

THE EXAMINED LIFE:  
ASSESSMENT AND THE ENDS OF GENERAL EDUCATION

Peter Ewell

AAC&U GENERAL EDUCATION AND ASSESSMENT CONFERENCE  
March 6, 2004

I am particularly pleased to have the opportunity to participate in this gathering on assessment and general education because it's been so long in coming. AACU was one of the few lonely voices calling for this conjunction some twenty years ago in *Integrity in the College Curriculum* and has since been advancing it more concretely through *Greater Expectations*. The vision then and now demands essentially three things, that like the components of liberal education itself, are intertwined and inseparable.

1. A clearly-articulated, collective, and detailed vision of the attributes and abilities of an educated person.
2. Coherence and connectedness in the educational experiences that are intended to yield these ends.
3. Defensible evidence of their accomplishment.

Others have said similar things from the beginning of this period—most notably, the members of the Study Group on the Conditions of Excellence in Higher Education who authored *Involvement in Learning* in 1984. Unfortunately, many more disagreed with these arguments, or simply ignored them, and institutions did little to systematically change their instructional practices. And the reasons are not hard to see. Change was hard and the need for it unclear. More importantly, perhaps, in the press for implementation we lost the integration that made the

original vision so compelling. Parts of the agenda were implemented in some places—comprehensive assessment, active learning strategies, or service learning and learning communities—but not in the coherent fashion the vision intended.

Twenty years is a long time to keep the faith about something, and I am personally glad to surround myself with others of like mind about the proper scope and role of assessment—if only for mutual consolation. But I feel that we are at a particularly important crossroad in this conversation because of the intersection of two opposing forces that are increasingly shaping our actions today.

- The “bad news” (as the write-up to this session implies)—is that a particular—and forceful—ideology of assessment linked to accountability is being advanced in Washington and elsewhere as a recipe for “fixing” education at all levels. This vision, of course, presents an “Examined Life” which is emphatically *not* worth living, but we ignore it at our peril.
- The “good news” is that we are finally developing confidence in a language and experience with a set of approaches that are capable of making scholarly evidence of student academic achievement beyond grades both credible and academically acceptable. But in the light of the “bad news” we have little time to lose in pressing to make this vision a reality.

With appropriate emphasis on the “good news,” it’s the conjunction of these two forces that I want to unpack with you today in the particular context of general education.

Let me begin with a brief discussion of external forces—the “bad news” story—because understanding how we got here is critical to moving forward. And, in fact, the account of how we got here is pretty complicated, and contains some important cautionary tales.

When the national “assessment movement” began in the mid-1980s, those of us who were involved with it shared some implicit beliefs about assessment’s potential to transform teaching and learning. At the risk of presumption, I believe the most important of these were the following:

- That a focus on learning outcomes would lead faculty and institutional leaders to rethink—and ultimately restructure—curricula and pedagogy.
- That evidence-based inquiry about teaching and learning would be construed as responsible scholarship—faculty would want to engage in it and institutions would welcome the chance to get beyond “management by anecdote.”
- That the results of such inquiry would be used for continuous improvement and the resulting “feedback loop” would constitute a kind of institutional learning that would be iterative and ongoing.
- That “assessment” as a body of technique would only have meaning in the context of a larger process of institutional metamorphosis involving such things as establishing high expectations for students, promoting engaging pedagogies, and transforming faculty work.

These core beliefs were, I think, implicit in many of the early writings about assessment in the 1980s. In retrospect, they were pretty naïve. What happened instead is that many (if not most)

institutions were captured by a model of assessment that was excessively mechanical and process-centered—one that does *not* embody the kind of ongoing institutional learning that we hoped would develop. There were a couple of reasons for this, I think.

First, through too early emphasis on technique, we let assessment get excessively distanced from the day-to-day business of teaching and learning. Not seeing the connection—and rightly fearing an awful lot of work—many faculty started avoiding the topic entirely. Second and not unrelated, assessment simultaneously became something that external bodies—state governments and accreditation organizations—started forcefully asking institutions to do. The fact that an activity—not seen by faculty as intrinsically important or engaging in any case—seemed only to be done at the behest of somebody else exacerbated the first tendency.

Assessment increasingly became the “A-word”—an *add-on* to teaching and learning, rather than something integrally related to it, and the province of bureaucrats and bean-counters.

The fact that so many institutions failed to seriously engage assessment under these conditions is perhaps understandable. But it has resulted in a serious shortfall in our current ability to answer legitimate public questions about results. And into this vacuum has come a powerful contrasting ideology about how to answer these questions in a different way, embodied in such initiatives as *No Child Left Behind*.

While it seems clear that the NCLB approach will not completely dominate conversations about the Higher Education Act Reauthorization Act as some of us feared at the outset, it does represent an extremely politically-popular approach that also characterizes the policy attitudes of the majority of state and federal lawmakers. For this reason, it requires a serious higher education response. And there are particular features of NCLB’s approach to policy that make it

so politically compelling—features that we will have to do our best to match in any alternative we advance.

- First, it sets clear and visible standards of academic achievement. Though we may not like an approach that is based almost entirely on external, standardized tests (and in fact such an approach is highly unlikely at this point for higher education), our alternatives must be equally public about our expectations for learning and the extent to which we have achieved them.
- Second, it is consequential about performance. Schools that do not perform well face substantial negative consequences as a result—a form of “accountability that, like hanging, surely “focuses the mind.” Our alternatives, while hopefully neither draconian nor negative, need to be crafted in such a way that they equally recognize publicly cases in which achievement may have fallen short, and what we are going to do about it.
- Third, this approach is deliberately comparative—driven by a single commensurable yardstick of performance to which all are held accountable. While we may in some ways legitimately object that the missions of colleges and universities differ, and that the resulting diversity is a significant strength of higher education in America, our alternatives need also to be able to respond to charges that we are simply using our variety to duck the fact that there are certain levels and types of attainment that ought to characterize the recipients of every baccalaureate degree.

Meeting these conditions in an alternative approach to assessment in general education will naturally be a tall order. But I think that we can make a start at it, building strongly on the

framework provided by *Greater Expectations* and on some of the techniques that we've seen demonstrated so effectively at this conference.

So for the bulk of my time with you I want to reflect on how we can reaffirm the values we started with in assessment twenty years ago—before it became the “A-word”—in the light of growing experience, new technique, and the rising legitimacy of the scholarship of teaching. And I want to do so in a manner that is strongly reminiscent of the architecture of *Greater Expectations*. So I will organize these reflections about linking assessment and general education around not one “A-word” but *four*—Abilities, Alignment, Assessment, and Action. And I want to contend that knitting these four together yields a final “A”—Accountability—*not* in the narrow sense of *No Child Left Behind*, but in the original professional sense of assuming genuine collective responsibility for the academy and its integrity.

I will begin with Abilities—by which I mean the clear specification of what a student ought to be able to know and do—as well as the qualities of educational intentionality and connectedness that we want her or him to embody—as a result of general education. Let me offer two initial observations here. First, in an admirable attempt to be comprehensive about what it means to be an “educated person” we tend to name an awful lot of them. There are, for instance, some twenty-one separate abilities named in *Greater Expectations*, all of them necessary and noble. But from the perspective of public communication alone, we need to get behind a few of them and really work on them as our primary common assessment agenda for general education—communications, diversity, or information and quantitative literacy, for example. The result might be to crystallize opinion and mobilize action as the National Education Goals did publicly some fifteen years ago when it called upon the nation's colleges to graduate students who could “think critically, communicate effectively, and solve problems.”

Second, it's often claimed that the ends of general education embodied in these named abilities are excruciatingly difficult to describe and assess—and are quite different from the outcomes of specialized knowledge or training. I'm not so sure. Admittedly, we are talking about an incredible range of things, some of which appear at first glance pretty ineffable. *Greater Expectations*, for example, notes a set of “intellectual and practical skills” that range from the relatively straightforward—like communication and problem-solving—through far more complexly-interpreted abilities like [my personal favorite] “intellectual agility and managing change.” At the same time, it embraces personal qualities like intellectual honesty, active civic participation, and self-understanding. But I think the case is compelling that these should not be seen as stand-alone qualities of a “generally educated” person. They are also aspects of personal and professional mastery as described in the marvelous literature describing the differences between novice and expert practice. Judith Ramaley reflected just this insight in the form of “Pasteur’s Quadrant”—noting the similar danger of seeing general education outside a what she called a “use-inspired” perspective.

What this says to me is quite simple—though I see the conclusion violated in lots of places. The assessment of general education is not about examining the consequences of any particular body of coursework that may carry these words as a label, but of the undergraduate experience as a whole. This simple conclusion also has a real implication for practice. It's not enough to examine general education in the class-based settings of core or distribution requirements, however embedded and authentic these assessments may be. We must also examine how these abilities infuse and inform expert practice. The assessment of general education, in short, must be integrally linked to the major.

In the same vein, I'd like to offer two additional lessons that I think I've learned about the slippery matter of articulating abilities. First, it's most useful to start with the *actual practice* of the ability and not what's stated as the outcome. Phrases like "intellectual agility" have great charm, but they mean very little in the absence of an actual student performance that might demonstrate it. Formal assessment design, as stated in the textbooks, demands ever more detailed verbal specifications of outcomes or competencies to be developed as guides for constructing assessment techniques. But with my respected colleague Grant Wiggins, I often find it more helpful to go the other way. Given a broad general descriptor like "intellectual agility," can you imagine a *very concrete* situation in which somebody might display this ability and how they might actually behave? Better still, can you quickly specify the parameters of an *assignment or problem* that might demand a particular level of this ability in order to perform successfully? Grant and I often stir things up in assessment conversations by saying that we are pure "operationists" in this manner, in the classic scientific sense of the term: the performance that the student exhibits on the assessment is the definition of the ability itself—it has no independent existence.

Starting with the concrete in this manner can also quickly uncover lurking inconsistencies in the way each of us thinks about a given ability that might lead to real trouble later on. Nowhere—at least for me—has this been more apparent than when we try to communicate what we mean by particular general-education-related abilities to external stakeholders. For example, I've found that in conversations about attributes like "critical thinking" with business leaders and academics in the same room, they quickly concur that graduates ought to have some it. But when you describe the ability in action—what its exercise actually looks like in a real situation—they may have quite different views about what it actually means. Academics, for instance, tend to anchor

the ability on things like recognizing flaws in reasoning in a presented case, examining evidence, and comparing/contrasting two lines of argument. Stakeholders, in contrast, tend to emphasize performance in very ill-structured situations—finding the “problem to solve,” for example—and knowing how to work around missing information or when to stop after identifying a “good enough” solution. These are quite different abilities, masquerading under a common name.

From an accountability standpoint, the implication is clear. Rather than saying simply that our graduates achieve at “level 4” in Critical Thinking (or some such ability), we need to succinctly describe exactly what they were asked to do and how well they performed. A good illustration here is the way the National Adult Literacy Survey was first constructed some fifteen years ago. Its designers did not start with a set of abstract abilities and then try to build an assessment to determine if people had them. They instead began with a set of very concrete tasks that citizens ought to be able to perform, like read a bus schedule or interpret a graph, and built them into a compelling set of tasks that any stakeholder could instantly see were important. By doing so, they essentially proclaimed, “we hold these tasks to be self-evident.” This strategy worked very well for public communication. And, as I will describe later, it has clear implications for the kinds of assessments we construct.

Let me now move on to some reflections on the second “A-word” in my mantra—Alignment. Perhaps the greatest disconnect that I see out there in this realm is the fact that too often our statements of desired outcomes for things like general education—however they are defined—are not meaningfully connected to curricular structure or pedagogy. In fact, a major irony I’ve run into on too many campuses is that statements of intended outcomes for general education were invented primarily to guide *assessment*, not teaching and learning practice. This curious condition, of course, is in part a product of assessment being done at the behest of others.

Accreditation agencies require “assessment,” and this appears predicated on the development of outcomes statements. They ask for them, so we have them. But there are few accreditation requirements as yet that they be connected to anything. [Recently, as embodied in the report of *Greater Expectations*’ Project on Accreditation and Assessment, some accreditors have recognized this problem and are beginning to fix it.]

The popular move toward more “embedded” forms of assessment in general education is also beginning to address this issue because it forces the institution to “map” the places where particular abilities are taught—and therefore where they might be assessed already—in a much more explicit way than before. And this exercise also starts to address a major gap in our knowledge about most general education curricula—the fact that we don’t really know exactly *how* they are working. Note that this is not the same thing as knowing *how well* they are working with respect to fostering desired abilities. It’s asking a far more basic question about the extent to which the curriculum is being acted out by students consistent with the “design for learning” that the faculty originally intended. Faculty are generally great designers when it comes to general education, and some of the curricular schemes I’ve seen out there are masterpieces of complexity and beauty. But there is often very little information available about the ways students are making course choices and with what consequences. Studies along these lines can be extremely fruitful in general education and at our Center we have ended up doing quite a few. Moreover, as Judith Ramaley reminded us from this podium earlier in the conference, the problem of the behavioral curriculum is now even more complicated by the growing incidence of “multiple pathways” students.

A first point of advice here is begin to sort out the many purposes of what we call general education—and the specific roles that particular learning experiences are intended to play. Let me note at least four of these purposes here, which are really quite distinct:

1. Prerequisite-skills development that will be needed for later work (e.g., in communications skills or in math).
2. Development of cross-cutting abilities, like critical thinking or problem-solving, that should be continuously reinforced throughout the entire curriculum.
3. General knowledge about particular disciplines and experience with different modes of inquiry.
4. Collegiate socialization—learning how to “do college” by acquiring skills like knowing how to use a library (or the Web), or how to plan and carry out an independent intellectual project.

Distinctions among these various purposes ought to dictate where in the curriculum particular learning experiences are located. We saw several examples of this recognition here as part of campus assessment approaches. One was UNC-Asheville’s careful longitudinal look at how students were developing particular attributes over time. Another was Hampshire College’s recognition that bowing entirely to self-selected student choices was not working well in the first year of study. A final example, at the opposite end of the policy continuum, was the need for a clear distinction between lower-division and upper-division general education abilities in the session on state-level transfer policies in Missouri.

I've seen enough data on student coursetaking patterns at many institutions to know that students don't always follow our prescriptions—choosing to postpone or avoid math or writing, for example. Such analyses of how students “act out” general education also frequently reveal an astounding lack of coherence because courses taken sequentially aren't really connected to one another. One consequence is there is often a lot of “re-work” needed—faculty being forced to readdress topics in a subsequent course that they presumed students had learned in a prerequisite. Such analyses may also reveal that a particular course is not doing the job assigned to it in the curriculum for one reason or another. For example, if we are successful in our application to NSF, NCHEMS will be helping the Mathematics Association of America undertake an in-depth look at College Algebra to determine its downstream consequences at about a hundred institutions over the next two years.

Finally, taking alignment seriously has implications for *how* things are taught, not just curricular structure and content. An early lesson on this for me in assessment was an attempt at the College of William and Mary to assess critical thinking through a systematic cross-section of student writing across courses that required it. Faculty there found that they simply couldn't evaluate critical thinking in a majority of cases because the assignments generating the student writing didn't ask for it. This kind of systematic sampling of student work can be very effective in determining if we are actually addressing many aspects of general education that we hope are being taught out there. For example, Truman State University recently used the technique to look for the incidence of spontaneous inter-disciplinary connection-making occurring in student work. I could cite many other examples of this kind. But all of them emphasize the point that “assessment” in general education should address issues of alignment explicitly and directly by generating evidence of student behavior, not just examine outcomes.

This last point, of course, leads me to the third anticipated “A-word, “Assessment, which remains resolutely focused on outcomes—the extent to which abilities are attained. Many excellent examples of assessment technique have been amply covered at this conference, so I want to confine my reflections here to some major choices in design and approach—and how these choices may be particularly conditioned by the need to provide credible evidence of student academic achievement for external constituencies. And to begin with here, it is useful to distinguish two quite different overall approaches. With apologies to the Biologists, let me for sake of convenience name them after two predominant animal body plans—“exoskeletal” and “endoskeletal.” *Both* metaphors, I want to emphasize, may be helpful for developing effective assessment in the context of General Education.

The “exoskeletal” approach is probably the most easily explained, as it resembles the classic early picture of assessment that most campuses adopted in the mid-1980s—usually in response to a state mandate or emerging accreditation requirements. [Elsewhere, I have labeled this approach the “Evaluative Paradigm.”] Under this approach, assessment opportunities are created *outside* and *alongside* the regular curriculum. They are intended to be special and infrequent occasions when student learning, usually in the aggregate, can be examined with greater deliberation and precision than in the standard classroom grading situation. Indeed, it was precisely because grades had become suspect as reliable indicators of student learning that this more “scientific” alternative was evolved twenty years ago. Its early manifestation in General Education settings tended to center on externally-obtained examinations and “Assessment Days.” But more authentic versions now center on techniques like sophomore writing exercises and portfolios. These newer approaches are far more compelling and faculty-centered than simply having students take the Academic Profile or ACT-CAAP, but they still stand outside the regular

day-to-day process of teaching and learning. We have had several good and diverse examples of such approaches demonstrated here, ranging from standard assessments of information literacy at James Madison University to Hampshire's use of student advisors as assessors.

Several respected and venerable curricular designs in General Education employed an "exoskeletal" approach to student assessment. Classics include the Core at the University of Chicago in the 1930s which featured general examinations in various fields (developed by Benjamin Bloom of "Bloom's Taxonomy," among others), and the experience of many small liberal arts colleges through the 1950s that administered ETS' long-defunct "Area Exams" to all students in addition to a culminating piece of scholarship in the major. And it's important to stress that although these testing occasions occurred in addition to classes, they were considered an integral part of the design for general education. Students expected them, and performance on them was consequential. There is a lot of wisdom of these historic approaches that would be good to imitate.

The "endoskeletal" approach to assessment has evolved more recently, largely in an attempt to escape the worst consequences of the "add-on" mentality. Here, the attempt is to as fully as possible "embed" occasions for assessment seamlessly throughout a given student's learning experience as she or he progresses. Under one version of this approach, specific assignments that explicitly address particular abilities are strategically laced throughout General Education coursework and are intended to do "double duty." They are a regular expectation of students and are graded as such in context. But the standards and methods used to examine the resulting student work are keyed not just to what the individual faculty member expects. Instead, they are disciplined by a set of collectively agreed-upon set of rubrics or grading guides that both align faculty judgments and enable the institution to determine aggregate patterns of strength and

weakness with respect to the ability in question. Other versions of this approach relax the constraint of common rubric-based grading by allowing faculty to grade in their own way for the record, scoring only a sample of student work products at a later point to garner assessment information. Writing is the ability most frequently examined in this way. The advantage here, of course, is that assessment activity is much closer to the business at hand and is entirely transparent to students. And creating the design in the first place helps ensure curricular coherence. But the challenge, as the experience of San Jose State presented at this conference demonstrates, is to round up all these judgments and be consistent.

Like all metaphors, the skeletal one is imperfect and there are many sound “intermediate forms” in between. The point is simply to be quite deliberate about deciding questions of assessment as an integral part of curricular design. Let me try to illustrate this point with two examples drawn from quite different institutions that I’ve had the opportunity to work with in just the last month—one a major public research university and one a well-known liberal arts college. Both are in the process of developing a new approach to general education, with the resulting opportunity to “engineer” assessment into their designs from the outset. And both are very much aware of the need for the results of any assessment to be publicly-credible. While both these conversations are at an early stage of development, some interesting ideas have already emerged. At the research university, expectedly, there are enormous concerns about faculty workload. So with an unusually complex new General Education curriculum to implement, most of the conversation so far has centered on the notion of identifying a small set of embedded assignments in diverse General Education courses mapped to previously-defined abilities. Part of the key to making this work will have to be a new approach to transcribing student performance on these assignments so they can be quickly aggregated to show how and at what

level students have met agreed-upon standards. While many campuses have attempted the embedded design, only a few have really thought through the challenges involved in aligning and aggregating the substantial volume of faculty judgment that this entails. Fortunately, technology is proving useful here, as illustrated at this conference by Hampshire's use of "The Hub"—a central site where faculty can record and aggregate complex judgments about individual student performance in a grade-less environment.

At the small liberal arts college, the proposed curricular design is more radical, but both size and tradition may be inclining it toward a more "exoskeletal" approach. Indeed, one member of the Commission that is evolving the new design who had attended the College in the 1950s recalled an "assessment regimen" (though nobody would have called it that then) that was pretty remarkable. At mid-point everybody took the ETS "Area Exams." And as juniors, they participated in a faculty-designed scenario-based exercise that was astoundingly creative. As he described it, it included tasks like being placed in a room with period furniture and a dining room table set, and being asked to determine what historical period you were in, what music might be playing, and what might be talked about over dinner that night. Finally, as seniors, they were required to complete and publicly present a thesis in their major field. Stirred by that tradition, the college is exploring the development of similar authentic, scenario-based cases, benchmarked to bodies of content and developmental goals given at different points. This might culminate in a common "senior problem" that all would contribute toward solving—both individually and in small work teams associated with their majors. The chosen topic would be deliberately "real-world," consistent with the new curriculum's emphasis on experiential learning communities—for example, implementing the Kyoto Protocol or rebuilding postwar Iraq.

These two institutional examples, of course, are quite different and reflect the particular circumstances and values of their institutional settings. But both are thinking appropriately about how to integrate from the outset clear collective goals for student learning cast in “ability” terms, a curricular structure designed explicitly and visibly to foster these abilities, and a mechanism to generate appropriate and credible evidence of student achievement—precisely the three “A-words” of my prior discussion.

Which brings us at long last to the final “A-word”—Action. Here I have always believed that the principal problems surrounding assessment—like the far larger challenge of evolving effective pedagogies that are consistent with what we know about learning—is not a matter of knowledge but of the will to act. Under current circumstances, we have little choice but to be more aggressive in presenting our own response to legitimate public questions about whether we know anything collectively about student achievement. If we do not do so soon, we will lose our place in the conversation. I do believe we have some time and opportunity to do this—but admittedly not much time—to evolve the two things we most need: a set of effective institutional solutions as examples to illustrate that such solutions are possible and that faculty will seriously embrace them, and a common policy voice that leads with an alternative, not a protest.

The first can be founded on some of the many rich and effective individual campus approaches that we’ve seen at this conference. But the second, I believe, will have to include at least some elements of an “exoskeletal” approach based on common culminating demonstrations of student achievement benchmarked to national standards. Part of the latter may come in the form of initiatives like RAND/CAE’s “Value Added” project based on authentic task-based assessments administered across institutions that was featured in a recent issue of *Peer Review*—today’s equivalent of the ETS Area Tests. Others may come in the form of the Pew Forum on College-

Level Learning’s five-state demonstration project, which uses some of these same instruments to establish aggregate benchmarks at the state level for policymaking, but avoids institutional comparisons. This approach has the advantage of establishing and tracking clear public standards of achievement but allowing each campus to make its own contribution to the work regardless of selectivity.

But clearly, our main ally in this public policy effort will be the accountability mechanism that remains most true to what we want to accomplish—regional accreditation. Although institutions sometimes given them little credit for it, regional accrediting organizations have undergone their own significant transformation in the last five years. And a lot of these changes (as institutions have complained) are in the realm of assessment. But I firmly believe that what the accreditors are asking for is exactly what we need to make a public case. Since we are in California, consider how the Standards presented in the WASC’s Senior Commission’s *2001 Handbook of Accreditation* embodies both the “exoskeletal” and “endoskeletal” imperatives.

CFR 2.1: All degrees...should be defined in terms of levels of achievement that represent more than just the accumulation of courses or credits.

CFR 2.6: The institution demonstrates that its graduates consistently attain its stated levels of achievement and that its expectations for student learning are embedded in the standards faculty use to evaluate student work.

These are pretty strong requirements, and they were put in place largely because the Commission wanted to set a bold, learning-centered tone from the outset—and grow into it. But in the background to these changes—as well as similar new directions being taken by other accrediting

bodies—is their need to walk an increasingly perilous line between taking a greater role in promoting authentic, learning-centered approaches to assessment and meeting federal obligations to “inspect” institutional quality in their assigned role as “gatekeepers” for federal financial aid dollars. Because of the latter role, accrediting organizations are likely to be under increasing pressure to be hard on institutions as the Reauthorization process unfolds. We desperately need to support the accreditation community in this effort by recognizing that what they are asking of us is what we should be expecting of ourselves. If institutions resist or respond mechanically to accreditors’ requirements that we seriously document and align our efforts to generate evidence of student academic achievement, as I see too many doing, we only cheat ourselves. More importantly, we deprive ourselves of the best public weapon we have to make our case.

Abilities, Alignment, Assessment and Action only *together* will give us the right kind of Accountability. And though the imperative to act is in part spurred by necessity, if we act right we may at the same time attain the institutional version of the “examined life” which, consistent with the ends of General Education, we work so diligently to foster in our students. So as you make your way home from Long Beach, I urge you to reflect on how you can knit these four “A-words” together on your own campuses to deepen the integration of assessment throughout general education, and how we can increasingly make common cause to get this message heard. Thank you, and the job continues Monday!