

Different Methods for Assessing Preservice Teachers' Instruction: Why Measures Matter

Journal of Teacher Education
2024, Vol. 75(2) 168–185
© 2023 American Association of
Colleges for Teacher Education
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/00224871231200279
journals.sagepub.com/home/jte



Arielle Boguslav¹  and Julie Cohen¹ 

Abstract

Teacher preparation programs are increasingly expected to use data on preservice teacher (PST) skills to drive program improvement and provide targeted supports. Observational ratings are especially vital, but also prone to measurement issues. Scores may be influenced by factors unrelated to PSTs' instructional skills, including rater standards. Yet we know little about how these measurement challenges play out in the preservice context specifically. Here, we investigate the reliability and sensitivity of two observational measures. We find measures collected during student teaching are especially prone to measurement issues; only 3% to 4% of variation in scores reflects consistent differences between PSTs, while 9% to 17% of variation can be attributed to the mentors with whom they work. When high scores stem not from strong instructional skills, but instead from external circumstances, we cannot use them to make consequential decisions about PSTs' individual needs or readiness for independent teaching.

Keywords

field experiences, performance assessment, practice-based teacher education, preservice teacher education

Over the past two decades, teacher preparation programs (TPPs) have become increasingly responsible for the collection and analysis of data on PST knowledge, skills, and dispositions (Bastian et al., 2016). Several states have implemented accountability systems that require TPPs to provide data on their graduates, share these data as a measure of program effectiveness, and even levy penalties on programs deemed ineffective (Levine, 2006; Texas State Legislature, 2009; U.S. Department of Education, 2011). The theory of action behind these policies is TPPs can use PST data to provide more targeted support based on PSTs' needs, as well as adjust programmatic experiences based on the degree to which such experiences promote PSTs' development (Bastian et al., 2018; Davis & Peck, 2020).

Most TPPs collect observational ratings of PSTs' instructional practice during clinical placements, among other measures, (Feuer et al., 2013). In capturing observable skills—what PSTs *do* in interactions with students—these measures are vital because TPPs are ultimately charged with preparing novices who engage in productive and supportive interactions with children. These measures are also particularly challenging to implement in reliable ways (Bell et al., 2012; Gitomer, 2009; Kane & Staiger, 2012). Ratings of teaching practices are often influenced by factors that may not be related to a teacher's instructional skills, including rater standards and the characteristics of students in a

classroom (Bartanen & Kwok, 2021; Campbell & Ronfeldt, 2018). As a result, TPPs risk erroneously attributing differences between external factors to differences in PSTs' instructional skills. For example, if observational scores are substantially influenced by supervisors' rating standards, the PSTs with the lowest scores may not be those most in need of targeted support, but instead, may be those rated by supervisors with harsher standards (Bartanen & Kwok, 2021). Moreover, unlike inservice teachers, PSTs are not solely responsible for their own classrooms, but instead teach in the context of mentors' classrooms, whose characteristics and teaching skills may further influence PST scores on observational measures.

Figuring out how to measure instructional quality has been a long-standing puzzle for education researchers, but few have explored the challenges particular to doing so in the context of preservice preparation, when time is in short supply and candidates are typically observed in contexts that are not fully "theirs." TPPs use PST observation scores to meet accreditation reporting requirements and guide

¹University of Virginia, Charlottesville, USA

Corresponding Author:

Arielle Boguslav, Annenberg Institute at Brown University, Box 1985, Providence, RI 02912, USA.
Email: arielle_boguslav1@brown.edu

programmatic decisions (American Association of Colleges for Teacher Education [AACTE], 2018; Council for the Accreditation of Educator Preparation [CAEP], 2022). If these measures suffer from the challenges described above, then we must re-evaluate what conclusions can reasonably be drawn from those data.

In this article, we investigate these measurement challenges using two distinct observational measures of PSTs' instructional skills, implemented as part of one TPP's efforts to generate robust data on such skills. The first measure is the Classroom Assessment Scoring System (CLASS) (Pianta & Hamre, 2009), designed to provide a holistic view of PSTs' instruction during clinical placements. The second measure provides a more focused view of PSTs' implementation of discrete instructional skills culled from simulation-based instructional activities (Lampert & Graziani, 2009).

In investigating these challenges, we focus on two key concepts: reliability and sensitivity. Following several prior studies of the reliability of observational measures, we define reliability according to the generalizability theory framework (Bartanen & Kwok, 2020; Briggs & Alzen, 2019). Under generalizability theory, a PST's score on an observational rubric consists of their "true score," the signal, and measurement error, the noise. Rather than treating measurement error as a single, monolithic construct, as classical test theory does, generalizability theory decomposes error into distinct sources. For PSTs, we conceptualize the signal as information a TPP wants to learn about a PST's instructional skills. Crucially, because PSTs could go on to teach in a wide variety of contexts, TPPs are ultimately concerned with making inferences about individual PSTs' *transferable* skills rather than their skills in a particular classroom context working alongside a particular mentor (Bell et al., 2012). To provide reliable information about PSTs' skills, observation scores must *generalize* beyond the context of the observation. In other words, under generalizability theory, a perfectly reliable measure will result in the same, stable score for an individual PST regardless of the context in which a PST is observed or the rater assessing their skills.

In practice, of course, no measure is perfectly reliable, which is why generalizability theory distinguishes stable signal about PSTs from sources of measurement error, thereby providing an estimate of how much of the variation in scores reflects signal we care about. Demonstrating reasonably strong reliability (a.k.a. generalizability) is a necessary but insufficient precursor for demonstrating predictive validity. If observation scores primarily reflect rater standards or classroom context, then we can hardly expect them to provide TPPs with predictive information about PSTs' instructional skills as teachers of record. The reliability analyses we conduct in this article are concerned with understanding the extent to which individual PSTs' observation scores are stable or generalizable across different contexts and observation conditions.

Sensitivity is a practical challenge, often created by low reliability. We conceptualize sensitivity as the extent to which measures can detect statistically significant and practically meaningful differences when making comparisons between individual PSTs, groups of PSTs, or aspects of instructional quality (M. Allen & Coble, 2018). The less reliable a measure is, the less likely it is to distinguish among differences in PSTs' instructional skills, because those differences may be drowned out by the external factors that influence scores. To provide PSTs with targeted support, TPPs need measures that are sensitive enough to identify which PSTs may have less-developed skills.

Based on the measures themselves and how they were implemented, we have hypotheses about the reliability and sensitivity of CLASS scores and simulation-based scores. First, we hypothesize that both researcher-created measures have the potential to be more reliable than more typical observations conducted by clinical supervisors or mentors using home-grown rubrics, and thus better able to detect consistent differences among PSTs (Bartanen & Kwok, 2021; Bastian et al., 2018). Second, although CLASS may be more reliable than typical observational measures completed by individuals who know the PSTs, we hypothesize scores may be substantially influenced by external factors that do not reflect PSTs' instructional skills, including mentor teachers. Third, we hypothesize that raters will be a source of some measurement error for standardized simulation-based observations, though we anticipate this error will be small given the standardized nature of these measures and the rating procedures employed. Fourth, we hypothesize the stronger reliability of CLASS and simulation scores (relative to commonly used measures in teacher preparation) will contribute to greater sensitivity to differences between individual PSTs, groups of PSTs, their learning experiences, and instructional skills.

We use data collected by the TPP to test these hypotheses by answering the following research questions:

Research Question 1 (RQ1): To what extent does each measure capture consistent differences between PSTs? (Reliability)

Research Question 2 (RQ2): To what extent does variation in CLASS scores reflect differences between external factors, namely mentor teachers? (Reliability)

Research Question 3 (RQ3): To what extent does variation in simulation scores reflect differences between external factors, namely raters? (Reliability)

Research Question 4 (RQ4): How well can each measure differentiate between individual PSTs, groups of PSTs, and facets of instruction? (Sensitivity)

In this way, we can identify both the strengths of these measures and the ongoing challenges for learning about PST practice.

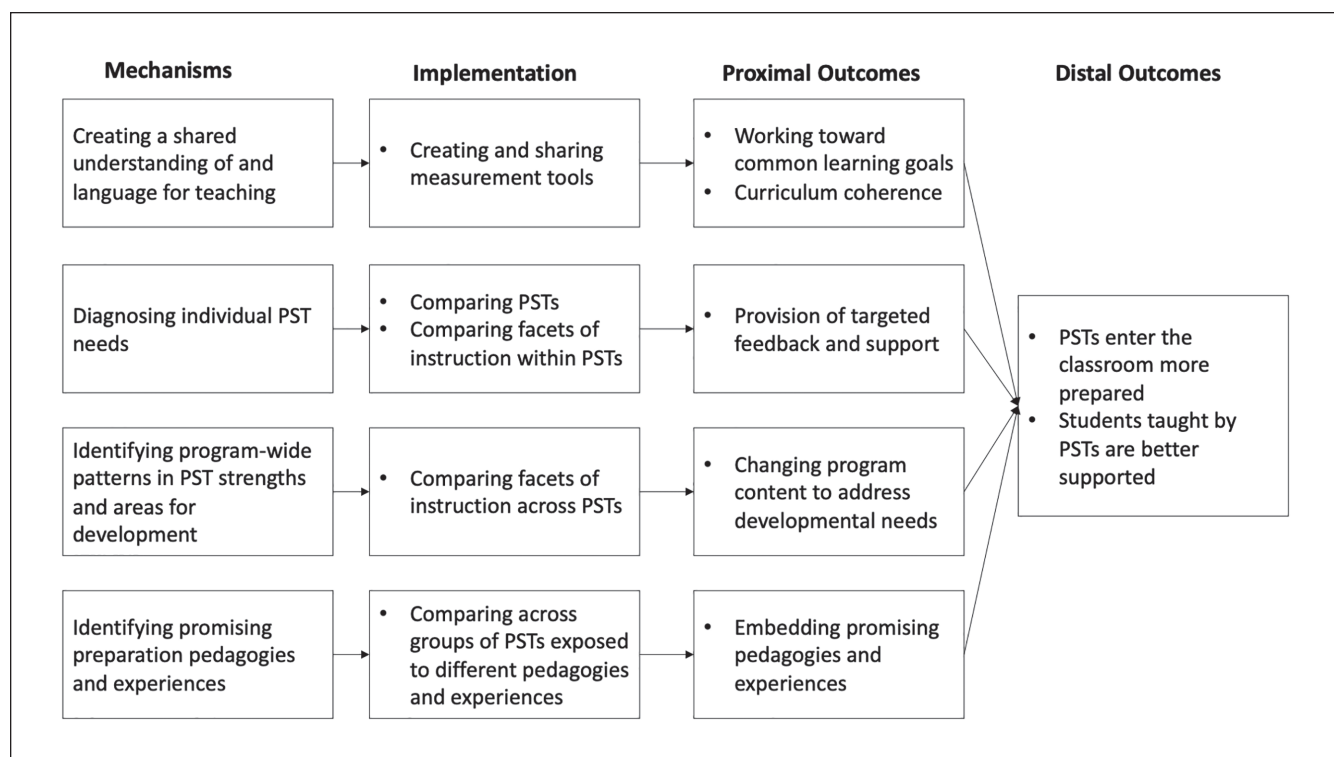


Figure 1. Conceptual Framework Articulating How Data on PST Knowledge, Skills, and Dispositions can Contribute to Improvements in Teacher Preparation.

Note. PST = preservice teacher.

The preservice period is both brief and formative. TPPs have a very short window to provide experiences and supports that position candidates for success as teachers of record. Despite the urgency of preparing skilled PSTs, we know very little about the extent to which the data TPPs currently collect can inform this work, or what changes might be necessary if they cannot. This article addresses this gap in two ways. First, we empirically explore the affordances and constraints of two measures designed to make inferences about PSTs' instructional skills. Second, we identify specific steps that TPPs should consider to enhance the reliability and sensitivity of the data they collect, as well as the tradeoffs such steps entail.

Background and Conceptual Framework

Why Observe PSTs' Instruction?

Many argue TPPs would benefit from having more insight about the degree to which preparation experiences support PSTs' development (e.g., Davis & Peck, 2020; Goldhaber, 2019). In our conceptual framework (Figure 1), we articulate a theory of action for how data can result in improvements in teaching when PSTs enter the classroom. In the first column, we highlight four mechanisms by which PST data may

support TPP improvement. In the remaining three columns, we indicate how TPPs may implement these mechanisms and highlight the proximal and distal outcomes.

The first mechanism highlights how measurement tools can facilitate a shared understanding and common language among TPP faculty, supervisors, mentors, and PSTs (Davis & Peck, 2020). This enables teacher preparation stakeholders to work toward the same learning goals and generate coherence across preparation experiences.

The remaining three mechanisms highlight how specific analyses can improve the supports and experiences TPPs provide. First, analyzing PST data can allow TPPs to diagnose and respond to PSTs' individual learning needs (M. Allen & Coble, 2018; Bastian et al., 2018). Second, sufficiently sensitive PST data can help TPPs identify program-wide areas for development, which can inform changes to course content and program curricula (M. Allen & Coble, 2018; Peck & McDonald, 2013). By comparing PSTs' mastery of different facets of instruction (e.g., relative strengths in classroom management vs. facilitating discussions), TPPs can tailor learning experiences to meet observed needs (e.g., practice facilitating student discourse). Moreover, TPPs can compare skill development across cohorts or licensure tracks (e.g., elementary vs. secondary) to identify areas of relative weakness and redesign or supplement program content in those areas. Finally, a TPP can compare the scores of PSTs

that participated in different preparation experiences, such as comparing scores between PSTs who completed rehearsals of core teaching practices and PSTs who did not (McDonald et al., 2013). In this way, TPPs can identify experiences that are more promising for supporting PSTs' development (Hill et al., 2020; Peck & McDonald, 2013).

What Makes a Measure Useful?

Realizing improvements from PST data depends on what the data can tell us. Here, we highlight two key measurement properties that influence the conclusions TPPs can draw: reliability and sensitivity.

Reliability. In line with generalizability theory, we conceptualize reliability as the extent to which differences in scores reflect differences between PSTs (Bell et al., 2012; Ho & Kane, 2013). When reliability is low, differences in scores are influenced primarily by differences in the conditions of assessment, rather than differences between PSTs. Measures requiring human judgment, including observational measures, tend to be less reliable than tests with a single, correct answer (Bell et al., 2019; Hill et al., 2012). This is because it is very difficult to ensure raters assign scores in the same way (Bell et al., 2015). Prior work suggests lesson content, time of year, and student demographics also contribute to relatively low reliability of inservice teacher observation scores (Bell et al., 2012; Casabianca et al., 2015; Ho & Kane, 2013).

There is limited research on the reliability of PST observation scores. Bartanen and Kwok (2021) applied a generalizability framework to decompose variation in PST observation scores from a large TPP in Texas, where PSTs were rated by clinical supervisors. Typical to many TPPs in the United States, supervisors were primarily tasked with supporting PST development but were also asked to rate PSTs at four timepoints during student teaching, using a proprietary observation rubric adapted from the state's teacher evaluation framework. The authors found only 20% of the variation in PST scores reflected consistent differences between PSTs, whereas 55% of the variation in scores could be attributed to differences in the stringency of supervisors' rating standards. The authors also suggest the potential for mentor effects: Since PSTs are typically observed in their mentor's classroom, the classroom climate may be shaped by the mentor's instructional practices rather than those of the PST. To our knowledge, all other existing work on the reliability of PST data focuses on nonobservational measures or on rater agreement, without considering other sources of statistical noise (e.g., Bastian et al., 2018).

The literature highlights two primary strategies for increasing measure reliability. First, TPPs can reduce the influence of assessment conditions by standardizing them. This includes providing raters with extensive training and feedback to ensure consistency, providing PSTs with a

standardized lesson plan or learning objective, and observing PST instruction in the context of standardized teaching simulations (Cohen et al., 2020; Cohen & Goldhaber, 2016). Second, TPPs can isolate the differences between PSTs that are consistent across observations by averaging PST scores over multiple observations under a variety of conditions (Kane & Staiger, 2012). In this case, randomly assigning conditions, such as raters or lesson objectives, for each observation is especially helpful (van der Lans et al., 2016).

Sensitivity. Although sensitivity is a less-developed empirical concept in education research than reliability, it connotes the degree to which a measure can detect differences a TPP might want to understand. For example, which PSTs are most skilled, if some groups of PSTs are comparably more skilled, or if some aspects of instruction are comparatively stronger across all PSTs in a TPP (M. Allen & Coble, 2018; Hill et al., 2020; Mancenido, 2022). Data-driven decision-making is unlikely to lead to improvements in PSTs' preparedness if measures are not sufficiently sensitive to these differences. Identifying which PSTs need additional support implicitly requires comparing PSTs' scores. Understanding the impacts of specific preparation experiences requires comparing scores between the group of PSTs that completed an experience and those who did not. Similarly, identifying specific instructional skills for which PSTs need more support requires comparing scores from one skill to another. To accurately interpret these comparisons, the measures TPPs use must be sensitive enough to detect these differences (Cohen & Goldhaber, 2016; Weisberg et al., 2009). When measures are not sufficiently sensitive, TPPs risk concluding that there are no differences, when in fact the measures may be unable to identify them.

Measure sensitivity can be influenced by a variety of factors. First, the less reliable a measure is, the less sensitive it is likely to be. When only a small portion of the variation in scores reflects consistent differences between PSTs, then these differences can be drowned out by the other external factors (raters, lesson context, etc.) that influence scores. Scoring procedures and the range of scores available also influence sensitivity. A measure with only four possible scores provides coarser differentiation among PSTs' skills than a measure with 10 possible scores, assuming scores are reliably measured and PSTs receive the full range of scores (Kane & Staiger, 2012). Unfortunately, evidence suggests PSTs generally do not receive the full range of scores, with most scores clustering at the highest end (Kraft & Gilmour, 2017).

The scope and granularity of a measure can also influence sensitivity (Janssen et al., 2015). Measures that provide a broader picture of PSTs' skills, like "instructional support," tend to provide less detailed areas of strength or improvement than a fine-grained measure, such as the specific instances where PSTs engage in "feedback loops" with students (Hill et al., 2020; Mancenido, 2022). Because

of reliability issues and logistical constraints, providing finer-grained information necessarily requires trading off breadth for depth. It is not feasible, for example, to ask raters to individually score hundreds of finer-grained indicators of PST skill (Bell et al., 2015). If the main differences between PSTs, however, are more nuanced, for example based on the frequency and substance of their feedback loops, then broader measures of instructional support will be less sensitive because they encompass much more than granular measures.

Common measures of PSTs' instruction are likely to be limited in sensitivity because of these issues. In addition to suffering from low reliability, these measures generally provide a broad picture of PSTs' skills, providing little information about finer-grained details of PST development (Hill et al., 2020; Mancenido, 2022). In addition, these measures often use a limited score range of 3 to 5 points (Bartanen & Kwok, 2021; Henry et al., 2013). We are not aware of any literature that focuses on the sensitivity of PST observation scores. From Bartanen and Kwok's (2021) analyses, however, we can see that the difference in baseline scores and growth for a PST in the 16th percentile and 84th percentile (a 2 standard deviation [*SD*] difference) is statistically significant, suggesting some ability to detect differences between PSTs. Bastian et al.'s (2018) identification of four distinct PST profiles based on Educative Teacher Performance Assessment (edTPA) scores also suggests some ability to detect differences between PSTs, although we note edTPA is a portfolio assessment not an observational measure.

Data and Method

Context

We utilize data from a TPP at a large, public university in the southeastern United States, which we call Lambeth University. Lambeth offers multiple pathways for PST licensure, enrolling approximately 120 PSTs each year. PSTs at Lambeth are typically white, speak English as a first language, below 27 years, and attended high schools with students of middle or high socio-economic status (Table A3).

Lambeth takes a practice-based approach to preparing PSTs (Forzani, 2014). Coursework emphasizes the development of robust content and pedagogical knowledge, while explicitly linking theory to pedagogies of practice (Grossman et al., 2009). In addition to role-play exercises, many instructors provide opportunities for PSTs to teach digital student avatars in mixed-reality simulations, where PSTs can practice specific skills, techniques, and approaches before working with real children (Cohen et al., 2020; Dieker et al., 2014). This allows faculty to observe and assess PSTs' teaching in ways that are otherwise difficult to replicate in a university classroom.

PSTs also complete clinical experiences each semester focused on application of course content. Early clinical

experiences vary by program but range from one-on-one reading tutoring to interning in a classroom for 15 hours a week. Across programs, the final semester features a full-time student teaching placement designed to afford experience with all aspects of teaching. During each clinical experience, PSTs complete multiple coaching cycles, following a modified version of the evidence-based *My Teaching Partner* (MTP) program (J. Allen et al., 2015). In each cycle, PSTs record a lesson in their clinical placement. Field supervisors analyze video segments and provide PSTs with reflective prompts before meeting with the PST to discuss PSTs' strengths and areas for growth.

Measures

Lambeth dedicated substantial resources to a robust data collection system for the purposes of accreditation, research, and program improvement. Lambeth systematically collects data on PST instructional practice at multiple timepoints using the CLASS, designed to capture broad features of classroom climate and instructional support in their placement (Pianta & Hamre, 2009), and finer-grained measures of teaching in standardized mixed-reality simulations (Cohen et al., 2020). Lambeth uses CLASS data for accreditation purposes and to research PSTs' preparedness to teach. Unlike the CLASS, the simulation measures were designed to identify differences between PSTs who received instructional coaching between simulation sessions and those who self-reflect to discern which method of scaffolding better-supported PSTs' skill development (Cohen et al., 2020). To date, such measures have not been used to compare individual PSTs' skills. However, Lambeth is interested in using these measures moving forward to answer questions about PSTs' instructional skills, including growth over time, differences between groups of PSTs who have different learning experiences (e.g., courses with different instructors), and program-wide strengths and weaknesses.

For both measures, the university employs rating procedures that are aligned with "best practice" in the classroom observation literature. Raters have no relationship with PSTs they rate. They also complete formal trainings, are required to pass a rater certification test, and receive ongoing feedback at weekly scoring calibration meetings (Park et al., 2015). Finally, observation videos are randomly assigned to different raters at each observation timepoint (Kane & Staiger, 2012).

CLASS Scores. PSTs select up to four videos submitted for MTP coaching for external scoring. Videos are independently scored by Lambeth using the CLASS framework (Pianta & Hamre, 2009), an observation protocol emphasizing the value of positive relationships, focusing on the tenor of interactions between a teacher and students and among students, and treating the *classroom* as the unit of analysis. To date, the framework has been used for measuring

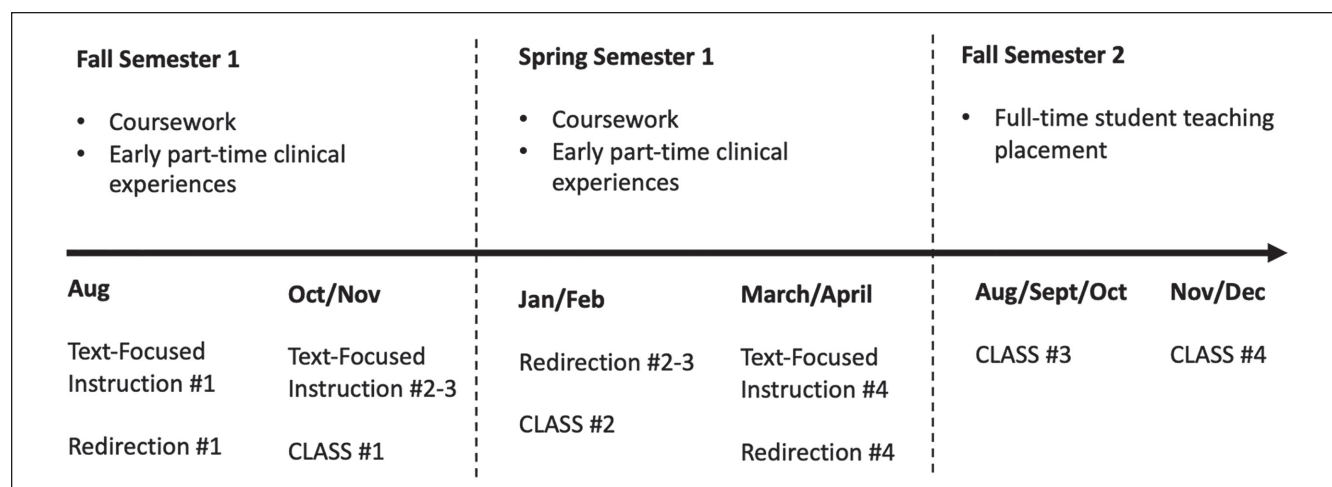


Figure 2. Data Collection Timeline for Each Simulation-Based Measure and CLASS Scores.

Note. This timeline reflects the collection of CLASS scores for the cohort that entered Lambeth in 2017 to 2018 and the collection of simulation scores for the cohort that entered Lambeth in 2018 to 2019 to be most reflective of our analytic sample. While the data collection timelines for other years were similar, there were some differences. For example, only the first two CLASS scores were collected in 2018 to 2019 and 2019 to 2020 because of logistical constraints and the coronavirus 2019 (COVID-19) pandemic, respectively. Similarly, 2017 to 2018 was a pilot year for the simulation-based measures that did not include the initial baseline data collection in August. In addition, the fourth simulation-based observations in March/April did not occur in 2019 to 2020 because of the COVID-19 pandemic. The timeline for coursework and clinical experiences, however, was similar across all 3 years, except for a small subset of PSTs who entered Lambeth in 2018 to 2019 or 2019 to 2020 and completed a 1-year MAT in secondary education. These students began coursework in the summer before the first fall semester, completed coursework and early part-time clinical experiences in the first fall semester, and completed a full-time student teaching placement in the first spring semester. For these students, the first CLASS observation was collected from their early clinical experience and the second was collected during their full-time student teaching placement. For other students who entered in 2018 to 2019 and 2019 to 2020, all available CLASS observations were collected during their early part-time clinical experiences. CLASS = Classroom Assessment Scoring System; PST = preservice teacher.

classroom quality for research purposes, guiding coaching conversations, and evaluating the quality of early child care programs (J. Allen et al., 2015; Araujo et al., 2016; Bassok et al., 2021). CLASS provides a high-level view of three broad domains: Emotional Support, Classroom Organization, and Instructional Support, each comprised of 3 to 5 dimensions (Hafen et al., 2015). For example, one dimension within the Classroom Organization domain is Productivity. One dimension within Instructional Support is Quality of Feedback. Raters score each dimension on a 7-point scale, using dimension-specific indicators to distinguish between teacher–student and student–student interactions that are low (1–2 points), mid (3–5 points), or high quality (6–7 points).

Prior research documents the reliability, sensitivity, and validity of CLASS when used to assess classroom quality for inservice teachers. When scores from multiple observations—each rated by a different rater—are averaged together, more than 60% of the variation in scores can be attributed to differences between teachers (Kane & Staiger, 2012). CLASS scores have also been used to detect differences between teachers and document changes in teaching over time (Bassok et al., 2021; Kane & Staiger, 2012; La Paro et al., 2004). Finally, higher CLASS scores are associated with other measures of classroom quality and stronger student outcomes, including academic performance and student engagement (e.g., J. Allen et al., 2013; Araujo et al., 2016;

Kane & Staiger, 2012; La Paro et al., 2004). However, a growing body of literature is raising concerns about the influence of contextual factors on CLASS scores and our ability to use these scores to explore change in teacher practice over time (reliability) and detect differences between teachers (sensitivity) (Briggs & Alzen, 2019; Casabianca et al., 2015; Gitomer et al., 2014; Wallace et al., 2020).

Simulation-Based Scores. Beginning in 2017 to 2018, all Lambeth PSTs also participated in two standardized simulations as part of their general method courses. In these simulations, PSTs teach virtual students voiced by a trained actor, providing an opportunity to engage in “approximations of practice” (Grossman et al., 2009). The simulations also serve as a standardized platform to assess skill development. In the first simulation scenario, referred to as “Redirection,” PSTs practiced redirecting standardized off-task student behaviors in a whole-class discussion (Cohen et al., 2020). In the second simulation scenario, referred to as “Text-Focused Instruction,” PSTs facilitated a discussion and responded to standardized student responses that were supported by textual evidence to different degrees (Cohen et al., 2023). PSTs completed each simulation scenario four times over the course of the program, though with different interval spacing than CLASS observations. Figure 2 illustrates the timeline for each simulation-based observation and when PSTs submitted

videos for CLASS scoring. As part of a larger research study, PSTs were randomly assigned to a short coaching session or a short self-reflection protocol in between simulation sessions. Roughly half the PSTs received coaching for the Redirection scenario and the other half for the Text-Focused Instruction scenario. For more detailed descriptions of the simulation measures and protocols, see Cohen et al. (2020).

The measures used to assess PSTs' instructional skill in the simulations are specific to each scenario. The rubric for the Redirection simulation was based on the Responsive Classroom framework that local K-12 schools used (Charney, 1993; Responsive Classroom, 2014). The rubric for the Text-Focused Instruction simulation was designed to reflect high-quality instructional practices in the relevant literature (Castles et al., 2018; Reznitskaya et al., 2009). The simulation rubrics assess more granular practices than the CLASS. Whereas the Redirection simulation rubric focuses entirely on behavioral redirection, for example, CLASS treats "effective redirection of misbehavior" as a component of the dimension of Behavior Management, which is, in turn, a component of the Classroom Organization domain. Nevertheless, there is conceptual overlap between the granular simulation rubrics and broad CLASS rubrics.

From a measurement standpoint, both the CLASS and simulation scores are criterion-referenced rubrics that require raters to select a specific score along a continuum from low to high. CLASS uses a slightly restricted score range of 1 to 7 compared with the simulation score range of 1 to 10. In addition, while raters select a single score for each simulation rubric, scores for Classroom Organization and Instructional Support are generated as an average of rater scores across multiple subdimensions.

Sample

For each measure, we selected a sample that maximized sample size, while also ensuring a similar number and timing of observations, to avoid bias from missing data. For CLASS scores, our sample consists of the 83 PSTs who entered Lambeth in 2017 to 2018 and for whom four observations are available, two from each of two clinical placements. In total, this sample includes 135 mentors¹ from 62 schools across more than 10 counties in one state. Schools are primarily elementary schools serving mostly students who are white and not eligible for free or reduced-price meals. For simulation scores, our sample consists of the 60 PSTs who entered Lambeth in 2018 to 2019 and for whom eight observations (four for each simulation scenario) are available. Appendix A includes a detailed discussion of the full dataset and rationale for sample restrictions. Though the samples for each measure are different, a covariate balance test (Table A3) suggests that there are few demographic differences between PSTs from the 2017 to 2018 cohort (CLASS analyses) and PSTs from the 2018 to 2019 cohort (simulation analyses).

Methods

Given our focus on cross-measure comparisons, we analyze CLASS scores for the two domains most closely aligned to the constructs measured by simulation rubrics: Classroom Organization (aligned with Redirection) and Instructional Support (aligned with Text-Focused Instruction). For each research question, we conduct analyses separately for each simulation scenario and CLASS domain. Below, we provide a general overview of our methods. Detailed explanations of the statistical models for each research question are included in Appendix B.

RQ1: To What Extent Does Each Measure Capture Consistent Differences Between PSTs?

We draw on generalizability theory (Bartanen & Kwok, 2020; Briggs & Alzen, 2019) to decompose variation in observation scores into distinct sources, with the goal of distinguishing between (a) variation that reflects consistent differences between PSTs, and (b) measurement error that results from differences in the conditions and context of measurement. To do so, we estimate PST observation scores as a function of the number of months since the first observation using a multilevel model with PST random effects. We then calculate the proportion of the variation that reflects consistent differences, relative to the overall variation in scores. Please see Appendix B for additional details.

RQ2 and RQ3: To What Extent Does Variation in Scores Reflect Differences Between Raters and Mentors?

We decompose variation in observation scores into contextual sources we hypothesize might influence scores without contributing to our understanding of PSTs' instruction. For CLASS scores, we separate consistent differences between PSTs from variation between mentors and any remaining measurement error. For simulation scores, we separate consistent differences between PSTs from variation between raters and any remaining measurement error. To do so, we augment the multilevel models employed in RQ1 to include mentor random effects (for CLASS scores) or rater random effects (for simulation scores). Please see Appendix B for additional details.

RQ4: How Well Can Each Measure Differentiate Between Individual PSTs, Groups of PSTs, and Facets of Instruction?

To explore the sensitivity of CLASS and simulation scores to individual differences, we review the estimates from the multilevel models used in RQ1 with the aim of understanding the extent to which we can detect between-PST

variation in scores and trajectories. Specifically, we use log-likelihood tests to compare models with and without PST intercept and slope random effects to determine whether there is significant between-PST variation in baseline scores and growth trajectories. To explore the sensitivity of scores to group differences, we review results from RQ1 for simulation scores, where a control for participation in coaching or self-reflection allows us to evaluate whether simulation scores are sufficiently sensitive to detect differences between PSTs exposed to these preparation experiences. To explore sensitivity to different facets of instruction, we graphically compare scores on the two CLASS domains and scores for the two simulation-based measures. We also estimate the correlations between scores, comparing the two CLASS domains (Classroom Organization and Instructional Support) to each other and the two simulation-based measures to each other (Redirection and Text-Focused Instruction). Please see Appendix B for additional details.

Limitations

Lambeth is not representative of all TPPs. In particular, Lambeth PSTs are largely white, female, and have college-educated parents. We see our analyses as proof of measurement issues that can surface across a range of TPP contexts, but we are not making any claims here about the degree to which these findings are generalizable. There are also several limitations of the measures we employ, which necessarily reflect specific conceptualizations of core components of high-quality instruction. Many crucial aspects of teaching, including culturally and linguistically responsive and sustaining pedagogies, are not included in these measures (Pacheco, 2009; Paris & Alim, 2014). We also do not yet know the extent to which CLASS scores from clinical placements or simulation-based scores provide meaningful information about PSTs' instruction when they become teachers of record. Though CLASS scores have been used to evaluate PSTs' instructional skills during student teaching (e.g., Malmberg et al., 2010), we are not aware of any studies that evaluate the predictive validity of CLASS scores for instructional skills down the road. Such evidence is also lacking for simulation scores, though such research is currently underway. Unfortunately, generating this evidence using data from Lambeth is virtually impossible because of the lack of state-wide longitudinal data systems that connect PSTs to their later employment and teaching outcomes. Nonetheless, these kinds of predictive validity analyses are only possible when measures primarily reflect differences between PSTs. When measures are heavily influenced by contextual characteristics, such as raters or mentors, they cannot tell us much about differences between PSTs' skills, let alone predict what PSTs will do when they begin full-time teaching. We, therefore, argue that our analyses raise important considerations for TPPs and provide a proof-of-concept relevant to any

measure of PSTs' skills, including those that capture other aspects of teaching.

Our analyses are also limited by unavoidable deviations from the ideal design for a generalizability study, which requires each PST be observed by every rater, in every kind of classroom context, and teaching every kind of lesson (Briggs & Alzen, 2019). Under these conditions, we could directly measure how each contextual characteristic influences scores. Random assignment of rater, classroom context, and lesson provide a more feasible alternative to estimate the average effects of each contextual characteristic. Because mentors are not randomly assigned to PSTs in our data, effects attributed to mentors may instead reflect differences between PSTs. This would be the case if PSTs with weaker instructional skills tend to be intentionally assigned to more skilled mentors. The estimated between-mentor variation in scores would then reflect both these initial differences in PSTs' instructional skills and the influence of mentors on PST scores.

There are two ways mentors may influence PST scores, each of which has different implications for the interpretation of our results. First, mentors may influence PST scores by supporting PST skill development. Indeed, a growing body of literature highlights the important role mentors play in supporting PSTs' instructional skills as teachers of record and feelings of preparedness (Bastian et al., 2022, 2023; Goldhaber, Krieg, Naito, & Theobald, 2020; Goldhaber, Krieg, & Theobald, 2020; Ronfeldt et al., 2018, 2020, 2021). In this case, differences between mentors' ability to support PST skill development would be included in our estimates of between-mentor variation in scores, even though they conceptually reflect "true" differences between PSTs that TPPs would want to discern. At the same time, mentors may also influence PST scores in ways that do not reflect PSTs' instructional skills. This would be the case, for example, if PSTs earn higher Classroom Organization scores when observed in classrooms taught by mentors with robust classroom routines the PST neither established nor maintained. Conceptually, we would want to remove this kind of variation in scores before drawing conclusions about between-PST differences in management skills. Unfortunately, we are unable to separate these two kinds of mentor influences in our data. One way this has been done in prior work involves estimating reliability separately at different timepoints to capitalize on the expectation that mentors' developmental effects increase over time, the desired "signal," while their indirect effects on PSTs' scores, the contextual "noise," would likely decrease as PSTs take over more responsibility in the classroom (Bartanen & Kwok, 2021). However, our data do not allow us to estimate reliability for individual timepoints or separately by semester.² Rather than providing precise estimates of the proportion of variation in CLASS scores attributable to true differences in PSTs' instructional skills, our results provide an initial estimated range of variation, while clearly problematizing the

Table 1. Proportion of the Variation That Reflects Consistent Differences Between PSTs Relative to the Overall Variation in Scores⁴.

Model	CLASS: Instructional Support	CLASS: Classroom Organization	SIM: Redirection	SIM: Text-Focused Instruction
Overall variation, assuming no growth in scores over time	0.07	0.03	0.17	0.11
Overall variation, allowing for growth in scores over time	0.04	0.03	0.22	0.21
Variation in baseline scores	0.04			0.21
Variation in growth over time	0.15			0.15

Note. Estimates are calculated using Equations 1 to 4 described in Appendix B. Estimates for fixed effects, random effects, and variance components from these equations can be found in Tables C1 to C3. CLASS = Classroom Assessment Scoring System; SIM = simulation-based scores.

use of raw scores in drawing inferences about differences in PST skill.

Our analyses are also limited by the lack of access to data on all contextual characteristics of interest. While we have access to rater information for simulation scores, we do not have access to rater information for CLASS scores. Any attempt to decompose the variation in CLASS scores, therefore, suffers from the problem that not all potential contextual factors are accounted for in the model. This means that our estimates serve as an upper bound, as the true proportions would be lower if any unobserved factors influence scores (Briggs & Alzen, 2019).

Finally, we acknowledge several challenges with making direct comparisons between CLASS and simulation results. First, we drew on different samples, raising the possibility that observed differences in the scores' measurement properties could stem from sample rather than measure differences. However, we see little evidence of systematic differences between the samples on observable characteristics (Table A3). Moreover, the rating and scoring procedures were highly standardized and consistent across both AY2017 to 2018 and AY2018 to 2019. In addition, we cannot rule out other potential differences between simulation and CLASS scores including characteristics of the rubrics themselves or unknown differences in the implementation of scoring procedures, including CLASS raters, that could theoretically influence their measurement properties. However, our goal is not to make conclusive claims about the cause of measurement issues we observe. Instead, we are surfacing factors that are within TPPs' control and that prior literature has suggested might play a role, including raters and mentors.

Results

RQ1: To What Extent Does Each Measure Capture Consistent Differences Between PSTs?

The proportion of variation in CLASS scores that reflects differences between PSTs is low relative to the overall variation in scores. Specifically, we estimate that 3% to 4% of the variation in individual Instructional Support and

Classroom Organization scores reflects consistent differences between PSTs (Table 1). This means that 96% to 97% of the variation in scores reflects measurement error. Figure 3 shows the remaining variation in scores that reflects differences between PSTs, after we account for measurement error. It is very limited, especially for Classroom Organization scores.³ In the second row of Figure 3, we plot the growth trajectories for all PSTs to illustrate how baseline scores and growth vary over time. Here, the y-axis reflects each PSTs' simulation or CLASS score, and the x-axis reflects the observation timepoint. We find 15% of the variation in growth in Instructional Support scores reflects consistent differences between PSTs over time (Table 1). The estimate for growth is higher because PST growth is calculated using all four scores, each of which contains signal of PST skill, while the proportion for individual scores is calculated using only a single score.

Considering the many strengths of the CLASS measure and our first hypothesis that CLASS scores would be more reliable than the home-grown measures documented in prior literature, our results are surprising and dismaying. Bartanen and Kwok (2021), for example, found that 20% of the variation in scores reflected consistent differences between PSTs even though their data did not rely on an established observational rubric and scores were generated by supervisors that knew the PSTs they rated and received much less robust rater training.

We also find a greater proportion of the variation in simulation scores reflects consistent differences between PSTs. Specifically, about 20% of the variation in scores reflects consistent differences between PSTs, as compared with the 3% to 4% for CLASS scores (Table 1), about a five-fold difference. When compared with the prior literature, these results are somewhat more encouraging. In addition to being consistent with Bartanen and Kwok's (2021) findings, these results are consistent with prior work on the reliability of inservice teacher observation scores under conditions more favorable for reliability⁵ (Briggs & Alzen, 2019; Ho & Kane, 2013; Kane & Staiger, 2012).

The estimates for variation in growth over time are similar for CLASS and simulation scores (Table 1).

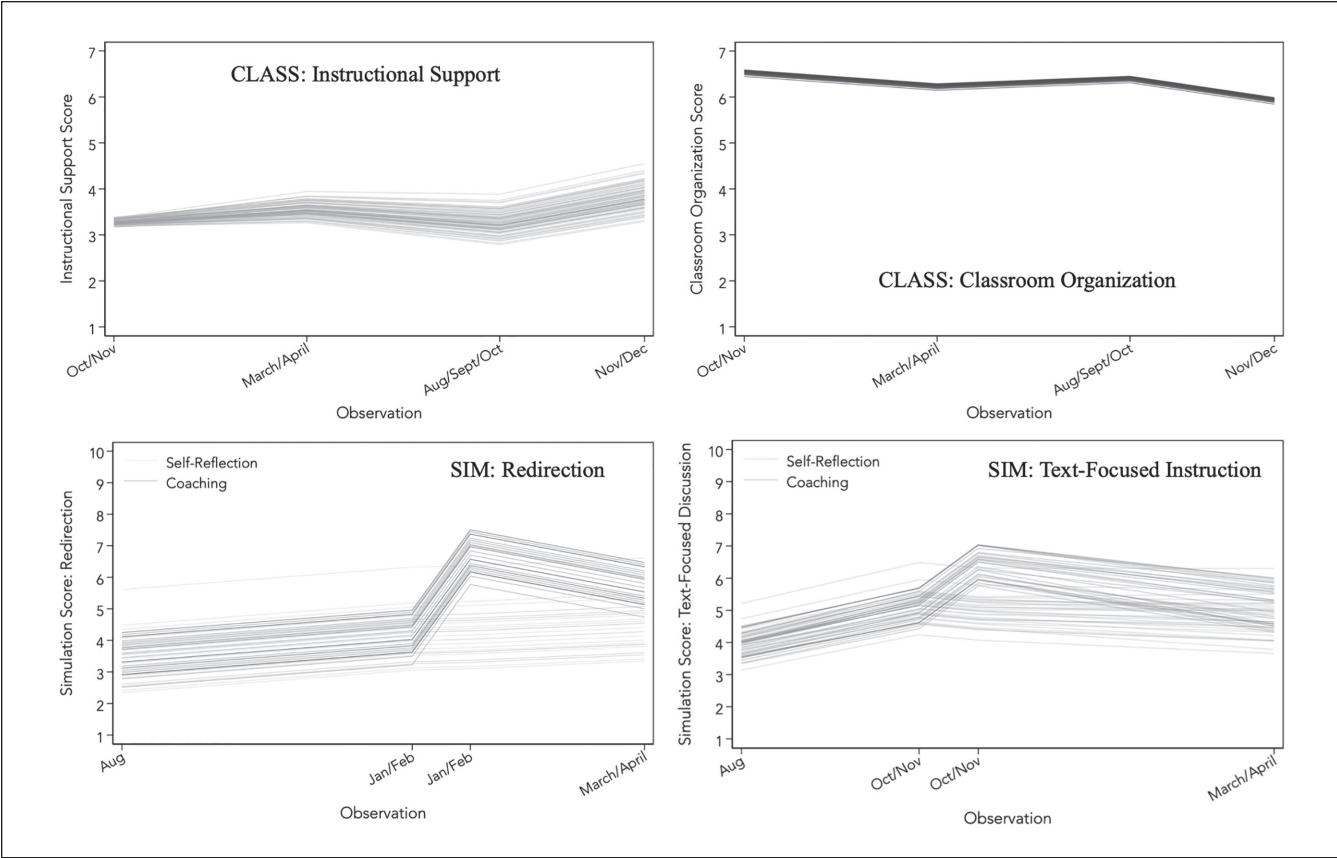


Figure 3. Predicted PST Scores After Accounting for Measurement Error to Isolate Between-PST Variation in Scores
Note. PST = preservice teacher; CLASS = Classroom Assessment Scoring System; SIM = simulation-based scores.

Directly comparing these two estimates is misleading, however, since the growth estimates are scaled by the intervals between observations, which are far greater for CLASS scores (see Appendix B). To allow more direct comparison, we analyze the proportion of variation in growth if CLASS scores and simulation scores were collected at the same intervals, by recalculating the estimate for simulation scores using the scaling factor for CLASS scores. If simulation scores were separated by the same intervals as CLASS scores, about 40% of the variation in growth on Text-Focused Instruction scores would reflect consistent differences between PSTs. This is about two and a half times larger than the estimate for CLASS Instructional Support.

RQ2: To What Extent Does Variation in CLASS Scores Reflect Differences Between External Factors, Namely Mentor Teachers?

When our models isolate between-mentor score variation, we find that 9% to 17% of the variation in CLASS scores can be explained by the mentor in whose classroom a PST is observed (Table 2).⁶ Once we account for variation between

mentors, the proportion of variation that reflects differences between PSTs falls to effectively zero. These results differ from Bartanen and Kwok's (2021) analysis, where they find little evidence mentors contributed to variation in scores but are consistent with our second hypothesis that CLASS measures would likely capture mentor effects.

A key question here is how to interpret our results. Some between-mentor variation may reflect differences in the classroom context (e.g., classroom routines previously established by the mentor) unrelated to differences in PSTs' instructional skills. At the same time, mentors may also differ in their ability to support PST skill development, in which case between-mentor variation in scores may reflect true differences in PST skill. In addition, because mentors were not randomly assigned to PSTs, between-mentor variation may also reflect nonrandom sorting of more skilled PSTs to certain mentors. In the latter two cases, removing between-mentor variation in scores would understate the true differences between PSTs.

While we cannot fully disentangle these sources of between-mentor score variation, we can engage in a bounding exercise. In the RQ1 models, where we don't account for between-mentor variation, any variation in scores stemming

Table 2. Proportion of Score Variation That Reflects Consistent Differences Between PSTs and Proportion of Score Variation That Reflects Consistent Differences Between Mentors (for CLASS Scores) and Raters (for Simulation Scores).

Estimate	CLASS scores		Simulation-based scores	
	Instructional Support	Classroom Organization	Redirection	Text-Focused Instruction
PST baseline	0.0002	0.00	0.23	0.21
PST growth	0.0004	n/a	n/a	0.15
Rater/Mentor	0.09	n/a	0.03	0.01
Sources of variation	Mentors	Mentors	Raters	Raters
	PSTs	PSTs	PSTs	PSTs
	Residual error	Residual error	Residual error	Residual error

Note. Estimates are calculated following the methodology described in Appendix B under RQ2 and RQ3. CLASS = Classroom Assessment Scoring System; PST = preservice teacher.

from mentors' developmental effects or PST sorting are absorbed as part of the estimated between-PST variation in scores. The estimates from Table 1, therefore, provide an upper bound for how much variation in PST scores may stem from PST sorting, mentors' developmental effects, and other differences in PSTs' instructional skills. When we decompose variation in scores into only between-mentor variation and measurement error, however, we find that 15% to 17% of the variation in scores can be explained by mentors alone, more than three times the upper bound provided by the Table 1 estimates (3%–4%). This discrepancy suggests at least some of variation in raw PST scores stems from differences between the mentor context, above and beyond the contributions from PST sorting, mentor developmental effects, and PSTs' underlying differences in skill. In other words, observing PSTs in different mentor classrooms appears to contribute to the large proportion of the variation in CLASS scores indicated as measurement error (96%–97%) in our initial RQ1 analyses.

RQ3: To What Extent Does Variation in Simulation Scores Reflect Differences in External Factors, Namely Raters?

Unlike mentors, who are not randomly assigned to candidates, raters *were* randomly assigned to observe specific PSTs at specific timepoints. Thus, we are confident that between-rater variation in scores does not result from systematic sorting of PSTs to raters (Kane & Staiger, 2012). We also have no reason to believe that raters, who do not know or interact with these PSTs, would have any developmental effects on PSTs' instructional skills like the supervisors who often score PSTs (Bartanen & Kwok, 2021). Therefore, we can interpret any observed between-rater variation in scores as stemming from differences in rater standards, which TPPs would want to account for before making inferences about PSTs' skills. In line with our third hypothesis, raters explain only 1% to 3% of score variation (Table 2) after separating out between-rater variation in scores. This

suggests differences in simulation scores between PSTs are not substantially influenced by differences in rater standards, in contrast to the influence of supervisors on PSTs' ratings in Bartanen and Kwok's (2021) work.

RQ4: How Well Can Each Measure Differentiate Between Individual PSTs, Groups of PSTs, and Facets of Instruction?

Sensitivity to Differences Between PSTs. The lower the proportion of variation in scores that reflects differences between PSTs, the less likely it is for a measure to be sensitive enough to detect those differences. Our fourth hypothesis was the hoped-for higher reliability of CLASS, and simulation scores would correspond with enhanced sensitivity of those measures. Though reliability for neither measure was ideal, we return to Figure 3 to understand the differences that CLASS and simulation scores can detect between PSTs, isolating only between-PST differences after accounting for measurement error, which is the most conservative approach. Results from additional analyses to quantify these differences and evaluate their statistical significance are included in Appendix D.

Differences between PSTs are small for CLASS scores. The difference in baseline CLASS scores between a PST at the 16th percentile and a PST at the 84th percentile (2 standard deviations) corresponds to 0.2 to 0.3 points (out of 7) for both Instructional Support and Classroom Organization (Table D1). This difference is only statistically significant for Instructional Support. There is also a statistically significant difference in growth rate for Instructional Support, corresponding to a difference of 0.03 points between a PST at the 16th percentile and a PST at the 84th percentile. Relative to the 7-point CLASS scale, these differences represent at most 5% of the maximum possible difference between PSTs.

Consistent with the greater reliability of simulation scores, differences between PSTs on simulation scores are statistically significant and larger. The difference in baseline simulation scores between a PST at the 16th percentile and 84th percentile is 1.11 points for Text-Focused Instruction and

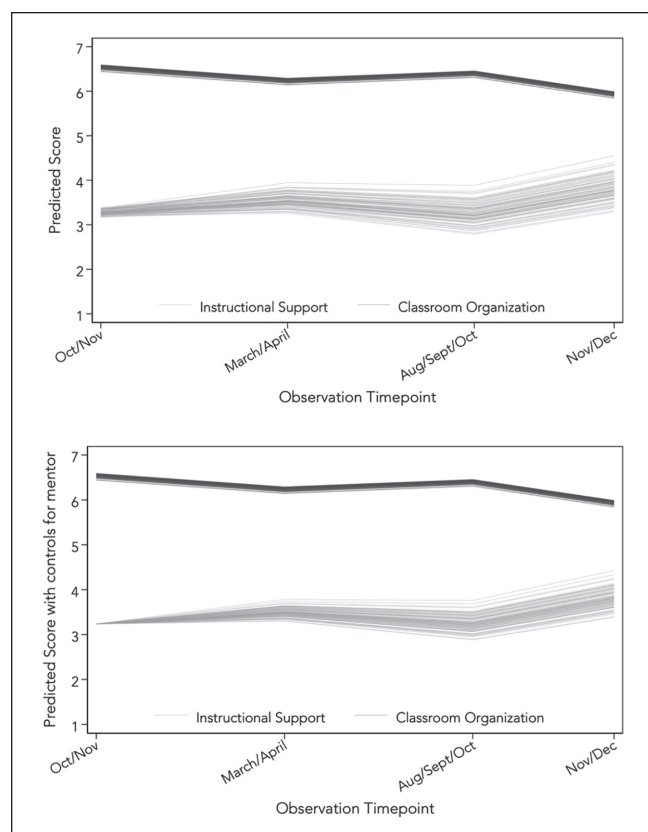


Figure 4. Comparison Between Instructional Support and Classroom Organization Scores Using Predicted Scores.

1.71 for Redirection (Table D2). Relative to the 10-point simulation scale, this represents 12% to 19% of the maximum possible difference in scores, more than double the estimate for CLASS scores. These differences are similar in magnitude to the differences Bastian et al. (2018) observed, on average, between PSTs on the edTPA.⁷ We acknowledge the possibility these results stem from PSTs having similar instructional skills, rather than a lack of measurement sensitivity. However, anecdotal evidence from Lambeth teacher educators and prior work documenting large differences between PSTs once they enter the classroom suggest this is unlikely to be the case (e.g., Atteberry et al., 2015; Boyd et al., 2008).

Sensitivity to Differences Between Groups-Simulation Scores. To explore score measurement sensitivity to different learning experiences, we compare simulation scores between PSTs who participated in coaching versus self-reflection. PSTs who received coaching for the Text-Focused Instruction simulation score 1.5 points higher immediately after coaching and 0.5 points higher when observed 5 months later, though the latter estimate is not significant in some models (Table C3). PSTs who received coaching for the Redirection simulation score 2.5 points higher immediately after coaching and 1.3 points higher 2 months later.

These results highlight that simulation-based measures are sensitive enough to detect differences between PSTs who participated in different preparation experiences immediately following such experiences and over time. In addition, these results reinforce the feasibility of comparing groups of PSTs, rather than individual PSTs, even when a large portion of the variation in scores reflects measurement error. If measurement error and the influence of contextual characteristics are the same across groups of PSTs, we can safely make comparisons across those groups.

Sensitivity to Differences Between Facets of Instruction—CLASS and Simulation Scores. Here, we compare PSTs' scores between Classroom Organization and Instructional Support to explore sensitivity to differences between facets of instruction. In Figure 4, using predicted scores to account for measurement error, we graph scores over time for all PSTs to see whether they perform similarly across CLASS domains. Across both graphs, Classroom Organization scores are consistently higher than Instructional Support scores, with a difference of 1.00 to 3.25 points out of 7. This suggests PSTs in our sample are, on average, considerably stronger in management skills than providing instructional support. These results provide suggestive evidence that CLASS scores are sufficiently sensitive to identify program-wide patterns in areas of relative strength.

Table 3 shows that scores between domains are moderately correlated with one another (0.20–0.45). This means PSTs that receive higher scores on Instructional Support also tend to receive higher scores on Classroom Organization. This does not necessarily contradict the previous findings. However, it raises the possibility that there may be differences between PSTs that CLASS scores are not able to detect. This would be the case, for example, if raters perceive Instructional Support as more challenging—resulting in lower average scores—but also form a general impression of each PSTs' overall skills (i.e., halo effects), rather than considering each domain individually (Cohen & Goldhaber, 2016). Under these conditions, PSTs with relatively higher Instructional Support scores could potentially be given high Classroom Organization scores, even if their underlying management skills were weaker. Unfortunately, we cannot determine whether this is the case at Lambeth. It is equally possible that PSTs at Lambeth simply have similar relative strengths and areas of improvement.

Results for simulation scores follow a different pattern. In Figure 5, we graph Redirection and Text-Focused Instruction scores over time for all PSTs and see more overlap in scores across the two facets of instruction within each group of PSTs (coaching vs. self-reflection). This suggests simulation scores may not be very helpful in understanding program-wide patterns in PSTs' relative strengths and areas of improvement. At the same time, correlations between Redirection and Text-Focused Instruction scores are quite weak (Table 3), providing empirical evidence that each

Table 3. Spearman's Rank Correlations Comparing Predicted CLASS and Simulation Scores Across the Different Facets of Instruction.

Timepoint	CLASS scores		Simulation scores	
	Without controls for mentor	With controls for mentor	Without controls for rater	With controls for rater
Observation 0	0.45***	0.45***	0.25*	0.25*
Observation 1	0.44***	0.41***	0.24*	0.24*
Observation 2	0.43***	0.41***	-0.14	-0.14
Observation 3	0.42***	0.41***	-0.01	-0.01

Note. CLASS Observations 0 to 4 are scored on both Classroom Organization and Instructional Support. Sim Observations 0 to 1 are scored on Redirection and 2 to 3 on Text-Focused Instruction. CLASS = Classroom Assessment Scoring System.

* $p < .1$. ** $p < .05$. *** $p < .01$.

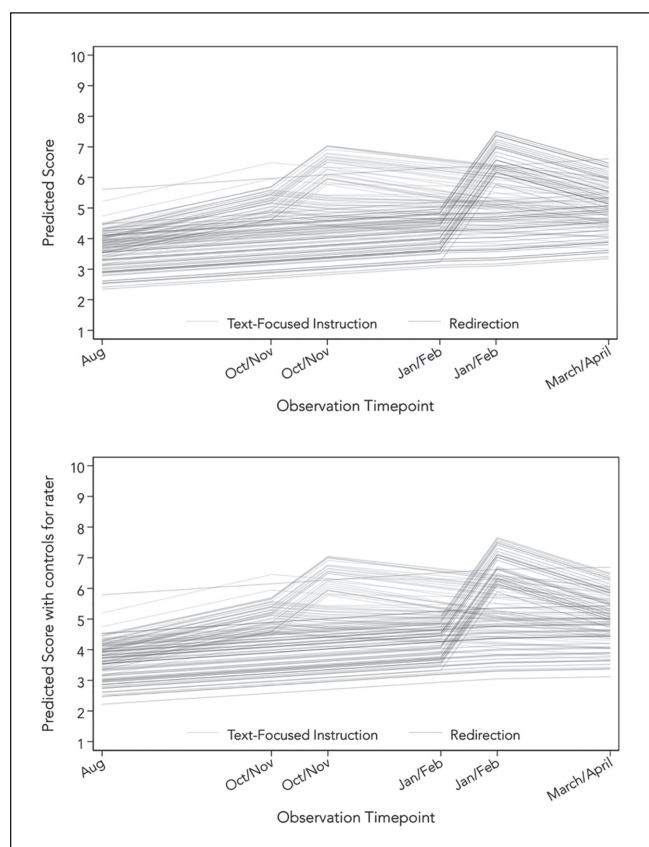


Figure 5. Comparison Between Text-Focused Instruction and Redirection Scores Using Predicted Scores. Both Models Include Controls for Coaching Treatment Status.

measure captures a different facet of PSTs' instructional skills and suggesting simulation scores can help illuminate relative strengths and areas of improvement for individual PSTs.

Discussion and Implications

In many ways, Lambeth University is at the forefront of measuring PSTs' instructional skills. Instead of relying on

problematic supervisor ratings (Bartanen & Kwok, 2021), the university employs trained and certified raters who do not know the PSTs they rate, resulting in high inter-rater reliability. Furthermore, raters are randomly assigned to videos to avoid systematic bias in scores from differences in rater standards. The university also draws on a validated observation protocol (CLASS) and innovative researcher-developed, simulation-based measures that allow the university to standardize the lesson context in which a PST is observed. The university's data collection procedures are well-aligned with established best practices for inservice teacher observations (Ho & Kane, 2013), representing substantial improvements over typical approaches to measuring PST practice (Bartanen & Kwok, 2021; Mancenido, 2022).

Nevertheless, our results raise serious concerns about the conclusions that can be drawn from Lambeth's carefully collected data. At most, only 20% of the variation in CLASS or simulation scores represents consistent differences between individual PSTs. At worst, effectively none of the variation in scores represents consistent differences between PSTs. Our results reinforce the challenges raised by Bartanen and Kwok (2021) and suggest they might extend beyond supervisor ratings. Drawing conclusions about individual PSTs' skills based on these scores is risky. When scores stem not from instructional skills but instead from statistical error and/or the systematic conditions under which an observation was conducted (e.g., working with a particular mentor), it would be misguided to use them to make consequential decisions about specific supports for PSTs or an individual's readiness to enter the classroom.

What then are the implications of our results for TPPs? If observational measures of PSTs' teaching provide little information about differences between PSTs, then one possible implication is that we should stop collecting such data. Doing so, however, leaves TPPs to "fly blind" when it comes to understanding how the experiences TPPs provide help PSTs develop their instructional skills. Instead, our results highlight several promising strategies that TPPs should consider when implementing such measures.

First, our results suggest the benefit of supplementing global assessments of "instructional quality" with

finer-grained measures of finer-grained aspects of teaching. With broad measures like CLASS, differences between PSTs in their use of “feedback loops” (a behavioral indicator in the dimension of “Quality of Feedback”) for example, are likely to be drowned out by the myriad skills captured in a domain score for Instructional Support. Simulation-based measures of fine-grained skills might be better able to detect differences between individual PSTs than CLASS scores. TPPs can also provide more nuanced feedback to PSTs about finer-grained aspects of teaching like providing timely behavioral redirections than the far more amorphous and multifaceted CLASS dimension of “Positive Climate” let alone the domain score for Emotional Support (Hill & Grossman, 2013; Wylie, 2020).

Second, TPPs can consider strategies for reducing the influence of contextual conditions on PSTs’ scores. The higher reliability of Lambeth’s simulation scores provides suggestive evidence that raters randomly assigned to observations have little systematic influence on PST scores. This greatly reduces the likelihood that observed differences in scores or growth between PSTs stem from differences in rater standards, a serious issue when PSTs are rated by a single supervisor or mentor (Bartanen & Kwok, 2021).

None of this is to say that supervisors or mentors should not *also* assess and support PSTs, using their knowledge of the context in which a PST is working to inform their understanding of a PST’s capacities, contingent on the students with whom they are working. Indeed, mentors and supervisors may be better able to provide formative, context-specific feedback and support for PSTs’ skill development if they do not also rate PST performance (Papay, 2012). Instead, we argue for a system that would decouple context-specific support systems for PSTs from systems designed to provide precise and actionable insights to TPPs. This approach is costly, however. Lambeth invested in having PSTs video-record their observations and paying additional personnel to score each observation, instead of leveraging supervisor or mentor ratings. Alternatively, TPPs could ask supervisors to collect videos and then randomly assign supervisors to videos. Moreover, we see these costs diminishing dramatically in future years, given rapid advancements to video capture, automated transcription, and even automated scoring of teaching using the textual data culled from such transcripts (Anglin et al., 2021). While it requires up-front investments, minimizing the impact of raters on PSTs’ scores could pay dividends in terms of the accuracy of information gleaned from these observations.

TPPs can also standardize the conditions of observations as an alternative to randomization, reducing the likelihood that differences between scores reflect differences between circumstantial conditions. Our results comparing scores from standardized simulations with CLASS scores provide suggestive evidence of the value of this approach. Though simulations are, by design, artificial and do not reflect the full complexity of the classroom, our results suggest that

they can be helpful for identifying differences in how candidates enact their knowledge and skills when required to face common problems of practice. Importantly, PSTs complete simulations without the aid of a mentor. While we recognize the importance of observing how PSTs enact their knowledge and skills in real classrooms, we argue that it’s equally important for some of the measures we use to capture what PSTs can do on their own, instead of what they do when assisted by mentors.

In practice, clinical observations cannot feasibly be standardized to the same extent as simulations, and so are likely to continue to suffer from more severe reliability and sensitivity issues. Our results indicate these challenges impact multiple CLASS domains. Standardizing would require observing PSTs teach the same lesson within the same mentor’s classroom, a logistically infeasible and disruptive approach for PSTs, mentors, and students. Randomization may be logistically more feasible but would limit TPPs’ ability to intentionally match PSTs with specific geographic areas, grade levels, content areas, school contexts, or mentor characteristics.⁸

To be clear, our results are about highlighting tradeoffs with these measurement decisions. We do not suggest designing observational systems that produce highly reliable and sensitive data at the expense of PSTs’ learning in clinical experiences and working alongside skilled mentors (Goldhaber et al., 2022; Ronfeldt, 2012). In addition to reflecting more authentic instructional contexts, observations during clinical placements can use existing TPP systems, whereas simulation-based measures require substantial up-front technology and infrastructure investments. Instead, our argument is that TPPs need to understand the affordances and constraints of the different measures they use, and ideally build a suite of measures with distinct strengths for distinct purposes.

Our results also make clear that mentors matter not just for PSTs’ learning (Goldhaber et al., 2022; Ronfeldt et al., 2018), but also for TPPs’ understanding of PSTs’ development. More work was needed to understand how a mentor’s classroom influences PST scores and to develop strategies for TPPs to minimize this influence. Indeed, we are not aware of any other work that highlights this issue or provides potential solutions. At a minimum, TPPs should consider implementing some additional standardization, such as a set of standardized instructional activities to complete at set times during clinical placements (e.g., facilitating a discussion about a word problem or orchestrating an analysis of a historical text). TPPs should also consider exploring new or modified measures whose indicators focus squarely on PST practices to minimize the potential for mentor effects. Measures focused on *classroom* quality that attend to the actions of all adults in the room (like the CLASS) might be less useful in the preservice context.

Finally, we highlight two potentially helpful strategies for addressing issues of reliability and sensitivity for clinical observations and simulation-based measures. First, our results suggest estimates of PST growth over time will likely

be more reliable than a score at a single point in time, because growth estimates necessarily incorporate multiple observation scores. This suggests TPPs would be well-served by collecting multiple scores of PST skill over the course of a program, not just at the end. It also suggests the value of TPPs using measures of growth rather than end-of-program average scores for accreditation and program approval purposes to provide a clearer and more accurate sense of how the TPP supports PST skill development.

Second, statistical adjustments using multilevel models can theoretically correct for low reliability and isolate consistent differences between PSTs. In practice, however, this kind of statistical adjustment requires substantial methodological expertise and represents a nontrivial departure from the status quo. Currently, most TPPs use raw, unadjusted scores for internal data analysis purposes, external reporting to accreditation bodies, and communicating results to PSTs and other stakeholders. Shifting from raw scores to statistically adjusted scores would require TPPs to develop the capacity to conduct such analyses, interpret adjusted scores in their own internal decision-making, and support PSTs and other stakeholders with understanding and interpreting these scores.

TPPs face complex decisions and tradeoffs in managing their PST data systems. Systematic collection of PST data requires substantial resources and time investment on the part of TPPs and teacher educators. Altering these systems to address the measurement challenges highlighted here requires even more time and resources, especially when new technologies or more intensive data collection efforts are required. In addition, TPPs must navigate potential tensions between improving the reliability and sensitivity of PST data and ensuring preparation experiences continue to support PST learning. Standardizing lesson objectives, for example, may improve reliability, but is also challenging when PSTs work in a wide range of grade levels and school contexts.

While our results raise serious concerns about using observational measures of PSTs' instructional practice to draw conclusions about individual PSTs' instructional skills, they also highlight other kinds of inferences we can make, even in the absence of high reliability. The simulation scores, for example, can successfully detect differences between groups of PSTs who had different learning experiences, such as coaching or self-reflection supports. They are sensitive enough to detect these differences immediately after those experiences, as well as months down the road. For programs trying out new courses, clinical experiences, or other kinds of learning supports, such measures could be invaluable in determining the degree to which such programmatic shifts are associated with corresponding shifts in PSTs' skills. CLASS scores are sensitive enough to detect program-wide patterns in PSTs' relative strengths and weaknesses across multiple facets of instruction, in this case, considerably stronger Classroom Organization skills than Instructional Support skills, though we caution both measures also capture a good deal of information about PSTs' mentors, too.

TPPs must decide what data will support the inferences they wish to make. Measures that can detect differences between groups of PSTs, for example, may look different than measures that can detect differences between individual PSTs. Obtaining reliable estimates of PST growth over time requires a different observation schedule than obtaining reliable estimates of PSTs' skills at specific moments. Prior work on inservice teacher observations suggests that expecting a single measure to serve several different purposes is unwise (Hill & Grossman, 2013; Papay, 2012). Instead, TPPs should use distinct measures to draw distinct conclusions about PSTs. This means that TPPs must be crystal clear about what conclusions they wish to draw when making decisions about what measure(s) to use, especially if logistical and financial constraints prevent the use of multiple observational measures.

TPPs stand to reap large benefits when they develop systems that enable data-driven programmatic decision-making. TPPs can better decide how to allocate limited resources to support the PSTs and areas of instructional practice with the greatest need, as well as how to evaluate the effects of specific preparation experiences on PST learning and skill development. However, the details of the data matter if we want TPPs to improve the quality of teacher preparation and not just complete a compliance exercise for program accreditation. This is especially true in the preservice period where PSTs are likely to exhibit smaller differences in skill than inservice teachers with a range of experience and when contextual factors, such as mentors, may influence assessments of PSTs' skills. To ensure that all TPP graduates enter the classroom ready to do the important and complex work of supporting K-12 students, we ultimately need PST data systems that are (a) reliable enough to identify consistent differences between individual PSTs, (b) sensitive enough to detect differences between PSTs and facets of instruction, and (c) allow TPPs to generate empirically backed conclusions about PST learning and development. Our findings suggest we need considerably more work on all fronts.

Acknowledgements

The authors wish to thank the members of the TeachSIM lab at the University of Virginia for their invaluable support in collecting these data. We also wish to thank Brendan Bartanen for feedback on initial results and methodological approaches. We are also grateful for feedback on earlier versions of this manuscript from Vivian Wong, Allison Atteberry, and Beth Schueler. All errors are those of the authors.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: The research reported here was supported by the Institute of Education

Sciences, U.S. Department of Education, through Grant #R305B140026 and Grant #R305D190043 to the Rectors and Visitors of the University of Virginia; the National Academy of Education/Spencer Foundation post-doctoral fellowship; the Jefferson Trust through Grant #DR02951; and the Bankard Fund through Grant #ER00562. The opinions expressed here are those of the authors and do not represent views of the Institute of Education Sciences, the U.S. Department of Education, or any other funding body.

ORCID iDs

Arielle Boguslav  <https://orcid.org/0000-0002-4055-5743>

Julie Cohen  <https://orcid.org/0000-0001-6794-7993>

Supplemental Material

Supplemental material for this article is available online.

Notes

1. Our sample includes more mentors than PSTs, because PSTs work with one mentor during early clinical experiences and a different mentor for their formal student teaching experience.
2. We cannot estimate reliability at a single point in time because most mentor teachers only work with one PST at a single time-point, so we cannot separate out between-mentor variation. We can only estimate between-mentor variation in scores when using data from both semesters together, because PSTs change mentors from the first semester to the second. However, this does not allow us to estimate reliability separately by semester.
3. Figure 3 also illustrates some surprising trends in PSTs' instructional practice over time. In particular, Classroom Organization scores appear to decrease over time. Unfortunately, we do not know why this is the case. We also see evidence that PSTs' Instructional Support scores dip between the spring of their first year and the beginning of the second year, which we hypothesize might result from the summer break away from the classroom.
4. For Classroom Organization and Redirection scores, all PSTs effectively grow at the same rate, so we cannot estimate what proportion of this "growth" variation reflects consistent differences between PSTs.
5. The data used included double the number of observations per teacher, and closer spacing between observations. Furthermore, because the data come from observations of inservice teachers there are no concerns about measurement error stemming from mentor teachers.
6. Because of challenges estimating the proportion of variation in scores between mentors for Classroom Organization, we also estimate an alternative model where we only consider differences between mentors, ignoring any additional variation that stems from differences between PSTs. Under this model, we find that 15% of the variation in Instructional Support scores reflects consistent differences between mentors. Similarly, we find that 17% of the variation in Classroom Organization scores reflects consistent differences between mentors.
7. Unfortunately, it is difficult to compare our results to Bartanen and Kwok's (2021) because they use standardized scores.
8. In theory, TPPs could exert some control over these issues by first dividing PSTs into groups based on preferences for

geographic area, grade level, content area, school context, and/or mentor characteristics. However, this requires a sufficiently large number of PSTs and mentor classrooms within each grouping to allow for random assignment. The more characteristics a TPP wants to influence, the more groups would be required and the smaller the size of each group. While it may be feasible, therefore, for a TPP to randomly assign PSTs to mentors within geographic areas, it may not be feasible for them to randomly assign PSTs to mentors within content area, grade levels, and geographic areas.

References

- Allen, J., Gregory, A., Mikami, A., Lun, J., Hamre, B., & Pianta, R. (2013). Observations of effective teacher-student interactions in secondary school classrooms: Predicting student achievement with the classroom assessment scoring system—Secondary. *School Psychology Review*, 42(1), 76–98.
- Allen, J., Hafen, C. A., Gregory, A. C., Mikami, A. Y., & Pianta, R. (2015). Enhancing secondary school instruction and student achievement: Replication and extension of the My Teaching Partner-secondary intervention. *Journal of Research on Educational Effectiveness*, 8(4), 475–489.
- Allen, M., & Coble, C. (2018). Creating a data culture in educator preparation: The role of the states. In E. Mandinach & E. Gummer (Eds.), *Data for continuous programmatic improvement* (pp. 35–67). Routledge.
- American Association of Colleges for Teacher Education (AACTE). (2018). *A pivot toward clinical practice, its lexicon, and the renewal of educator preparation* (Technical Report). <http://www.nysed.gov/common/nysed/files/cpc-aactecpreport.pdf>
- Anglin, K. L., Wong, V. C., & Boguslav, A. (2021). A natural language processing approach to measuring treatment adherence and consistency using semantic similarity. *AERA Open*, 7, 23328584211028615.
- Araujo, M. C., Carneiro, P., Cruz-Aguayo, Y., & Schady, N. (2016). Teacher quality and learning outcomes in kindergarten. *The Quarterly Journal of Economics*, 131(3), 1415–1453.
- Atteberry, A., Loeb, S., & Wyckoff, J. (2015). Do first impressions matter? Improvements in early career teacher effectiveness. *AERA Open Access Journal*, 1(4), 1–23.
- Bartanen, B., & Kwok, A. (2020). *Pre-service teacher quality and workforce entry* (EdWorkingPaper No. 20-223). Annenberg Institute at Brown University.
- Bartanen, B., & Kwok, A. (2021). Examining clinical teaching observation scores as a measure of preservice teacher quality. *American Educational Research Journal*, 58(5), 887–920.
- Bassok, D., Magouirk, P., & Markowitz, A. J. (2021). Systemwide quality improvement in early childhood education: Evidence from Louisiana. *AERA Open*, 7(1), 1–19.
- Bastian, K. C., Henry, G. T., Pan, Y., & Lys, D. (2016). Teacher candidate performance assessments: Local scoring and implications for teacher preparation program improvement. *Teaching and Teacher Education*, 59, 1–12.
- Bastian, K. C., Lys, D., & Pan, Y. (2018). A framework for improvement: Analyzing performance-assessment scores for evidence-based teacher preparation program reforms. *Journal of Teacher Education*, 69(5), 448–462.
- Bastian, K. C., Lys, D. B., & Whisenant, W. R. L. (2023). Does placement predict performance? Associations between student

- teaching environments and candidates' performance assessment scores. *Journal of Teacher Education*, 74(1), 40–54.
- Bastian, K. C., Patterson, K. M., & Carpenter, D. (2022). Placed for success: Which teachers benefit from high-quality student teaching placements? *Educational Policy*, 36(7), 1583–1611.
- Bell, C. A., Dobbelaer, M. J., Klette, K., & Visscher, A. (2019). Qualities of classroom observation systems. *School Effectiveness and School Improvement*, 30(1), 3–29.
- Bell, C. A., Gitomer, D. H., McCaffrey, D. F., Hamre, B. K., Pianta, R. C., & Qi, Y. (2012). An argument approach to observation protocol validity. *Educational Assessment*, 17(2–3), 62–87.
- Bell, C. A., Qi, Y., Croft, A. J., Leusner, D., McCaffrey, D. F., Gitomer, D. H., & Pianta, R. C. (2015). Improving observational score quality: Challenges in observer thinking. In T. Kane, K. Kerr, & R. Pianta (Eds.), *Designing teacher evaluation systems: New guidance from the measures of effective teaching project* (pp. 50–97). John Wiley.
- Boyd, D., Lankford, H., Loeb, S., Rockoff, J., & Wyckoff, J. (2008). The narrowing gap in New York City teacher qualifications and its implications for student achievement in high-poverty schools. *Journal of Policy Analysis and Management*, 27(4), 793–818.
- Briggs, D. C., & Alzen, J. L. (2019). Making inferences about teacher observation scores over time. *Educational and Psychological Measurement*, 79(4), 636–664.
- Campbell, S. L., & Ronfeldt, M. (2018). Observational evaluation of teachers: Measuring more than we bargained for? *American Educational Research Journal*, 55(6), 1233–1267.
- Casabianca, J. M., Lockwood, J. R., & McCaffrey, D. F. (2015). Trends in classroom observation scores. *Educational and Psychological Measurement*, 75(2), 311–337.
- Castles, A., Rastle, K., & Nation, K. (2018). Ending the reading wars: Reading acquisition from novice to expert. *Psychological Science in the Public Interest*, 19(1), 5–51.
- Charney, R. S. (1993). *Teaching children to care: Management in the responsive classroom*. Northeast Foundation for Children.
- Cohen, J., Erickson, S., Krishnamachari, A., & Wong, V. (2023). Experimental evidence on the robustness of coaching supports in teacher education. *Educational Researcher*.
- Cohen, J., & Goldhaber, D. (2016). Building a more complete understanding of teacher evaluation using classroom observations. *Educational Researcher*, 45(6), 378–387.
- Cohen, J., Wong, V., Krishnamachari, A., & Berlin, R. (2020). Teacher coaching in a simulated environment. *Educational Evaluation and Policy Analysis*, 43(2), 208–231.
- Council for the Accreditation of Educator Preparation (CAEP). (2022). 2022 Initial Level Standards. <https://caepnet.org/~media/Files/caep/standards/2022-initial-standards-1-pager-final.pdf?la=en>
- Davis, S. C., & Peck, C. A. (2020). Using data for program improvement in teacher education: A study of promising practices. *Teachers College Record*, 122(3), 1–48.
- Dieker, L. A., Rodriguez, J. A., Lignugaris/Kraft, B., Hynes, M. C., & Hughes, C. E. (2014). The potential of simulated environments in teacher education: Current and future possibilities. *Teacher Education and Special Education*, 37(1), 21–33.
- Feuer, M., Floden, R., Chudowsky, N., & Ahn, J. (2013). *Evaluation of teacher preparation programs: Purposes, methods, and policy options*. National Academy of Education.
- Forzani, F. M. (2014). Understanding “core practices” and “practice-based” teacher education: Learning from the past. *Journal of Teacher Education*, 65(4), 357–368.
- Gitomer, D., Bell, C., Qi, Y., McCaffrey, D., Hamre, B. K., & Pianta, R. C. (2014). The instructional challenge in improving teaching quality: Lessons from a classroom observation protocol. *Teachers College Record*, 116(6), 1–32.
- Gitomer, D. (Ed.). (2009). *Measurement issues and assessment for teaching quality*. SAGE.
- Goldhaber, D. (2019). Evidence-based teacher preparation: Policy context and what we know. *Journal of Teacher Education*, 70(2), 90–101.
- Goldhaber, D., Krieg, J., Naito, N., & Theobald, R. (2020). Making the most of student teaching: The importance of mentors and scope for change. *Education Finance and Policy*, 15(3), 581–591.
- Goldhaber, D., Krieg, J., & Theobald, R. (2020). Effective like me? Does having a more productive mentor improve the productivity of mentees? *Labour Economics*, 63, Article 101792.
- Goldhaber, D., Ronfeldt, M., Cowan, J., Gratz, T., Bardelli, E., & Truitt, M. (2022). Room for improvement? Mentor teachers and the evolution of teacher preservice clinical evaluations. *American Educational Research Journal*, 59(5), 1011–1048.
- Grossman, P., Compton, C., Igra, D., Ronfeldt, M., Shahan, E., & Williamson, P. (2009). Teaching practice: A cross-professional perspective. *Teachers College Record*, 111(9), 2055–2100.
- Hafen, C. A., Hamre, B. K., Allen, J. P., Bell, C. A., Gitomer, D. H., & Pianta, R. C. (2015). Teaching through interactions in secondary school classrooms: Revisiting the factor structure and practical application of the classroom assessment scoring system-secondary. *The Journal of Early Adolescence*, 35(5–6), 651–680.
- Henry, G. T., Campbell, S. L., Thompson, C. L., Patriarca, L. A., Luterbach, K. J., Lys, D. B., & Covington, V. M. (2013). The predictive validity of measures of teacher candidate programs and performance: Toward an evidence-based approach to teacher preparation. *Journal of Teacher Education*, 64(5), 439–453.
- Hill, H., Charalambous, C. Y., & Kraft, M. A. (2012). When rater reliability is not enough: Teacher observation systems and a case for the generalizability study. *Educational Researcher*, 41(2), 56–64.
- Hill, H., & Grossman, P. (2013). Learning from teacher observations: Challenges and opportunities posed by new teacher evaluation systems. *Harvard Educational Review*, 83(2), 371–384.
- Hill, H., Mancenido, Z., & Loeb, S. (2020). *New research for teacher education* (EdWorkingPaper No. 20-252). Annenberg Institute at Brown University.
- Ho, A. D., & Kane, T. J. (2013). *The reliability of classroom observations by school personnel*. Bill & Melinda Gates Foundation.
- Janssen, F., Grossman, P., & Westbroek, H. (2015). Facilitating decomposition and recomposition in practice-based teacher education: The power of modularity. *Teaching and Teacher Education*, 51, 137–146.
- Kane, T., & Staiger, D. (2012). *Gathering feedback for teaching: Combining high-quality observations with student surveys and achievement gains*. Bill & Melinda Gates Foundation.
- Kraft, M. A., & Gilmour, A. F. (2017). Revisiting the widget effect: Teacher evaluation reforms and the distribution of teacher effectiveness. *Educational Researcher*, 46(5), 234–249.

- Lampert, M., & Graziani, F. (2009). Instructional activities as a tool for teachers' and teacher educators' learning. *The Elementary School Journal*, 109(5), 491–509.
- La Paro, K. M., Pianta, R. C., & Stuhlman, M. (2004). The classroom assessment scoring system: Findings from the prekindergarten year. *The Elementary School Journal*, 104(5), 409–426.
- Levine, A. (2006). *Educating school teachers* (No. 2). The Education Schools Project. http://www.edschools.org/teacher_report_release.htm
- Malmberg, L.-E., Hagger, H., Burn, K., Mutton, T., & Colls, H. (2010). Observed classroom quality during teacher education and two years of professional practice. *Journal of Educational Psychology*, 102(4), 916–932.
- Mancenido, Z. (2022). *Impact evaluations of teacher preparation practices: Challenges and opportunities for more rigorous research* (EdWorkingPaper No. 22–534). Annenberg Institute at Brown University.
- McDonald, M., Kazemi, E., & Kavanagh, S. S. (2013). Core practices and pedagogies of teacher education: A call for a common language and collective activity. *Journal of Teacher Education*, 64(5), 378–386.
- Pacheo, A. (2009). Mapping the terrain of teacher quality. In D. Gitomer (Ed.), *Measurement issues and assessment for teaching quality* (pp. 168–178). SAGE.
- Papay, J. (2012). Refocusing the debate: Assessing the purposes and tools of teacher evaluation. *Harvard Educational Review*, 82(1), 123–141.
- Paris, D., & Alim, H. S. (2014). What are we seeking to sustain through culturally sustaining pedagogy? A loving critique forward. *Harvard Educational Review*, 84(1), 85–100. <https://doi.org/10.17763/haer.84.1.9821873k2ht16m77>
- Park, Y. S., Chen, J., & Holtzman, S. L. (2015). Evaluating efforts to minimize rater bias in scoring classroom observations. In T. Kane, K. Kerr, & R. Pianta (Eds.), *Designing teacher evaluation systems: New guidance from the measures of effective teaching project* (pp. 381–414). John Wiley.
- Peck, C. A., & McDonald, M. (2013). Creating “cultures of evidence” in teacher education: Context, policy, and practice in three high-data-use programs. *The New Educator*, 9(1), 12–28.
- Pianta, R. C., & Hamre, B. K. (2009). Classroom processes and positive youth development: Conceptualizing, measuring, and improving the capacity of interactions between teachers and students. *New Directions for Youth Development*, 121, 33–46.
- Responsive Classroom. (2014). *The responsive classroom approach: Good teaching changes the future*. https://www.responsiveclassroom.org/sites/default/files/pdf_files/RC_approach_White_paper.pdf
- Reznitskaya, A., Kuo, L.-J., Clark, A.-M., Miller, B., Jadallah, M., Anderson, R. C., & Nguyen-Jahiel, K. (2009). Collaborative reasoning: A dialogic approach to group discussions. *Cambridge Journal of Education*, 39(1), 29–48.
- 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., Bardelli, E., Truwit, M., Mullman, H., Schaaf, K., & Baker, J. C. (2020). Improving preservice teachers' feelings of preparedness to teach through recruitment of instructionally effective and experienced cooperating teachers: A randomized experiment. *Educational Evaluation and Policy Analysis*, 42(4), 551–575.
- Ronfeldt, M., Brockman, S. L., & Campbell, S. L. (2018). Does cooperating teachers' instructional effectiveness improve preservice teachers' future performance? *Educational Researcher*, 47(7), 405–418.
- Ronfeldt, M., Matsko, K. K., Greene Nolan, H., & Reininger, M. (2021). Three different measures of graduates' instructional readiness and the features of preservice preparation that predict them. *Journal of Teacher Education*, 72(1), 56–71.
- Texas State Legislature. Texas Senate Bill 174. (2009).
- U.S. Department of Education. (2011). *Our future, our teachers: The Obama administration's plans for teacher education reform and improvement*.
- van der Lans, R. M., van de Grift, W. J., van Veen, K., & Fokkens-Bruinsma, M. (2016). Once is not enough: Establishing reliability criteria for feedback and evaluation decisions based on classroom observations. *Studies in Educational Evaluation*, 50, 88–95.
- Wallace, T. L., Parr, A. K., & Correnti, R. J. (2020). Assessing teachers' classroom management competency: A case study of the classroom assessment scoring system—secondary. *Journal of Psychoeducational Assessment*, 38(4), 475–492.
- Weisberg, D., Sexton, S., Mulhern, J., Keeling, D., Schunck, J., Palcisco, A., & Morgan, K. (2009). *The widget effect: Our national failure to acknowledge and act on differences in teacher effectiveness*. New Teacher Project.
- Wylie, E. C. (2020). Observing formative assessment practice: Learning lessons through validation. *Educational Assessment*, 25(4), 251–258.

Author Biographies

Arielle Boguslav is a senior research associate at the Annenberg Institute at Brown University where she designs and conducts research on teacher professional learning with the Research Partnership for Professional Learning (RPPL). Her research focuses on the design, implementation, and effects of teacher preparation and professional learning programming and experiences, inspired by her experiences as a teacher. Prior to joining Annenberg, she completed her PhD in Education Policy at the University of Virginia. She previously taught middle school English Language Arts in Texas and English as a foreign language for adults in Russia.

Julie Cohen is the Charles S. Robb associate professor of Curriculum and Instruction at the University of Virginia School of Education and Human Development. Her research focuses on teachers and teaching, with a focus on the policies and practices that support the development of effective instruction across content areas. She is an elementary teacher educator and former elementary and middle school teacher.