

Assessing Assessment: The SAUM Evaluator's Perspective

Peter Ewell
*National Center for Higher Education
Management Systems (NCHEMS)*
peter@nchems.org

The SAUM project took place within a broader context of assessment in American higher education. Faculty teams in mathematics departments experienced in microcosm what their colleagues in many other disciplines were simultaneously experiencing, and their actions were shaped by larger forces of politics and accountability affecting their institutions. At the same time, their efforts to develop and document viable department-level approaches to assessment in mathematics helped inform the national assessment movement—a field badly in need of concrete, discipline-level examples of good practice. Evaluation of SAUM helped bridge these two worlds.

In my personal role as project evaluator, I continued to participate in national conversations about assessment's purposes and prospects throughout the three-year grant period. But watching SAUM participants struggle with the day-to-day reality of crafting workable assessment approaches in their own departments helped keep me honest about what could and could not be accomplished. Similarly, the participant experiences that were revealed through the evaluation information we compiled often mirrored what was happening to other “early adopters” elsewhere.

The first section of this chapter sets the wider stage for SAUM by locating the project in a national context of assessment. A second section reflects on my role as project evaluator, and describes the kinds of evaluative information we collected to examine the project's activities and impact. A third section presents some of what we learned—focused primarily on what participants told us about how they experienced the project and the challenges they faced in implementing assessment initiatives back home.

SAUM in a National Context

The so-called “assessment movement” in higher education began in the mid-1980s with the confluence of two major forces. One originated inside the academy, prompted by growing concerns about curricular coherence and the conviction that concrete information about how and how well students were learning could be collectively used by faculty to improve teaching and learning (NIE, 1984). This version of “assessment” was low-stakes, incremental, faculty-owned, and guided by a metaphor of scholarship. The other driving force for assessment originated outside the academy prompted by policymakers' growing concerns about the productivity and effectiveness of colleges and universities (NGA, 1986). This version of “assessment” was high-stakes, publicly visible, accountability-oriented, and infused with the urgency of K–12 reform embodied in *A Nation at Risk* (USDOE, 1983).

Although fundamentally contradictory, both these forces were needed to launch and sustain a national movement. External authorities—first in the guise of states and later in the guise of regional accrediting organizations—served to constantly keep assessment at the forefront of institutional attention. But because these external requirements were at first fairly benign—and because academic leaders quickly saw the need to protect the academy’s autonomy by developing locally-owned processes that might actually be useful—internal preferences for diverse evidence-gathering approaches aimed at institutional improvement served to discharge accountability as well for many years (Ewell, 1987).

The environment within which the SAUM project was launched was shaped by fifteen years of growing institutional experience with steering between these contradictory poles of assessment. By 2000, virtually every institution could claim that it “did assessment,” at least in the sense that it had developed learning outcomes goals for general education and that it periodically surveyed its students and graduates. Most could also point to the beginnings of an institution-level organizational infrastructure for assessment—a coordinator operating out of the academic affairs office perhaps, or a faculty-staffed institutional assessment committee. About a third could lay claim to more sophisticated efforts including testing programs in general education, portfolios assembled by students and organized around general learning outcomes like “effective communication” or “critical thinking,” or specially-designed assignments intended to both grade students individually and provide faculty with broader information about patterns of student strength and weakness in various abilities. Indeed, as revealed by the programs at such gatherings as the annual Assessment Forum hosted by the American Association for Higher Education (AAHE), there was a steady increase in the sophistication of institutional assessment efforts with respect to method and approach throughout this period, and equally steady progress in faculty acceptance of the fact that assessment was a part of what colleges and universities, for whatever reason, had to do.

By 2000, moreover, the primary reason why institutions had to “do assessment” had become regional accreditation. State mandates for assessment in public institutions, instituted in the wake of the National Governors Association’s *Time for Results* report in the mid-1980s, had lost a lot of steam in the recession that appeared about 1990. States had other things to worry about and there were few resources to pursue existing mandates in any case. Accreditors, meanwhile, were under mounting pressure by federal authorities to increase their focus on student learning outcomes.

Regional accrediting organizations must be “recognized” by the U.S. Department of Education in order for accredited status to serve as a gatekeeper for receipt of federal funds. The federal recognition process involves a regular review of accreditation standards and practices against established guidelines. And since 1989 these guidelines have emphasized the assessment of learning outcomes more forcefully each time they have been revised by the Department. Accreditors are still accorded the leeway to allow institutions to develop their own learning outcomes and to assess them in their own ways. But by 2002, when SAUM was launched, it was apparent that accreditors could no longer afford to allow institutions to get by with little or no assessment—which had up to then essentially been the case—if they hoped to maintain their recognized status. The result was growing pressure on institutions to get moving on assessment, together with growing awareness among institutional academic leaders that a response was imperative.

But even at this late date, assessment remained something distant and faintly “administrative” for the vast majority of college faculty. It was rarely an activity departments engaged in regularly outside professional fields like engineering, education, business, or the health professions where specialized accreditation requirements made assessment mandatory. And even in these cases, the fact that deans and other academic administrators were front and center in the process, complete with the requisite guidelines, memos, schedules, and reports—all written in passive prose—made it likely that faculty in departments like mathematics would keep their distance. At the same time, despite their growing methodological prowess, few institutions were able to effectively “close the loop” by using assessment results in decision-making or to improve instruction. Periodic assessment reports were distributed, to be sure, but most of them ended up on shelves to be ritually retrieved when external visitors inquiring about the topic arrived on campus. Much of the reason for this phenomenon, in hindsight, is apparent. Assessment findings tend to be fine-grained and focused, while institutional decisions remain big and messy. *Real* application required smaller settings, located much closer to the teaching and learning situations that assessment could actually inform.

In this context, the notion of grounding assessment in the individual disciplines where faculty professionally lived and worked made a great deal of sense. For one thing, assessment practitioners had discovered that methods and approaches ought appropriately to vary substantially across fields and that such purportedly “generic” academic abilities as “critical thinking” and “communication” were manifest (and thus had to be assessed) very differently in differ-

ent disciplinary contexts. At the disciplinary level, moreover, learning outcomes were generally much more easily specified than at the institutional level where of necessity they had to be so broadly cast that they often lost their meaning. More importantly, faculty tended to listen to one another more carefully in disciplinary communities bound by common languages and familiar hierarchies of respect. Even when assessment leaders on campus were faculty instead of administrators, their obvious background in methods derived from education and the social sciences often distanced them from colleagues in the sciences, humanities, fine and performing arts—as well as mathematics. For all these reasons, anchoring assessment in individual disciplinary communities was critical if it was to become a meaningful activity for faculty.

But why mathematics? In my view, mathematics became an “early adopter” of assessment for at least three reasons. First, the discipline is embedded in multiple aspects of teaching and learning beyond its own major at most institutions. Like colleagues in writing—but unlike those in physics, philosophy, and French—mathematics faculty had to staff basic skills courses in general education. As a result, both their course designs and pedagogies in such offerings as calculus and statistics must be closely aligned with a range of client disciplines including the sciences, engineering, business, and the social sciences. As a “basic skill,” moreover, mathematics is generally assessed already at most institutions in the form of placement examinations, so at least some members of every mathematics department have experience with test construction and use. Where developmental mathematics courses are offered, moreover, they are often evaluated directly because the question of effectiveness is of broad institutional interest—a condition not enjoyed by, say, a course in Chaucer. All these factors meant that at least some members of any institution’s mathematics department have at least some familiarity with broader issues of testing, evaluation, and pedagogy.

Second, mathematics has more explicit connections than most other disciplines with the preparation of elementary and secondary school teachers. Even if mathematics faculty are not explicitly located in mathematics education programs, many at smaller colleges and universities that produce large numbers of teachers are aware of pedagogical and assessment issues through this connection, and this knowledge and inclination can translate quickly to the post-secondary level.

Finally, at least at the undergraduate level, learning outcomes in mathematics are somewhat more easily specified than in many other disciplines. Although I have learned through the SAUM project that mathematicians are as apt to

disagree about the nuances of certain aspects of student performance as any other body of faculty—what constitutes “elegance,” for instance, or an effective verbal representation of a mathematical concept—they can certainly come to closure faster than their colleagues elsewhere on a substantial portion of what undergraduate students ought to know and be able to do in the discipline. For all these reasons, mathematics was particularly well positioned as a discipline in 2002 to broaden and deepen conversations about assessment through a project like SAUM.

Evaluating SAUM: Some Reflections.

Serving as SAUM’s external evaluator provided me with a personally unmatched opportunity to explicitly test my own beliefs and assumptions as an assessment practitioner. On the one hand, I have spent almost 25 years advocating for assessment, helping to develop assessment methods and policies, and working with individual campuses to design assessment programs. One cannot do this and remain sane unless one is at some level convinced of assessment’s efficacy and benefit. Yet evaluation is an empirical and unforgiving exercise. SAUM’s central premise was that it is possible to create a practice-based infrastructure for assessment that departments of mathematics could adapt and adopt for their own purposes, and thus improve teaching and learning. On the larger stage of institutional and public policy, this premise has been the basis for my professional career. The opportunity to “assess assessment” as it was acted out by one important discipline—and to reflect on what I found—was both exciting and sobering.

On a personal note, I also came to strongly value my role as “participant-observer” in the project and the opportunities that it provided me (and I hope to the project’s participants) to see beyond customary professional boundaries. For my own part, I was gratified to witness many of the lessons about how to go about assessment that I had been preaching to Provosts and Deans for many years confirmed in microcosm among mathematics faculty at the departmental level. But I also saw (at times to my chagrin) the many differences in perception and failures of communication that can occur when such organizational boundaries are crossed.

As one telling example, at one of the SAUM department-level workshops I encountered a departmental team that reported a particularly frustrating bureaucratized approach to assessment at its institution being undertaken in response to an upcoming accreditation review—an institution that I knew from another source was being cited as a “model” of

flexible and creative assessment implementation by the accreditor in question. I like to think that such insights, and they occurred throughout the project, helped keep me humble in the balance of my work in assessment.

At the same time, I like to think that my boundary-spanning role helped participants achieve some of the project's objectives. An instance here, as the previous example suggests, was my considerable ongoing work with accrediting organizations, which allowed me to interpret their motives and methods for SAUM participants, and perhaps set a broader context for their local assessment efforts.

Like many large, multi-faceted projects, SAUM presented many evaluation dilemmas. Certainly, it was perfectly straightforward to conduct formative data-collection efforts intended to guide the future implementation of project activities. For example, we collected participant reactions from the sixteen SAUM workshops conducted at MAA section meetings and used them to focus and improve these sessions. Responses from section meeting participants early in the process stressed the need for concrete examples from other mathematics departments that faculty could take home with them. Participating faculty observed that they often learned as much from interaction with other participants as from the material presented. These lessons were steadily incorporated into sessions at later section meetings (as well as into the design of the SAUM department-level workshops that were beginning to take shape at that time) and participant reactions steadily improved. Similarly, we learned through a follow-up of SAUM department-level workshop participants that a three-meeting format was superior to a two-meeting format and returned to the former in the project's last phase.

But determining SAUM's effectiveness in a more substantive way posed significant challenges. The most important of these was the fact that the bulk of the project's anticipated impact on mathematics faculty and departments would occur (if it did) well after the project was over. (As good an example of this dilemma as any is the fact that the publication in which this essay appears is one of the project's principal products; yet it reaches your hands as a reader only after the conclusion of the formal project evaluation!) Determining SAUM's effect on assessment practice in mathematics departments thus had to be largely a matter of following the experiences of project participants—particularly those mathematics faculty who attended the multi-session department-level workshops—when they returned to their home departments to apply what they learned. We did this primarily through email surveys given to participants ten months to a year after the conclusion of their workshop experience. The multi-session format of the

department-level workshops also helped the evaluation because at each of the workshop's concluding sessions we were able to explicitly ask participants about their experiences between workshop sessions. What we learned about the experience of assessment at the departmental level is reported in the following section.

We also set the stage for a more formal evaluation of SAUM's impact by conducting an electronic survey of mathematics departments early in the project's initial year. This was intended to provide baseline information about existing department-level assessment practices. A similar survey of departmental assessment practices will be undertaken at the conclusion of the project in the fall of 2005. The baseline survey was administered via MAA departmental liaisons to a sample of 200 mathematics departments stratified by size, institutional type, and location. 112 responses were received after three email reminders sent by the MAA, yielding a response rate of 56%. Questions on the electronic survey were similar to questions that we also posed to 316 individuals who attended SAUM workshop sessions at section meetings, which constituted another source of baseline information.

Justifying the project's potential impact, both sets of baseline data suggested that in 2002 most mathematics departments were at the initial stages of developing a systematic assessment approach. About 40% of department liaisons (and only 20% of participants at section meetings) reported comprehensive efforts in which assessment was done regularly in multiple areas, and another 35% (and 31% of participants at section meetings) reported that assessment was done "in a few areas." About 10% percent of department liaisons (and 21% of participants at section meetings) reported that assessment was "just getting started," and 15% percent (and more than a quarter of participants at section meetings) reported that there was "no systematic effort." Respondents from research universities reported somewhat lower levels of activity than other types of institutions. Differences in responses between department liaisons and the regular mathematics faculty who presumably attend section meetings are notable and reflect the pattern of reporting on institutional assessment activities typical of the late 1980s: in these surveys, administrators routinely reported higher levels of institutional engagement in assessment than was apparent to faculty at their own institutions (El-Khawass, 1987).

Baseline survey results also revealed that the mathematics major is the most popular target for assessment activities, with almost three quarters of responding departments indicating some activity here. About half of the departmental liaisons surveyed indicated that assessment takes place

in either general service or remedial courses, and about a third reported that assessment takes place in courses for prospective teachers and in placement and advising. Not surprisingly, community colleges were somewhat more likely to report assessment in remedial and developmental courses, and less likely to report assessment of the major. Masters degree granting universities were more likely to be engaging in assessment of general service courses and courses for prospective teachers. Doctorate-granting research universities were somewhat less likely than others to be undertaking assessment in any of these areas.

Survey results suggested that mathematics departments are using a wide variety of assessment methods. The most popular method was faculty-designed examinations, which 62% of departments reported using. 53% of departments reported using standardized tests, which is not surprising, but more than 40% of departments reported employing so-called “authentic” approaches like work samples, project presentations (oral and written), or capstone courses. About 40% also reported using surveys of currently-enrolled students and program graduates. Standardized examinations tended to be used slightly more by departments engaging in assessment of remedial and developmental courses, while projects and work samples were somewhat more associated with assessing general service courses.

Finally, the departmental baseline survey asked respondents about their familiarity with *Assessment Practices in Undergraduate Mathematics* (Gold et al., 1999), which had then been in print for several years. Some 19% reported that they had consulted or used the volume, while another 35% noted that they were aware of it, but had not used it. The balance of 46% indicated that they were not aware of the volume. As might be expected, awareness and use were somewhat related to how far along a department felt it was with respect to assessment activity. About 61% of respondents from departments reporting comprehensive assessment programs in place said they were at least aware of *Assessment Practices* and about 25% had actively used it. Only about a third of those reporting no systematic plans had even heard of the volume and none had used it. Certainly, these baseline results leave plenty of room for growth and it will be instructive to see if three years of SAUM have helped move the numbers.

Emerging Impacts

Despite the fact that most of the SAUM project’s impact will only be apparent after the publication of this volume, evaluation results to date suggest some emerging impacts. The majority must be inferred from responses to the follow-

up surveys administered to department-level workshop participants about a year after they attended, focused on their continuing efforts to implement assessment projects in their own departments. Many of these results parallel what others have found in the assessment literature about the effective implementation of assessment at the institutional level, and for disciplines beyond mathematics.

Collegueship. Like any change effort in the academy, implementing assessment can be a lonely business because its faculty practitioners are dispersed across many campuses with few local colleagues to turn to for practical advice or support. Indeed, one of the most important early accomplishments of the national assessment movement in higher education was to establish visible and viable networks of institution-level assessment colleagues through such mechanisms as the AAHE Assessment Forum (Ewell, 2002a).

Results of the evaluation to date indicate that SAUM is clearly fulfilling this role within the mathematics community. A first dimension here is simply the fact that SAUM is a network of *mathematicians*, not just “people doing assessment.” As one faculty member told us, “history, agriculture, and even physics have different flavors of assessment from mathematics” and the opportunity to work with other mathematicians on mathematics topics in assessment was critical in grounding effective departmental efforts. Another dimension is simply the reassurance for individual mathematicians who first get involved in local projects that assessment is a going concern. Here it was useful for SAUM participants to learn that many mathematicians are already involved in assessment—more than many realized—and that assessment is not a peripheral activity that only a few mathematics departments are involved in. These points were seen by SAUM participants as particularly important in “selling” assessment to other faculty when they returned to their home departments.

The team basis for participation in SAUM workshops meant automatic collegueship and mutual support. Simply being away together with colleagues, far from the pressures of everyday campus work was also deemed helpful. At the same time, working with other campuses at the workshop in multiple encounters helped build a feeling among SAUM participants of being part of a larger “movement” that had momentum. This was especially important for faculty who felt, in the words of one, that they had “been thrust into a leadership position on assessment” with little real preparation for this role. Knowing that others were in the same position and sharing approaches about what to do about it was seen as especially important.

The same was true of learning about more specific assessment approaches where, as expected, SAUM depart-

ments borrowed liberally from one another. But by far the most important impact of having colleagues was the stimulus they provided to keep participating departments moving. The need to present departmental progress periodically and publicly was important to this dynamic: teams at the workshops knew that they were going to have to report to their peers, so worked hard to have something to say. As one representative noted, “if we had run out of time and didn’t accomplish what we intended, it probably wouldn’t have had any consequences on campus—we were already doing more than most departments—but because we had to have presentations ready at different workshops we were pushed to follow through on plans and to discuss and revise our activities.” This “peer stimulation” effect was a particularly important dynamic in SAUM, and parallels similar lessons learned about collegiality in other assessment-related change projects (e.g., Schapiro and Levine, 1999).

On a more sobering note, however, early evaluation results also suggest the difficulty of maintaining collegiality absent the explicit framework of a project or a visible network to support it. Few departmental representatives reported contacting other participants on their own, largely due to pressures of time. Again confirming lessons of the assessment movement more generally, an infrastructure for sustaining assessment in mathematics must be actively *built*; it will not just happen as a result of peoples’ good experiences.

Learning By Doing. Another lesson of the assessment movement nationally was the importance of early hands-on experience and practice. Evaluation results to date suggest that SAUM is strongly replicating this finding within mathematics departments. Rather than looking for the best “model program” and planning implementation down to the last detail, SAUM departments, like their colleagues in many disciplines, learned quickly that time invested in even a messy first effort trumped similar investments in “perfecting” design. As one participant told us, “the most important lesson that I learned was to just get started doing something.” Another said, “begin with a manageable project, but begin.” Small projects can not only illustrate the assessment process in manageable ways with only limited investments of resources but can also quickly provide tangible accomplishments to show doubting colleagues. The importance of this insight was reinforced by the fact that most participants also discovered that assessment was a good deal more time-consuming than they had first imagined, even for relatively simple things.

At the same time, participating departments learned the importance of finding their own way *in* their own way, and that local variations in approach are both legitimate and

effective. Like their colleagues elsewhere in assessment, they also learned tactical lessons about implementation that could only be learned by doing. As one department reported, “for us, designing an assessment program means finding a balance between getting good information ... and not increasing faculty workload too much.” Additional comments stressed the importance of knowing that “one size does not fit all” and that good assessment should be related to local circumstances.

Finally, participants encountered aspects of local departmental culture that could not be addressed through formulaic methods. One of them summarized this condition nicely: “there are rules at my institution about how we have to do assessment even though those rules are unwritten, unarticulated (except when violated), and specific to my institution and the larger community. I used to think that these rules were to be found somewhere in the literature ... now I know that I’m dealing with the unknown and with rules that are likely being made up as we go. This makes me much more confident in my own ideas instead of backing down when I am told that something is ‘not allowed.’”

Growing Maturity. As their projects evolved, most participating departments reported a growing maturity with assessment. Several departments doing program-level assessment, for example, had replicated their assessment models in another related department or program (e.g., computer science), or had been working with other departments at the institution to help them develop an assessment approach. Others implemented or regularized activities that they had planned or experimented with at the workshop. Most indicated that they had expanded their departmental assessment efforts to become more systematic and comprehensive—adding new assessment techniques and applying them to more courses or involving more faculty. As one participating department reported, “We believe one of our greatest accomplishments is to have engaged a significant proportion of the department (more than half the faculty) in assessment in one way or another.”

Growing maturity is also apparent in organizational and motivational dimensions. With regard to the former, several departments reported that they had discovered the importance of having a departmental advocate or champion for assessment who could set timelines, enforce deadlines, and provide visibility. A few also reported “regularizing” assessment activities—in one case allowing the original project leader to hand off assessment activities to a newly interested faculty member to coordinate or lead. As one departmental representative put it, participating “got us off to a great start and developed a sense of confidence that we are in a better posture with assessment than other departments on our campus.”

Motivational shifts were more subtle, but reflected a shift toward internal instead of external reasons for engaging in assessment. Mirroring experience in other fields, many mathematicians first heard about assessment through accreditation or their administration's desire to "create a program." But as their participation in SAUM progressed, many also reported new attitudes toward assessment. As one faculty member put it, "Earlier [activities] were about responding to outside pressure...[later activities] were about doing this for ourselves." Another noted, "Most people [at my institution] are not as advanced in assessment as we are in the mathematics department. ... The task still seems to most people like a necessary activity conducted for external reasons, rather than an activity that has intrinsic value to improve their own work." This shift requires time to accomplish and findings from other fields emphasizes the fact that outside pressures or occasions are important to start things moving on assessment at the institutional level (Ewell, 2002b). But SAUM participants began to recognize also that sustaining assessment requires the kind of internal motivation that can only be developed over time and through collective action.

Changing Departmental Culture. Twenty years after the emergence of assessment as a recognizable phenomenon in higher education, it has yet to become a "culture of use" among faculty in disciplines that lack professional accreditation. Many reasons for this have been advanced, ranging from alien language to lack of institutional incentives for engagement, but by far the most prominent is the imposition of assessment requirements by external authorities (Ewell 2002a). Consistent with national experience in other disciplines, SAUM workshop participants thus returned to their own departments determined to make a difference, but they faced an uphill battle to change their colleagues' attitudes about assessment and, in the longer term, to begin to transform their department's culture.

A first milestone here was the fact that participation in SAUM itself helped legitimize the work of developing assessment. Being part of a recognized, NSF-funded project was important in convincing others that the work was important. So was the clear commitment of workshop participants to working on their projects. As one faculty member told us, "Because [the participants] were genuinely interested, ... that interest and enthusiasm has been acknowledged by others." Several also mentioned the value of knowing the "justifications for assessment" in communicating with fellow faculty members.

But SAUM participants also tended to end up being the "assessment people" in their departments—accorded legitimacy for their activities to be sure, but not yet joined by sig-

nificant numbers of colleagues. As one wryly stated, "it has probably made more work for me as when I share an idea of something that we can do, I usually get put in charge of doing it." Another doubted that he and his SAUM teammate had gained much stature in the department because of their participation, "but I guess at least more people recognize what we have done."

As national experience suggests, moreover, wider impacts on departmental culture with respect to assessment require time to develop—more than the three years of engagement most SAUM participants have to date enjoyed. Most indicated that their colleagues were in general more informed about assessment as a result of SAUM and were therefore more willing to agree that it might be beneficial for their departments or institutions. So despite little groundswell of enthusiasm, most did report slow progress in changing departmental attitudes. One participant captured the typical condition succinctly when he reported that his colleagues "remain largely indifferent to assessment...they are in favor of improving programs as long as it doesn't bother them." Another described this condition as follows: "I think there is still a degree of skepticism about all of this, but at least we don't run into outright hostility or claims that this is all a great waste of time and effort." Echoing these comments, a third reported that "the department has definitely become more open to the idea of assessment ... for one thing, they have finally realized that it is not going away ... for another, if there is someone willing to do the work, they will cooperate." These are no small achievements. But the overall pattern of impact to this point remains one of increased awareness and momentum for assessment among SAUM departments with only a few early signs of a changed departmental culture.

"Assessing assessment" through the SAUM evaluation remains an ongoing activity. Like assessment itself in many ways, the task will never be finished. But it is safe to conclude at this point that mathematics has built a resource through SAUM that if maintained, will be of lasting value. On a personal note, I have learned much from my colleagues in mathematics and have been grateful for the opportunity to work with them on a sustained basis. And from a national perspective, I can say without reservation that they are making a difference.

References

- El-Khawas, Elaine (1987). *Campus Trends: Higher Education Panel Report #77*. Washington, DC: American Council on Education (ACE).
- Ewell, Peter T. (1987). *Assessment, Accountability, and Improvement: Managing the Contradiction*. Washington, DC: American Association of Higher Education (AAHE).

- Ewell, Peter T. (2002a). An Emerging Scholarship: A Brief History of Assessment. In T.W. Banta and Associates, *Building a Scholarship of Assessment*. San Francisco: Jossey Bass, 3–25.
- Ewell, Peter T. (2002b). A Delicate Balance: The Role of Education in Management. *Quality in Higher Education*, 8, 2, 159–172.
- Gold, Bonnie, Sandra Z. Keith, and William A. Marion, eds. (1999) *Assessment Practices in Undergraduate Mathematics*. MAA Notes, Vol. 49. Washington, DC: Mathematical Association of America.
www.maa.org/saum/maanotes49/index.html.
- National Governors Association (1986). *Time for Results: The Governors' 1991 Report on Education*. Washington, DC: National Governors Association.
- National Institute of Education, Study Group on the Conditions of Excellence in American Higher Education (1984). *Involvement in Learning: Realizing the Potential of American Higher Education*. Washington, DC: U.S. Government Printing Office.
- Schapiro, Nancy S. and Levine, Jodi H. (1999). *Creating Learning Communities: A Practical Guide for Winning Support, Organizing for Change, and Implementing Programs*. San Francisco: Jossey-Bass.
- U. S. Department of Education, National Commission on Excellence in Education (1983). *A Nation at Risk: The Imperative for Educational Reform*. Washington, DC: U.S. Government Printing Office.