Making the Grade: How Learner Engagement Changes After Passing a Course

David Lang
Stanford University
dnlang@stanford.edu

Alex Kindel
Princeton University
akindel@princeton.edu

Ben Domingue
Stanford University
bdomingu@stanford.edu

Andreas Paepcke
Stanford University
paepcke@stanford.edu

ABSTRACT
Understanding how individuals interact with a course after receiving a passing grade could have important implications for course design. If individuals become disengaged after passing a class, then this may raise questions about optimal ordering of content, course difficulty, and grade transparency. Using a person-fixed effects model, we analyze how individuals who obtained passing grades subsequently behaved within a course. These learners were less likely to complete videos and more likely to watch videos faster after receiving notice of a passing grade in the class. These learners were also less likely to reattempt items they initially got wrong.

Keywords
Video-interactions; grading schemes; learning analytics; MOOCs;

1. INTRODUCTION
Grades are a key component of online courses. However, there is a great deal of heterogeneity in the downstream effects of grading and grading schemes. For instance, female students who received an ‘A’ in their introductory economics courses were substantially more likely to major in the subject than individuals who received a B but had similar scores in the class [1]. On the other end of the spectrum, research suggests that pass-fail grading schemes may be beneficial in terms of student stress in high-stakes environments [2]. Other work suggests that the presence of pass-fail grading discourages student performance[3].

MOOCs offer a unique opportunity to understand how grading affects within-course behavior. First, clickstream data can document subtle changes in behavior that are reasonable proxies for engagement and effort (e.g. video consumption, video interactions, multiple attempts on items). Second, compared to traditional courses, grading in MOOCs is much more salient and immediate. Grades are recomputed instantaneously, and solutions are presented after every single problem.

Understanding how individuals interact with a course after receiving a passing grade could have important implications for course design. If individuals disengage after passing a class, then it may make sense to structure a course such that final grades are not revealed until all problems have been attempted. Alternatively, if individuals exert more effort in a class after reaching passing status, then perhaps courses should be designed with gamification/scaffolding in mind such that a learner is continually working for a new certificate/badge.

2. DATA
The dataset used in this analysis was an introductory course in Statistical Learning administered multiple times via Stanford’s Lagunita Platform. 55,000 individuals enrolled in the class. Of that population, 11,301 individuals interacted with both course videos and with assignments related to the course at least once. Of these individuals, 2,485 achieved certification.

The course includes 77 videos. The cumulative length of these videos is 15.3 hours. We used the clickstream created by learners who viewed the course via the Lagunita platform. Clickstream events are generated each time a video is loaded, finished, played or paused, fast forwarded or rewound. Other clickstream activities include changes to the media player’s playback speed to one of six settings (0.5X, .75X, 1.0X, 1.25X,1.5X, and 2.0X). These activities were aggregated on a user-video level. In total, there were 126,799 learner-video observations.

2.1 Course Items

The course assignments consisted of 103 multiple-choice, short-responses, and fill-in-the-blank items. Learners who answered at least 50% of all items correctly received a certificate. Individuals who obtained a score of 90% or more received a certificate of distinction. We limited the dataset to include only individuals who attempted at least a simple majority of items.

3. ANALYSIS
In this course, learners are keenly aware of the grading cutpoints. The distribution of learners’ scores show substantial jumps in density at just above 50% and just above 90% (red lines), as seen in Figure 1. In an educational context, such jumps usually indicate a bias on the part of graders to give students with marginal scores the benefit of the doubt [4]. In this instance, though, all exams are graded electronically, and this type of manipulation by a grader is not possible. Instead, this heaping likely reflects a subset of learners who are extremely motivated by the certificate, and cease attempts after obtaining it. In this case, we identified that approximately 5% of students stopped attempting items shortly after they hit the 50% threshold. Formal evaluation via the McCrary Density[1] tests rejects continuity of the density function

[1] The McCrary Density Test estimates the continuity of exam scores at the cutoff using local linear regression. If the left and

at the cutoff scores with a t-statistic of 7.1 and 5.1 at the 50% percent and 90% percent thresholds [5].

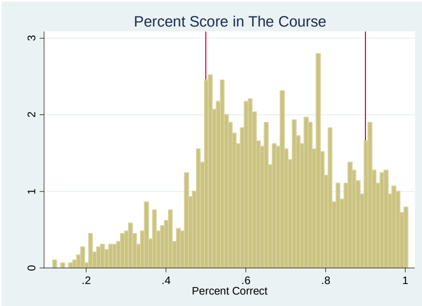

Figure 1 Histogram of Course Scores

Given how pronounced and precise this heaping was, we examined the grade-reporting interface. If a user clicks on a link to their progress, a report is generated with a user's score on each exam, as well as their overall status with the course, indicating whether they have currently passed the course. Figure 2 depicts a mock-up of these reports. There are several noteworthy features of this reporting format. First, these grading thresholds are very clearly identified by their shading. Light grey depicts the region that is considered passing (>50%) and dark grey depicts the region that is considered passing with distinction (>90%). A learner's grade is communicated by their total score bar (right most column). If this bar is at 50% or more, they will be able to observe the top of the total score bar in the light grey region, indicating that they passed. If the total score bar is in the dark grey region, this indicates the learner has earned a certificate of distinction. On top of these features, the total score is computed and displayed in percentages terms, making the learner's grade relative to the passing threshold eminently clear. In this artificial example, the learner obtained a 100% on every item but stopped almost immediately after obtaining a passing grade in the course.

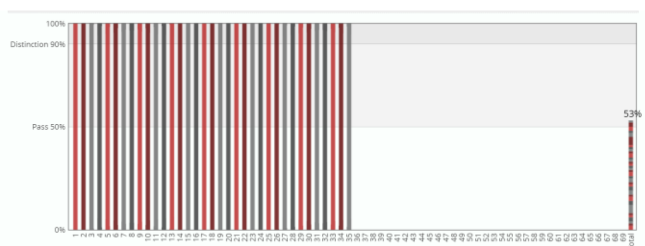

Figure 2 Example Grading Report

This reporting format could help explain the popularity of grade checking behavior in the course. Ninety-eight percent of learners checked their grades at least once, and the median user checked their grade 32 times during the course of the class.

3.1 When Passing Occurs

There is considerable variation in when a learner passes a course. Our identification strategy leverages within-learner variation before and after they became aware they passed the course. Figure 3 shows that of the 2,485 learners who passed the course, the median individual tends to do so within the first 70 items. This leaves almost a third of the course and its items to serve as a behavioral contrast. We also exploit variation of when students become aware they have passed the course. Approximately 70% of individuals, checked their grade on the day that they passed a course. Others realized this information at a later date.

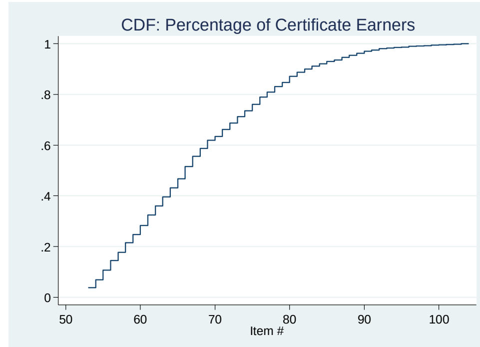

Figure 3 Item on which an Individual Obtains a Certificate

3.2 Impact of Passing on Engagement

We estimate user engagement by analyzing video interactions before and after a learner receives notification that they have passed the course via person fixed-effect regression. The specification is below:

$$UserBehavior_{ij} = B_1 PassNotification_{ij} + \Gamma_i + e_{ij}$$

The Γ_i denotes the person-fixed-effect and the user behavior/pass notification refers to the ith person's performance on their jth video. Outcomes include playback speed, fast forwarding, and video completion. For the purposes of this analysis, we define video completion as a student completing 90% of a video. This threshold was chosen as these videos often contain summaries, production details, and end titles in the last minute or so of content.

Our first analysis suggests that individuals sped up after passing a course. The first column of Table 1 corresponds to a univariate regression model of playback speed on pass notification. The second column corresponds to a regression model of playback speed on pass notification and a person-fixed-effect. The third column also includes a time trend that accounts for how many days a student has been enrolled in a course at the time of their video interactions. After accounting for person-fixed effects, our preferred regression model (Column 2) finds individuals speed up on average about 1%. Given that playback speed has six discrete speeds (0.5X, 0.75X, 1.0X, 1.25X, 1.5X, 2.0X) this speed-up reflects a subset of learners adjusting their playback speed on a subset of videos that they interacted with rather than a gradual shift across all videos. Depending on how early a learner obtained a passing grade for the course, this speedup represents as much as a 10-minute reduction on time spent watching videos over the remainder of the course. In terms of effect size, this increase corresponds to roughly an increase of .05 of a standard deviation.

right hand-side estimates produce substantially different estimates, it would suggest manipulation or selection into one of the two groups.

Table 1 Effect of Pass Notification on Playback Speed

	(1) Univariate	(2) Person Effects	(3) Time Trend
Pass Notice	0.0184***	0.0107***	0.00607**
	(3.68)	(5.17)	(3.04)
Log Days			0.00412***
			(5.29)
Constant	1.080***	1.082***	1.070***
	(298.23)	(2318.28)	(438.87)
Observations	126799	126799	126799
Adjusted R^2	0.002	0.776	0.776

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Other video behaviors suggest that individuals may be less engaged in the course after receiving certification. Modeling the effect of receiving a passing grade on fast forwarding behavior suggests that passing notification is associated with a 4-5% percentage point reduction in fast forwarding and a 3-4% percentage point reduction in rewinding. A decrease in fast forwarding behavior may be seen as a form of increased engagement by some. However, it should be noted that fast forwarding and rewinding are symmetric actions (The concordance within video between rewinding and fast forwarding is 73%).

When answering a question on an assignment, a very common learner strategy is to review prior material. If a user is searching a video for a particular statement or graph, a learner is unlikely to skip to exactly the right point in time. Even if they were, learners may like to check the immediately preceding and following slides for context or clarifying information. In these cases, one would expect to see both fast forwarding and rewinding. Most of the reduction in rewinding and fast forwarding seems to come from cases like these. In terms of total effect size, these reductions correspond to a .10 reduction in fast forwarding and a .06 reduction in rewinding.

Table 2 Effect of Pass Notification on Fast Forwarding (Top) and Rewinding (Bottom)

	(1) Univariate	(2) Person Effects	(3) Time Trend
Pass Notice	-0.0430***	-0.0419***	-0.0496***
	(-7.59)	(-11.20)	(-12.03)
Log Days			0.00682***
			(4.07)
Constant	0.259***	0.259***	0.239***
	(64.27)	(307.34)	(48.20)
Observations	126799	126799	126799
Adjusted R^2	0.002	0.180	0.181

	(1) Univariate	(2) Person Effects	(3) Time Trend
Pass Notice	-0.0258***	-0.0303***	-0.0418***
	(-4.08)	(-7.05)	(-9.10)
Log Days			0.0102***
			(5.73)
Constant	0.393***	0.394***	0.364***
	(78.88)	(407.29)	(68.39)
Observations	126799	126799	126799
Adjusted R^2	0.000	0.200	0.201

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We also note the percentage of videos that are completed decreases after pass notification. Here we find that individuals are less likely to complete videos after passing a course by approximately five percentage points. This corresponds to approximately .15 of a standard deviation.

Table 3 Effect of Pass Notification on Video Completion

	(1) Univariate	(2) Person Effects	(3) Time Trend
Pass Notice	-0.0226***	-0.0498***	-0.0425***
	(-5.12)	(-14.69)	(-11.99)
Log Days			-0.00646***
			(-4.45)
Constant	0.858***	0.864***	0.882***
	(303.85)	(1130.74)	(200.95)
Observations	126799	126799	126799
Adjusted R^2	0.001	0.182	0.182

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Finally, we examine the number of attempts individuals made to answer items. We find that individuals who passed the course were subsequently less likely to make multiple attempts on incorrect items.[2] Before passing, there were an average of 1.11 attempts. After passing, this declined to 1.07 attempts. This corresponds to an effect size of approximately of .07.

Table 4 Effect of Pass Notification on Item Attempts

	(1) Univariate	(2) Person Effects	(3) Time Trend
Pass Notice	-0.0361***	-0.0394***	-0.0316***
	(-7.60)	(-6.98)	(-4.73)
Log Days			-6.702*
			(-2.41)
Constant	1.114***	1.114***	67.54*
	(427.81)	(1888.66)	(2.45)
Observations	113562	113562	113562
Adjusted R^2	0.001	0.203	0.203

t statistics in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

3.3 Limitations to Analysis

This study was conducted on a single MOOC. It should also be noted that this MOOC was a terminal course. This course was not part of a broader sequence and its content was not necessary for other courses available within the platform. A such, our findings that users disengaged in course material after passing the course may not generalize.

4. DISCUSSION

On balance, our findings suggest that passing notification discourages subsequent engagement for at least a subset of users. We see increases in playback speed and less video completion.

These findings are consistent with evidence from the educational psychology and behavioral economics literature, which has suggested that receipt of a certificate or badges can discourage intrinsic motivation in individuals [6][7]. Earlier work in MOOCs

[2] Observations differ in this specification because it is based on person-item level data rather than person-video level data.

also found that individuals who obtain certificates in courses actually skipped nearly a quarter of a course's video content [8]. We have documented several learner behaviors that are relevant to the design of MOOCs, and likely the design of online teaching more generally.

With respect to grading schema, there is substantial evidence that individuals act in a more engaged manner before passing a course than after they have received a notification of passing. We also see this strategic behavior in that there are substantially more students just above the passing threshold than just below it.

One policy implication of these findings is how quickly learners should be notified about their overall success in a course. Currently many courses notify learners instantaneously, daily, or on a near weekly basis when these events occur. For courses with a well-defined end date, it may make sense to not notify users of their final grades until the course is completed.

A second consideration is how transparent instructors should be in terms of grading. Learners could not manipulate their grades as easily if they did not know the exact threshold for passing. Using language that describes *approximate* cutpoints may discourage learners from conflating certification and completion while allowing for more rigorous causal inference.

Lastly there is the question of course structure, if individuals put forth less effort after passing a class, then perhaps a more traditional instructional environment of weekly assignments with a summative final project or exam may yield more total learning.

5. FUTURE STEPS

We found that notification of a passing grade decreased subsequent effort in the *same* course. An equally intriguing question is how individuals who are enrolled in multiple classes behave after this notification. If these individuals are solely interested in accumulation of credentials or certificates, presumably we would see effort shift to courses where learners have yet to obtain certificates.

ACKNOWLEDGMENTS

The research reported here was supported in part by the Institute of Education Sciences, U.S. Department of Education, through Grant (#R305B140009). We also acknowledge the support of Stanford University's CAROL (Center for Advanced Research through Online Learning), the Spencer Foundation, and CEPA (Center for Education and Policy Analysis).

6. REFERENCES

[1] A. Owen, "Grades, gender, and encouragement: A regression discontinuity analysis," *J. Econ. Educ.*, 2010.

[2] D. Rohe, P. Barrier, M. Clark, and D. Cook, "The benefits of pass-fail grading on stress, mood, and group cohesion in medical students," *Mayo Clin.*, 2006.

[3] R. Gold, A. Reilly, and R. Silberman, "Academic achievement declines under pass-fail grading," *J.*, 1971.

[4] T. Dee, B. Jacob, and J. Rockoff, "Rules and discretion in the evaluation of students and schools: The case of the New York regents examinations," *Sch. Res. Pap.*, 2011.

[5] J. McCrary, "Manipulation of the running variable in the regression discontinuity design: A density test," *J. Econom.*, 2008.

[6] M. Lepper, D. Greene, and R. Nisbett, "Undermining children's intrinsic interest with extrinsic reward: A test of the" overjustification" hypothesis.," *J. Personal.*, 1973.

[7] M. Hanus and J. Fox, "Assessing the efects of gamification in the classroom: A longitudinal study on intrinsic motivation, social comparison, satisfaction, effort, and academic," *Comput. Educ.*, 2015.

[8] P. Guo and K. Reinecke, "Demographic differences in how students navigate through MOOCs," *Proc. first ACM Conf.*, 2014.