JOHNS HOPKINS
SCHOOL of EDUCATION

Study of Knewton Online Courses for Undergraduate Students: Examining the Relationships Among Usage, Assignment Completion, and Course Success

Rebecca Wolf, Ph.D.
Clayton Armstrong, B.A., B.S.
Steven M. Ross, Ph.D.

Johns Hopkins University
Center for Research and Reform in Education

August 2018

Table of Contents

Johns Hopkins School of Education

Executive Summary

Knewton provides online and adaptive learning courses for undergraduate students in chemistry, economics, mathematics, and statistics. Knewton contracted with the Center for Research and Reform in Education (CRRE) at Johns Hopkins University to conduct an independent study of the relationship between student use of Knewton's online courses and subsequent success in the course. The study addressed the following research questions:

1. To what degree does completion of an assignment or learning objective predict student performance outcomes?
2. How does usage relate to assignment completion and other student performance outcomes?
3. To what degree does completion of previous assignments predict completion of subsequent assignments and course dropout?
4. How do the above relationships vary across students with different ability levels or different courses?
5. What are the profiles of students who complete courses successfully, are on the borderline of course success, or are dropouts or failures?

Method

The study used Knewton data from the 2017 fall semester, which yielded a number of usage and student performance variables, including:

- Average score on online tests and quizzes
- Proportion of online assignments completed
- Potential course dropout
- Numbers of adaptive items, learning objectives, and assignments attempted

Two proxies of student ability also were derived from the data. Hierarchical linear modeling with students nested within courses determined the relationship between usage of the Knewton online learning platform and outcomes, while controlling for student ability.

Study Limitations

One limitation of this study is that the online components of the course may have been optional; therefore, it was impossible to distinguish in the data which assignments were required. As a result, the proportion of assignments completed may be biased downward if some students participated in optional assignments. The proportion of assignments completed also may have been confounded with course dropouts. If, for example, students dropped out of the course, then they had lower rates of assignment completion.

Another limitation identified in this study is that instructors may not have used Knewton's online assessments; average quiz/test scores were available for only 32% of students in the sample. A final limitation is that the estimated relationship between usage of Knewton and student performance may have been somewhat confounded with student ability. In other words, the two proxies of student ability were not sufficient to fully account for student ability in the statistical analyses. Nonetheless, these limitations are unavoidable given data and usage properties yet, in the strong opinion of the evaluators, do not preclude obtaining reasonable evidence addressing the major research questions.

Findings

Assignment completion predicted student performance on online assessments:

- A 10 percentage point increase in the proportion of assignments completed (of those offered in the course) was associated with an increase in average student performance of 1.4 percentage points.
- A 10 percentage point increase in the proportion of assignments completed (of those attempted by students) was associated with an increase in average student performance of 1.2 percentage points.

Completion of a single learning objective predicted student performance on online assessments:

- Completion of a single learning objective was associated with a 6.6 percentage point increase in the average score for all quiz/test items related to the learning objective.

Usage predicted student performance on online assessments:

- Attempting an additional 250 adaptive items (beyond the mean) was associated with higher average test and quiz scores by approximately 1.4 percentage points.
- Attempting an additional 10 assignments beyond the mean was associated with an improved average quiz/test score by 3.5 percentage points.

Usage predicted assignment completion:

- Attempting an additional 250 adaptive items (beyond the mean) was associated with an increase in the proportion of assignments completed by 14 percentage points.

Completion of previous assignments predicted completion of subsequent assignments:

- When students completed an additional 10% (beyond the mean) of the previous 25% of assignments in a course, completion in the subsequent 25% of assignments increased by 6–11 percentage points.

Assignment completion predicted dropout:

- When students completed an additional 10% (beyond the mean) of the first 25% of assignments in the course, students were more likely to remain engaged in the work throughout the duration of the course by 4 percentage points.

Results were similar for students of different ability levels or in courses in different subjects:

- Results were generally consistent across courses in different subjects, and for the most part, across students of different ability levels. The one exception was that students of higher ability completed assignments faster than lower ability students, which is expected given the adaptive nature of Knewton.

Conclusion

Students who engaged with more content on the Knewton online platform outperformed peers on online tests and quizzes, compared with peers of the same ability who used the platform to a lesser extent. Increased usage of Knewton also was associated with higher rates of assignment completion, and assignment completion was positively associated with higher average scores on tests and quizzes. Assignment completion earlier in the course predicted subsequent assignment completion, as well as whether the student remained engaged in the work throughout the duration of the course.

Students of all ability levels were able to successfully complete assignments, and students of all ability levels had similar rates of assignment completion. One potential explanation of this finding is that Knewton's adaptive platform allows students of all ability levels to complete assignments by providing low-ability students more items to master the content, as needed. In addition, assignment completion was more strongly correlated with usage of the Knewton platform than student ability, while student performance on online assessments was more strongly correlated with student ability than with usage. These findings suggest that students of all ability levels were able to successfully complete assignments and that while increased usage of Knewton yielded higher average quiz/test scores, performance on online assessments was explained more by student ability than by usage of the Knewton platform.

Across all measures, Knewton appeared to influence outcomes similarly for students of different ability levels. The only exception was that high-ability students completed assignments faster than low-ability students. Given the adaptive nature of the Knewton online platform, this finding is expected. Results were also consistent across different course subjects (e.g., chemistry, economics, mathematics, and statistics), but statistics courses used the Knewton online platform the least.

Overall, Knewton appears to be a useful tool for students. This study suggests a positive correlation among usage of Knewton, assignment completion, and performance on online assessments.

Introduction

Knewton provides online and adaptive learning courses for undergraduate students. In the four subject areas of mathematics, chemistry, statistics, and economics, each student interacts with adaptive assignments that assess level of mastery after each interaction. Adaptive models estimate when students complete a learning objective, and when several learning objectives related to a broader assignment have been completed, students then complete the assignment. As an option, course instructors may assign supplemental quizzes and tests administered to students through the platform.

Knewton contracted with the Center for Research and Reform in Education (CRRE) at Johns Hopkins University to conduct an independent study of the relationship between student use of Knewton's online courses and subsequent success in the course. The study used Knewton data from the 2017 fall semester to address the following research questions:

1. To what degree does completion of an assignment or learning objective predict student performance outcomes?
2. How does usage relate to assignment completion and other student performance outcomes?
3. To what degree does completion of previous assignments predict completion of subsequent assignments and course dropout?
4. How do the above relationships vary across students with different ability levels or different courses?
5. What are the profiles of students who complete courses successfully, are on the borderline of course success, or are dropouts or failures?

The remaining sections of this report review the methods, findings, and conclusion from the independent study conducted by CRRE.

Methods

Data

Data were from the fall 2017 semester and were at the item (adaptive, quiz, or test) level. Each item was linked to the institution and class identification numbers and to a particular learning objective and assignment. For each learning objective, student ability on the objective was estimated from custom item response theory models after each interaction. Also, provided after each interaction was the status of student learning on a particular learning objective (e.g., not started, in progress, or struggling). The data also included the time when the model estimated that students had "mastered" the core learning objectives in an assignment and thus completed the assignment, as well as the time of day of each interaction.

Student performance outcomes. Because course grades were not available, proximal outcomes for student performance were derived from the data and include:

- **Performance on course assessments**—Average quiz/test score, or the proportion of items answered correctly on quizzes and tests, aggregated overall, as well as at the learning objective level.[1]
- **Assignment completion**—Proportion of assignments completed calculated by dividing the number of assignments completed by the number of possible assignments in the course.
- **Assignment completion for assignments attempted only**—Calculated by dividing the number of assignments completed by the number of assignments attempted. This outcome is different from the previous one because it shows completion rate only for assignments that the student attempted. This completion rate may be artificially high, however, if the student attempted very few assignments.
- **Completion of a single learning objective**—Dummy variable indicating whether the student had completed the learning objective according to a statistical model (1=yes, 0=no).[2]
- **Potential course dropout**—Dummy variable indicating whether the student did not attempt any of the last 25% of assignments offered in the course (1=yes, 0=no).[3]

Table 1 shows the average student performance outcomes by course subject and the percentage missing. Notably, the majority of students were missing either a quiz or test score. In addition, students attempted and completed the majority of assignments offered in online courses. Moreover, as expected, assignment completion rates were higher when restricting to assignments that students had attempted, as opposed to all assignments offered in the course. Finally, across all course subjects, approximately one-fifth of students did not remain engaged in the online platform through the end of the course.

---

[1] Note that for a very small percentage of students, their overall test/quiz average included scores from pre-quizzes and pre-tests. Additionally, we did not examine performance on course assessments at the assignment level because the vast majority of assignments could not be linked with any type of assessment.

[2] We first isolated the progress for the target learning objective linked to the item. We then assumed that learning objectives were completed if the model-estimated student ability for the objective was greater than or equal to .99.

[3] Note that assignments may have been optional for students.

Table 1: *Average student performance outcomes by course subject*

	CHEM (%)	***ECON (%)***	***MATH (%)***	***STAT (%)***	***Missing (%)***
Online assessments					
Quiz item score	84	65	76	74	86
Test item score	74	62	71	75	78
Quiz and test item score (combined)	75	64	73	74	68
Assignments					
Completion rate	64	64	74	63	0
Percent attempted	71	72	81	76	0
Completion rate if attempted	82	85	87	77	0
Potential course dropout	26	28	28	16	0

We also derived variables to explore to what degree completion of previous assignments predicted completion of subsequent assignments. We first grouped assignments by course into the first 25% of assignments, the second 25% of assignments, the third 25% of assignments, and the last 25% of assignments. For each category, we then calculated the percentage of assignments that each student completed but restricted the calculation to include only the assignments that students attempted to avoid conflating assignment completion with course dropout in this analysis.[4]

Figure 1 also shows that the vast majority of students completed assignments when they attempted them. In addition, rates of assignment completion generally decreased over the duration of the course and in most courses, increased again towards the end of the course. This finding was true for students of all ability levels.[5] One plausible explanation is that student interest in a course waned over time, but students who were concerned with their course grades made a final push at the end of the semester.

[4] By definition, course dropout is synonymous with not participating in any of the last 25% of assignments in a course.
[5] See the proxies of student ability section for more information about how student ability was derived.

Figure 1: *Assignment completion (for assignments attempted) over duration of course by student ability level and course subject*

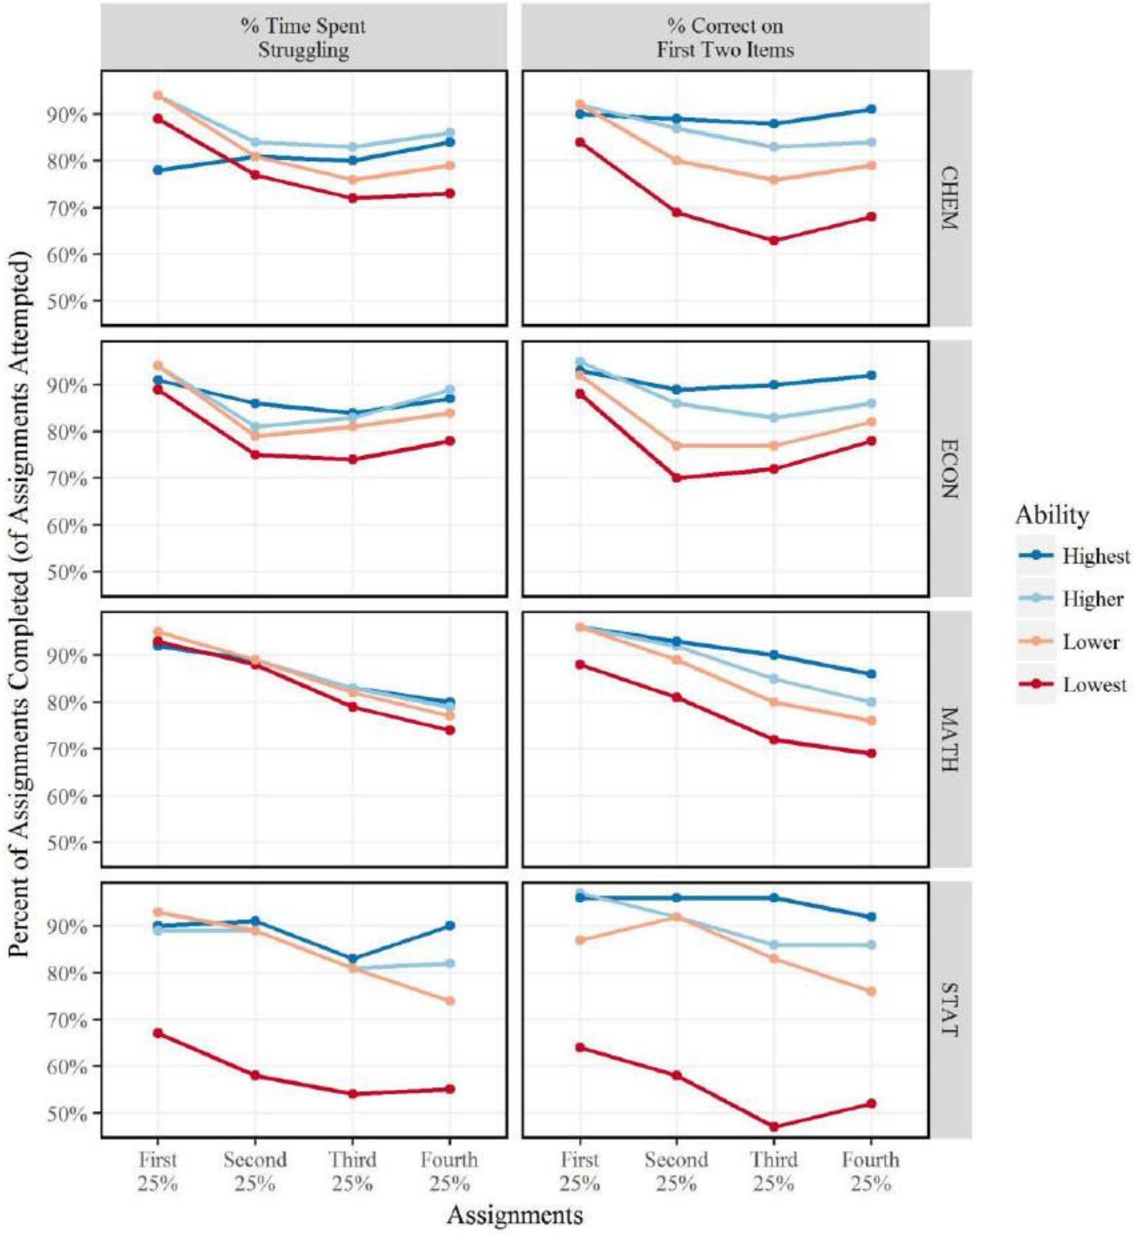

Course usage of Knewton. Use of Knewton's online learning may have been optional for course instructors, and instructors were able to customize online content to meet their needs. As a result, courses widely varied to the extent that they used components on Knewton online learning, even within the same academic subject. Table 2 outlines descriptive information (e.g., the minimum, mean, and maximum) of the number of Knewton learning objectives[6] and assignments covered in each course, the number of assessments administered online in each course, the duration (in days) of each course, and the number of students enrolled in the course.

[6] Note that some learning objectives may have been remedial and offered to only a subset of students in the course.

Statistics courses appeared to use Knewton the least, compared with chemistry, economics, and mathematics courses. In addition, some courses did not include online quizzes or tests.

Table 2: *Descriptive information about courses by course subject*

	Min	Mean	Max
Chemistry			
Number of learning objectives	6	107.6	210
Number of assignments	2	29.1	61
Number of quizzes	0	2.6	35
Number of tests	0	1.8	10
Duration of online portion of course in days	28	98.8	142
Number of students per course	5	55.9	353
Economics			
Number of learning objectives	4	87.5	148
Number of assignments	1	28.2	50
Number of quizzes	0	6.5	32
Number of tests	0	1.5	11
Duration of online portion of course in days	36	89.7	121
Number of students per course	3	46.3	269
Mathematics			
Number of learning objectives	4	128.6	291
Number of assignments	1	34.9	80
Number of quizzes	0	0.8	37
Number of tests	0	2.9	138
Duration of online portion of course in days	21	102.2	130
Number of students per course	5	34	160
Statistics			
Number of learning objectives	17	52	144
Number of assignments	3	12.3	33
Number of quizzes	0	0.7	2
Number of tests	0	0.7	2
Duration of online portion of course in days	49	93.2	125
Number of students per course	15	33.7	85

Student usage variables. The extent to which students used Knewton online courses was also derived from the data using the following variables:

- **Number of adaptive items attempted**—The number of all unique adaptive items that a student was exposed to in the online platform. Adaptive items comprised the vast majority of all items of any type (e.g., adaptive, content, test, quiz).
- **Number of learning objectives attempted**—The number of all unique learning objectives that a student was exposed to in the online platform.[7]
- **Number of assignments attempted**—The number of all unique assignments that

[7] This was calculated on the basis of "adaptive" items only. Additionally, some learning objectives may have been remedial and offered to only a subset of students in the course.

a student was exposed to in the online platform.[8]

Table 3 outlines the average student values for these usage variables by course subject. Across all courses and on average, students in the sample attempted 739 adaptive items and were exposed to 77 unique learning objectives throughout 26 unique assignments. In addition, students attempted an average of 12 adaptive items to complete a learning objective, and 35 adaptive items to complete an assignment.

Table 3: *Average student usage by course subject*

	CHEM	***ECON***	***MATH***	***STAT***
Number of adaptive items attempted	724	942	658	397
Number of learning objectives attempted	73	66	88	39
Number of assignments attempted	23	22	29	12
Number of adaptive items per learning objective	11	17	9	13
Number of adaptive items per assignment	36	50	26	39

NOTE. The numbers are greater in Table 2 than in Table 3 because the numbers in Table 2 were calculated for all students in the course, whereas the numbers in Table 3 were calculated for individual students. In other words, all students in a course will engage with more adaptive items, learning objectives, and assignments than any individual student.

Proxies for student ability. We designed this study to explore the relationship between use of Knewton's online courses and student performance outcomes while accounting for differences in students' ability levels. To best proximate student ability, we used two different variables simultaneously:

- We used as one measure of student ability the model-based estimates that flagged a student as struggling on a particular learning objective and calculated the proportion of interactions where a student was flagged as struggling.[9] Students who struggled the least were identified as high-ability students, and students who struggled the most were identified as low-ability students.
- We also used average student performance on the first two adaptive items for all new learning objectives. In theory, a student's performance on these initial items was one proxy for student ability prior to use of the Knewton online platform.

These two proxies for student ability were negatively correlated at $\rho = -.46$, as expected, given their definitions. While the two proxies were related, they each captured a unique portion of students' true ability.

One limitation of these proxies was that they may overestimate student ability if students only engaged in the online platform earlier in the course, and earlier assignments were easier than subsequent ones. Given expected positive relationships between these proxies of student

[8] This was calculated on the basis of "adaptive" items only.

[9] We first isolated the struggling status for the target learning objective. We then calculated the proportion of interactions where a student was flagged as struggling by dividing the number of instances of "struggling" by the sum of the instances of "struggling" and "in progress."

ability and student performance on online assessments, however, it appears that these proxies were reliable.

Students were also categorized into one of four ability categories (highest, higher, lower, and lowest) on the basis of each ability variable and by course subject. Students were categorized according to quartiles, and the 25% lowest ability students in the course subject constituted the lowest ability group, and so on. Figures 2 and 3 show the average values of the two proxies of student ability by ability category and by course subject.

Figure 2: *Average percentage of time spent "struggling" for students of different ability levels*

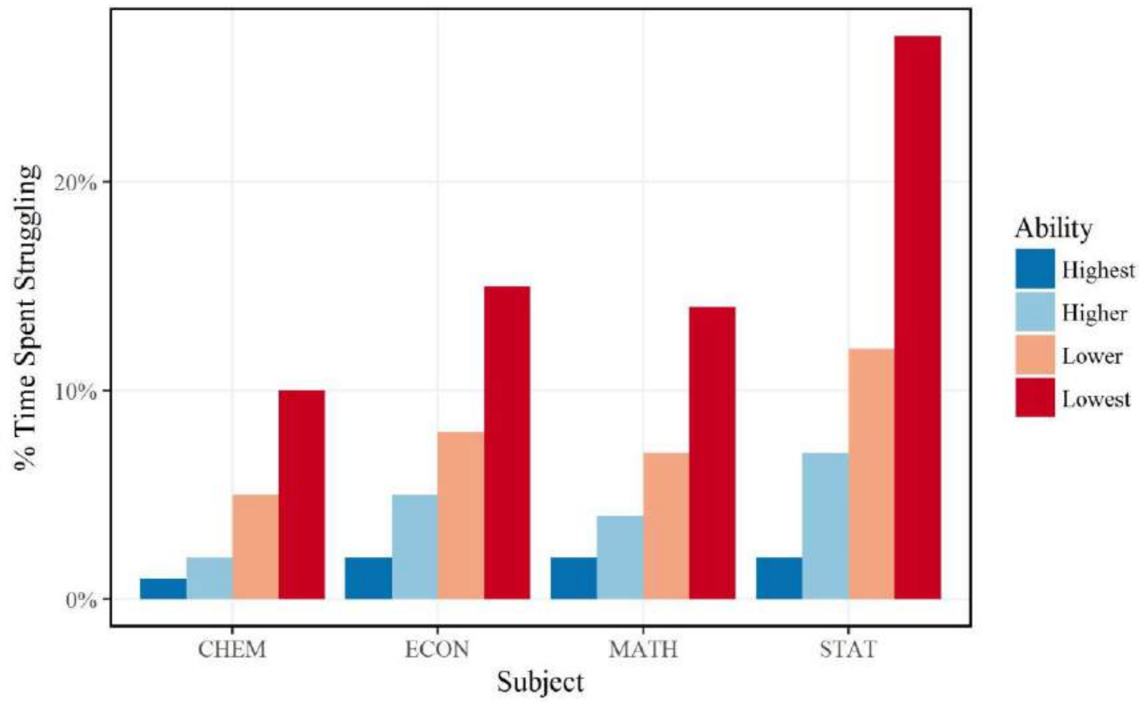

Figure 3: *Average percentage correct on first two items per learning objective for students of different ability levels*

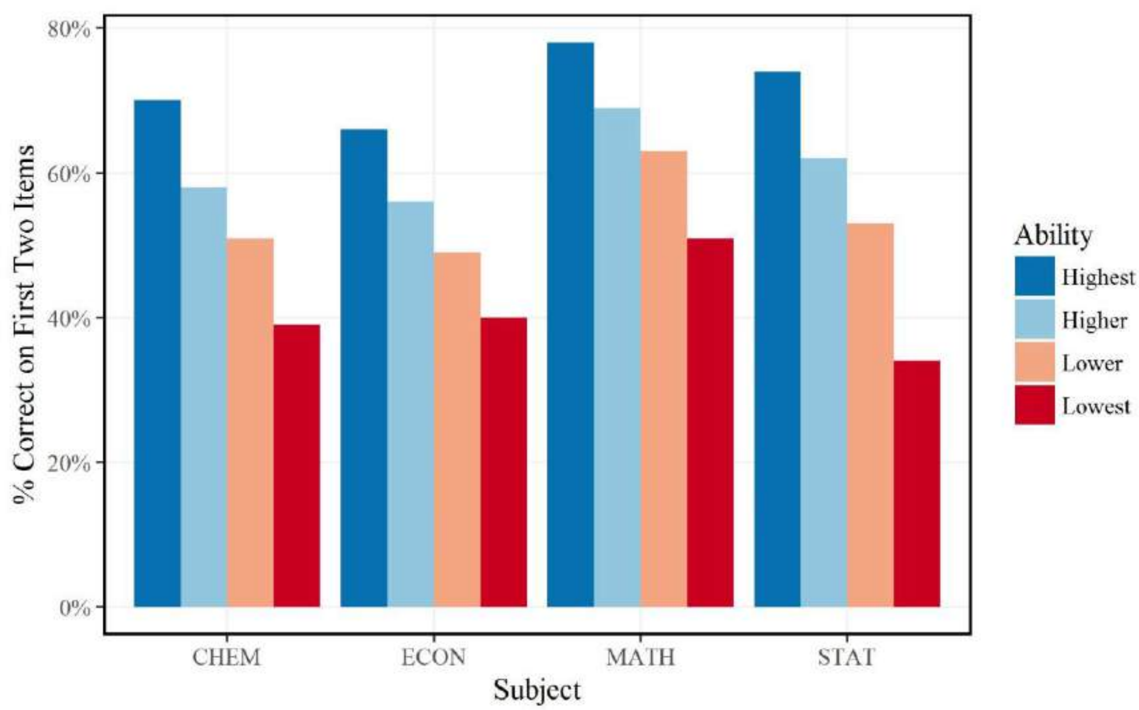

To determine whether overall study results were similar for students of all ability levels, subgroup analyses were conducted for students in each ability quartile separately. We used the proportion of time the student was deemed as "struggling" as the proxy for student ability for the subgroup analyses, although using the other proxy for student ability (e.g., average score on first two items) produced similar results.

Sample Restrictions

To understand how participation in Knewton's online learning related to student performance outcomes, we restricted the analyses to courses and students where students appeared to have been exposed to the online learning to a non-trivial degree. To do so, we excluded courses and students where:[10]

- The number of students in the course who used the online platform was one.[11]
- The duration of the online platform component of the course was less than or equal to 10 days.[12]
- Students used Knewton for only a 10-day period or shorter, individually or on average for a course.
- The total number of adaptive items (for the course or student) was zero.
- The total number of assignments for the course was zero.

Ultimately, the majority of courses and students were retained in the analytic sample, and all students in the sample attempted at least one assignment. Table 4 shows the numbers of higher education institutions, courses, and students in the analytic sample.

Table 4: *Sample sizes by course subject*

	CHEM	***ECON***	***MATH***	***STAT***
Number of institutions	20	23	23	5
Number of courses	29	44	104	9
Number of students[13]	1,226	1,783	3,018	251

[10] We fully acknowledge that these decisions are subjective and arbitrary. However, these decisions were made in conjunction with reviewing patterns in the data. We attempted to exclude courses and students who did not appear to meaningfully engage in Knewton's online learning.
[11] This decision was iterative in that if other exclusions caused the class size to be reduced to one, the student and course then were dropped.
[12] This was calculated by taking the difference between the first and last timestamp for the course.
[13] There were a few students who were enrolled in more than one course participating in Knewton online learning and thus were counted more than once. The numbers in Table 4 represent the unique users by user identification number, institution number, and class number.

Analytic Approach

We used hierarchical linear modeling to explore the relationship between usage of Knewton and student performance outcomes, while accounting for the nested nature of the data (e.g., students within courses).[14] The model to estimate program effects can be written generally as:

$\mathrm{Y}_{ij} = \gamma_{00} + \gamma_{01} struggle_{ij} + \gamma_{02} first\ two_{ij} + \gamma_{03} usage_{ij} + \gamma_{0\mathrm{k}} \sum college\ dummy\ indicators_j + u_{0j} + \mathrm{r}_{\mathrm{ij}}$ where:

Y_{ij}: Student performance outcome for student i in class j
γ_{00}: Grand mean
γ_{01}: Regression coefficient for struggle
$struggle_{ij}$: Proportion of time student was deemed "struggling" by model for student i in class j
γ_{02}: Regression coefficient for first two item score
$first\ two_{ij}$: Proportion of the first two adaptive items per learning objective the student answered correctly for student i in class j
γ_{03}: Regression coefficient for usage variable
$usage_{ij}$: Usage amount (defined in multiple days) for student i in class j
$\gamma_{0\mathrm{k}}$: Regression coefficients for the k college dummy indicators to account for college effects
$\sum college\ dummy\ indicators_j$: Vector of dummy indicators to account for college effects for class j
u_{0j}: Random school effect for class j
r_{ij}: Residual for student i in class j

The model above was adapted in cases when Y_{ij} was a binary variable (e.g., potential dropout) to a multilevel mixed-effects logistic model.[15] The model also was adapted to explore relationships at the learning objective level—the model was estimated with learning objectives nested within courses.[16] For some models, quadratic terms of the student ability and usage variables also were added to the model to best explain the data (see the Technical Appendix for a full list of predictor variables for each model estimated). Finally, to facilitate interpretation, the predictor variables were grand-mean centered, with the exception of the dummy variable indicating learning objective and assignment completion.[17]

To determine the relationship between usage and outcomes for students of different prior ability or in different subjects, we re-estimated the models for specific student subgroups.[18] We created four categories of prior student ability by using the quartiles of the proportion of time spent "struggling" within each subject (e.g., Chemistry, Economics, Mathematics, and Statistics). These four categories were highest ability, higher ability, lower ability, lowest ability. By design, each category contained approximately 25% of all students in the sample. We also separately re-

[14] It did appear that there was meaningful variation between classes that would be unaccounted for otherwise.
[15] For these logistic models, the regression coefficients were in terms of log-odds.
[16] This model also used a multilevel mixed-effects logistic model due to violations of the assumptions of the hierarchical linear model.
[17] Enders, C. K., & Tofighi, D. (2007). Centering predictor variables in cross-sectional multilevel models: a new look at an old issue. *Psychological methods*, *12*(2), 121.
[18] When re-estimating the model for a specific student ability subgroup, we did not include additional proxies of student ability in the model.

estimated the models for each course subject: Chemistry, Economics, Mathematics, and Statistics.

Study Limitations

One limitation of this study is that the online components of the course may have been optional; therefore, it was impossible to distinguish in the data which assignments were required. As a result, the proportion of assignments completed may be biased downward, if some students participated in optional assignments. The proportion of assignments completed also may be confounded with course dropout. If, for example, students dropped out of the course, then they had lower rates of assignment completion.

Another limitation is that instructors may not have used Knewton's online assessments; on average, quiz/test scores were available for only 32% of students in the sample. Moreover, for a very small percentage of students, the overall test/quiz score average included scores from pre-tests and quizzes. Analyses of student performance on online assessments, however, provided some indication of how usage of Knewton related to student performance in the course.

A final limitation is that the estimated relationship between usage of Knewton and student performance may have been somewhat confounded with student ability. In other words, the two proxies of student ability were not sufficient to fully account for student ability in the statistical analyses. Considering all limitations, the correlational findings obtained in the study cannot support causal claims, yet they do provide meaningful suggestive evidence of positive associations between usage of Knewton and student performance outcomes.

Findings

Assignment Completion and Student Performance

To determine the degree to which assignment completion predicted student performance outcomes, we examined the association between a student's average score on all quizzes and tests and the proportion of assignments completed—of those offered in the course or of those attempted by the student—while controlling for student ability. Results showed a positive trend, where completing a greater proportion of assignments positively correlated with improved student performance on quizzes and tests. Specifically, a 10 percentage point increase in the overall assignment completion rate was associated with an increase in average student performance of 1.4 percentage points. Restricting to assignments attempted only, a 10 percentage point increase in the assignment completion rate was associated with an increase in average student performance of 1.2 percentage points.

Figure 4 shows to what extent the completion of an additional 10%, 20%, or 30% of assignments of those offered in the course was associated with improved student performance on quizzes and tests. Figure 4 also shows to what extent completion of an additional 10% of assignments of those attempted was associated with improved student performance; the average assignment completion rate was already very high for assignments attempted.

Figure 4: *Relationship between assignment completion and average student performance*

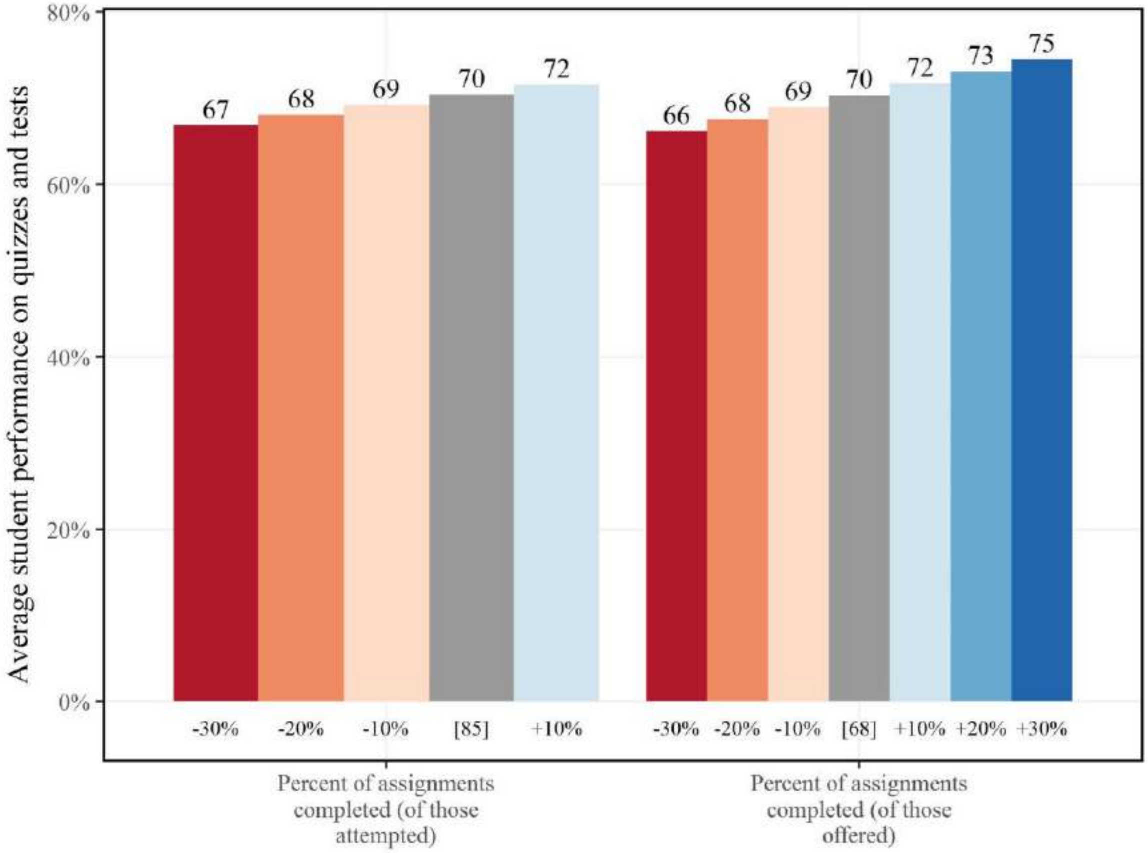

NOTE—The average assignment completion rate for all assignments offered in the course was 68%, and the average assignment completion rate for assignments attempted by students was 85%. These averages were calculated for students who were not missing an average score on quizzes and tests.

Learning Objective Completion and Student Performance

We also examined the relationship between completion of a single learning objective and student performance on paired quiz and test items that could be linked to the particular learning objective. We found that completion of a single learning objective was indeed associated with a 6.6 percentage point increase in the average score on all related quiz and test items. Given that the average student was exposed to 77 unique learning objectives, a 6 percentage point gain for completion of a single learning objective may be impactful.

Usage and Student Performance

Next, we examined how usage related to student performance outcomes and assignment completion, while controlling for student ability. We found that the more the student used the Knewton online platform—in terms of number of adaptive items, learning objectives, and assignments attempted—the higher the student's average quiz/test score. For example, attempting an additional 250 adaptive items beyond the mean was associated with higher average test and quiz scores by approximately 1.4 percentage points. Similarly, attempting an additional 10 assignments beyond the mean was associated with an improved average quiz/test score by 3.5 percentage points. Finally, attempting an additional 25 learning objectives beyond the mean was associated with an improved average quiz/test score by three percentage points. Figure 5 demonstrates how average student performance on tests and quizzes was associated with increased usage of Knewton at various levels of usage.

Figure 5: *Relationship between usage and average student performance*

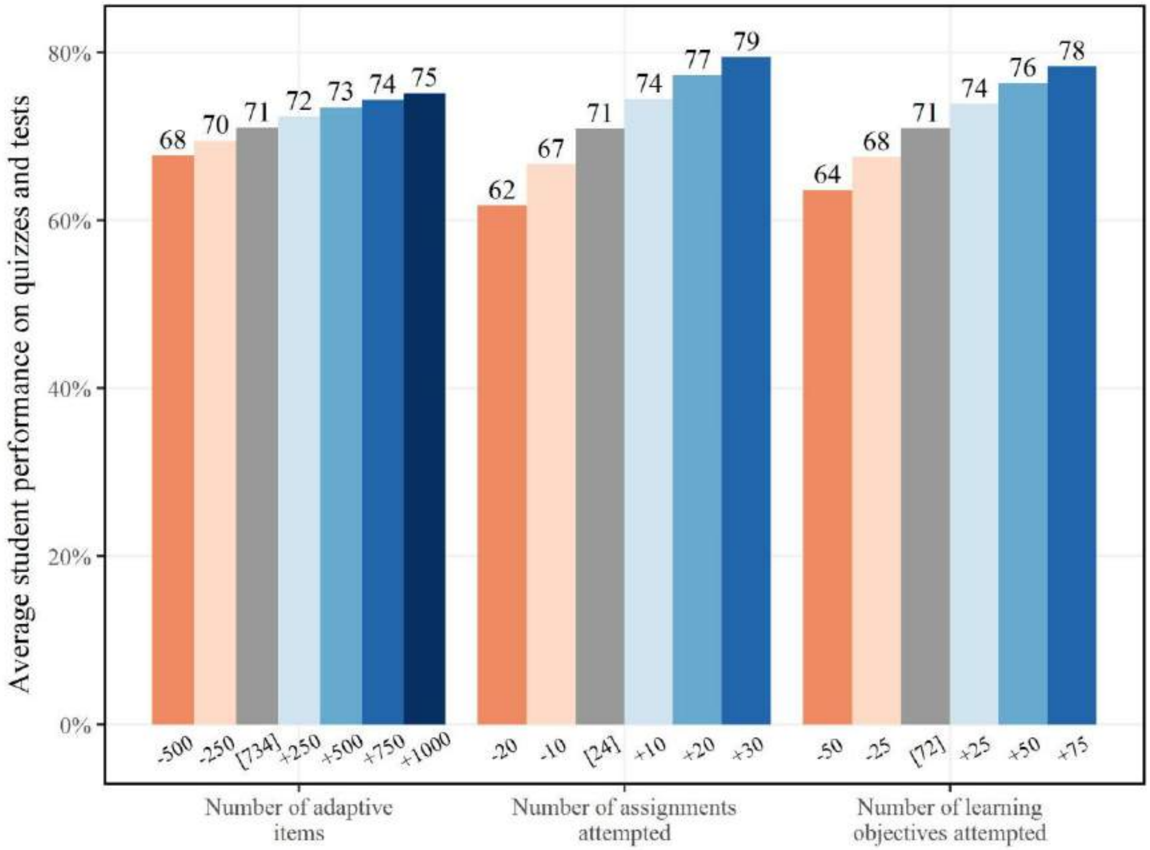

NOTES—

1. The mean number of adaptive items attempted across all courses was 734, the 25th percentile was 279 items, the 50th percentile was 650 items, and the 75th percentile was 951 items.
2. The mean number of assignments attempted across all courses was 24, the 25th percentile was 11 assignments, the 50th percentile was 26 assignments, and the 75% percentile was 33 assignments.
3. The mean number of learning objectives attempted across all courses was 72, the 25th percentile was 33 objectives, the 50th percentile was 77 objectives, and the 75th percentile was 102 objectives.
4. These descriptive statistics were calculated for students who were not missing an average score on tests and quizzes.

For all models including usage variables (e.g., number of adaptive items, assignments, or learning objectives attempted), there was a diminishing return as usage increased. Technically, the relationship between the usage variable and the outcome was not linear because both the usage variable and the square of the usage variable were statistically significant. The regression coefficient of the square of the usage variable was negative, albeit close to zero, which explained the slight diminishing return of usage.

Usage also related to assignment completion. The more adaptive items a student attempted, the higher the proportion of assignments completed (of those offered in the course). Specifically, attempting an additional 250 adaptive items beyond the mean was associated with an increase in the proportion of assignments completed by 14 percentage points. Relatedly, the more learning objectives a student was exposed to, the higher the assignment completion rate. Exposure to an additional 25 learning objectives beyond the mean was associated with a greater proportion of assignments completed by 19 percentage points. Figure 6 shows the relationship between usage and assignment completion.

Figure 6: *Relationship between usage and assignment completion (of those offered in the course)*

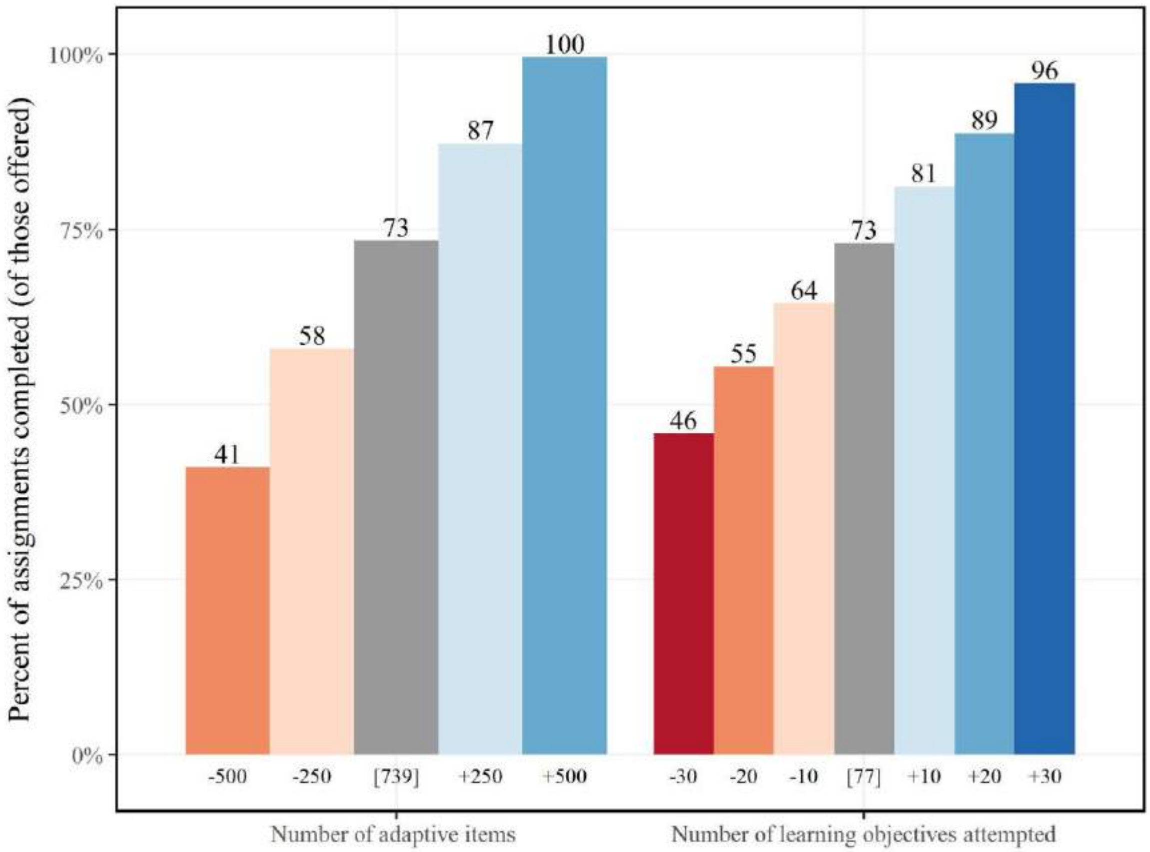

NOTES—

1. The mean number of adaptive items attempted across all courses was 739, the 25th percentile was 328 items, the 50th percentile was 628 items, and the 75th percentile was 942 items.
2. The mean number of learning objectives attempted across all courses was 77, the 25th percentile was 43 objectives, the 50th percentile was 82 objectives, and the 75th percentile was 105 objectives.
3. These descriptive statistics were calculated for all students in the sample.

Considering Figures 5 and 6 together, it appears that there was a stronger relationship between usage (defined multiple ways) and assignment completion than between usage and student performance on online assessments. Examining unadjusted pairwise correlations confirms this hypothesis; correlations showed that usage variables and assignment completion variables were strongly correlated, and usage variables and the test/quiz score average were weakly correlated. Moreover, student ability (in terms of time spent struggling) was modestly correlated with the test/quiz score average and weakly correlated with assignment completion rates. Thus, student performance on online assessments appeared to be more correlated with student ability than usage, while assignment completion rates appeared to be more correlated with usage than ability. These findings indicate that students of all ability levels successfully completed online assignments and that increased usage of Knewton yielded higher average quiz/test scores, but student ability was still a strong predictor of how well students did on assessments, regardless of time spent in the online platform.

Completion of Previous and Subsequent Assignments

We determined the degree that completion of previous assignments predicted completion of subsequent assignments, also controlling for student ability. For this analysis, we restricted to assignments that students had attempted to understand how early success with course content related to success later in the course in terms of completing assignments.

When students attempted assignments, assignment completion rates were generally high. In the first 25% of assignments, the average completion rate was 94%. In the second 25% of assignments, the average completion rate was 88%; when students completed an additional 10% of assignments in the second set, their completion rate in the third 25% of assignments increased by an average of 6 percentage points. In the third 25% of assignments, the average completion rate was 84%; when students completed an additional 10% of assignments in the third set, their completion rate in the fourth 25% of assignments increased by 11 percentage points, on average. Thus, completion of prior assignments predicted completion of subsequent assignments, even when controlling for student ability. Figure 7 shows how completion in the previous 25% of assignments related to completion in the subsequent 25% of assignments.

Figure 7: *Relationship between previous and subsequent assignment completion (of those attempted)*

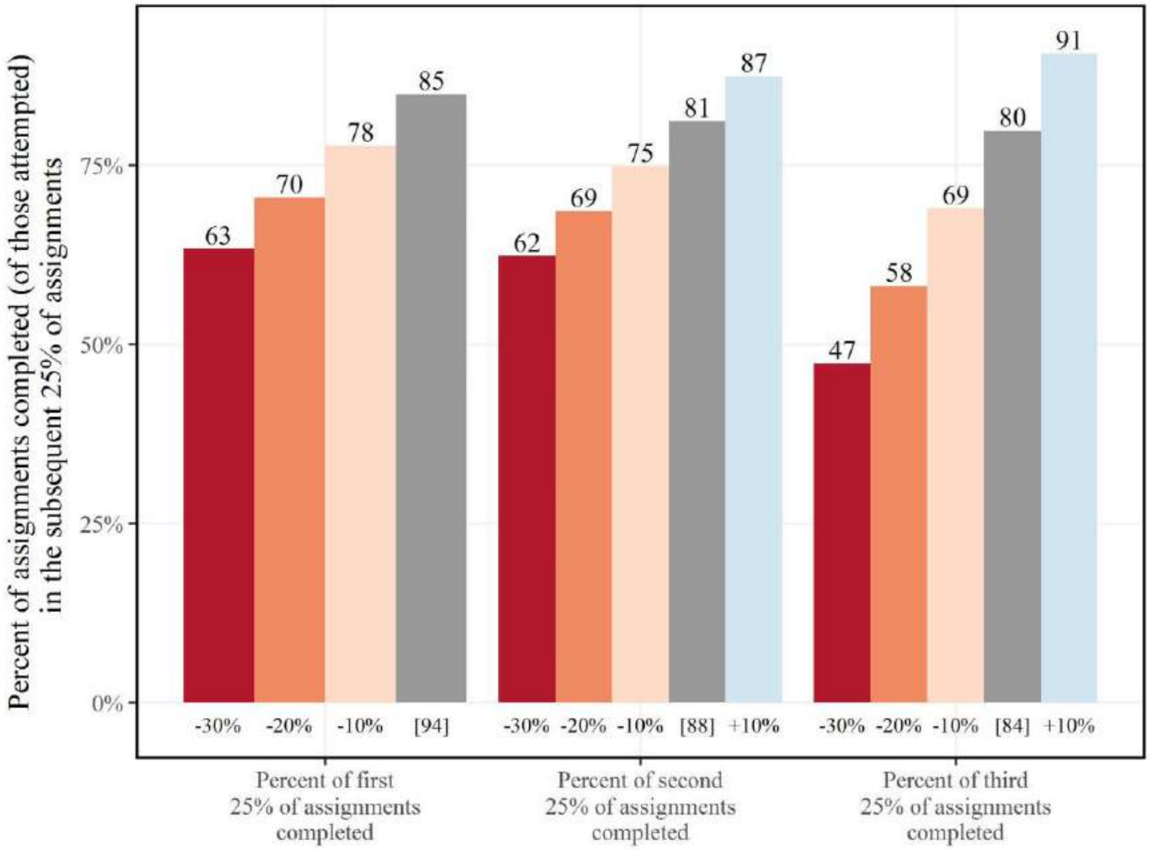

NOTE—The average assignment completion rate for the first 25% of assignments was 94%, and it was 88% for the second 25% of assignments, and 84% for the third 25% of assignments. These averages were calculated for students who were not missing a completion rate for both the previous and subsequent 25% of assignments in each analysis.

The decreasing effects of previous assignment completion predicting subsequent assignment completion over the duration of the course was partly attributed to potential course dropout. Students who dropped out of the course were present in the earlier analyses and were excluded from the later analyses due to missing data. Thus, it is likely that the decreasing effects are explained by this phenomenon, as opposed to the earlier assignments being more impactful than the later ones.

Course Dropout

Another outcome of interest is whether students dropped out of the course. Although students may drop out of the course for personal reasons unrelated to Knewton, students may also drop out if they feel unsuccessful. We explored to what extent assignment completion at various points in the course was related to potential course dropout. Roughly 22% of students were potential course dropouts, given that they failed to complete any of the last 25% of assignments in the Knewton platform.

Of assignments attempted, completing an additional 10% (beyond the mean) of the first 25% of assignments in a course was related to lower course dropout by four percentage points. Completing an additional 10% of the second 25% of assignments was related to lower course dropout by 3 percentage points. Completing an additional 10% of the third 25% of assignments was related to lower course dropout by 1.5 percentage points. Hence, successful completion of assignments earlier in the course was related to whether the student remained engaged throughout the duration of the course.[19] Figure 8 demonstrates the relationship between dropout and assignment completion throughout the course.

Figure 8: *Relationship between assignment completion (of those attempted) and course dropout*

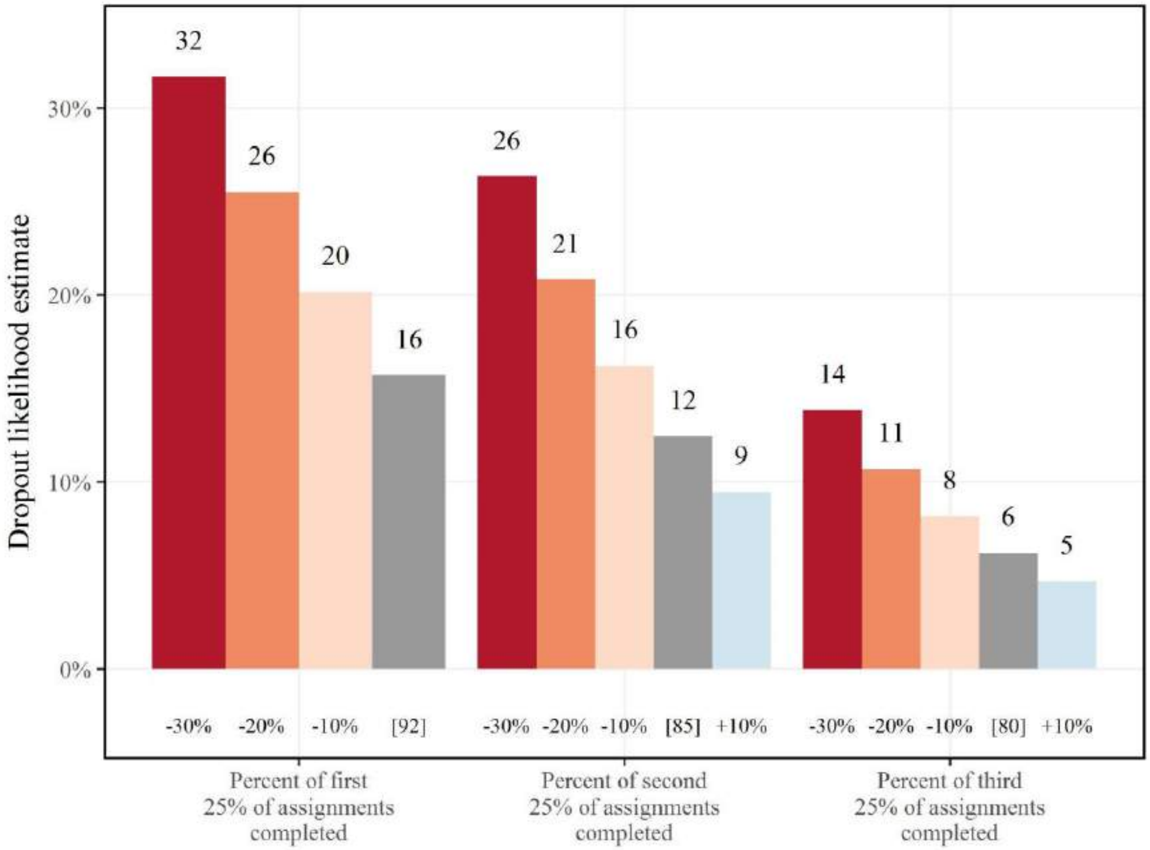

NOTE—The average assignment completion rate for the first 25% of assignments was 92%, and it was 85% for the second 25% of assignments, and 80% for the third 25% of assignments. For each analysis, these averages were calculated for students who were not missing the assignment completion rate. By definition, dropout was not missing for any student.

Results for Students of Different Ability Levels or in Different Courses

We examined whether findings varied for students with different ability levels (e.g., highest, higher, lower, or lowest) and generally found similar trends for students of different ability levels in the relationships among usage, assignment completion, and performance on tests

[19] Again, the decreasing effects over the duration of the course were confounded with course dropout.

and quizzes.[20] The only exception was that higher-achieving students completed assignments faster (with fewer adaptive items) than lower-ability students. This finding is expected, given the adaptive nature of the Knewton online platform. Lower-achieving students were presented with additional questions until the platform determined that the student had completed the assignment.

Figure 9 displays the relationship between assignment completion (of those offered in the course) and usage for students of different ability levels, given the same amounts of usage. This figure shows that the higher-ability students completed assignments faster (with fewer adaptive items) than lower-ability students because the slopes of the lines were greater for students with higher ability. Figure 9 also demonstrates that higher-ability students completed more assignments (of those offered in the course) than the highest-ability students. Higher-ability students completed 73% of assignments, on average, whereas the highest-ability students completed 64% of assignments. However, higher-ability students also attempted more items (633), on average, than the highest-ability students (416).

Figure 9: *Relationship between assignment completion (of those offered in course) and usage by student ability*

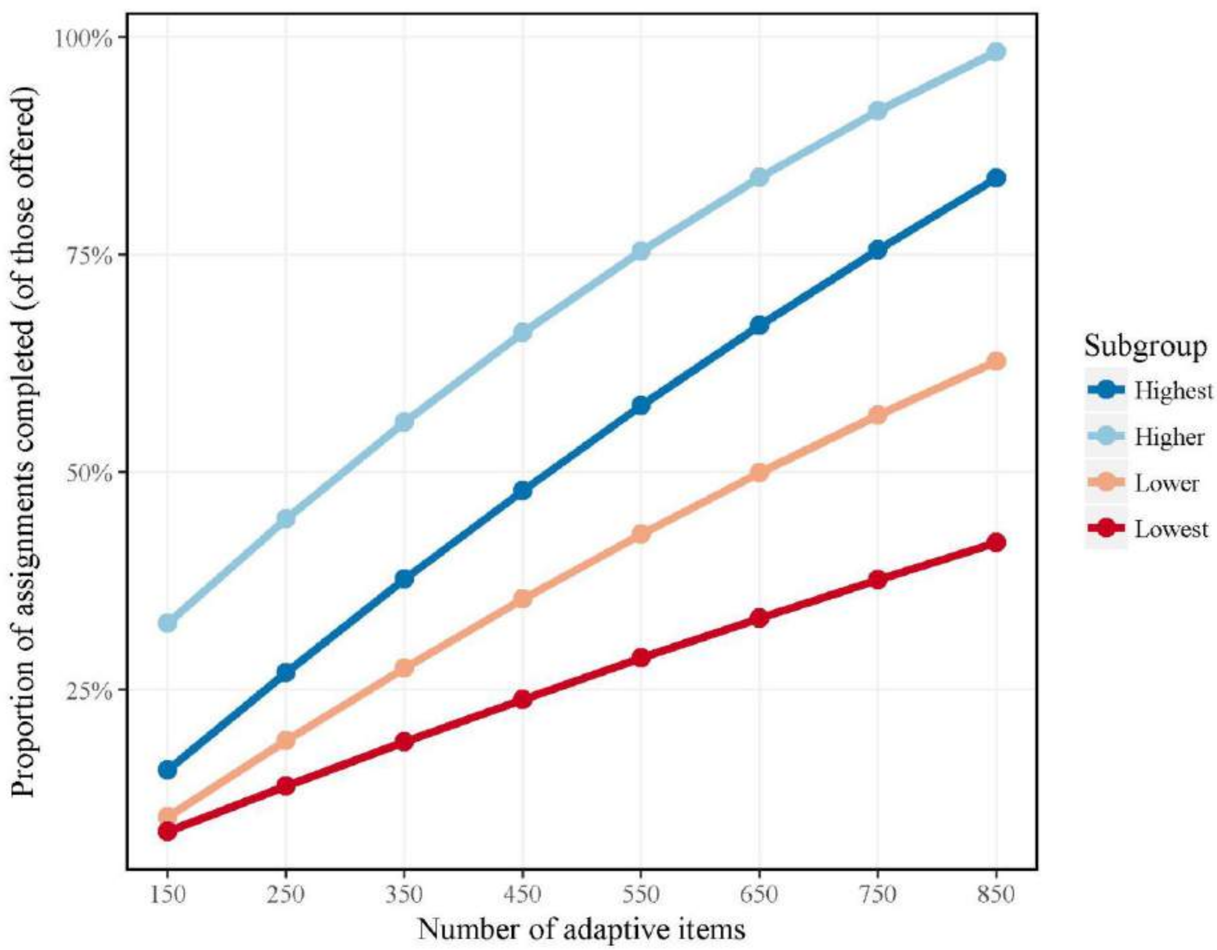

[20] We used the proportion of time deemed "struggling" as the proxy for student ability for the subgroup analyses.

There was also a slightly higher return on the number of adaptive items for higher-ability students in preparing the students for quizzes and tests. Because higher-ability students needed fewer adaptive items to complete assignments, higher-ability students scored higher than lower-ability students with the same level of usage. Figure 10 shows the relationship between average test/quiz score and usage for students of different ability levels, given the same amounts of usage.

Figure 10: *Relationship between average test/quiz score and usage by student ability*

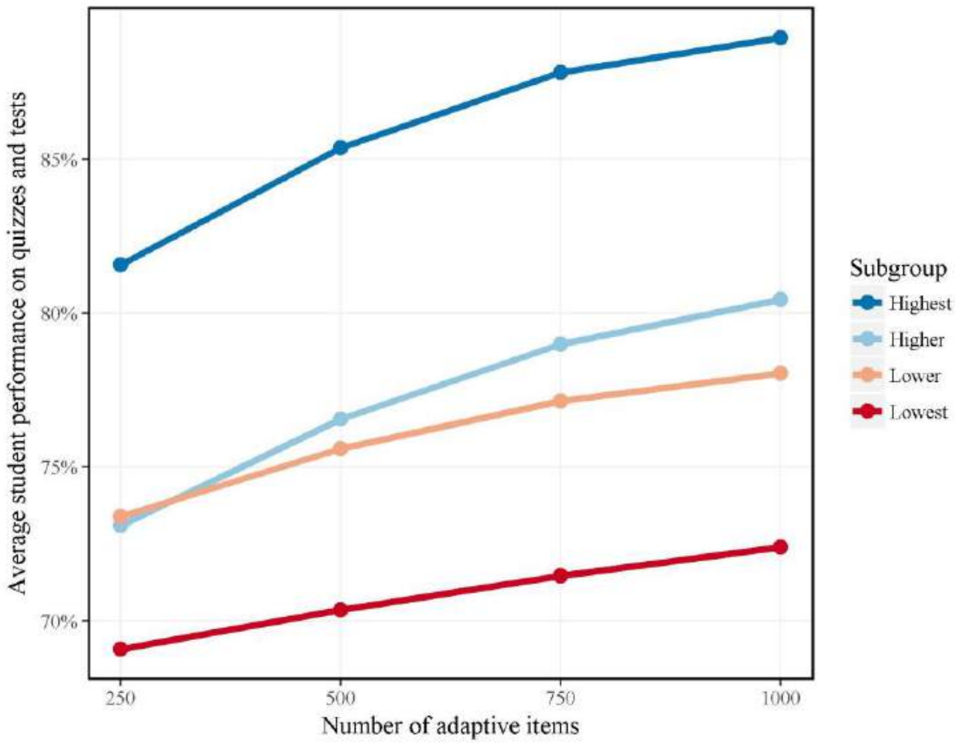

We also examined differences in findings across courses in the four subject areas, chemistry, economics, mathematics, and statistics. We found similar patterns in the relationships among usage, assignment completion, and performance across the different course subjects, when restricted to plausible ranges of student usage by course subject. The statistics courses differed from courses in other subjects, however, due to their limited use of the Knewton platform.

Student Profiles

To determine profiles of students who completed courses successfully, were on the borderline of success, or were dropouts or failures, we created categories of course success using student performance and assignment completion. Specifically, we examined characteristics of students by categories of (a) average grade on all tests and quizzes, (b) proportion of assignments completed (of those offered in the course), and (c) potential course dropout.

Average grade on tests and quizzes. Consistent with earlier findings and in general, students who had greater exposure to the Knewton online platform had higher average performance on tests and quizzes, as well as higher completion rates of learning objectives and assignments (of those offered in the course). The one exception was that students who scored an A or B average on quizzes and tests attempted fewer adaptive items than students who scored worse; however, the number of adaptive items was correlated with student ability where lower achieving students had to interact with more items to complete assignments than did higher achieving students. Moreover, when we controlled for student ability in the previous analyses, we found a positive relationship between the number of adaptive items attempted and average performance on quizzes and tests. Figure 11 outlines the characteristics of students by average grade on tests and quizzes. Note also that student ability was correlated with average score on quizzes and tests.

Figure 11: *Characteristics of students by average grade on tests and quizzes*

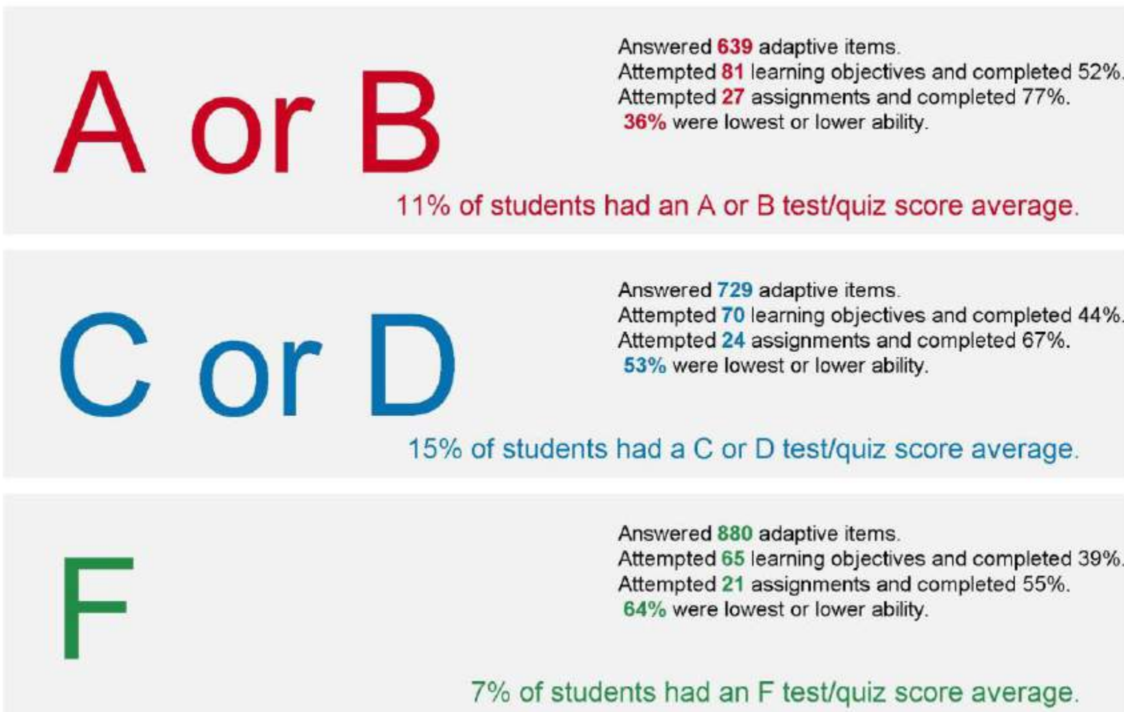

NOTE–68% of students were enrolled in courses that did not utilize online tests or quizzes.

Proportion of assignments completed. Second, students who had greater exposure to Knewton also had higher rates of assignment completion (of those offered in the course). However, exposure to a greater number of learning objectives and assignments did not appear to be as related to assignment completion rates as students' ability to successfully complete individual learning objectives, which can be accomplished by attempting more items. Also of note is that assignment completion did not appear to be as correlated with student ability as student performance on quizzes and tests. Thus, it appears that assignment completion was possible for all students with adequate usage of Knewton, regardless of student ability. Figure 12 outlines the characteristics of students by proportion of assignments completed.

Figure 12: *Characteristics of students by percent of assignments completed (of those offered in the course)*

Potential course dropout. Finally, as expected, students who may have dropped out of the course used Knewton less and completed fewer assignments and learning objectives, compared with students who remained engaged in the Knewton platform throughout the duration of the course. Unexpectedly, however, students who remained engaged throughout the course had lower ability, on average, than those who may have dropped out of the course. Although there was an expected positive relationship between performance on tests/quizzes and the proxy for student ability, this finding calls into question to what extent course dropout was accurately captured in this study, given data constraints. Figure 13 outlines the characteristics of students for potential course dropouts and for students who remained engaged in the online platform throughout the duration of the course.

Figure 13: *Characteristics of students by potential course dropout*

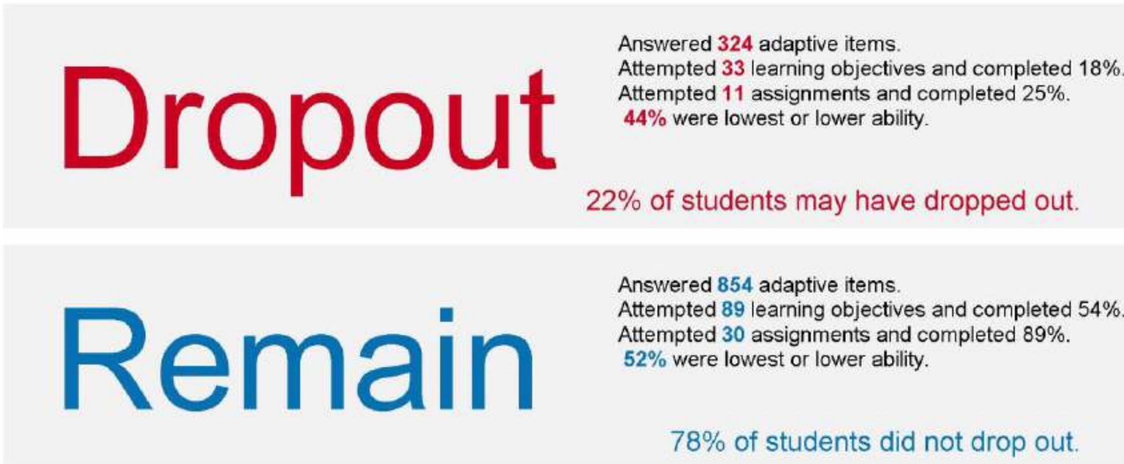

NOTE – By definition, dropout is synonymous with not participating in any of the last 25% of assignments in a course.

Conclusion

This study explored relationships between usage of the Knewton online platform, completion of learning objectives and assignments, and student performance on assessments. Students who engaged with more content on the Knewton online platform outperformed peers on online tests and quizzes, compared with peers of the same ability who used the platform to a lesser extent. Increased usage of Knewton was also associated with higher rates of assignment completion, and assignment completion was positively associated with higher average scores on tests and quizzes. Assignment completion earlier in the course predicted subsequent assignment completion, as well as whether or not the student remained engaged in the work throughout the duration of the course.

Students of all ability levels were able to successfully complete assignments, and students of all ability levels had similar rates of assignment completion. One potential explanation of this finding is that Knewton's adaptive platform allows students of all ability levels to complete assignments by providing low-ability students more items to master the content, as needed. In addition, assignment completion was more strongly correlated with usage of the Knewton platform than student ability, while student performance on online assessments was more strongly correlated with student ability than usage. These findings suggest that students of all ability levels were able to successfully complete assignments and that while increased usage of Knewton yielded higher average quiz/test scores, performance on online assessments was explained more by student ability than by usage of the Knewton platform.

Across all measures, Knewton appeared to influence outcomes similarly for students of different ability levels. The only exception was that higher-ability students completed assignments faster than lower-ability students. Given the adaptive nature of the Knewton online platform, this finding is expected. Results also were consistent across different course subjects (e.g., chemistry, economics, mathematics, and statistics), but statistics courses used the Knewton online platform the least.

Overall, Knewton appears to be a useful tool for students. This study suggests a positive correlation among usage of Knewton, assignment completion, and performance on online assessments. Given the limitations noted regarding data availability and potential confounds for the present analyses, additional study is recommended, particularly to adequately account for student ability. Capturing students' viewpoints of Knewton also may be insightful. Future work, for example, could solicit feedback from students, including what students liked and did not like about Knewton, ease of use and suggestions for improvement, and to what extent students believed that Knewton improved their learning.

Technical Appendix

This technical appendix contains the regression estimates from the estimated hierarchical and multilevel mixed-effects logistic models. Regression estimates for the dummy variables indicating the institution were not included for simplicity. The first set of tables presents the regression coefficients for the full analytic student sample. The second set of tables presents the regression coefficients for the subgroup analyses.

Full Sample Results

Table 1

Model estimates for assignment and learning objective completion predicting average quiz/test score

Outcome variable	Predictor variables	Estimate		SE	P-value
Average quiz/test score	Proportion of assignments completed (of those attempted) (gmc)	0.12	***	0.02	0.000
	Struggle (gmc)	-0.67	***	0.15	0.000
	Struggle^2 (gmc)	1.22	*	0.54	0.024
	First two (gmc)	0.25	***	0.03	0.000
	Intercept	0.70	***	0.01	0.000
	Student N	2036			
	Class N	70			
Average quiz/test score	Proportion of assignments completed (of those offered) (gmc)	0.14	***	0.01	0.000
	Struggle (gmc)	-0.69	***	0.14	0.000
	Struggle^2 (gmc)	1.29	*	0.52	0.014
	First two (gmc)	0.23	***	0.03	0.000
	Intercept	0.70	***	0.01	0.000
	Student N	2036			
	Class N	70			
Average quiz/test score related to learning objective^{log}	Learning objective completed	0.44	***	0.02	0.000
	Struggle (gmc)	-1.20	***	0.25	0.000
	First two (gmc)	2.40	***	0.13	0.000
	Intercept	1.28	***	0.04	0.000
	Learning Objective N	83322			
	Class N	63			

NOTES—1) (gmc) indicates that the predictor variable was grand-mean centered; 2) SE=standard error of the estimate; 3) ***p<.001, **p<.01, *p<.05; and 4) ^{log} indicates that a logarithmic transformation was needed, and the estimate and standard error of the estimate are in terms of log-odds.

Table 2

Model estimates for usage predicting average quiz/test score

Outcome variable	**Predictor variables**	**Estimate**		**SE**	**P-value**
Average quiz/test score	Number of adaptive items (gmc)	0.00	***	0.00	0.000
	Number of adaptive items^2 (gmc)	0.00	***	0.00	0.000
	Struggle (gmc)	-0.90	***	0.16	0.000
	Struggle^2 (gmc)	1.73	**	0.57	0.003
	First two (gmc)	0.50	***	0.14	0.000
	First two^2 (gmc)	-0.17		0.12	0.152
	Intercept	0.70	***	0.01	0.000
	Student N	2036			
	Class N	70			
Average quiz/test score	Number of learning objectives attempted (gmc)	0.00	***	0.00	0.000
	Number of learning objectives attempted^2 (gmc)	0.00	*	0.00	0.035
	Struggle (gmc)	-0.80	***	0.15	0.000
	Struggle^2 (gmc)	1.41	**	0.53	0.008
	First two (gmc)	0.26	***	0.03	0.000
	Intercept	0.70	***	0.01	0.000
	Student N	2036			
	Class N	70			
Average quiz/test score	Number of assignments attempted (gmc)	0.01	***	0.00	0.000
	Number of assignments attempted^2 (gmc)	0.00	*	0.00	0.048
	Struggle (gmc)	-0.75	***	0.15	0.000
	Struggle^2 (gmc)	1.33	*	0.53	0.012
	First two (gmc)	0.25	***	0.03	0.000
	Intercept	0.70	***	0.01	0.000
	Student N	2036			
	Class N	70			

NOTES—1) (gmc) indicates that the predictor variable was grand-mean centered; 2) SE=standard error of the estimate; 3) ***p<.001, **p<.01, *p<.05; and 4) the regression coefficients of adaptive items, learning objectives, and assignments are non-zero but very small in magnitude.

Table 3

Model estimates for usage predicting assignment completion (of those offered)

Outcome variable	**Predictor variables**	**Estimate**		**SE**	**P-value**
Proportion of assignments completed	Number of adaptive items (gmc)	0.00	***	0.00	0.000
	Number of adaptive items^2 (gmc)	0.00	***	0.00	0.000
	Struggle (gmc)	-1.51	***	0.10	0.000
	Struggle^2 (gmc)	1.74	***	0.29	0.000
	First two (gmc)	1.51	***	0.11	0.000
	First two^2 (gmc)	-0.68	***	0.09	0.000
	Intercept	0.73	***	0.01	0.000
	Student N	6278			
	Class N	182			
Proportion of assignments completed	Number of learning objectives attempted (gmc)	0.01	***	0.00	0.000
	Number of learning objectives attempted^2 (gmc)	0.00	***	0.00	0.000
	Struggle (gmc)	-0.33	***	0.03	0.000
	First two (gmc)	0.08		0.08	0.278
	First two^2 (gmc)	0.11		0.06	0.070
	Intercept	0.73	***	0.01	0.000
	Student N	6278			
	Class N	182			
Proportion of learning objectives completed	Number of adaptive items (gmc)	0.00	***	0.00	0.000
	Number of adaptive items^2 (gmc)	0.00	***	0.00	0.000
	Struggle (gmc)	-0.95	***	0.04	0.000
	First two (gmc)	0.59	***	0.02	0.000
	Intercept	0.49	***	0.01	0.000
	Student N	6278			
	Class N	182			

NOTES—1) (gmc) indicates that the predictor variable was grand-mean centered; 2) SE=standard error of the estimate; 3) ***p<.001, **p<.01, *p<.05; and 4) the regression coefficients of adaptive items and learning objectives are non-zero but very small in magnitude.

Table 4

Model estimates for prior assignment completion ((of those attempted) predicting subsequent assignment completion

Outcome variable	**Predictor variables**	**Estimate**		**SE**	**P-value**
Proportion of second 25% of assignments completed	Proportion of first 25% of assignments completed (gmc)	0.72	***	0.02	0.000
	Struggle (gmc)	0.66	***	0.16	0.000
	Struggle^2 (gmc)	-1.36	*	0.61	0.025
	First two (gmc)	1.59	***	0.18	0.000
	First two^2 (gmc)	-0.95	***	0.15	0.000
	Intercept	0.85	***	0.00	0.000
	Student N	5743			
	Class N	172			
Proportion of third 25% of assignments completed	Proportion of second 25% of assignments completed (gmc)	0.63	***	0.01	0.000
	Struggle (gmc)	0.44	**	0.14	0.001
	Struggle^2 (gmc)	-0.90		0.47	0.054
	First two (gmc)	0.38	***	0.04	0.000
	Intercept	0.81	***	0.00	0.000
	Student N	5309			
	Class N	179			
Proportion of fourth 25% of assignments completed	Proportion of third 25% of assignments completed (gmc)	1.08	***	0.01	0.000
	Struggle (gmc)	0.01		0.05	0.809
	First two (gmc)	-0.04		0.02	0.124
	Intercept	0.80	***	0.00	0.000
	Student N	4909			
	Class N	182			

NOTES—1) (gmc) indicates that the predictor variable was grand-mean centered; 2) SE=standard error of the estimate; and 3) ***p<.001, **p<.01, *p<.05.

Table 5

Model estimates for assignment completion (of those attempted) predicting course dropout

Outcome variable	**Predictor variables**	**Estimate**		**SE**	**P-value**
$Dropout^{log}$	Proportion of first 25% of assignments completed (gmc)	-3.04	***	0.21	0.000
	Struggle (gmc)	-15.89	***	2.12	0.000
	Struggle^2 (gmc)	36.20	***	8.00	0.000
	First two (gmc)	-14.44	***	2.15	0.000
	First two^2 (gmc)	9.87	***	1.81	0.000
	Intercept	-1.68	***	0.09	0.000
	Student N	6005			
	Class N	170			
$Dropout^{log}$	Proportion of second 25% of assignments completed (gmc)	-3.08	***	0.15	0.000
	Struggle (gmc)	-11.41	***	2.22	0.000
	Struggle^2 (gmc)	20.12	*	7.92	0.011
	First two (gmc)	-10.15	***	2.43	0.000
	First two^2 (gmc)	7.26	***	2.03	0.000
	Intercept	-1.95	***	0.10	0.000
	Student N	5746			
	Class N	175			
$Dropout^{log}$	Proportion of third 25% of assignments completed (gmc)	-2.97	***	0.17	0.000
	Struggle (gmc)	-9.14	***	2.30	0.000
	Struggle^2 (gmc)	13.15		6.83	0.054
	First two (gmc)	-5.78	*	2.74	0.035
	First two^2 (gmc)	4.27		2.38	0.072
	Intercept	-2.72	***	0.10	0.000
	Student N	5115			
	Class N	168			

NOTES—1) (gmc) indicates that the predictor variable was grand-mean centered; 2) SE=standard error of the estimate; 3) ***p<.001, **p<.01, *p<.05; and 4) ^{log} indicates that a logarithmic transformation was needed, and the estimate and standard error of the estimate are in terms of log-odds.

Subgroup Results

Table 6

Model estimates for usage predicting assignment completion (for those offered in course) by student ability

Outcome variable	Predictor variables	Estimate		SE	P-value
Proportion of assignments completed (highest ability)	Number of adaptive items	0.00	***	0.00	0.000
	Number of adaptive items^2	0.00	***	0.00	0.000
	Intercept	-0.02		0.07	0.791
	Student N	1567			
	Class N	176			
Proportion of assignments completed (higher ability)	Number of adaptive items	0.00	***	0.00	0.000
	Number of adaptive items^2	0.00	***	0.00	0.000
	Intercept	0.13		0.08	0.096
	Student N	1572			
	Class N	172			
Proportion of assignments completed (lower ability)	Number of adaptive items (gmc)	0.00	***	0.00	0.000
	Number of adaptive items^2 (gmc)	0.00	***	0.00	0.000
	Intercept	-0.04		0.09	0.677
	Student N	1570			
	Class N	172			
Proportion of assignments completed (lower ability)	Number of adaptive items	0.00	***	0.00	0.000
	Number of adaptive items^2	0.00	***	0.00	0.000
	Intercept	0.01		0.06	0.924
	Student N	1569			
	Class N	179			

NOTES—1) SE=standard error of the estimate; 2) ***p<.001, **p<.01, *p<.05; and 3) the regression coefficients of adaptive items are non-zero but very small in magnitude.

Table 7

Model estimates for usage predicting average test/quiz score by student ability

Outcome variable	Predictor variables	Estimate		SE	P-value
Average quiz/test score (highest ability)	Number of adaptive items	0.00	***	0.00	0.000
	Number of adaptive items^2	0.00	*	0.00	0.015
	Intercept	0.76	***	0.02	0.000
	Student N	526.00			
	Class N	62.00			
Average quiz/test score (higher ability)	Number of adaptive items	0.00	**	0.00	0.002
	Number of adaptive items^2	0.00	*	0.00	0.024
	Intercept	0.69	***	0.05	0.000
	Student N	494.00			
	Class N	65.00			
Average quiz/test score (lower ability)	Number of adaptive items	0.00	*	0.00	0.014
	Number of adaptive items^2	0.00	*	0.00	0.030
	Intercept	0.71	***	0.03	0.000
	Student N	501.00			
	Class N	65.00			
Average quiz/test score (lowest ability)	Number of adaptive items^2	0.00	**	0.00	0.003
	Intercept	0.00	**	0.00	0.001
	Student N	0.68	***	0.04	0.000
	Class N	515.00			

NOTES—1) SE=standard error of the estimate; 2) ***p<.001, **p<.01, *p<.05; and 3) the regression coefficients of adaptive items are non-zero but very small in magnitude.