4), 351–380.
- Glewwe, P., N. Ilias, and M. Kremer (2010). Teacher incentives. *American Economic Journal: Applied Economics* 2(3), 205–227.
- Greenwald, B. (1986). Adverse selection in the labour market. *Review of Economic Studies* 53, 325–47.
- Hanushek, E. A. (2011). The economic value of higher teacher quality. *Economics of Education*

Review 30(3), 466–479.

- Hanushek, E. A., J. F. Kain, and S. G. Rivkin (2004). Why public schools lose teachers. *Journal of Human Resources* 39(2), 326–354.
- Hanushek, E. A. and S. G. Rivkin (2010). Generalizations about using value-added measures of teacher quality. *American Economic Review* 100(2), 267–271.
- Hu, L. and C. Taber (2011). Displacement, asymmetric information, and heterogeneous human capital. *Journal of Labor Economics* 29(1), 113–152.
- Imberman, S. A. and M. F. Lovenheim (2015). Incentive strength and teacher productivity: Evidence from a group-based teacher incentive pay system. *Review of Economics and Statistics* 97(2), 364–386.
- Ingersoll, R. M., C. M. Hoxby, and A. F. Scrupski (2004). Why some schools have more under-qualified teachers than others. *Brookings papers on education policy* (7), 45–88.
- Jackson, C. K. (2018). What do test scores miss? the importance of teacher effects on non-test score outcomes. *Journal of Political Economy* 126(5), 2072–2107.
- Jackson, C. K., J. E. Rockoff, and D. O. Staiger (2014). Teacher effects and teacher-related policies. *Annual Review of Economics* 6(1), 801–825.
- Jacob, B. A. (2007). The challenges of staffing urban schools with effective teachers. *The Future of Children*, 129–153.
- Jacob, B. A. and L. Lefgren (2008). Can principals identify effective teachers? evidence on subjective performance evaluation in education. *Journal of Labor Economics* 26(1), 101–136.
- Kahn, L. B. (2013). Asymmetric information between employers. *American Economic Journal: Applied Economics* 5(4), 165–205.
- Kane, T. J. and D. O. Staiger (2008). Estimating teacher impacts on student achievement: An experimental evaluation. Technical report, National Bureau of Economic Research.
- Kennan, J. and J. R. Walker (2011). The effect of expected income on individual migration decisions. *Econometrica* 79(1), 211–251.
- Kraft, M. A., E. J. Brunner, S. M. Dougherty, and D. J. Schwegman (2020). Teacher accountability reforms and the supply and quality of new teachers. *Journal of Public Economics* 188, 104212.
- Laing, D. (1994). Involuntary layoffs in a model with asymmetric information concerning worker ability. *The Review of Economic Studies* 61(2), 375–392.
- Lang, K. and M. D. Palacios (2018). The determinants of teachers’ occupational choice. Technical report, National Bureau of Economic Research.
- Lankford, H., S. Loeb, and J. Wyckoff (2002). Teacher sorting and the plight of urban schools: A descriptive analysis. *Educational evaluation and policy analysis* 24(1), 37–62.
- Laverde, M., E. Mykerezi, A. Sojourner, and A. Sood (2025). Gains from alternative assignment? evidence from a two-sided teacher market.
- Lavy, V. C. (2002). Evaluating the effect of teachers’ group performance incentives on pupil achievement. *Journal of Political Economy* 110(6), 1286–1317.
- Lazear, E. P. (1986). Salaries and piece rates. *Journal of Business*, 405–431.

- Loeb, S., L. C. Miller, and J. Wyckoff (2015). Performance screens for school improvement: The case of teacher tenure reform in new york city. *Educational Researcher* 44(4), 199–212.
- Mahvar, T., M. A. Farahani, and A. Aryankhesal (2018). Conflict management strategies in coping with students’ disruptive behaviors in the classroom: Systematized review. *Journal of Advances in Medical Education & Professionalism* 6(3), 102.
- Mansfield, R. K. (2015). Teacher quality and student inequality. *Journal of Labor Economics* 33(3), 751–788.
- McFadden, D. (1978). Modelling the choice of residential location in a. karlqvist, l. lundqvist, f. snickars and j. weibull, eds, spatial interaction theory and planning models. *Amsterdam: North Holland* 75, 96.
- Menzel, K. (2015). Large matching markets as two-sided demand systems. *Econometrica* 83(3), 897–941.
- Milgrom, P. and S. Oster (1987). Job discrimination, market forces, and the invisibility hypothesis. *The Quarterly Journal of Economics* 102(3), 453–476.
- Morgan, A., M. Nguyen, E. A. Hanushek, B. Ost, and S. Rivkin (2026). Attracting and retaining highly effective educators in hard-to-staff schools. *American Economic Journal: Economic Policy*. Forthcoming.
- Muralidharan, K. and V. Sundararaman (2011). Teacher performance pay: Experimental evidence from india. *Journal of Political Economy* 119(1), 39–77.
- National Education Association (2023). Teacher tenure & due process protections for educators. <https://www.nea.org/resource-library/teacher-tenure-due-process-protections-educators>. State survey showing that Florida, North Carolina, and Wisconsin have effectively eliminated traditional tenure protections for new teachers.
- Neal, D. (2011). The design of performance pay in education. *Handbook of the Economics of Education* 4, 495–550.
- Pinkston, J. C. (2009). A model of asymmetric employer learning with testable implications. *The Review of Economic Studies* 76(1), 367–394.
- Riordan, M. H. and R. W. Staiger (1993). Sectoral shocks and structural unemployment. *International Economic Review*, 611–629.
- Rivkin, S. G., E. A. Hanushek, and J. F. Kain (2005). Teachers, schools, and academic achievement. *Econometrica* 73(2), 417–458.
- Rockoff, J. E., D. O. Staiger, T. J. Kane, and E. S. Taylor (2012). Information and employee evaluation: Evidence from a randomized intervention in public schools. *American Economic Review* 102(7), 3184–3213.
- Rodriguez, L. A., W. A. Swain, and M. G. Springer (2020). Sorting through performance evaluations: The influence of performance evaluation reform on teacher attrition and mobility. *American Educational Research Journal* 57(6), 2339–2377.
- Scafidi, B., D. L. Sjoquist, and T. R. Stinebrickner (2007). Race, poverty, and teacher mobility. *Economics of Education Review* 26(2), 145–159.

- Schönberg, U. (2007). Testing for asymmetric employer learning. *Journal of Labor Economics* 25(4), 651–691.
- Staiger, D. O. and J. E. Rockoff (2010). Searching for effective teachers with imperfect information. *Journal of Economic Perspectives* 24(3), 97–118.
- Steinberg, M. P. and L. Sartain (2015). Does teacher evaluation improve school performance? experimental evidence from chicago’s excellence in teaching project. *Education Finance and Policy* 10(4), 535–572.
- Stinebrickner, T. R. (2001a). Compensation policies and teacher decisions. *International Economic Review* 42(3), 751–780.
- Stinebrickner, T. R. (2001b). A dynamic model of teacher labor supply. *Journal of Labor Economics* 19(1), 196–230.
- Tincani, M. M. (2021). Teacher labor markets, school vouchers, and student cognitive achievement: Evidence from chile. *Quantitative Economics* 12(1), 173–216.
- Tippett, N. and D. Wolke (2014). Socioeconomic status and bullying: A meta-analysis. *American Journal of Public Health* 104(6), e48–e59.
- Waldman, M. (1984). Job assignments, signalling, and efficiency. *The Rand Journal of Economics* 15(2), 255–267.
- Waldman, M. and O. Zax (2016). An exploration of the promotion signaling distortion. *The Journal of Law, Economics, and Organization* 32(1), 119–149.
- Wiswall, M. (2007). Licensing and occupational sorting in the market for teachers. *Unpublished manuscript, Department of Economics, New York University*.
- Wu, A. H. (2025). Reveal or conceal? employer learning in the labor market for computer scientists. Technical report, Working Paper.
- Zhang, Y. (2007). Employer learning under asymmetric information: The role of job mobility. Available at SSRN 1058801.

A Appendix

A.1 Equilibrium Objects: $P_s(\cdot)$, $h_s(\cdot)$ and $\bar{B}_s(\cdot)$

It follows from teacher’s optimization that *if receiving an offer from school s* , a teacher with (x, q, v, s_0) would accept this offer with probability

$$P_s(x, q, v, s_0) =$$

$$\left\{ \begin{array}{l} \frac{\exp\left(\frac{U_s(x, q, v, s_0)}{\sigma_\epsilon}\right)}{\exp\left(\frac{U_s(x, q, v, s_0)}{\sigma_\epsilon}\right) + \sum_{s' \in \{0, \dots, S\} \setminus s} O_{s'}(\Gamma_{s'}(x, q, v, s_0)) \exp\left(\frac{U_{s'}(x, q, v, s_0)}{\sigma_\epsilon}\right)} \text{ for } s_0 > 0, \\ \frac{\exp\left(\frac{U_s(x, q, v, 0) + \beta EV_1(x, q, v, s_0 = s, \epsilon')}{\sigma_\epsilon}\right)}{\exp\left(\frac{U_s(x, q, v, 0) + \beta EV_1(x, q, v, s_0 = s, \epsilon')}{\sigma_\epsilon}\right) + \sum_{s' \in \{0, \dots, S\} \setminus s} O_{s'}(\Gamma_{s'}(x, q, v, 0)) \exp\left(\frac{U_{s'}(x, q, v, 0) + \beta EV_1(x, q, v, s_0 = s', \epsilon')}{\sigma_\epsilon}\right)} \text{ for } s_0 = 0. \end{array} \right. \quad (11)$$

From a school's perspective, the acceptance probability based on its information $\Gamma_s(\cdot)$ is $h_s(\Gamma_s(x, q, v, s_0)) = E[P_s(x, q, v, s_0) | \Gamma_s(x, q, v, s_0)] =$

$$\begin{cases} P_s(x, q, v, s_0) & \text{if } \Gamma_s(x, q, v, s_0) = \{x, q, v, s_0\}, \\ \int P_s(x, q', v', s_0) dF(q', v' | x, A(q), s_0) & \text{if } \Gamma_s(x, q, v, s_0) = \{x, A(q), s_0\}, \\ \int P_s(x, q', v', s_0) dF(q', v' | x, s_0) & \text{if } \Gamma_s(x, q, v, s_0) = \{x, s_0\}, \end{cases} \quad (12)$$

where we use a prime to highlight a variable that needs to be integrated out. The first two rows of (12) are about district incumbent teachers $s_0 > 0$. Before ETI, Row 1 applies to a teacher's employer and Row 2 to other schools; after ETI, Row 1 applies to all schools. Row 3 of (12) is about potential entrant teachers ($s_0 = 0$), applicable to all schools both before and after ETI. Similarly, the inferred value of the teacher conditional on her acceptance is:

$$\bar{B}_s(\Gamma_s(x, q, v, s_0)) = E[B(x, q, v, z_s) | \Gamma_s(x, q, v, s_0), \text{accept}] =$$

$$\begin{cases} B(x, q, v, z_s) & \text{if } \Gamma_s(x, q, v, s_0) = \{x, q, v, s_0\} \\ \frac{1}{h_s(\Gamma_s(x, q, v, s_0))} \int B(x, q, v', z_s) P_s(x, q', v', s_0) dF(q', v' | x, q, s_0) & \text{if } \Gamma_s(x, q, v, s_0) = \{x, A(q), s_0\} \\ \frac{1}{h_s(\Gamma_s(x, q, v, 0))} \int B(x, q', v', z_s) P_s(x, q', v', s_0) dF(q', v' | x, s_0) & \text{if } \Gamma_s(x, q, v, s_0) = \{x, s_0\}. \end{cases} \quad (13)$$

A.2 Detailed Functional Forms

A.2.1 Teachers' Payoff from the Outside Option

Teachers' outside value may differ for district incumbents and potential entrants:

$$u_0(x, q, v, s_0) = I(s_0 = 0)(\varphi_{10} + \varphi_{20}x_1 + \varphi_{30}x_2 + \varphi_{40}\mathbb{I}(x_1 \geq 30) + \varphi_{50}q + \varphi_{60}v) + \\ I(s_0 > 0)(\varphi_{11} + \varphi_{21}x_1 + \varphi_{31}x_2 + \varphi_{41}\mathbb{I}(x_1 \geq 30) + \varphi_{51}q + \varphi_{61}v)$$

A.2.2 Schools' Preferences for Teachers

$$B(x, q, v, z_s) = xb_0 + \exp(b_1 z_{s1}) \ln(q - \underline{q}) + \exp(b_2 + b_3 z_{s1}) \ln(v - \underline{v}).$$

Schools' preferences for q and v are assumed to be concave (in logs) with coefficients given by $\exp(b_1 z_{s1})$ and $\exp(b_2 + b_3 z_{s1})$, respectively; a teacher's value is normalized to be in units of $(\log) q$. Together with $b_0 \geq 0$, such parameterization guarantees that $B(\cdot)$ is weakly increasing. We subtract q by a fixed value \underline{q} to ensure that $\ln(q - \underline{q}) > 0$; we treat v in the same way.⁴¹ As such, all teachers yield non-negative values, which implies that schools' capacities are binding in the equilibrium. This re-scaling does not affect the relative ranking of $q(v)$ among teachers.

⁴¹ \underline{q} and \underline{v} are set at 0.5 standard deviations below the sample min of q and v , respectively.

Table A1: Other Parameter Estimates: Outside Option Values

Period: Information:	Post-ETI		Pre-ETI			
	Full Info		q Known		ASPIRE Known	
Outside option for incumbents						
Constant	110.33	(8.00)	98.71	(4.25)	99.26	(9.05)
Experience	1.84	(0.25)	1.77	(0.26)	1.76	(0.34)
Retirement age	6.73	(2.86)	2.06	(0.81)	1.87	(2.32)
Graduate degree	5.00	(2.89)	0.43	(1.31)	0.44	(5.86)
q	17.25	(3.65)	16.22	(3.08)	16.15	(4.84)
v	4.02	(2.24)	4.91	(2.21)	4.92	(2.29)
Outside option for potential entrants						
Constant	53.16	(1.79)	52.00	(17.97)	52.03	(2.97)
Experience	0.93	(0.22)	3.03	(4.67)	2.95	(0.60)
Graduate degree	2.26	(1.79)	0.93	(1.09)	0.97	(5.34)
q	16.87	(3.96)	18.36	(1.81)	18.36	(1.28)
v	9.63	(1.76)	6.57	(2.66)	6.57	(4.52)

Each column shows estimates of parameters governing the outside option under alternative information environments. Std errors (in parentheses) are derived numerically via the Delta Method. The columns labeled “ q known” shows the case when $A(q) = q$. The columns labeled “ASPIRE Known” shows the case when $A(q)$ gives the ASPIRE reward outcomes.

Table A2: Model Fit: OLS of Teacher-School Match (post-ETI)

Teacher’s Choice Set	Inferred Offer Set ^a		All Schools ^b	
	Data	Model	Data	Model
Wage	0.001	0.001	0.0004	0.0002
Funding	0.00005	0.00008	0.00001	-0.00001
$I(s = s_0)$	0.839	0.887	0.708	0.662
$I(s = s_0) \times \text{experience}$	0.0002	-0.002	0.0031	0.006
z_1	0.004	0.004	0.002	0.0005
$z_1 \times q$	-0.002	-0.001	0.0007	0.002
$z_1 \times v$	-0.002	-0.002	-0.00009	0.0008
$I(s = 0)$	0.264	0.256	0.236	0.230
$I(s = 0) \times \text{experience}$	-0.004	-0.001	-0.003	-0.002
$I(s = 0) \times \text{retirement age}$	0.104	0.094	0.088	0.106
$I(s = 0) \times \text{grad degree}$	0.034	0.037	0.034	0.033
$I(s = 0) \times q$	-0.002	0.003	-0.001	-0.013
$I(s = 0) \times v$	-0.027	-0.026	-0.048	-0.036
#Obs.	139,452	-	275,400	-

^{a(b)}: OLS specified in Aux 1a (1b), with teacher fixed effects: data vs model, post-ETI.

Table A3: Model Fit: Outcome Sorting, Entry, Exit

A. Average Employee Characteristics post-ETI (Aux 2)							
School Group by z_1	Experience		q		v		
	Data	Model	Data	Model	Data	Model	
Quartile 1	8.00	7.90	-0.46	-0.45	-0.24	-0.25	
Quartile 2	8.85	8.87	-0.18	-0.17	-0.15	-0.15	
Quartile 3	9.99	9.67	0.01	0.07	-0.09	-0.03	
Quartile 4	10.49	10.56	0.24	0.26	0.03	0.06	

B. Entry, Exit, and Switching post-ETI (Aux 3)					
	Entrant (x, q, v)		Entrant $(x, q, v) \times z_1$		
	Data	Model	Data	Model	
Experience	4.54	4.51	3.14	3.24	
Grad Deg.	0.23	0.20	0.15	0.14	
q	-0.34	-0.36	-0.20	-0.22	
v	-0.53	-0.55	-0.38	-0.38	

	Data	Model
$E(z_1 entrants)$	0.70	0.72
$\frac{\#Entrants}{\#Stayers}$	23.9%	22.8%
Exit rate	18.8%	20.7%
Switch Rate	6.6%	6.4%

Panel A shows auxiliary model 2: data vs model, post-ETI.

Panel B shows auxiliary model 3: data vs. model, post-ETI.

Table A4: Model Validation: Untargeted OLS of Teacher-School Match (pre-ETI)

	Data	Model	
		q Known	ASPIRE Known
Wage	-0.0001	-0.0001	0.0001
Funding	-0.0001	0.00001	0.00001
$I(s = s_0)$	0.8660	0.7739	0.7840
$I(s = s_0) \times \text{experience}$	0.0024	0.0076	0.0072
z_1	0.0001	0.0010	0.0012
$z_1 \times q$	0.0006	0.0023	0.0017
$z_1 \times v$	-0.0001	0.0007	0.0005
$I(s = 0)$	0.0959	0.0968	0.1052
$I(s = 0) \times \text{experience}$	-0.0028	-0.0025	-0.0025
$I(s = 0) \times \text{retirement age}$	0.0820	0.0624	0.0641
$I(s = 0) \times \text{grad degree}$	-0.0210	-0.0111	-0.0094
$I(s = 0) \times q$	-0.0195	-0.0237	-0.0186
$I(s = 0) \times v$	-0.0097	-0.0125	-0.0109

$a(b)$: OLS specified in Aux 1b, teacher fixed effects included: data vs model, pre-ETI, with re-estimated outside value parameters. The column “ q Known” shows the case when $A(q) = q$. The column “ASPIRE Known” shows the case when $A(q)$ gives the ASPIRE reward outcomes.

Table A5: Model Validation: Untargeted Outcome Sorting (pre-ETI)

School Group by z_1	q			v		
	Data	q Known	ASPIRE Known	Data	q Known	ASPIRE Known
Quartile 1	-0.18	-0.27	-0.26	-0.08	-0.11	-0.11
Quartile 2	0.08	0.08	0.09	-0.11	-0.08	-0.07
Quartile 3	0.11	0.10	0.09	-0.02	-0.05	-0.05
Quartile 4	0.14	0.18	0.16	0.03	0.03	0.02

School Group by z_1	Experience		
	Data	q Known	Aspire Known
Quartile 1	10.88	10.84	11.01
Quartile 2	10.40	10.42	10.47
Quartile 3	10.33	10.40	10.35
Quartile 4	10.74	10.54	10.36

Auxiliary model 2: data vs model, pre-ETI, with re-estimated outside value parameters.

The column “ q Known” shows the case when $A(q) = q$. The column “ASPIRE Known” shows the case when $A(q)$ gives the ASPIRE reward outcomes.

Table A6: Model Validation: Pre-ETI Sorting with All Post-ETI Parameter Estimates

School Group by z_1	Relative q			Relative v		
	Data	q Known	Aspire Known	Data	q Known	ASPIRE Known
Quartile 1	-0.22	-0.30	-0.28	-0.03	-0.08	-0.05
Quartile 2	0.04	0.02	0.03	-0.06	-0.03	-0.03
Quartile 3	0.07	0.08	0.08	0.03	0.01	0.01
Quartile 4	0.10	0.19	0.16	0.07	0.09	0.07
Average School	0.040	0.023	0.020	-0.043	-0.172	-0.169

Relative teacher quality: data vs model, pre-ETI. The last row shows the teacher quality in an average school, data versus model. Rows 1-4 show the average teacher quality (q, v) for schools in a given quartile of z_1 relative to that in an average school. The column “ q Known” shows the case when $A(q) = q$. The column “ASPIRE Known” shows the case when $A(q)$ gives the ASPIRE reward outcomes.