ESE Policy Brief • October 2017

Teacher Equity Gaps in Massachusetts

Effective teachers make a real difference for student learning. But research shows that both in Massachusetts and nationwide, academically struggling students and those from historically low performing subgroups are less likely to be assigned to the teachers who are most likely to generate strong results. This results in missed opportunities to close achievement gaps and increase educational outcomes for all students.

This policy brief provides an overview of how effective teachers are identified, summarizes research from around the nation, and analyzes Massachusetts data to address several important questions:

- How much difference can an effective teacher make?
- How do researchers measure teacher effectiveness?
- Which teacher characteristics are associated with stronger student outcomes?
- Are there gaps in access to effective teachers in Massachusetts? If so, how consequential are those gaps likely to be for disadvantaged students?
- What are the sources of inequity in teacher assignments in Massachusetts?
- What policies can Massachusetts districts and schools adopt that show evidence of increasing teachers' effectiveness or increasing equitable access to effective educators?
- What additional resources are available from the Department of Elementary and Secondary Education about access to effective educators?

Written by: James Cowan, *American Institutes for Research;* Dan Goldhaber, *American Institutes for Research and University of Washington*; and Roddy Theobald, *American Institutes for Research*

MASSACHUSETTS DEPARTMENT OF ELEMENTARY AND SECONDARY EDUCATION

www.doe.mass.edu

Key Findings About Teacher Equity Gaps in Massachusetts

Effective teachers make a real difference for student learning. But research shows that, both in Massachusetts and nationwide, academically struggling students and those from historically low performing subgroups are less likely to be assigned the teachers who generate the strongest results. This results in missed opportunities to close achievement gaps and increase educational outcomes for all students.

In this brief, we document several key facts about teacher equity gaps in Massachusetts:

- Compared to the average Massachusetts teacher, the 60th percentile teacher raises student achievement by the equivalent of an additional four weeks of learning per year. The 75th percentile teacher improves the achievement of their students by about 13 to 15 weeks of learning.
- Assigning a Massachusetts student to a 60th percentile teacher every year from fourth to eighth grade corresponds to about two additional months of learning in math over those five years, compared to assigning that student to an average teacher every year.
- Students assigned a teacher earning an exemplary evaluation accrue about nine to ten additional weeks of student learning per year relative to those assigned a proficient teacher. The difference between an exemplary teacher and an unsatisfactory one is even greater, equivalent to about 18 to 24 additional weeks of learning.
- In Massachusetts, the average low income student is assigned to a teacher who generates two fewer weeks of learning in mathematics and four weeks fewer in English language arts per year than the teachers assigned to non-low income students.
- Low income students in Massachusetts are 31 percent more likely to be assigned to teachers with less than three years of experience and more than twice as likely to be assigned to a teacher who earns an evaluation of unsatisfactory or needs improvement, as compared to non-low income students.
- In Massachusetts, inequitable access to effective teachers for low income students increases achievement gaps by up to three weeks of learning in mathematics and six weeks in English language arts between fourth and eighth grade.
- Three-quarters of the teacher equity gap for low income students is explained by the fact that low income students are disproportionately enrolled in districts with lower average teacher effectiveness.

This policy brief provides an overview of how effective teachers are identified, summarizes research from around the nation, and analyzes Massachusetts data to address the important issue of access to effective educators. It also provides connections to resources available to Massachusetts schools and districts working to eliminate equity gaps.

How much difference can an effective teacher make?

Empirical evidence has consistently shown classroom instruction to be one of the most important in-school factors affecting student learning. Teachers also affect students' long-term academic and economic outcomes, including educational attainment and earnings (Chetty et al., 2014b).

Compared to an average teacher, the most effective 40 percent of teachers in Massachusetts increase student achievement on standardized tests by the equivalent of about one month of learning per year.

Looking across different states, grade levels, and standardized tests, researchers have found that the most effective teachers have educationally meaningful effects on student learning. The most common way that researchers measure teachers' effectiveness is by examining their impact on students' standardized test scores. (In the next section we will discuss the details of how this is calculated, along with several other measures of effectiveness.) Students who are assigned to the most effective teachers as measured by student test score gains experience substantially larger increases in learning than other students. In Massachusetts, the top 40 percent of teachers raise student achievement by *at least* the equivalent of four weeks of learning relative to the average teacher.[1] The top 25 percent of teachers improve the achievement of their students by 13 to 15 weeks of learning in both math and ELA, or more than one-third of a nine-month school year.[2]

Effective educators also improve other outcomes beyond test scores.

In addition to improving student achievement, teachers influence a variety of other important educational outcomes. Researchers found that in New York City public schools, even relatively small differences in teacher effectiveness had long-term effects on student outcomes. For instance, their results suggest that just moving a student from the median teacher in Massachusetts to a 60th percentile teacher in *one year* would increase college attendance by age 20 by about 1 percent and earnings at age 28 by about $120 (Chetty et al., 2014b). These differences quickly add up: the 60th percentile teacher increases the present value of lifetime earnings by about $65,000 for *each* classroom taught.[3]

Teachers can also affect students' future success even if they don't improve their test scores. Researchers have estimated teacher impacts on students' attendance and classroom behavior and found that some teachers are more effective at improving these "non-cognitive" skills than they are at increasing achievement on standardized tests (Gershenson, 2016; Jackson, 2016). Moreover, these effects on short-term, non-test outcomes predict students' likelihood of high school completion and their college plans. Results from one study suggest that assigning students to the teacher at the 60th percentile of the "non-cognitive" effectiveness distribution rather than the median reduces the dropout rate by about 3 percent (Jackson, 2016).

The impacts of effective teachers on student outcomes accumulate over time. Assigning a Massachusetts student to a 60th percentile teacher every year from fourth to eighth grade corresponds to about eight additional weeks of learning over those five years, compared to assigning that student an average teacher every year.

Being assigned even one effective teacher can make a difference for students, and research also shows that those impacts persist over time. Looking at students' learning results when they are assigned to a less effective teacher the following year, researchers have found that about one-half of teacher effects persist onto the next year's standardized test and about one-fifth persist five years later (Chetty et al., 2014b). Because the effects of instruction are cumulative, having an effective teacher in multiple years improves student achievement even more than assignment to an effective teacher in a single year.

Figure 1 below displays the impact of highly effective teachers over time. We define a highly effective teacher as the 60th percentile of the distribution of teachers, or one who improves student learning relative to the

average teacher by about an additional four weeks per year. We then show typical student results over five years, from grade 3 to grade 8, with three scenarios compared to students with average teachers every year:

- a highly effective teacher in all five years
- a highly effective teacher in fourth grade only
- a highly effective teacher every other year

The immediate impact of a highly effective fourth grade teacher is that students gain a month of learning compared to those assigned an average teacher. As student study new material, or forget some of what they learned previously, some of the effects of past teachers fade out with time. Thus, although assignment to highly effective teachers in fifth through eighth grade continues to increase student learning compared to students with average teachers, it does so at a slower rate. By eighth grade, the total additional effect of consistent access to highly effective teachers corresponds to about eight weeks of learning, or about twice the effect of having a highly effective teacher in a single year alone.

On the other hand, even occasional access to the most effective teachers can have long-run impacts on student achievement. Students assigned to a highly effective teacher only in fourth grade will tend to retain about a fifth of those results by eighth grade, or nearly one week of additional learning. Students who have an effective teacher every other year would reach eighth grade with gains amounting to about five weeks of additional learning, or about 20 percent more than having an effective teacher in eighth grade alone.

Figure 1 also shows how achievement gaps among students can develop over time. We know that there are consistent differences in educational outcomes for groups of disadvantaged students such as low income or African American students. If the effects of highly effective teachers are cumulative and advantaged students tend to have more highly effective teachers, then Figure 1 demonstrates how these achievement gaps will be exacerbated over time. Correspondingly, schools that make a concerted effort to assign disadvantaged students a sequence of highly effective teachers may progressively reduce achievement gaps.

Figure 1. The Effects of Access to Highly Effective Teachers Accumulate Over Time

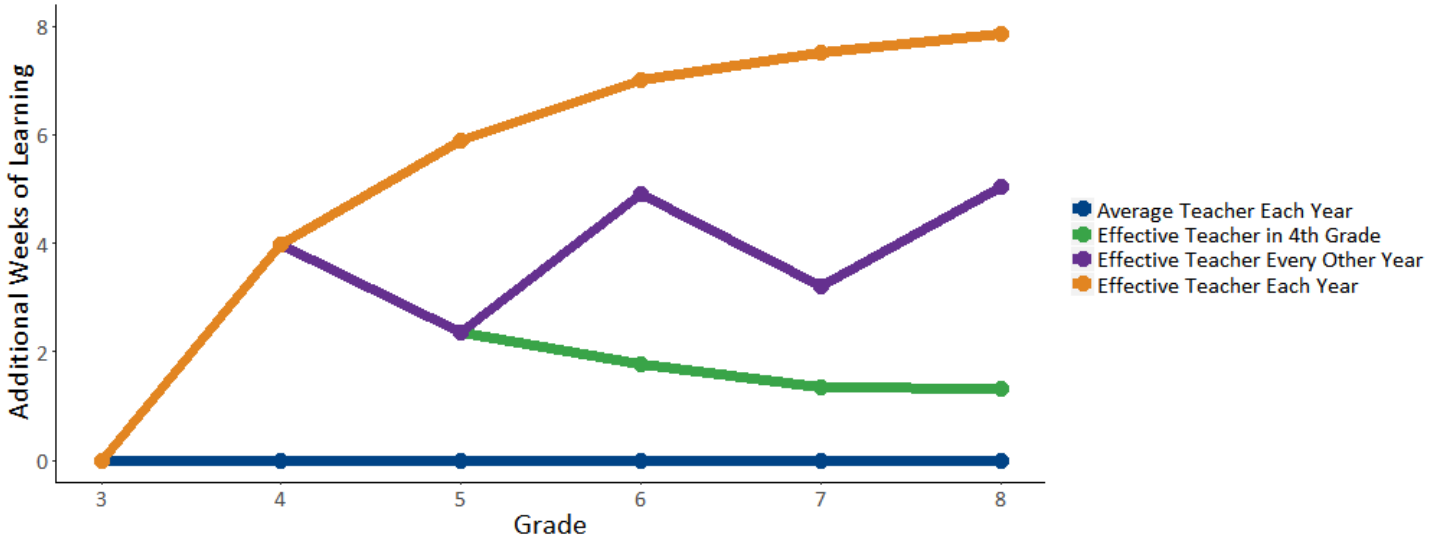

How do researchers measure teacher effectiveness?

Education researchers have used a variety of performance measures to assess teacher effectiveness. We focus on two kinds of measures: those based on direct assessments of student outcomes, such as standardized tests, and those based on qualitative assessments of specific teaching competencies. Researchers have generally found that both outcomes-based measures, such as teacher value-added, and assessments of various facets of teaching practice, such as classroom observations, predict which teachers will most improve student achievement.

Measures of effectiveness that use assessment data typically focus on differences in the change in student test scores over time.

Measures that incorporate student achievement data have become common in education research. Because they are based on direct measures of student learning, they provide the best predictions of future student performance on achievement tests (Mihaly et al., 2013). Test-based measures of teacher effectiveness also correlate with teachers' impacts on important student outcomes, such as college attendance and student earnings (Chetty et al., 2014b).

Researchers and policymakers are interested in how much teachers improve test scores, regardless of where students start at the beginning of the year. It is not easy to measure this accurately, and much effort has been put into developing statistical models to accomplish it. These statistical models are generically called *value-added models* (VAMs).[4] Because these measures depend on student data, their interpretation as indicators of teacher effectiveness warrants some caution. Other factors about students' lives may influence how quickly they improve academically independently of the impact of their teachers, and in ways that are difficult to measure. For instance, high income students may receive personal tutoring outside the classroom, or have better nutrition, housing, or other economic resources that affect their ability to learn. The statistical methods of VAMs partially control for these differences. Nonetheless, if these adjustments fail to capture the effects of all the resources available to non-low income students, then we may not see the full impact an effective teacher could have on low income students. That is, teachers assigned non-low income students may appear more effective than they would be if they were assigned less advantaged students. Overall, recent research suggests that the importance of these student background factors is likely to be minimal (Bacher-Hicks et al., 2017; Chetty et al., 2014a).

Qualitative assessments of teaching practice provide additional information about teachers' contributions to student learning.

Evaluations of teacher performance based on qualitative evidence are sources of information about teachers' practice that can provide a meaningful assessment of teachers' effects on student learning. Using data from other states, researchers have found that classroom observations contain useful information about teachers' contributions to student learning (Garrett & Steinberg, 2015; Grossman et al., 2013; Kane et al., 2011). Observational measures also predict a variety of other important outcomes, including teachers' content and pedagogical knowledge and professional contributions (Harris & Sass, 2014).

Consistent with prior research, qualitative teacher evaluations in Massachusetts predict student achievement gains. Teachers in Massachusetts are assigned a summative performance rating based on evaluators' professional judgments incorporating information from observations and artifacts of practice, measures of student learning, other evidence such as student or staff feedback, and educator goal attainment. Statewide in 2016, 84.2 percent of Massachusetts teachers were rated proficient, while 11.5 percent were exemplary, 3.9 percent were needs improvement, and 0.4 percent were unsatisfactory. In Figure 2, using data from 2014 and 2015, we estimate the average differences in student improvement produced by teachers by summative performance rating.[5]

The bars in each figure compare average student achievement gains by teacher rating category to Massachusetts teachers receiving the proficient rating. Students assigned an exemplary teacher accrue about 8 to 10 additional weeks of learning relative to those assigned a proficient teacher. The difference between an exemplary teacher and an unsatisfactory one is even greater, equivalent to about 18 to 24 additional weeks of learning.

As with value-added measures, qualitative results may be affected by factors unrelated to teachers. For instance, high achieving students may better model the learning behavior that observers attempt to measure. Evaluators may also have implicit or explicit biases about the performance of minority students and thus rate the teacher lower in an observation (Gershenson et al., 2016). Either case conflates characteristics of students with the effectiveness of the classroom teacher. Researchers are still exploring the importance of these factors for qualitative measures of teacher practice, but some studies have found that teachers assigned higher achieving students perform better on classroom observations of teacher practice (Steinberg & Garrett, 2016; Whitehurst et al., 2014). Similar to the value-added case, biases in these measures would tend to work against teachers assigned low-achieving students, and tend to overstate true differences in access to effective teaching.

Figure 2. Teacher Performance Ratings Predict Student Improvement

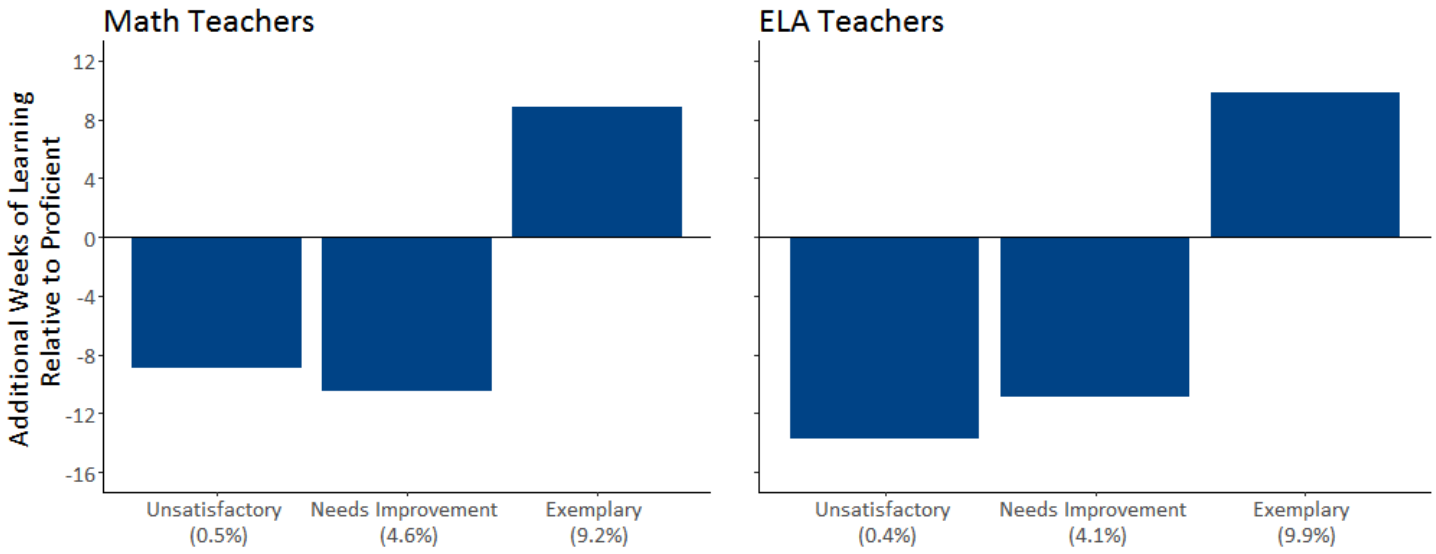

Which teacher characteristics are associated with stronger student outcomes?

So far we have focused on direct measures of teacher effectiveness: their impact on student learning, their performance ratings, and students' perceptions of their practice. Another way to analyze effectiveness is to identify *characteristics* of teachers that are related to effectiveness. If a characteristic is correlated with student outcomes, we can use that characteristic as a proxy for teacher effectiveness. This can be useful for district or school leaders, who can then make student assignments by the known characteristics of effective teachers.

Teachers improve rapidly during their first few years in the classroom. In Massachusetts, the typical third-year teacher produces one additional month of learning per year as compared to the typical novice teacher.

Empirical evidence from many sources consistently finds that teachers improve dramatically during the first few years in the classroom. Teaching is challenging work, and it takes time for novice teachers to master the many dimensions that make them become effective teachers.

Research done in other states shows that on average, teachers in their fourth year in the classroom impart about 4 to 12 weeks more learning per year than novice teachers (Atteberry et al., 2015). Additional years of teaching experience provide less additional benefit to student achievement. While teachers do generally continue to improve over the course of their careers, recent evidence suggests that the first five years explain 60% to 70% of a teacher's lifetime gains from additional experience (Papay & Kraft, 2015; Wiswall, 2013).

Massachusetts evidence also shows that teachers improve quickly during their first few years in the classroom, as shown in Figure 3.[6] The largest improvements occur in the first three years of teaching. By their third year in the classroom, teachers produce achievement gains equivalent to about four additional weeks of learning than first-year teachers, and by their fifth year in the classroom, they produce the equivalent of five to six weeks of additional learning per year.

Figure 3. Teachers Improve Rapidly Early in Their Careers

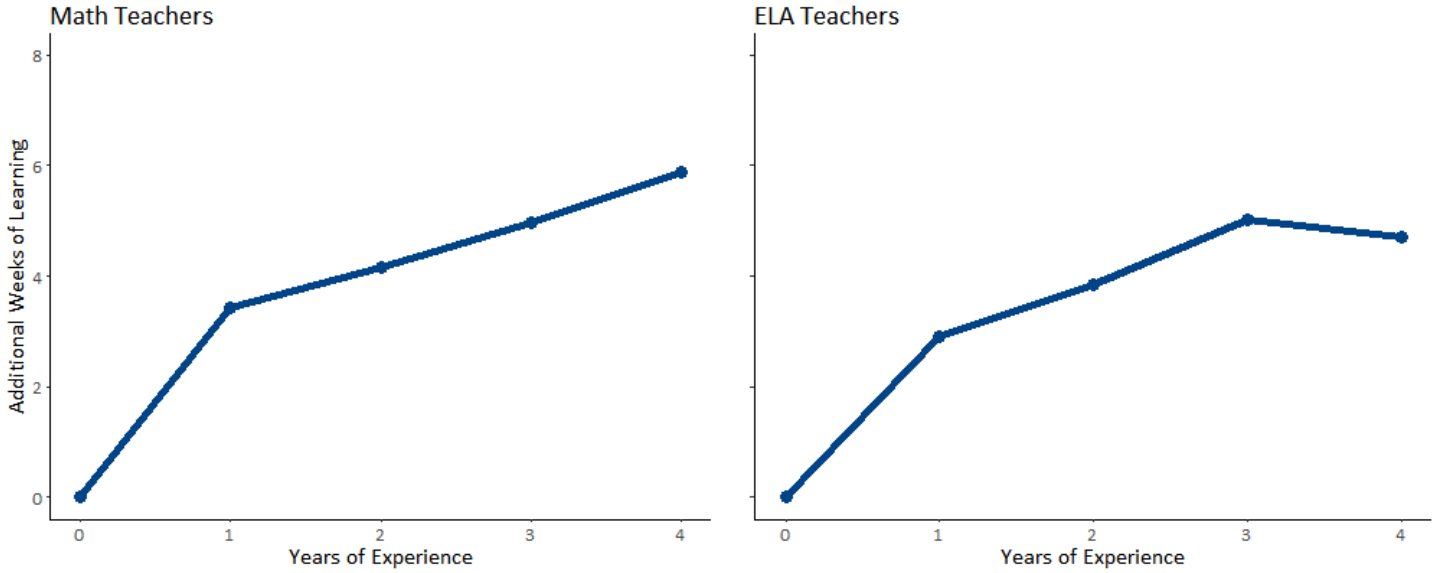

Licensure and teaching in-field have mixed evidence of impact on student outcomes.

State and federal regulations put a premium on licensure and field-specific qualifications, but evidence of their impact on teacher effectiveness is not strong or consistent. A number of randomized controlled trials have found minimal differences in student impact between teachers who enter the classroom through traditional and alternative licensure routes (Constantine et al., 2009; Glazerman et al., 2006). Some studies have found that teachers with subject-specific certifications are more effective than those who lack it, but some of the difference may be attributable to the fact that those that lack subject-specific certifications often also lack general pre-service preparation (Darling-Hammond et al., 2001; Goldhaber & Brewer, 2000). Studies have also found that teachers who improve student outcomes in

one subject tend to improve outcomes when teaching other subjects. However, similar to the benefit of experience in general, teachers do benefit from gaining additional teaching experience in the same subject regardless of subject matter license (Cook & Mansfield, 2016; Goldhaber et al., 2013).

Researchers have also generally failed to find a relationship between teachers with advanced degrees and student learning outcomes, and additional subject-specific content courses or degrees do not generally predict student learning (Aaronson et al., 2007; Harris & Sass, 2011; Monk, 1994).

Researchers have found a more consistent connection between student achievement and teachers' performance on the tests required for licensure. Teachers who receive higher scores on tests of pedagogy or content knowledge are more successful at improving student outcomes (Clotfelter, 2007, 2010; Goldhaber, 2007; Hill et al., 2005). Researchers have documented a connection between teacher scores and student outcomes for other portfolio-based assessments as well, such as the edTPA (Goldhaber et al., 2017).

Teachers with poor attendance records generate less improvement in student achievement.

Not surprisingly, students perform better on standardized tests when their teachers have better attendance records. Conservative estimates suggest that a full year of short-term substitutes would reduce student learning by about 3.5 months of learning (Clotfelter et al., 2009; Miller et al., 2008; Herrmann & Rockoff, 2012). These studies compare individual teachers' performance in years in which they have different attendance records, so the relationships are not driven by lower-performing teachers having worse attendance.

Substitutes may be less effective because they have less experience, or receive minimal lessons plans, or simply are less familiar with individual students and less capable of tailoring their instruction to students' particular needs. These issues are lessened with long-term substitutes, who have about two-thirds of the negative impact of a similar period of short-term substitutes.

Are there gaps in access to effective teachers in Massachusetts? If so, how consequential are those gaps likely to be for disadvantaged students?

We now turn to student assignment to teachers, to look for evidence of whether all students—and particularly those who are disadvantaged—are assigned equitably to the most effective educators. We define these discrepancies in the allocation of teacher assignments across different student types as *teacher equity gaps*. In this brief, we focus on equity gaps by students' socioeconomic status, comparing low income students (defined as those who received free or reduced price lunch) to non-low income students.[7]

In Massachusetts, the average low income student is assigned to a teacher who generates 1.7 fewer weeks of learning in mathematics and 3.6 weeks fewer in English language arts per year than the teachers assigned to non-low income students.

When researchers have looked at teacher effectiveness and student assignment in a number of states and districts, they have found equity gaps. The extent of these gaps varies by location and over time. Two recent studies illustrate the range of equity gaps found in contemporary school systems. In a national study of 26 large school districts, researchers found that students qualifying for free- or reduced-price lunches had teachers who produced, on average, about half a week less learning per year than higher-income students (Isenberg et al., 2016). On the other hand, other researchers using statewide data from North Carolina and Washington State found teacher equity gaps of about 2 to 3 weeks of learning (Goldhaber et al., 2016b).

Researchers have also considered teacher characteristics and observed that low income, minority, and low-achieving students are less likely to be assigned to experienced or fully licensed teachers. In North

Carolina, for instance, African American students are about 40 to 50 percent more likely to have a novice teacher than white students (Clotfelter et al., 2005). In New York State, low income students are about 20 percent more likely to have a teacher without any prior experience and about 30 percent more likely to have a teacher lacking certification (Lankford et al., 2002).

We explored teacher equity gaps in Massachusetts using similar methods as in the studies referenced above, using data from 2011-2015.[8] As Figure 4 shows, non-low income students are assigned teachers that are more effective at raising student achievement than the average Massachusetts teacher, and low income students are assigned teachers less effective than average, in both math and ELA. The differences in average teacher quality correspond to about 1.7 weeks of student learning in math and about 3.6 weeks in ELA. These differences are statistically significant and similar in magnitude to those found in other states.

Figure 4. Low Income Students Are Less Likely to be Assigned to Effective Teachers

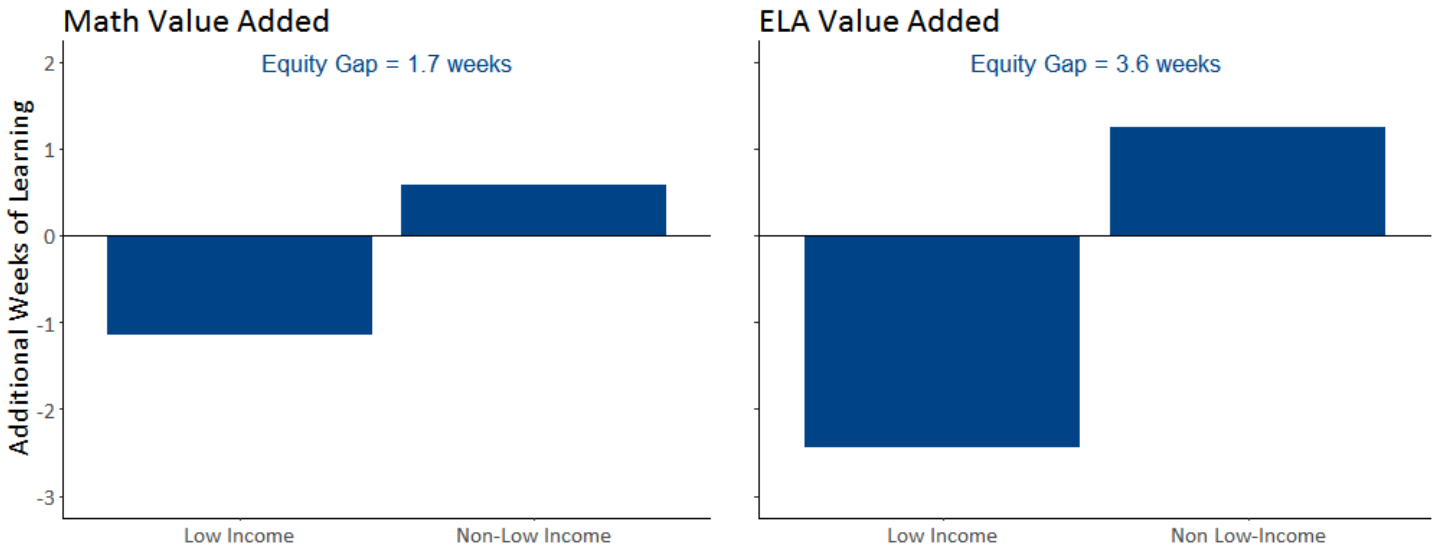

Low income students in Massachusetts are 31 percent more likely to be assigned to teachers with less than three of experience and more than twice as likely to be assigned to a teacher with low performance ratings, as compared to non-low income students.

Another way to think about equity gaps is to compare differences across student groups in access to types of teachers that are more effective on average: for example, to compare the likelihood that low income students are assigned to teachers with less than three years of experience to the likelihood for non-low income students. In this case, we express the equity gap as a *risk ratio.* The risk ratio is calculated by dividing the percentage of low income students assigned to teachers with a particular characteristic to the percentage of non-low income students assigned to those teachers. For example, if 20 percent of low income students are assigned a teacher with less than three years of experience, and 16 percent of non-low income students are, the corresponding risk ratio is 1.25 (20 divided by 16). This would indicate that low income students were about 1.25 times, or 25 percent, more likely to be assigned to teachers with less than three years of experience.

In Figure 5, we plot percentages of low income and non-low income students' teacher assignments by teacher evaluations and teacher experience. In each plot, we also show the teacher equity gap expressed as a ratio as described above, so they can be compared.

In general, low income students in Massachusetts are assigned to less effective teachers than non-low income students. The percentage of low income students assigned teachers receiving unsatisfactory or needs improvement ratings on their educator evaluations is more than twice the percentage for non-low income students. The general pattern reflects the teacher equity gaps in Figure 4: low income students tend to have less effective teachers overall.

Low income students in Massachusetts also tend to have less qualified teachers than non-low income students. Low income students are 31 percent more likely to have an inexperienced teacher—one with fewer than three years in Massachusetts public schools—than non-low income students in the state.

Figure 5. Low Income Students Tend to Be Assigned to Less Qualified Teachers

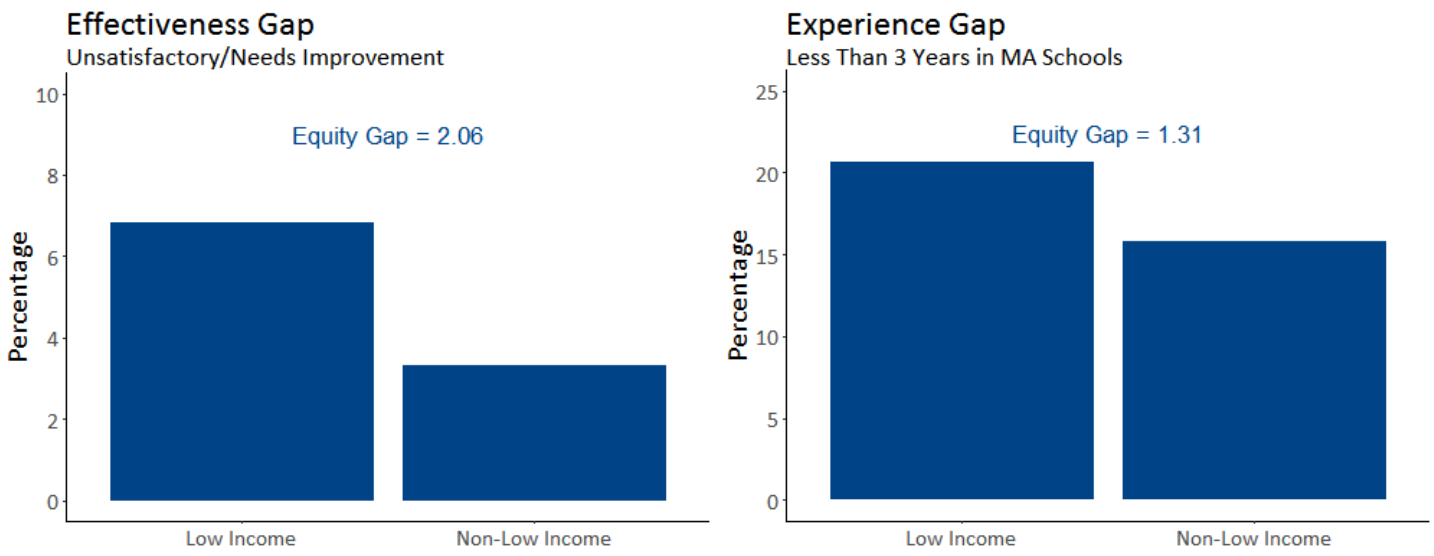

From fourth to eighth grades, inequitable access to effective teachers for Massachusetts' low income students increases achievement gaps by about 3 weeks of learning in mathematics and 6 weeks in English language arts.

Earlier, we showed that the effects of teachers accumulate over time and that teacher equity gaps become more consequential as students progress through school. We showed that having a highly effective math teacher each year would increase achievement gaps from about four weeks in fourth grade to about eight weeks by eighth grade. In ELA, the equivalent cumulative increase is about seven weeks. Based on observed equity gaps in Massachusetts, we would therefore expect that inequitable access to effective teachers between fourth and eighth grade increases achievement gaps by the equivalent of about three weeks of learning in math and about six weeks in ELA. In other words, the cumulative effect of inequitable access to effective teaching in Massachusetts among high and low income students amounts to about an additional month of instruction over those years.

Figure 6. Differences in Effectiveness across Districts Explain Most of the Equity Gaps in Massachusetts

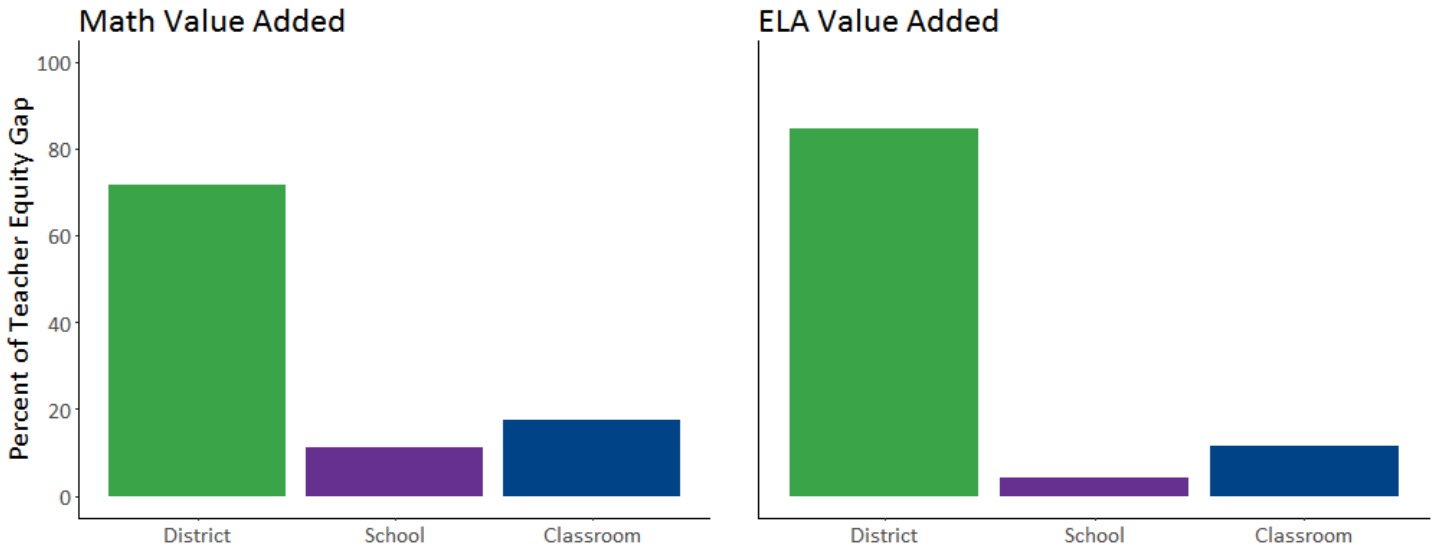

What are the sources of inequity in teacher assignments in Massachusetts?

The statewide data suggests that students have inequitable access to effective teachers in Massachusetts. But why? If we can understand more about what contributes to these inequities, we will have better insight into how to reduce them.

Three-quarters of the teacher equity gap for low income students is explained by the fact that low income students are disproportionately enrolled in districts with lower average teacher effectiveness.

One important consideration is whether teacher equity gaps are greatest across districts, schools or classrooms. If low income students disproportionately attend districts with lower average teacher effectiveness, then *districts* will explain a large share of the teacher equity gap. In that case, district-level factors, such as teacher compensation or recruitment polices, may be important policy levers. On the other hand, if they disproportionately attend *schools* with lower average teacher effectiveness, this would reflect student segregation by socioeconomic status across schools within districts. Finally, if the same is true mostly for *classrooms* this would indicate that schools have similar teaching staffs and that disparities in access to effective teaching are likely due to tracking students within schools and assigning less effective teachers to tracks with more disadvantaged students.

In Massachusetts, in both math and ELA, differences across districts are by far the largest source of inequity in teacher quality (see Figure 6). In other words, disadvantaged students disproportionately attend districts with lower average teacher quality. Differences in average teacher value added across districts explains about 72% of the total teacher equity gap in math and about 85% of the gap in ELA. Differences across schools and classrooms explain the remainder of the gap.

Figure 7. Districts in Towns and Rural Areas Have Less Effective Teachers

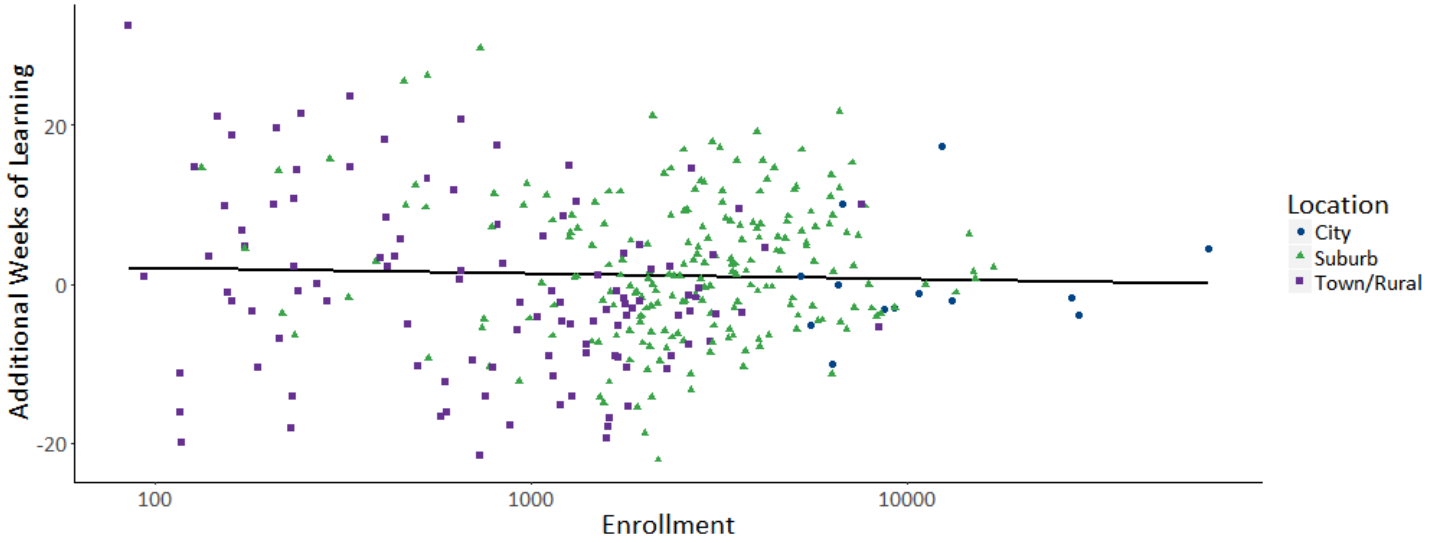

Students in rural districts tend to have less effective teachers, by nearly a month on average, than those in suburban school districts.

Because differences across districts explain so much of low income students' access to effective educators, we explore additional characteristics of districts that may be correlated with the effectiveness of their teaching staff. In Figure 7, we show the relationship between district size, location, and teacher effectiveness. We plot enrollment on the logarithmic scale (in powers of 10) in order to accommodate the large differences in district size in Massachusetts. Overall, district enrollment is not associated with average teacher value added. However, districts in towns and rural areas tend to have less effective teachers than urban or suburban districts.[9] These districts tend to have a number of obstacles to recruiting effective teachers, such as low funding, high educational costs, and a lack of highly educated workers (Monk, 2007). Accordingly, the average value added in rural districts is lower than suburban districts by about the equivalent of 3.5 weeks of learning.

Figure 8. Districts That Serve More Low Income Students on Average Have Fewer Effective Teachers

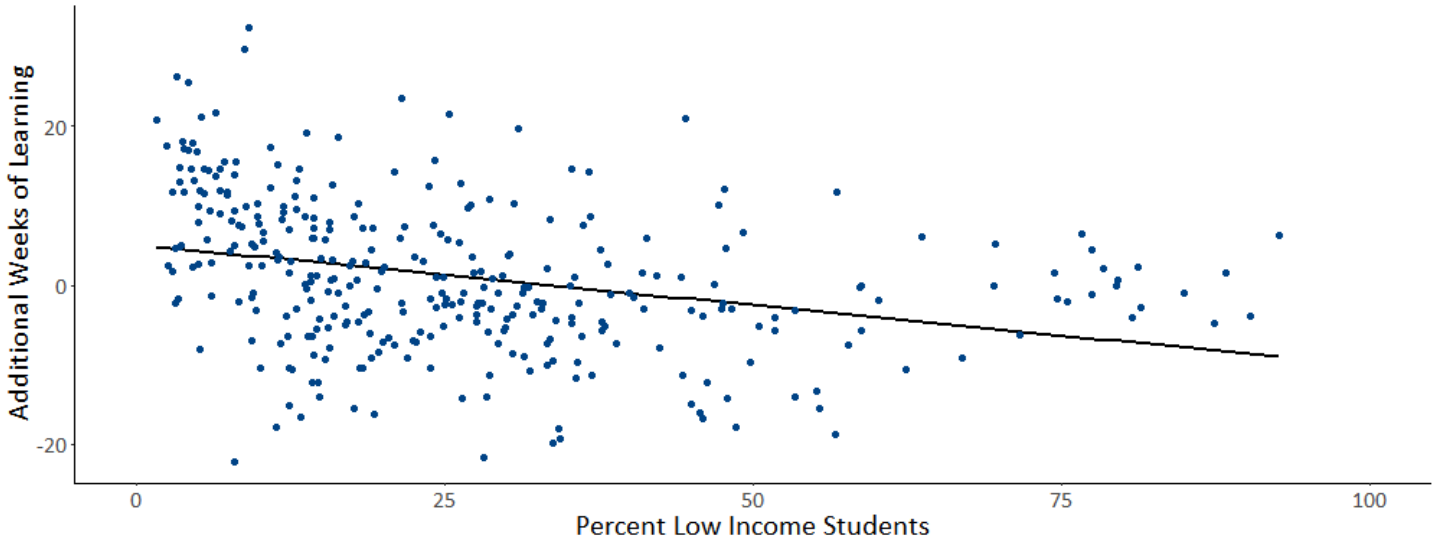

Teachers in higher income school districts tend to be more effective.

The large teacher equity gaps across districts we observed in the previous section are driven in part by the relationship between district demographics and teacher effectiveness. In Figure 8, we show this relationship explicitly by plotting the percentage of low income students in a district and the average value added of their teachers. As indicated by Figure 6, districts with more low income students have fewer effective teachers: a 10 percentage point increase in low income student enrollment is associated with a reduction in average teacher effectiveness equivalent to about 1.5 weeks of learning. These patterns suggest that district policies might influence equity gaps. In the next section, we therefore consider some potential policies that prior research has found to influence the distribution of teacher effectiveness.

What policies can schools and districts adopt that show evidence of increasing teachers' effectiveness or increasing equitable access to effective educators?

Schools and districts have a number of policy options to improve disadvantaged students' access to effective teachers. We discuss two categories of policy levers: policies that improve the recruitment and retention of effective teachers and policies that improve the effectiveness of incumbent teachers. Policies in the first category increase the likelihood that the most effective teachers work with disadvantaged students, while those in the second improve the effectiveness of all teachers in a school or district. Either category of policy may increase the likelihood that disadvantaged students have strong teachers.

Policies that affect the distribution of teachers across schools

Early hiring can help teachers transition into new positions.

Districts may be able to improve the effectiveness of their first-year teachers by ensuring that open positions are filled earlier in the year. One team of researchers found that more than a third of new teachers in California and Florida are hired after the school year has begun (Liu & Johnson, 2006). In Massachusetts, 29 percent of novice teachers were hired after September 15 in 2014 and 2015. About 1.7 percent of students—and about 2.0 percent of low income students—were taught by teachers who were hired late. Researchers have found that students gain about three additional weeks of learning in both math and reading when they are assigned teachers who were hired before the school year begins relative to teachers who were hired after the school year began (Papay & Kraft, 2016). Accelerated hiring timelines may help low achieving or low income schools hire the best available applicants (Peske & Haycock, 2006).

Some instruments can identify effective teachers even before they enter the classroom.

As researchers have connected richer measures of teacher skills to student achievement data, it has become apparent that considering measures beyond years of experience and educational attainment can improve administrators' ability to screen teachers even before they enter the classroom. As we discussed above, some traditional characteristics, such as licensure test results, predict classroom effectiveness. But it also appears that the kind of information that districts collect during the hiring process may be useful for identifying the most promising candidates. One study of an urban district's hiring process in Washington found that evaluations of candidates' written applications and references predicted classroom effectiveness (Goldhaber et al., 2016a). The researchers found that a one-standard-deviation increase in the principal's screening score (e.g., an increase from the 50th percentile to the 66th percentile of screening scores) corresponded to about five weeks of student learning. Further, higher scoring teachers were significantly less likely to leave the district after their first year of teaching. Another team of researchers, studying the hiring process in Washington, DC, found that administrators' reviews of candidates' essays, in-person interviews, and demonstration lessons predicted in-service teacher evaluations (Jacob et al., 2016).

Strikingly, they also found that the highest scoring candidates were not more likely to be hired, which suggests that a more structured hiring process could improve the overall effectiveness of the teacher workforce.

Financial incentives can improve hiring and increase the retention of effective teachers.

Districts can help schools serving disadvantaged students recruit and retain effective teachers by offering financial incentives to effective teachers to work in their schools. Some states and districts pay bonuses to all teachers working in hard-to-staff schools. Others specifically focus on teachers with high performance evaluations or those with certain teaching credentials. The latter policies may be more effective at raising the average level of teacher effectiveness in hard-to-staff schools. One national randomized experiment found that offering a one-time bonus of $20,000 to highly effective teachers willing to work in low-achieving schools increased the effectiveness of new hires and improved student achievement in the short run, but ultimately did not improve teacher retention (Glazerman et al., 2013).[10] Studies of smaller bonuses offered continuously while teachers work in hard-to-staff schools have found that they improve the hiring and reduce the turnover of effective teachers (Clotfelter et al., 2008; Springer et al., 2016)

Policies that improve the effectiveness of teachers

Districts might additionally consider policies that improve the practice of existing teachers, especially those working with traditionally disadvantaged student groups. Although identifying effective methods of professional development is difficult (Yoon et al., 2007), some policies have proven successful at raising student achievement.

Feedback and personalized mentoring can enhance the effectiveness of professional development.

Several studies show that providing more specific and more frequent feedback to teachers may help them improve their skills. Two studies of the introduction of new teacher evaluation systems in Chicago and Cincinnati suggest that teachers who participated in more structured evaluations that used classroom observation to provide specific feedback on strengths and weaknesses improved by about the equivalent of nine weeks of additional student learning (Steinberg & Sartain, 2015; Taylor & Tyler, 2012). Similarly, a recent experiment in Tennessee matched struggling teachers to mentors who were evaluated as high performing in the specific areas in which the mentees were relatively weak. At the end of the school year, students taught by the mentees scored 0.12 standard deviations higher in schools randomly assigned to the mentoring program, a difference of about ten weeks of learning (Papay et al., 2016).

Maintaining stable teaching assignments helps teachers specialize their instruction.

Just as teaching experience matters generally, teachers improve their practice when they teach the same grades and subjects in multiple years (Cook & Mansfield, 2016; Ost, 2014). This finding should not be surprising: taking on a new teaching assignment requires becoming familiar with a new curriculum and developing new instructional content. For instance, one study estimated that an elementary teacher who remained in the same grade for her first three years in the classroom would outperform a teacher who switched grades by about two to three weeks of learning in both math and reading (Ost, 2014). Teachers who stay in their same assignments are also less likely to leave the profession than those who frequently change grades (Ost & Schiman, 2015). Furthermore, recent research suggests that disadvantaged students are more likely to be assigned teachers with less experience in their specific assignment (Atteberry et al., 2017). Thus, reducing the within-school turnover of teachers may disproportionately benefit disadvantaged students.

Preparing new teachers through high quality student teaching experiences may improve the applicant pool.

One recent study found that about 40 percent of novice teachers in Washington State found teaching positions

in the school or district in which they completed their student internship (Krieg et al., 2016). This is consistent with a larger literature on teachers' job search that suggests that prospective teachers have a strong preference for taking positions near where they attended school. Teachers are twice as likely to take positions within five miles of their hometowns as positions just 20 miles away (Boyd et al., 2005). Additional research, although correlational, has begun to identify characteristics of student internship sites that appear to be beneficial for student teachers. For instance, teachers who student-taught in schools that had low staff turnover or that were better at increasing student achievement became more effective teachers and were less likely to leave the profession (Ronfeldt 2012, 2015). District investments in the student teaching process may therefore result in more effective novice teachers down the road.

What additional resources are available from the Department of Elementary and Secondary Education about access to effective educators?

The evidence in this brief demonstrates that traditionally underserved students in Massachusetts are less likely to have the teachers that are most effective at fostering student learning. In particular, non-low income students are assigned teachers that produce student learning gains that correspond to about two to three weeks of additional learning than the teachers typically assigned low income students. Over time, these learning gains can contribute meaningfully to achievement gaps between high and low income students.

School and district policy can improve students' access to effective educators. As a first step, districts and policymakers may wish to evaluate their own data on equitable access to effective teachers to learn where these gaps arise. They may also want to make use of the resources the Department has developed to help close equity gaps statewide.

Where can I learn more about access to excellent teachers in my district?

The Massachusetts Student Learning Experience (SLE) reports, found in the Edwin Analytics data platform, provide more information on students' access to effective teachers in Massachusetts. The SLE reports include several additional teacher characteristics as well as other student categories beyond what is described in this brief. They also contain information on teacher equity gaps at the district and school level. Using these reports, district and school leaders can examine questions like:

- Do gaps in equitable access exist in their district?
- How large are the gaps?
- How do they compare across schools?
- How do they compare across groups of students (e.g., economically disadvantaged, special education, English language learners)?
- Are the student groups with teacher equity gaps the same as those that also show gaps in student outcomes?

The SLE reports are available through Edwin Analytics, the Department's secured online data platform. Contact your district's directory administrator for access to Edwin and the SLE reports.

What resources does ESE offer to help close teacher equity gaps?

1. The Massachusetts Plan for Equitable Access to Excellent Educators describes statewide initiatives to improve historically disadvantaged student groups' access to excellent teachers, along with analysis of statewide data about the equity gap. The report shows:

- Schools in the top quartile of populations of economically disadvantaged students and students of color (i.e., high poverty and high minority schools) employ novice teachers at more than twice the rate of schools in the lowest quartile.
- The turnover rate for teachers rated exemplary or proficient in high poverty schools is twice as high as in low poverty schools.
- High poverty schools are 85% more likely than more affluent schools to have a first-year school leader.

The report also outlines the state's strategy to eliminate equity gaps, including: improving educator preparation; supporting effective implementation of the Educator Evaluation Framework; focusing on the student learning experience; promoting more inclusive practice and placements for working with students with IEPs, and using professional learning networks to pilot new approaches to equity.

2. The Educator Effectiveness Guidebook for Inclusive Practice provides tools and rubrics for a comprehensive approach to strengthening inclusive practices in classrooms, including self-assessment, goal-setting, conducting observations, analyzing artifacts and collecting feedback.
3. The Department's Massachusetts Equity Playbook II, a 4-page brochure outlining issues, data and suggestions about equitable access to effective educators, provides information and can serve as a facilitation guide for leadership discussions. A video that accompanied the first version provides an introduction to using the Playbook.
4. When the Department published the first Equity Plan in 2015, staff set up the Educational Equity Professional Learning Network. With districts in the network, they developed programs, toolkits and strategies that may be useful to other districts, including:
    - Partnerships with Educator Preparation Programs
    - Calibration and Inclusion Toolkit
    - Strategies for Cultural Proficiency

For more information, please visit:
http://www.doe.mass.edu/research/
or contact research@doe.mass.edu.

MASSACHUSETTS DEPARTMENT OF ELEMENTARY AND SECONDARY EDUCATION

References

Aaronson, D., Barrow, L., & Sander, W. (2007). Teachers and student achievement in the Chicago Public High Schools. *Journal of Labor Economics*, *25*(1), 95–135.

Atteberry, A., Loeb, S., & Wyckoff, J. (2015). Do first impressions matter? Predicting early career teacher effectiveness. *AERA Open*, *1*(4), 2332858415607834.

Atteberry, A., Loeb, S., & Wyckoff, J. (2017). Teacher churning: Reassignment rates and implications for student achievement. *Educational Evaluation and Policy Analysis*, *39*(1), 3–30.

Bacher-Hicks, A., Chin, M. J., Kane, T. J., & Staiger, D. O. (2017). *An evaluation of bias in three measures of teacher quality: Value-added, classroom observations, and student surveys* (No. 23478). Cambridge, MA: National Bureau for Economic Research.

Bloom, H. S., Hill, C. J., Black, A. R., & Lipsey, M. W. (2008). Performance trajectories and performance gaps as achievement effect-size benchmarks for educational interventions. *Journal of Research on Educational Effectiveness*, *1*(4), 289–328.

Boyd, D., Lankford, H., Loeb, S., & Wyckoff, J. (2005). The draw of home: How teachers' preferences for proximity disadvantage urban schools. *Journal of Policy Analysis and Management*, *24*(1), 113–132.

Chetty, R., Friedman, J. N., & Rockoff, J. E. (2014). Measuring the impacts of teachers I: Evaluating bias in teacher value-added estimates. *The American Economic Review*, *104*(9), 2593–2632.

Chetty, R., Friedman, J. N., & Rockoff, J. E. (2014). Measuring the impacts of teachers II: Teacher value-added and student outcomes in adulthood. *The American Economic Review*, *104*(9), 2633–2679.

Clotfelter, C. T., Glennie, E., Ladd, H., & Vigdor, J. (2008). Would higher salaries keep teachers in high-poverty schools? Evidence from a policy intervention in North Carolina. *Journal of Public Economics*, *92*(5-6), 1352–1370.

Clotfelter, C. T., Ladd, H. F., & Vigdor, J. (2005). Who teaches whom? Race and the distribution of novice teachers. *Economics of Education Review*, *24*(4), 377–392.

Clotfelter, C. T., Ladd, H. F., & Vigdor, J. L. (2009). Are teacher absences worth worrying about in the United States? *Education Finance and Policy*, *4*(2), 115–149.

Clotfelter, C. T., Ladd, H., & Vigdor, J. (2007). Teacher credentials and student achievement: Longitudinal analysis with student fixed effects. *Economics of Education Review*, *26*(6), 673–682.

Clotfelter, C. T., Ladd, H., & Vigdor, J. (2010). Teacher credentials and student achievement in high school: A cross-subject analysis with student fixed effects. *The Journal of Human Resources*, *45*(3), 655–681.

Constantine, J., Player, D., Silva, T., Hallgren, K., Grider, M., & Deke, J. (2009). *An evaluation of teachers trained through different routes to certification: Final report* (No. NCEE 2009-4043). Washington, D.C.: National Center for Education Evaluation and Regional Assistance, National Center for Education Statistics, U.S. Department of Education.

Cook, J. B., & Mansfield, R. K. (2016). Task-specific experience and task-specific talent: Decomposing the productivity of high school teachers. *Journal of Public Economics*, *140*, 51–72.

Darling-Hammond, L., Berry, B., & Thoreson, A. (2001). Does teacher certification matter? Evaluating the evidence. *Educational Evaluation and Policy Analysis*, *23*(1), 57–77.

Garrett, R., & Steinberg, M. P. (2015). Examining teacher effectiveness using classroom observation scores: Evidence from the randomization of teachers to students. *Educational Evaluation and Policy Analysis*, *37*(2), 224–242.

Gershenson, S. (2016). Linking teacher quality, student attendance, and student achievement. *Education Finance and Policy*, *11*(2), 125–149.

Gershenson, S., Holt, S. B., & Papageorge, N. W. (2016). Who believes in me? The effect of student–teacher demographic match on teacher expectations. *Economics of Education Review*, *52*, 209–224.

Glazerman, S., Mayer, D., & Decker, P. (2006). Alternative routes to teaching: The impacts of Teach for America on student achievement and other outcomes. *Journal of Policy Analysis and Management*, *25*(1), 75–96.

Glazerman, S., Protik, A., Teh, B.-R., Bruch, J., & Max, J. (2013). *Transfer incentives for high-performing teachers: Final results from a multisite randomized experiment* (No. 2014-4003). Washington, D.C.: National Center for Education Evaluation and Regional Assistance, Institute of Education Sciences, U.S. Department of Education.

Goldhaber, D. (2007). Everyone's doing it, but what does teacher testing tell us about teacher effectiveness? *The Journal of Human Resources*, *42*(4), 765–794.

Goldhaber, D., & Brewer, D. J. (2000). Does teacher certification matter? High school teacher certification status and student achievement. *Educational Evaluation and Policy Analysis*, *22*(2), 129–145.

Goldhaber, D., Cowan, J., & Theobald, R. (2017). Evaluating prospective teachers: Testing the predictive validity of the edTPA. *Journal of Teacher Education*, *68*(4), 377–393.

Goldhaber, D., Cowan, J., & Walch, J. (2013). Is a good elementary teacher always good? Assessing teacher performance estimates across subjects. *Economics of Education Review*, *36*, 216–228.

Goldhaber, D., Grout, C., & Huntington-Klein, N. (2016a). Screen twice, cut once: Assessing the predictive validity of applicant selection tools. *Education Finance and Policy*, *12*(2), 197–223.

Goldhaber, D., Quince, V., & Theobald, R. (2016b). *Reconciling different estimates of teacher quality based on value added* (No. Policy Brief 14). Washington, DC: National Center for Analysis of Longitudinal Data in Education Research.

Grossman, P., Loeb, S., Cohen, J., & Wyckoff, J. (2013). Measure for measure: The relationship between measures of instructional practice in middle school English language arts and teachers' value-added scores. *American Journal of Education*, *119*(3), 445–470.

Harris, D. N., & Sass, T. R. (2011). Teacher training, teacher quality and student achievement. *Journal of Public Economics*, *95*(7-8), 798–812.

Harris, D. N., & Sass, T. R. (2014). Skills, productivity and the evaluation of teacher performance. *Economics of Education Review*, *40*, 183–204.

Herrmann, M. A., & Rockoff, J. E. (2012). Worker absence and productivity: Evidence from teaching. *Journal of Labor Economics*, *30*(4), 749–782.

Hill, H. C., Rowan, B., & Ball, D. L. (2005). Effects of teachers' mathematical knowledge for teaching on student achievement. *American Educational Research Journal*, *42*(2), 371–406.

Isenberg, E., Max, J., Gleason, P., Johnson, M., Deutsch, J., & Hansen, M. (2016). *Do low income students have equal access to effective teachers? Evidence from 26 districts* (No. NCEE 2017-4007). Washington, DC: National Center for Education Evaluation and Regional Assistance, Institute of Education Sciences, U.S. Department of Education.

Jackson, C. K. (2016). *What do test scores miss? The importance of teacher effects on non-test score outcomes* (No. 22226). Cambridge, MA: National Bureau of Economic Research.

Jacob, B., Rockoff, J. E., Taylor, E. S., Lindy, B., & Rosen, R. (2016). *Teacher applicant hiring and teacher performance: Evidence from DC Public Schools* (No. 22054). Cambridge, MA: National Bureau of Economic Research.

Kane, T. J., Taylor, E. S., Tyler, J. H., & Wooten, A. L. (2011). Identifying effective classroom practices using student achievement data. *The Journal of Human Resources*, *46*(3), 587–613.

Koedel, C., Mihaly, K., & Rockoff, J. E. (2015). Value-added modeling: A review. *Economics of Education Review*, *47*, 180–195.

Krieg, J. M., Theobald, R., & Goldhaber, D. (2016). A foot in the door: Exploring the role of student teaching assignments in teachers' initial job placements. *Educational Evaluation and Policy Analysis*, *38*(2), 364–388.

Lankford, H., Loeb, S., & Wyckoff, J. (2002). Teacher sorting and the plight of urban schools: A descriptive analysis. *Educational Evaluation and Policy Analysis*, *24*(1), 37–62.

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

Liu, E., Johnson, S. M., & Peske, H. G. (2004). New teachers and the Massachusetts Signing Bonus: The limits of inducements. *Educational Evaluation and Policy Analysis*, *26*(3), 217–236.

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

Miller, R. T., Murnane, R. J., & Willett, J. B. (2008). Do teacher absences impact student achievement? Longitudinal evidence from one urban school district. *Educational Evaluation and Policy Analysis*, *30*(2), 181–200.

Monk, D. H. (1994). Subject area preparation of secondary mathematics and science teachers and student achievement. *Economics of Education Review*, *13*(2), 125–145.

Monk, D. H. (2007). Recruiting and retaining high-quality teachers in rural areas. *The Future of Children*, *17*(1), 155–174.

Ost, B. (2014). How do teachers improve? The relative importance of specific and general human capital. *American Economic Journal. Applied Economics*, *6*(2), 127–151.

Ost, B., & Schiman, J. C. (2015). Grade-specific experience, grade reassignments, and teacher turnover. *Economics of Education Review*, *46*, 112–126.

Peske, H. G., & Haycock, K. (2006). *Teaching inequality: How poor and minority students are shortchanged on teacher quality*. Washington, DC: Education Trust.

Papay, J. P., & Kraft, M. A. (2015). Productivity returns to experience in the teacher labor market: Methodological challenges and new evidence on long-term career improvement. *Journal of Public Economics*, *130*, 105–119.

Papay, J. P., & Kraft, M. A. (2016). The productivity costs of inefficient hiring practices: Evidence from late teacher hiring. Unpublished manuscript.

Papay, J. P., Taylor, E. S., Tyler, J. H., & Laski, M. (2016). *Learning job skills from colleagues at work: Evidence from a field experiment using teacher performance data* (No. 21986). Cambridge, MA: National Bureau of Economic Research.

Ronfeldt, M. (2012). Where should student teachers learn to teach? Effects of field placement school characteristics on teacher retention and effectiveness. *Educational Evaluation and Policy Analysis*, *34*(1), 3–26.

Ronfeldt, M. (2015). Field placement schools and instructional effectiveness. *Journal of Teacher Education*, *66*(4), 304–320.

Springer, M. G., Swain, W. A., & Rodriguez, L. A. (2016). Effective teacher retention bonuses: Evidence from Tennessee. *Educational Evaluation and Policy Analysis*, *38*(2), 199–221.

Steinberg, M. P., & Garrett, R. (2016). Classroom composition and measured teacher performance: What do teacher observation scores really measure? *Educational Evaluation and Policy Analysis*, *38*(2), 293–317.

Steinberg, M. P., & Sartain, L. (2015). Does teacher evaluation improve school performance? Experimental evidence from Chicago's Excellence in Teaching Project. *Education Finance and Policy*, *10*(4), 535–572.

Taylor, E. S., & Tyler, J. H. (2012). The effect of evaluation on teacher performance. *The American Economic Review*, *102*(7), 3628–3651.

Whitehurst, G. J., Chingos, M. M., & Lindquist, K. M. (2014). *Evaluating teachers with classroom observations*. Washington, D.C.: Brown Center on Education Policy, Brookings Institution.

Wiswall, M. (2013). The dynamics of teacher quality. *Journal of Public Economics*, *100*, 61–78.

Yoon, K. S., Duncan, T., Lee, S. W.-Y., Scarloss, B., Shapley, & L, K. (2007). *Reviewing the evidence on how teacher professional development affects student achievement* (No. REL 2007-033). Washington, D.C.: U.S. Department of Education, Institute of Education Sciences, National Center for Education Evaluation and Regional Assistance, Regional Educational Laboratory Southwest.

---

[1] In this report, we convert performance on standardized tests to measures of weeks of learning using the benchmarks reported by Bloom et al. (2008) for grades 4 to 8. Over these grades, four weeks of learning in corresponds to about 0.047 standard deviations on math tests and about 0.035 standard deviations on ELA tests.

[2] These estimates are based on models that estimate teacher value added over the 2011-2015 school years. For more details on how we estimated teacher value added, see Chetty et al. (2014a). We find that a top 20 percent teacher raises student achievement by about 0.23 standard deviations in math and about 0.20 standard deviations in ELA.

[3] We estimate these effects using the figures reported in Chetty et al. (2014b). They estimate the present value of lifetime earnings assuming a 3 percent discount rate to age 12 and a class size of 28.2 students.

[4] Some specific implementations of VAMs go by different names. For instance, Massachusetts uses a specific type of model called *student growth percentiles* (SGPs). All value-added models use statistical models to estimate the effect of the classroom teacher on student improvement on state tests. Although the differences between traditional VAMs and SGPs are important in some contexts,

they tend to provide very similar assessments of the effectiveness of individual teachers. We refer interested readers to Koedel et al. (2015) for more information about the differences between these two measures.

[5] We estimate value-added models using specifications that are standard in the research literature. In particular, we control for a cubic polynomial in lagged math and ELA achievement, student demographics and program participation, the means of these variables at the classroom and school level, and year and grade fixed effects.

[6] We estimate the effects of teacher experience using the same general model discussed in Note (3) above and used elsewhere in this brief. We make a few additional adjustments to estimate the effects of teacher experience: we include controls for each of the first five years of teacher experience, an indicator for six or more years of experience, and teacher fixed effects.

[7] In 2015, Massachusetts began classifying *economically disadvantaged* students by matching enrollment records to administrative data on participation in the Supplemental Nutrition Assistance Program, Transitional Aid to Families with Dependent Children, MassHealth, or foster care. We use the traditional low income measure so that socioeconomic status is defined consistently across years in our data.

[8] We also estimated average value added for students who are economically disadvantaged under the new state measure. This information is only available in 2015 and the results are similar to these using the older low income measure.

[9] We use the geographic classifications of school districts reported by the National Center for Education Statistics. Towns are incorporated areas in small urban areas of less than 50,000 people. Rural districts are in areas recognized as rural by the Census Bureau. About one-third of Massachusetts districts are in towns or rural areas.

[10] See also Lui et al. (2004), who studied a Massachusetts program awarding bonuses to new teachers and found that teachers did not report that the incentive influenced their decision to stay in teaching.