RESEARCH ARTICLE

Do Early-Offers Equal Better Teachers?

Caitlyn Keo and Kristine West

St. Catherine University

Lesley Lavery

Macalester College

Napat Jatusripitak and Elton Mykerezi

University of Minnesota

Christopher Moore

Minneapolis Public Schools

We examine the impact of human resource changes in a large urban district that accelerated the hiring timeline for promising teacher candidates. After an initial screening, outstanding candidates were granted early offers and placed into the district's internal labor market to interview for open positions alongside incumbent teachers. We find that early-offer candidates were more likely to apply, interview and be hired into positions that were less desirable to incumbents. After their first year or two in the classroom, early-offer teachers had higher average scores than their peers who accepted standard offers on student surveys. Value-added and formal observation metrics differences were not statistically significant.

Keywords: selection, human resources, hiring, teacher quality

Over a decade ago, reports by the New Teacher Project (TNTP), a non-profit organization, challenged longstanding school district human resource practices by suggesting that hiring practices, particularly those of large urban districts serving primarily low-income and minority student populations, impede school systems' ability to hire high-quality teacher applicants. Between 2003 and 2013, large urban districts from Memphis to Houston, New Orleans and Washington, D.C. developed programs to accelerate hiring timelines (Van Cleef, 2013). Starting in 2014, Minneapolis Public Schools (MPS) joined this group and began to screen and make offers to new teachers in two well-defined stages. In the first stage, the district extended early offers to outstanding teacher candidates in April and May. These early-offer candidates were placed in the

district's internal labor market to interview for open positions alongside incumbent candidates. In the second stage, the district extended typical new hires offers attached to specific, remaining positions.

Though Minneapolis and others have responded to critiques of delayed hiring timelines, scholars know little about whether and how changes in personnel policy influence candidate pools, teacher quality, and student outcomes. Does the accelerated timeline help districts fill open positions with high-quality candidates? Are early-offer teachers better than those hired later in the cycle? To help address these questions, we combine two detailed administrative datasets from MPS to compare early-offer candidates to new teachers hired on traditional hiring timelines and to incumbents.

We test two hypotheses: (1) that the accelerated timeline improved the candidate pool for hard-to-staff schools and (2) that early-offer teachers outperform typical new hires. For the first hypothesis, we find that the program did provide a better pool of candidates available to principals at hard-to-staff sites. To measure "hard-to-staff" we characterized positions by the number of incumbent applicants -- a direct measure of how desirable a position is. Indirect measures of hard-to-staff included school-level demographics (share of students non-white, share of students proficient in reading and math, share of non-white faculty, average years of experience and the pupil-teacher ratio). Early-offer candidates were more likely to apply, interview, and be hired into positions that were less desirable to incumbents as well as to positions with school-level demographics that indicate less advantaged settings. For the second hypothesis, we see some evidence that after their first year or two in the classroom, early-offer teachers outperformed their peers who accepted traditional cycle offers on all three of the district's teacher evaluation metrics: student surveys, formal observation metrics and value-added. However, only the student survey results reach statistical significance so further research is needed before strong claims can be substantiated.

Below, we review the relevant literature on teacher hiring and discuss the Minneapolis context, then present methods and results of a three-pronged - interview/placement/performance - analysis. We conclude with a discussion of the results and comment on the policy implications.

LITERATURE REVIEW

Existing research highlights the importance of intentional recruitment and selection to secure high-caliber applicants (Bruno & Strunk, 2018; Dee & Goldhaber, 2017; Goldhaber, Grout, & Huntington-Klein, 2017). This work demonstrates that hiring practices can be designed to make a district more competitive in the local labor market and improve overall market efficiency. This work also suggests that the current allocation of teachers is suboptimal. The status quo disadvantages large urban districts that serve larger shares of at-risk students who would have a great marginal benefit from highly effective teaching.

Researchers at The New Teacher Project (TNTP) drew policymakers' attention to pervasive late-hiring, documenting both the scope of rushed practices and asserting a link between the timing of a hire and pre-hire credentials. Specifically, TNTP's Levin and Quinn (2003) found that when districts failed to make job offers until mid-to-late summer, 31-61% of candidates withdrew their applications. Withdrawn candidates had higher GPAs and were 40% more likely to have a degree in their teaching field than the remaining pool of applicants. To the extent that GPA and field of degree are good predictors of teacher quality, this suggests the timing of hiring

has a negative impact on a district's ability to hire excellent teachers. In a 2007 follow-up study, using data from Chicago Public Schools, TNTP found that more-credentialed teacher candidates had more options available to them and were more likely to withdraw their applications due to drawn out hiring timelines than those with fewer credentials, leaving a weaker pool of candidates to fill the district's classrooms.

Others have presented more nuanced results. In 2009 and 2012, Engel further explored the link between hiring delays and credentials, armed with data from the nationally representative 1999-2000 Schools and Staffing Survey (Engel, 2009; Engel, 2012). Her work confirmed that urban and socioeconomically disadvantaged districts hire the majority of their teachers late. In urban districts, one fifth of hires are made after the first day of school! Contrary to the TNTP reports, however, Engel (2012) found no difference in early and late hires' credentials including certification, Master's degree, and the selectivity of undergraduate institution.

Bypassing the need to infer quality from credentials, Papay & Kraft (2016) directly explored the relationship between late hiring and teacher effectiveness, measured via individual teachers' student test scores and found that hiring after the school year has started leads to lower student achievement. They also found that in math, but not reading, gaps in teacher effectiveness persist beyond the first year. This suggests that, in particular, hiring delays may cause a district to miss out on promising math teachers.

Most of the existing literature focuses on extremely late hiring. The TNTP studies and Engel (2009 & 2012) consider job offers made late in the summer and Papay & Kraft (2016) and Engel (2012) look at hires made after the school year has already begun. Hiring when classes are about to begin or have already begun is clearly far from ideal -- either teachers have to scramble to prepare or students have to start the year with long-term substitutes. In the current study, we examine a program that extended offers to outstanding candidates in April and May. Thus, rather than comparing a standard hiring timeline to an emergency hire, we study a more policy relevant comparison between an accelerated hiring timeline and a standard hiring timeline.

Context

We study Minneapolis Public Schools (MPS), an urban district in Minnesota that generally hires in excess of 300 new teachers per year. In 2014, when the MPS early offers program began, MPS included nearly 80 schools, employed approximately 3,750 teachers, and served more than 35,000 students. Of these students, 67% were students of color, 65% were eligible for free/reduced price meals, and 10% were homeless or highly mobile. The demographics in the district changed very little over the course of the time period studied.

In 2014, the district hired a new senior director of talent acquisition who was convinced that delays in the hiring timeline negatively impacted district recruiting efforts and had detrimental consequences for student achievement, especially in hard to staff schools and subject areas like special education and ESL/bilingual education. To combat this, the director implemented two important district-wide changes to facilitate early hiring. First, the district began conducting centralized and structured interviews. In addition, the district began offering early employment contracts to the most-promising candidates who were then allowed to participate in the internal labor market. The changes allow Human Resources staff to screen, evaluate and offer employment (but not a specific position) to external candidates, which contrasts with decentralized hiring in which a principal evaluates and offers a specific position.

The district's internal labor market is referred to as "interview and select." Interview and select was first laid out in a Memorandum of Understanding (MOU) attached to the 2007-2009 teachers collective bargaining agreement. This MOU replaced a system that was entirely dictated by seniority with a system that let principals and site-based hiring teams interview up to ten candidates -- the five most senior candidates in the applicant pool and any five other candidates of their choosing.

Prior to 2014, internal candidates interviewed for open positions before any external candidates were considered. Only after all incumbents were placed were external candidates interviewed. External candidates submitted applications to the district's central human resources (HR) office who then sent principals and site-based hiring teams a long list of external candidates who met minimal criteria (i.e. correct licensure and legal standing for employment). From the district's perspective, while teachers' collective bargaining agreement did not mandate late hiring, in practice, waiting for the internal market to clear first often led to a drawn out process.

Starting in the spring of 2014, the district's central HR office implemented more rigorous screening of external candidates and began making use of the "contract unassigned" part of the MOU. Resume markers such as GPA, licensure, and experience were used to select promising external candidates for preliminary screening interviews. HR staff conducted these short interviews with high mark candidates both in-person at large job fairs and via telephone, ranking them in response. Candidates who rated highly, signed early offer contracts with the district during "Early Contract Meet and Greet" events at the end of April and May.

These early offer contracts were not guarantees of employment nor were they linked to a specific position. Rather, these early offer contracts gave the candidates official employee status so that they could be placed into the district's internal labor market and could interview for open positions alongside incumbents. In 2014, after the policy change, there was no significant change in the overall number of new teachers; however, approximately 25% of new teachers were offered early contracts, and thus their placement was accelerated in a way that had not previously been possible.

Once incumbents and early-offer teachers applied for open positions, principals and site based hiring teams interviewed the five most senior candidates and up to five additional candidates of their choosing to fill out their pool. The additional candidates can be incumbents and/or early offer teachers. There is no formal flag in the centralized system that indicates to principals and/or hiring teams that a candidate is an incumbent teacher or an early-offer teacher, but presumably, hiring committees can deduce the information from candidates' resumes (early-offer teachers generally lack teaching experience and would not list an MPS school as their most recent employment). Anecdotal evidence collected from interviews with HR personnel and school principals suggests that there was some informal information sharing that happened outside of the centralized system where HR personnel alerted principals to early-offer candidates they thought would be a good match for their site. HR provided only basic support to the early offer candidates as they navigated the internal labor market. The candidates received information about the process at the Early Contract Meet and Greet events but were not systematically encouraged to apply to any specific sites.

After the interviews were completed, schools ranked their top four choices and the centralized system made an offer to the top ranked candidate. The system automatically moved to the second ranked candidate if the top ranked candidate declined and so on. Once this process cleared (all positions were filled or all ranked candidates had been offered a position), a second round repeated the process from the beginning. The second round filled positions vacated by

incumbents who moved in the first round as well as any other positions that became open in the meantime. At the end of the second round, any candidate who was still in need of a placement and any school that was still in need of a candidate all met in a large space for a "speed-dating" type final round to try to make a few final matches. It was possible for an early offer candidate to remain unmatched even through this final round. In this case, the district could release an early offer candidate with an early contract the same way they could release any probationary teacher. The early contract offer was not, therefore, a guarantee of employment, but rather an opportunity to compete alongside incumbents. Table 1 provides an overview of the hiring and placement timeline for early-offer candidates.

TABLE 1
Early Offer Hiring and Placement Timeline

Task and Activity	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug
Resumes submitted	X	X	X	X				
Screenings by HR at job fairs and/or phone interviews		X	X	X				
Early offer contracts issued at "Meet and Greet" events				X	X			
Internal labor market				X	X			
Internal labor market final matching round. Unmatched early offer contracts terminated						X		
Standard external hiring process until all positions filled						X	X	X

METHODOLOGY

Data

We relied on two administrative datasets to assess the early hiring process. Variable selection was governed largely by what was collected and available in the internal HR systems. The first dataset contained information gathered during the interview and select process for spring 2014 and spring 2015. This dataset was position-applicant level (i.e. a teacher can be an applicant to multiple open positions). We observed which positions early-offer candidates and incumbent teachers applied to, whether they were granted an interview, and, if so, whether they were ranked and what that ranking was. The second dataset combined administrative HR data regarding position and personnel characteristics (specifically school, subject, step/lane, seniority, age) with teacher evaluation metrics for three school years: 2013-14, 2014-15 and 2015-16. This dataset was teacher-level. The analytic sample for the second dataset included all teachers - all incumbents regardless of whether they applied for a new position and all new teachers, both early-offers and others hired on the more standard timeline - with both human resource data and evaluation scores. Our analytical sample is constrained to only those who have both outcome data and HR data. Therefore, we acknowledge that the comparisons are not necessarily representative of all of the district's teachers. This is largely due to the fact that outcome data disproportionately picks up teachers from certain subjects and grade levels.

MPS evaluated teacher effectiveness with four distinct metrics: value-added math scores, value-added reading scores, student surveys, and formal classroom observations. This multi-faceted approach is in line with best practices outlined by the Measures of Effective Teaching Project (Mihaly, McCaffrey, Staiger, & Lockwood, 2013). All of the evaluation scores have been normalized to a mean of 0 and a standard deviation of 1, and should be interpreted as standard deviations from the mean of the district's overall teacher workforce. A teacher who scores a 1 is one standard deviation above the mean and thus outperformed 84% of her peers in that year. In some models, we pooled value-added math and reading to reduce noise, increase sample size, and improve statistical precision. For all the metrics, we calculated lifetime-averages using inverse-variance weights over the 2013-14 to 2015-16 school years. The lifetime averages included all the data available for a given teacher, thus increasing accuracy. The inverse variance is a direct index of the precision of the effect size. Inverse-variance weighting is used to combine the results from independent measurements. The inverse-variance weighted average has the least variance among all weighted averages, which are calculated as:

$$D^2(\hat{y}) = \frac{1}{\Sigma_i \quad 1/\sigma_i^2}.$$

In 2014, the district hired 419 new teachers between the start of April and the end of September. Of those, 322 were hired on the standard timeline and 97 were early offers (however, only 96 early-offer teachers had at least one outcome score). In the second year of the early offers program, 338 new teachers were hired, 269 on a standard timeline and 69 early offers. Our main analytic sample pooled the two cohorts by lifetime averages giving us a total of 165 unique early hires. In both years there is some "melt," that is, people who are given early offers and then never show up in our dataset. This happens for three reasons (1) they are given an early offer but then take a position elsewhere and withdraw their application before being placed into a school, (2) they continue with the early offer process but are unmatched through the internal labor market and are released, (3) they are successful in the internal labor market and are placed into a position but were not evaluated.

Analysis

We assessed whether and how changes in the MPS hiring process influenced candidate pools, teacher placements and student outcomes via three primary comparisons. First, we compared the interview patterns for early-offer teachers and incumbent teachers who participated in interview and select. Next, we compared early-offer first-year teachers to other first-year teachers hired on the standard timeline and to all incumbents in the district along two dimensions: the characteristics of the school where they worked and their individual evaluation scores.

Recall that our first hypothesis was that the early offers program impacted the candidate pool for hard-to-staff schools and subject areas. Without the early offers program, the pool of candidates available during interview and select would have been only incumbents so the relevant comparison is early-offer teachers to incumbent teachers who applied to one or more open position during interview and select. We tested for differences in age, race and gender -- the demographic markers available in the dataset. Race was coded as white or non-white and gender was coded as female or non-female. For these binary variables, the mean shows the share white and the share female, which were multiplied by 100 to be converted into percentages. .

We report differences in means/percentages between early-offers and incumbents and whether the differences were statistically significant using two-tailed, two-sample t-tests with pooled variance. We tested the result of this simple comparison of means to more sophisticated methodologies including multivariate regression and kernel density plots but in all cases found that the basic t-test was sufficient, and thus favored this methodology given the ease of exposition.

$$H_o : \mu_{early} = \mu_{incumbent}$$
$$H_A : \mu_{early} \neq \mu_{incumbent}$$

μ_{early} is the mean for early-offer teachers who participated in interview and select and $\mu_{incumbent}$ is the mean for incumbent teachers who participated in interview and select.

In addition to teacher-level demographics, we characterized teachers by the number of positions they applied to and we characterized positions by the number of incumbent applicants, the latter being a direct measure of how desirable a position is and thus how "hard-to-staff" it might be. Indirect measures of "hard-to-staff" included school-level demographics (share of students non-white, share of students proficient in reading and math, share of non-white faculty, average years of experience and the pupil-teacher ratio). These school characteristics described the types of positions that early-offer teachers favored relative to the types of positions favored by incumbents.

We were able to infer the quality of the candidates by noting how they fared at each stage of the interview and select process. Here, quality was measured by whether the hiring team was impressed enough with their materials to grant an interview and impressed enough with their interview to rank them for an offer. For each stage we calculated the probability of success. For example, the probability of an incumbent teacher being invited to interview was the count of all incumbent interviews divided by the total number of incumbent applicants. Additionally, since the interview and select process mandated that the most senior teachers be granted interviews, we also calculated probability of being invited to interview excluding these automatic interviews. On one hand, the fact that it is impossible to know whether these most senior teachers would have been invited to interview in the absence of automatic interviews argues for including these teachers in the calculation. On the other hand, the fact that these interview slots were not the purview of the site hiring committee argues for excluding automatic interviews. We report both.

Next, we used the district's main human resource database to assess the differences in placement for incumbent and first year teachers by comparing the placements for each group across several dimensions including experience, zone, and subject. Here we made two sets of comparisons:

$$H_o : \mu_{early} = \mu_{incumbent} \qquad\qquad H_o : \mu_{early} = \mu_{othernew}$$
$$H_A : \mu_{early} \neq \mu_{incumbent} \qquad\qquad H_A : \mu_{early} \neq \mu_{othernew}$$

As before, μ_{early} was the mean for early-offer teachers and $\mu_{incumbent}$ was the mean for incumbents. We also compared $\mu_{othernew}$, the mean for other new teachers who were hired on the more standard timeline (these teachers were not allowed to participate in interview and select, so they were not in the previous data set). The teacher demographics we observed were age, advanced degree (i.e. Master's or higher) and years of experience. We also employed a range of characteristics about the job where the teacher was placed including subject and grade level, where the school was located, and whether it was considered a "focus" or "priority" school. All of these

variables gave us a sense for how the early offers program changed the district's hiring outcomes across different types of job postings.

Finally, we compared mean value-added and student survey evaluation scores for early-offer, other first year, and incumbent teachers. Using all three evaluation metrics described above (value-added, student surveys, and formal observations), we could address our second research hypothesis and investigate whether early-offer teachers were better hires than teachers hired on a more standard timeline.

RESULTS

We present results in three parts, each corresponding to the analyses described above. First, we report on how early-offer teachers fared in the internal labor market. We see the schools to which early-offer teachers chose to apply, whether they were granted interviews, and, if so, whether they were ranked highly by the principal and/or hiring team at each school. Our second set of results illustrates differences in the placements of incumbent and first year teachers. This is important because the early offers program was intended to help fill hard to staff vacancies such as special education and ESL/bilingual. Finally, we show the ex-post effectiveness measured by lifetime averages of value-added, student surveys and formal observations for early-offer teachers compared to incumbent teachers and other new teachers.

Comparing Teacher Preferences

First, we assessed whether early-offer teachers had different preferences than incumbents. The relationship between early-offer and incumbent preferences over school characteristics is theoretically ambiguous. On one hand, teacher candidates who commit to a large urban district early in the hiring season may be disproportionately interested in working for schools that serve large numbers of low-income and non-white students. On the other hand, evidence from previous work shows that, holding salary constant, teachers prefer to work at schools that serve higher achieving, higher-income and whiter student bodies, and early-offer teachers may be no different.

Table 2 compares early-offer teachers to the incumbent teachers who participated in the internal labor market, pooling data from 2014 and 2015. The early-offers were a sizable component of the internal labor market, comprising fully 15% of all teachers who applied for open positions in these two years. It is important to remember that the comparison group here was not all incumbent teachers but rather incumbents who self-selected into the interview and select system, indicating that they were interested in changing positions. Demographically, early-offer teachers and incumbents were equally likely to be female and non-white. Early-offer teachers were much younger than incumbents. The average early-offer teacher was 13.65 years younger than the average incumbent who participated in interview and select.

TABLE 2
Summary Statistics of Early-Offer and Incumbent Teachers in Internal
Interview & Select System, 2014-2015 Pooled School Years

	(1) Incumbent	(2) Early	Difference in Means (Early-incumbent)
% Female	75.00	75.00	0.00
Avg. Age	42.42	28.77	-13.65 ***
% Non-White	15.00	16.00	1.00
Avg. Number of Applications	6.88	14.01	7.13 ***
Number of incumbents in pool	10.71	8.90	-1.81 ***
School Attributes			
% Non-white Teachers	12.96	14.03	1.08 **
% Non-white Students	67.05	76.19	9.14 ***
Avg. Years of Teacher Exp.	13.35	12.46	-0.89 ***
% Students Proficient in Reading	52.7	46.67	-6.04 ***
% Students Proficient in Math	46.12	37.94	-8.18 ***
Pupil-teacher Ratio	16.15	15.17	-0.98 ***
Interview			
Prob. Interview	0.52	0.41	-0.11 ***
Prob. Interview (excluding automatic interviews)	0.15	0.41	0.26 ***
Prob. Ranked | Interview	0.38	0.47	0.09 ***
Prob. Ranked First | Interview	0.18	0.21	0.04 ***
Maximum Unique Teachers	879	160	
Maximum Unique Applications	5,932	2,454	

Note: *** $p < .01$, ** $p < .05$, * $p < .1$.

Early-offer teachers applied more widely than incumbent teachers. On average, early-offer teachers applied for 14.01 positions while incumbents applied for only 6.88. Figure 1 compares the distribution of applications by teacher type. The modal incumbent applied for one or two positions while the modal early-offer candidate applied for 10 or more. This reflects the fact that incumbents generally have an existing position to fall back on so they only apply for open positions they find more desirable than their current job while early-offer teachers are effectively unemployed since they have no confirmed position and thus apply for any job they would find minimally acceptable.
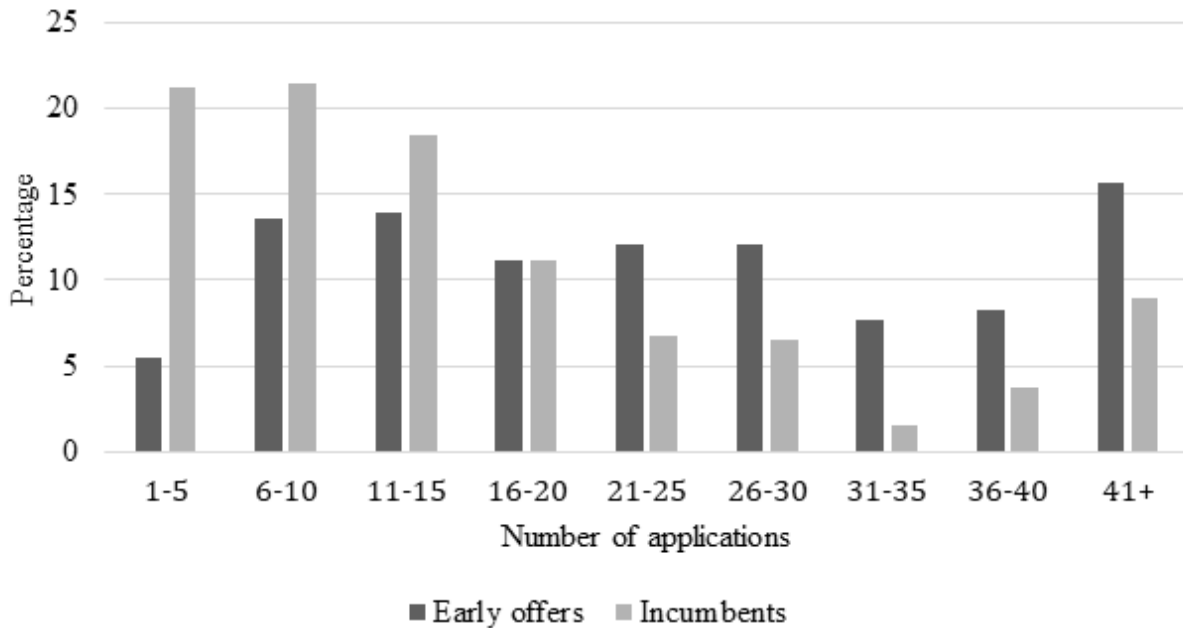
Figure 1. Number of applications by incumbents and early offers, 2014-2015 pooled school years.

Table 2 also describes the types of schools that early-offer teachers chose to apply to and how this compared to incumbents. For teachers who apply to positions at more than one school, we took the average across all their applications. On average, early-offer teachers applied to schools that had more non-white staff and students and where teachers had fewer years of experience. Rather than use school characteristics to indirectly measure desirability, we used a more direct measure of interest in schools: the size of the applicant pool. The number of incumbent teachers who applied to a given position is a good indicator of how attractive the position is to teachers. In the first panel of Table 2, we see that incumbent teachers apply to positions that attracted an average of 1.81 more incumbent applicants indicating that early-offer teachers were overrepresented for positions with fewer applicants.

The third panel of Table 2 describes how early-offer teachers fared with their applications. We found that early-offer teachers were significantly less likely to be granted an interview than incumbent teachers; 52% of incumbent teachers who applied were granted an interview while only 41% of early-offer teacher applications were. Much of this was driven by the fact that the four most senior teachers who applied to a position were automatically given an interview. When we excluded these automatic interviews, the findings reversed and early-offer teachers actually got more interviews than incumbents - fully 41% of early-offer teachers were granted interviews compared to only about 15% of incumbent teachers.

Furthermore, when early-offer teachers progressed to the interview stage, they were more likely to be ranked as one of the top four candidates. Post-interview, we found that 38% of incumbents were ranked compared with 47% of early-offer teachers. When we focused on teachers who were ranked first for a given position, and thus had the first right of refusal, early-offer teachers again edged out incumbents 21% to 18%. We further characterized how well they fared by how many offers they received. Figure 2 shows the distribution of offers in any given round of the internal labor market process. More than half of those granted offers received only one offer

in a given round; however, a few teachers received many offers. Early-offer teachers were more likely to receive multiple offers than were incumbents. Approximately 13.5% of early-offer teachers received three or more offers in a given round; the same figure for incumbents was 4.7%.
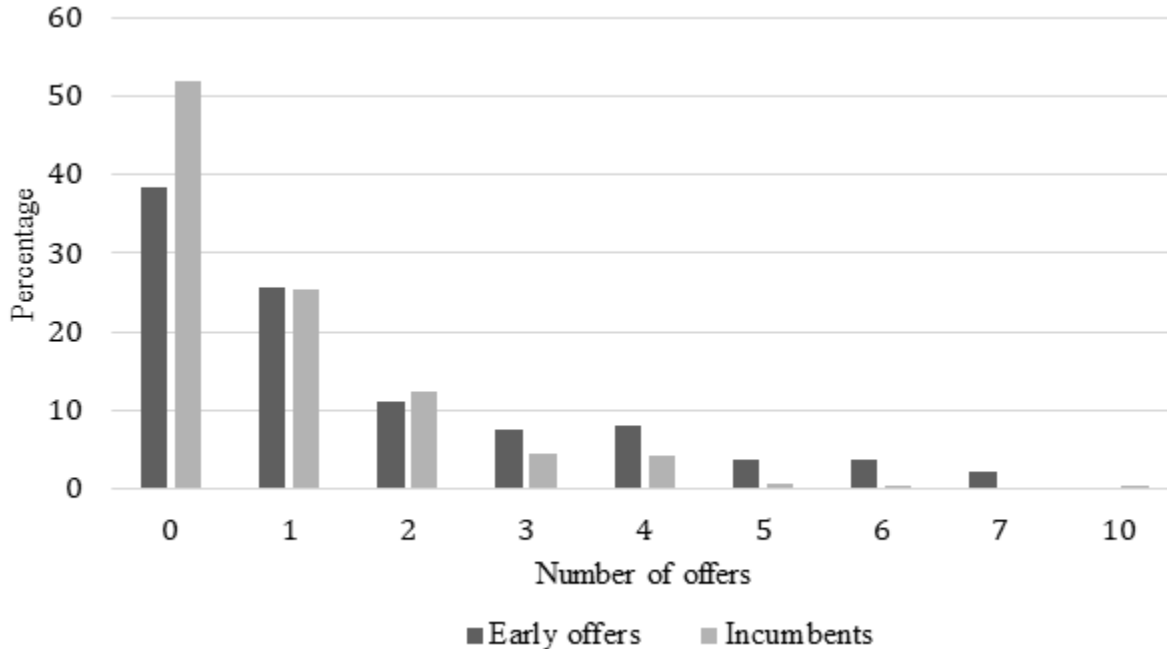

Figure 2. Number of job offers during any round by incumbents and early offers, 2014-2015 pooled school years.

Not all early-offer teachers secured a position through the interview and select process, and some remained unmatched as the internal labor market iterated through offers. We found that 7.5% of early-offer teachers were not picked up by a school during either of the two interview rounds. These teachers then participated in the "speed-dating" style attempt to connect unmatched teachers and open positions. The share of early-offer teachers who remained unmatched after that final opportunity was not statistically different than the 5.6% of incumbents who were also unmatched.

Comparing Teacher Placements

The early offers program was intended to secure outstanding teacher candidates and to facilitate placement of these teachers into classrooms where they were needed most. We characterized need by demographics (low income, low test scores) and in subject area (special education, English second language, science). Table 3 summarizes the positions for all district teachers broken out by incumbents and first-year teachers who are further divided into early-offer teachers and other new teachers. For consistency with the next table that focuses on evaluation scores, we limited the analytic sample to only those teachers with at least one lifetime average outcome score. We found that early-offer teachers were a younger, less experienced and less academically credentialed group than both incumbents and other new teachers.

TABLE 3
Placements for District Teachers with Evaluation Scores
2014-2015 Pooled School Years

| | (1) Incumbent | First year teachers | | Difference in Means | |
		(2) Early	(3) Other new	Early – other	Early - incumbent
Age	44.70	28.62	35.06	-6.44 ***	-16.08 ***
Advanced degree (lane)	0.74	0.36	0.44	-0.08 *	-0.37 ***
Experience (step)	17.00	3.40	6.69	-3.30 ***	-13.61 ***
Zone, Priority, Focus					
Zone1	0.29	0.40	0.42	-0.03	0.10 ***
Zone2	0.29	0.27	0.23	0.03	-0.02
Zone3	0.30	0.32	0.27	0.06	0.01
Priority School	0.12	0.28	0.21	0.08 **	0.17 ***
Focus School	0.29	0.38	0.27	0.11 **	0.09 ***
Subject					
Elementary	0.34	0.36	0.27	0.09 **	0.02
Secondary (Math, English, Science, Soc.)	0.21	0.27	0.20	0.07 *	0.07 **
Math	0.04	0.12	0.05	0.07 ***	0.07 ***
Science	0.04	0.05	0.04	0.00	0.01
ESL/Bilingual	0.07	0.14	0.09	0.05 *	0.07 ***
Special Ed	0.23	0.13	0.19	-0.06 *	-0.10 **
K12 Specialists	0.13	0.08	0.19	-0.11 ***	-0.05 **
Placement					
Elementary	0.59	0.68	0.57	0.11 ***	0.10 **
Middle Schools	0.11	0.09	0.16	-0.06 *	-0.02
High Schools	0.20	0.22	0.21	0.01	0.01
Unique teachers	2,491	162	642		

Notes. Experience (step) was not available for 2015 new teacher hires. *** $p < .01$, ** $p < .05$, * $p < .1$.

The second panel of Table 3 shows the distribution of teachers by zone, priority school status and subject. The district was divided into three zones. One way to understand the differences in zones is to consider state-identified schools with low student achievement or large achievement gaps. Minneapolis had 21 state-defined "Focus" schools and 11 "Priority" schools. The state Department of Education releases a list of Priority and Focus schools every three years. Our data reflects a Department of Education list a year prior to the start of the Early Offers policy, 2013-14. A Priority school is a school identified in the most persistently low-performing 5% of Title I schools. The 10% of schools with the largest achievement gaps statewide are declared Focus schools. Looking at Focus schools by zones, there were six Focus schools in Zone 1, ten Focus

schools in Zone 2, and five Focus schools in Zone 3. Nine of the district's "Priority" schools were in Zone 1. This zone was home to the most students of color, the most students eligible for free/reduced price meals, and the highest number of students who were homeless or highly mobile. More first-year teachers were placed in Zone 1 than either of the other two zones. This zone had 40% and 42% of early-offer and other new teachers, respectively, compared to 29% of incumbent teachers. While the difference between early-offer and other new teachers in Zone 1 was not statistically significant, the difference in proportions of early-offers and incumbents in this zone was significant (p < .01). Likewise, both early-offer teachers and other new teachers were more likely than incumbents to be placed in focus or priority schools. Here, however, we saw that early-offer teachers were even more likely than other new teachers to be in high-need schools. We observed that 28% of early-offer teachers were placed in priority schools, compared to 21% of other new teachers. This difference was significant (p < .05). Moreover, 38% of early-offer teachers and 27% of other new teachers were placed in focus schools compared to 30% of incumbents, and these differences were also statistically significant.

In the bottom two panels of Table 3, we look at the distribution of teachers by subject and grade level. Early-offer teachers were more likely than other new teachers to work in elementary schools, secondary subject areas (math, English, science, and social studies), and English language learner or bilingual classrooms. Looking at math, 12% of early-offers were math teachers compared to only 5% of other new teachers, but early-offers and other new teachers were just as likely to be science teachers. Early-offer teachers were less likely to be special education teachers or K-12 specialists. K-12 specialists include physical education teachers, business, family and consumer science, health, world languages, reading, industrial technology, art, media, music, technology, and theater/dance. This pattern repeated when we compared early-offer teachers to incumbent teachers.

Comparing Teacher Outcomes

In this section, we compare the measured outcomes of early-offer teachers to other new teachers and to incumbents. Table 4 reports the lifetime average effectiveness of early-offer, other new, and incumbent teachers using the district's three metrics: value-added, student surveys and classroom observations.

We found that, on nearly every metric, incumbents outperformed early-offer teachers and other new teachers hired on a standard timeline, and in several cases (pooled value-added, value-added math and formal observations) differences were large and statistically significant. Early offer teachers, however, outscored other new teachers on only student surveys. While the means for early offer teachers are consistently higher than other new teachers, there is insufficient evidence to reject the null hypothesis of no difference.

Recall that some early-offer teachers were ranked highly right away, others were not picked up even through the final matching process. When we focused on the early-offer candidates who were given one or more specific job offers in the first round of the internal labor market interview process we found that these teachers outperformed other new teachers on both the student surveys and the formal observations (results are available upon request). The fact that the formal observation difference was then statistically significant indicates that the characteristics that make these candidates stand out in an interview setting are the same characteristics which make them stand out when observed in the classroom.

TABLE 4
Summary Effectiveness Outcomes of the District's Evaluated Teachers
2014-2015 Pooled School Years

Effectiveness Measure	First Year Teachers			Difference in Means	
	(1) Incumbent	(2) Early	(3) Other new	Early – other	Early - incumbent
VA pooled (N)	1385	102	251		
Mean	0.00	-0.11	-0.23	0.12	-0.11 *
VA reading (N)	1185	82	207		
Mean	-0.02	-0.01	-0.18	0.17	0.01
VA math (N)	934	77	176		
Mean	0.00	-0.22	-0.38	0.17	-0.22 **
Student surveys (N)	2226	141	476		
Mean	0.00	-0.06	-0.25	0.18 *	-0.06
Formal obs. (N)	2855	161	631		
Mean	0.00	-0.57	-0.71	0.14	-0.58 ***
Unique teachers	2855	162	642		

Note: *** p < .01, ** p < .05, * p < .1.

Robustness and Limitations

One potential worry is that even if the early offer program successfully identified strong candidates, these candidates might very well have been hired in the absence of the early offers program. That is, they may have persisted with the standard hiring timeline, in which case the time and energy invested in the screenings could have been better spent on other HR initiatives. We were able to test this possibility for a subset of the candidates for whom we had full screening interview information. We observed 52 candidates who were not invited to participate in the early offers hiring process but ended up being hired by the district later in the summer on the more standard timeline. These teachers, as a group, had slightly lower, but statistically indistinguishable, effectiveness scores than early-offer teachers.

A second worry could be that the early offers program was differentially effective and that average outcomes were masking important variance. Contrasting means can hide differences in the tails of a distribution, and since the goal of extending early offers was to identify teachers who are highly effective, one could also deem the early offers program successful if it was able to secure teachers at the very top or screen out teachers at the very bottom of the effectiveness distribution. As a robustness check, we also compared teachers using kernel density plots. Kernel density plots estimate the underlying distribution of effectiveness using nonparametric methods to estimate the probability of observing teachers in any given range of effectiveness. We found no evidence that early-offer teachers were more (less) likely to be at the top (bottom) end of the average lifetime effectiveness distribution.

DISCUSSION

In sum, the early offers program was intended to secure outstanding teacher candidates and to facilitate placement of these teachers into classrooms where they were needed most. The district was able to increase the size of the applicant pool at less desirable sites by adding early-offer teachers to the internal labor market. For those granted interviews, the early-offer candidates fared better than the incumbent candidates. This was well-aligned with the goal of giving early offers to candidates who the central office felt would be desired by multiple principals and have little trouble finding a position.

In our analysis of teacher placements, we found that early-offer teachers were more likely than incumbents to work in the schools in Zone 1, home to the most challenging conditions in the district, but so were other new teachers. Early-offer teachers were also more likely than incumbents to work in focus and priority schools, and here they did also differ from other new teachers. The program had mixed results in high need subject areas. The early offers program was able to secure an above average number of math and ESL/bilingual teachers. It was not, however, able to secure an above average number of special education teachers.

In our analysis of teacher outcomes, we saw that among first-year teachers, early-offer teachers outperformed those hired on a standard timeline, but this result was only statistically significant for student surveys. The suggestive evidence demands further research. More data is necessary to discern whether the differences between early-offer teachers and other new teachers are robust. It may be that student surveys are an important indicator if students are able to observe learning both within themselves and amongst their peers before that learning can be measured on standardized tests or made clear to observers who only see one or two lessons. Given that screening interviews were looking for teachers who had a strong growth mentality and who were receptive to performance feedback, it could also be that the teachers who were identified for early offers have a steeper growth trajectory and in subsequent years they will outperform the teachers hired on the standard timeline by wider margins. If this is the case, differences registered after the first year or two of teaching would be smaller than differences detectable mid-career. We leave this for future research.

Conclusion

MPS sought to identify and hire outstanding teachers early in the hiring process. To do this, they implemented centralized and standardized screening interviews and extended offers in April and May to standout candidates. The early-offer teachers fared well in the internal labor market where they were more likely to be chosen for interviews (once controlling for default rules advantaging the four most senior teachers who apply for a position), more likely to be ranked highly, and more likely to be offered positions than incumbent teachers who applied for openings. This was despite the fact that early-offer teachers were younger, had less experience, and were less likely to hold advanced degrees than incumbent teachers.

We also found that the district had some success getting these teachers to the classrooms where they were needed most. On average, early-offer teachers were more likely to apply to and be hired into schools that were less desirable to incumbents. These schools had higher shares of non-white students as well as more students below proficiency in math and/or reading. Early-offer teachers were more likely to be hired into "focus" or "priority" schools and for ESL/bilingual and

math positions. Early-offer teachers, however, were not more likely to be special education teachers, another hard to staff area. Though securing staff for difficult positions is notable, it remains to be seen whether the early-offer teachers remain in these positions in the long run.

The program's promising interview/placement patterns are encouraging and consistent with the conclusion that the early offers program changed hiring patterns for this large urban district. The best test of the program's success, however, is to see if these stand out candidates turned out to be stand out teachers. Here we have suggestive but not conclusive evidence. The cohorts studied had been in the classroom for one or two years. In that short tenure, we found that they outperformed other new teachers on most measurable outcomes but that most the differences rose to the conventional levels of statistical significance. This leads us to believe that these teachers may be above average teachers, but more data is needed before we can make a strong claim. A more conservative, and entirely appropriate, conclusion is that early-offer teachers are indistinguishable from their peers hired on a more traditional timeline. Some of these candidates may have been hired anyway later in the summer. Even if this is the case, the early placement is beneficial if it allows for better onboarding and preparation during the summer months.

The early offers program is an example of a low-cost intervention to improve the efficiency in teacher labor markets. That said, even a low-cost intervention requires investment of time and district HR resources. It also required a good amount of political capital in negotiation with the teachers' union. The main cost to the program is thus the opportunity cost of staff time and buy-in from senior teachers who now have to compete for positions against early-offer candidates. If the early-offer teachers turn out to be no better or only slightly better than other new hires, the benefits of the accelerated timeline may not be large enough to justify the opportunity cost.

Another concern is that the early offer program did not successfully secure more special education teachers nor did it produce a more diverse pool of candidates. As the early offer program has evolved, the district has placed more emphasis on using it for hard to staff areas and also as a tool to secure high performing student teachers. Future research will follow the initial cohort of early-offer teachers to assess retention and longer-term effectiveness outcomes as well as newer cohorts where the program was more targeted and also more formally defined by the collective bargaining contract.

One final policy implication is important to note. Our examination has allowed us to explore the early offer program's operation and consequences within a single urban district. That is, this paper is a firm-specific, rather than market-level analysis. State-level policymakers interested in reforming teacher hiring standards and practices to address teacher shortages would need to pay attention to how one district's reforms may influence a neighboring district's overall supply of labor and distribution of teacher talent and future research should address these larger equilibrium issues.

REFERENCES

Bruno, P., & Strunk, K.O. (2018). *Making the Cut: The Effectiveness of Teacher Screening and Hiring in the Los Angeles Unified School District*. CALDER Working Paper No. 184. Available at: https://files.eric.ed.gov/fulltext/ED583631.pdf.

Dee, T. S., & Goldhaber, D. (2017). Understanding and addressing teacher shortages in the United States. *The Hamilton Project*, (April), 1–28.

Engel, M. (2009). *Principals' and districts' hiring practices: An exploration of the demand side of the teacher labor market*. Northwestern University.

Engel, M. (2012). The timing of teacher hires and teacher qualifications: Is there an association? *Teachers College Record*, *114*(12), 1-29.

Ferguson, R. F., & Danielson, C. (2015). How framework for teaching and tripod 7Cs evidence distinguish key components of effective teaching. *Designing teacher evaluation systems: New guidance from the measures of effective teaching project*, 98-143.

Goldhaber, D., Grout, C., & Huntington-Klein, N. (2017). Screen twice, cut once: Assessing the predictive validity of applicant selection tools. *Education Finance and Policy*, 12(2), 197-223.

Jatusripaitak N., Mykerezi E., Sojourner A., & West K. (2018). School and teacher preferences: Evidence from a multi-stage internal labor market. Working paper.

Johnson, S. M., Kraft, M. A., & Papay, J. P. (2012). How context matters in high-need schools: The effects of teachers' working conditions on their professional satisfaction and their students' achievement. *Teachers College Record*, 114(10), 1-39.

Johnson, S. M., & Marietta, G. (2009). Taking human resources seriously in Minneapolis. Public Education Leadership Project at Harvard University, PEL No.-055.

Levashina, J., Hartwell, C. J., Morgeson, F. P., & Campion, M. A. (2014). The structured employment interview: Narrative and quantitative review of the research literature. *Personnel Psychology*, *67*(1), 241-293.

Levin, J.,& Quinn, M. (2003). Missed opportunities: How we keep high quality teachers out of urban classrooms. Washington, DC: New Teacher Project.

Liu, E., & Johnson, S. (2006). New teachers' experiences of hiring: Late, rushed, and information-poor. *Educational Administration Quarterly* 42(3), 324-360.

Mihaly, K., McCaffrey, D. F., Staiger, D. O., & Lockwood, J. R. (2013). A composite estimator of effective teaching. *Bill & Melinda Gates Foundation: Seattle, WA*.

Minnesota Department of Education. (2013). Content Archive: State's 2013 Accountability Results Show Significant Progress for Priority and Focus Schools. Retrieved from http://education.state.mn.us/MDE/ContentArchive/052702

The New Teacher Project. (2007). Hiring, assignment, and transfer in Chicago Public Schools. Report from The New Teacher Project, July 2007. Retrieved from http://www.tntp.org/files/TNTPAnalysis-Chicago.pdf

Papay, J. P., & Kraft, M. A. (2016). The productivity costs of inefficient hiring practices: Evidence from late teacher hiring. *Journal of Policy Analysis and Management*, 35(4), 791-817.

Van Cleef, V. (2013). "Missed Opportunities, 10 Years Later" The New Teacher Project (November 18). Available at: https://tntp.org/blog/post/missed-opportunities-10-years-later. Accessed on June 6, 2018.