

The Course Design Survey

May 11, 2014 by David Dirlam

In a recent post to the ASSESS listserv Ephraim Schechter proposed an elegant solution to the problem of public disclosure impacting assessment bias. 'real accountability includes the also reporting the data's impact on planning. Even when the answer is "we didn't change anything because the data said we shouldn't," making that explicit still shows that our planning includes paying attention to how we're doing.' In short, programs can be proud when their assessment results help them make discoveries about learning. What follows is a way to ensure that they will make discoveries: i.e., given a reasonable sample size (and our college of 1400 is plenty big) the probability of not discovering something reduces nearly to zero.

There were three steps from two different institutions in creating an impossible-to-avoid-discovery design. First, I came across this bias-in-reporting problem several years ago at Hebrew Union College when we implemented a Learning Outcomes Network (LON). A LON involves evaluating every student in every course in a program using the same multidimensional rubric in which every level of every dimension (we use Beginning, Fundamental, Practical, and Inspiring) describes a unique learning outcome. With this data it was possible to calculate both a reliability score and an impact score for all but capstone courses. Both calculations require a comparison across predecessor and successor instructors. If an instructor rates most of his or her students higher on a dimension than all predecessor instructors, then there are two interesting possibilities for successor raters. On the one hand, if the successors rate the students the same as the predecessors (meaning lower than the instructor in question), then either the instructor had too rosy an idea of the student progress or the learning that was used for the rating was not sustained. On the other hand, if the successor instructors agreed with the higher ratings, then the course in question had a high impact on learning within that dimension. The trouble with reporting impacts, however, was what happens when a course had no impact? My solution was to get permission to report the impact results only to the instructor of the course in question. I was granted that permission and carried the problem to my next place of employment, Virginia Wesleyan College, where I was granted the same permission. This is the same problem, on an individual level as sharing assessment results, on line.

The second and third steps for creating the impossible-to-discover-nothing design occurred at VWC. One of the things that attracted me to the college was that the faculty had very recently undergone a wholesale curriculum revision from five three-credit courses to four four-credit courses and for every course change they had identified which of eleven "enhancements" (plus "other") would account for the additional credit hour. After a year of working toward Learning Outcomes Networks a faculty committee identified that we could solve the problem of reporting course impacts by focusing instead on educational enhancement practices that were used across courses. We could calculate the impact of practices rather than the impact of courses. When a practice was used multiple times and found to have no impact, instructors would be much less defensive than if their courses were found to have no impact. They could keep the course and change the practice--exactly the kind of outcome that excites assessment researchers.

However, a third problem became immediately apparent. One committee member, our Director of General Studies, had helped to create the list of enhancements and criticized it as being mostly "seat-of-the-pants" and requiring a more careful look. George Kuh's "high impact practices" were certainly interesting in this regard, but most of them were in the list that the committee found unsatisfactorily abrupt. The solution was prompted by Robert Zemsky's sage advice in his Checklist for Change: "It is advantageous to disaggregate the traditional instructional format into a set of more or less discrete activities."

We in the assessment community have been disaggregating learning for decades, but few of us have systematically disaggregated instruction. I set about identifying six dimensions with a few levels of each: (1) locations, (2) social contexts, (3) instructor roles, (4) student resources, (5) student objectives, and (6) student preparation strategies. I showed the form to our faculty committee just yesterday and they not only came up with a name "The Course Design Survey", they enriched it to six or seven categories plus "other" for each dimension. If instructors identify which of 5 levels of emphasis (from major to none) for each course design strategy, there will be 10^{27} possible patterns of strategies--certainly better than 11. The huge number of patterns is equivalent to the number of grams of mass in the earth. We can look for high probability patterns of the 40 components across any or all of the programs in the college. Given the rich data that we get from our LONs, the odds of us discovering some approaches that work better than others are astronomically good.

The Course Design Survey leaves faculty free to design courses as they see fit and to change course designs from one term to the next. Given the power of the novelty effect in educational research, we should not expect that our solutions would often be permanent or universal. But the survey takes a minor fraction of an hour, and the LON ratings only one or two minutes per student. Both are small fractions of the time it takes to write a syllabus or to compile final grades. And the solutions should be useful not only to us, but to other institutions.

The key to public disclosure, as Ephraim pointed out, is discovery. It needs to happen and we need to share it. Combining LONs with Course Design Surveys provides a powerful method for enhancing both.
