



Lessons Learned From Early Implementations of Adaptive Courseware

April 2016

SRI Education[™]

A DIVISION OF SRI INTERNATIONAL

Contract #24133, Work Order #27; Contract #28981, Work Order # 29
SRI Project Nos. 21987 and 22997

Prepared for:

Yvonne Belanger
Senior Program Officer
Measurement, Learning and Evaluation
Education, Postsecondary Success, U.S. Programs
Bill & Melinda Gates Foundation
P.O. Box 23350
Seattle, WA 98102
yvonne.belanger@gatesfoundation.org

Submitted by:

Louise Yarnall, Barbara Means, and Tallie Wetzel
SRI Education
333 Ravenswood Avenue
Menlo Park, CA 94025
Phone: 650-859-3181

Executive Summary

To address the urgent need to improve student outcomes in developmental and general education courses, higher education institutions are turning to new learning technologies. Prominent among these is adaptive learning courseware that uses computer algorithms to parse learning analytic data collected as students interact with online learning environments. These adaptive algorithms then determine the student's next learning activity and can be used to generate personalized feedback, study reminders, content recommendations, and real-time progress dashboards that both students and instructors may review.

The Bill & Melinda Gates Foundation initiated the Adaptive Learning Market Acceleration Program (ALMAP) to advance evidence-based understanding of how adaptive learning technologies such as adaptive courseware could improve opportunities for low-income adults to learn and to complete postsecondary credentials. Over three academic terms, from summer 2013 through winter 2015, the ALMAP grant program provided 14 higher education institutions with seed funding to incorporate nine adaptive learning products into 23 courses and to conduct quasi-experiments (side-by-side comparisons with comparison groups) to measure their effects on student outcomes, and to gather data on cost impacts and instructor and student satisfaction.

The foundation asked grantees to use the adaptive technologies to address two common hurdles to college completion: success and mastery in gateway general education courses and developmental (remedial) courses. These courses were chosen as targets for improvement because student success in gateway general education courses and developmental courses paves the way to persistence and success in the first two years of college, a time when many students fall off track for completing

an associate's or bachelor's degree. The foundation asked grantees to evaluate the impact of their adaptive courseware implementations on student learning, course grades, and probability of course completion. Additionally, in view of trends showing college costs outpacing general inflation since 1979,¹ the foundation directed grantees to explore both the costs and potential savings associated with adaptive courseware implementation.

To strengthen the consistency and credibility of the evidence grantees gathered, the foundation also contracted with SRI International to aggregate and analyze the ALMAP data. SRI's evaluation looked at the learning and cost impacts of adaptive courseware implementations both separately and together. We did not expect to find one answer for the many forms of adaptive instruction and diverse institutions of higher education in the United States, but the ALMAP study attempts to provide information that will help local decision makers identify approaches that might work for them.

SRI assembled learning impact, cost, and satisfaction findings across the portfolio of ALMAP grantee product evaluations. This synthesis of product evaluations encompassed data collected from over 19,500 unique students in classes taught by more than 280 unique instructors. The resulting ALMAP evaluation report provides a glimpse into the state of the art of adaptive learning courseware implementations across the range of U.S. institutions of higher education—from research universities to colleges focused on undergraduate education, and from public community colleges to private online colleges.

¹Bureau of Labor Statistics. (2010, September). *Back to college: Spotlight on statistics*. Washington, DC: Bureau of Labor Statistics. <http://www.bls.gov/spotlight/2010/college/>

Portfolio Description

The ALMAP grantees implementing one or more of the nine adaptive learning products were 10 bachelor's degree programs and 4 associate's degree programs. ALMAP grantees used adaptive courseware in 15 gateway general education courses and 7 developmental education courses. The gateway course subjects in which adaptive courseware was used included psychology, biology, business, marketing, and economics. The developmental education courses focused on the mathematics or English language arts proficiencies needed to succeed in college-level work.

ALMAP grantees used the adaptive courseware to make different kinds of changes in course delivery. SRI organized these changes into three use case categories:

- Blended Adaptive vs. Lecture— adaptive courseware was used as part of a shift from traditional lecture to blended instruction;
- Online Adaptive vs. Online— adaptive courseware was used as an enhancement to existing online courses; and
- Blended Adaptive vs. Blended—adaptive courseware was swapped into face-to-face courses already using blends of classroom-based and online approaches to support learning.

Key Findings

In reviewing the learning, cost, and satisfaction outcomes for the ALMAP portfolio, it is important to keep in mind the relative immaturity of the field of adaptive learning technology. Both technology capacity and ways to support instruction and learning with technology are evolving rapidly, and these results should be viewed as snapshots in time. Nevertheless, the impact estimates included in SRI's synthesis all passed screens for adherence to generally accepted research standards, and the resulting data set is one of the largest of its kind for commercially available adaptive courseware products. Major findings were as follows:

EFFECTS ON STUDENT LEARNING AND COURSE COMPLETION

- Some adaptive courseware implementations (4 of the 15 with a data set adequate for analysis) resulted in slightly higher average **course grades**, but the majority had no discernible impact on grades.
- Overall, in the 16 grantee-provided data sets appropriate for estimating courseware impacts on **course completion**, the odds of successfully completing a course were not affected by the use of adaptive courseware.
- Only seven controlled side-by-side comparisons of scores on common **learning assessments** were available; the average impact of adaptability for these seven was modest but significantly positive.
- The impacts of adaptive courseware varied by use case:
 - Switching from a lecture format to adaptive blended instruction had a positive impact on student learning as measured by posttests.
 - Moving from nonadaptive to adaptive learning systems in fully online courses had a small positive effect on course grades.
 - There were too few cases contrasting blended adaptive versus blended nonadaptive courses to draw any conclusions about impacts.
 - None of the use case analyses found a significant average impact on course completion rates; only 2 of the 16 side-by-side comparisons of completion rates found a significantly positive impact on the odds of course completion.

- Courseware products with adaptivity at a micro level (individual lesson or learning object) produced stronger student outcomes than those with adaptivity at a more macro level.
- The size of the adaptive courseware effect did not vary significantly for different academic disciplines, but impacts appeared to be larger for mathematics and biology courses, a trend that might have attained statistical significance with a larger sample.
- Impacts for Pell Grant students were equivalent to those for students overall. The ALMAP study provided a limited amount of data on the impact of adaptive courseware on outcomes for Pell Grant (low-income) students, but the data that were available provided no indication that adaptive courseware was either more nor less advantageous for Pell Grant students. This finding contrasts with some previous studies that have suggested that online and blended learning approaches put low-income students at a disadvantage.

Effects on Course Costs

- Comparisons of per-student costs for both adaptive courseware and comparison versions of a course found that in most cases costs went up during the first term of adaptive courseware implementation. Course costs are driven largely by instructor labor, and a number of the adaptive products were platforms into which instructors inserted content.
- The adaptive courseware was associated with lower ongoing costs in 7 of the 10 cases with cost data for second and third implementations of adaptive courseware.
- There were only eight cases for which we had both learning impact estimates and comparative cost data. Five of these eight cases had reduced costs but only one of those five produced improved learning outcomes at the same time. In the other four cases of cost reduction, there was no significant change in learning outcomes.

INSTRUCTOR AND STUDENT PERCEPTIONS OF ADAPTIVE LEARNING EXPERIENCES

- Among ALMAP adaptive courseware instructors who responded to SRI's survey, 74% reported satisfaction with the adaptive courseware they used.
- More developmental course instructors (67%) than gateway general education course instructors (49%) planned to use adaptive courseware in the future.
- In short written responses, instructors endorsed the adaptive courseware's real-time progress dashboards as useful for informing their teaching.
- The major concern expressed by instructors was getting students to use the adaptive courseware frequently enough.
- SRI's analysis of responses to student surveys administered by ALMAP grantees indicated that 2-year college students had more positive views of the adaptive courseware than did students at 4-year colleges and universities.
 - Around half (51%) of 4-year college students, compared to over three-quarters (77%) of 2-year college students, reported that they had made positive learning gains with the adaptive courseware.
 - In addition, 56% of 2-year college students reported satisfaction with their adaptive courseware experience compared with only 33% of bachelor's degree students.
- Developmental course students reported higher rates of engagement (60%) and learning gains (95%) with the courseware than did gateway course students (25% and 35%, respectively).

Implications for Future Work

Future research into blended learning technology implementation and efficacy in higher education is badly needed in a field awash in marketing claims. The ALMAP evaluation showed that adaptive courseware is not a homogeneous category that can be established as “effective” or “ineffective.” The diversity of products, use cases, and outcomes for the ALMAP implementations yielded important lessons for postsecondary institutions and faculty, courseware suppliers, researchers, and funders.

FOR POSTSECONDARY INSTITUTIONS

- Postsecondary institutions planning large-scale adoptions of adaptive courseware should conduct their own internal analyses of student outcomes with that courseware compared to other alternatives.** Even when the identical adaptive courseware product is used in different institutions or academic terms, learning outcomes can differ markedly depending on how it is used. To provide a valid answer to questions about the relative effectiveness of the adaptive courseware, these analyses need to establish the baseline equivalence of students in the courseware and comparison course sections being compared.
- Baseline equivalence is essential for justifying claims about courseware effects,** but the common practice in higher education institutions is to simply compare course success rates without any data on student characteristics (baseline equivalence). ALMAP analyses found that student characteristics and prior learning often vary markedly from course section to section and across terms within the same institution.
- Adaptive courseware is unlikely to reduce per-student course costs in its initial implementation.** Even when a college does not need to make infrastructure enhancements to accommodate the courseware, there are costs entailed for instructors learning how to use the courseware and in many cases, for instructor insertion or modification of content. First-term cost reductions were

observed in only a minority of cases, most of which involved vendor-developed rather than instructor-generated content.

- Subsequent implementations of adaptive courseware have stronger prospects for cost reductions.** If the same courseware product is implemented for multiple terms, costs often drop considerably from those incurred in the initial term. Institutional decision making around large-scale courseware adoption, should incorporate multi-year projections of both costs and impacts (such as reduced need for repeating developmental courses or reductions in attrition).

FOR POSTSECONDARY INSTRUCTORS

- Instructors in 2-year colleges and those teaching developmental education courses would do well to consider adaptive courseware options.** The ALMAP evaluation found only minor enhancement of course grades associated with adaptive courseware, but results for the relatively few cases with direct measures of student learning (posttests), were encouraging. Moreover, instructors and students in 2-year colleges and developmental courses reported high levels of satisfaction with their adaptive courseware experiences.
- Adoptions of adaptive courseware in 4-year institutions should include planning to make sure the courseware’s benefits are apparent to students.** Only a third of the students in 4-year colleges responding to surveys about their experiences with adaptive courseware expressed overall satisfaction. It is not clear from the available data whether qualities of the courseware they were using, the way in which the courseware was implemented, or a sense of loss of instructor personal attention was most responsible for their lack of enthusiasm. Instructors of 4-year college gateway courses are encouraged to attend to all of these issues when transitioning to use of adaptive courseware.

- **Instructors can make valuable contributions to student success by sharing their insights about adaptive courseware** with each other and with the field more broadly, including with vendors. Course instructors are in a good position to understand more about sources of college students' satisfaction and dissatisfaction with adaptive courseware as well as ways in which it could better support learning. It is quite possible that courseware designers or instructors can improve courseware effectiveness and quality by changing courseware settings or implementing courseware somewhat differently.

FOR SUPPLIERS OF ADAPTIVE COURSEWARE

- **Courseware providers can leverage partnerships with their higher education clients to obtain better data for use in product improvement.** Courseware developers already analyze the use data available from their learning systems, but typically do not have other important kinds of data regarding how their courseware is being used, student characteristics, and student outcomes. Courseware vendors should seek opportunities to partner more deeply with their field sites to better understand the aspects of course implementation that influence the learning and cost impacts of their products. Data-driven insights can inform continuous improvement cycles so that courseware quality and user experience are enhanced over time.
- **Courseware providers can work with their institutional partners and researchers to articulate and validate implementation guidelines.** As effectiveness and user satisfaction data are accumulated across multiple implementations in different institutional contexts, those data can be analyzed to derive empirically based recommendations for how the courseware should be used, including the total amount of time and spacing of use.

- **Courseware software developers can make sure that it is easy to pull user data revealing the key interactions students have with the courseware.**

Courseware log file data can reveal areas where students appear to get stuck, whether all of the major components of the courseware are being used, and indications that students are trying to bypass learning opportunities in order to move through the courseware as quickly as possible. Student “click stream” data need to be aggregated to a level that is interpretable by instructional designers, researchers, and faculty. It should be possible to query and export data for all users within a given class within a given period of time.

FOR RESEARCHERS AND EVALUATORS

- **Analyses of adaptive courseware effectiveness should take into account the specifics of the way in which the courseware was used in a particular implementation.** Researchers should help courseware users understand that learning efficacy is not a trait of a product per se or simply a matter of matching the right product to the right subject matter. Rather, multiple factors affect learning outcomes and to make sense of student outcomes, analyses need to incorporate student characteristics, specifics of how the adaptive courseware is used, aspects of the course beyond the courseware product, and the way learning is measured to make sense of student outcomes.²

² Means, B., Bakia, M., & Murphy, R. (2014). *Learning online: What research tells us about whether, when and how*. New York: Routledge.

- **Course grades and course completion rates are less than ideal as outcome measures for efficacy studies.** These measures reflect more than what students have learned or can do; in many cases, they also reflect student attendance, timely completion of assignments, and class participation. More precise measures of learning, such as common tests, and other more detailed behaviors, such as student assignment completion, should be tracked as outcomes in evaluations of the impact of introducing new kinds of courseware. Especially when comparing outcomes for course sections taught by multiple instructors with different or unknown grading policies, using these measures as student outcomes can bias impact estimates. However, grades and course credits are critical for students and to inform institutional policies so they should be captured in evaluation studies along with more direct measures of student learning.
- **Satisfaction surveys, whether of instructors or students, are insufficient as the only outcome in a courseware evaluation.** Instructor and student perceptions as expressed on surveys can be useful information, but positive survey responses and impacts on learning do not always go hand-in-hand. Moreover, these surveys are dogged by low response rates and selection bias. Use of more frequent but very brief surveys embedded in the courseware could help raise response rates and lower concerns about sampling bias, but a certain amount of judicious observation and interviewing is recommended to provide insights to complement learning data and reports from survey samples.
- **More research is needed to develop cost effective ways to capture the changes in instructional practice associated with implementations of adaptive courseware.** The ALMAP evaluation focused on time devoted to lecture and presentations, but future work should examine (1) how adaptive courseware affects the relative balance between low-level and high-level content interactions between instructors and students and (2) how the automated dashboards in adaptive courseware affect instructors' sensitivity to individual and whole-class learning needs.

FOR FUNDERS

- **Funders seeking to contribute to the knowledge base around courseware effectiveness should foster the use of controlled studies of courseware impacts.** Only 16 of the 64 implementations of courseware as part of ALMAP generated any adaptive courseware and comparison course section data incorporating measures to control for any pre-existing differences between the student groups being compared. The ALMAP evaluation revealed that even higher education institutions with a commitment to innovation and adoption of adaptive learning software are, by and large, unaccustomed to performing controlled studies of course outcomes. Further, adoption of common analytic approaches and common learning outcome measures for developmental and gateway courses would support aggregation of findings across institutions.

- **Postsecondary institutions can support student success by sharing anonymized data sets from side-by-side comparisons of adaptive courseware and other instructional approaches.** As more and more institutions turn to adaptive courseware, sharing data sets linking information about implementation, student administrative data, courseware system data, and student outcomes can build the empirical base needed to answer many of the questions left unanswered by the ALMAP data sets. Linked data sets need to be screened to insure they do not contain personally identifying information, but FERPA (Family Educational Rights and Privacy Act) compliant processes are readily available.
- **Grantees should be encouraged to develop unit-cost formulas that permit fair comparisons between adaptive courseware and comparison course sections.** Per-student costs depend greatly on the size of each course section, and major differences in class size (such as the transition from large lecture classes to blended instruction) can overwhelm any efficiencies or inefficiencies related to courseware per se. Cost analyses should not narrowly focus on whether costs increase or decrease, but rather should take into account the tradeoffs between changes in per-unit cost associated with a unit of learning gain.
- **Funders should encourage modeling of cost effectiveness over the longer term, not just for an individual course.** It is necessary to take a 3-5 year perspective on student outcomes to capture the monetary savings associated with lower odds of needing to retake a developmental course or higher persistence rates. It is also useful to separate the cost ingredients of up-front development, preparation, and infrastructure for initial implementation from those of ongoing instructional delivery.

In conclusion, although ALMAP grantees varied in their capacity for institutional research, the ALMAP grantees made good-faith efforts to address the challenge of measuring learning impacts, costs, and student and instructor perceptions. We acknowledge them not only for their leadership in innovating with new courseware but also for their willingness to share their experiences and findings with the field at large. We view this sharing of results as a first step toward cultivating a critical community of inquiry around the usefulness of new courseware in supporting student learning and course completion in institutions of higher education.