

(Dis)organization and success in an economics MOOC

By ABHIJIT V. BANERJEE AND ESTHER DUFLO *

Massive Open Online Courses have experienced a fast cycle of hype and disappointment. In November 2012, the New York Times proclaimed 2012 to be the year of the MOOC. The excitement was based on the promise that the world can now have access to a wide variety of courses from the best universities¹. However, on December 12, 2013, a Washington Post article headline was “Are MOOCs already over?”².

The main problem that MOOCs face is that few users actually complete the class. For example, a widely publicized study by the Graduate School of Education at UPenn based on a million participants in the first seventeen UPenn courses offered on the Coursera platform finds that only 5% of those who register for a class actually complete the course (Ahmad et al., 2013). The number is somewhat misleading, since a substantial fraction of those who register are never active on the site (in the sense of viewing one lecture or completing one problem set). Across all UPenn courses, only half of those who registered viewed at least one lecture within their selected course. This might just reflect a version of the shopping around for classes as is common in brick and mortar colleges. It is very possible that in the case of MOOCs, poten-

tial students first register in every class that might be of interest, and then decide which to complete.

More tellingly though, even among those who actually start the course, a large majority eventually drop out; across the UPenn Coursera courses, between 2% and 14% of those who actually started the course showed any activity in the last week of the course. In their analysis of outcomes from MIT-based edX course “Circuits and Electronics,” Bergner et al. find that 76% of students account for only 8% of the time spent on the course while 7% of the participants accounted for 60% of the time spent on the course (Bergner et al., 2010).

This may just be a reflection of the mismatch between what people expect from a MOOC and what they discover in practice. In particular, since MOOCs are typically modeled on semester-long academic classes at top institutions, they tend to be quite demanding. The typical semester course is a 12-unit class, which strictly speaking means that it requires 12 hours of work each week, a time commitment many MOOC students may not have. Actual requirements vary widely: 5 of the UPenn courses required watching 1 to 3 videos a week for a total of less than 40 minutes, while 6 of the courses had more than 7 videos per week. Some courses had no homework, while others had 8 assignments. It is true that even the more demanding of these did not officially require 12 hours a work a week, but then those assessments of hours were based on students in the universities where the course was originally taught, who are probably better prepared. It is certainly true that the drop-out rates were much larger for courses that required more weekly hours and had more homework assignments.

If this is all that there is to it, then the high drop out rate is not really a failure of MOOCs, and there is no particular reason

* This project received no outside funding. It received IRB approval from MIT, number 1309005892. We thank Noam Angrist for more than outstanding research assistance. His energy and drive made this happen and he is the first one who spotted the discontinuity. We also thank Angela Ambroz and Ting Mao for their great work on the course, the MITx and EdX teams for support, and Madeline Duhon for excellent research assistance.

¹Pappano, Laura. “The Year of the MOOC.” *The New York Times*. November 2012. <http://www.nytimes.com/2012/11/04/education/edlife/massive-open-online-courses-are-multiplying-at-a-rapid-pace.html>.

²Strauss, Valeri. “Are MOOCs Already Over?” *The Washington Post*. December 2013. <http://www.washingtonpost.com/blogs/answer-sheet/wp/2013/12/12/are-moocs-already-over/>.

to be worried. It is efficiently high at the moment because users are learning about what MOOCs are like and the suppliers are learning what the users want. There would be a natural process of convergence to an equilibrium where students go in with a much clearer sense of what they want and expect, and therefore drop out less frequently, or where the course offerings will be shorter and/or less demanding.

But it is also plausible that the specific nature of the MOOCs leads to more dropouts and failure than would be observed in a more traditional classroom environment, not because of user selection, but because of the nature of teaching in an online environment. In that case there may well be modifications to the platform as a whole that would ameliorate this problem.

In this paper, we focus on one particular factor that plausibly places the MOOC user at a disadvantage with respect to a classroom participant: the need for self-discipline and focus without the benefit of a peer group or a structured study time. We show that students whose behavior suggests that they are not organized are significantly less likely to succeed in a MOOC, even after controlling for a number of factors that are associated with success in the course. This is entirely driven by their failure to complete assignments on time, rather than by their performance conditional on completing them. Of course, we have not performed the same study in a traditional setting, and therefore this does not actually prove that this is more of a problem in MOOCs. However, these results do suggest that there is scope for improving MOOC performance by providing more structure for students. In future research, we are hoping to pursue this idea in a series of randomized experiments.

I. 14.73x: The Challenges of Global Poverty

As of December 2013, economics and finance classes represented 11% of all past and present classes in the three largest platforms (edX, Coursera, and Udacity). Among the humanities, economics classes

are particularly amenable to the platform, since machine graded assignments (multiple choice or numerical problems) are appropriate.

In the spring of 2013, MIT launched its first economics course on the EdX platform, 14.73x, “The Challenges of Global Poverty.” It was taught by the two of us, with the support of Angela Ambroz (of J-PAL) as the course manager, and two amazing teaching assistants (Ting Mao and Noam Angrist). The class was taught simultaneously to the general public and to MIT students, in a “flipped classroom” format. The MIT students got all the lectures and homework assignments from 14.73x, and the class time was devoted to the preparation and the presentation of case studies in small groups. This paper focuses on the non-MIT students. The course had 10 sections, each with roughly 3 hours of lectures (recorded from a previous version of the class taught at MIT in a lecture format, and cut in 6-10 minute segment separated by short “finger exercise type” quizzes), accompanied by a set of slides, and (except for the intro section) a homework assignment (which was entirely machine graded). To get a certificate, students had to have a score of at least 50% on 8 of 9 homeworks plus the finger exercises.

The teaching assistants moderated a student forum for this class, from which they culled the best questions that were then sent to us for answering. They also answered questions themselves, and held weekly office hours on Skype, the transcripts of which were posted.

While for an MIT economics course, 14.73x was unusually non-technical, the requirements of watching 3 hours of video each week and doing a problem set for every unit were higher than almost all the UPenn classes mentioned above. Given that, the retention rate was reasonably high by the standard of most MOOCs. 42,314 registered for the course. 26,140 ever viewed any page (they are usually referred to as the “starters”). 12,947 were ever “active”, in that they attempted any assignment or finger exercise (even if they got a zero on the that assignment). Of those, 4,597 earned a

certificate. The retention rate is thus 11% of registrants, 18% of starters, and 35% of active participants.

As befits a course on global poverty, the enrollment was very diverse. According to self-reported demographics data collected on 4,600 students towards the middle of the course, we had just over 50% male participants, coming from 194 countries. The top few countries were the US (28%), India (10%), United Kingdom (5%), Canada (3%), Brazil (3%), Germany (3%) and Nigeria (2%). We had an approximately equal proportion of men and women. 110 languages were represented, and English was the first language for 50% of the students. The mean age was 30 years old. The students came from a relatively advantaged background: 99% had computers at home (this reflects the fact that the class is mainly taken at home, not in a college or school environment or computer labs), 87% had more than 25 books. 98% of students had a high school degree or higher, 82% had a bachelor's degree or higher, and over 75% had parents who have completed at least some college.

II. Organization and success

While there was a deadline for enrollment, in practice, enrollment remained open much after the deadline passed (as is the case on the EdX platform). Figure 1 shows that people did continue to enroll after the deadline, although 73% of the enrollment happened on or before the deadline, and there was some bunching on the deadline day itself (4%).

Figure 2 presents the main result of the paper. It plots for active users the overall grade in the class as a function of enrollment date (normalized such that 0 is the last “on time” enrollment date). The figure clearly shows a discontinuity at the enrollment deadline. Students who enrolled even one day late did significantly worse than those who enrolled by the deadline. Among active users, the grade they received does not depend on when the student registered as long as they did so before the deadline. The fact that performance is flat

until that point is reassuring in light of the bunching in registrations at the discontinuity shown in figure 1; while many people register on the last day, there appears to be nothing systematically different about them. However, there is something significantly different about those who enroll a day late. Those who enroll later get even lower grades.

Table 1 presents a formal regression discontinuity analysis of the “deadline effect,” focusing on active users in a narrow window of 15 days either side of the discontinuity. In the specification presented in this table, we control for a cubic function of the enrollment date, estimated separately on both side of the discontinuity. Column (4) presents the impact on the overall grade (as in figure 1). Columns (1) to (3) presents the results on a dummy of obtaining a certificate. Students who enrolled one day late were less likely to get a certificate (a reduction of 16.6 percentage points, compared to an average of 45% in this sample, which represents a 37% reduction). Their grades were 10.7 percentage points lower, compared to an average grade of 42.5% (a 25% reduction).

The discontinuity clearly does not reflect a causal effect of enrolling late, as enrolling one day late had no impact on timing of access to the material (which was available a few days after the enrollment deadline). Even enrolling a few days late would have left enough time to access the first unit of material on time, especially as the introductory unit was easy and did not have a homework. Therefore, the discontinuity is likely to reflect the “type” of the student, and in particular their ability to force themselves to do complete tasks on time.

We present more evidence to support the interpretation that the discontinuity probably reflects disorganization, rather than something else that is different about those students.

First, the discontinuity appears to persist after controlling for demographic variables. Second, in table 2, we show that there is no discontinuity in the demographic characteristics, suggesting that the discontinuity does not pick up some particular observable

“type” of students. The one possible exception is that the point estimate for “English language” (i.e an English speaker) is fairly large, and although insignificant here, is significant when using a larger window for enrollment.

Second, in table 3 we investigate whether the “enrolling late” variable actually simply picks up motivation rather than organization. While there is no perfect way to rule this out, we make use of the fact that we collected (self-reported) data on a number of motivation and self-efficacy variables. Students who enroll one day late are indeed significantly less likely to say that they intend to complete the course. Controlling for these variables halves the estimates, but it remains negative and fairly large (9 percentage point decline in the probability to earn a certificate).

Finally, since homework represents the most important part of the grade, columns (5) and (6) decompose homework grades between missed homework and grade conditional on attempting. The deadline for completing homework was binding, as is typical is for a MOOC. The solutions were posted and the homework disabled immediately after the deadline expired, and a student was not able to complete a late homework. They could still pass the class since the grading scheme dropped the lowest homework scores (including missed ones) and the required average grade was 50%. However, this would certainly place them at a disadvantage. Column (5) shows that compared to a comparable student who enrolled just on time, a student who enrolled one day late missed on average 1.15 homework assignments (out of 9 possible homework, and just about 4 actually completed!). However, conditional on completing the homework (although we are mindful that the sample is differentially selected on either side of the discontinuity), their grade was the same. In regressions unreported, we also find no difference on performance on the finger exercises. Thus, lateness does not seem to pick up ability: rather, it likely picks up that those students have a harder time committing to complete the work on time.

III. Future research

While this data suggests that “disorganization” is negatively correlated with performance in the MOOC, the natural next step (short of disallowing late registration, which would certainly mechanically improve retention rates, but without much social gain) is to think about ways in which MOOCs could provide more structure to help students deal with their self control issues. In the spring of 2014, we are planning to conduct a series of A/B testing experiments conjoint with the next installment of the 14.73x class. One of them is to offer some randomly selected students the option to set aside regular study times. This will be reinforced by two subtreatments. In one treatment, we will remind them that the course staff can see when they log in (and thus verify their study time). In a third treatment, we will additionally offer them the option to receive nagging reminders from the course staff. This is part of a series of experiments we intend to run on whether some of the “standard” features of brick and mortar courses can be combined with MOOCs, and whether those would make a significant difference to course success.

References

- Ahmad, Seher, Robert Boruch, Chad Evans, Laura Perna, Alan Ruby, Janie Scull, and Nicole Wang (2013). “The Life Cycle of a Million MOOC Users.” *MOOC Research Initiative Conference*.
- Bergner, Yoav, Isaac Chuang and Piotr Mitros and David E. Pritchard and Daniel T. Seaton, “Who Does What in a Massive Open Online Course?” *International Journal of Human Computer Studies*, 68 (2010): 223-241.

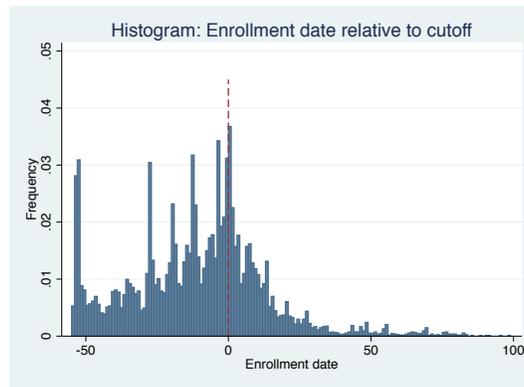


Figure 1. : Enrollment date relative to “cutoff”

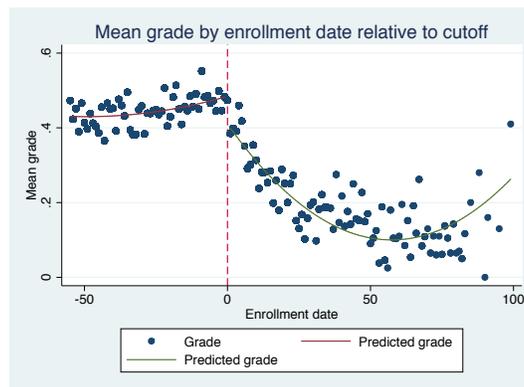


Figure 2. : Mean grade by enrollment date relative to “cutoff”

Table 1: Effect of Late Enrollment on Performance

	(1)	(2)	(3)	(4)	(5)	(6)
	Certificate	Certificate	Certificate	Grade	Number HW Missed	HW Score
Enrolled late	-0.166*** (0.0521)	-0.173*** (0.0546)	-0.0964 (0.0601)	-0.107*** (0.0377)	1.146*** (0.356)	-0.0439 (0.0331)
Observations	5,500	4,974	1,929	5,506	5,414	5,414
Mean	0.451	0.451	0.451	0.425	4.176	0.505
Demographic Controls	No	Yes	Yes	No	No	No
Motivation Controls	No	No	Yes	No	No	No

Notes:

1. Includes enrollment, enrollment squared, and enrollment cubed terms, each estimated for enrollment within 15 days on both sides of the discontinuity.
2. Demographic control: Gender, age, level of education
3. Motivation controls: Confidence in basic skills, confidence in mastering skills, think will do well, intend to complete, aspire to do development economics in future, importance of certificate for future

Table 2: Demographic Indicators

	(1)	(2)	(3)	(4)	(5)
	Gender	Age	Level of Education	Parent in Economics or Policy	Primary Language English
Enrolled late	0.0245 (0.0541)	-1.249 (1.104)	-0.0156 (0.0408)	0.000193 (0.0695)	-0.0827 (0.0863)
Observations	5,161	5,116	5,051	2,150	2,142
Mean	0.446	30.62	0.839	0.189	0.415
Controls	No	No	No	No	No

Notes:

1. Includes enrollment, enrollment squared, and enrollment cubed, each estimated for enrollment within 15 days on both sides of the discontinuity.

2.pdf

Table 3: Motivation Indicators

	(1)	(2)	(3)	(4)
	Intend to Complete	Think Will Do Well	Confident Can Master Basic Skills	Confident Can Master Skills
Enrolled late	-0.208* (0.117)	-0.0875 (0.177)	0.0378 (0.148)	0.102 (0.186)
Observations	2,226	2,212	2,216	2,207
Mean	6.700	5.739	6.369	5.671
Controls	No	No	No	No

Notes:

1. Includes enrollment, enrollment squared, and enrollment cubed, each estimated for enrollment within 15 days on both sides of the discontinuity.

3.pdf