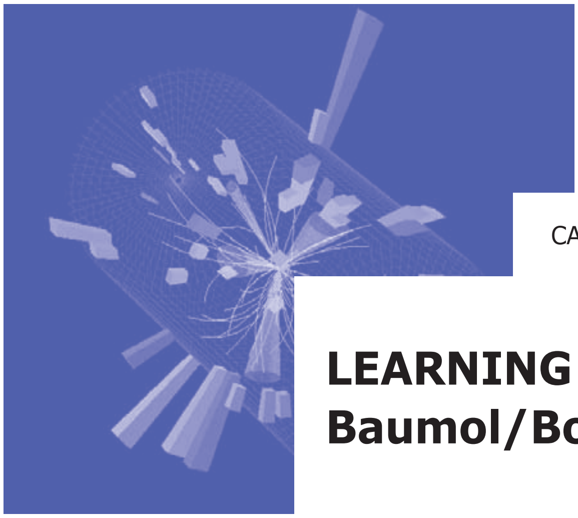

CANDACE THILLE AND JOEL SMITH, CARNEGIE MELLON UNIVERSITY

LEARNING UNBOUND: Disrupting the Baumol/Bowen Effect in Higher Education

In the United States, from 1982 to 2006, the cost of higher education increased 439%, far outstripping the consumer price index, which increased 106% over the same period.[1] Reports such as the one prepared for the UNESCO 2009 World Conference on Higher Education detail the fact that, despite some progress, the supply of tertiary education remains far from meeting the global demand.[2] Cost is clearly a key factor in providing greater access. Explanations of the high cost of higher education abound, ranging from well-meaning efforts to improve service to students and the professional lives of faculty[3] to poor management practices and new requirements for complying with government regulations[4] to the increased capital equipment costs associated with teaching increasingly complex topics using more expensive technology.[5] Of particular interest is the analysis of price pressures in all service industries first described by William Baumol and William Bowen in 1965[6] and again by Baumol in 1967, when he explicitly identified instruction as one service subject to seemingly uncontrollable upward price pressures.[7] Their explanation of these rising costs has come to be known as Baumol and Bowen's "cost disease." Candace Thille, Director of the Open Learning Initiative (OLI) at Carnegie Mellon University, and Joel Smith, Vice Provost and Chief Information Officer for Carnegie Mellon and Director of the Office of Technology for Education there, describe Carnegie Mellon's Open Learning Initiative, which presents promising possibilities for mitigating Baumol/Bowen's cost disease in higher education.

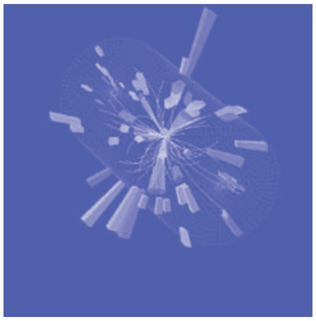

BAUMOL AND BOWEN'S COST DISEASE

The cost disease in service industries is produced by a number of interacting economic factors. The first factor comes from those sectors of the economy that produce goods rather than services. Goods-producing sectors have been able to significantly increase productivity over the years through technological innovations. Leveraging new technologies often requires a work force with more sophisticated skills, but the associated productivity increases are usually sufficient to increase salaries while reducing the costs of goods or, at least, holding cost increases within the rate of inflation. Increases in salaries for workers in the goods-producing sectors of the economy create pressures to increase salaries in the services sectors. If the services sectors don't keep up with rising salaries, they will lose

- Baumol and Bowen's cost disease assumes that improvements in productivity in service industries such as higher education are almost guaranteed to reduce the quality of the service.
- Yet it is widely acknowledged that the traditional labor-intensive process of having every instructor design his or her own course is incredibly inefficient; what may be less obvious is that the traditional course design and delivery process is often ineffective as well.
- Carnegie Mellon's Open Learning Initiative (OLI) courses are developed by teams composed of learning scientists, faculty content experts, human-computer interaction experts, and software engineers in order to make best use of multidisciplinary knowledge for designing effective instruction.
- OLI courses work. Results show that students who took OLI-Statistics, for example, learned a full semester's worth of material in half as much time and performed as well or better than students in traditional courses. Further, there were no significant differences in retention between OLI students and traditional students in follow-up tests given one or more semesters later.

qualified labor to the goods sectors. However, in the services sectors, there are far fewer opportunities for technological advances to lead to increases in productivity. Hence, prices for the same level of service increase as they have in higher education and, likewise, lower-income consumers are priced out of those markets.

Baumol and Bowen identified the role played by labor as the fundamental difference between goods and services. In the production of goods, human labor is a means to the end of producing a commodity that the consumer purchases. Consumers have no concern for how labor is deployed to produce a quality product as long as they are able to purchase what they want at a price they are willing to pay. If the productivity of the labor used to produce a good is increased, lowering the cost without a diminution in quality, all the better for the consumer. But in the service sector, labor is what Baumol describes as an "end in itself," and this makes all the difference. In his words:

We emphasize our belief that taking a scientifically based approach to course design and improvement has been and will continue to be key to creating instructional interventions that increase productivity without sacrificing quality.

> … there are a number of services in which the labor is an end in itself, in which quality is judged directly in terms of amount of labor. Teaching is a clear-cut example, where class size (number of teaching hours expended per student) is often taken as a critical index of quality. Here, despite the invention of teaching machines and the use of closed circuit television and a variety of other innovations, there still seem to be fairly firm limits to class size. We are deeply concerned when elementary school classes grow to 50 pupils and are disquieted by the idea of college lectures attended by 2,000 underclassmen. Without a complete revolution in our approach to teaching there is no prospect that we can ever go beyond these levels (or even up to them) with any degree of equanimity. An even more extreme example is one I have offered in another context: live performance. A half hour horn quintet calls for the expenditure of 2.5 man hours in its performance, and any attempt to increase productivity here is likely to be viewed with concern by critics and audience alike.[8]

In this passage, Baumol identifies the aspect of the cost disease analysis that most concerns us in this article: the assumption that improvements in productivity in service industries, like higher education, are almost guaranteed to reduce the quality of the service. Indeed, it can be argued that some of the comparatively small steps taken by higher education to increase productivity as a way to contain costs have reduced its quality. A widely used technique to increase instructional productivity has been to increase class sizes with concomitant attempts to mitigate the impact on quality through the use of recitation sections led by student teaching assistants and tutors. Introductory class sections at many universities are now lecture classes enrolling from 200 to more than 1,000 students. Anyone who has lectured to such a section knows that the number of students who have opportunities for meaningful interaction with the material in such a setting is extremely small, and that the quality of the interaction pales in comparison to the range of learning interventions that are possible in a freshman seminar with 15 to 20 students.

Baumol and Bowen seem to present higher education with an impossible dilemma with respect to increasing productivity by reducing the cost of instruction per student. On the one hand, increasing productivity in this sense will lower the quality of the education provided. On the other hand, if we do not increase productivity, the Baumol/Bowen's cost disease guarantees that we will price quality education out of the reach of more and more people in the world. The analysis leads to the conclusion that maintaining quality dooms us to reducing access due to increasing costs.

A DEFEASIBLE DILEMMA FOR HIGHER EDUCATION

We believe that Baumol and Bowen's dilemma is defeasible in higher education: that access to and the quality of higher education can be increased at the same time as we increase productivity. In his 1967 article, Baumol seemed pessimistic about technology making a significant difference in teaching. Unfortunately, the history of the use of educational technologies in higher education over the last three decades has justified much of his pessimism.

As work began on using computers as tutoring systems on mainframe computers in the 1960s, hopes were high that feedback-rich computer-based tutoring systems could both improve instruction and increase instructional productivity. There was no "complete revolution" as a result. But while progress in the use of computerized tutoring systems was much slower than many had hoped, there were successes. Notable among the successes were the intelligent tutoring systems, known as "Cognitive Tutors," developed at Carnegie Mellon University and based on John Anderson and his colleagues' theories of cognition.[9] But the impact of these successes was limited by a number of factors. The scholars engaged in studying artificial intelligence, cognitive science, and tutoring systems were, properly, focused on research results, and not on products that would use those results to improve learning outcomes in colleges and universities. Resistance from educators to fundamental change necessary to bring to bear results from the emerging disciplines of learning sciences meant few saw the creation of products based on scientific methodologies as an attractive market to develop.

The advent of the personal computer, the Internet, and the World Wide Web led many to focus on delivery of traditional materials through these new channels as the key to using technology to address the problem of access. Many colleges and universities rushed to provide an "online presence" with little consideration of how online materials would be used to create an effective learning experience or how they would actually meet the skyrocketing world demand for quality education.[10] Even today, many interested in finding ways that information technologies can be used to mitigate Baumol/Bowen's cost disease seem to look for scaling solutions without seeking guidance from the learning sciences community about "what works" based on what we now know about how humans learn.

One current candidate for cost-effective scaling of instruction is recording every lecture and making them available as an educational resource to both matriculated students and the world at large. Surely providing 7x24 web access to lectures is a possible path for lowering the cost per student because more students can be provided the same service of listening to a lecture at only the incremental cost of recording and webcasting or podcasting the lecture. The problem is that technology is being used to scale the service of lecturing but that is not, ultimately, the service that needs to be scaled. Rather than the performance of the lecture, the service that needs to be scaled is the collection of activities that effectively change the knowledge state of the learner. All of our understanding about human learning throughout the last 20 years of learning science research tells us that learning is an active, not passive process. As Grant Wiggins said: "It's not teaching that causes learning. Attempts by the learner to perform cause learning, dependent upon the quality of feedback and opportunities to use it." [11]

In reality, most colleges and universities are using class recordings under the assumption it will be a powerful additional tool for reviewing lectures. Whether students will use this tool in ways to improve their learning is, not surprisingly, dependent on whether they use it in ways that learning scientists have shown improve retention and understanding, namely as a source to enrich their interaction with the material through note taking.[12]

OPEN LEARNING INITIATIVE: IMPROVING PRODUCTIVITY *AND* QUALITY

Carnegie Mellon University has been employing a different approach to produce technologies (course development, evaluation, and improvement methodologies and specific web-based learning interventions) that *simultaneously* improve productivity and the quality of instruction. Using intelligent tutoring systems, virtual laboratories, simulations, and frequent opportunities for assessment and feedback, the university's Open Learning Initiative (OLI) builds learning environments that enact the kind of dynamic, flexible, and responsive instruction that fosters learning. Moreover, all student learning activities in OLI courses and labs are, with the student's permission, digitally recorded in considerable detail to monitor student activity and capture data that informs further course revisions and improvements.

In the remainder of this article, we describe how OLI uses the web to deliver online instruction that instantiates course designs based on research from the learning sciences, and how data from student use of these courses contributes back to the underlying design principle or learning theory and facilitates a system of continuous improvement for the student, the instructor, the course designers, and for learning science. Throughout, we emphasize our belief that taking a scientifically based approach to course design and improvement has been and will continue to be key to creating instructional interventions that increase productivity without sacrificing quality.[13]

OLI is an open educational resources project that began in 2002 with a grant from The William and Flora Hewlett Foundation. Like many open educational resources projects, OLI makes its courses openly and freely available. However, OLI courses are much more than collections of material created by individual faculty to support traditional instruction. While OLI courses are often used by instructors to support classroom instruction, the goal of OLI has been to create courses that *enact instruction*, that is, they offer structure, information, activities, practice, and feedback—all arranged so that students can learn even if they do not have the benefit of an instructor or classmates.

We know that the traditional labor-intensive process of having every instructor design his or her own course is incredibly inefficient; what may be less obvious is that the traditional course design and delivery process is often ineffective.

Translating scientific results from the learning sciences into effective instruction requires significant design and assessment efforts. Such an effort by one faculty member for a single class is rare and, even when done, typically has an impact on comparatively few students. Such efforts made by a team for online virtual learning environments with the specific goal of creating virtual learning environments that can be used in standard introductory courses in various disciplines produces effective materials that can be used by many faculty and learners. Each OLI course is developed by a team composed of learning scientists, faculty content experts, human-computer interaction experts and software engineers in order to make best use of multidisciplinary knowledge for designing effective instruction.

OLI course development begins with a study of the teaching and learning challenges in the domain under development. The study includes literature reviews, reviews of existing artifacts of student learning, classroom observations, lab studies and/or classroom-based studies. The design team articulates a set of student-centered measurable learning objectives, then designs the instructional environment to support students to achieve the articulated objectives. The instructional activities include small amounts of explanatory text and many activities that capitalize on the computer's capability to display digital images and simulations and promote interaction. Many of the courses include virtual lab environments that encourage flexible and authentic exploration and problem solving. OLI courses are guided by principles from current cognitive and learning theory, and each course attempts to genuinely reflect the core epistemic structure of the domain it represents. A hallmark of all OLI courses is the frequent opportunities for students to assess their own learning and receive context-specific and targeted feedback on their work.

OLI benefits from inheriting some of the best work done in the area of online tutoring by Carnegie Mellon and University of Pittsburgh faculty. Many OLI courses feature Cognitive Tutors and "mini-tutors" that give students feedback in the problem-solving context. A Cognitive Tutor is a computerized learning environment whose design is based on cognitive principles and whose interaction with students is based on that of a human tutor, that is, making comments when the student errs, answering questions about what to do next, and maintaining a low profile when the student is performing well. This approach differs from traditional computer aided instruction in that traditional instruction gives didactic feedback to students on their final answers, whereas the Cognitive Tutors and "mini-tutors" provide context-specific assistance during the problem-solving process.

FEEDBACK MECHANISMS IN OLI COURSES

In a mini-tutor from the OLI engineering statics course, the student is presented with a graphical representation of three force vectors and asked to give the direction and magnitude of the sum of the forces. If the student is unsure of the procedure for summing concurrent forces, the first hint provides a link which, when clicked, expands the tutor into the first set of steps needed to solve the problem. The tutor provides scaffolding hints and feedback to support the student to learn the steps of the procedure as needed. As the student works through each step of the problem, the hints and feedback given by the tutor change depending on which part of the exercise the student is attempting. There are multiple levels of hints for each sub-step: the student may continue to ask for hints by clicking the 'get next hint' link at the bottom of the hint window until the final hint that gives the answer for that sub-step and allows the student to continue working on the larger problem. The system tracks the student's use of the scaffolding hints and feedback. When the student successfully completes the task after having used the scaffolding, the system suggests the student try another problem. A new problem statement, graphical representation, hints, feedback and answers are dynamically generated by the system. The student can work through the task multiple times, receiving a different problem each time, until he or she is confident about understanding the concept and has developed fluency with the procedure. The student is given virtually unlimited opportunities for supported practice.

One of the most powerful features of OLI learning environments is that they embed ongoing formative assessment and feedback not only in tutored problem solving but into virtually every instructional activity. In many domains digital learning environments are used to make features, processes and causal relationships visible to students in ways that are not possible in text books or in the traditional classroom. However, often when students are provided with such simulations created by the instructors or by the textbook publishers, they do not engage sufficiently with the materials or the simulation does not provide sufficient support for the student to learn the target concept or skill from the engagement. In such activities, there is a risk that students will develop or strengthen incorrect or surface knowledge, and/or create or strengthen inappropriate connections among concepts. Students need sufficient support so that they do not discover, practice and encode incorrect or surface knowledge. In OLI courses, learners experiment with the parameters and see the effects of their experimentation in interactive guided simulations into which assessments are embedded. Students are asked questions and prompted to reflect on their observations as they explore animations and simulations. They receive immediate feedback in two forms: directed observation of the causal interactions depicted by the simulations, and explanatory text in response to their answers.

OLI Feedback to students includes corrections, suggestions and cues that are tailored to the individual's current performance, and encourages the student to revise and refine their performance. Many learning studies have shown that students' learning improves and their understanding deepens when they are given timely and targeted feedback on their work.[14] Regarding the timing and frequency of feedback, the best learning outcomes occur when feedback comes immediately after the students' response but not before the student is ready to revise his or her understanding.[15]

Evaluators internal and external to the project have conducted numerous studies of the effectiveness of OLI environments in supporting student learning. Several studies collected empirical evidence of the instructional effectiveness of the OLI courses in stand-alone mode, as compared to traditional instruction. In all cases, in-class exam scores showed no significant difference between students in the stand-alone OLI course and students in the traditional instructor-led course. In a study of the OLI logic course at a large state university[16] the OLI online-only students covered more material than the traditional students and there was no significant difference in performance between the two groups on exams. The most interesting result from this study is that attrition in the online course was almost non-existent, whereas the in-class condition had very high attrition (although at typical levels for this large, rigorous course at this institution). In the OLI online-only condition 84 students started the course and 83 students successfully competed it, showing a 99% retention rate. In the traditional face-to-face condition, 259 students started the course and only 105 students successfully completed it, showing a 41% successful completion rate. This result is important for two reasons: First, this differential attrition rate likely produces very significant performance biases in favor of the in-class condition on the final exam and yet the in-class condition did not perform better on the final exam. Second, it shows that another potential benefit of using OLI in the online-only mode is that it can significantly reduce attrition rates.

The OLI environments are used not only to provide virtually unlimited supported practice to develop procedural fluency in problem solving but also to support students to build a conceptual structure of knowledge that is authentic to the domain.

COHERENCE AS AN ELEMENT OF OLI COURSE DESIGN

The OLI environments are used not only to provide virtually unlimited supported practice to develop procedural fluency in problem solving but also to support students to build a conceptual structure of knowledge that is authentic to the domain. The conceptual structure of knowledge in a given domain is usually obvious to experts in that domain but not to novices. Introductory courses tend to overwhelm students with what seems to be a set of isolated facts, lacking in connective structure.[17] [18] OLI courses seek to promote coherence by teaching students how the discrete skills they are learning fit together into a meaningful conceptual big picture.

Much of college level chemistry is often taught out of context as a set of abstract mathematical skills. Students employ learning strategies to solve typical text book problems and perform well on traditional chemistry exams but often fail to see either the relationship between the mathematical procedures and the chemical phenomena those procedures represent, or the relationship between the chemical phenomena and the real world.

The OLI chemistry course is designed to address both of these educational challenges. It addresses the challenge of connecting the mathematical procedures used in chemistry to chemical phenomena by replacing traditional textbook problems with problems to be constructed and solved in the virtual chemistry lab. In the virtual chemistry lab, students are presented with unstructured problems that require flexible application of procedural knowledge. The course addresses the second challenge of connecting the procedures of chemistry to the real world by employing scenario-based learning. It situates the learning of chemistry in an authentic investigation that addresses questions that are significant to the domain of chemistry and to real world problems. The OLI Chemistry course unit on stoichiometry, for example, is situated in a real world problem of arsenic contamination of the water supply in Bangladesh. The course opens with a video that gives the student an overview of the real world problem, the distribution of arsenic in Bangladesh, and the health effects of arsenic poisoning, the difficulty of discouraging people from drinking the contaminated well water because it is clear and the arsenic is odorless, colorless and tasteless, and the need for inexpensive and easy ways to test the level of contamination in each well and remove a sufficient amount of arsenic from the water to make it safe to drink.

The system collects data as the students work, and automatically analyzes and organizes the data to present the instructor with the students' current "knowledge state." The instructor reviews the information and adapts instruction accordingly.

Following the video, the student engages in a process designed to solve the problem as a chemist would. At each step of exploring a solution to the arsenic contamination problem, the student is introduced to and practices one of the target stoichiometric concepts or skills. For example, in the very first step, determining the level of arsenic contamination in a sample of well water, the student uses the Chemistry Virtual Lab to analyze a well water sample and compare the level of arsenic found in the sample to the acceptable levels set by the World Health Organization (WHO). The challenge the student confronts is that virtual lab experiment gives them the concentration of arsenic in units of moles/liter. The WHO gives its safety standard as 10 micrograms of arsenic per liter. The student must be able to convert the results from the lab to evaluate the concentration of elemental arsenic in units of micrograms per liter. In order to evaluate the safety of the water, the student must understand the concept of the mole and apply dimensional analysis, composition stoichiometry and solution stoichiometry. If the student does not already understand these concepts or can't demonstrate mastery of these procedures in the context of solving the problem, he or she is directed into an instructional sequence that includes demonstrations, worked examples, and mini-tutors.

An analysis of the data logs of student use from a study conducted on the OLI stoichiometry course revealed that the number of engaged actions with the Virtual Lab is strongly correlated with success: The degree of interaction explains about 48% of the variation observed in the post-test scores for students taking OLI. The number of interactions with the Virtual Lab outweighed all other factors, including gender and SAT score, as the predictor of positive learning outcome.[19]

ASSESSMENT OF OLI COURSE EFFECTIVENESS

The initial motivation of the OLI was to develop exemplars of high quality, online courses that support individual learners in achieving the same learning goals as students enrolled in similar courses at Carnegie Mellon University. Although originally designed to support individual learners, OLI courses are increasingly used by instructors inside and outside of Carnegie Mellon to complement their instructor-led courses and address the challenges they confront due to the increasing variability in their students' background knowledge, relevant skills, and future goals.

Creating an effective feedback loop to instructors using the OLI courses is an ongoing area of investigation. The process of using an OLI course to support the classroom goes something like this: The instructor assigns students to work through a segment of the OLI course. The system collects data as the students work, and automatically analyzes and organizes the data to present the instructor with the students' current "knowledge state." The instructor reviews the information and adapts instruction accordingly. The richness of the data collected about student use and learning provides an unprecedented opportunity for keeping instructors in tune with the many aspects of students' learning.

ROBUST LEARNING OUTCOMES, HALF THE CLASS TIME: EXPERIMENTS WITH THE OLI STATISTICS COURSE

OLI evaluation efforts have investigated OLI courses' effectiveness not only in stand-alone mode, but also in an instructor-led "accelerated learning" mode. This type of study owes its origins to Ben Bloom's mastery learning concept and the subsequent accelerated schools program. The most common dependent measure used in such studies is time, i.e. the time it takes a learner to complete a particular amount of material, with proper assessment of equivalent learning outcomes. In these studies of OLI courses, we have demonstrated accelerated learning by showing that a learner can complete a semester-long course in significantly less than a semester and/or that a learner can complete significantly more than a semester's worth of material within a semester's time.

Results showed that OLI-Statistics students learned a full semester's worth of material in half as much time and performed as well or better than students in traditional instruction.

Two studies conducted at Carnegie Mellon tested whether learners using the OLI course in hybrid mode—that is, students meeting with instructors regularly, but less frequently than in traditional courses, while also using the online modules and assignments of OLI- Statistics—would learn the same amount of material in a significantly shorter time than students in traditional class formats. Results exceeded expectations: OLI-Statistics students completed the course in 8 weeks with 2 class meeting per week, while traditional students completed the course in 15 weeks with 4 class meetings per week. Significantly, student logs showed that the OLI students spent no more total time studying statistics outside of class than the traditional students. Yet the OLI students demonstrated as good or better learning outcomes than the traditional students. Further, there was no significant difference in retention between OLI students and traditional students in tests given 1+ semesters later.[20] Usually, that kind of effectiveness or efficiency effect would be expected only as the result of individualized, human-tutored instruction. And yet in this case, students who met for less than two hours of class per week demonstrated phenomenal performance.

Why does this hybrid approach work so well? The accelerated OLI-Statistics students used their out-of-class time much more effectively and attended their class meetings much better prepared than traditional students. As opposed to skimming (or skipping) the reading before a traditional lecture or floundering through homework, these accelerated students prepared for class by actively engaging with the material and receiving targeted and timely feedback as they needed it. As a result, students came to class ready to make best use of their time with the instructor.

Equally significant, the instructor came to class better prepared to teach. Thanks to OLI's automatically generated instructor reports, the instructor was able to access reports on student progress and performance. With this information in hand, the selected discussion topics and example problems targeted the topics with which the students were struggling. Class time was spent with students actively engaged in working on the material that was most likely to need more supported practice or a novel explanation from the instructor. The faculty member who led the course reported that it was one of the most satisfying experiences he had had in his 15 years of teaching.

It is this combination of focused preparation by both the students and the instructor—informed by real-time feedback loops in well-designed learning environments—that underpins the success of the accelerated hybrid mode. The potential for addressing the Baumol/Bowen's cost disease should also be clear: under the condition tested, a professor can teach what is now a year-long sequence of introductory statistics in one semester, with roughly one-half the amount of time commitment to class meetings and other forms of in-person student support.

The student log data is also used to evaluate and iteratively refine parts of the OLI course. For example, by examining the data from students working through the Causal and Statistical Reasoning course, the development team observed that students were engaging in all of the learning activities but still failed on a target skill of "building causal response structures." The team constructed six additional learning activities and mini-tutors designed to support students' understanding and increase their practice of this target skill. The following semester the team analyzed the student data again to confirm that students were using the new activities, and that those activities resulted in the students learning the target skill.

Some OLI courses also serve as part of the research environment for the Pittsburgh Science of Learning Center (PSLC). In 2004, the National Science Foundation funded the PSLC, one of the two NSF-funded national centers that study the nature of human learning. The PSLC dramatically increases the ease and speed with which learning researchers can create the rigorous, theory-based experiments that pave the way for an understanding of robust learning. The PSLC makes use of advanced technologies to facilitate the design of experiments that combine the realism of classroom field studies and the rigor of controlled theory-based laboratory studies. Learning researchers affiliated with the PSLC embed experimental manipulations in OLI courses to test specific learning theories. The researchers then analyze the data collected by the OLI logging service using the PSLC "Datashop" tools. The PSLC Datashop has created a number of tools specifically designed to generate meaningful displays of student learning data. OLI Learning environments build on what is known about learning and also serve as a platform in which new knowledge about human learning can be developed and further refined.

CONCLUSION

OLI is much more than a technology. It is a set of strategies for course design, development, delivery and evaluation. OLI development teams use learning science research results to inform course design and use learning science research methods both to unpack the cognitive tasks associated with learning and to design appropriate instructional interventions. OLI courses instantiate the research and hold the collective memory of what works and what doesn't work, so that time and resources are well-spent on refinement rather than wasted on reinventing existing successes or failures. OLI collects data to provide feedback loops to students, instructors, and course design teams, and supports the implementation of learning science for continuous evidence-based improvement.

Baumol said: "Without a complete revolution in our approach to teaching there is no prospect that we can ever go beyond (current) levels [of productivity] (or even up to them) with any degree of equanimity."[21]

We believe that, using OLI strategies to develop, assess, deliver, and iteratively improve courses, we have in place the key elements to the revolution Baumol sought. This revolution will allow us to disrupt the Baumol/Bowen effect and make higher education less expensive and more accessible without sacrificing quality.

ENDNOTES

1 *Measuring Up, The National Report Card on Higher Education*, 2008, p. 8. http://measuringup2008.highereducation.org/

2 *Trends in Global Higher Education, Tracking and Academic Revolution*, 2009. www.unesco.org/en/wche2009/resources/global-reports/

3 See the discussion on the "administrative lattice" and "academic ratchet" in, Robert Zemsky, Gregory R. Wenger, and William F. Massy, *Remaking the American University: Market-Smart and Mission-Centered* (Rutgers University Press, 2006)

4 Malcolm Getz and John J. Siegfried, "Cost and Productivity in American Colleges and Universities," in Charles Clotfelter, Ronald Ehrenberg, Malcolm Getz, and John J. Siegfried (eds.), *Economic Challenges in Higher Education* (University of Chicago Press, 1991), pp. 261-392.

5 Robert Archibald and David Feldman, "Why Do Higher Education Costs Rise More Rapidly than Prices in General?," *Change Magazine*, May/June 2008, pp. 25-31.

6 William J. Baumol and William G. Bowen, "On the Performing Arts: The Anatomy of their Economic Problems." *The American Economic Review*, (1965), pp. 495-502.

7 William Baumol ,"Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis," *The American Economic Review*, (1967), pp. 415-426.

8 Ibid. , p. 416.

9 John Anderson, Albert Corbett, Ken Koedinger, and Ray Pelletier, "Cognitive Tutors, Lessons Learned," *The Journal of the Learning Sciences*, (1995), pp. 167-207.

10 Zemsky, Robert and Massy, William. Thwarted Innovation, What Happened to e-Learning and Why, The Learning Alliance, 2004.

11 Wiggins, G., *Educative Assessment: Designing Assessments to Inform and Improve Student Performance.* (Jossey-Bass: San Francisco, 1998).

12 McKinney, Dani and Dyck, Jennifer and Luber Elise, "iTunes University and the Classroom: Can Podcasts Replace Professors?," *Computers and Education*, (2009) pp. 617-623.

13 We emphasize that our philosophy in this regard both owes a debt to and is consistent with efforts made over the last decade by Carol Twigg and the National Center for Academic Transformation she leads. It is equally influenced by the long history of scientifically based approaches to improving teaching started by Herbert Simon and continued by Susan Ambrose and the Eberly Center for Teaching Excellence at Carnegie Mellon.

14 D. L. Butler and P. H. Winne, "Feedback and self-regulated learning: A theoretical synthesis." *Review of Educational Research*, (1995), pp. 245-281

15 A. T. Corbett and J.R. Anderson, "Locus of feedback control in computer-based tutoring: Impact on learning rate, achievement and attitudes," *Proceedings of ACM* (2001), pp. 245-252.

16 C. D. Schunn and M. Patchan, "An evaluation of accelerated learning in the CMU Open Learning Initiative course 'Logic & Proofs,'" *Technical Report by Learning Research and Development Center*, University of Pittsburgh, (2009).

17 M. T. H. Chi, "Common Sense Conceptions of Emergent Processes: Why some misconceptions are robust," *Journal of the Learning Sciences* (2005) , 161-199.

18 A.A. diSessa, "Coherence versus fragmentation in the development of the concept of force," *Cognitive Science* (2004), 843-900.

19 K. Evans, D. Yaron, G. Leinhardt, "Learning stoichiometry: a comparison of text and multimedia formats." *Chemistry Education Research and Practice* (2008) pp. 208 – 218.

20 M. Lovett, O. Meyer, & C. Thille, C., "The Open Learning Initiative: Measuring the effectiveness of the OLI statistics course in accelerating student learning," *Journal of Interactive Media in Education* (2008), http:// jime.open.ac.uk/2008/14/

21 Baumol, "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis," (op. cit.)

CANDACE THILLE *is director of the Open Learning Initiative (OLI) at Carnegie Mellon University, a position she has held since the program's inception in 2002. She is also co-director of OLnet, an open educational research network, a collaboration between Carnegie Mellon and the Open University, UK. She has published and presented more than eighty conference proceedings, workshops, articles and book chapters on open educational resources and effective web-based learning environments. Thille serves as a redesign scholar for the National Center for Academic Transformation. She can be reached at cthille@cmu.edu.*

JOEL SMITH *is vice provost and chief information officer for Carnegie Mellon, Computing Services and also directs the Office of Technology for Education. He has been involved in technology enhanced learning projects for two decades. Smith is currently the principle investigator in charge of the Open Learning Initiative, a project at Carnegie Mellon funded by the William and Flora Hewlett Foundation. He serves on advisory boards for Intel Corporation, Apple Computer,* Campus Technologies *magazine, and Harvard University. Smith can be reached at joelms@andrew.cmu.edu.*