Promoting Equitable Access to Effective Teachers

GEORGIA
POLICY LABS

Promoting Equitable Access to Effective Teachers

Tim R. Sass and M. Cade Lawson

Metro Atlanta Policy Lab for Education

March 2024

Background and Motivation

Motivation

Teachers are the most important school-based input to a student's education, with research showing that exposure to more effective teachers during childhood provides benefits that persist well into adulthood.[1] However, the question of whether all students have equal access to effective teachers has proven difficult to answer. If disparities do exist, as some prior evidence from Florida, North Carolina and Washington suggests,[2] uneven access to effective teachers may be a driving force behind differences in academic achievement and later-life outcomes between advantaged and disadvantaged students. As such, understanding the way teacher quality is distributed has important implications for teacher recruitment and retention policies at the state and school-district level.

In this report, we use student-level achievement data to provide descriptive evidence on the distribution of teacher quality across a large school district in the metro-Atlanta area. We discuss where disparities exist and how interventions to attract and retain teachers could potentially reduce these differences. Our analysis focuses on the distribution of teacher quality both geographically and over time, comparing differences in access to effective teachers between students in high- and low-income areas and assessing how the distribution of teacher effectiveness has changed since the start of the COVID-19 pandemic.

Background

A handful of studies from other states have examined the distribution of teacher quality across sub-groups of students, leading to an emerging consensus that students from households experiencing low income and students of color frequently have less access to highly effective teachers than do their more-advantaged peers. These differences hold regardless of how teacher effectiveness is defined. For example, Black and Hispanic students are 50% more likely to be taught by a novice teacher during elementary school than their White peers.[3] Because teachers improve during their first few years on the job,[4] this disproportionate exposure to novice teachers implies unequal access to effective teaching across racial and ethnic groups.

Evidence from North Carolina affirms this finding, expanding the definition of "teacher quality" to include a teacher's National Board certification status and licensure exam performance in addition to years of experience in the classroom. White students have substantially greater access to better-credentialed teachers than their Black and Hispanic peers–with the largest discrepancies occurring in the most populous counties.[5]

Studies evaluating the distribution of teacher effectiveness based on a teacher's impact on student test scores (i.e., "value-added"), which is the metric we use in this work, have yielded mixed results.[6] Evidence from Florida and North Carolina indicates that teachers in high-poverty schools have both lower effectiveness on average and higher variation in quality than teachers in low-poverty schools.[7] An analysis using data from both North Carolina and Washington finds similar disparities in teacher quality between advantaged and disadvantaged students, with these "teacher quality gaps" persisting regardless of how either "student disadvantage" or "teacher quality" are defined.[8] The strongest evidence against the existence of teacher quality gaps comes from a national study of 26 school districts, which finds teacher quality differences between high- and low-income students and between Black, White, and Hispanic students that are very modest in magnitude compared to prior studies.[9] This finding implies that access to effective teachers may vary depending on geographic location, further motivating the need to study teacher effectiveness in a different context.

Research Questions

We address the following research questions:

1. Is teacher effectiveness distributed equally between high- and low-poverty regions in the district?
2. Did the distribution of teacher effectiveness in the district change after the onset of the COVID-19 pandemic?

Data and Methods

Data

Our analysis uses student-level administrative data from a large school district in the metro-Atlanta area between school years (SY) 2017–18 and 2021–22. These data include demographic information such as a student's race, identified disability and free or reduced-price meals (FRPM) status, English Learner status, school, and grade. Crucially, the student data also include fall and winter formative assessment scores for students in the district during the five-year period of study. We construct a student-subject-year panel, which allows us to link each student to their teachers in math and reading for a given year. After considering only student-subject-year observations which are usable in our value-added models,[10] we are left with a sample of roughly 195,000 student-subject-year observations. From these observations, we construct a total of about 4,500 teacher-subject value-added estimates for the district's math and reading teachers. We focus our analysis on students in grades K–8 because nationally-normed formative assessments were only introduced at the high school level during SY 2021–22.

Our analytical focus is a student's winter formative assessment score. We express this variable as a z-score (equal to the scale score divided by the standard deviation of scores within a subject and grade), rather than a raw scale score, to account for differences in scales across grades and subjects. The z-scores are computed using the national distribution of scores in SY 2018–19, so scores in each year are relative to national pre-pandemic achievement levels. This allows us to interpret our value-added estimates in standard deviation units or "effect sizes," facilitating comparisons across grades, subjects, and time periods. We also present some supplementary results using the national percentile ranking as the measure of student achievement.

Empirical Methods

We quantify teacher effectiveness by estimating value-added models for math and reading teachers in the district. The value-added approach uses multivariate regression models that seek to isolate a teacher's impact on student test score growth over time.[11] In our analysis, it involves predicting a student's end-of-term test score as a function of their beginning-of-term test score and other observable characteristics of the student, class, and school. The difference between the student's predicted and actual end-of-term score is then attributable to the contribution—or "value-added"—of the teacher. Value-added models have been shown to provide unbiased estimates of a teacher's ability to promote test score growth.[12]

Our preferred specification of the value-added model controls for student gender, race/ethnicity, FRPM status, identified disability status, English Learner status, and grade level. We additionally control for class size, the number of instructional days between the two exam dates, and the square of a student's initial test score. The model also includes an indicator for the school year (also known as a school year "fixed effect"), which means that comparisons are made across teachers within a year.[13] Including these observable characteristics in the model helps to ensure that the estimated teacher value-added reflects their effectiveness as an instructor rather than any possible differences associated with teaching particular students or in a certain environment. We estimate models separately for math and reading teachers so that we can analyze differences between the two subjects.

We estimate average teacher value-added during the fall semester (between fall and winter assessments) over the period SY 2017–18 through SY 2021–22. To allow for differences in teacher effectiveness during the pandemic, we also separately estimate value-added over a pre-pandemic period (SY 2017–18 through SY 2019–20) and a pandemic period (SY 2020–21 through SY 2021–22). This later period includes the fall semester in which hybrid and remote learning occurred (SY 2020–21) and the first year after the return to in-person instruction (SY 2021–22).

Because there can be random fluctuations in a student's test scores (e.g., due to illness on the exam day or simply being lucky at guessing answers on a multiple-choice exam), estimates of value-added can be noisy measures of a teacher's effectiveness. To compensate for this, we use an Empirical Bayes shrinkage method to stabilize our estimates. Empirical Bayes shrinkage involves adjusting especially noisy value-added estimates, particularly those from teachers who we observe teaching very few students.[14]

To study the distribution of teacher quality by region, we compare value-added between teachers in parts of the district with varying levels of poverty. By comparing teacher value-added between regions, which we refer to as the Higher-Poverty Region (HPR) and the Lower-Poverty Region (LPR), we can better understand how access to effective teachers varies along racial and socioeconomic lines.

Finding I: Differences in Average Teacher Quality by Region

Average teacher effectiveness is lower in the Higher-Poverty Region than in the Lower-Poverty Region. The difference is equal to about 7% of a standard deviation in math and reading scores over the course of a semester. Over a full school year, this would be roughly double the typical difference between the performance of a first-year teacher and a teacher with 3–5 years of experience.

There is a statistically significant difference in average teacher quality between the Lower- and Higher-Poverty Regions, meaning we can state with high confidence that average teacher effectiveness is higher in the Lower-Poverty Region than in the Higher-Poverty Region. The average teacher in the HPR generates 0.07 fewer standard deviations in test score growth over the course of a semester than the average teacher in the LPR. When student performance is measured by national percentile rankings, this difference in teacher effectiveness is equal to 1.54 national percentile points in math and 1.78 national percentile points in reading. For comparison, the difference in average effectiveness between teachers in the Lower and Higher-Poverty Region is equal to about twice the difference in effectiveness between a rookie teacher and a teacher with 3–5 years of experience over the course of a full school year. This gap persists across both math and reading.

Figures 1 and 2 graphically illustrate the differences in average teacher effectiveness between the two regions. The heights of the bars represent the estimated value-added. Because these are statistical estimates, they carry some degree of uncertainty. We can be 95% sure that the (unknown) true value of teacher value-added lies within the range indicated by the black vertical bars. In both math and reading and whether using national percentiles or standard deviation units, we can be 95% sure that the true difference in average teacher effectiveness between teachers in the regions is not zero.

Figure 1. Average Math Teacher Value-Added, by Region.

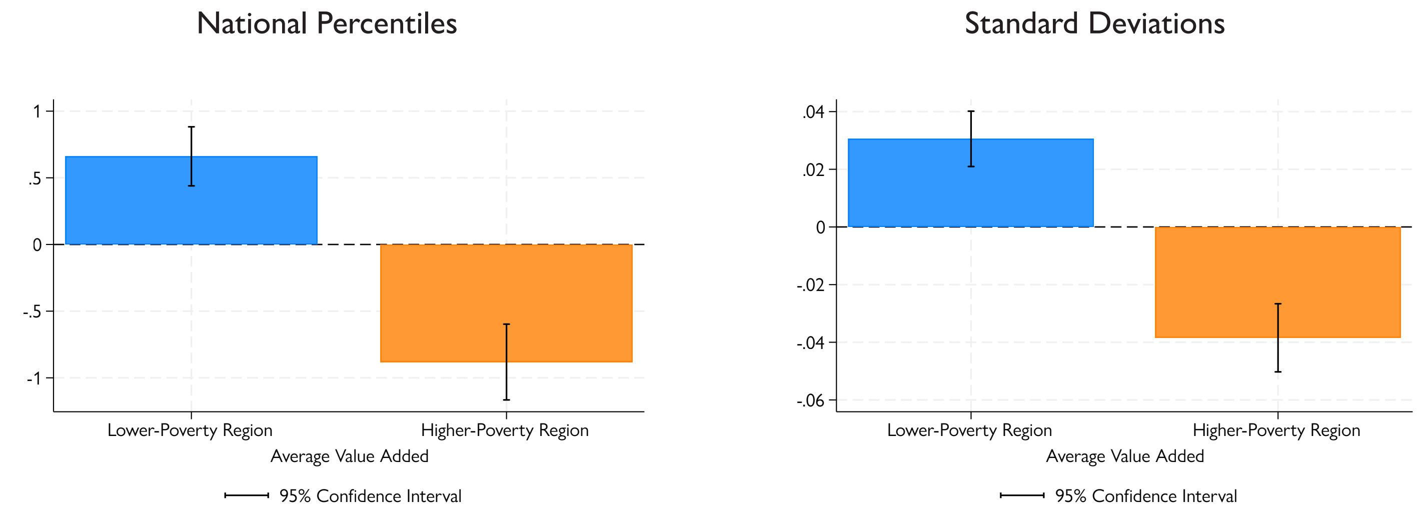

Notes. The left panel displays teacher effectiveness measured in national percentiles and the right displays effectiveness measured in standard deviations. Value-added is estimated using our full sample of data from SY 2017–18 to SY 2021–22 for students in grades K–8 in the appropriate subject. Reported estimates are adjusted using Empirical Bayes shrinkage prior to averaging.

Figure 2. Average Reading/EL Teacher Value-Added, by Region

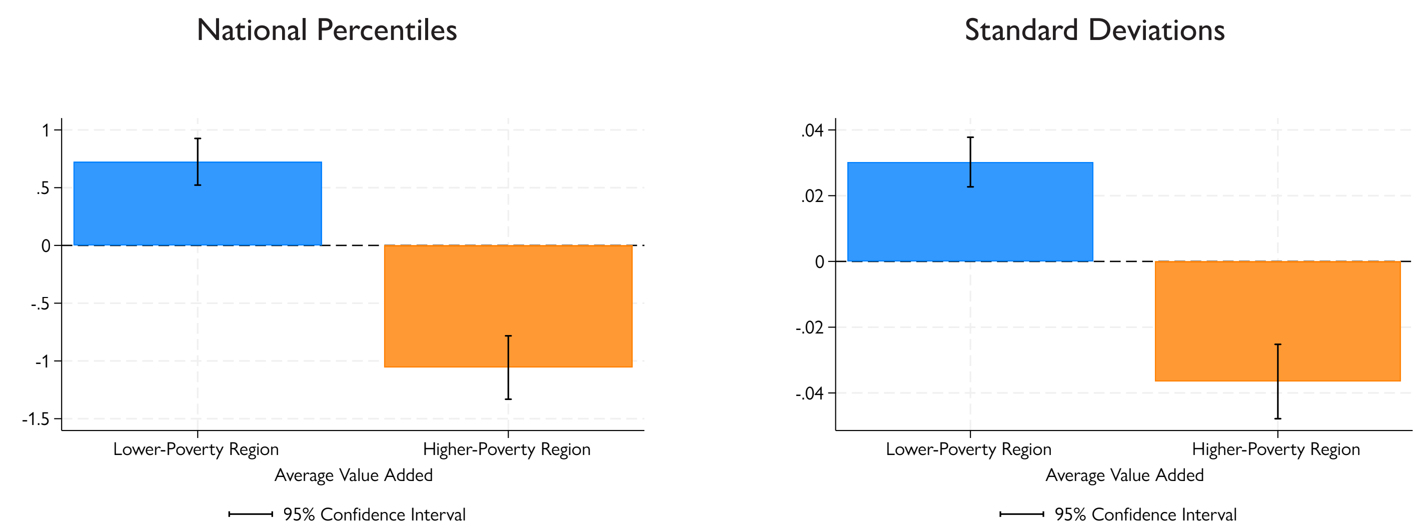

Notes. The left panel displays teacher effectiveness measured in national percentiles and the right displays effectiveness measured in standard deviations. Value-added is estimated using our full sample of data from SY 2017–18 to SY 2021–22 for students in grades K–8 in the appropriate subject. Reported estimates are adjusted using Empirical Bayes shrinkage prior to averaging.

Finding 2: Uneven Variation in Teacher Quality by Region

There is generally higher variance in teacher quality in the Higher-Poverty Region. This difference is largely driven by a much higher share of teachers in the bottom decile of teacher effectiveness working in the Higher-Poverty Region than in the Lower-Poverty Region, particularly in reading.

In Figures 3A and 3B, we present comparisons of the overall distribution of teacher quality between the HPR and LPR in math and reading, respectively. There is greater variability in teacher effectiveness in the Higher-Poverty Region, implying that teacher effectiveness is more "spread out" in the Higher-Poverty Region and relatively more centered on the average in the Lower-Poverty Region.

There is substantial overlap in teacher quality between the two regions. While average effectiveness may be lower in the HPR, there are many teachers who are similarly effective in both regions. This includes many highly effective teachers in the Higher-Poverty Region. However, the large vertical distance between the HPR and LPR quality distributions to the left of zero implies that a greater proportion of teachers in the Higher-Poverty Region demonstrate relatively low effectiveness in boosting student test scores. We explore this further in Figures 4A and 4B.

As Figure 3 suggests, the HPR has a much higher share of teachers in the lower deciles of value-added than LPR. This result is most pronounced at the first decile in reading: 16.8% of reading teachers in the Higher-Poverty Region are in the bottom 10% of teacher effectiveness district-wide, compared to just 5.8% in the Lower-Poverty Region. There is also a greater share of top-decile teachers in the Lower-Poverty Region, but the discrepancies are not as large in magnitude as in the first decile. The nearly identical share of teachers at the fourth through sixth deciles in math and at the fourth and fifth deciles in reading between the Higher- and Lower-Poverty Regions—the middle of the distribution—reinforces that fact that middling-quality teachers are distributed relatively equally across the two regions. The geographic differences in teacher effectiveness are largely driven by concentration of the most- and least-effective teachers (as measured by value-added) in the LPR and HPR, respectively.

Figure 3A. Distribution of Teacher Effectiveness in Math, by Region

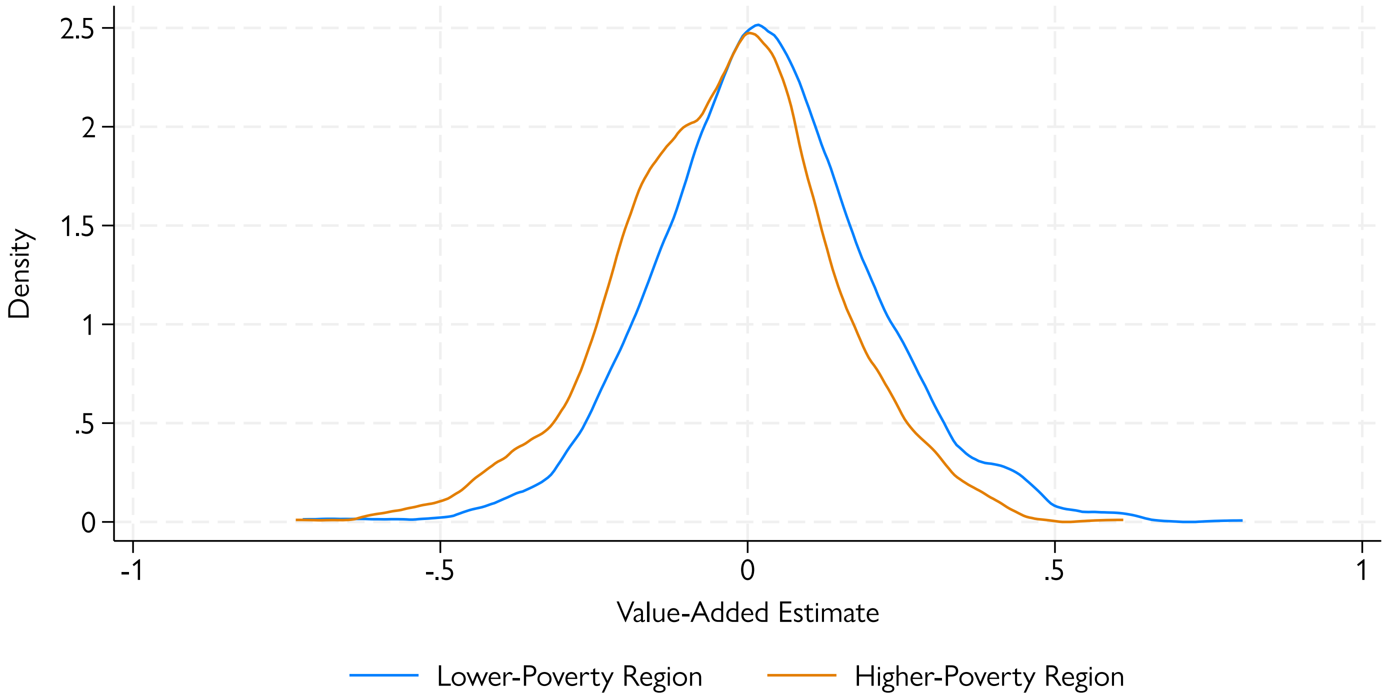

Notes. Value-added, measured in standard deviations, is estimated using our full sample of data from SY 2017–18 to SY 2021–22 for students in grades K–8 in the appropriate subject. Reported estimates are adjusted using Empirical Bayes shrinkage.

Figure 3B. Distribution of Teacher Effectiveness in Reading/ELA, by Region

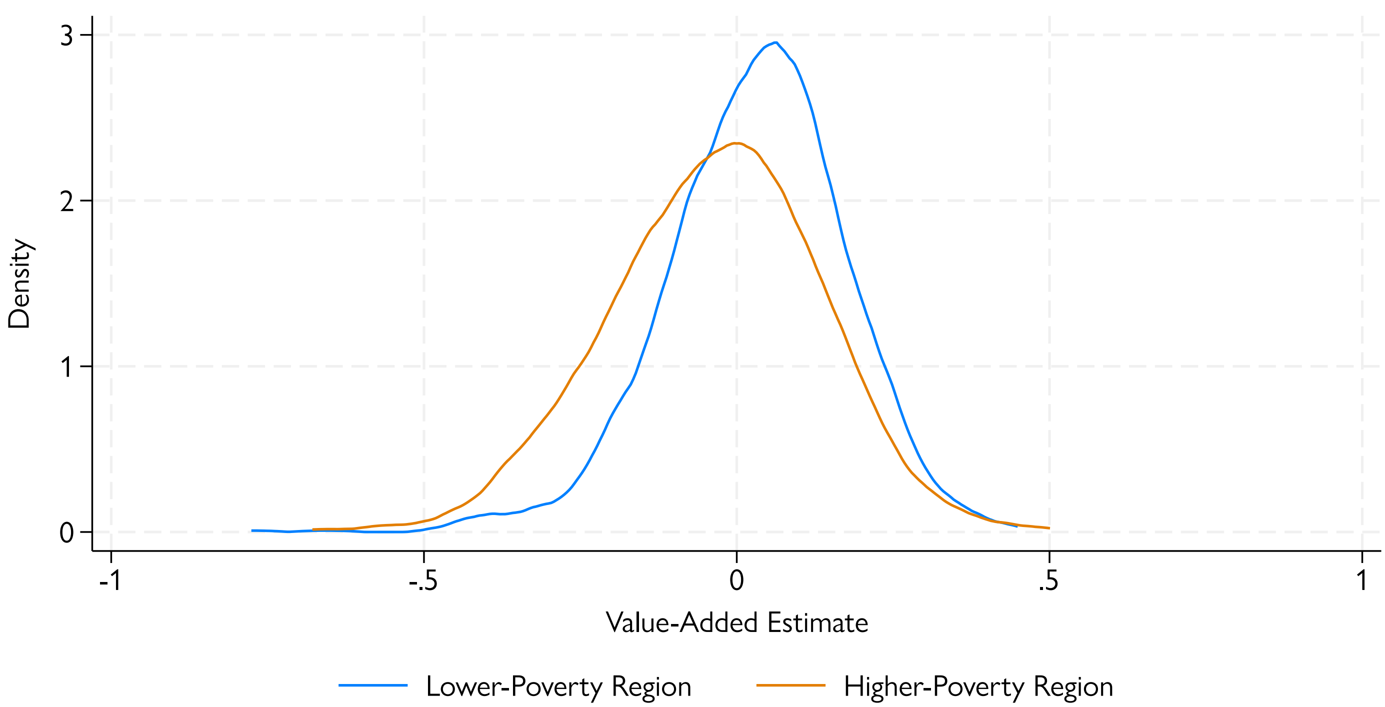

Notes. Value-added, measured in standard deviations, is estimated using our full sample of data from SY 2017–18 to SY 2021–22 for students in grades K–8 in the appropriate subject. Reported estimates are adjusted using Empirical Bayes shrinkage.

Figure 4A. Percent of Math Teachers in Each Decile of Teacher Effectiveness, by Region

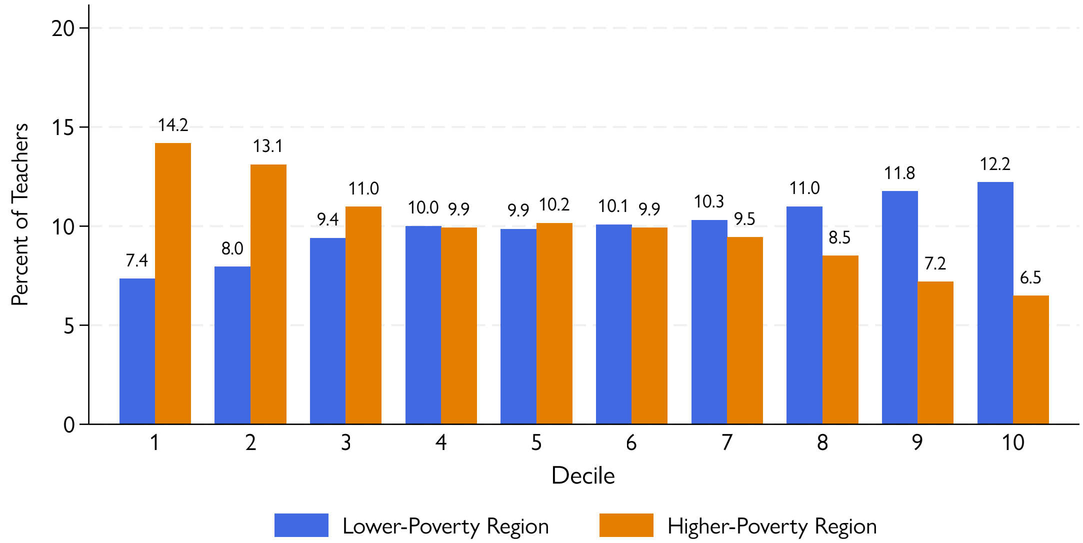

Notes. The figure is created by calculating teacher effectiveness deciles across both the Lower- and Higher-Poverty Regions, then determining the percent of all teachers in each region who fall into each decile. If the effectiveness distributions between the HPR and LPR were identical, every bar would be at 10% (by definition of decile). Value-added, measured in standard deviations, is estimated using our full sample of data from SY 2017–18 to SY 2021–22 for students in grades K–8 in the appropriate subject. Reported estimates are adjusted using Empirical Bayes shrinkage prior to averaging.

Figure 4B. Percent of Reading/ELA Teachers in Each Decile of Teacher Effectiveness, by Region

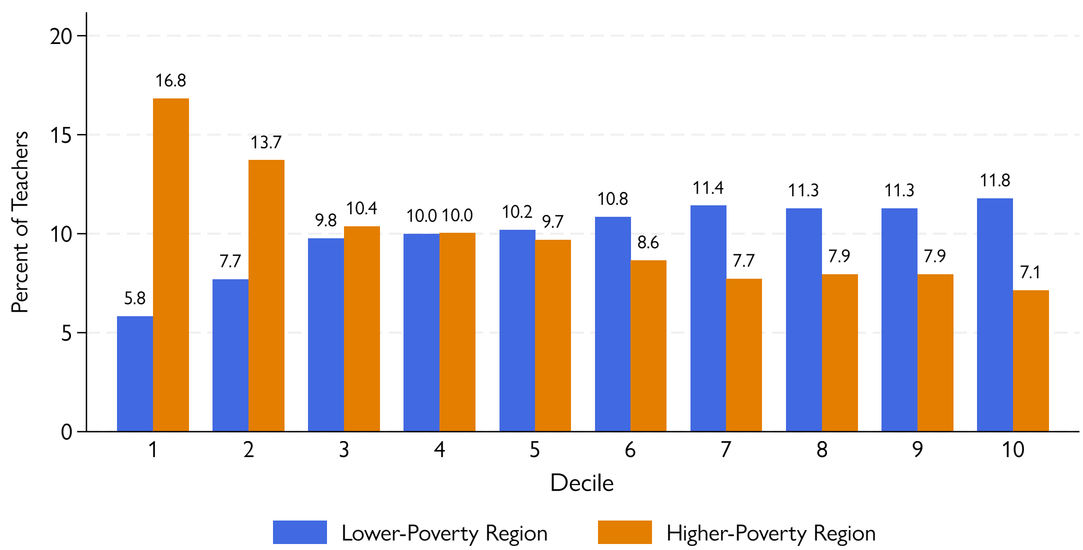

Notes. The figure is created by calculating teacher effectiveness deciles across both the Lower- and Higher-Poverty Regions, then determining the percent of all teachers in each region who fall into each decile. If the effectiveness distributions between the HPR and LPR were identical, every bar would be at 10% (by definition of decile). Value-added, measured in standard deviations, is estimated using our full sample of data from SY 2017–18 to SY 2021–22 for students in grades K–8 in the appropriate subject. Reported estimates are adjusted using Empirical Bayes shrinkage prior to averaging.

Finding 3: Differences in Teacher Effectiveness Over Time

Variation in teacher quality increased after the COVID-19 pandemic across the entire district. However, the increased variability in teacher quality was greatest among teachers in the Lower-Poverty Region.

The onset of the COVID-19 pandemic created unprecedented disruption to the U.S. public education system, but little is known about the effects of the pandemic and remote instruction on the relative performance of teachers. At the beginning of SY 2020–21, all teachers and students in the district were engaged in remote learning. Over the course of the fall semester, students and teachers transitioned back to in-person learning. Consequently, estimates of teacher effectiveness in SY 2020–21 (which are based on student achievement growth between the fall and winter assessments) primarily reflect teacher performance in remote instruction, while estimates of teacher effectiveness in SY 2021–22 measure teacher effectiveness in an in-person learning environment.

For the two-year pandemic-era period (SY 2020–21 and SY 2021–22), we find greater variation in teacher effectiveness across the district than in the pre-pandemic period. This is driven by a redistribution of teachers from the middle of the effectiveness distribution to the tails. As illustrated in Figure 5, the proportion of teachers with value-added scores high enough to rank in the top 10% of teachers in the pre-pandemic period rose to 12.2%. At the other end of the spectrum, the pandemic-era proportion of teachers with value-added scores—which would have placed them in the bottom 30% of teachers in the pre-pandemic era—rose to 33.3%.

A similar pattern holds in Figure 6, where we analyze the pre-pandemic and pandemic distributions of teacher effectiveness at the region-level. While variation in teacher effectiveness increased in both the Higher- and Lower-Poverty Regions, the changes were more pronounced in the Lower-Poverty Region. For example, in the pandemic-era, the proportion of teachers whose value-added scores would have placed them in the bottom 10% of teachers in the pre-pandemic era was 14.9%, and the proportion who would been among the top 10% of teachers in the pre-pandemic period was 12.7%. In contrast, for the Higher-Poverty Region, the proportion in the bottom decile of the pre-

Figure 5. Percent of Teachers in Each Pre-Pandemic Value-Added Decile, Pre-Pandemic vs. Pandemic Period

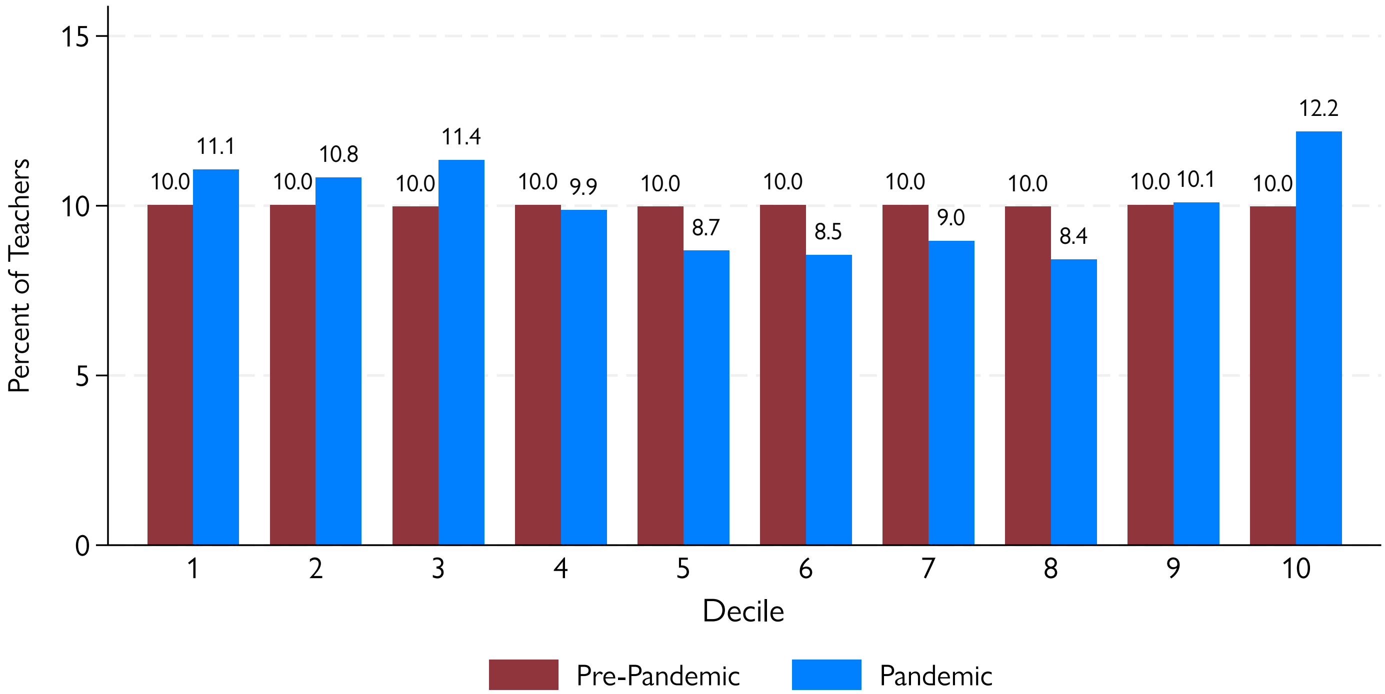

Notes. The red bars represent the proportion of teachers in each decile of the pre-pandemic distribution of teacher effectiveness. By definition, this equals 10% in each decile. The blue bars show the proportion of teachers in the pandemic period that would have been in the 10 groups defined by the pre-pandemic teacher effectiveness distribution deciles. Value-added, measured in standard deviations, is estimated using data from SY 2017–18 to SY 2019–20 for our pre-pandemic models and from SY 2020–21 to SY 2021–22 for our pandemic-era models. All models use data for students in grades K–8 in the appropriate subject.

pandemic distribution increased to only 11.9%, while the proportion in the top decile rose to only 11.4%.[15]

There are two potential explanations for the increased variability in teacher performance in the pandemic era. One possibility is that some teachers excelled in remote instruction (relative to their performance during in-person instruction), and some teachers performed relatively worse, which would increase the variability in relative performance across teachers.[16] An alternative explanation is that the pandemic led to a change in the quality of the teachers employed by the district. For example, if difficulties in hiring new teachers during the pandemic lowered the quality of new hires, that could explain the increase in relatively ineffective teachers. Similarly, if the rebounding labor market in calendar years 2021 and 2022 led teachers in the middle of the quality distribution to leave for more lucrative opportunities elsewhere, it would be consistent with the reduction in the proportion of teachers that would be ranked in the middle deciles of the pre-pandemic distribution.

Figure 6A. Percent of Teachers in Each Pre-Pandemic Value-Added Decile, Pre-Pandemic vs. Pandemic Period, Lower-Poverty Region

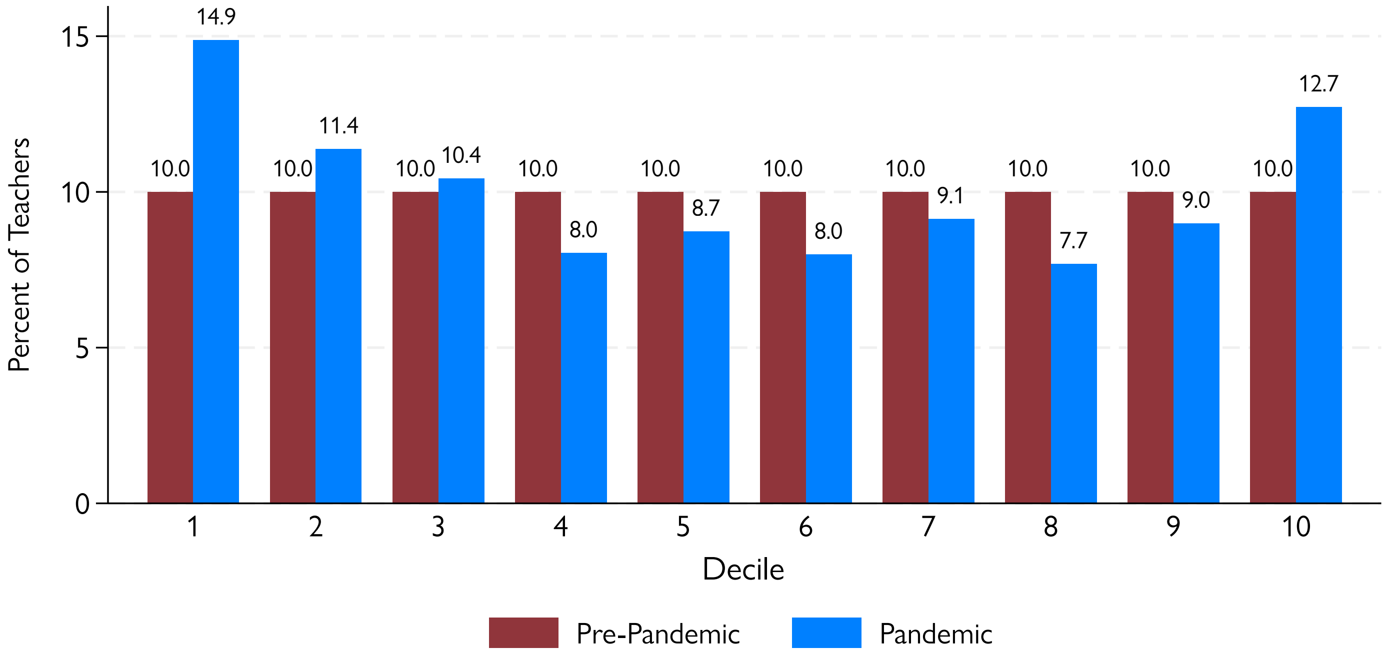

Figure 6B. Percent of Teachers in Each Pre-Pandemic Value-Added Decile, Pre-Pandemic vs Pandemic Period, Higher-Poverty Region

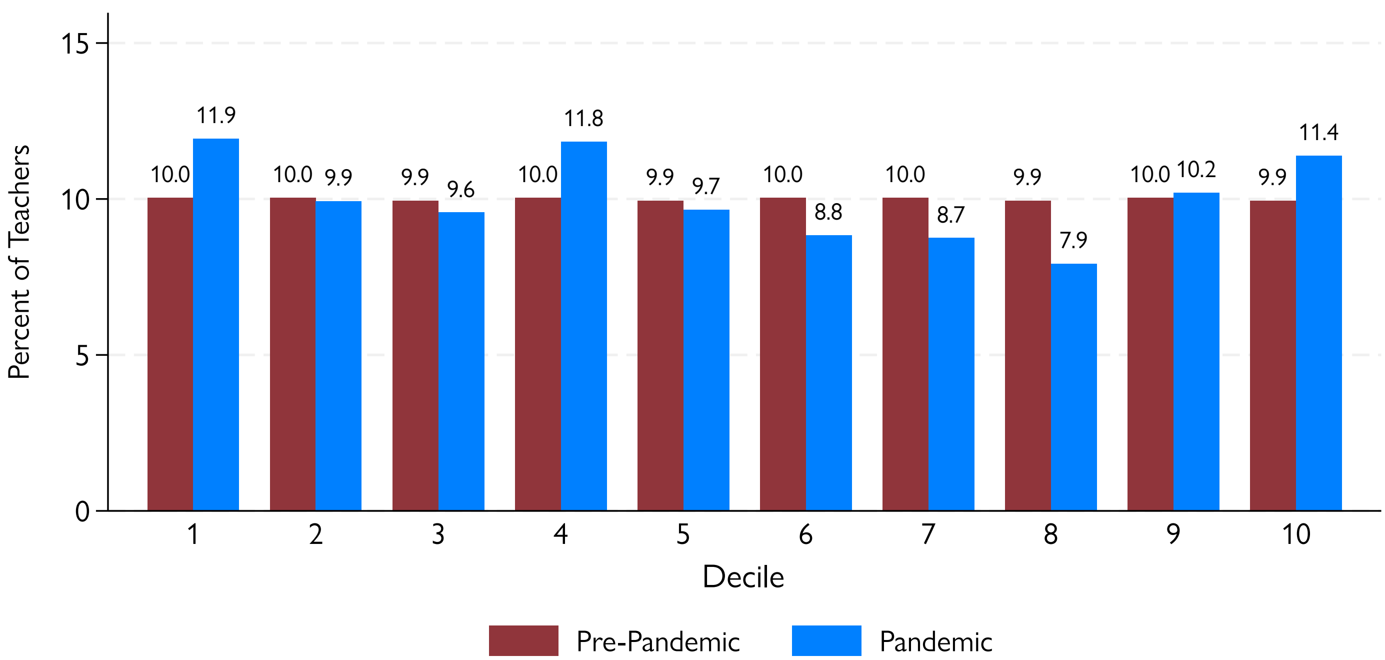

Notes. The red bars represent the proportion of teachers in each decile of the Lower-Poverty (Figure 6A) or Higher-Poverty (Figure 6B) Region's pre-pandemic distribution of teacher effectiveness. In some cases, the proportion in each decile is not exactly equal to zero due to the sample size not being exactly divisible by 10 (i.e, an additional teacher in one decile bin throws off the percent when rounded to the nearest 0.1). The blue bars show the proportion of teachers in the pandemic period that would been in the 10 groups defined by the pre-pandemic teacher effectiveness distribution deciles. Value-added, measured in standard deviations, is estimated using data from SY 2017–18 to SY 2019–20 for our pre-pandemic models and from SY 2020–21 to SY 2021–22 for our pandemic-era models. All models use data for students in grades K–8 in the appropriate subject.

Discussion

We find differences in teacher effectiveness between the Lower-Poverty and Higher-Poverty Regions in the district equaling about 7% of a standard deviation but observe substantial overlap of the teacher effectiveness distribution between the two regions. The difference in average quality is largely driven by a high concentration of teachers from the lowest decile of effectiveness, as measured by test score growth, teaching in the Higher-Poverty Region. These results for our district of study are consistent with prior research that finds students in high-poverty schools do not always have the same level of access to effective teachers as their peers in low-poverty schools.

We also find that variation across the district in teacher effectiveness increased after the pandemic began, with the increased variability in teacher effectiveness being more pronounced in the Lower-Poverty Region than in the Higher-Poverty Region. This pattern could reflect changes in relative teacher performance associated with remote instruction and/or pandemic-era changes in the composition of the teacher workforce. Future work will untangle the potential causes of this trend by analyzing the factors associated with teacher effectiveness in remote instruction. Understanding the relationship between instructional mode and teacher effectiveness could provide valuable information on how best to staff remote learning programs and which students are most likely to be successful in online learning.

There are a variety of potential policy options that could make the distribution of teacher quality more equal. However, attempting to get highly effective teachers to move from the Lower-Poverty Region to the Higher-Poverty Region would be costly and likely would not lead to significant numbers of teachers switching schools. Prior research finds that few highly effective teachers are willing to move to "hard-to-staff" schools, even when offered bonuses up to $20,000 over a two-year period.[17] A more cost-effective approach would be to focus on reducing the proportion of relatively ineffective teachers in higher-poverty schools. This could potentially be accomplished by improvements in the teacher selection process, more stringent reviews of early-career teachers in the Higher-Poverty Region, and financial incentives to new teachers who agree to begin their careers in the Higher-Poverty Region. These strategies could be coupled with retention incentives that are tied to teacher performance. There is evidence that such effectiveness-related bonuses for teaching in hard-to-staff schools can promote teacher retention.[18]

To determine the best combination of policies, future work should aim to understand the sources of differences in teacher effectiveness between higher- and lower-poverty schools. For example, comparing the candidate pools between higher- and lower-poverty schools would provide valuable information on whether policies should focus on candidate selection and early-career mentoring in higher-poverty schools (if pools are similar) or providing financial incentives to work in higher-poverty schools to attract more promising candidates (if pools are different).

Endnotes

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

2. Sass, T. R., Hannaway, J., Xu, Z., Figlio, D. N., & Feng, L. (2012). Value added of teachers in high-poverty schools and lower poverty schools. *Journal of Urban Economics*, *72*(2–3), 104–122.

Goldhaber, D., Quince, V., & Theobald, R. (2018). Has it always been this way? Tracing the evolution of teacher quality gaps in US public schools. *American Educational Research Journal*, *55*(1), 171–201.

3. James, J., & Wyckoff, J. H. (2022). *School Segregation, Teacher Sorting, and the Distribution of Teachers*. EdWorkingPaper No. 22-659. Annenberg Institute for School Reform at Brown University.

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

Ladd, H. F., & Sorensen, L. C. (2017). Returns to teacher experience: Student achievement and motivation in middle school. *Education Finance and Policy*, *12*(2), 241–279.

5. Clotfelter, C.T., Ladd, H.F., & Clifton, C.R. (2021). *White Advantage in Teacher Assignment*. (EdWorkingPaper: 21-455). Annenberg Institute at Brown University.

6. The value-added approach compares expected student test scores (based on prior scores and observed student demographics) with actual scores and attributes the difference to the effectiveness of the teacher. We provide additional details on the estimation of such value-added models in the Methods section of this report.

7. Sass et al. (2012).

8. Goldhaber et al. (2018).

9. Isenberg, E., Max, J., Gleason, P., & Deutsch, J. (2022). Do low-income students have equal access to effective teachers? *Educational Evaluation and Policy Analysis*, *44*(2), 234–256.

10. To avoid bias in our value-added estimates, we exclude observations corresponding to the following situations: students who had more than one teacher in the same subject in a single school year, students who were observed in multiple classes with the same teacher and subject

in the same year, students who joined a class after September of the fall semester, classes with fewer than 10 formative assessment scores, and classes with multiple instructors of record.

11. Further details on the specification and estimation of the value-added model are presented in the Appendix.

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

13. In addition to our preferred model, we also estimated five alternative specifications. Results from these robustness checks are presented in Appendix Table A1. All five alternative specifications yield statistically significant differences in average teacher effectiveness between the two regions.

14. The Empirical Bayes methodology is discussed further in the Appendix.

15. This result is confirmed in Appendix Figures A1 and A2, which display average teacher effectiveness before and after the pandemic in the Lower vs. Higher-Poverty Regions across reading and math.

16. We are currently working on a separate project to estimate how teacher effectiveness varies with instructional mode and the factors associated with relatively strong performance in remote instruction.

17. Protik, A., Glazerman, S., Bruch, J. & Teh, B. (2015), Staffing a Low-Performing School: Behavioral Responses to Selective Teacher Transfer Incentives. *Education Finance and Policy*, *10*(4), 573–610.

18. 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.

About the Authors

Tim R. Sass

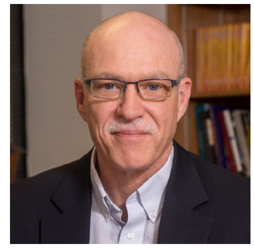

Tim R. Sass is a Distinguished University Professor in the Department of Economics at Georgia State University and the W. J. Usery Chair of the American Workplace in the Andrew Young School of Policy Studies. He is also the faculty director of the Metro Atlanta Policy Lab for Education (MAPLE). His research interests include the teacher labor supply, the measurement of teacher quality, and school choice. His work has been published in numerous academic journals and has been supported by several federal and philanthropic grants. He has acted as a consultant to school systems across the country. He is also a senior researcher at the Center for Analysis of Longitudinal Data in Education Research (CALDER).

M. Cade Lawson

M. Cade Lawson is a graduate research assistant with the Georgia Policy Labs. He is currently pursuing a Ph.D. in economics at Georgia State University, where his research is supported by the National Science Foundation's Graduate Research Fellowship. He received his B.S. in economics and M.S. in data analytics from the Georgia Institute of Technology. His research area of interest is the application of data science and machine learning techniques to issues in public economics.

About the Georgia Policy Labs

The Georgia Policy Labs is an interdisciplinary research center that drives policy and programmatic decisions that lift children, students, and families—especially those experiencing vulnerabilities. We produce evidence and actionable insights to realize the safety, capability, and economic security of *every* child, young adult, and family in Georgia by leveraging the power of data. We work alongside our school district and state agency partners to magnify their research capabilities and focus on their greatest areas of need. Our work reveals how policies and programs can be modified so that every child, student, and family can thrive.

Housed in the Andrew Young School of Policy Studies at Georgia State University, we have three components: the Metro Atlanta Policy Lab for Education (metro-Atlanta K–12 public education), the Child & Family Policy Lab (supporting children, families, and students through a cross-agency approach), and the Career & Technical Education Policy Exchange (a multi-state consortium exploring high-school based career and technical education).

Learn more at gpl.gsu.edu.