Value-Added and Growth Models in Education

Cory Koedel
K-12 Education
Last Updated: December 2025

Introduction

Value-added models (VAMs), also referred to as growth models, are used in education research and policy to estimate the impacts of districts, schools, and teachers on student outcomes. Achievement is often the focal outcome of a VAM but not always. The "value-added" terminology is borrowed from the long-standing production literature in economics. In this literature, this terminology refers to the amount by which the value of a product is increased at each stage of the production process. In education applications, the idea is that if we know each student's level of achievement up to some point—e.g., by the conclusion of the prior school year—then we can estimate how the inputs applied during the current school year contribute to achievement growth.

The process through which value added is estimated can be thought of as occurring in two steps (although, in practice, estimation often occurs all at once). In the first step, the VAM uses available data from an education system (e.g., a state system) to make a prediction about each student's academic performance based on other similar students. Next, the deviation from the prediction is calculated for each student. These student-level deviations are then aggregated up to the teacher, school, or district level. Teachers, schools, and districts whose students consistently outperform their predictions are identified as having "high value added," and vice versa for teachers, schools, and districts whose students consistently underperform their predictions.

Key to the success of any VAM is the quality of the predictor variables in the model. By far, the most important predictor is prior academic achievement, which is where the "value-added" terminology comes from. Other common predictors include indicators for student gender, race/ethnicity, poverty status, language status, disability status, and student mobility across schools. Some or all of these additional predictor variables are excluded in some policy applications of VAMs, primarily for political reasons or due to a misunderstanding of statistics.

This article summarizes key findings from research on value added and discusses the supporting evidence. It is not intended to provide the technical information necessary to construct a VAM.[1]

Key findings

Research on value added has concentrated primarily on teachers. There are also smaller bodies of research on school value added and principal value added. There is virtually no research on district value added. The key findings described in this chapter are divided by whether they are from teacher research (5 findings), school research (1 finding), or principal research (1 finding).

Teacher value added

Key finding #1: *Estimates of teacher value added are unbiased, on average, and provide useful information about teacher effectiveness.*

There is widespread agreement that value added is a valid methodology for measuring the contributions of teachers, at least on average. Although there is not absolute consensus among researchers, this common view has contributed to rapid growth in the use of teacher value added in research applications.

Key finding #2: *Teacher value added is better than alternative effectiveness indicators at predicting future value added.*

Teacher value added is more predictive of future value added than resume qualifications (e.g., teacher experience and education levels), evaluations based on classroom observations, and evaluations based on student surveys. Notably, teacher evaluations based on classroom observations—by far the predominant tool used to evaluate teachers in U.S. public schools—have been shown to exhibit biases not found in value added.

Key finding #3: *While value added has many desirable properties as a measure of teacher effectiveness, it has weaknesses that have limited its use in teacher policy applications.*

Value added does not unfairly penalize teachers who have particular characteristics, work with particular types of students, or are in particular types of schools, but this does not mean it is an error-free measure. The best evidence suggests that errors in teachers' value-added estimates are not systematic (and, hence, there is no bias); however, like any statistical estimate, there are errors.

Additionally, teacher evaluation systems that incorporate value added have encountered practical problems. For instance, value added cannot be reliably estimated for teachers outside of tested grades and subjects, making it difficult to use value added in districtwide or statewide policies meant to apply to all teachers. Value added has also been criticized for not being a transparent measure of teacher effectiveness and because it does not provide information on how to improve.

Key finding #4: *Teachers add value to more than just test scores.*

VAMs are used most commonly to estimate teacher effects on student test scores. However, an emerging body of research is adapting value-added tools to estimate teacher contributions to nontest outcomes such as attendance, high school graduation, grade point average, and socioemotional well-being. Teacher effects on these outcomes are as large as teacher effects on test scores. However, teacher value added to test and nontest outcomes is not highly correlated within teachers. This finding supports the view that teacher effectiveness is best conceptualized as a set of skills and cannot be summarized by effectiveness in a single domain.

Key finding #5: *When information about value added is available, actors in the teacher labor market respond to it.*

In settings where information about value added becomes available to employers and teachers, job opportunities for high-value-added teachers increase. Teachers with low value added are more likely to exit or move to districts that do not collect and/or report information on value added.

School value added

Key finding #6: *Estimates of school value added are unbiased, on average, and provide useful information about school effectiveness.*

The literature on school value added is much smaller than that on teacher value added. However, the evidence that is available suggests VAMs are an effective methodology for identifying the impacts of schools on student achievement.

Principal value added

Key finding #7: *There is no consensus in research on how to estimate value added for school principals.*

One finding from research on principals is clear: school value added is NOT an accurate measure of principal value added.

Understudied topics: The literature on teacher value added is rich. We know much about how to design VAMs, and the validity of value-added estimates has been tested in many settings. The development of high-quality, alternative measures of teacher effectiveness is an adjacent area where more work is needed. It would be of great value to develop new measures of teacher effectiveness that are unbiased and predictive of future performance, like value added, but that overcome some of its limitations, such as being feasible to estimate only for teachers in tested grades and subjects. Validation tools developed in VAM research could be applied to rigorously test new alternatives.

In contrast, the literature on school value added is thin, even though school value added is used more frequently than other types of value added in U.S. education policy. For example, many states use value-added measures in their school accountability systems. The small body of evidence on the validity of school value added is encouraging, but surely there is more to learn.

There is virtually no research on the value added of school districts, although some states incorporate value added into their district accountability systems. Evidence on the validity of teacher and school value added suggests that district value added is likely to be a useful measure of district effectiveness, but I am not aware of any direct evidence.

Evidence

Teacher Value Added

Key finding #1: *Estimates of teacher value added are unbiased, on average, and provide useful information about teacher effectiveness.*

An unbiased estimate of teacher effectiveness is an estimate that has an expected value equal to the actual effectiveness of the teacher. Put differently, factors that are not attributable to the teacher—e.g., those related to the background of the students in the teacher's classroom, observed or unobserved—do not distort the estimate.

One way in which researchers have tested whether teacher value added is unbiased has been to estimate it in one setting, Setting A, and to then see how well it predicts student achievement when the teacher is observed in a different setting, Setting B. If teacher value added is unbiased, then an

estimate in Setting A should forecast student achievement gains in Setting B accurately, on average. This test has been conducted in several experimental studies, where teacher value added—estimated using data from business-as-usual classroom assignments—is used to predict student achievement when students are randomly assigned to teachers. Research has shown that business-as-usual estimates of value added accurately predict student achievement under random assignment.[2]

A related strand of research conducts tests in a similar spirit in nonrandom settings. In these studies, researchers estimate teacher value added and then test whether student achievement changes as value added would predict when teachers switch schools and grades. These tests are more complicated than their experimental counterparts due to the potential for contaminating factors. However, well-designed tests in several settings have shown that teacher value added is an unbiased predictor of student achievement.[3]

An alternative approach used to test for bias in VAMs adds information to the model that is typically unavailable in standard education datasets. This type of test is a response to the concern that a VAM estimated with standard data does not include information about all of the factors that determine student test scores. If the factors that are not included are correlated with teacher value added, then the model would lead to biased estimation. This is an example of *omitted variable bias*, which is a common type of bias in statistical and econometric analyses.

The most notable example of this type of test was conducted by researchers Raj Chetty, John Friedman, and Jonah Rockoff, who linked school data to detailed tax records for students in a large, urban school district over many years.[4] They estimated a typical VAM using standard data and then another VAM that included information about students' families from the tax data. The key variables from the tax data are the mother's age at the student's birth, indicators for parental 401(k) contributions and home ownership, and an indicator for parental marital status interacted with a quartic in household income. While the information in the tax data strongly predicts student achievement, they find its absence from the "standard-data" VAM causes almost no bias. The authors give two reasons for this result. First, the controls in their standard-data VAM (lagged test scores, poverty status as measured by free and reduced-price lunch, etc.) capture much of the variation in the tax-data variables. Second, the variation in the tax-data variables that is otherwise unaccounted for in the standard-data VAM is essentially uncorrelated with teacher value added. Put differently, conditional on the information in standard education datasets, students are not systematically sorted to high- or low-value-added teachers along the dimensions measured by these variables.

The evidence on the validity of value-added estimates for individual teachers strongly implies that value added can also be used to assess differences in effectiveness for different groups of teachers. For example, researchers have used VAMs to compare teachers from different preparation pathways and to compare teachers who perform differently on screening tools used by school districts during hiring, among other applications.[5]

While the results of most research are consistent with the claim that value added is an unbiased estimate of teacher effectiveness, there is not complete consensus on this point. Interested readers may benefit from reading the exchange between several prominent researchers in this area—Raj Chetty, John Friedman, Jonah Rockoff, and Jesse Rothstein.[6] The exchange gives contrasting viewpoints on the quality of evidence showing that teacher value added is unbiased. It is also important to give the caveat that the accuracy of estimates of teacher value added depends on the quality of the model. Andrew Bacher-Hicks and I provide a technical discussion on model design.[7]

Key finding #2: *Teacher value added is better than alternative effectiveness indicators at predicting future value added.*

Teacher value added is generally less biased and a stronger predictor of future teacher value added than available alternative measures of teacher effectiveness. One set of alternative measures includes teachers' observable qualifications—i.e., "resume characteristics," such as experience, the teacher's own education level and degree field, certification test scores, and the prestige of the undergraduate institution. The ability of these types of indicators to identify effective teachers has been widely studied. To summarize this research, teacher certification, teacher education, and other observable teacher qualifications are positively but only weakly related to teacher effectiveness. It follows that they are poor predictors of future teacher value added.

Among in-service measures, teacher effectiveness indicators based on direct observation in the classroom by school principals are the most common. Classroom observations of teachers are usually positively correlated with teacher value added, which, given the strong evidence of the validity of value added, suggests that they contain useful information about teacher effectiveness. However, the correlation between classroom observations and value added is low.[8] This low correlation could be the result of several factors, such as (a) classroom observations pick up dimensions of teacher effectiveness missed by value added; (b) classroom observations (and value added) are noisy; and (c) classroom observations are biased. I am not aware of any direct evidence for explanation (a). However, research shows that teacher evaluations based on classroom observations are noisier than value added (point b), and two measures can have a low correlation even if they are measuring the same construct if they are noisy.[9] Perhaps more importantly, there is also evidence that teacher evaluations based on classroom observations are biased (point c).[10] Teachers who teach less socioeconomically advantaged students and students with lower prior achievement systematically receive lower observation-based evaluation ratings after their actual performance is accounted for. There is also evidence that teachers with particular characteristics—most notably, male teachers and teachers of color—receive classroom observation scores that are negatively biased.

Finally, there is a small body of research on the use of student surveys to measure teacher effectiveness. There is evidence that student surveys contain useful information along dimensions that value added does not capture. However, survey-based measures of teacher effectiveness are weak and imprecise predictors of future teacher value added, and they may also be biased.[11]

Key finding #3: *While value added has many desirable properties as a measure of teacher effectiveness, it has weaknesses that have limited its use in teacher policy applications.*

As a measure of teacher effectiveness, value added has several weaknesses. First, it can measure effectiveness only in terms of the outcome used in the statistical model of student performance. The most common VAMs in research and policy applications are based on student achievement, which raises concerns about whether and how VAMs can be used to capture teacher effectiveness along other dimensions. A relatively recent innovation in VAM research is the development of VAMs for nontest outcomes, including attendance, high school graduation, grade point average, and socioemotional well-being.[12] The literature on these "alternative-outcome" VAMs is not as robust as the achievement-based literature, but the evidence available so far suggests optimism about our ability to use VAMs to capture teacher effects along many dimensions of student performance. I elaborate on the findings from research on alternative-outcome VAMs below.

A more problematic feature of teacher value added is that like any statistical estimate, it is subject to estimation error. In this case, the estimation error problem is nontrivial because statistically, the number of students taught by a typical teacher is small. It is appropriate to describe individual estimates of teacher value added (from a well-specified model) as unbiased but imprecise. That is, the expected value of a value-added estimate is the true effectiveness of the teacher, but the estimate has a wide error band. For questions about the full population of teachers—e.g., "How large are differences between teachers over the distribution of teacher effectiveness?" or "Do teachers who undergo training improve?"—we can remove the influence of the errors in the individual teacher estimates and provide good answers. However, if we want to know the effectiveness of an individual teacher in a particular year, then we must recognize that our estimate of value added will include estimation error.

In a clever illustration of estimation error in teacher value added, one group of researchers uses a VAM to estimate the "effects" of teachers on student height.[13] Of course, we do not believe that teachers affect student height—this is an example of a placebo test, where the expected effect of teachers is zero. However, the researchers find nonzero teacher "effects" on height, which raises concerns about the efficacy of VAMs. The researchers go on to show that the apparent value added of teachers to height is driven by idiosyncratic, nonpersistent estimation error in the individual teacher estimates. This error can be removed in well-designed research studies, but in some policy applications, it is difficult to remove.

In applications where we want to, for example, identify individual teachers as "effective" or "ineffective," it is unclear how much estimation error is acceptable. On the one hand, no indicator of teacher effectiveness is error free, even if it is hard to measure the level of imprecision. For example, classroom observations by school principals are error prone, but it is difficult to quantify the amount of error. On the other hand, the fact that we know that value added is estimated with error and can precisely quantify the magnitude of the error understandably causes trepidation with regard to overreliance on value added in consequential personnel decisions.

Research generally finds that despite the estimation error, individual estimates of teacher value added have enough information to be useful for informing personnel decisions. For instance, the year-to-year stability of teacher value added is similar to the stability of performance ratings used in other occupations to inform high-stakes decisions.[14] There is also research on the likely impacts of using value added to reward effective teachers and dismiss ineffective teachers, inclusive of the estimation error. These studies are purely theoretical, but they generally find that personnel decisions made based on value added would improve workforce quality (although studies vary widely in terms of how much improvement is possible).[15] One general lesson from this research is that if value added is to be used to identify the effectiveness of individual teachers, then it makes the most sense to focus on identifying the highest- and lowest-performing teachers. Differentiating between teachers within the middle of the performance distribution is much more difficult.

Although no state or school district has attempted to evaluate teachers entirely based on value added, several large school districts have implemented teacher evaluation systems that include value added among several indicators of effectiveness. The efficacy results from these systems are mixed. In some instances, workforce quality has improved, but in others, meaningful gains have failed to materialize.

The development of these evaluation systems has also brought to light logistical challenges associated with incorporating value added into personnel policies for teachers. For instance, value added can be credibly estimated only for teachers in tested grades and subjects, but it is politically

difficult for a school district to implement an evaluation system for only some teachers. In practice, districts have used substitute measures for teachers outside of tested grades and subjects, but the quality of these substitute measures is suspect. Another practical problem is that value added is not considered a transparent indicator of performance. That is, because value added is the product of a complex statistical model, it cannot be readily understood or deconstructed by the teachers who are being evaluated. It is unclear why other effectiveness indicators—e.g., a principal's evaluation—are viewed as more transparent than value added, but intractability is a common complaint about value added and has been raised in court cases about the use of value added in formal teacher evaluations.[16]

Key finding #4: *Teachers add value to more than just test scores.*

Researchers are increasingly designing VAMs to estimate teacher effects on nontest student outcomes. Teacher value added has been estimated for the following nontest outcomes: course grades, attendance, suspensions, grade repetition, growth mindset, grit, effort, self-efficacy, happiness, and in-class behavior.[17]

Nontest VAMs typically use a host of lagged-outcome controls, similarly to how test-based VAMs control for lagged test scores. Some of the most compelling studies conduct validity tests along the lines of those described above for traditional, test-based VAMs. That is, they use instances of teachers switching between schools and grades to test the predictive validity of non-test-based value added. Using these tests, researchers have repeatedly shown that it is possible to recover unbiased estimates of teacher value added to nontest outcomes.

Research on non-test-based VAMs shows that teachers have impacts on nontest outcomes that are comparable in magnitude to their effects on test scores. However, value added is only weakly correlated across test and nontest outcomes within teachers. This finding supports the view that teacher effectiveness is best conceptualized as a set of skills and cannot be summarized by effectiveness in a single domain.

Key finding #5: *When information about value added is available, actors in the teacher labor market respond to it.*

The research community has largely (but not entirely) reached the consensus that value added is a credible methodology for measuring teacher effectiveness. However, estimates for individual teachers are noisy. When information about value added becomes available, do on-the-ground actors in the teacher labor market—i.e., employers (districts), supervisors (principals), and teachers themselves—care about it?

One way to answer this question is to see whether aspects of the labor market change when information about value added becomes available. Do districts and school principals try to hire teachers with higher value added? Do high-value-added teachers have more career opportunities when their value added becomes known? If so, do they leverage this newfound advantage, and if so, how do they do so?

The answers to these questions are most interesting in settings where the use of value added in personnel decisions is not dictated from above. Put differently, if a district implements a new policy that requires principals to consider value added when they make personnel decisions, then it is less interesting to know that principals consider the information. Alternatively, if information about value added is made available without prescriptive guidance regarding its use, then we can monitor

changes in behavior in the labor market to gain insight into the perceived "labor market value" of value added.

The literature on this topic is relatively small, but it suggests that the labor market responds to information about teacher value added. For instance, in one experiment, researchers randomly assigned access to information about value added to school principals.[18] Principals who received the information incorporated it into their annual evaluations of teachers, leading to systematically higher principal ratings for high-value-added teachers. Low-value-added teachers in mathematics were also more likely to exit their schools after the information on value added was provided. The researchers interpreted their findings as showing that school principals value and act on information about teacher value added.

Another study examined what happened in the teacher labor market in North Carolina when two large school districts began to report teacher value added to school principals and teachers themselves.[19] When this happened, teachers with higher value added became more likely to change schools within their districts and, in particular, to move to higher-performing schools. This result is consistent with school principals valuing—and acting on—information about value added and with high-value-added teachers leveraging this information to improve their teaching placements.[20] The study also found that teachers with low value added exited these districts at a higher rate and moved to districts that do not make information about value added available to principals and teachers.

Although research shows that the labor market responds to information about value added and the cost of providing this information is relatively low, the estimated effects of the information are small. Hence, while the market response suggests that the information conveyed by value added is useful, value added does not seem to be relied upon too heavily. Given the limitations in the use of value added to estimate the effects of individual teachers described above, perhaps this response to value added is appropriate.

School value added

Key finding #6: *Estimates of school value added are unbiased, on average, and provide useful information about school effectiveness.*

There is a small body of research on the prospects for using value added to measure school effectiveness. As in the teacher literature, the most prominent question addressed by this research is whether school value added is an unbiased estimate of school effectiveness. Recall from above that an unbiased estimate is an estimate where factors that are not attributable to the school—e.g., those related to the background of the students who attend the school, observed or unobserved—are not included in the estimate. The strong evidence on validity from research on teachers suggests that school value added is also likely to have little to no bias. Direct tests are uncommon, but some have been conducted. The most compelling tests leverage the random assignment of students to schools to validate estimates of school value added.

One such study uses school choice lottery data from Charlotte-Mecklenburg Schools (CMS), where students are guaranteed admission to their local neighborhood schools but can also apply to attend other CMS schools.[21] When schools are oversubscribed (i.e., when there are more applicants than open seats), enrollment is determined by random lottery. Leveraging data from randomly assigned winners and losers, the study shows that school value added—estimated nonexperimentally via data from business-as-usual school assignments—accurately predicts student achievement gains for students who are admitted by lottery into these schools, on average.

Another study that leverages lottery-based school assignments in Boston reaches a similar conclusion.[22] It finds that school value added forecasts student achievement accurately, on average, for students who are randomly assigned to schools. Although this study shows that there can be prediction errors for individual schools (and develops some advanced methodological procedures to reduce the influence of the errors when some lottery data are available), it ultimately concludes that conventional estimates of school value added provide useful information about school effectiveness.

A limitation of both studies is that their tests for unbiasedness are not especially well powered, which makes it difficult to fully rule out bias in estimates of school value added.

Principal value added

Key finding #7: *There is no consensus in research on how to estimate value added for school principals.*

Unlike in the teacher literature, there is considerable disagreement among researchers over whether VAMs can isolate the effects of principals. Some researchers are optimistic about the prospects of applying VAMs to this problem; other researchers are pessimistic.[23]

It is much more difficult to estimate principal value added than to estimate teacher value added. Why? The most important reason is that schools have only a single principal at a time, which makes it difficult to separate the effect of the principal from the effect of the school. Another problem is that principals turn over often, with the typical principal spell in a school lasting for only a few years. Compounding this problem, principal effects may not coincide exactly with the timing of a principal's exposure to students. For instance, a principal who is new to a school will inherit the staff, procedures, and work environment established by his or her predecessor. Changes implemented by the new principal could take years to come to fruition and may persist beyond the principal's tenure.

One takeaway from the research on principals is that we are very far from being able to estimate principal value added in a way that can be useful for policy—e.g., to identify the most and least effective principals in a statewide or large-district workforce. The methods that are the most promising for estimating principal value added might allow us to answer important research questions, such as how much principals matter overall or whether specific characteristics of principals are predictive of effectiveness, but they are not well suited to estimate individual principal effects. We also know that school value added is a poor proxy for principal value added. The reason is that there are many factors that contribute to the variation in student outcomes across schools and within schools over time, and very little of the total variation is explained by principals.[24]

Endnotes and references

[1] For a technical overview of value-added estimation, see Bacher-Hicks, Andrew, and Cory Koedel. 2023. Estimation and Interpretation of Teacher Value Added in Research Applications. In Handbook of the Economics of Education. Volume 6. Edited by Eric A. Hanushek, Stephen Machin and Ludger Woessmann. Amsterdam: Elsevier Science. 93–134.

[2] Kane, Thomas J., and Douglas O. Staiger. 2008. Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation (Working Paper No. 14607). National Bureau of Economic Research; Kane, Thomas J., Daniel F. McCaffrey, Trey Miller, and Douglas O. Staiger. 2013. Have We Identified Effective Teachers? Validating Measures of Effective Teaching Using Random Assignment. Bill and Melinda Gates Foundation; Bacher-Hicks, Andrew, Mark J. Chin, Thomas J. Kane, and Douglas O. Staiger. 2019. An Experimental Evaluation of Three Teacher Quality Measures: Value-Added, Classroom Observations, and Student Surveys. Economics of Education Review 73: 1–15.

[3] Bacher-Hicks, Andrew, Thomas J. Kane, and Douglas O. Staiger. 2014. Validating Teacher Effect Estimates Using Changes in Teacher Assignments in Los Angeles (Working Paper No. 20657). National Bureau of Economic Research; Chetty, Raj, John N. Friedman, and Jonah E. Rockoff. 2014a. Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates. American Economic Review 104(9): 2593–2632; Cullen, Julie Berry, Cory Koedel, and Eric Parsons. 2021. The Compositional Effect of Rigorous Teacher Evaluation on Workforce Quality. Education Finance and Policy 16(1): 7–41.

[4] Chetty, Friedman, and Rockoff (2014a).

[5] The literature on this topic is too large to list all relevant studies. Several examples include the following: Backes, Ben, Dan Goldhaber, Whitney Cade, Kate Sullivan, and Melissa Dodson. 2018. Can UTeach? Assessing the Relative Effectiveness of STEM Teachers. Economics of Education Review 64: 184–198; Bruno, Paul, and Katharine O. Strunk. 2019. Making the Cut: The Effectiveness of Teacher Screening and Hiring in the Los Angeles Unified School District. Educational Evaluation and Policy Analysis 41(4): 426–460; von Hippel, Paul T., and Laura Bellows. 2018. How Much Does Teacher Quality Vary Across Teacher Preparation Programs? Reanalyses from Six States. Economics of Education Review 64: 298–312.

[6] Rothstein, Jesse. 2017. Measuring the Impacts of Teachers: Comment. American Economic Review 107(6): 1656–1684; Chetty, Raj, John N. Friedman, and Jonah E. Rockoff. 2017. Measuring the Impacts of Teachers: Reply. American Economic Review 107(6): 1685–1717.

[7] Bacher-Hicks and Koedel (2023).

[8] Kane, Thomas J., and Douglas O. Staiger. 2012. Gathering Feedback for Teaching: Combining High-Quality Observations with Student Surveys and Achievement Gains. Bill and Melinda Gates Foundation. Correlations between principals' overall ratings of teachers and value added are similarly low. See Harris, Douglas N., and Tim R. Sass. 2014. Skills, Productivity, and the Evaluation of Teacher Performance. Economics of Education Review 40: 183–204; Jacob, Brian A., and Lars Lefgren. 2008. Can Principals Identify Effective Teachers? Evidence on Subjective Performance Evaluations in Education. Journal of Labor Economics 26(1): 101–136.

[9] Bacher-Hicks, Chin, Kane, and Staiger (2019).

[10] Steinberg, Matthew P., and Rachel Garrett. 2016. Classroom Composition and Measured Teacher Performance: What do Teacher Observation Scores Really Measure? Educational Evaluation and Policy Analysis 38(2): 293–317;
Campbell, Shanyce L., and Matthew Ronfeldt. 2018. Observational Evaluation of Teachers: Measuring More than We Bargained For? American Educational Research Journal 55(6): 1233–1267; Bartanen, Brendan, and Andrew Kwok. 2021. Examining Clinical Teaching Observation Scores as a Measure of Preservice Teacher Quality. American Educational Research Journal 58(5): 887–920; Grissom, Jason A., and Brendan Bartanen. 2021. Potential Race and Gender Biases in High-Stakes Teacher Observations. Journal of Policy Analysis and Management 41(1): 131–161.

[11] Bacher-Hicks et al. (2019).

[12] Backes, Ben, James Cowan, Dan Goldhaber, and Roddy Theobald. Forthcoming. How to Measure a Teacher: The Influence of Test and Nontest Value-Added on Long-Run Student Outcomes. Journal of Human Resources; Blazar, David, and Matthew A. Kraft. 2017. Teacher and Teaching Effects on Students' Attitudes and Behaviors. Educational Evaluation and Policy Analysis 39(1): 146–170; Gershenson, Seth. 2016. Linking Teacher Quality, Student Attendance, and Student Achievement. Education Finance and Policy 11(2): 125–149; Jackson, Kirabo. 2018. What Do Test Scores Miss? The Importance of Teacher Effects on Non-Test Score Outcomes. Journal of Political Economy 126(5):

2072–2107; Kraft, Matthew A. 2019. Teacher Effects on Complex Cognitive Skills and Social-Emotional Competencies. Journal of Human Resources 54(1): 1–36; Liu, Jing, and Susanna Loeb. 2021. Engaging Teachers: Measuring the Impact of Teachers on Student Attendance in Secondary School. Journal of Human Resources 56(2): 353–379.

[13] Bitler, Marianne, Sean P. Corcoran, Thurston Domina, and Emily K. Penner. 2021. Teacher Effects on Student Achievement and Height: A Cautionary Tale. Journal of Research on Educational Effectiveness 14(4): 900–924.

[14] McCaffrey, Daniel F., Tim R. Sass, J. R. Lockwood, and Kata Mihaly. 2009. The Intertemporal Variability of Teacher Effect Estimates. Education Finance and Policy 4(4): 572–606; Glazerman, Steven, Susanna Loeb, Dan Goldhaber, Douglas O. Staiger, Stephen Raudenbush, and Grover J. Whitehurst. 2010. Evaluating Teachers: The Important Role of Value-Added. Brown Center on Education Policy at Brookings.

[15] Staiger, Douglas O., and Jonah E. Rockoff. 2010. Searching for Effective Teachers with Imperfect Information. Journal of Economic Perspectives 24(3): 97–118; Hanushek, Eric A. 2011. The Economic Value of Higher Teacher Quality. Economics of Education Review 30(3): 466–479; Rothstein, Jesse. 2015. Teacher Quality Policy When Supply Matters. American Economic Review 105(1): 100–130; Winters, Marcus A., and Joshua M. Cowen. 2013. Would a Value-Added System of Retention Improve the Distribution of Teacher Quality? A Simulation of Alternative Policies. Journal of Policy Analysis and Management 32(3): 634–654.

[16] Paige, Mark A., and Audrey Amrein-Beardsley. 2020. "Houston, We Have a Lawsuit": A Cautionary Tale for the Implementation of Value-Added Models for High-Stakes Employment Decisions. Educational Researcher 49(5): 350–359.

[17] Backes et al. (forthcoming); Blazar and Kraft (2017); Gershenson (2016); Jackson (2018); Kraft (2019); Liu and Loeb (2021).

[18] Rockoff, Jonah E., Douglas O. Staiger, Thomas J. Kane, and Eric S. Taylor. 2012. Information and Employee Evaluation: Evidence from a Randomized Intervention in Public Schools. American Economic Review 102(7): 3184–3213.

[19] Bates, Michael. 2020. Public and Private Employer Learning: Evidence from the Adoption of Teacher Value Added. Journal of Labor Economics 38(2): 375–420.

[20] A similar pattern of mobility was also found in Houston when the district introduced an evaluation reform that included the use of value added in combination with other measures of effectiveness. See Cullen, Koedel, and Parsons (2021).

[21] Deming, David J. 2014. Using School Choice Lotteries to Test Measures of School Effectiveness. American Economic Review 104(5): 406–411.

[22] Angrist, Joshua D., Peter D. Hull, Parag A. Pathak, and Christopher R. Walters. 2017. Leveraging Lotteries for School Value-Added: Testing and Estimation. Quarterly Journal of Economics 132(2): 871–919.

[23] Bartanen, Brendan, Aliza N. Husain, and David D. Liebowitz. 2024. Rethinking Principal Effects on Student Outcomes. Journal of Public Economics 234: 105–115; Cullen, Julie Berry, Eric A. Hanushek, Gregory Phelan, and Steven G. Rivkin. 2024. Performance Information and Personnel Decisions in the Public Sector: The Case of School Principals. Journal of Human Resources 59(1): 109–140; Grissom, Jason A., Demetra Kalogrides, and Susanna Loeb. 2015. Using Student Test Scores to Measure Principal Performance. Educational Evaluation and Policy Analysis 37(1): 3–28.

[24] Bartanen, Husain, and Liebowitz (2024).

Suggested citation

Koedel, Cory (2025). "Value-Added and Growth Models in Education," *in Live Handbook of Research on Education Finance and Policy*, in Douglas Harris (ed.), Association for Education Finance and Policy, viewed MM/DD/YYYY, <WEB LINK>.