

Randomized
trial compares
hybrid learning
to traditional
course

Online Learning in Higher Education

Higher education in the United States, especially the public sector, is increasingly short of resources. States continue to cut appropriations in response to fiscal constraints and pressures to spend more on other things, such as health care and retirement expenses. Higher tuition revenues might be an escape valve, but there is great concern about tuition levels increasing resentment among students and their families and the attendant political reverberations. President Obama has decried rising tuitions, called on colleges and universities to control costs, and proposed to withhold access to some federal programs for colleges and universities that do not address “affordability” issues.

Costs are no less a concern in K–12 education. Until the 2008 financial crisis and the subsequent slowdown in U.S. economic growth, per-pupil expenditures on elementary and secondary education had been steadily rising. The number of school personnel hired for every 100 students more than doubled between 1960 and the first decade of the 21st century. But in the past few years, local property values have stagnated and states have faced intensifying fiscal pressure. As a result, per-pupil expenditures have for the first time in decades shown a noticeable decline, and pupil-teacher ratios have begun to shift upward (see “Public Schools and Money,” *features*, Fall 2012). With the rising cost of teacher and administrator pensions, the squeeze on school districts is expected to continue.

A subject of intense discussion is whether advances in information technology will, under the right circumstances, permit increases in productivity and thereby reduce the cost of instruction. Greater, and smarter, use of technology in teaching is widely seen as a promising way of controlling

costs while reducing achievement gaps and improving access. The exploding growth in online learning, especially in higher education, is often cited as evidence that, at last, technology may offer pathways to progress (see Figure 1).

However, there is concern that at least some kinds of online learning are of low quality and that online learning in general depersonalizes education. It is important to recognize that “online learning” comes in a dizzying variety of flavors, ranging from simply videotaping lectures and posting them online for anytime access, to uploading materials such as syllabi, homework assignments, and tests to the Internet, all the way to highly sophisticated interactive learning systems that use cognitive tutors and take advantage of multiple feedback loops. Online learning can be used to teach many kinds of subjects to different populations in diverse institutional settings.

Despite the apparent potential of online learning to deliver high-quality instruction at reduced costs, there is very little rigorous evidence on learning outcomes for students receiving instruction online. Very few studies look at the

by WILLIAM G. BOWEN, MATTHEW M. CHINGOS, KELLY A. LACK, and THOMAS I. NYGREN

use of online learning for large introductory courses at major public universities, for example, where the great majority of undergraduate students pursue either associate or baccalaureate degrees. Even fewer use random assignment to create a true experiment that isolates the effect of learning online from other factors.

Our study overcomes many of the limitations of prior studies by using the gold standard research design, a randomized trial, to measure the effect on learning outcomes of a prototypical, interactive online college statistics course. Specifically, we randomly assigned students at six public university campuses to take the course in a hybrid format, with computer-guided instruction accompanied by one hour of face-to-face instruction each week, or a traditional format, with three to four hours of face-to-face instruction each week. We find that learning outcomes are essentially the same: students in the hybrid format pay no “price” for this mode of instruction in terms of pass rates, final-exam scores, or performance on a standardized assessment of statistical literacy. Cost simulations, although speculative, indicate that adopting hybrid models of instruction in large introductory courses has the potential to reduce instructor compensation costs quite substantially.

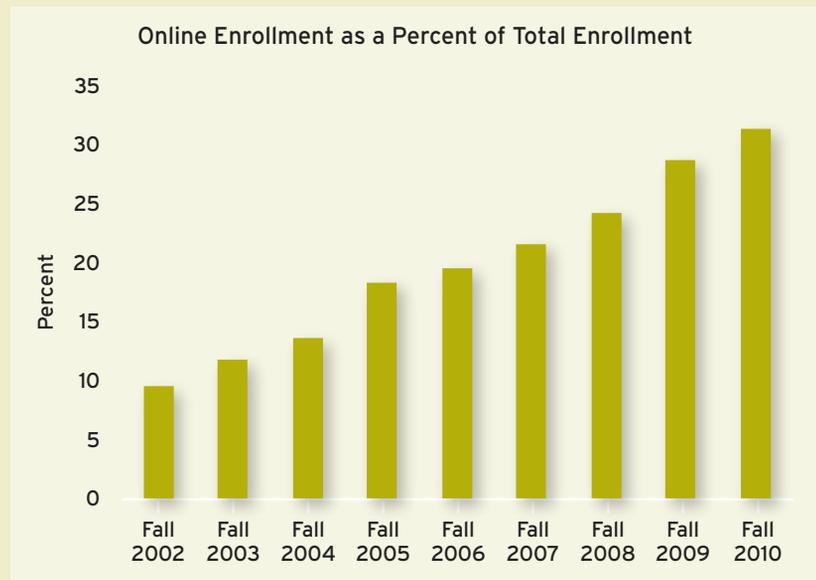
Research Design

Our study assesses the educational outcomes generated by what we term interactive learning online (ILO), highly sophisticated, web-based courses in which computer-guided instruction can substitute for some (though usually not all) traditional, face-to-face instruction. Course systems of this type take advantage of data collected from large numbers of students in order to offer each student customized instruction, as well as to enable instructors to track students’ progress in detail so that they can provide more targeted and effective guidance.

We worked with seven instances of a prototype ILO statistics course at six public university campuses (including two separate courses in separate departments on one campus). The individual campuses include, from the State University of New York (SUNY): the University at Albany and SUNY Institute of Technology; from the University of Maryland: the University of Maryland, Baltimore County,

Online Takes Hold (Figure 1)

The proportion of higher education students taking at least one online course has increased steadily over the last decade.



SOURCE: "Going the Distance: Online Education in the United States, 2011," Babson Survey Research Group

and Towson University; and from the City University of New York (CUNY): Baruch College and City College.

We examine the learning effectiveness of a particular interactive statistics course developed at Carnegie Mellon University (CMU), considered a prototype for ILO courses. Although the CMU course can be delivered in a fully online environment, in this study most of the instruction was delivered through interactive online materials, but the online instruction was supplemented by a one-hour-per-week face-to-face session in which students could ask questions or obtain targeted assistance.

The exact research protocol varied by campus in accordance with local policies, practices, and preferences, but the general procedure followed was 1) at or before the beginning of the semester, students registered for the introductory statistics course were asked to participate in our study and offered modest incentives for doing so; 2) students who consented to participate filled out a baseline survey; 3) study participants were randomly assigned to take the class in a traditional or hybrid format; 4) study participants were asked to take a standardized test of statistical literacy at the beginning of the semester; and 5) at the end of the semester, study participants were asked to take the standardized test of statistical literacy again, as well as to complete another questionnaire.

Of the 3,046 students enrolled in these statistics courses in the fall 2011 semester, 605 agreed to participate in the study and to be randomized into either a hybrid- or traditional-format section. An even larger sample size would have been desirable, but the logistical challenges of scheduling at least two sections (one hybrid section and one traditional section) at the same time, to enable students in the study to attend the statistics course regardless of their (randomized) format assignment, restricted our prospective participant pool to the limited number of “paired” time slots available. Also, student consent was required in order for researchers to randomly assign them to the traditional or hybrid format. Not surprisingly, some students who were able to make the paired time slots elected not to participate in the study. All of these complications notwithstanding, our final sample of 605 students is in fact quite large in the context of this type of research.

The baseline survey administered to students included questions on students’ background characteristics, such as socioeconomic status, as well as their prior exposure to statistics and the reason for their interest in possibly taking the statistics course in a hybrid format. The end-of-semester survey asked questions about their experiences in the statistics course. Students in study-affiliated sections of the statistics course took a final exam that included a set of items that was identical across all the participating sections at that campus. The scores of study participants on this common portion of the exam were provided to the research team, along with background administrative data and final course grades of all students (both participants and, for comparison purposes, nonparticipants) enrolled in the course.

very similar characteristics overall. The 605 students who chose to participate in the study also have broadly similar characteristics to the other students registered for introductory statistics. The differences that do exist are quite small. For example, participants are more likely to be enrolled full-time but only by a margin of 90 versus 86 percent. Their outcomes in the statistics course are also comparable, with participants earning similar grades and being only slightly less likely to complete and pass the course than nonparticipants.

An important limitation of our study is that while we were successful in randomizing students between treatment and control groups, we could not randomize instructors in either group and thus could not control for differences in teacher quality. Instructor surveys reveal that, on average, the instructors in traditional-format sections were much more experienced than their counterparts teaching hybrid-format sections (median years of teaching experience was 20 and 5, respectively). Moreover, almost all of the instructors in the hybrid-format sections were using the CMU online course for either the first or second time, whereas many of the instructors in the traditional-format sections had taught in this mode for years.

The “experience advantage,” therefore, is clearly in favor of the teachers of the traditional-format sections. The questionnaires also reveal that a number of the instructors in hybrid-format sections began with negative perceptions of online learning, which may have depressed the performance of the hybrid sections. The hybrid-format sections were somewhat smaller than the traditional-format sections, however, which may have conferred some advantage on the students randomly assigned to the hybrid format.

We examine the impact of assignment to the hybrid format, relative to the traditional format, on students’ probability of passing the course, their performance on a standardized test of statistics, and their score on a set of final-exam questions. We find no clear differences in learning outcomes.

The participants in our study are a diverse group. Half come from families with incomes less than \$50,000 and half are first-generation college students. Less than half are white, and the group is about evenly divided between students with college GPAs above and below 3.0. Most students are of traditional college-going age (younger than 24), enrolled full-time, and in their sophomore or junior year.

The data indicate that the randomization worked properly in that traditional- and hybrid-format students in fact have

Learning Outcomes

Our analysis of the experimental data is straightforward. We compare the outcomes for students randomly assigned to the traditional format to the outcomes for students randomly assigned to the hybrid format. In a small number of cases—4 percent of the 605 students in the study—participants attended a different format section than the one to which they were randomly assigned. In order to preserve the randomization procedure, we associated students with the

section type to which they were randomly assigned. This is sometimes called an “intent to treat” analysis, but in this case it makes little practical difference because the vast majority of students complied with their initial assignment.

Our analysis controls for student characteristics, including race/ethnicity, gender, age, full-time versus part-time enrollment status, class year in college, parental education, language spoken at home, and family income. These controls are not strictly necessary, since students were randomly assigned to a course format. We obtain nearly identical results when we do not include these control variables, just as we would expect given the apparent success of our random assignment procedure.

We first examine the impact of assignment to the hybrid format, relative to the traditional format, on students’ probability of passing the course, their performance on a standardized test of statistics, and their score on a set of final-exam questions that were the same in the two formats. We find no clear differences in learning outcomes between students in the traditional- and hybrid-format sections. Hybrid-format students did perform slightly better than traditional-format students on the three outcomes, achieving pass rates that were about 3 percentage points higher, standardized-test scores about 1 percentage point higher, and final-exam scores 2 percentage points higher, but none of these differences is statistically significant (see Figure 2).

It is important to note that these non-effects are fairly precisely estimated. This precision implies that if there had been pronounced differences in outcomes between traditional-format and hybrid-format groups, it is highly likely that we would have found them. In other words, we can be quite confident that the actual effects were in fact close to zero, and therefore differ from a hypothetical finding of “no significant difference” that may result from excessively noisy data or an insufficiently large sample.

We also calculate results separately for subgroups of students defined in terms of various characteristics, including race/ethnicity, gender, parental education, primary language spoken, score on the standardized pretest, hours worked for pay, and college GPA. We do not find any consistent evidence that the hybrid-format effect varies by any of these characteristics. There are no groups of students that benefited from or were harmed by the hybrid format consistently across multiple learning outcomes.

In addition, we examine how much students liked the hybrid format of the course, and find that students gave the hybrid format a modestly lower overall rating than their counterparts gave the traditional-format course (the rating was about 11 percent lower). By similar margins, hybrid students report feeling that they learned less and that they found the course more difficult. But there were no notable differences in students’ reports of how much the course raised their interest in the subject matter.

We also asked students how many hours per week they spent *outside of class* working on the statistics class. Hybrid-format students report spending 0.3 hours more each week, on average, than traditional-format students. This difference implies that in a course where a traditional section meets for three hours each week and a hybrid section meets for one hour, the average hybrid-format student would spend 1.7 fewer hours each week in *total time* devoted to the course, a difference of about 25 percent. This result is consistent with nonexperimental evidence that ILO-type formats can achieve the same learning outcomes as traditional-format instruction in less time, which has potentially important implications for scheduling and the rate of course completion.

Potential Savings

In other sectors of the economy, the use of technology has increased productivity, measured as outputs divided by inputs, and often increased output as well. Our study shows that a leading prototype hybrid-learning system did not increase outputs (student learning) but could potentially increase productivity by using fewer inputs.

It would seem to be straightforward to compare the side-by-side costs of the hybrid version of the statistics course and the traditional version. The problem, however, is that contemporaneous comparisons can be nearly useless in projecting long-term costs, because the costs of doing almost anything for the first time are very different from the costs of doing the same thing numerous times. This is especially true in the case of online learning, where there are substantial start-up costs that have to be considered in the short run but are likely to decrease over time. For example, the development of sophisticated hybrid courses will be a costly effort that would only be a sensible investment if the start-up costs were either paid for by others (foundations and governments) or shared by many institutions.

There are also transition costs entailed in moving from the traditional, mostly face-to-face model to a hybrid model that takes advantage of more sophisticated ILO systems employing computer-guided instruction, cognitive tutors, embedded feedback loops, and some forms of automated grading. Instructors need to be trained to take full advantage of such systems. On unionized campuses, there may also be contractual limits on section size that were designed with the traditional model in mind but that do not make sense for a hybrid model. It is possible that these constraints would be changed in future contract negotiations, but that too will take time.

We address these issues by conducting cost simulations based on data from three of the campuses in our study. Our

basic approach is to start by looking, in as much detail as possible, at the actual costs of teaching a basic course in traditional format (usually, but not always, the statistics course) in a base year. Then, we simulate the prospective, steady-state costs of a hybrid version of the same course. These exploratory simulations are based on explicit assumptions, especially about staffing, which allow us to see how sensitive our results are to variations in key assumptions.

We did exploratory simulations for two types of traditional teaching models: 1) students taught in sections of roughly 40 students per section, and 2) students attending a common lecture and assigned to small discussion sections led by teaching assistants. We focus on instructor compensation because these costs comprise a substantial portion of the recurring cost of teaching and are the most straightforward to measure. We compare the current compensation costs of each of the two traditional teaching models to simulated costs of a hybrid model in which most instruction is delivered online, students attend weekly face-to-face sessions with part-time instructors, and the course is overseen by a tenure-track professor.

These simulations are admittedly speculative and subject to considerable variation depending on how a particular campus organizes its teaching, but they suggest that significant cost savings are possible. In particular, we estimate savings in compensation costs for the hybrid model ranging from 36 percent to 57 percent compared to the all-section traditional model, and 19 percent compared to the lecture-section model.

These simulations confirm that hybrid learning offers opportunities for significant savings, but that the degree of cost reduction depends (of course) on exactly how hybrid learning is implemented, especially the rate at which instructors are compensated and section size. A large share of cost savings is derived from shifting away from time spent by expensive professors toward both computer-guided instruction that saves on staffing costs overall and time spent by less-expensive staff in Q and A sessions.

Our simulations substantially underestimate the savings from moving toward a hybrid model in many settings because we do not account for space costs. It is difficult to put a dollar figure on space costs because capital costs are difficult to apportion accurately to specific courses, but the difference in face-to-face meeting time implies that the hybrid course requires 67 to 75 percent less classroom use than the traditional course.

In the short run, institutions cannot lay off tenured faculty or sell or demolish their buildings. In the long run, however, using hybrid models for some large introductory courses would allow institutions to expand enrollment without a commensurate increase in space costs, a major savings relative to what institutions would have to spend to serve the same number of students with a traditional model of instruction. In other words, the hybrid model need not just “save money”; it can also support an increase in access to higher education. It serves the access

goal both by making it more affordable for the institution to enroll more students and by accommodating more students because of greater scheduling flexibility. This flexibility may be especially important for students who have to balance family and work responsibilities with course completion, as well as for students who live far from campus.

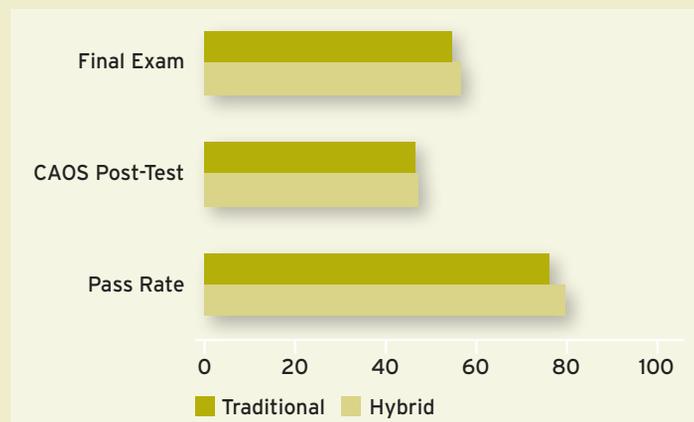
Conclusions

In the case of online learning, where millions of dollars are being invested by a wide variety of entities, we should perhaps expect that there will be inflated claims of spectacular successes. The findings in this study warn against too much hype. To the best of our knowledge, there is no compelling evidence that online learning systems available today—not even highly interactive systems, which are very few in number—can in fact deliver improved educational outcomes across the board, at scale, on campuses other than the one where the system was born, and on a sustainable basis.

This is not to deny, however, that these systems have great potential. Our study demonstrates the potential of

Student Success (Figure 2)

Results from the two course formats were essentially equivalent for all three learning outcomes.



SOURCE: Authors' calculations

truly interactive learning systems that use technology to provide some forms of instruction, in properly chosen courses, in appropriate settings. We find that such an approach need not affect learning outcomes negatively and conceivably could, in the future, improve them, as these systems become ever more sophisticated and user-friendly.

pressure than others to alter their teaching methods. ILO systems can be helpful not only in curbing cost increases (including the costs of building new space), but also in improving retention rates, educating students who are place-bound, and increasing the throughput of higher education in cost-effective ways.

It is hard to exaggerate

the importance of confronting the cost problems

facing American public education at all levels.

All of higher education has a stake in addressing this problem.

It is also entirely possible that by reducing instructor compensation costs for large introductory courses, such systems could lead to more, not less, opportunity for students to benefit from exposure to modes of instruction such as independent study with professors, if scarce faculty time can be beneficially redeployed.

What would be required to overcome the barriers to adoption of even simple online learning systems—let alone more sophisticated systems that are truly interactive? First, a system-wide approach will be needed for a sophisticated customizable platform to be developed, made widely available, maintained, and sustained in a cost-effective manner. It is unrealistic to expect individual institutions to make the up-front investments needed, and collaborative efforts among institutions are difficult to organize, especially when nimbleness is needed. In all likelihood, major foundation, government, or private-sector investments will be required to launch such a project.

Second, as ILO courses are developed in different fields, it will be important to test them rigorously to see how cost-effective they are in at least sustaining and possibly improving learning outcomes for various student populations in a variety of settings. Such rigorous testing should be carried out in large public university systems, which may be willing to pilot such courses. Hard evidence will be needed to persuade other institutions, and especially leading institutions, to try out such approaches.

Finally, it is hard to exaggerate the importance of confronting the cost problems facing American public education at all levels. The public is losing confidence in the ability of the higher-education sector in particular to control costs. All of higher education has a stake in addressing this problem, including the elite institutions that are under less immediate

We do not mean to suggest that ILO systems are a panacea for this country's deep-seated education problems. Many claims about "online learning" (especially about simpler variants in their present state of development) are likely to be exaggerated. But it is important not to go to the other extreme and accept equally unfounded assertions that adoption of online systems invariably leads to inferior learning outcomes and puts students at risk. We are persuaded that well-designed interactive systems in higher education have the potential to achieve at least equivalent educational outcomes while opening up the possibility of freeing up significant resources that could be redeployed more productively.

Extrapolating the results of our study to K–12 education is hardly straightforward. College students are expected to have a degree of self-motivation and self-discipline that younger students may not yet have achieved. But the variation among students within any given age cohort is probably much greater than the differences from one age group to the next. At the very least, one could expect that online learning for students planning to enter the higher-education system would be an appropriate experience, especially if colleges and universities continue to expand their online offerings. It is not too soon to seek ways to test experimentally the potential of online learning in secondary schools as well.

William G. Bowen is senior advisor to Ithaca S+R (the strategy and research arm of ITHAKA). Matthew M. Chingos is senior research consultant at Ithaca S+R and a fellow at the Brookings Institution's Brown Center on Education Policy. Kelly A. Lack is a research analyst at Ithaca S+R, where Thomas I. Nygren is a former business analyst.