Equilibrium in the Market for Public School Teachers: District Wage Strategies and Teacher Comparative Advantage (Online Appendix)

Barbara Biasi, Chao Fu and John Stromme*

April, 2022

B1 Algorithms

Teachers' decision rule implies that if District d makes an offer to the teacher, the teacher's acceptance probability is given by

$$h_d\left(x, c, d_0\right) = \frac{\exp\left(\frac{V_d(x,c,d_0)}{\sigma\epsilon}\right)}{\exp\left(\frac{V_d(x,c,d_0)}{\sigma_\epsilon}\right) + \sum_{d' \in D \backslash d} o_{d'}\left(x, c, d_0\right) \exp\left(\frac{V_{d'}(x,c,d_0)}{\sigma\epsilon}\right)}. \tag{1}$$

We assume that districts make decisions based on a simplified belief, given by

$$\begin{aligned} \widetilde{h}_d\left(x, c, d_0 | \overline{w}\left(x, c\right), \sigma_w\left(x, c\right)\right) &= \frac{1}{1 + \exp\left(f\left(x, c, d_0, w_d, q_d, \lambda_d\right)\right)}, \\ \text{with } f\left(\cdot\right) &= x\zeta_1 + \zeta_2 \frac{c_1 + c_2}{2} + \zeta_3 \left(\frac{w_d - \overline{w}\left(x, c\right)}{\sigma_{w(x,c)}}\right) + \zeta_4 q_d + \zeta_5 e^{\lambda_d} + \zeta_6 \lambda_d c_1 \\ &\quad + \left(1 - I\left(d_0 = 0\right)\right)\left[I\left(d \neq d_0\right)\left(\zeta_7 + \zeta_8 x_1\right) + \zeta_9 I\left(z_d \neq z_{d_0}\right)\right], \end{aligned} \tag{2}$$

*Biasi: Yale School of Management and NBER; Fu: University of Wisconsin and NBER, cfu@ssc.wisc.edu; Stromme: University of Wisconsin.

where $\overline{w}(x,c)$ and $\sigma_w(x,c)$ are the mean and standard deviation of wages across all districts for a teacher with (x,c), i.e.,

$$\overline{w}(x,c) \equiv \frac{1}{D}\sum_d w_d(x,c;\omega_d) \tag{3}$$

$$\sigma_{w(x,c)} \equiv \sqrt{\frac{1}{D-1}\sum_d \left(w_d(x,c;\omega_d) - \overline{w}(x,c)\right)^2}. \tag{4}$$

An equilibrium requires beliefs $\widetilde{h}_d(x,c,d_0)$, and in particular the vector ζ and the wage statistics $\left\{\overline{w}(x,c), \sigma_{w(x,c)}\right\}_{x,c}$, to be consistent with decisions made by teachers and districts.

B1.1 Estimation Algorithm

The estimation algorithm involves an outer loop searching for the parameter vector Θ and an inner loop solving the model for each given Θ. This inner loop does not require finding the fixed point for all components in $\{\zeta, \overline{w}(\cdot), \sigma_w(\cdot)\}$: Assuming that data were generated from an equilibrium, $\{\overline{w}(\cdot)\}$ and $\{\sigma_w(\cdot)\}$ can be derived directly from the observed district wage schedules $\{\omega_d^o\}_d$, where the superscript o denotes "observed." For estimation, one only needs to find the fixed point for ζ; the observed equilibrium wage statistics $\{\overline{w}^o(\cdot), \sigma_w^o(\cdot)\}$ can be plugged directly into the belief function (2). Given a parameter vector Θ, the inner loop of the estimation algorithm involves the following steps.

1. Search for $\zeta^*(\Theta)$

   (a) Guess ζ, which, together with $\overline{w}^o(\cdot)$ and $\sigma_w^o(\cdot)$, implies a belief $\left\{\widetilde{h}_d\left(\cdot|\zeta, \overline{w}^o(\cdot), \sigma_w^o(\cdot)\right)\right\}$ as defined in (2).

   (b) Given $\widetilde{h}_d\left(\cdot|\zeta, \overline{w}^o(\cdot), \sigma_w^o(\cdot)\right)$, solve for the optimal job offers $o_d^*(\cdot;\omega_d^o)$ under the observed ω_d^o for each district d.

   (c) Given the job offers and the wages implied by $\{o_d^*(\cdot;\omega_d^o), \omega_d^o\}_d$, calculate each teacher's acceptance probabilities $h_d(\cdot)$ for each d, as in (1), and the distance $\left\|h(\cdot) - \widetilde{h}\left(\cdot|\zeta, \overline{w}^o(\cdot), \sigma_w^o(\cdot)\right)\right\|$.

   (d) Repeat Steps 1a-1c until $\left\|h(\cdot) - \widetilde{h}\left(\cdot|\zeta, \overline{w}^o(\cdot), \sigma_w^o(\cdot)\right)\right\|$ is below a tolerance level; the associated ζ is the consistent belief parameter vector $\zeta^*(\Theta)$.

2. Given job offers $\{o_d^*(\cdot;\omega_d^o)\}_d$ under $\widetilde{h}_d\left(\cdot|\zeta^*(\Theta), \overline{w}^o(\cdot), \sigma_w^o(\cdot)\right)$ and wages implied by $\{\omega_d^o\}$, each teacher chooses the most preferred among their received offers. The implied teacher-district matches will be compared with the observed matches in the outer loop.

3. Given $\widetilde{h}_d\left(\cdot|\zeta^*(\Theta), \overline{w}^o(\cdot), \sigma_w^o(\cdot)\right)$, each district makes optimal decisions on its wage schedule $\omega_d^*(\Theta)$.[1] The resulting $\{\omega_d^*(\Theta)\}_d$ will be compared with the observed $\{\omega_d^o\}_d$ in the outer loop.

B1.2 Solving for the Equilibrium

Both the teacher-specific wage statistics $\left\{\left(\overline{w}(x,c), \sigma_{w(x,c)}\right)\right\}_{x,c}$ and the wage rules $\{(\omega_{d1}, \omega_{d2})\}_d$ that govern these statistics are high-dimensional objects. However, notice that districts' wages are given by

$$w_d(x,c;\omega) = \begin{cases} \underline{w} \text{ if } \omega_1 W_d^0(x) + \omega_2\left[\lambda_d c_1 + (1-\lambda_d)c_2\right] < \underline{w} \\ \overline{w} \text{ if } \omega_1 W_d^0(x) + \omega_2\left[\lambda_d c_1 + (1-\lambda_d)c_2\right] > \overline{w} \\ \omega_1 W_d^0(x) + \omega_2\left[\lambda_d c_1 + (1-\lambda_d)c_2\right] \text{ otherwise} \end{cases}, \tag{5}$$

where the pre-reform wage schedule $W_d^0(x)$ is a linear function of experience categories (x_1) and the MA dummy (x_2). It follows that the mean wage is a linear function of the following form governed by some parameter vector θ^1

$$\widetilde{\overline{w}}(x,c) = \begin{cases} \underline{w} \text{ if } \sum_n \theta_{1n}^1 I(x_1 = n) + \theta_2^1 x_2 + \theta_3^1 c_1 + \theta_4^1 c_2 < \underline{w} \\ \overline{w} \text{ if } \sum_n \theta_{1n}^1 I(x_1 = n) + \theta_2^1 x_2 + \theta_3^1 c_1 + \theta_4^1 c_2 > \overline{w} \\ \sum_n \theta_{1n}^1 I(x_1 = n) + \theta_2^1 x_2 + \theta_3^1 c_1 + \theta_4^1 c_2 \text{ otherwise.} \end{cases} \tag{6}$$

Similarly, the cross-district wage standard deviation for a teacher will be the square root of a quadratic function (Q), governed by some parameter vector θ^2, and bounded from above by the largest possible wage spread, i.e.,

$$\widetilde{\sigma}_{w(x,c)} = \min\left\{\sqrt{\max\left\{Q\left(x_1, x_2, c_1, c_2; \theta^2\right), 0\right\}}, \overline{w} - \underline{w}\right\}. \tag{7}$$

Instead of searching for fixed points of $\left\{\{h_d(x,c,d_0)\}_{x,c}, (\omega_{d1}, \omega_{d2})\right\}_d$, one can search for parameter vectors ζ, θ^1, and θ^2 in (2), (6) and (7) to guarantee equilibrium consistency. Note that ζ, θ^1, and θ^2 are not structural parameters; rather, they serve to summarize the equilibrium under a given policy scenario and are policy dependent. We now describe the algorithm we use to simulate the equilibrium outcomes, for a given policy environment.

[1] We assume that changing a single district's wage for Teacher i has a negligible effect on wage statistics $\left(\overline{w}^o(x_i, c_i), \sigma_w^o(x_i, c_i)\right)$, i.e., the mean and standard deviation of Teacher i's wage across the 411 districts in our sample.

B1.2.1 Equilibrium Algorithm

We draw M economies, each with D districts and N teachers. All economies share the same observable teacher and district characteristics as those in the data, but each economy is assigned a different realization of wage-choice-specific shocks $\{\{\eta_{d\omega}\}_{\omega}\}_{d}$, drawn from the i.i.d. extreme value distribution, with the scaling parameter σ_{η}. The expected equilibrium outcomes are calculated as the average outcomes across M economies. For each economy m, we apply the following procedure.

1. Guess parameters ζ, θ^1, and θ^2, which imply $\left\{\widetilde{\overline{w}}(x,c), \widetilde{\sigma}_{w(x,c)}, \widetilde{h}_d\left(x,c,d_0|\widetilde{\overline{w}}(x,c), \widetilde{\sigma}_{w(x,c)}\right)\right\}$ from (2), (6) and (7).

2. Given $\left\{\widetilde{h}_d\left(x,c,d_0|\widetilde{\overline{w}}(x,c), \widetilde{\sigma}_{w(x,c)}\right)\right\}$, each district d chooses its optimal wage and offer policies $\{\omega_d, O(\omega_d)\}$.

3. Given $\{\omega_d, O(\omega_d)\}_d$, compute teacher acceptance probabilities $h_d(\cdot)$ from their decision rules (1), the mean wage $\overline{w}(x,c)$ based on (3), and standard deviation $\sigma_{w(x,c)}$ based on (4).

4. Calculate the distance between $\left\{\widetilde{\overline{w}}(x,c), \widetilde{\sigma}_{w(x,c)}, \widetilde{h}_d\left(x,c,d_0|\widetilde{\overline{w}}(x,c), \widetilde{\sigma}_{w(x,c)}\right)\right\}$ and $\left\{\overline{w}(x,c), \sigma_{w(x,c)}, h_d\left(x,c,d_0|\left(\overline{w}(x,c), \sigma_{w(x,c)}\right)\right)\right\}$.

5. Repeat Step 1 to Step 4 and search for $\{\zeta^*, \theta^{1*}, \theta^{2*}\}$ that bring the distance in Step 4 below a tolerance level. The vector $\{\zeta^*, \theta^{1*}, \theta^{2*}\}$ renders the consistent belief (2). Equilibrium outcomes in economy m consist of the decisions made by districts and teachers under this consistent belief.

B2 Data Details

B2.1 Sample Construction

We construct our samples as follows. For estimation and empirical analysis, we focus on full-time Grades 4-6 math teachers employed in Wisconsin school districts in 2014 (411 districts and 6,625 individuals).[2] We exclude 3 teachers from the sample, whose schools did not report test scores. We also exclude 22 teachers with missing information on years of experience. This leaves us with 6,600 teachers and 411 districts in the final estimation sample.

For the validation sample, we focus on 6,751 full-time Grades 4-6 math teachers employed in

[2] Wisconsin had 424 school districts in 2014, 11 of which did not have any elementary school, and 2 of which did not have any full-time Grades 4-6 math teachers.

411 districts in 2010. We exclude 10 teachers with missing information on years of experience. This leaves us with 6,741 teachers and 411 districts in the final validation sample.

B2.2 Teacher's Previous District

Our model requires identifying the district where each teacher was working at the beginning of the model period (d_{i0}). For the estimation sample, which is based on 2014 data, we define d_{i0} as follows. If the teacher never moved or moved only once between 2011 and 2014, d_{i0} is the district where she was employed in 2011. If a teacher moved more than once between 2011 and 2014, we set d_{i0} to be the last employer she worked for before 2014. For example, if teacher i worked in District A in 2011 and 2012, and District C in 2013 and 2014, then $d_{i0} = C$. If teacher i worked in District A in 2011 and 2012, in District B in 2013, and in District C in 2014, then $d_{i0} = B$.
For the validation sample, based on data from 2010, we obtain teachers' d_{i0} following the same procedure as above, using a teacher's employment history between 2007 and 2010.

B2.3 Teacher Effectiveness

Students were tested on math and language in the Wisconsin Knowledge and Concepts Examination (WKCE, 2007-2014) and Badger test (2015-2016); we focus on their math scores. The WKCE was administered in November of each school year, whereas the Badger test was administered in March. To account for this change, for the years 2007–2014 we assign each student a score equal to the average of the standardized scores for the current and the following year. The test score data also include individual characteristics of test takers, such as gender, race and ethnicity, socioeconomic (SES) status, migration status, English-learner status, and disability status.

Our data allow us to link students and teachers up to the school-grade level, rather than the classroom level. To account for this data structure, we estimate two student achievement models and derive teacher effectiveness measures from each of them. In the following, we first describe the achievement model used in our empirical analysis, and its estimation and identification. The distribution of effectiveness measures estimated with this achievement model is summarized in Tables B8, B9 and Figure B3. Next, we describe the alternative model, and its estimation and identification. Finally, we show that the effectiveness measures we obtain from both models are strongly correlated and that our auxiliary models used in our structural estimation are robust to the choice of effectiveness measures.

B2.3.1 Achievement Model 1 (Main)

The effectiveness measures used in our empirical analysis are estimated using the following achievement model:

$$A_{kt} = \gamma Z_{kt}^{s} + \sum_{i:SG_{kt}=SG_{it}^{T}} \sum_{n=1}^{2} I\left(\tau_k = n\right)\left(\rho_n x_{it} + v_{in}\right) + \varepsilon_{kt} \tag{8}$$

$$= \gamma Z_{kt}^{s} + \sum_{i:SG_{kt}=SG_{it}^{T}} \sum_{n=1}^{2} I\left(\tau_k = n\right)\rho_n x_{it} + \varphi_{kt} \tag{9}$$

where A_{kt} is achievement (measured as the standardized Math test score) of student k in year t. The vector Z_{kt}^{s} contains the following: a cubic polynomial of previous year's test scores, interacted with grade fixed effects; a cubic polynomial of previous year's average test scores for k's cohort in the school, interacted with grade fixed effects; a set of student characteristics, including gender, race and ethnicity, disability status, English-language status, and socioeconomic status; the same average characteristics for student k's cohort; cohort size; grade-by-school fixed effects; and year fixed effects. The variable ε_{kt} is an i.i.d. unobservable component of achievement, idiosyncratic to each student and year. SG_{kt} (SG_{it}^{T}) denotes the school-grade of student k (teacher i) in year t. The variable τ_k equals 1 for low-achieving students and 2 for high-achieving ones; we consider a student to be low-achieving if their test score in the previous year was below the grade-specific median in the state, and high-achieving otherwise. The contribution of teacher i to the achievement of a student of type $n \in \{1, 2\}$ is $\rho_n x_{it} + v_{in}$, where x_{it} denotes i's education and experience in year t and v_{in} is the part unexplained by x_{it}.

The achievement model in (8) assumes that all teachers in a given school-grade contribute to the achievement of all students in the same school-grade. We make this choice to be able to allow x_{it} to directly enter teacher effectiveness (since experience has been shown to affect teacher effectiveness (Wiswall 2013), especially in the first years of a teacher's career (Rockoff 2004)), even if we do not observe all the teacher-student classroom links in the data. Model (8) allows us to identify the component of teacher effectiveness that depends on a teacher's experience and education.

Constructing our measures of effectiveness (c_{i1}, c_{i2}) requires estimating v_{in} and ρ_n for $n \in \{1, 2\}$. We make the following two assumptions:

A1. ε_{kt} is i.i.d. with mean 0 and variance σ_{ε}^{2}.

A2. $Cov(\varepsilon_{kt}, v_{in}) = 0 \ \forall k, i, t, n : SG_{it}^{T} = SG_{kt}$. This implies that there is no sorting on unobservables of teachers across school-grades within a district. Although there is no direct test of this assumption, in Section B3.3 we combine the approaches of Chetty et al. (2014) and Rothstein (2010) and we do not find evidence of non-random sorting.

Estimation Procedure: Model 1

1. Given A1 and A2, we estimate γ and ρ_n via OLS on equation (8), to obtain $\hat{\gamma}$ and $\hat{\rho}_n$.

2. With the estimated $\hat{\gamma}$ and $\hat{\rho}_n$, we can then estimate v_{in} using an empirical Bayes estimator similar to the one of Kane and Staiger (2008) which we adapt to take into account the structure of our data.

    (a) Let

    $$\widehat{\varphi}_{kt} = A_{kt} - \hat{\gamma} Z^s_{kt} - \sum_{i:SG_{kt}=SG^T_{it}} \sum_{n=1}^{2} \hat{\rho}_n x_{it} I(\tau_k = n). \quad (10)$$

    The quantity $\widehat{\varphi}_{kt}$ is an estimate for φ_{kt}, i.e.,

    $$\varphi_{kt} \equiv \sum_{i':SG_{kt}=SG^T_{i't}} \sum_{n=1}^{2} v_{i'n} I(\tau_k = n) + \varepsilon_{kt}.$$

    Let $K_{SG^T_{it}n}$ be the number of achievement type-n students in the school-grade that i belongs to. For each teacher i we define, for $n \in \{1, 2\}$

    $$\widehat{v}_{int} = \frac{1}{K_{SG^T_{it}n}} \sum_{k:SG_{kt}=SG^T_{it}} \widehat{\varphi}_{kt} I(\tau_k = n) \quad (11)$$

    which is an estimate of

    $$\sum_{i':SG^T_{i't}=SG^T_{it}} v_{i'n} + \frac{1}{K_{SG^T_{it}n}} \sum_{k:SG_{kt}=SG^T_{it}} \varepsilon_{kt}.$$

    This quantity corresponds to the average test score residuals of type-n students in teacher i's school-grade in year t, conditional on observables Z^s_{kt} and the characteristics x of all teachers in the same school-grade in t.

    (b) We form a weighted average of the residuals $\{\widehat{v}_{int}\}_t$ by weighting each $\widehat{v}_{int}$ by $\varpi_{int} = \frac{K_{SG^T_{it}n}}{\sum_t K_{SG^T_{it}n}}$, so that residuals corresponding to more observations receive more weight:

    $$\bar{v}_{in} = \sum_t \varpi_{int} \widehat{v}_{int} \quad (12)$$

Note that assumption A1 implies

$$E(\bar{v}_{in}) = v_{in} + \sum_{t} \varpi_{int} \sum_{i' \neq i : SG^T_{i't} = SG^T_{it}} v_{i'n}$$

Taking the limit of this expectation as t approaches infinity yields

$$\lim_{t \to \infty} E(\bar{v}_{in}) = v_{in} + \lim_{t \to \infty} \sum_{t} \varpi_{int} \sum_{i' \neq i : SG^T_{i't} = SG^T_{it}} v_{i'n}$$

It follows that a requirement for the estimator $\bar{v}'_{in}$ to be asymptotically unbiased is that $\lim_{t\to\infty} \sum_t \varpi_{int} \sum_{i' \neq i : SG^T_{i't} = SG^T_{it}} v_{i'n} = 0$. In words, the weighted sum of the effects of all teachers in i's school-grade over time has to approach 0 as the number of periods grows large. This requirement is met because 1) the teacher effect v_{in} is defined as a residual component of standardized test scores conditioning on grade-by-school fixed effects (which implies that, across time, the mean of v_{in} is zero within each school-grade) and 2) Assumption A2 guarantees that there is no sorting of teachers on unobservables across school-grades over time.

(c) Armed with $\bar{v}_{in}$, we can construct the empirical Bayes estimator of v_{in} by multiplying $\bar{v}_{in}$ by the shrinkage factor, a measure of the reliability of the estimator defined as the ratio between the estimated variance of the quantity to be estimated, $\hat{\sigma}_{vn} = Var(v_{in})$, and the variance of the estimator:

$$\hat{v}_{in} = \bar{v}_{in} \left(\frac{\widehat{\sigma}^2_{vn}}{Var(\bar{v}_{in})} \right),$$

where, given assumptions A1 and A2, we can estimate $\hat{\sigma}^2_{vn}$ as

$$\widehat{\sigma}^2_{vn} = \frac{Cov(\widehat{v}_{int}, \widehat{v}_{int-1})}{J_{SG^T_{it,t-1}}}$$

and $J_{SG^T_{it,t-1}} = \sum_{i'} I(SG^T_{i't} = SG^T_{it}) I(SG^T_{i't-1} = SG^T_{it-1})$ is the number of teachers who are in the same school-grade as i in both t and $t-1$.

Identification: Model 1 The identification of teacher effects v_{in} leverages teacher turnover across school-grades over time. Our identifying assumption is that turnover of teachers across school-grades, within a district, is unrelated to v_{in} (Assumption A2). Importantly, this assumption allows for the endogenous sorting of teachers across districts based on v_{i1} and v_{i2}, as is the case in our model. In the estimation of v_{in}, this type of sorting is accounted for by the school-grade fixed effects included in Z^s_{kt}.

Teacher turnover across school-grades allows us to identify v_{in} from $\bar{v}_{in}$ for all i and n. In particular, we can stack all the equations (12) for all I teachers and $n = 1, 2$, forming a system of $2I$ equations (where I is the total number of teachers) in $2I$ unknowns ($\{v_{in}\}_{i,n\in\{1,2\}}$). Identification is achieved if the rank condition of the system is satisfied, i.e., if the coefficient matrix of the system is full-rank.

In practice, this requires that the set $\{i' : SG^T_{i't} = SG^T_{it} \forall t\}$ is empty for all i, which means that there are no two teachers who teach the same school-grade in all t. When this is the case, the system (and the v_{in} for all i and n) is perfectly identified. In our data, $\{i' : SG^T_{i't} = SG^T_{it} \forall t\}$ is empty for 75% of teachers, for whom we can precisely estimate (v_{i1}, v_{i2}) . For the remaining 25% of teachers, $\{i' := SG^T_{i't} = SG^T_{it} \forall t\}$ is non-empty, and our estimated v_{in} is the average of $v_{i'n}$ for $i' : SG^T_{i't} = SG^T_{it} \forall t$.

B2.3.2 Achievement Model 2 (Alternative)

An alternative model would feature the assumption that each teacher contributes only to the achievement of the students in her classroom, while also assuming that teacher effectiveness is fixed over time. These assumptions have been used extensively in the value-added literature (e.g. Rockoff, 2004; Aaronson et al., 2007; Kane and Staiger, 2008).[3] The achievement model in this case would be:

$$A_{kt} = \gamma Z^s_{kt} + \sum_{n=1}^{2} I\left(\tau_k = n\right) v_{i(kt)n} + \varepsilon_{kt} \tag{13}$$

$$= \gamma Z^s_{kt} + \varphi_{kt} \tag{14}$$

where $i(kt)$ denotes student k's teacher in year t, i.e., k is in teacher i's classroom in year t. The contribution of teacher i to the achievement of a student of type $n \in \{1, 2\}$ is simply v_{in}. To estimate this quantity, we add the following assumption to A1 and A2:

A3. The variable $j_{int} = K_{int}/K_{SG^T_{it}n}$ is i.i.d. with mean $1/J_{SG^T_{it}n}$, where K_{int} is the number of students of type n in the classroom of teacher i in year t and $J_{SG^T_{it}}$ is the number of teachers in school-grade SG^T_{it} in t. Furthermore, $Cov(j_{int}, v_{i'n}) = 0 \ \forall i, i', t$. That is, class size is unrelated to teacher effectiveness within each school-grade.

Estimation: Model 2 With A1-A3, we can adapt the estimation procedure as follows.

1. We estimate γ via OLS on equation (13) to obtain $\widehat{\gamma}$.

[3] Besides assuming that teacher effectiveness is fixed over time, these studies assume that teacher effectiveness is one-dimensional, rather than student-type-specific.

2. We construct

$$\widehat{\varphi}'_{kt} = A_{kt} - \hat{\gamma} Z^s_{kt} \tag{15}$$

which is an estimate for $\sum_{n=1}^{2} v_{i(kt)n} I(\tau_k = n) + \varepsilon_{kt}$. For each teacher i, we define, for $n \in \{1, 2\}$

$$\widehat{v}'_{int} = \frac{1}{K_{SG^T_{it}n}} \sum_{k:SG_{kt}=SG^T_{it}} \widehat{\varphi}'_{kt} I\left(\tau_k = n\right) \tag{16}$$

$$\text{which is an estimate of} \quad \sum_{i':SG^T_{it}=SG^T_{i't}} j_{i'nt} v_{i'n} + \frac{1}{K_{SG^T_{it}n}} \sum_{k:SG_{kt}=SG^T_{it}} \varepsilon_{kt} \tag{17}$$

3. We form a weighted average of $\{\widehat{v}'_{int}\}_t$, with the same weights ϖ_{int} as before:

$$\bar{v}'_{in} = \sum_t \varpi_{int} \widehat{v}'_{int}$$

Assumption A1. implies

$$E(\bar{v}'_{in}) = v_{in} \sum_t \frac{\varpi_{int}}{J_{SG^T_{it}}} + \sum_t \frac{\varpi_{int}}{J_{SG^T_{it}}} \sum_{i':SG^T_{it}=SG^T_{i't}} v_{i'n}$$

Taking the limit of this expectation as t approaches infinity implies

$$\lim_{t\to\infty} E(\bar{v}'_{in}) = v_{in} \sum_t \frac{\varpi_{int}}{J_{SG^T_{it}}} + \lim_{t\to\infty} \sum_t \frac{\varpi_{int}}{J_{SG^T_{it}}} \sum_{i':SG^T_{it}=SG^T_{i't}} v_{i'n}$$

It follows that the estimator

$$\bar{\bar{v}}'_{in} = \frac{1}{\sum_t \frac{\varpi_{int}}{J_{SG^T_{it}}}} \bar{v}'_{in} \tag{18}$$

is asymptotically unbiased if $\lim_{t\to\infty} \sum_t \frac{\varpi_{int}}{J_{SG^T_{it}}} \sum_{i':SG^T_{it}=SG^T_{i't}} v_{i'n} = 0$. As before, this requirement implies that the weighted average of the effects of all teachers in i's school-grade over time has to approach 0 as the number of periods grows large. Assumption A2 and the fact that we are conditioning on school-grade fixed effects guarantees that this is the case asymptotically.

4. Finally, we construct the empirical Bayes estimator for v_{in} as

$$\hat{v}'_{in} = \bar{\bar{v}}'_{in} \left(\frac{\widehat{\sigma}^{2\prime}_{vn}}{Var(\bar{\bar{v}}'_{in})} \right)$$

Table B1: Correlation of Teacher Effectiveness between Model 1 and Model 2

	Estimation Sample (2014)		Validation Sample (2010)	
experience	$corr(c_{i1}, \hat{v}'_{i1})$	$corr(c_{i2}, \hat{v}'_{i2})$	$corr(c_{i1}, \hat{v}'_{i1})$	$corr(c_{i2}, \hat{v}'_{i2})$
$= 0$	0.91	0.98	0.86	0.90
$\in [1, 2]$	0.85	0.87	0.86	0.90
$\in [3, 4]$	0.88	0.93	0.88	0.91
$\in [5, 9]$	0.85	0.91	0.85	0.87
$\in [10, 14]$	0.85	0.86	0.86	0.88
≥ 15	0.86	0.87	0.84	0.86

and we can estimate the variance of v_{in}, $\widehat{\sigma}^{2\prime}_{vn}$, as

$$\widehat{\sigma}^{2\prime}_{vn} = J_{SG^T_{it}} J_{SG^T_{it-1}} \frac{Cov(\widehat{v}'_{int}, \widehat{v}'_{int-1})}{J_{SG^T_{it,t-1}}}$$

Identification: Model 2 The identification of this alternative model also relies on within-district school-grade turnover as in Model 1. Equation (18) represents a system of $2I$ equations (where I is the total number of teachers) in $2I$ unknowns, where the unknowns are $\{v_{in}\}_{i,n\in\{1,2\}}$. Teacher effectiveness v_{in} is perfectly identified for teachers for whom there are at least two periods t and t' with $SG^T_{it} \neq SG^T_{it'}$.

B2.3.3 Teacher Effectiveness: Model 1 vs Model 2

Correlation of Teacher Effectiveness Measures Table B1 displays the correlations between (c_{i1}, c_{i2}), the measures of teacher effectiveness we use in our preferred model (Model 1), and $(\hat{v}'_{i1}, \hat{v}'_{i2})$, estimates of teacher effectiveness obtained with the alternative model (Model 2). We report these for both the estimation sample (2014) and the validation sample (2010). Teacher effectiveness measures estimated from the two models are highly correlated.

Inferred Offer Sets As discussed in the identification section of the paper, an important step of our estimation is to infer subsets of the offers received by each teacher from the observed teacher-district matches (we denote these as O^s_i). To show that the model estimates are robust to using $(\hat{v}'_{i1}, \hat{v}'_{i2})$ in place of (c_{i1}, c_{i2}), we re-constructed the inferred offer (sub)sets using $(\hat{v}'_{i1}, \hat{v}'_{i2})$, denoted by $\widetilde{O}^s_i$. Comparing O^s_i with $\widetilde{O}^s_i$ for each of the 6,600 teachers in our estimation sample, we find that 1) $O^s_i = \widetilde{O}^s_i$ for 27% of teachers, 2) $O^s_i \supset \widetilde{O}^s_i$ for 23% of teachers, 3) $O^s_i \subset \widetilde{O}^s_i$ for 21% of teachers, and 4) for the rest 28% of teachers, there are some districts in O^s_i but not in $\widetilde{O}^s_i$ and some districts in $\widetilde{O}^s_i$ but not in O^s_i. For the robustness of

Table B2: OLS of Teacher-District Match

	Aux 1a		Aux 1b	
Achievement	Model 1	Model 2	Model 1	Model 2
wage	0.002 (0.0002)	0.001 (0.0001)	-0.000005 (0.000002)	-0.000005 (0.000002)
e^{λ_d}	-0.004 (0.009)	-0.003 (0.005)	-0.0001 (0.0001)	-0.0002 (0.0001)
$c_1 \times \lambda_d$	0.52 (0.29)	0.32 (0.19)	-0.02 (0.006)	-0.02 (0.015)
$I(d \neq d_0)$	-0.72 (0.02)	-0.73 (0.02)	-0.80 (0.01)	-0.80 (0.01)
$I(d \neq d_0) \times x_1$	-0.008 (0.001)	-0.008 (0.001)	-0.008 (0.0005)	-0.008 (0.0005)
$I(z_d \neq z_{d_0})$	-0.06 (0.006)	-0.06 (0.005)	-0.0006 (0.0001)	-0.0006 (0.0001)
q_d : urban	0.01 (0.002)	0.01 (0.002)	0.003 (0.0002)	0.003 (0.0002)
q_d : suburban	0.01 (0.002)	0.01 (0.002)	0.001 (0.0001)	0.001 (0.0001)
q_d : large metro	0.11 (0.03)	0.09 (0.02)	0.01 (0.002)	0.01 (0.002)
# Obs	57,068	60630	2,712,600	

Standard errors are in parentheses.

teacher preferences under O_i^s in place of $\widetilde{O}_i^s$, case 1) is ideal, and cases 2) and 3) are not concerning, because we only need subsets of offers to infer teacher preferences Fox (2007). These three cases account for 72% of teachers.

Auxiliary Models A key source of identification comes from our auxiliary models Aux 1a and Aux 1b that characterize teacher-district matches via regressions,

$$y_{id} = \beta_1^m w(x_i, c_i|\omega_d) + I(d_{0i} > 0) \begin{bmatrix} I(d \neq d_{0i}) x_i \beta_2^m \\ +\beta_3^m I(z_d \neq z_{d_{0i}}) \end{bmatrix} + q_d \beta_4^m + \beta_5^m e^{\lambda_d} + \beta_6^m c_{1i} \lambda_d + \psi_i + \varepsilon_{id}^m.$$

In Aux 1a, i's are all the teachers whose inferred subsets of offers O_i^s contain more than one district, and an observation (i, d) is a teacher-district pair in these inferred subsets. In Aux 1b, an observation is any teacher-district pair, with $I \times D$ total observations.

In Table B2, we compare Aux 1a and Aux 1b when a teacher is characterized by (x, c) (Model 1) against their counterparts when a teacher is characterized by $(x, \hat{v}')$ (Model 2). Between the two cases, regression coefficients in Aux 1a are very similar, and those in Aux 1b are almost identical.

B2.3.4 Teacher Effectiveness: Two-Dimensional vs One-Dimensional

To check whether allowing teacher effectiveness to vary by student type provides gains in terms of explaining the overall variation in test scores, we estimate a counterpart of Model

Table B3: Sum of squared test score residuals under (c_1, c_2) and under c

Effectiveness measure	c	(c_1, c_2)	% difference
Student type			
all students	0.1680	0.1370	22.61%
$\tau_k = 1$	0.1922	0.1552	23.87%
$\tau_k = 2$	0.1438	0.1189	20.97%

(8) with one-dimensional rather than two-dimensional teacher effectiveness and compare it with Model (8). Table B3 compares the average sum of squared test score residuals $\hat{\varphi}_{kt}$, by student type, obtained from each model. Our two-dimensional teacher effectiveness model explains approximately 20% more variation in test scores compared to its one-dimensional effectiveness counterpart.

B2.3.5 Teacher Effectiveness: Race

Previous studies suggest that the match between the teacher's race and the student's race can matter for achievement. In comparison, we focus on teachers' comparative advantages in teaching students with different prior achievement types. We make this choice for two reasons. First, as shown in Table B4, if we add teacher race and the interaction of teacher and student race to our achievement model (student race is already included in our achievement model), almost none of the added terms are significant. Second, if we add a teacher's race and gender and their interactions with the district's racial and gender composition of students to our Aux 1a (Column 1 of Table 2 in the main text), the R^2 is barely improved (from 0.68 to 0.681).

B2.4 Wage Schedules

B2.4.1 Pre-Reform Wage Schedules

We obtain $W_d^0(x_i)$ as the predicted values from the following regression, estimated using data from 2007 to 2011:

$$w_{it}^0 = \delta_d^0 + Exp_{it}\delta_{g(i)}^e + MA_{it}\delta_{g(i)}^m + \varepsilon_{it}, \tag{19}$$

where i and t refer to teacher and year, respectively; w_{it}^0 is the wage of teacher i in year t; Exp_{it} is a vector of indicators for six classes of years of experience: 0, $[1, 2]$, $[3, 4]$, $[5, 9]$, $[10, 14]$, and $[15, +\infty)$; and MA_{it} is an indicator for having a Master's degree (MA) or a

Table B4: Estimates of achievement model in equation (13), obtained controlling for teachers' (T) and students' (S) race/ethnicity indicators and their interactions

	$\tau = 1$	$\tau = 2$
	(1)	(2)
Black S	-0.056***	-0.067***
	(0.003)	(0.003)
Hisp S	-0.007**	-0.022***
	(0.003)	(0.003)
Asian S	0.053***	0.081***
	(0.004)	(0.004)
Black T	-0.001	0.0001
	(0.005)	(0.005)
Black T * Black S	-0.008	-0.019*
	(0.006)	(0.010)
Hisp T	-0.010*	-0.006
	(0.005)	(0.005)
Hisp T * Hisp S	0.007	0.008
	(0.007)	(0.009)
Asian T	0.003	0.004
	(0.007)	(0.008)
Asian T * Asian S	0.015	0.022
	(0.017)	(0.016)
Observations	3,360,517	3,635,942

higher degree. The parameter δ_0 can be interpreted as the average wages for teachers with zero experience and without a MA; with $\delta^e_{g(i)}$ normalized to 0 for those with zero experience, $\delta^e_{g(i)}$ is the average wage premium for teachers in each of the higher experience category, relative to those with zero experience with the same education; and δ^m is the wage premium for teachers who have a MA.

We estimate the intercept δ^0_d separately for each district. Trading off the accuracy of our wage schedules with power, we estimate the coefficients δ^e and δ^m by groups of districts, defined as follows:

1. For the 35 large districts (i.e., those with at least 10 teachers in each experience and education category), each group corresponds to a district.

2. For the remaining 356 districts, we construct groups based on the similarity in their salary schedules. To do so, we proceed as follows.

   (a) For each district, we calculate the following summary statistics for their salary schedules: (i) wages for teachers with 0 years of experience and $MA_{it} = 0$ (i.e., the lowest possible wage category); (ii) wages for teachers with over 15 years of experience and $MA_{it} = 0$ (i.e., the highest possible wage category for those without MA); (iii) average salary difference between a teacher with more than 15 years of experience and a MA, and one with the same experience and no MA.

   (b) We check whether each district is above or below the median of the cross-districts distribution for each of the three statistics.

   (c) We form eight groups based on the statistics (i), and (ii), and (iii), and assign each district to a group as follows:

Group	(i)	(ii)	(iii)
1	≥median	≥median	≥median
2	≥median	≥median	<median
3	≥median	<median	≥median
4	<median	≥median	≥median
5	<median	<median	≥median
6	<median	≥median	<median
7	≥median	<median	<median
8	<median	<median	<median

Table B5 summaries the point estimates from Equation (19). In particular, it reports the cross-district means and standard deviations of the estimated vectors δ. Figure B1 shows a binned scatter plot of $W_d^0(x_i)$ and data wage w_{it}^0 in 2010. The former predicts the latter remarkably well, with a correlation coefficient of 0.93 (significant at 1 percent).

B2.4.2 Districts' Choice Set of Wage Schedules

A district's wage rule is given by

$$w_d(x, c|\omega) = \max\left\{\min\left\{\omega_1 W_d^0(x) + \omega_2(\lambda_d c_1 + (1-\lambda_d)c_2), \overline{w}\right\}, \underline{w}\right\}. \tag{20}$$

A district chooses (ω_1, ω_2) from a discrete set Ω, the grid points of which are chosen as follows.

Table B5: Cross-district Summary of Pre-Reform Wage Schedules

		Cross-district Mean	Cross-district Std Dev.
δ^0		34,686.8	3,286.1
δ^e:	$[1, 2]$	1,719.2	598.3
	$[3, 4]$	3,939.1	1,103.3
	$[5, 9]$	8,227.8	1,536.6
	$[10, 14]$	14,644.0	2,348.5
	≥ 15	21,235.4	3,063.4
δ^m(MA)		7,008.5	2,456.6

Figure B1: Relationship between $W_d^0\,(x_i)$ and w_{it}^0

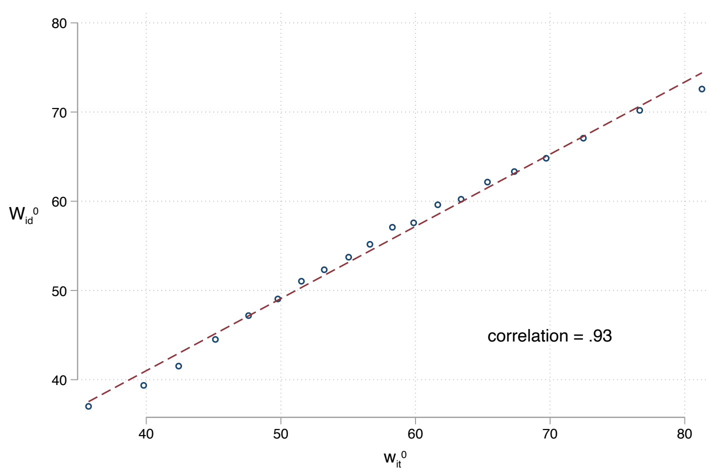

Note: Binned scatterplot of $W_d^0\,(x_i)$ and w_{it}^0 using wage data from 2010.

1. We start by estimating the parameters $(\widetilde{\omega}_{d1}, \widetilde{\omega}_{d2}) \geq 0$ separately for each district from

$$w_i = \widetilde{\omega}_{d1} W_d^0(x_i) + \widetilde{\omega}_{d2} TC(c_i, \lambda_d) + \varepsilon_i^w, \text{ for } i : d(i) = d$$

where w_i is the observed 2014 wage for teacher i working in district d $(i : d(i) = d)$, $W_d^0(x_i)$ is defined as in Section B2.4.1, and teacher contribution $TC(c_i, \lambda_d)$ is given by

$$TC(c_i, \lambda_d) = \lambda_d c_{i1} + (1 - \lambda_d) c_{i2}.$$

2. Based on the estimated $\{(\widetilde{\omega}_{d1}, \widetilde{\omega}_{d2})\}_d$, we choose a set of equally spaced grid points that provides a good coverage of the empirical distribution in the data:

$$\Omega^o = \{0.9, 0.95, 1, 1.05, 1.1\} \times \{0, 10, 30, 50, 75, 100, 200\}.$$

3. We assign each district the wage schedule $(\omega_{d1}^o, \omega_{d2}^o) \in \Omega^o$ that best summarizes the distribution of teacher wages in that district $\{i : d(i) = d\}$, i.e.,

$$(\omega_{d1}^o, \omega_{d2}^o) = \arg \max_{(\omega_1, \omega_2) \in \Omega^o} \sum_{i:d(i)=d} (w_i - w_d(x_i, c_i; \omega))^2,$$

$$s.t.\ w_d(x_i, c_i; \omega) = \begin{cases} \underline{w} \text{ if } \omega_1 W_d^0(x_i) + \omega_2 TC(c_i, \lambda_d) < \underline{w} \\ \overline{w} \text{ if } \omega_1 W_d^0(x_i) + \omega_2 TC(c_i, \lambda_d) > \overline{w} \\ \omega_1 W_d^0(x_i) + \omega_2 TC(c_i, \lambda_d) \text{ otherwise} \end{cases},$$

where $\underline{w}$ $(\overline{w})$ is 0.3 standard deviations below (0.2 standard deviations above) the observed 1st (99th) wage percentile in the sample.

   - The $(\omega_{d1}^o, \omega_{d2}^o)$ selected with this procedure predicts teachers' actual salaries quite well: 1) the absolute percentage deviation of predicted wages from actual wages in 2014, i.e., $\left|1 - \frac{w_d(x_i, c_i; \omega)}{w_i}\right|$, is less than 10% for 95% of teachers in our sample; and 2) regressing w_i on $w_d(x_i, c_i; \omega)$ yields a slope coefficient of 0.98 (with a standard error of 0.001) and an R^2 of 0.99.

4. Finally, we expand the grid range to allow for the possibility that district choices may go out of the empirical range in counterfactual scenarios. The choice set in the model is given by

$$\Omega = \{0.9, 0.95, 1, 1.05, 1.1, 1.15\} \times \{0, 10, 30, 50, 75, 100, 200, 225\}.$$

Figure B2: Relationship between deviations of true wages from $w_d(x, c|\omega)$, obtained using rules (20) and (21)

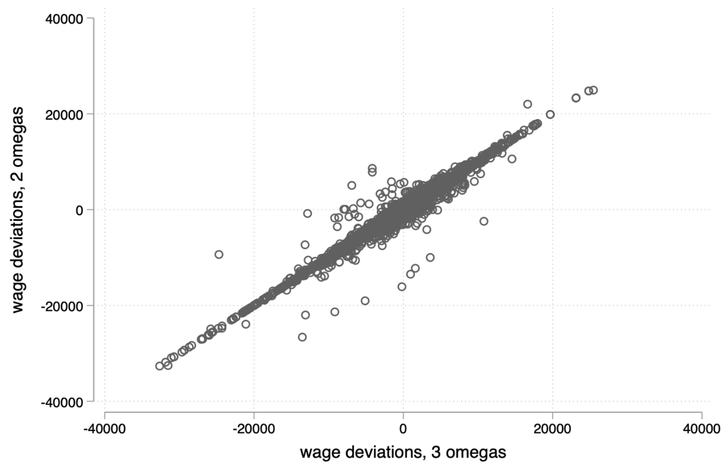

Note: Binned scatterplot of the difference between true 2014 teacher wages and $w_d(x, c|\omega)$, calculated using (20) (vertical axis) and (21) (horizontal axis).

where both $\omega_1 = 1.15$ and $\omega_2 = 225$ are outside of Ω^o.

B2.4.3 An Alternative Wage Rule with Three ω's

We have also tried to allow for a more flexible alternative wage schedule as follows

$$w_d(x, c|\omega) = \max\left\{\min\left\{\omega_1 W_d^0(x) + \omega_2 \lambda_d c_1 + \omega_3 (1 - \lambda_d) c_2, \overline{w}\right\}, \underline{w}\right\}. \tag{21}$$

Wage rule (20) we use in the paper is a special case of (21) with $\omega_2 = \omega_3$. We repeat the exercise as in Section B2.4.2, but under the three-ω specification (21). This procedure yields the triplet $(\omega'_{d1}, \omega'_{d2}, \omega'_{d3})$ that best summarizes the observed distribution of teacher wages in each district d. Figure B2 compares the predicted wage under rule (20) and that under rule (21). The two predicted wages are nearly indistinguishable from each other, indicating the absence of large predictive gains associated with the use of (21) instead of (20).

B3 Across-District vs Within-District Variation

In our model we abstract from within-district competition for teachers, focusing on competition across districts. Here we show that cross-district variation clearly dominates within-district, cross-school variation in terms of both teacher wages and the share of low-achieving students.

Table B6: Variation in salaries across and within districts, 2013-2016

Specification	sqrt(MSE)	R^2	Δsqrt(MSE) from Baseline
Baseline: Experience, Education, c_1,c_2	6,856	0.69	–
+ District FE	4,711	0.86	31.3%
+ School FE	4,523	0.87	34.0%

B3.1 Wages

Table B6 shows the adjusted R^2 and the root mean-squared error (MSE) of a regression of post-Act 10 salaries on c_1, c_2, experience and education (first row). It then shows how the R^2 and MSE change as we sequentially add district fixed effects (second row) and school fixed effects (third row). Adding district fixed effects reduces the root MSE by 31.3%; this implies that differences across districts explain 31.3% of the residual variation in salaries, conditional on teacher characteristics. Adding school fixed effects instead only explains an additional 2.7% of the root MSE. We can conclude that the main source of variation in wages is across districts, not across schools within districts.

B3.2 Student Composition

The cross-district variation in the share of low-achieving students (λ in our model) largely dominates the within-district, cross-school variation. We provide evidence of this in three different ways.

1. Estimates from an OLS student-level regression of an indicator for being low-achieving, to which we progressively add district and school fixed effects, indicates that districts explain 8.7% of the variation in this probability whereas schools only explain an additional 2.7%.
2. The estimated R^2 of an OLS regression of the school-level share of low-achieving students on district fixed effects, weighted by enrollment, indicates that 74% of the variation in the school-level share is explained by the district.
3. For each school, we calculate the absolute difference between the school-level and the district-level shares of low-achieving students. This absolute difference has a mean of 0.05 and a standard deviation of 0.06. The 25th, 50th and 75th percentile of this absolute difference are 0.01, 0.03 and 0.07 respectively.

B3.3 Teacher Assignment Across School-Grades Within a District

The identification of c_1 and c_2 in our achievement model relies on the assumption of random sorting of teachers across school-grades within each district, conditional on all the covariates described in Appendix B2.3. To test for the presence of non-random sorting, in Table B7 we combine the approaches of Chetty et al. (2014) and Rothstein (2010). In columns 1 and 2 we follow Chetty et al. (2014) and estimate the slope of the relationship between changes in students' test score residuals (obtained from a regression of test scores on all the covariates in equation (13)) and changes in c_1 and c_2. As in Chetty et al. (2014), we control for school-by-grade and school-by-year fixed effects. These tests, shown in columns 1 and 2, reveal a slope coefficient that is statistically indistinguishable from one, indicating that our estimates of (c_1, c_2) are forecast unbiased for (c_1, c_2).

In columns 3 and 4 of Table B7 we combine the above empirical design with the test proposed by Rothstein (2010) and estimate the relationship between changes in (c_1, c_2) and changes in *lagged* test score residuals. If the estimates in this specification were significant, they would indicate non-random sorting of teachers across grade-schools. Reassuringly, the slope coefficients in columns 3 and 4 are smaller than those in columns 1 and 2 and statistically indistinguishable from zero.

Table B7: Test for Non-Random Teacher Sorting Across Grade-Schools (Rothstein 2010)

	Residuals		Lagged residuals	
	(1)	(2)	(3)	(4)
Δc_0	1.204***		0.365	
	(0.072)		(0.250)	
Δc_1		0.905***		0.394
		(0.164)		(0.286)
School-by-year FE	Yes	Yes	Yes	Yes
N	6448	1269	1518	298
# school-grades	1950	694	582	174

B4 Robustness Checks

In this section, we conduct two sets of robustness checks with respect to the two maintained assumptions underlying our identification strategy:

A1: (x, c) are observable to all districts. With our data, it is difficult to separate preferences from information friction; we abstract away from the latter.

A2: Districts cannot discriminate among teachers by factors other than (x, c).

As a partial test for the robustness of our results with respect to A1, we conduct the following exercise: Instead of (c_1, c_2), districts observe $(c_1 + err_1, c_2 + err_2)$ and make wage and job offer decisions based on these noisy measures. Assuming that $err_k \sim N\left(0, \sigma^2_{err_k}\right)$ are i.i.d. random noises and considering values of σ_{err_k} equal to one, two, or four times the standard deviation of c_k, for $k = 1, 2$, we repeat the procedure described in Section 4.1.2 of the main text to construct sub-offer sets using the observed matches. Column 1 of Table B11 reports the baseline estimates of Aux 1a, which are key for the identification of teachers' preferences. Columns 2-4 show estimates obtained assuming that both teachers' and districts' decisions are based on $(c_1 + err_1, c_2 + err_2)$, while the researcher observes (c_1, c_2). Columns 5-7 show the corresponding estimates assuming that districts' decisions are based on $(c_1 + err_1, c_2 + err_2)$, while teachers' decisions are based on (c_1, c_2). Throughout these exercises, the estimates of Aux 1a are robust.

To investigate robustness to a violation of A2, we consider the possibility that some ineffective teachers may have been hired for reasons other than (x, c). Table B12 compares our auxiliary model Aux 1a with its counterpart that does not use observed teacher-district (i, d) matches to infer offers for other teachers if i's effectiveness with either low- or high-achieving students is below the 10th percentile among all teachers. Doing so has a significant impact on the number of inferred offers for other teachers; yet Aux 1a remains robust.

It should be noted that although our robustness checks give some comfort that simple violations of A1 and A2 may not seriously affect our inference, they are no proof that these assumptions (maintained throughout) are innocuous.

B5 The Impact of Changes in Parameter Values on Auxiliary Models

Following Einav et al. (2018), we provide more evidence on the mapping between data and parameters via a perturbation exercise. We adjust each parameter one at a time and measure responses of the predicted auxiliary models we use for estimation.

To be specific, letting $\left\{\widehat{\theta}_n\right\}_{n=1}^{20}$ be the vector of estimated structural parameters and $\{\widehat{\sigma}_{\theta_n}\}_{n=1}^{20}$ be the vector of their standard errors, we re-simulate our model 20 times. In the n^{th} simulation, we use the parameter vector $\left\{\widehat{\theta}_1, \widehat{\theta}_2, ..., \widehat{\theta}_{n-1}, \widehat{\theta}_n + \widehat{\sigma}_{\theta_n}, \widehat{\theta}_{n+1}, ..., \widehat{\theta}_{20}\right\}$, where the n^{th} parameter is perturbed by its standard error, and obtain new estimates of the auxiliary models. We then compute the percent change in absolute terms for each auxiliary model (regression coefficient or moment). This exercise produces a matrix of dimension 115×20 (number of auxiliary models $\times$ number of parameters). To ease exhibition, we take simple averages within sub-blocks of this matrix. Specifically, we split the auxiliary models into five groups as specified in the paper (Aux 1a, Aux 1b, Aux 2, Aux 3, and Aux 4) and split parameters into three groups (teacher preference parameters, district preference parameters, and wage-setting resistance cost parameters). This results in the 5 x 3 summary matrix shown in Table B13. Each cell in Table B13 shows the average percent change across auxiliary models and parameter permutations within a given sub-block.

Column 1 of Table B13 shows that teacher preference parameters primarily affect the sub-offer and all-offer regression models (Aux 1a and Aux1b), as well as the regression coefficients that link districts' wage choices to their pre-determined conditions (Aux 3). It is unsurprising that Aux 1a and Aux 1b are closely related to teachers' preferences, as these regressions are designed to mimic a conditional logit model of teachers' choices. Additionally, as teachers' preferences change, districts change their wage schedules in order to attract their preferred teachers; such responses are captured by changes in Aux 3.

Column 2 shows that district preference parameters mostly affect the regression coefficients that link wages to districts' pre-determined conditions (Aux 3) and the all-offer regression (Aux 1b), but they also affect the sub-offer regression (Aux 1a). As we argued in our identification section, Aux 3 should be informative of districts' preferences as districts can use wage choice to push or pull teachers; moreover, the *difference* between Aux 1a and 1b are also informative of districts' preferences.

Finally, Column 3 shows that the wage-setting resistance cost parameters affect the wage regressions and cross-district wage moments (Aux 3 and Aux 4). This is unsurprising as these two auxiliary models directly summarize wage choices. Notice that, by design, resistance cost parameters should have zero impact on Aux 1a, Aux 1b, and Aux 2, because these auxiliary models are obtained while holding wage schedules at the observed equilibrium levels.

B6 Additional Tables and Figures

Table B8: Estimated parameters of teacher effectiveness

	$\hat{\rho}_1$	$\hat{\rho}_2$
exp $= 0$	0	0
exp $\in [1, 2]$	0.0068	0.0009
exp $\in [3, 4]$	0.0154	0.0057
exp $\in [5, 9]$	0.0117	0.0028
exp $\in [10, 14]$	0.0117	0.0049
exp $\in [15, +\infty)$	0.0112	0.0038
R^2	0.677	0.625

Table B9: Distribution of teacher effectiveness

	c_1	c_2
min	-0.1398	-0.1988
p1	-0.0630	-0.0779
p5	-0.0345	-0.0417
p10	-0.0225	-0.0278
p25	-0.0049	-0.0075
median	0.0115	-0.0108
mean	0.0116	0.0110
p75	0.0282	0.0300
p90	0.0454	0.0503
p95	0.0582	0.0664
p99	0.0894	0.0978
max	0.1532	0..2362

Table B10: Teacher and District Characteristics (2010)

A. Teacher Characteristics	All	$x_1 < 3$	$x_1 \geq 10$
x_1: Experience	15.6 (9.6)	1.6 (0.5)	20.2 (7.7)
x_2: MA or above	0.55 (0.50)	0.05 (0.22)	0.66 (0.48)
$10c_1$	0.11 (0.25)	0.07 (0.27)	0.11 (0.25)
$10c_2$	0.12 (0.30)	0.06 (0.32)	0.12 (0.29)
Corr (c_1, c_2)	0.65	-	-
# Teachers	6,741	391	4,675
B. District Characteristics	**All**	**λ_d 1st Quartile**	**λ_d 4th Quartile**
Urban	0.04	0.02	0.03
Suburban	0.15	0.34	0.09
λ_d	0.50 (0.12)	0.34 (0.07)	0.64 (0.06)
Capacity	16.4 (30.7)	18.4 (16.2)	15.1 (46.2)
Budget/Capacity ($1,000)	52.4 (6.1)	54.3 (6.7)	51.2 (5.7)
Characteristics of District Incumbent Teachers ($d_0 = d$)			
Average experience	17.5 (5.1)	16.6 (4.6)	18.0 (5.6)
Share w/MA or above	0.52 (0.26)	0.57 (0.26)	0.48 (0.28)
Average $10c_1$	0.10 (0.10)	0.10 (0.09)	0.09 (0.13)
Average $10c_2$	0.11 (0.13)	0.11 (0.11)	0.09 (0.15)
# Districts	411	103	103

Means and std. deviations (in parentheses) of teacher (Panel A) and district (Panel B) characteristics.

Table B11: Estimates of Aux 1a Assuming Noisy Measures of (c_1, c_2)

	Baseline[a]	For teachers and districts			For districts only		
		σ_{err_k}	2*σ_{err_k}	4*σ_{err_k}	σ_{err_k}	2*σ_{err_k}	4*σ_{err_k}
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
wage	0.0015	0.0017	0.0018	0.0027	0.0017	0.0018	0.0027
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
e_d^λ	-0.0038	-0.0177	-0.0170	-0.0121	-0.0261	-0.0271	-0.0448
	(0.0088)	(0.0093)	(0.0096)	(0.0077)	(0.0115)	(0.0149)	(0.0156)
$c_1 \times \lambda_d$	0.5295	0.9761	0.8342	0.8377	0.9704	0.6981	0.8483
	(0.2923)	(0.3027)	(0.3137)	(0.2835)	(0.2963)	(0.3095)	(0.2389)
$d \neq d_0$	-0.7138	-0.7162	-0.7073	-0.7020	-0.7163	-0.7074	-0.7017
	(0.0166)	(0.0168)	(0.0170)	(0.0175)	(0.0168)	(0.0170)	(0.0175)
$d \neq d_0 \times$ exp	-0.0079	-0.0076	-0.0079	-0.0078	-0.0076	-0.0079	-0.0078
	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0006)
$z_d \neq z_{d_0}$	-0.0628	-0.0641	-0.0671	-0.0709	-0.0641	-0.0671	-0.0711
	(0.0060)	(0.0060)	(0.0061)	(0.0069)	(0.0060)	(0.0061)	(0.0069)
urban	0.0112	0.0242	0.0227	0.0214	0.0242	0.0226	0.0210
	(0.0022)	(0.0027)	(0.0026)	(0.0031)	(0.0027)	(0.0026)	(0.0031)
suburban	0.0126	0.0110	0.0110	0.0041	0.0111	0.0111	0.0041
	(0.0023)	(0.0022)	(0.0022)	(0.0025)	(0.0022)	(0.0022)	(0.0025)
large metro	0.1111	0.1053	0.1070	0.0962	0.1061	0.1080	0.1006
	(0.0270)	(0.0271)	(0.0263)	(0.0297)	(0.0271)	(0.0263)	(0.0298)
N	57068	52439	53310	46906	52439	53310	46906

[a]: Estimates of Aux 1 used in the main text (Column 1 of Table 2). Robust standard errors in parentheses.

Table B12: OLS of Teacher-District Matches (Aux 1a): Baseline and Robustness

Teacher's Choice Set	Baseline Inferred Offer Set[a]	Robustness Inferred Offer Set[b]
wage	0.002 (0.0002)	0.002 (0.0003)
e^{λ_d}	-0.004 (0.009)	-0.024 (0.014)
$c_1 \times \lambda_d$	0.53 (0.29)	1.12 (0.42)
$I(d \neq d_0)$	-0.72 (0.02)	-0.67 (0.02)
$I(d \neq d_0) \times$ experience	-0.008 (0.001)	-0.009 (0.001)
$I(z_d \neq z_{d_0})$	-0.06 (0.006)	-0.08 (0.008)
q_d : urban	0.01 (0.002)	0.004 (0.003)
q_d : suburban	0.01 (0.002)	0.01 (0.003)
q_d : large metro	0.11 (0.03)	0.09 (0.03)
# Obs	57,068	33,053

a and b: OLS specified in Aux 1a, teacher fixed effects included; 2014 data.

a: The auxiliary model used in the main text (Column 1 of Table 2).

b: Estimates obtained ignoring teacher-district matches (i, d) for teachers with c_{1i} or c_{2i} below the 10th percentile of their respective distribution when inferring matches.

Robust standard errors are in parentheses.

Table B13: Parameter Permutation Exercise: Change in Estimates of Auxiliary Models from Parameter Perturbation

	Parameter Group		
Auxiliary Model	Teacher Preferences	District Preferences	Wage-Setting Resistance Costs
Aux 1a	27.30%	0.45%	0.00%
Aux 1b	18.63%	2.41%	0.00%
Aux 2	0.40%	0.03%	0.00%
Aux 3	56.53%	11.17%	244.53%
Aux 4	1.33%	0.38%	29.58%

Figure B3: Distribution of teacher effectiveness

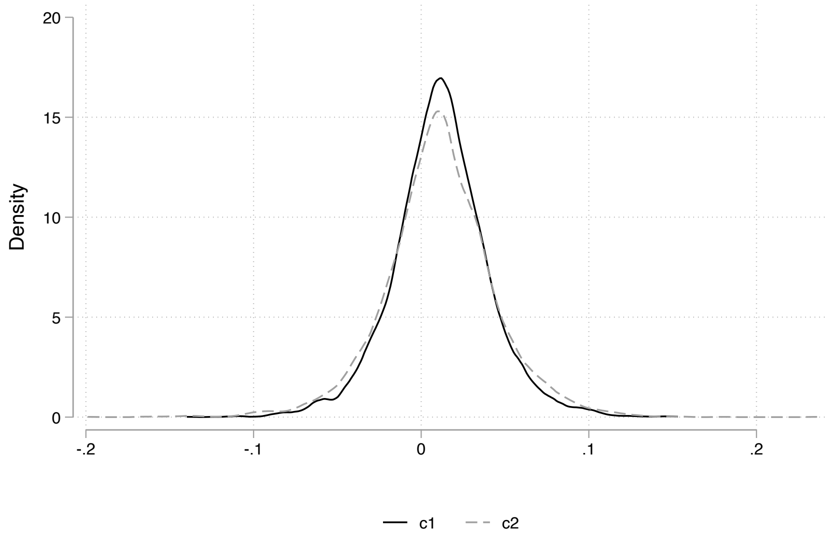

References

Aaronson, D., L. Barrow, and W. Sander (2007). Teachers and student achievement in the chicago public high schools. *Journal of labor Economics 25*(1), 95–135.

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.

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.

Kane, T. J. and D. O. Staiger (2008). Estimating teacher impacts on student achievement: An experimental evaluation. *NBER Working Paper*.

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

Rothstein, J. (2010). Teacher quality in educational production: Tracking, decay, and student achievement. *The Quarterly Journal of Economics 125*(1), 175–214.