Model-Based Performance Assessment

CSE Technical Report 465

Eva L. Baker
CRESST/University of California, Los Angeles

January 1998

MODEL-BASED PERFORMANCE ASSESSMENT

Eva L. Baker

CRESST/University of California, Los Angeles

There is concern in many quarters about the type and level of knowledge our children are acquiring in schools. For example, on national (National Assessment of Educational Progress) and international (Second International Mathematics and Science Study; Third International Mathematics and Science Study) assessments, many students do poorly. Further, there is a growing concern in American business and industry (U.S. Department of Labor, 1991, 1992) that young people entering the workforce are not adequately prepared for the world of work. These concerns are a major source of what is now almost a decade of effort to restructure what students are taught and the ways in which we can accurately assess their learning. For example, government policy makers have attempted to address these issues, beginning with the 1989 National Governors' Conference in Charlottesville, Virginia, followed by the work of the National Education Goals Panel (1991) and the National Council on Education Standards and Testing (1992), and finally by the passage of the *Goals 2000: Educate America Act* by Congress in 1994.

The National Education Goals Panel, responsible for monitoring achievement towards the Goals (see *Goals* 2000) on a national and state level, also advised that standards of learning be set so that students' progress towards the Goals could be determined. Goal 3, which states that "all students will leave grades 4, 8, and 12 having demonstrated competency over challenging subject matter . . .", was a particular focal point for the development of standards. The National Education Goals Panel also pointed to the need to define more specifically what constitutes "challenging subject matter" and "competency" in it—for instance, the kind of learning characterized by higher order thinking skills, deep content knowledge within and across subject areas, problem-solving ability—and the need to determine how that competency would be measured.

Standards now have been developed for eight of the nine subject areas listed in Goal 3—English, mathematics, science, foreign languages, civics and

government, arts, history, and geography—by groups of experts, such as the National Council of Teachers of Mathematics (1989), the National Research Council (*National Science Education Standards*, 1996), and the National Council of Teachers of English (1996). Many states, for example, California, Kentucky, and Maryland, and school districts and schools also have created, adapted, and adopted individual sets of standards (National Education Goals Panel, 1996).

But goals for national educational achievement and resulting standards to describe those goals will not tell us how well our children are doing unless we also measure their progress in learning the content of the standards. To do this, new kinds of tests are being created, called performance assessments, in which students engage in tasks that may require significant amounts of time and in which they are asked to "communicate their understanding of content, of process and strategy, and of the results obtained" (Baker & O'Neil, 1994, p. 11).

Performance assessment also has been described by its proponents as a major strategy to assist teachers to improve the learning of their students. This piece will describe both the values ascribed to performance assessment and the major criticisms of assessment that have developed in the last few years of exploration. One approach, model-based performance assessment, will be described as a way to remedy and to avoid criticisms of performance assessment.

Why Performance Assessment?

Part of the argument for performance assessment is that it is supposed to focus on tasks calling for complex thinking, deep understanding of subject matter, and open-ended responses, in contrast to assessment by multiple-choice measures, which are well suited, for example, to testing memorized factual knowledge. Performance assessment can direct the attention of teachers and learners to the important forest and not the trivial trees. Students can concentrate on worthy, long-term goals and integrate their learning so that they can solve problems, analyze text, and communicate their understanding in useful ways. A particular support for performance assessment grows out of the current educational reform that encourages states and districts to set challenging standards in particular subject matters (or interdisciplinary learning) and depends upon appropriately challenging assessment tasks to determine the attainment of standards (National Education Goals Panel, 1996). Another source of support for performance assessment is research from educational psychology

that documents that many students learn best when asked to construct their solution, rather than when they merely choose a solution from a set of provided options (Mayer, 1996). If performance-based approaches more closely adhere to the way students effectively acquire skills and understanding, tests should match this process. Up to this time, much formal testing has relied on multiple-choice tests. As the public continues to pressure teachers and schools to raise test scores, it makes sense to change tests so that there is greater congruence between effective teaching practices and methods of evaluating school and student progress (Baker, O'Neil, & Linn, 1993). In fact, many versions of performance-based assessment are underway in various educational settings, and their results are used for different purposes—among them classroom feedback, parent conferences, and judging the quality of important educational programs.

The Dark Side

But there are numerous concerns raised about performance assessment. Many questions arise simply because performance assessment is new and has been promoted with high expectations for its impact. Other issues have surfaced from the experiences of users of performance assessment (Stecher et al., 1996). These concerns need to be examined and solutions to them should be considered. A quick review of prominent doubts includes the following problems:

1. Performance assessments are difficult and expensive to develop.

2. Performance assessments, because they are open-ended, require trained judges to evaluate students' efforts and cost far more than other approaches.

3. Many teachers are not prepared to teach in the way performance assessments imply.

4. Many parents believe that performance assessment is a less rigorous method to evaluate students than more familiar multiple-choice tests.

These complaints can be validly levied against many forms of performance assessment. However, the criticisms are most credible when performance assessments have been created in a particular way, without a strong model or tight specifications. For example, if one is trying to assess the quality of students' understanding of integrated science, it is possible to create an infinite number of

performance tasks to measure such a general goal. Students could be asked to develop environmental projects, to solve problems in teams involving weather forecasting, or to search the Internet and make a presentation about health and nutrition. Each of these tasks would then require the development of a special way to judge performance, a scoring rubric constructed for the particular task to be rated, and would apply only to the particular topics investigated. Teachers would have to be trained to use each rubric, and the technical quality of each task, its fairness, and its reliability would have to be determined if it were to be used for formal testing purposes. Most importantly, performance assessments that have been developed as stand-alone or one-at-a-time propositions get used up. If students are asked to forecast weather for a particular geographical and seasonal position, that task cannot easily be used again on a test that will count for a grade or to evaluate students in a district or state assessment. The very qualities of performance assessment that make it desirable—a task takes time and thought to complete—also make the task memorable and inappropriate for reuse. So we can agree that performance assessments, if they are developed to be unique, may very well cost too much to implement.

Furthermore, there is the question of who should develop performance assessments for what purposes. Whereas many initial efforts involving performance assessment used teams of teachers in development, the fact was that relatively few of the draft assessments survived technical quality standards. This phenomenon occurred for a variety of legitimate reasons: (a) Teachers may think in terms of their own student groups and create tasks and scoring schemes that are less generally usable; (b) teachers have not been trained to develop performance tasks, so many of their efforts may be more closely related to interesting instructional experiences than to reasonable measures of student accomplishments; (c) teachers may have major gaps in their understanding of particular topics in school subject matter—they may be strong in Civil War history, weak in the Colonial period, excellent in geometry, weaker in algebra. The unevenness of teachers' content knowledge may affect their choice of performance tasks as well as the quality of particular tasks.

Another downfall of creating a wide range of performance tasks in a "hand-crafted" way is the unreasonable expectation this approach requires of teaching itself. Assume that each performance task is separately designed and developed. On the face of it, teachers would need to understand how the tasks were

developed and clearly have in mind the standards used to score each task so that students could be well prepared for success. Imagine that, in a fifth-grade class, there were six such tasks in science, five in mathematics, seven in language arts, and four in social studies. A single teacher would need to be aware of the scoring rubrics for 22 tasks. Multiply these tasks by the numbers of students who will need special assistance and attention, and it is clear that this approach presents teachers with an impossible load. Furthermore, tasks developed very separately ignore one of the most desirable objectives of teaching, teaching so that students can apply their learning to new settings, sometimes called teaching for application or transfer (Mayer & Wittrock, 1996). Unless performance tasks provide teachers with guidance on the selection of other similar tasks for instructional purposes, there is little likelihood that a teacher will be able to develop reasonable instructional experiences for students. Under that condition, two contingencies occur: Either the teacher will be left to teach the particular performance task on the test, an unfair and unethical decision; or the teacher will be unlikely to prepare students to attain desired goals, opening, however inappropriately, the teacher and the school to charges of incompetence.

A last complaint involves the suspicion of parents and policy makers that performance assessments are inevitably less rigorous than multiple-choice tests. Multiple-choice tests of course have the benefit of objectivity—they can be automatically scored, and their marking will not be biased for or against any particular group of students (or teachers). Some concerns about performance assessments have been voiced because the content of the assessment (the subject matter knowledge to be learned) has been subsumed to the process the assessment is measuring. A good example is in the area of writing, where many parents would like to see student essays that are composed carefully, observe standard English conventions, and deal with a particular subject matter. Parents may become concerned when assessments emphasize the way students developed ideas for the essay, and the use of drafting and editing cycles as major criteria. Furthermore, parents may attend to superficial characteristics of tasks—such as whether students are collecting objects and ordering them according to a particular scheme. Such tasks may appear to be more like play than learning. Finally, parents have objected to tasks that have been designed to increase student attention and engagement. The use of controversial content, test directions that call for students to express their own opinions, and instructions to draw examples from the student's own life have suggested that particular

ideology is being promoted, that any opinion suffices as a "good" answer, and that the privacy of students and their families is being invaded. Although it is unlikely that these inferences are valid, the fact they have been aired repeatedly by many parent groups requires some level of attention.

Model-Based Performance Assessment

A different development approach has been taken by some explorers of performance assessment. This approach focuses on deciding what type of learning is to be accomplished and then embeds the learning in the subject matter content (or content standard) that the assessment is supposed to address. Then a set of specifications is developed to provide guidance for task developers, for teachers looking to create comparable instructional experiences, and for those responsible for renewing the system. These specifications include the content to be assessed, the kinds of information and other content given to students, and the scoring approaches to be used (Baker, Aschbacher, Niemi, & Sato, 1992; Baker, Aschbacher, Niemi, Sato, & Redfield, 1992). In a very real way, the model is a general blueprint about what is expected, delimiting the nature of the assessment.

Performance-Assessment Learning Models

Researchers at the National Center for Research on Evaluation, Standards, and Student Testing (CRESST) have developed models in five areas of learning to measure students' understanding of content knowledge, their representation of knowledge, their teamwork, their problem-solving skills, and their metacognitive abilities, as illustrated in Figure 1 (Baker, Aschbacher, et al., 1996). These models are then implemented in each subject matter area. Particular tasks may assess more than one of the learning types. In those cases, specifications for learning areas are combined. There are multiple ways to assess any of these areas, and the examples below illustrate some alternatives.

Let's start with an example. Assume we are developing performance assessments in the field of history. We are interested in acquiring evidence that students can understand complex historical content. Building a set of specifications, we first decide what content standard or standards our assessment is to measure. We choose one that says:
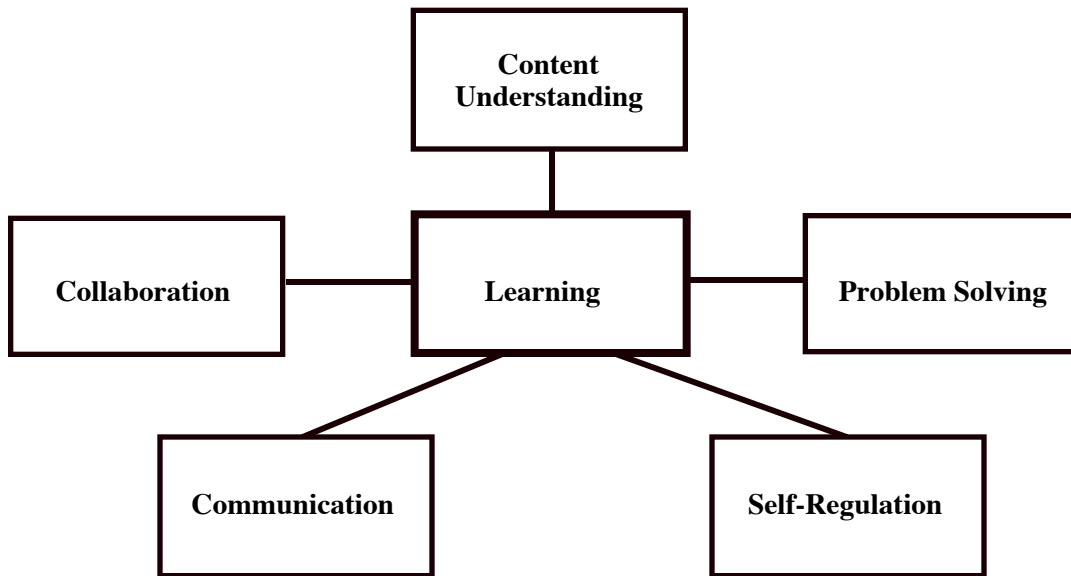
```
                    ┌──────────────────┐
                    │     Content      │
                    │  Understanding   │
                    └────────┬─────────┘
                             │
┌──────────────┐   ┌─────────┴────────┐   ┌──────────────────┐
│ Collaboration│───│     Learning     │───│ Problem Solving  │
└──────────────┘   └───┬──────────┬───┘   └──────────────────┘
                      ╱            ╲
          ┌──────────────┐   ┌──────────────────┐
          │ Communication│   │ Self-Regulation  │
          └──────────────┘   └──────────────────┘
```

Figure 1. Areas of learning.

History

> Students are to demonstrate their knowledge and ability
> to analyze important events from American history.

Most standards would then go on to elaborate the historical eras of most importance, for example, the Colonial period, the Westward Expansion, the Cold War period.

The second step is to determine the type of performance assessment of interest. We could choose from tasks requiring short (15-minute) analyses, an on-demand task taking a few hours to complete, tasks where students collected a wide variety of information and created a term paper or project over a period of weeks, or a portfolio containing different examples of student work accumulated over a semester or more. We select an on-demand task, meaning that the student is "tested" without having complete choice over the content and in a manner that attempts to keep testing procedures comparable from classroom to classroom to ensure fairness.

At this point we select the set of specifications we want to use to assess the content in history. Let us provide one example created by a team of assessment specialists and teachers.

Sample Specifications for Content Understanding in History

Students will be asked to read primary source material exemplifying important events in an historical period. The primary source materials may be text, maps, or pictures. They need to be presented in their original language, without expurgation. They should be reasonably short (taking 15 minutes or so to read) and have been directed by their author(s) at a common rather than technical audience. Text may consist, for example, of letters, speeches, or newspaper articles of the period under assessment. Two selections exemplifying different points of view in the particular era or about the particular event will be given to the student. Prior to reading the texts students will be asked to complete either a brief short-answer test or a short multiple-choice test focused on key facts and principles relevant to the period. This test has two purposes—to help the students recall information they have previously learned that will facilitate their understanding the new primary source texts; and to provide a measure of students' "prior knowledge" in order to determine whether inadequate student performance on a subsequent essay task was caused by lack of factual information, lack of understanding of principles, or inability to synthesize the primary source materials into a coherent argument. Following the reading of the texts, students will be asked to create an essay in the form of a letter to a relative explaining the alternative positions taken by the authors and the meaning they have for understanding the particular event or period.

The actual texts provided in one task involved one of the Lincoln-Douglas debates (Baker, Freeman, & Clayton, 1991). Students were asked to write a letter to a cousin after having attended the debate in Springfield, Illinois, and explain the differences in position of the debaters and the importance of their points of

view. To summarize, the model so far includes content standards, type of performance, and specifications (Figure 2).

Any number of primary source materials can replace the Lincoln-Douglas debate, creating a template for the rapid development of comparable tests in American history. For example, speeches from the McCarthy era could be used. Documents about the use of atomic weapons to end World War II could be used. Students could read excerpts from the Declaration of Independence and writings of a Loyalist with strong ties to England. Furthermore, the task may be simplified by choosing shorter or more appropriate texts for use by younger students, for example, 5th- or 7th-grade students. In one case, students were given maps and letters and asked to write about the topic of the Westward Expansion. In another, students in Hawaii were given excerpts from the diary of Queen Lili'uokalani and speeches by Senator Sanford B. Dole to help them discuss the annexation of the Hawaiian Islands (Baker, Niemi, et al., 1996). The same type of specifications have been used to generate interdisciplinary tasks assessing geography and history, economics, and political science.

We have also used approaches to deal with complex language (18th-century prose, for example) and have provided students with glossaries or mini-dictionaries of difficult words in provided texts (Baker, Niemi, et al., 1996).

- Content Standard(s) Referenced by Assessment: Students are to demonstrate their knowledge of and ability to analyze important events from American history

- Type of Performance Assessment: On Demand

- Specification

  – Two opposing positions conveyed in text, map, or other representation intended for common use, taking about 15 minutes each to read

  – 15-item multiple-choice or short-answer examination asking students to identify key principles (e.g., sectionalism), documents, or events (Dred Scott decision) relevant to understanding texts

  – Set of directions to the student about what is to be produced (e.g., write a letter)

Figure 2. Content understanding specifications (11th-grade history).

Scoring Criteria

What the tasks now need is an approach to scoring them that will remain reasonably consistent over different content areas (to foster good teaching and efficient scoring) but also takes into account the important issues of content.

After many years of research, CRESST created a scoring approach that was developed by asking experts (in content and teaching) as well as students to answer draft questions we had prepared. We determined that the experts addressed the problems using the same general approaches. They wrote from premises or important principles; they used knowledge they had acquired outside of the testing setting (that is, not solely from the primary sources provided in the assessments) to elaborate and explain their answers; they integrated information from the text carefully, selecting it to illustrate a point; and they constructed their argument in a form appropriate to their discipline, for example, contrastive analyses, discussion of cause and effect, or chronological narrative. In comparison, students who were relatively inexperienced wrote "flat" essays where all points received about equal weight, prepared essays that depended a great deal on the material presented in the testing session, used almost no prior knowledge, and demonstrated major misconceptions, for instance, mixing up figures from different centuries.

Analyses of the differences between experts and novices led to our creating a scoring rubric that incorporated major elements from each (Figure 3).

- Overall Content Understanding—A holistic score related to the appropriateness of the essay and the quality of understanding displayed.

- Use of Principles—Incorporates important and appropriate principles as organizing ideas

- Use of Prior Knowledge—References knowledge to explain particular event and to justify analysis

- Use of Text—Text references to illustrate or extend argument

- Integrated Argument—Using a style of argument that is integrative and appropriate to subject area

- Misconceptions—Number and type

Figure 3. CRESST content understanding scoring rubric elements.

The sections of the rubric specifically related to the application of principles, prior knowledge, and so on in combination with the multiple-choice or short-answer test on prior knowledge are especially useful for teachers as they attempt to pinpoint areas of classroom instruction requiring more attention and effort. Of course, a more formal and detailed rubric, with illustrations of papers and descriptions of score point values, is available as are model training procedures (Baker, Aschbacher, Niemi, & Sato, 1992).

The rubric as created makes strong demands on teacher knowledge of subject matter. Teachers must be able to recognize key principles and prior knowledge as well as point out misconceptions. The rubric was constructed to permit students a variety of ways to demonstrate their understanding, but the gatekeeper on its use is a high level of teacher expertise. When we have used topics that are relatively new to the curriculum and for which there is little expectation that teachers possess much knowledge, we have augmented the rubric with examples of specific principles and prior knowledge relevant to a particular era or topic. We intend, however, to avoid creating a rubric that generates a cookie-cutter right answer to a complex topic, particularly because we want to avoid the possibility that such "right answers" could be taught or memorized by those wishing to game the system.

We have then taken our specifications and our rubric and applied them to multiple topics within history, appropriate for different grade levels. We also have used the model to generate performance assessments in science, where the primary source materials may be a description or video of a single experiment, in mathematics, where the primary source material is a series of solved problems and the students are asked to explain the method of solution, and in interdisciplinary areas such as the humanities (Aschbacher & Herman, 1991). In each of these cases, the scoring rubric adapts to the particular subject matter area: Principles in mathematics, in science, and in history are very different from one another, but they are still principles. Similarly, each area will require the use of background or prior knowledge, although the nature of prior knowledge vastly changes with the content. What the model does, however, is to focus the student's and teacher's attention on the elements of constructing knowledge in a content area: the relation of factual knowledge to "big ideas" and to prior understanding, and the use of strategies (such as argumentation) that are appropriate for each discipline.

This approach has been subjected to many research studies; it has been used in the Hawaii State Assessment and is now forming the core of the Los Angeles Unified School District Performance Assessment. In these settings, we have added another category to our rubric focused on the use of writing conventions to get double duty from the assessments.

Conclusions

We believe that we can argue that the use of model-based assessments can mitigate the criticisms leveled at performance assessments: (a) Specifications provide a relatively efficient way to develop items and a strategy to translate standards into assessments. (b) Specifications provide guidance about important attributes of instruction. (c) General scoring rubrics reduce the load on teachers, for in the multiple-subjects classroom found in many elementary schools, teachers must remember only one set of criteria (principles, prior knowledge, etc.) that can be applied to different topics. (d) General scoring rubrics reduce the cost of scoring since common rubrics are used in multiple topics.

Model-based performance assessments have the capability to be used for both summative (formal) assessments, used by districts and states, and formative assessments, used at the classroom level by teachers. If a common model can generate both classroom and district or state assessments, students and teachers will experience greater coherence in their instructional experiences. Finally, model-based assessments, because they depend upon application in subject matter, may overcome some parents' doubts about the rigor of performance assessments. Simply asking parents to try to write such an essay gives them a real feel for the challenges put before their children.

There still remain numerous questions about performance assessments, primarily related to how best to schedule them (as they are time consuming). Research is underway at CRESST and elsewhere to attempt to augment ways to score them using computer systems so as to reduce the cost of assessments (Baker & O'Neil, 1995; Chung, O'Neil, Herl, & Dennis, 1997; Herl, O'Neil, Chung, & Schacter, 1997). In addition, explorations of how to involve teachers in a reasonable and cost-efficient way in the scoring of the assessments have been undertaken by many groups.

At CRESST, we have developed models for performance assessment, in each of the learning areas depicted at the outset—problem solving,

communication, collaboration, metacognition, and content understanding—and in some areas we have more than one approach. Our strategy makes some trade-offs. It emphasizes comparability among different assessments, reasonable cost, technical quality, fairness, and utility for instruction. It gives up on a wide, anything-goes approach to assessment and focuses on deeper assessment of fewer interpretations of types of learning. We believe it is important to align content standards, classroom assessment, and external assessment in a practical way, particularly when assessment is used for policy purposes. Our model-based assessment is one way to do it.

References

Aschbacher, P. R., & Herman, J. (1991, Fall). *Final report of the Humanitas program evaluation, 1990-91*. Los Angeles: University of California, Center for the Study of Evaluation.

Baker, E. L., Aschbacher, P. R., Herman, J. L., Chin-Chance, S., Aguirre-Muñoz, Z., & Niemi, D. (1996, April). *Panel discussion: Transferring validated performance assessment models to state/district school reform*. Presented at the annual meeting of the American Educational Research Association, New York.

Baker, E. L., Aschbacher, P. R., Niemi, D., & Sato, E. (1992). *CRESST performance assessment models: Assessing content area explanations*. Los Angeles: University of California, National Center for Research on Evaluation, Standards, and Student Testing.

Baker, E. L., Aschbacher, P., Niemi, D., Sato, E., & Redfield, D. (1992, April). *Issues in large-scale performance assessment*. Invited minicourse presentation at the annual meeting of the American Educational Research Association/National Council on Measurement in Education, San Francisco.

Baker, E. L., Freeman, M., & Clayton, S. (1991). Cognitive assessment of history for large-scale testing. In M. C. Wittrock & E. L. Baker (Eds.), *Testing and cognition* (pp. 131-153). Englewood Cliffs, NJ: Prentice-Hall.

Baker, E. L., Niemi, D., Herl, H., Aguirre-Muñoz, A., Staley, L., & Linn, R. L. (1996). *Report on the content area performance assessments (CAPA): A collaboration among the Hawaii Department of Education, the Center for Research on Evaluation, Standards, and Student Testing (CRESST) and the teachers and children of Hawaii* (Final Deliverable). Los Angeles: University of California, National Center for Research on Evaluation, Standards, and Student Testing.

Baker, E. L., & O'Neil, H. F., Jr. (1994). Performance assessment and equity: a view from the USA. *Assessment in Education, 1*(1), 11-26. [Reprinted in M. B. Kane (Ed.), *Implementing performance assessment: Promises, problems, and challenges.* Hillsdale, NJ: Lawrence Erlbaum Associates, 1996.]

Baker, E. L., & O'Neil, H. F., Jr. (1995). Computer technology futures for the improvement of assessment. *Journal of Science Education and Technology, 4*(1), 37-45.

Baker, E. L., O'Neil, H. F. Jr., & Linn, R. L. (1993). Policy and validity prospects for performance-based assessments. *American Psychologist, 48*, 1210-1218.

Chung, G. K. W. K., O'Neil, H. F., Jr., Herl, H. E., & Dennis, R. A. (1997, March). Use of networked collaborative concept mapping to measure team processes and team outcomes. In H. F. O'Neil, Jr. (Chair), *An integrated simulation approach to assessment*. Symposium presented at the annual meeting of the American Educational Research Association, Chicago.

Goals 2000: Educate America Act. P. L. 103-227, March 31, 1994.

Herl, H. E., O'Neil, H. F., Jr., Chung, G. K. W. K., & Schacter, J. (1997, March). Feasibility of an on-line concept mapping construction and scoring system. In H. F. O'Neil, Jr. (Chair), *An integrated simulation approach to assessment*. Symposium presented at the annual meeting of the American Educational Research Association, Chicago.

Mayer, R. E. (1996). Learners as information processors: Legacies and limitations of educational psychology's second metaphor. *Educational Psychologist, 3*, 151-161.

Mayer, R. E., & Wittrock, M. (1996). Problem-solving transfer. In D. C. Berliner & R. C. Calfee (Eds.), *Handbook of educational psychology* (pp. 47-62). New York: Simon & Schuster Macmillan.

National Council of Teachers of English. (1996). *Standards for the English language arts*. Urbana, IL/Newark, DE: National Council of Teachers of English/International Reading Association.

National Council of Teachers of Mathematics. (1989). *Curriculum and evaluation standards for school mathematics*. Reston, VA: Author.

National Council on Education Standards and Testing. (1992). *Raising standards for American education. A report to Congress, the Secretary of Education, the National Education Goals Panel, and the American people*. Washington, DC: U.S. Government Printing Office.

National Education Goals Panel. (1991). *The National Education Goals report, 1991: Building a nation of learners*. Washington, DC: U.S. Government Printing Office.

National Education Goals Panel. (1996). *The National Education Goals report, 1996. Executive summary. Commonly asked questions about standards and assessments*. Washington, DC: U.S. Government Printing Office.

National science education standards. (1996). Washington, DC: National Academy Press.

Stecher, B. M., Rahn, M., Ruby, A., Alt, M., Robyn, A., & Ward, B. (1996). *Using alternative assessments in vocational education* (MR-836-NCRVE/UCB). Santa Monica, CA: RAND.

U.S. Department of Labor, Secretary's Commission on Achieving Necessary Skills. (1991). *What work requires of schools. A SCANS report for America 2000*. Washington, DC: (SCANS), U.S. Government Printing Office.

U.S. Department of Labor, Secretary's Commission on Achieving Necessary Skills. (1992). *Skills and tasks for jobs: A SCANS report for America 2000*. Washington, DC: U.S. Government Printing Office.