

Finding Out If Your Course Is Working for Your Students

When teaching a new course, offering a new learning experience within the course, or teaching for students with different backgrounds than in the past, it is a good idea to do some checking to make sure the course is working for students. This is a little different than just gathering evidence that the course was effective overall. Educational measurement folks refer to efforts to find out how well a course worked generally as *summative* assessment. They refer to efforts to find out what needs to be tuned in a course as *formative* assessments. This note is about formative assessment.

We do formative assessment all the time in daily life. For example, when driving down the road, we continually look at where the car is going and make small corrections to our steering to keep the car on the road. That works best when we apply what we see quickly. One way that experienced drivers are different from novice drivers is that they make small corrections continually, based on quick glances. They may never know perfectly how much to turn the wheel to stay centered in a lane, but they glance and correct often enough that their “formative” assessments of how well they are staying in lane don’t need to be all that perfect.

Successful enterprises do similar continual steering assessments. A good example is hotel chains. Often, if you stay in a hotel run by a large chain, you will get an email afterwards asking you to just rate the quality of your accommodation on a 1-5 scale. You might think this to be meaningless, and you are partly right. People actually are pretty nice on such surveys, and the hotel almost always will get mostly 4’s and 5’s back. That is not important to them. What is important is if the average rating suddenly changes. If that happens, good managers will do some quick checking to see if they can find evidence for what caused the problem. Sometimes, it will be an emergent event, like a power outage that inconvenienced a lot of guests. Other times, though, more checking will be required to find out what caused the change in satisfaction.

You can do the same thing in your teaching of a large course. Pitt has software licenses for several tools that make it easy to quickly and anonymously sample student opinion on how a course is going. It is possible to do such sampling with Top Hat, Qualtrics, or Canvas. Each has advantages and disadvantages. If you already use Top Hat to get student responses to questions during class, you might find it the most convenient. Also, compared to Canvas, Top Hat currently provides a stronger assurance for students that their responses to a survey labeled as anonymous really is anonymous.

With almost no work by you or your students, it would be easy to ask students, perhaps at the beginning of the first class session each week, to rate how well the course is allowing them to learn. You could use a five point scale that might look something like this:



With a little experience, you should have a good sense of what level of average response to expect. Then, if you hit a week where the average is lower than usual, you can quickly follow up. There are two ways you might follow up. One is to simply ask students to respond with what sorts of problems they had with the material just covered. Top Hat, for example, allows you to ask a long-response question, and students can just try to tell you what was problematic or where they are confused. This can allow a timely repair. Suppose you give the general 1-5 item and you get a lower than expected response level. You could have a “what went wrong” item prepared in advance and just enable it right away, get student inputs, and come to class the next time ready to better explain something students struggled with or provide a different exercise to help students work further to grasp a concept.

An alternative to an open-ended follow-up question might be to have students provide 1-5 ratings for each of the several learning activities you offered during the previous week. This is easier for students to do, although it does not give you any specific clues about where a problem might lie. Still, it can be helpful to narrow down student problems to a specific learning activity if there are multiple possible tasks that students might have found problematic.

The basic strategy for continual improvement of large-enrollment courses, then, looks like this:

1. Get regular feedback on how things are going for students, and watch for a drop in students' feeling of doing well in their learning.
2. If necessary, do some probing to see which specific learning activity is not working as well as needed.
3. Then, seek student input on how that activity might serve them better or, at least, where the specific problem might be.

Using this approach, you are, in essence, applying to teaching the same kinds of continual improvement strategy that successful enterprises use, and it is likely to produce good results, just as it does outside of academe.