



Using Common Assignments to Support Educator Effectiveness in Kentucky

Scott Marion, Center for Assessment¹

January 20, 2015

Executive Summary

Student performance results have become a central feature of reformed educator effectiveness systems. This is not to say that measures of teacher practice (e.g., classroom observations) are not a major focus as well, but classroom observations of teacher practices has been a part or, in fact, the only part of most previous educator evaluation designs. Validly incorporating the results of student assessment data into effectiveness determinations has been a challenge on many fronts, from the quality of the assessments to the type of analytic methods employed to transform those data into indicators in effectiveness systems. The challenge exists for teachers from both “tested” and “non-tested” subjects and grades (NTSG), but perhaps more so for the latter. A “tested” subject and grade is one with a state test score and a state test score from at least one prior grade (generally, the immediate prior grade) to allow for the use of complex statistical models (e.g., value-added models, student growth percentiles) to analyze the change in students’ scores. The Common Assignment Study (CAS) likely holds promise for helping to improve the quality of assessment data for use in the effectiveness determinations for teachers in both tested and non-tested subjects and grades.

The Common Assignment Study (CAS) is a research and development project in which teachers collaborate across districts and states to develop and implement high-quality curricular units to ensure student learning of meaningful disciplinary content and skills. Initiated in the 2013-14 school year, the Common Assignment Study is a three-year effort being led by the Colorado Education Initiative (CEI) and The Fund for Transforming Education in Kentucky (The Fund) with support from the Bill & Melinda Gates Foundation. In response to the demand for resources to provide instruction that captures both the rigor and content of the Kentucky Common Core

¹ Thanks to Andy Brownstein, Renee Boss, Rebecca Woosley, Ash Vasudeva, Christy Schneider, Jeri Thompson, and other members of the CAS team for helpful edits and suggestions.



Standards and the Colorado Academic Standards, the CAS was conceived to facilitate multi-district and multi-state collaboration among teachers within participating districts and states during the transition to new college and career-ready standards.

This brief describes how CAS units and assessments may be well-suited to provide much needed, high-quality learning and assessment information to contribute to educator effectiveness determinations. The CAS units are already designed around critical disciplinary concepts and skills (the “big ideas” of the discipline) and being able to link two or more CAS units throughout a school year via a common big idea will allow educators and others to determine how much students have progressed on such key concepts and skills. Those familiar with educator effectiveness methods for NTSG should be recognizing the similarity between CAS units and Student Learning Objectives (SLO) or Student Growth Goals (SGG). In fact, this brief presents an argument for using the structure of SGGs to provide the evaluative framework for using CAS units as part of educator effectiveness determinations.

Using CAS units eliminates several threats to the use of student performance information in educator effectiveness determinations. First, CAS units provide high quality assessment information in ways that are fully integrated with instruction and the daily work of teaching and learning. In other words, “extra” assessments are not needed to support effectiveness determinations when using CAS. Second, many NTSG approaches rely on a limited amount of assessment information (e.g., pretest and posttest). The SGG model in Kentucky advocates for the use of more information than just those two end points. Each CAS unit contains multiple assessment tasks, so using multiple CAS units to support effectiveness determinations would fit into the SGG model seamlessly. Third, CAS is built on a premise that collaboration is critical for producing high quality units and that the CAS process is an effective vehicle for supporting collaboration. This collaborative model can strengthen SGGs by encouraging teachers (and leaders) to work together to create stronger SGGs than individuals can create on their own. Finally, the CAS units provide a natural bridge between the evaluation of educator practices and student performance results in ways that can support more holistic determinations of



effectiveness. In short, CAS units are designed to support high quality curriculum and instruction related to the Common Core State Standards, which is a considerable benefit in its own right, but the CAS units and processes can provide rich and credible information to support educator effectiveness systems.

Introduction

State and district education leaders are rapidly reforming their educator effectiveness systems to include a greater emphasis on the results from student performance measures. There have been legitimate criticisms of basing teacher evaluations on the results of state test scores analyzed using value-added models or student growth percentiles. However, if used appropriately and thoughtfully, the results of student test scores can provide useful information that would enhance the quality of educator evaluations. However, only 25-30% or so of the teaching force nationally is in grades and subjects with at least two years of state test data — what is referred to as the “tested subjects and grades.” Kentucky requires more state, standardized tests than many other states, especially in science and social studies, so the percentage of teachers in tested grades and subjects is likely higher in Kentucky than the national average. One of the greatest challenges in designing and implementing reforms in the area of educator effectiveness determinations is documenting the contributions of the other 70% or so of educators in the “non-tested subjects and grades (NTSG).” There are several assessment and analytic approaches for incorporating measures of student learning into the evaluations of teachers from NTSG, but in reality, most approaches suffer from significant weaknesses (see Appendix A for an overview of commonly used analytic methods).

One of the major shortcomings of current educator evaluation and effectiveness systems is that the two major components—teaching observations and student performance—are treated as separate entities which risks introducing more incoherence into the system than necessary. We argue that evidence from the Common Assignment Study can help bridge this gap. This brief provides an overview of how the Common Assignment Study may offer some promise for strengthening teacher evaluations.

The Common Assignment Study

The Common Assignment Study (CAS) is a research and development project in which teachers collaborate across districts and states to develop and implement high-quality curricular units to ensure student learning of meaningful disciplinary content and skills. Initiated in the 2013-14 Marion. CAS and Teacher Effectiveness in Kentucky: January 23, 2015



school year, the Common Assignment Study is a three-year effort being led by the Colorado Education Initiative (CEI) and The Fund for Transforming Education in Kentucky (The Fund) with support from the Bill & Melinda Gates Foundation. Technical and implementation support for the project is being provided by multiple organizations including Stanford Center for Assessment, Leadership & Equity (SCALE), Westat, The Center for Assessment, and Research for Action.

In response to the demand for resources to provide instruction that captures both the rigor and content of the Kentucky Common Core Standards and the Colorado Academic Standards, the CAS was conceived to facilitate multi-district, multi-state collaboration among teachers within participating districts and states during the transition to new college and career-ready standards. The goals of the project include the following:

- **Unit Design and Use:** Explore how to successfully design, revise, and implement high quality Common Assignment units in participating districts and states with a focus on the use of student work on common assignments to inform instructional conversations;
- **Resources and Support:** Identify the resources, support and professional learning opportunities educators need to effectively design and implement the units in order to support professional learning; and,
- **Examining Evidence of Student Learning:** Determine how the evidence of student learning yielded through the common assignment units may contribute to more robust teacher evaluation systems, in both tested and non-tested grades and subjects².

A major focus of CAS project is to address some of these weaknesses and evaluate the feasibility of using the evidence of student learning generated through the common units as part of educator effectiveness determinations. The project will explore, among other goals, the extent to which student work samples, other classroom-based artifacts, and end-of-unit assessments provide

² While the current project includes subjects and grades that are “tested” in Kentucky, except perhaps for HS biology, the long-term plan is spread these approaches to most subjects and grades.

information related to educator effectiveness to complement, supplement or provide alternatives to traditional test-based measures of student performance. Some of the key questions and topics for study under this strand of the project are:

1. What are the most defensible approaches for analyzing data drawn from the common units in order to make inferences about teacher contributions to student learning?
2. How do inferences about student learning drawn from the common unit assessments relate to external or other measures of student learning, such as state tests?
3. What are the opportunities and challenges associated with inferences about student learning over time (from multiple units) compared with single or even dual data collection windows?
4. Does drawing evidence of student learning from sources of evidence that are fully integrated into the teaching and learning process ameliorate some of the potential unintended negative effects associated with accountability uses of assessment information?

Based on the CAS project, as well as considerable experience developing and implementing educator effectiveness systems around the country, this brief presents an overview of how common assignment information can be used to support educator effectiveness determinations in Kentucky. Common assignment information can be used to support the “student growth” component of the effectiveness determination, but also can provide important evidence related to teaching practices. The brief first discuss some concerns with current approaches for integrated student performance data into educator effectiveness methods and then illustrates how the use of common assignment information may ameliorate some of these challenges and support more credible effectiveness evaluations for Kentucky educators.

Issues and Challenges with Current Educator Effectiveness Approaches

As noted above, the NTSG issue is an obvious challenge that could be minimized through the use of common assignments. Less obvious, though, is the benefit that common units may offer

to evaluating teacher practices and connecting those effectiveness evaluations to student learning.

NTSG Challenges

We want to avoid some of the current misconceptions about the evaluation of NTSG, where unfortunately many states/districts are using analytic approaches that cannot support the intended inferences about educator effectiveness. This is most often played out in the ways that learning goals are operationalized that determine the targets for student “growth,” which are aggregated to targets for teacher “growth.” For example, one of the most common and egregious approaches involves simply subtracting a pretest score from a posttest score, when both scores actually represent different scales and perhaps different learning targets. Many of these concerns were discussed in a previous paper (e.g., Marion et al., 2012)³ about measurement problems with most of the analytic approaches currently used with NTSG. Another problem with some of these more simplistic approaches is that many they rely on only two sources of data (pretest/posttest). In other words, these systems generally preclude the use of the kinds of multiple measures that recognize the value of each measure. Fortunately, Kentucky’s approach for “developing quality growth goals” in its Teacher Professional Growth & Effectiveness System (TPGES) avoids many of the common problems associated approaches noted above in the way that the state encourages the use of status-based or conditional status-based goals.

Another, perhaps even more important, challenge to the successful implementation of educator effectiveness evaluation systems for teachers in NTSG is the disconnect between the evaluations and the actual day-to-day work of teaching and learning. Many see these additional tests administered as part of the effectiveness evaluation as “dropped from the sky,” not much differently from the way in which they see state tests. Even Student Learning Objectives (SGG) are not immune from these problems, depending on how they are implemented. For example, teachers and students in one state with strong state control over the SGG system have been heard

³ For more details on these issues, see Marion et al. (2012).
http://www.nciea.org/publication_PDFs/Measurement%20Considerations%20for%20NTSG_052212.pdf

saying things like, “We took the ‘SLO’ test today,” which clearly indicates that is seen as distinct from the regular teaching and learning system. In other words, the tests being used in this particular system are developed at the district level and teachers are required to administer the tests at specific times whether they fit the particular curriculum at that specific time. On the other hand, Kentucky has taken a more holistic approach for teachers in NTSG and uses existing assessments and other data rather than adding new tests into the mix.

A Role for CAS

Using assessment results drawn from common units offers the promise of ameliorating both of these concerns. The common units, by design, include multiple measures of student performance, hopefully at successive levels of understanding. Second, the rich performance assessments in the common units are fully integrated into the regular work of teaching and learning, and are not seen as extra, accountability assessments. Certainly, we still need to attend to the potential corrupting influences of accountability on the results from assessments included in the common units, but CAS offers promise that these more robust measures of teachers’ contributions to student learning will feel more trustworthy and there would be less temptation for corruption. Third, CAS was based on the premise, which has been borne out by initial evaluations, that collaboration among teachers was an important vehicle for increasing the quality and usefulness of units. Finally, common units and LDC modules offer a way to bridge the false dichotomy in teacher evaluation between teaching practices and student learning.

Incorporating CAS Results into Educator Effectiveness Systems

If evidence of student learning could be collected from multiple units — for example, one from the fall and one from the spring — we would still have to make sense of the data. Ideally, in the future, three or four units may be administered throughout the year. Before getting into the specifics, let’s back up and review the essential components of educator effectiveness evaluation determinations.

- ✓ **Assessment/Data**: These data include assessments administered to students, but could also include data from student survey and classroom observations as is being done with

TPGES. It is helpful to think broadly about data and to recognize that while high quality data are necessary for valid educator evaluations, they are far from sufficient. In other words, a very high quality pretest could be administered in the fall and a very high quality posttest could be administered in the spring, but the inferences in terms of educator effectiveness from these two tests might be poorly substantiated for reasons we discuss below.

- ✓ **Analytics**: There is a famous Seinfeld episode that takes place at a car rental counter where Jerry tries to explain the difference between “taking the reservation,” which the clerk keeps emphasizing that they have done, and “holding the reservation,” which the car rental company clearly has not done because there is no a car available. The analytics of accountability systems is the “holding of the reservation,” because it turns the assessment information into something useful. In other words, it is relatively easy to collect data (taking the reservation), but it is quite difficult to analyze the data in ways to support defensible inferences (holding the reservation). This includes the methods by which we turn raw data into accountability indicators. In terms of state tests, this may mean applying student growth percentile (SGP) methods to translate test scores across multiple years into SGPs for students and aggregate SGPs for educators. Similarly, an observation rubric is a tool for turning the raw material of classroom interactions into numeric quantities that are then usually summed or averaged. Given the many choices of methods for turning raw data into indicators (e.g., noticeably different observation rubrics), it is easy to see how the choice of analytic method can influence the inference about educator quality. The analytic challenge is generally greater for NTSG than for state tests or classroom observations, in part because of the small sample sizes and limited analytic capacity among school and district personnel.
- ✓ **Attribution**: Perhaps the least considered and most important aspect of using data for evaluation decisions is deciding which teachers deserve the “credit” or “blame” for which students. This is one of those issues that seems so simple to those outside of education, but the devil is truly in the details. Besides the overwhelming issue of data quality, there is a critical need to understand how schools really operate. Many elementary students,

for example, are nominally assigned to a single teacher, but in reality are moved around in grade-level teams for some forms of differentiated instruction or other activities. Some have tried to address this issue by measuring “dosage” which is the assignment of weights to the student score based on the proportion of time the child spent with each teacher. Attributing student learning results to teachers is more complex than simply counting seat time and such dosage approaches would require strong evidence to support such strong assumptions. Therefore, we have advocated certain types of “shared attribution” approaches in which all students in a given team of teachers might be “pooled” for the purposes of attributing the scores from student assessments — if that fits the theory of action and structures for the school. This can occur for both tested and non-tested subjects and grades.

Student Growth Goals⁴ as a Framework

The previous discussion should demonstrate that we cannot just throw assessment data at the wall and hope that something useful for educator evaluation sticks. We need a framework for contextualizing and analyzing the evidence of student learning that emerges from the common assignments. Student Growth Goals (SGG) provide such a framework in Kentucky.

SGG (SLO in many other states) have gained popularity as a means of attributing student performance results to educators in new forms of teacher evaluation systems for all teachers, but especially those in NTSG. SGGs can serve as a framework for incorporating the results from common assignments into educator evaluation systems. SGGs are content- and grade/course-specific measurable learning objectives that can be used to document student learning over a defined period of time, based on a SMART goal-setting approach. SGGs are designed to reflect and incentivize good teaching practices, such as setting clear learning targets, differentiating instruction for students, monitoring students’ progress toward these targets, and evaluating the extent to which students have met the targets.

⁴ Kentucky uses the term Student Growth Goals for what many others call Student Learning Objectives (SLO). The two terms are used interchangeably here.

Clear and Meaningful Learning Goal

SGGs should reflect the relevant content standards, skills, and the associated curriculum. They should describe what students will be able to do at the end of the course/grade, or at least over a reasonably long period of instruction (e.g., a semester). The learning goal(s) generally will be established by a collection of teachers in the same grade and/or subject area and overseen by the district or school leadership. The learning goal for an SGG should reflect an **“enduring skill” in the discipline** and may encompass several key content standards. Multiple teachers are often working with their students on the same learning goals and the results from the same SGG implemented by multiple teachers may be shared (i.e., shared attribution).

The following is an example of a learning goal from a Common Assignment Study middle school social studies:

Students will independently use primary and secondary sources to form evidence-based interpretation of historical events and ideas significant to westward expansion in the antebellum era. In doing so, students will consider and discuss multiple causes for historical events with a focus on how and why the US acquired Western lands. Further, students will understand how the use of authentic primary source documents enriches our understanding of historical events and ideas.

As can be seen, the learning goal represents key ideas in the study of history, but is still somewhat general. The addition of the assessments and student targets help make the SGG more specific in terms of performance expectations⁵.

Assessments used to evaluate students’ achievement of the learning goals

The assessments used to evaluate the degree to which students have achieved the learning goals should be of high quality; that is, they should be designed to provide evidence aligned to the specific learning goal. First, if the learning goals take the rich form described above—as they

⁵ Much more information about SGGs (SLOs) can be found in the Center for Assessment’s SLO Toolkit at www.nciea.org.

should if they are to support high quality instruction and deeper learning— it is unlikely that they will be measured well with just a single assessment. Multiple assessments will likely be required, and performance or other authentic assessments must be part of the assessment system designed to evaluate the learning goal. The TPGES expects teachers to collect “ongoing formative assessment” to determine whether students have achieved the goals. An example of a multiple tasks from the middle CAS history unit tied to the learning goal presented earlier follows:

Task 1: Considering the historical perspectives on the Mexican-American War, was President Polk justified when he declared war in 1846? After reading primary and secondary sources given in class, write an essay in which you address the question and argue whether Polk was justified in his decision to declare war with Mexico. Support your position with evidence from the text(s). Be sure to (acknowledge; refute) competing views.

Task 2: Part A. Explain in detail how the U.S. acquired one territory through war and one territory through diplomacy

Task 2: Part B. Choose two groups of people that were affected by westward expansion and explain how the expansion affected those groups.

The responses to these tasks will be scored with a multi-dimensional rubric focused on historical analysis, argumentative writing, use of sources, and quality of writing and presentation.

Targets for both student performance and aggregate targets for educator performance

Student Targets. The student target, also referred to as proficiency statements in Kentucky, is the expected outcome at the end of the instructional period. Targets in Kentucky are the same for all students in the class and should be appropriate, given the interval of instruction, for the whole class and for special populations (e.g., English Learners, Students with Disabilities). Those proposing SGGs should ensure that the student targets are both ambitious and realistic, which is quite a challenging design task, especially during the early years of SGG implementation.

Several researchers (e.g., Marion, et al., 2012; Lachlan-Haché, Cushing, & Bivona, 2012⁶) have suggested that teachers set targets using available baseline data to help contextualize the learning targets for individuals or groups of students. The TPGES process uses baseline data as a basis for both growth and achievement/proficiency goals for students. A description of a set of targets for the learning goal described above follows:

These targets are based on the rubric scores for the CAS Unit tasks described above, as well as the three (3) formal performance tasks administered in this course over the entire school year. This is not an average of scores, but rather an evaluation of students' demonstrated consistent improvement over time in the criteria of the rubric.

- *All students will either advance one performance level from the beginning of the year to the end of year. Students maintaining the highest level of performance from the beginning to the end of year will count as having met this requirement.*
- *Approximately 60% of the students will score at the “proficient” level on the majority of the rubrics’ criteria across the multiple summative tasks.*

This example presents one way of establishing SGG targets. Setting ambitious and reasonable targets for SGGs is one of the most challenging aspects of SGG design and implementation. This has been the focus of several recent national meetings as states wrestle with how to approach the issue in fair and valid ways.⁷

Teacher Targets. Teacher targets specify how the student aggregate scores (results) will be used to determine the degree to which the teacher has met the SGG targets and whether these results will be employed directly or transformed into an indicator for use in accountability determinations. In some cases, these targets are set by the state, but more typically, they are

⁶ Lachlan-Haché, L., Cushing, E., & Bivona, L. (2012). Student Learning Objectives as Measures of Educator Effectiveness: The Basics. Washington, DC: American Institutes for Research.

⁷ This topic could fill an entire paper, but for now see Marion, et al. (2012).

determined by the district in conjunction with school leaders. Ideally, school leaders will tailor the targets in consultation with teachers to account for specific classroom contexts. Typically, teacher targets and the corresponding performance rating are classified into three or four levels. For example, a teacher may be classified as “not meeting the SGG” if less than 50% of the students reach their target, “meeting” if 51-85% of the students reach their target, and “exceeding” if more than 85% of the students reach their target. Obviously, the appropriateness of these targets is contingent upon the learning goal, assessments, and student targets. It will take several years of data collection and analysis to evaluate the appropriateness of these targets.

Incorporating CAS into an SGG Framework

By now, those familiar with CAS should recognize the obvious connections between SGGs and CAS. These connections are explored in more detail in this section. If the multiple common units administered through the school year were tied to common “enduring skills” of the discipline, we could begin to describe student progress referenced to actual changes in knowledge and skills rather than just scale-based (i.e., points) descriptions of growth. The common assignments could provide the measurement information at key points along the trajectory implied by this learning goal. The TPGES expects ongoing data collection to evaluation goals and the CAS could contribute by providing high quality data to support effectiveness determinations. Given the current state of affairs in teacher evaluation, this could be one of the most important contributions of this project. In the following section, I relate the CAS approach to an SGG framework.

Learning Goal

Ideally, the learning goal(s) should be identified ahead of time so that the units are designed to support the learning and assessment of this goal. In the real world, this will likely be an iterative process between high-priority learning goals and available units. To provide credible evidence of changes in student learning over time, the multiple common units in a school year must be tied to the same big idea of the discipline. Of course, the question is “What idea?” The Literacy Design Collaborative (LDC) module incorporated into each unit is designed to focus on various

aspects of argumentative writing across the curriculum. There is no question that this is very important and argumentative writing is critical vehicle for communicating understanding in most disciplines, but the key (and specific) concepts and skills of science, social studies, or other content areas should be represented in the learning goals. More recent work in LDC has focused on ensuring that the disciplinary content is well-represented in the LDC task and scoring rubrics. Therefore, argumentative writing could represent one of the cross-unit learning goals while a more discipline-based idea (e.g., understanding how geographic and economic factors influence human settlement) could represent an additional cross-unit strand.

Assessments

A major advantage of using assessments from common assignment units for educator evaluation is that it avoids the typical impoverished assessment evidence common to many NTSG approaches. Unit assessments are designed as a system to yield high quality information about students' understanding of the big idea as well as supporting knowledge and skills. Further, the assessments across multiple units will be intentionally connected—not formally linked (psychometrically), but conceptually connected via the assessment targets and rubrics. This will require a focused effort to design and implement high quality assessments and rubrics that can support inferences related to the enduring skills. Importantly, these units are closer to the actual teaching and learning than is typically the case with assessments like end-of-course tests and as evidence accumulated through the CAS research that hopefully shows the relationship between the CAS assessments and these more distal assessments, states like Kentucky will feel confident about the validity of the interpretations from the CAS assessments.

Targets

As noted above, establishing targets for students is one of the most challenging aspects of SGG design and implementation. Establishing reasonable baselines to contextualize the end-of-year assessment results is a critical component of the SGG process. Depending on when the fall unit is implemented, the assessments from the fall unit could serve as the baseline for the SGG. However, if baseline information is available from earlier in the school year than implantation of

the fall unit, assessment information from the fall unit could serve as an interim check on the goal. The assessment information from the spring unit would be used to evaluate the degree to which students have met the targets associated with the learning goal. Using multiple assessments would be an improvement over most SGGs, but SGGs based on common assignments could be improved even further if the multiple units were tied to practice-based learning progressions. The progressions would then allow teachers to develop a better understanding of how students move toward deeper understanding of the big idea. The progressions can serve as explicit anchors for the unit assessments and other activities. Documenting student progress along these learning progressions can then serve as evidence for teacher effectiveness determinations.

Importantly, we should avoid simple reductionist approaches for determining and evaluating targets. The “rough conditioning” approach described previously (Marion, et al., 2012) and a forthcoming Reform Support Network (RSN) brief on target setting for SLOs offers a sensible alternative to establishing arbitrary numeric goals based on some distance between a pretest score and 100% (or some other percentage correct). The evidence from common assignments will allow us to more accurately document students’ learning trajectories in a discipline. It is beyond the scope of this brief and perhaps too decontextualized to offer very specific recommendations for establishing meaningful student targets based on CAS units, but I strongly recommend capturing qualitative changes in student understanding through a domain-specific rubric. Such rubric scores can then be used to aggregate student-level results to the classroom or school levels. The TPGES in Kentucky uses an approach for establishing targets aligned with the recommendations in Marion, et al. (2012) as well as those just noted.

Common Assignments as Evidence of Teaching Practices

Essentially all teacher evaluation and effectiveness systems being implemented in the United States count evidence of student learning and evidence of teacher practices as separate elements in overall evaluation scores. To be fair, there are sound arguments for such approaches. Many believe that observations and other indicators of teacher practice measure the process of

teaching, while assessment data provides evidence of the “product” of teaching. Further, nobody doubts that essentially all traditional evaluations have focused exclusively on process. Therefore, the theory of action for counting student learning results as a significant part of teacher evaluations is intended to demonstrate that student learning must be valued.

However, by treating the two major components as separate aspects of teacher evaluation, there is a risk of introducing more incoherence into the system than necessary. For example, most of the major tools for measuring teacher practice (e.g., Danielson’s Framework for Teaching)⁸ require evidence of how well teachers plan and execute instructional activities, but it would also make sense to include evidence related to what students learned from this instruction. This sounds obvious, but is rarely done. The evidence of student learning generally comes from assessments only peripherally related to classroom observations.

This disconnect is expected (but not a goal) for the “tested grades and subjects” because of the broad brush of the state tests⁹. On the other hand, we have the leeway with NTSG to design assessments closer to the specific activities in the course or classroom. For example, if the school was using common assignments to generate assessment information for SGGs, it would make sense to link classroom observations to the times when the teacher and students are working on the common units. This would provide the observer a way to triangulate multiple forms of evidence in order to draw more defensible inferences. This is a relatively easy step schools could take when implementing common assignments, but additional steps can be followed to bring even more coherence to educator evaluation systems.

The common units provide opportunities for generating and collecting data related to many aspects of teaching practice. In fact, the Kentucky Department of Education has already made similar linkages between LDC and the Danielson Framework. For example, the common units

⁸ See: <http://www.danielsongroup.org/article.aspx?page=frameworkforteaching>

⁹ This is not to say that value-added (VAM) or student growth percentile (SGP) models for large scale assessment data cannot provide useful contributions to teacher evaluation decisions especially when the results of such analyses of student growth are shared among multiple educators.

require considerable planning and thinking about instruction to ensure that students are provided with meaningful learning opportunities. Examine the strands in Danielson’s Domain 1 (see below), it is clear that common units could easily provide evidence for each of these six strands. The CAS unit planning templates, among other artifacts, will serve as meaningful evidence of each these six strands. I am using Danielson is being used as an example because of the direct relevance to Kentucky’s Framework for Teaching, but similar strands are found in other tools for capturing evidence of teaching practice (e.g., Marzano, Marshall).

From Danielson’s Framework for Effective Teaching

Domain 1: Planning and Preparation

- 1a. Demonstrating Knowledge of Content and Pedagogy*
- 1b. Demonstrating Knowledge of Students*
- 1c. Setting Instructional Outcomes*
- 1d. Demonstrating Knowledge of Resources*
- 1e. Designing Coherent Instruction*
- 1f. Designing Student Assessments*

More specific to assessment literacy, almost all of the major tools for evaluating practice include rubrics for teachers’ assessment literacy and the use of assessment results to improve learning. For example, Danielson includes criteria for “designing student assessments” as part of Strand 1f from Domain 1 (see above) as characterized by the following description of the proficient level in her rubric:

All the instructional outcomes may be assessed by the proposed assessment plan; assessment methodologies may have been adapted for groups of students.

Assessment criteria and standards are clear. The teacher has a well-developed strategy for using formative assessment and has designed particular approaches to be used¹⁰.

¹⁰ <http://www.danielsongroup.org/userfiles/files/downloads/2013EvaluationInstrument.pdf>

This is just one excerpt, but it is clear that Danielson and others draw attention to appropriate design and use of assessments for instruction and learning in many aspects of evaluating teaching practice. The CAS unit planning template and other planning tools, along with actual assessments and associated student work, can provide rich evidence to support this and related strands.

The Role of Collaboration and Examining Student Work

The CAS project has emphasized the critical role that close examination of student work can play in the development and evaluation of units as well as for helping teachers gain insights into students' strengths and needs. We argue that student work can serve as a credible bridge between the evaluation of teachers' practice and characterization of student performance and growth. If we cannot see it in the student work, it isn't there! In the case of educator effectiveness, student work is the perfect bridge between teacher practices and student performance results in that the student work can help shed light on the degree to which specific teacher practices were effective at improving student learning.

As noted earlier, CAS is based on the premise that collaboration among engaged educators can lead to higher quality instructional, curriculum, and assessment materials than would be the case if teachers were working individually. While there is certainly value in having individual teachers examine student work from their classrooms, the value increases exponentially when groups of teachers collaboratively examine student work from common and even unique (to each classroom) tasks. Generally, teachers are able to gain deeper insights into student learning and their own teaching practices through collaborative inquiry into student work than each could likely realize on their own.

While the SGG process used in Kentucky does not require that teachers work with others when establishing and evaluating their SGGs, we argue that common and collaboratively established SGGs for educators teaching the same courses can strengthen the SGG process and improve the

effectiveness of those educators involved in the collaborative work. Additionally, the Professional Responsibility dimension of the Danielson framework supports teachers working with colleagues to improve the profession and each teacher's growth. This is another example of how the CAS process provides a connection between the student performance and teacher practice aspects of educator effectiveness determinations.

Discussion

As illustrated throughout this paper, the common units at the core of CAS offer considerable opportunities for providing evidence in support of both student learning outcomes and teaching practices. An advantage in Kentucky is that there is considerable flexibility afforded to local educational leaders for designing assessments and methods for documenting changes in student achievement over time. Designing local educator evaluations is very difficult work. Educational leaders and others are often being asked to do things without the necessary expertise and experience to do so. The work of CAS can help support these district leaders and educators as they strive to build coherent systems that are both fully integrated into the real work of teaching and learning and serve as a vehicle for deepening this work. As project leaders, we will have to work with policy leaders in both states to understand the potential of CAS to support meaningful educator evaluations. Closer to the action, project leaders will need to work closely with district leaders responsible for educator evaluation so they have an understanding of how to credibly include CAS information in local evaluation decisions.

Appendix A: Summary of Major Analytic Approaches Used for Educator Evaluation¹¹

Analytic Approaches for Estimating Growth	Necessary Conditions	Procedures	Growth Interpretation
Growth Models	Pre-and post-test measures are available within the subject area of interest and exist on a common vertical scale	Calculate difference between pre and post-test performance on common scale	Gain (or Loss) in student performance between two points in time
Conditional Status Models (e.g., VAM-Models; SGP)	Pre-test data on one or more assessment(s) in or related to (i.e., correlated with) the subject area of interest and subject-specific post-test data are available	Condition on pre-test data (and potentially other covariates) as a means of evaluating post-test performance for a given student.	Difference between expected performance and observed performance given prior performance in the same or a related subject area.
Goal Setting/SGG Process	Process by which teachers use existing student performance data (of a variety of types) to establish learning goals for students in their class, and then evaluate student performance relative to those goals		Degree to which a student or group of students attained one or more specified learning goals
Shared Attribution	Any of a variety of techniques that involves the attribution of a common estimate of student growth, achievement or teacher impact on student learning — based on aggregation at the group, school or district level — to one or more teachers		Depends on nature of shared attribution approach

¹¹ From: Hall, E., Gagnon, D., Thompson, J., Schneider, C., & Marion, S. (2014). State Practices Related to the Use of Student Achievement Measures in the Evaluation of Teachers in Non-Tested Subjects and Grades. www.nciea.org.