

Exploring the Impact of the Default Option on Student Engagement and Performance in a Statistics MOOC

Emma Brunskill
Stanford University
Stanford, USA
ebrun@cs.stanford.edu

Dawn Zimmaro
Stanford University
Stanford, USA
dzimmaro@stanford.edu

Candace Thille
Stanford University
Stanford, USA
cthille@stanford.edu

ABSTRACT

Engagement and motivation are particularly important in optional learning environments, like educational games and massive open online courses. Providing some aspects of autonomy and choice to the student can yield significant benefits to learner motivation and persistence; yet there is also evidence that unsupported learners may not always automatically choose to allocate their learning time to pedagogical activities that are most known to be associated with better learning outcomes. We investigated the impact of choice on student engagement and learning in a Massive Open Online Course (MOOC) on introductory statistics and probability. We compared conditions in which students are given free choice over the practice problems completed to conditions in which students receive a full set of practice activities or no practice activities before completing a post-test. In all cases students were free to navigate to other sections of the course at any time. In one of the two topic sections that included personalized practice activities we found that students performed better in the condition in which they were prompted to complete all practice activities. Though more students in this condition dropped out before reaching the post-test, many more students completed the full set of practice activities in this section than those who did in the free choice condition. These results are still quite preliminary but suggest that providing a default encouraged opt in procedure can encourage students to do more problems than they would otherwise, and that doing such additional problems can yield learning gains.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous; See <http://acm.org/about/class/1998/> for the full list of ACM classifiers. This section is required.

Author Keywords

massive open online course; autonomy; motivation; student performance; choice

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

L@S 2018, June 26–28, 2018, London, United Kingdom

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-5886-6/18/06... 15.00

DOI: <https://doi.org/10.1145/3231644.3231692>

INTRODUCTION

Engagement and motivation are a key part of learning. They are particularly important in optional learning environments, like educational games and massive open online courses, that dominate the learning at scale space. Unlike in classroom settings which enforce some amount of social and authority pressure to participate, optional learning environments are mostly reliant on the individual to persist. Providing some aspects of autonomy and choice to the student can yield significant benefits to learner motivation and persistence (e.g. [1, 7]). However, simultaneously there is evidence that unsupported learners may not always choose the learning activities that are associated with the best learning outcomes: for example, Koedinger et al. [3, 4] demonstrated that students who did activities instead of watching and rewatching lecture videos had significantly better learning outcomes, though many students choose to watch videos. There are many reasons that learners may not choose to do learning activities that will yield the maximum effectiveness, such as: learners may lack information about what activities will best help their learning; learners may have different objectives than maximizing their learning outcomes; and learners may find some types of activities significantly more difficult or involve a higher cognitive load than other activities.

An alternative is to require students to perform the type and number of activities that the instructor (or an instructional policy derived using machine learning) believes to be most effective. This practice is common in many intelligent tutoring systems used in more structured settings, like K-12 education. However, in doing so it can be easy to lose the aspects of autonomy and choice that may help the learner to take ownership over his/her learning, and be deeply motivated and engaged for learning. Some research has looked at the potential influence of Open Learning Models on student performance and motivation (e.g. [6, 5]) but such research has particularly focused on whether providing students information about their stage of learning can help influence their learning choices, typically during classroom K-12 educational software studies. Other work has considered creating personalized nudges for students watching videos to engage in behavior that typically results in higher learning[2].

To explore these issues further in the context of optional learning environments, we investigated the impact of choice on students learning in a Massive Open Online Course (MOOC) on introductory statistics and probability. We are interested in comparing student engagement, as measured by the num-

ber of practice problems completed and whether the student completed post-test problems, and student test performance, between conditions in which students are required to perform a fixed number of problems before advancing to a post-test, and those where they are not. Unsurprisingly, we found that in one of the two topic sections included in the study students required to complete all practice problems before advancing to a post-test left the section before completing all problems and reaching the post test at significantly higher rates than in the free choice condition, where students could elect how many problems to do before completing a post test. Students who did complete all the required practice problems performed significantly better on the post-test than those who were allowed to choose how many problems to do. Interestingly, the proportion of students who did all the problems in the required condition was significantly higher than those who choose to do all problems in the free condition. Note that no choice was actually binding, in the sense that students were always free to navigate to another course section: it was only to advance to the post test (which was not required) that students had to complete all the problems in one condition. This suggests that nudges to complete all problems may be effective in encouraging students to complete all problems, and that doing so can have a significant impact on student outcomes. However, many open issues remain about how best to combine instruction-mediated pedagogical activity selection with autonomous choice.

METHODS

The Open Learning Initiative (OLI) Probability and Statistics Open and Free course is a free introductory probability and statistics course offered through Stanford Online. The course is designed to be equivalent to approximately one semester of an introductory university statistics class. The course is self-paced and no certificates of completion or other accreditations are provided as part of the course, which likely impacts the population of learners that pursue and complete it. The course itself offers a number of interactive elements and has been developed in conjunction with learning science experts. All learners who register for a Stanford Online course agree to participate in research studies.

In this work we adapted Stanford's instance of the Open edX platform to enable the platform to adaptively select problems for the learner: this adaptation could be based on fixed experimental conditions, features of the learner, and/or features of the learner's interactions with the platform so far. We introduced two additional sections to the course to explore this adaptivity: "Producing Data: Sampling Personalized Practice" and "Producing Data: Designing Studies Personalized Practice."

In this paper we focus on three of the four conditions in our experimental setting: the fourth condition had some implementation limitations so we exclude it here from additional analysis. The setting is identical for both topic sections, and so for simplicity we simply refer to the situation in a section. In all three conditions, a post-test was presented at the end of the section. The material in all parts of these new sections was optional (like all of the course) and participants were free

to navigate to other sections at any time. In one condition, which mimics the business-as-usual standard experience for this course, students were allowed free choice of what problems to complete during this section. We refer to this as the "free" condition. The user interface looks similar to the other (non modified) course sections, with one alteration. At the top of each practice page was a small note in red text that stated:

- When you feel you have done enough practice, please test your knowledge by clicking on the button below. Once you start the assessing your knowledge part, you cannot return to the practice problems in this section.

Below this note was a button labeled "Post assessment problems" that, when clicked, brings the learner to the post-test items. Below the button is the standard practice problem. Figure 1 shows an example of a practice problem interaction for the free choice condition.

In the second condition ("nudged all") the students were required to complete all practice problems in a section before being advancing to the post-test problems. Like in the free choice condition, there was a small note in red text at the top of a practice problem user interface, which now stated:

- Note: the order and selection of these questions are personalized for you. Please do not click on the horizontal navigation bar as it will simply redirect you back to the current problem. Use the "Next" button to go to the next problem (it will become available after you have selected the correct answer). At the end of the section you will get some problems to assess your knowledge. When you feel you have done enough practice, please test your knowledge by clicking on the button below. Once you start the assessing your knowledge part, you cannot return to the practice problems in this section.

In this condition students were presented with practice problems in a random order until all practice problems in this section had been presented. If a student attempted to navigate to another problem in the same section, the system automatically redirected the student to go back to the problem that had been selected for the student. No restrictions were made on navigating to other items outside of the current section, and the student was free to leave the section at any point. If the student completed the full set of all practice problems in this section, as presented automatically to him or her, the student was then automatically advanced to the post-test. At this point the student was shown a small number of post-test items. Figure 2 shows a sample screen shot of a post-test problem. Students were informed that if they tried to navigate back to prior problems in this section, they would be redirected back to the currently selected post-test item. Students also were not able to go back to post-test items once completed. The post-test interactions were identical for all three conditions presented in this paper. Students were free to leave the post-test problems to navigate to other sections in the course at any time. Note that students in this condition only complete post-test items if they had completed all practice items in this section.

Figure 1: Screenshot of sample practice problem interaction for the free choice condition.

In the third condition (no practice), the students received no practice problems and were immediately advanced to the post-test problems. This condition was included to help assess if additional practice leads to benefits over the practice provided in other sections. Students in the "free" and nudged all conditions were given one attempt at the practice problems and received immediate feedback on the correctness of their answer choice with a corresponding explanation. Students in all conditions were given one attempt at the post-test problems and received immediate feedback after submitting their answer.

Figure 2: Screenshot of sample post test item

PRELIMINARY RESULTS

Here we focus on analyzing how the conditions impacted student learning, as assessed by their performance on our post-

test items, and engagement, as measured by the proportion of students who persisted to complete at least one post-test item, and the number of problems completed by students in the section.

We first analyzed results for the “Producing Data: Sampling Personalized Practice” section. Somewhat surprisingly, we found almost no difference in average student performance in the post-test among the three conditions, even though in one condition students completed no practice problems. We then computed the correlation coefficient between the number of practice problems completed by a student, and the student’s performance on the post-test (as measured by post-test percent correct done across all items, or the number of items done correctly), across all conditions. In both cases the correlation was extremely low ($< |0.1|$). While we will investigate this issue further, two possible explanations are that students were already at ceiling performance before completing this section (e.g. the students had already mastered this material due to prior material provided to this section), or that the post-test items were not well aligned with the practice problems, which would mean additional practice problems do not improve post-test performance. Given that average student performance was less than 75% in all conditions, it seems unlikely that ceiling performance is the cause of this observation, and we will investigate the alignment further. Another possibility is that the amount of practice possible (14 problems in this section) was insufficient for students and students still require more

practice to master this section, and without that significant differences among the conditions are getting lost.

We next analyzed results for the second section, “Producing Data: Designing Studies Personalized Practice.” This section has a larger pool of practice problems (23 compared to 14 in the prior section). Note that students were randomly assigned to a condition for each section, so students could be one condition for the first section, and a different experimental condition in the second section.

We computed the mean performance for students who completed all 5 post-test items in the free condition, the do no practice condition, and the do all condition who completed all practice problems: 2.48 (no practice, $N=82$), 2.75 (free choice, $N=106$), 3.08 (nudged all, $N=84$). A t-test yielded significant differences between the nudged all condition and the free choice condition ($p < 0.05$). The estimated effect size of the improvement of the nudged all condition over the free choice is $d = 0.35$, which is between a small and medium effect size. This suggests that doing all the practice problems may be associated with stronger learning as assessed on the post-test.

Of course, a natural hypothesis is that the nudged condition only allows students to complete the post-test items if they complete all 23 practice problems, and many students may drop out before this point. Indeed, 77.9% of students in the nudged all condition completed at least one post-test item, compared to 94.3% of students in the free choice condition, and this difference was statistically significant. Therefore with the above analysis, we cannot infer that the condition in which students were required to do all practice problems before completing the post-test caused higher average test scores than the other two conditions. Going forward we are interested to leverage techniques from casual inference to estimate the treatment effect of the condition of nudging students to do all problems on those students’ resulting post-test performance.

However, though the nudged condition did result in more students dropping out before reaching the post-test, students did many more problems in the nudged condition. Students in the nudged all condition completed an average of 19.4 practice problems, compared to 7.5 problems in the free choice condition. In addition, the proportion of students in the nudged all condition who completed all the practice problems was 72.7% compared to only 9.8% in the free choice condition.

DISCUSSION AND NEXT STEPS

The above analysis is still quite preliminary, and many open questions remain.

In particular, here strong nudges to complete a set of problems resulted in those students completing more practice problems than when given free choice (and no nudges) of how many problems to complete, perhaps by resetting a norm of the amount of practice the students expected was required. However, students in this strong nudge condition also did not complete the within experiment post test as frequently. Note that here the post test occurred in the same section as the practice problems, and it was not part of the general course structure evaluation or assessment (more precisely, it formed no part

of the student’s resulting grade). We also do not currently know if such students merely exited the experimental section before completing the post test, or whether being part of this particular experimental condition had follow on impacts on the rest of the student’s course trajectory, including whether it was associated with such students dropping out of the course earlier. These potential impacts are very important: if strong nudges can be used to increase the amount of practice problems completed in one subtopic but lead to higher attrition or other broader negative effects, that should impact when and if such strong nudges should be employed.

While such issues are crucial to investigate, it is interesting that these early results suggest that, even outside a strongly regularized learning environment like K-12 classrooms, nudges to complete problems may result in significant behavior changes by learners. In particular, it is intriguing that the nudge to complete practice problems resulted in both more practice problems being completed, and for those that persisted to complete all posed problems, stronger learning outcomes in the resulting within-experiment assessment.

ACKNOWLEDGEMENTS

This work was supported in part by Google, a NSF CAREER award and a Microsoft Faculty Fellowship to Brunskill.

REFERENCES

1. Diana I Cordova and Mark R Lepper. 1996. Intrinsic motivation and the process of learning: Beneficial effects of contextualization, personalization, and choice. *Journal of educational psychology* 88, 4 (1996), 715.
2. Vania Dimitrova, Antonija Mitrovic, Alicja Piotrkowicz, Lydia Lau, and Amali Weerasinghe. 2017. Using learning analytics to devise interactive personalised nudges for active video watching. In *UMAP*. ACM, 22–31.
3. Kenneth R Koedinger, Jihee Kim, Julianna Zhuxin Jia, Elizabeth A McLaughlin, and Norman L Bier. 2015. Learning is not a spectator sport: Doing is better than watching for learning from a MOOC. In *L@S*. ACM, 111–120.
4. Kenneth R Koedinger, Elizabeth A McLaughlin, Julianna Zhuxin Jia, and Norman L Bier. 2016. Is the doer effect a causal relationship?: how can we tell and why it’s important. In *LAK*. ACM, 388–397.
5. Yanjin Long and Vincent Alevan. 2017. Enhancing learning outcomes through self-regulated learning support with an Open Learner Model. *User Modeling and User-Adapted Interaction* 27, 1 (2017), 55–88.
6. Yanjin Long, Zachary Aman, and Vincent Alevan. 2015. Motivational design in an intelligent tutoring system that helps students make good task selection decisions. In *AIED*. Springer, 226–236.
7. Maarten Vansteenkiste, Joke Simons, Willy Lens, Kennon M Sheldon, and Edward L Deci. 2004. Motivating learning, performance, and persistence: The synergistic effects of intrinsic goal contents and autonomy-supportive contexts. *Journal of personality and social psychology* 87, 2 (2004), 246.