# Is Faster Better? A Study of Video Playback Speed

David Lang
dnlang86@stanford.edu
Stanford University, United States

Guanling Chen
guanliang.chen@monash.edu
Monash University, Australia

Kathy Mirzaei
kmirzaei@stanford.edu
Stanford University, United States

Andreas Paepcke
paepcke@cs.stanford.edu
Stanford University, United States

## ABSTRACT

We explore the relationship between video playback speed and student learning outcomes. Using an experimental design, we present the results of a pre-registered study that assigns users to watch videos at either 1.0x or 1.25x speed. We find that students who consume sped content are more likely to get better grades in a course, attempt more content, and obtain more certificates. We also find that when videos are sped up, students spend less time consuming videos and are marginally more likely to complete more video content. These findings suggest that future study of playback speed as a tool for optimizing video content for MOOCs is warranted. Applications for reinforcement learning and adaptive content are discussed.

## CCS CONCEPTS

**• Human-centered computing → Human computer interaction (HCI)** .

## KEYWORDS

MOOCs, randomized controlled trials, clickstreams, video analytics, playback speed

**ACM Reference Format:**
David Lang, Guanling Chen, Kathy Mirzaei, and Andreas Paepcke. 2020. Is Faster Better? A Study of Video Playback Speed. In *Proceedings of the 10th International Conference on Learning Analytics and Knowledge (LAK '20), March 23–27, 2020, Frankfurt, Germany*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3375462.3375466

## 1 INTRODUCTION

Increasingly, education is becoming transformed by technology and innovation. This innovation offers the promise of more open and better solutions to distance education. In practice, many of these innovations have been viewed as failing to achieve these goals. The most recent example of this phenomena is the MOOC (Massively Open Online Course). Retrospective analyses of the period surrounding the rise and fall of the MOOC contend that these technologies are misunderstood, both in their primary application and their associated audiences [24]. While MOOCs may have the capacity to democratize education, they primarily have been used by college graduates to continue to advance their studies and as a form of self-study [6].

One of the intriguing consequences of the MOOC is that it allows for and creates the infrastructure for increasing understanding of learner behavior. The granularity of clickstream data and the fact that many MOOCs have enrollments in the thousands or even hundreds of thousands allows for researchers to understand the impact of relatively small considerations in course and platform design. Despite the great amount of learner and user behavior data these platforms have accrued, this technology has not always translated into improved learning outcomes. For example, course completion rates in MOOCs are much lower than traditional courses [10].

The most commonly cited reason for not completing a MOOC is that students claim they did not have sufficient time to complete the course [18]. In this paper, we address this claim by artificially shortening the duration of courses by increasing the speed of videos and examining the impact of shortened duration on grades and other related outcomes.

## 2 LITERATURE REVIEW

The motivation for this project draws from a variety of fields including cognitive science, linguistics, human-computer interaction, education, and economics.

### 2.1 Human Comprehension of Accelerated Speech

One of the primary concerns in initiating this study was that learners may be unable to comprehend accelerated content. Existing literature suggests that this concern may be overstated. Generally, the average human speaker's natural cadence is at a rate of 150 words per minute or 9 syllables per second [23]. With respect to listening and comprehension, literature reviews suggest that speech can be accelerated to nearly 250 words per minute before audio comprehension starts to deteriorate [21]. Studies across languages suggest that while the speaking rate of languages may vary, their information rates are remarkably constant, averaging 39 bits of information per second [8]. Other literature suggests that human comprehension may be further aided by using multiple channels of communication [16]. By distributing content via audio, video, and text components, learners may be able to infer context even if one channel of communication is unclear. MOOC videos are a


*LAK 2020, March 03–05, 2018, Frankfurt, Germany*

ACM ISBN 978-1-4503-9999-9/18/06…$15.00
https://doi.org/10.1145/3375462.3375466

promising candidate for an intervention based on this approach in that they contain not only video and audio but also annotated transcripts of the material.

## 2.2 Playback Speed Studies

Previous playback speed studies examined the context of short-duration interventions followed by an immediate assessment. These studies have typically found null or negative results [17]. For example, one study investigated how 59 students performed on a pre-post test assessment when students were assigned to either a 1.0x,1.5x, or 2.0x speed. This work found that students disliked lectures that had a speed of 2.0x but found little evidence that accelerated content had a negative impact on student performance. On the other side of the spectrum, research focused on college lectures found that instructors were perceived as less credible or authoritative when they were perceived as speaking at a relatively slow cadence [27]. Work that has focused on this type of intervention in time-intensive settings such as medical education have found deleterious effects on student learning outcomes when forced to view content in a sped state [28].

However, key limitations of previous research are two-fold. In addition to being relatively underpowered, these interventions have tended to be of short duration or assumed students will fully consume the content. In many contexts, students have elective choice and if they find a medium unpleasant, they may simply choose to not finish watching videos or dropout of a course. Moreover, many of these investigations have ignored the personalized use cases where students can choose to use other compensatory strategies such as rewinding or pausing videos when presented with complex or challenging content. Additionally, most of these interventions have been at a speed of 1.5x or more. Prior literature has thus left open question whether smaller increases in speed may induce time savings while not decreasing learner performance.

## 2.3 Video Interaction Technology

Despite previous studies of playback speed, there have been several meaningful innovations since the rise of MOOCs that suggest that findings from prior studies may not generalize to the latest iteration of video playback technology. First, the technology surrounding accelerated playback has changed. Originally, accelerating the playback speed of video altered the pitch and tempo of a speaker's voice. This makes comprehension more difficult as accelerated voices experience a phenomena called "audio chipmunking". Most course videos are now post processed such that a user's pitch will be maintained through different playback speeds [14]. Second, there has been a growing body of work that has attempted to intuit learner behavior from video interactions. One of these approaches involves looking at complex clickstream behaviors and mapping them to more interpretable grammars such as reviewing, skipping, and reflecting on content [5]. Log data of these video interactions can provide more useful measures of how accelerated playback speed could affect learner performance.

## 2.4 Video Design

The other feature that is worth noting in MOOC settings is that video design practices have changed compared to recording of traditional lectures [12]. One of these changes has many MOOC authors chunk content into short and precise videos that rarely exceed more than ten minutes. Benefits of this approach include increased consumption and easier review of content that is organized into these segments. Moreover, past studies utilized video content recorded in classrooms or lecture halls. More recently, MOOC video tends to be shot in studio settings with relatively little ambient noise. Students may be able to understand sped content in the latter state but may struggle with video content recorded in a traditional classroom. Moreover, many MOOCs have additional resources such as closed captions of lecture and high resolution recording that may aid learner comprehension in a sped state.

## 2.5 Time Use in Education

One of the puzzles that educational researchers and economists have been trying to solve is why there has appeared to be a decline in study time for college students from the 1960s to the early 2000s [1]. These studies speculate that part of the reason for the decline in study time may be some combination of increased productivity in education technology, increasing benefits associated with soft skills, and less pressure on time to completion at current academic institutions.

A related thread of literature suggests that increases in seat-time and increased instructional time are positively associated with student gains in mathematics [9] . While these two findings are not at odds with one another, they suggest that educational systems may currently be squeezed by dual constraints of reduced student time and pressure to increase instructional time. Identifying candidate solutions that directly address both of these issues would be an important step in improving educational outcomes.

## 2.6 Education Production Functions and Educational Technology

Educational innovations and their downstream effects on educational quality are often fraught with concerns about efficacy. Much of the promise of MOOCs were tempered by findings that documented high course attrition [10]. Similar innovations continue to face criticism because their adoption may have unintended effects. For example, video-taped lectured may open courses to more participants, but may at the same time decrease in-person attendance [26]. Such criticism has noted that whether or not this technology ends up being a complement or a substitute to a student's other learning approaches is ambiguous [11]. For example, while laptops are a tremendous source of educational productivity, they may be a distraction within the classroom [22]. Further, policy constraints on the use of technology can often have unintended and undesirable consequences. For instance, banning students from laptop use during lectures could have a deleterious effect on attendance. Moreover, even when schools are in a position to address these concerns with policies like mandatory attendance, the policies may be sub-optimal. Certain models of human capital accumulation have found that mandatory attendance policies can have a deleterious effect on student grades if mandatory attendance reduces students' ability to engage in self-study and review [4]. Given that platform designers can ease or increase the frictions associated with variable playback speed, understanding these considerations as well as the unintended consequences of their implementation is crucial.

## 2.7 Time Management

Amongst the most commonly explored interventions to improve course persistence and completion in MOOCs are strategies that help learners manage and address barriers to course completion. Past experiments that asked MOOC learners to explicitly schedule their study time found largely null results [2]. Attempts that nudged students to allocate more time towards studying generally resulted in null results [19].

Intriguingly, despite citing lack of time as the primary reason for not completing a MOOC, there has been relatively little research studying how innovations and time savings may help students. An economic theory argument would consider the decision to consume video content as a trade-off between consuming educational content and using this time for any other purpose. When a user watches videos faster, they have more time to allocate to educational and non-educational activities. In traditional economic parlance, this trade-off is known as an income effect. The other consequence of being able to consume more educational content more quickly is that the relative price of education becomes cheaper, suggesting that individuals will obtain more education. The only uncertainty in this assumption is whether the difficulty of learning changes when playback speed is modified. Asymptotically, we have little reason to believe that individuals will learn at either extremely slow or extremely fast speeds; however, there has been relatively little exploration of the region surrounding unity.

## 2.8 Gamification, Engagement, and Efficacy Tradeoffs

One of the most common tradeoffs that must be considered while designing educational content is whether or not the content is actually used. For example, one could imagine a product that is highly effective at teaching a certain skill but is unlikely to be used because of a user's preferences. Typically, platform designers face this problem and must choose which elements of content they will sacrifice in order to make content more engaging. Recent work highlighting this tradeoff includes the usage of flashcards versus educational chatbots to teach content. This research found that students learned better on a per minute basis with flashcards but users chose to study for longer periods when given the chatbot [25]. Other tools in this vein involve gamification. Previous work has found that students tend to spend more time engaging with a platform when it is gamified with badges and points [20]. This intervention has two risks. First, such interventions have the potential to undermine intrinsic motivation in learning [15]. Second, these types of interventions could cause learners to try to optimize for the gamified point or badge system rather than try to learn actual content [3].

Modifying playback speed represents a similar risk in that there may be a trade-off between fidelity of the learner's understanding and content coverage.

# 3 MODEL

We assume that users are optimizing their educational productivity subject to a budget constraint. A common approach to model an individual's productivity is to assume that it follows a Cobb-Douglas production function [7]:

$$K = F^{\alpha} I^{\beta} \tag{1}$$

In this equation $F$ denotes formal learning. This could represent time allocated to attending lecture, sessions, or other formal learning experiences. For simplicity, we are defining $F$ to only refer to watching lectures. $I$ represents informal learning such as self-study of lecture notes, independent review, and time working on supplementary problems. The parameters $\alpha$ and $\beta$ correspond to factors that influence how increases in formal and informal learning change when scaled.

We also assume that individuals are subject to temporal budget constraints such that individuals have a fixed allocation of hours for education but can distribute their time between formal and informal learning:

$$p_F F + p_I I = H \tag{2}$$

The coefficient $p_F$ is an individual's associated time to consume a unit of formal education while $p_I$ is the associated time for an individual to consume a unit of informal education. $H$ represents the total number of hours that an individual has to consume educational content.

We then formulate a learner's decision to maximize their learning as an optimization problem where learners try to maximize their human capital by choosing a bundle of formal and informal learning subject to their time constraints. The associated optimization can be expressed by the following problem:

$$\begin{aligned} \underset{F,I}{\text{maximize}} \quad & F^{\alpha} I^{\beta} \\ \text{subject to} \quad & H = p_F F + p_I I, \\ & p_F, p_I, F, I, H, \alpha, \beta \geq 0 \end{aligned} \tag{3}$$

This optimization problem can be expressed as a Lagrangian :

$$F^{\alpha} I^{\beta} + \lambda(H - p_F F - p_I I) \tag{4}$$

Equation 4 can be solved by taking partial derivatives with respect to $I$ and $F$. this problem suggests that the optimal consumption of educational goods would be

$$I* = (\frac{\beta}{\alpha + \beta})\frac{H}{p_I} \tag{5}$$

$$F* = (\frac{\alpha}{\alpha + \beta})\frac{H}{p_F} \tag{6}$$

Assuming this model of behavior is correct, we then interpret what would happen if the associated cost of acquiring formal capital were to decrease due to playback speed innovation.We illustrate this allocation in Figure 1 for an individual who has twenty-hours available for formal and informal learning. The solid and dashed blue line correspond to budget constraints before and after the change to playback speed. In response to the playback speed change, learners can choose larger bundles of formal and informal learning with the same amount of time. The solid and dashed red lines correspond to the indifference curves, bundles of learning that produce the

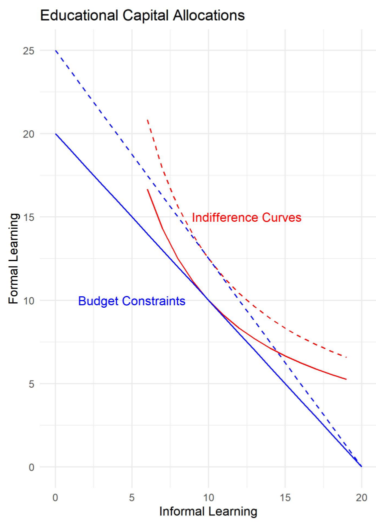


**Figure 1: Budget Constraints and Indifference Curves**

same educational capital.The intersection of the indifference curve and the budget constraints correspond to the optimal consumption bundle. This model suggests that as the associated temporal cost for formal learning decreases, individuals should invest more heavily in formal education. The only uncertainty is whether or not content becomes more difficult to consume as it is accelerated.

## 3.1 Playback Speed

With that consideration in mind, we also wish to remind readers of an important fact. There is an asymptotic relationship between playback speed and the amount of time to complete content. Assuming an individual completes the content, the amount of time required to consume the content can be represented by the following equation:

$$Time = \frac{VideoLength}{PlaybackSpeed} \quad (7)$$

One important implication of this relationship is that there are diminishing and decreasing returns to accelerating content. Figure 2 below illustrates this asymptotic property with a learner who consumes a twenty-hour course. If the individual learner were to increase their associated speed from 1.0x to 1.25x, this learner would take approximately sixteen hours to view the entirety of the course and save four hours of time. If the individual attempts to save additional time by increasing the speed another 1.5x, they would save approximately approximately seven hours. It may seem advantageous to go at the maximal possible speed but we also know that literature suggests that human comprehension of speech has an upper bound. Large increases in playback speed may not fall within the range of human comprehension of speech.Consequently, it may be wise to focus on relatively small increases in playback speed, where the time-savings are largest on a per-unit basis.

# 4 HYPOTHESES

Based on our model of human capital, there are several natural implications. As such, we generate the following hypotheses:

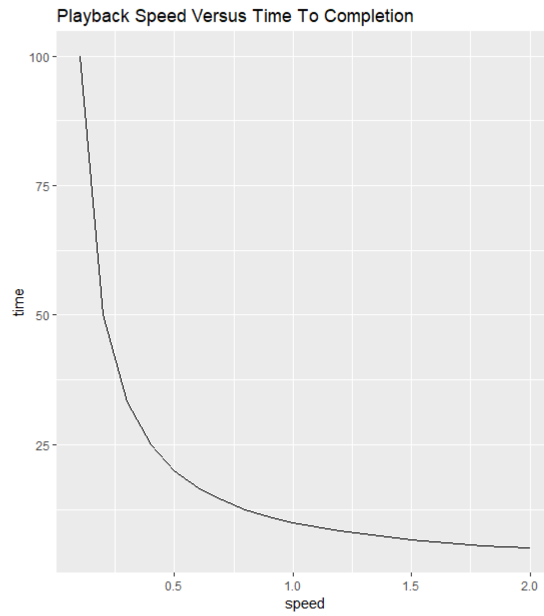


**Figure 2: Playback Speed versus Time**

***H1. Students who are exposed to video in a sped state will experience time savings.***

Based on our model of how individuals allocate their time to build human capital, learners should increase their consumption of video content but experience no time savings. In practice, however, we know that the number of hours of content that an individual can consume in a given course is fixed. As such, we believe that students who watch videos in a sped state will experience some time savings[1].

***H2. Students who are exposed to sped videos will consume more content.***

We know that MOOCs experience high dropout rates. One of the potential consequences of our intervention is that many students will still dropout but the associated time savings will result in students making more progress through the course and its associated videos before dropout occurs.

***H3. Students who experience videos in a sped state will exhibit different pausing and rewind behavior.***

The rationale for this hypotheses is two-fold. On the one hand, if content becomes more difficult as a result of this sped state, students may exhibit more self-regulatory behavior to better understand content. On the other hand, if students are disengaged by sped material, students will be less likely to engage and this may manifest in fewer clickstream activities.

***H4. Students who are in a sped state will get better grades.***

One of the impacts of students consuming videos more quickly is that students will have additional time to work on course exams and assignments. As a consequence of this fact, we hypothesize students will attempt more content and get better grades.

[1]This hypotheses was not part of the preregistered hypotheses.

# 5 DATA

The experimental data we are utilizing comes from six distinct courses in Stanford's Lagunita Platform. These courses were chosen due to past enrollment and the heterogeneity these courses represented in subject matter and length. All of these courses had enrollment of at least 1,000 learners in their prior administrations. The duration of these courses had an associated video component ranging from two to over twenty hours of video content (See Table 1) . On average, each video was approximately fifteen minutes long. The specific courses are:

(1) Introduction to SQL
(2) Introduction to Medical Statistics
(3) Introductory Computer Science
(4) Scientific Writing
(5) Introduction to Statistical Learning
(6) Introduction to Algorithms

**Table 1: Video Statistics**

| course | Avg Video Length | Videos | Hours |
|---|---|---|---|
| SQL | 13.723 | 10 | 2.287 |
| CS | 14.618 | 36 | 8.771 |
| ScientificWriting | 16.792 | 52 | 14.553 |
| StatLearning | 12.775 | 77 | 16.395 |
| Algorithms | 16.146 | 76 | 20.451 |
| Med Stats | 13.285 | 111 | 24.577 |

## 5.1 Observational Data

First, we present observational data with regards to the usage of playback speed modification on the platform prior to our intervention. Our analytical sample correspond to the same six courses as our experimental data but prior to implementation of the experiment. In the pretreatment cohort, playback speed modification is a comparatively rare event. Of the nearly 21, 835 learners enrolled in the courses, 4, 345 decided to modify their playback speed in any of their video interactions. The fact that approximately a fifth of users exhibit this behavior suggests not only that this population is somewhat common but also that a large percentage of users may be unaware of this option.

When users modified their playback speed, they tended to not go towards extrema. Users had a choice of up to six distinct playback speeds: 0.5x ,0.75x, 1.0x, 1.25x,1.5x, and 2.0x. Figure 3 shows the cumulative distribution of playback speed across individuals who modify playback speed. The median user tends to choose a playback speed between 1.25x and 1.5x.

Generally, few individuals choose playback speeds slower than the default of 1.0x.

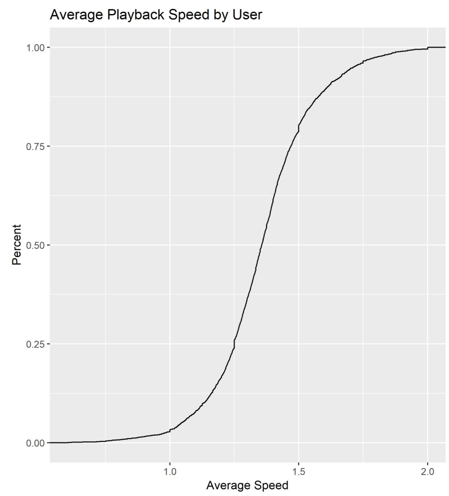


**Figure 3: Average Playback Speed (Pre-treatment Cohort)**

# 6 RESEARCH DESIGN

This experiment was preregistered at the Open Science Foundation https://osf.io/pm5b6/. The plan stated that all courses would run for at least six months before final results were calculated. The results as currently presented do not reflect the preregistered version in one important way. First, according to the preregistration, we planned on six courses with enrollment of at least 1000 students in each course, but one course experienced a technical failure. The intervention has been re-implemented in this course but subsequent enrollment in this course has been anemic (Scientific Writing).

## 6.1 Intervention

The nature of the intervention was to assign students to consume videos at either 1.25x or 1.0x speed. Individuals were randomized upon enrollment in the course. Students in the treatment group were assigned a playback speed of 1.25x. Each time a student in the treatment group pressed play on a video, students' videos were set to play at 1.25x speed. For students in the control group, playback speed was set back to 1.0x. Each time a student interacted with a new video, their playback speed was reset to their assigned speed. Students were able to manipulate the treatment but would have to actively manipulate it each and every time they interacted to not receive any treatment.

## 6.2 Measures

Translating our hypotheses to measures requires some modest assumptions, due to the fact that most of our outcomes require clickstream data to generate activity logs. For example, the log data we are using can detect when a user has started watching videos but it cannot detect whether or not a user has walked away from their computer screen or when a user closed out of a particular video. Moreover, we cannot detect any activity that the user engages in that does not explicitly utilize the platform. For example, if users save content to review outside of the platform, we can not measure the playback speed of that interaction.

With these caveats, we constructed the following three measures to track video consumption:

### 6.2.1 *Video Watching and Video Completion.*

(1) **Time Spent Watching Videos** We defined a user spending time watching video based on the time between video

interactions. If an individual were to spend more than thirty minutes between video interactions, we impute that activity to be zero. Based on the length of videos, we feel that this definition should be relatively precise. Videos that last longer than a half-hour are extraordinarily rare. This measure also assures that individuals were actively engaging between these intervals.

(2) **Video Completion** Another criterion was whether or not a user reaches the end of a video. This definition assures that a user's browser reached the end of a video but not whether or not the individual watched it.

(3) **Average Video Completion** We define this final outcome, in part, because many of the courses contain summations, course credits, and acknowledgements. We believe that many users may deliberately skip this content. This measure potentially captures partial consumption of videos.

*6.2.2 Regulatory Behaviors.* The goal behind identifying these regulatory behaviors is that individuals in sped states may engage in varying amounts of reflective and compensatory behavior.

(1) **Pausing** We define this as a count of the number of times a user pauses the video content. This measure could indicate when a user takes the opportunity to process content or reflect on prior learning.
(2) **Rewinding** We define this as any time a user engages in video seek behavior in a backward direction of more than one second. A key limitation of this measure is that the underlying user interface is somewhat imprecise. Users may rewind and fast forward repeatedly to identify a particular sequence.
(3) **Skipping** We define this as any time a user engages in video seek behavior in a forward direction of more than one second. Again, the imprecision of the underlying user interface is a limitation.

*6.2.3 Grades and Non-Video Behavior.*

(1) **Item Attempts** We look at this as a measure of how often learners interact with course content. We define this as the number of items a user interacts with, on at least one occasion.
(2) **Final Grades** We examine this as a measure of a learner's mastery of content throughout the course. Final grades represent the percentage of items that a learner correctly identified across all possible items. We assume the items that are not attempted get zero credit.
(3) **Certificates** Similar to course grades, we also look at measures of whether or not a learner ultimately ended up receiving a certificate for the course. In our sample, the pass criterion ranges from 50% to 90%.

## 7 ESTIMATION STRATEGY

The estimation strategy that we use is as follows. We regress our focal outcome $Y_{ij}$ as a function of individual $i$'s treatment status in course $j$. We define students assigned to fast content as being in our treatment cohort. $\Gamma_j$ represents a course-specific fixed effect. All errors are clustered on the course level.

$$Y_{ij} = \beta_1 Fast_{ij} + \Gamma_j + \epsilon_{ij} \tag{8}$$

### 7.1 Randomization Checks

Balance checks were performed in Table 2 and suggest that individual characteristics are relatively balanced across treatment and control[2]. Users could potentially manipulate their treatment status by repeatedly enrolling and unenrolling in the course. However, we found no evidence of such behavior.

## 8 FINDINGS

In this section we report the results for the three outcome areas of interest with regard to our playback speed intervention.

### 8.1 Confirmatory Analysis

*8.1.1 Video Consumption.* With respect to video consumption, we examined three measures to identify both the amount of time spent interacting with videos and the number of videos with which students interacted. As illustrated in Table 3, we find that students who are consuming faster content experience modest time-savings, supporting ***H1***. Students exposed to that treatment saved an average fourteen-hundred seconds in time savings, more than a third of an hour of time. This estimate is statistically significant with a p-value of less than 0.05. We find comparably less evidence that this playback speed intervention affected how much content user's consume. Using our measure of average video progress, we find null results. Using our coarser measure of whether or not students completed videos, we find marginal evidence that students in the treatment group were more likely to consume more content, suggesting moderate evidence in support of ***H2***.

*8.1.2 Self-Regulatory Behavior.* One of the concerns regarding this intervention is that it may materially affect how students use self-regulatory strategies and practice metacognition (See ***H3***). Table 4 shows that students are less likely to exhibit certain self-regulatory behaviors when exposed to accelerated content. Students are less likely to pause videos when given accelerated content. We also find that students are marginally more likely to rewind videos, suggesting mixed changes to self-regulatory behavior.

*8.1.3 Course Performance.* Finally, while video playback speed is the primary mechanism of our intervention, our intervention may also affect a student's performance on assessments. Our results in Table 5 shows that students who are exposed to accelerated lectures are more likely to obtain higher grades and attempt more course content. On average, students in a sped state get higher grades by two percentage points. In terms of the effect size, these effects are around .05 $\sigma$ . These effect sizes are quite similar to early work by Dobbie and Fryer that studied increases in instructional time [9]. Their work found that increasing seat-time by 25% had an similar magnitude increase in student test scores.Thus, we find strong evidence in support of ***H4***.

In part, this increase in grades stems from the fact that these students are more likely to attempt additional content. Students exposed to fast videos attempted an additional 2.3% of course content. These effects also translate into whether or not an individual

[2] TimeStamps for these demographic data were not available.

**Table 2: Balance Checks**

| | (1) Control | | (2) Fast | | T-test Difference |
|---|---|---|---|---|---|
| Variable | N/[Clusters] | Mean/SE | N/[Clusters] | Mean/SE | (1)-(2) |
| Male | 3252 | 0.605 | 3144 | 0.613 | -0.008 |
| | [6] | (0.029) | [6] | (0.026) | |
| Missing Gender | 3252 | 0.185 | 3144 | 0.178 | 0.007 |
| | [6] | (0.012) | [6] | (0.014) | |
| College or Higher | 3252 | 0.354 | 3144 | 0.368 | -0.014 |
| | [6] | (0.029) | [6] | (0.030) | |
| Domestic Learner | 3252 | 0.373 | 3144 | 0.371 | 0.002 |
| | [6] | (0.028) | [6] | (0.024) | |
| Age | 2498 | 29.405 | 2418 | 29.157 | 0.248 |
| | [6] | (0.954) | [6] | (1.095) | |
| Missing Age | 3252 | 0.232 | 3144 | 0.231 | 0.001 |
| | [6] | (0.013) | [6] | (0.005) | |

*Notes*: The value displayed for t-tests are the differences in the means across the groups. Standard errors are clustered at variable course_code. Fixed effects using variable course_code are included in all estimation regressions. ***, **, and * indicate significance at the 1, 5, and 10 percent critical level.

**Table 3: Video Consumption**

| | *Dependent variable:* | | |
|---|---|---|---|
| | Time Use | Video Progress | Video Completion |
| | (1) | (2) | (3) |
| Fast | $-1{,}425.099^{***}$ | 0.031 | $0.412^{*}$ |
| | (533.489) | (0.019) | (0.217) |
| Observations | 6,396 | 6,396 | 6,396 |
| $R^2$ | 0.029 | 0.025 | 0.043 |
| Adjusted $R^2$ | 0.028 | 0.024 | 0.042 |

*Note:* All models include a course fixed effect. Errors are clustered on the course level.

$^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01

**Table 4: Self-Regulatory Behaviors**

| | *Dependent variable:* | | |
|---|---|---|---|
| | Pauses | Video Rewinds | Video Skips |
| | (1) | (2) | (3) |
| fast | $-7.080^{**}$ | $0.293^{*}$ | $-0.017$ |
| | (3.090) | (0.152) | (0.252) |
| Observations | 6,396 | 6,396 | 6,396 |
| $R^2$ | 0.009 | 0.004 | 0.003 |
| Adjusted $R^2$ | 0.008 | 0.003 | 0.002 |

*Note:* All models include a course fixed effect. Errors are clustered on the course level.

$^{*}$p<0.1; $^{**}$p<0.05; $^{***}$p<0.01

obtains a certificate in their course of study. Students in the treatment group are nearly 2.6 percentage points more likely to obtain a certificate in their course of study. This effect translates to nearly a 10% increase relative to the base rate across courses. In absolute terms, the largest lift was in the SQL course with a nearly five percentage point increase (see Figure 4)[3].

## 8.2 Exploratory Analysis

While our main effects suggest that modifying playback speed may be an effective tool for improving student performance, there remain several outstanding questions with regards to the appropriateness of this tool. In particular, we suspect that there may be treatment heterogeneity where accelerated playback speed is more effective for certain courses or for certain types of students.

*8.2.1 Course Heterogeneity.* To assess course-level heterogeneity, we separately estimated standardized treatment effects across each course for selected outcomes, namely grades and time use. Figure 5 displays the standardized effect size on grades for each course. Based on the $I^2$ statistic, we find little evidence of heterogeneous treatments effects for grades by course [13].

With respect to time use, we see a somewhat different picture. Notably as displayed in Figure 6, we find that students who were

[3]No students obtained certificates in the Scientific Writing course

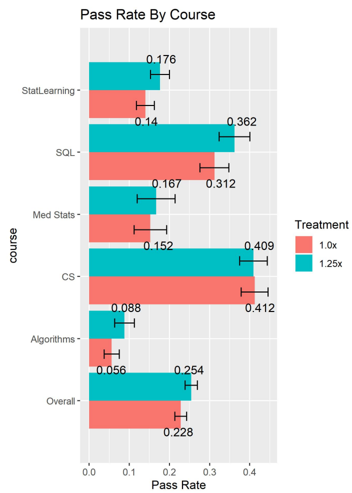


**Figure 4: Pass Rates**

**Table 5: Course Performance**

| | *Dependent variable:* | | |
|---|---|---|---|
| | Grades | Attempts | Certificate |
| | (1) | (2) | (3) |
| Fast | 0.020** | 0.023*** | 0.026*** |
| | (0.008) | (0.009) | (0.010) |
| Observations | 6,396 | 6,396 | 6,396 |
| $R^2$ | 0.187 | 0.175 | 0.092 |
| Adjusted $R^2$ | 0.187 | 0.174 | 0.091 |

*Note:* All models include a course fixed effect. Errors are clustered on the course level.

*p<0.1; **p<0.05; ***p<0.01

enrolled in the SQL course experienced out-sized time-savings relative to students who were enrolled in any other course, saving more than a fifth of a standard deviation. This finding is intriguing in that this course only had approximately two-hours of video content whereas all other courses had nearly an order of magnitude more content. With respect to the $I^2$ statistic, we report an estimate of 0.53 suggesting modest heterogeneity.

*8.2.2 Student Heterogeneity.* While we do not have a particularly rich set of demographic characteristics, we do know the reported age, educational status, and gender of students. To explore the presence of heterogeneity in this group of individuals, we utilize a technique called causal forests [29] This approach relies on building distinct regression trees for maximizing the split between treatment control and building a separate regression tree to estimate the effect. This process is done repeatedly with bootstrapped estimates to generate a causal forest.

One benefit of this approach is that we can calculate variable importance in the same manner as one would use for random forests. Moreover, this technique allows for post-hoc exploration of the data without introducing additional researcher degrees of freedoms. According to this technique, we find that the most important features for identifying heterogeneous treatment effects are a student's age, identification as male, and whether or not they identify as coming from the USA. The associated importance weights can be seen in Figure 7.

We then mapped these heterogeneous treatment effects into age bins as defined by the US Census. The plot in Figure 8 suggests that within these groups, treatment effects on grades are remarkably homogeneous, ranging from 0.045 to 0.059 deviations. We also find little heterogeneity with respect to user time savings behavior. Figure 9 illustrates that user time savings varies between -0.024 and -0.041 deviations.

## 9 DISCUSSION AND NEXT STEPS

This work presents several key findings with regards to course design and playback speed modification. Our intervention ,in effect, reduced the temporal costs of taking a MOOC course. This translated into higher course completion and persistence measures. This finding lends additional credence to the idea that part of the reason why MOOC persistence rates are so low is that MOOC learners simply did not have enough time to complete the course. Our findings are consistent with this hypotheses.

With respect to playback speed, this work documents that a non-trivial percentage of students engage with this playback-speed technology. Nearly a fifth of students used this tool before any of our interventions. We also found that through an experimental intervention, students who are assigned to watch accelerated content experience time savings and get better grades. The fact that we detected improved course performance and increased video consumption is consistent with our model of how individuals decide to build human capital. Nevertheless, these findings should be interpreted with some caution. Given that the associated time savings from these interventions are relatively small, we believe that these time savings and results may not translate into more traditional environments such as college lectures or other course work. The courses used in our analysis had high-drop out rates and low completion. If the gains from the increased playback speed are due to continued participation in the course, these benefits may not translate into environments where users are likely to persist.

Moreover, these courses are structured such that students receive short assessments after a handful of videos. For courses that have summative assessments, students may not retain knowledge when learned in a sped state.

Ultimately, what this work suggests is that playback speed is one of many potential parameters that educational platform designers may choose to calibrate in the future. These findings indicate that there may be some regions where instructional time can be reduced and learning can be improved with relatively little tradeoff. This suggests that future work in both reinforcement learning and contextual bandits could gain from incorporating playback speed as a tuning parameter.

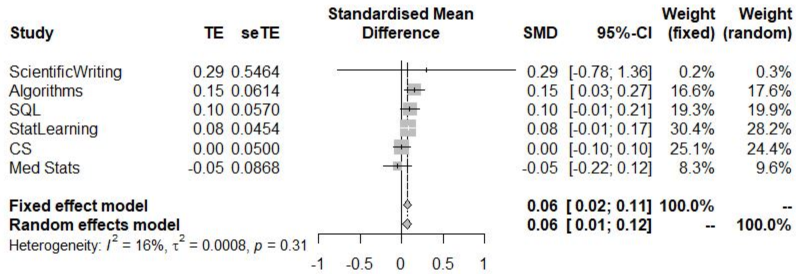


**Figure 5: Treatment Effects on Grades by Course**

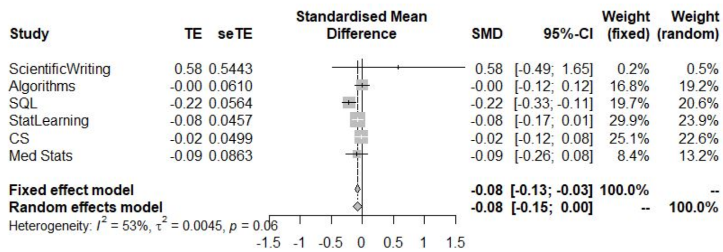


**Figure 6: Treatment Effects on Time Use by Course**

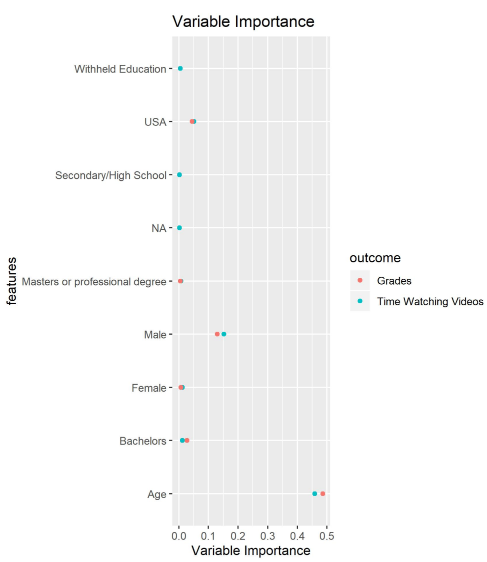


**Figure 7: Variable Importance**

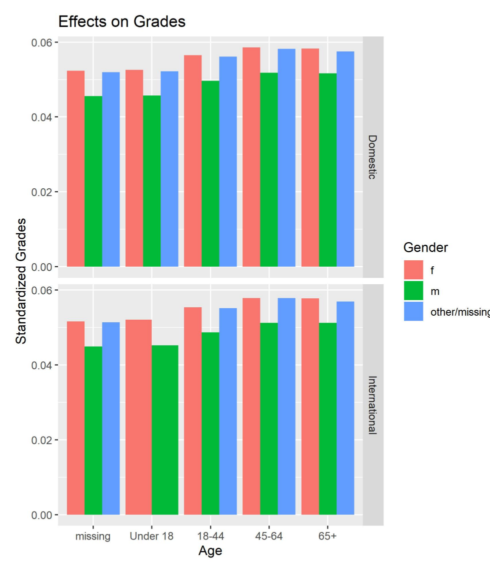


**Figure 8: Grade Heterogeneity**

## 9.1 Future Work

There are several key limitations that this work has yet to address. First, due to the structure of our intervention, users were largely constrained into their assigned playback speed. Replicating these findings with a softer nudge may make shifting default playback speeds more palatable to both users and platform designers.

Second, while our work has provided some evidence that 1.0x speed is not the optimal setting for all students or courses, it does not inform us what the local optimum for playback speed is. We are currently doing follow-up studies where we test additional speeds ranging from 0.75 x to 1.75x. Third, this work does not inform course designers or instructors about how rapidly they should speak. Subsequent work will look at the natural language characteristics of the transcripts to understand how lexical density may affect optimal playback speed.

Finally, this work looked at MOOCs. Colleges and universities are increasingly recording lecture and discussion to accommodate a wider range of student needs. Replicating this study in a traditional college environment would speak to the specific question of

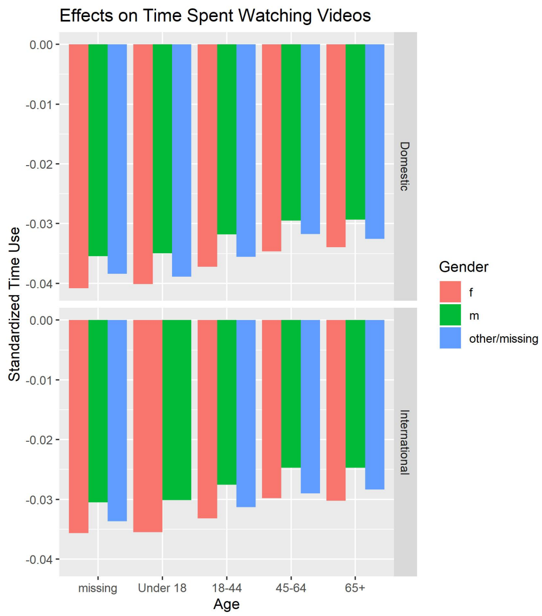


**Figure 9: Time Heterogeneity**

whether or not platforms should default to faster playback speeds in these environments as well.

## 10 ACKNOWLEDGEMENTS

The research reported here was supported by the Institute of Education Sciences, U.S. Department of Education, through Grant R305B140009 to the Board of Trustees of the Leland Stanford Junior University. The opinions expressed are those of the author and do not represent views of the Institute or the U.S. Department of Education.

## REFERENCES


[1] Philip S Babcock Mindy Marks, Kelly Bedard, Ted Bergstrom, Eli Berman, John Briggs, Steven Brint, David Fairris, Peter Kuhn, Valerie Ramey, Doug Steigerwald, Aman Ullah, Philip S Babcock, and Mindy Marks. 2010. *NBER WORKING PAPER SERIES THE FALLING TIME COST OF COLLEGE: EVIDENCE FROM HALF A CENTURY OF TIME USE DATA This work benefited from conversations with The Falling Time Cost of College: Evidence from Half a Century of Time Use Data.* Technical Report. http://www.nber.org/papers/w15954

[2] Rachel Baker, Brent Evans, and Thomas Dee. 2016. A Randomized Experiment Testing the Efficacy of a Scheduling Nudge in a Massive Open Online Course (MOOC). *AERA Open* 2, 4 (10 2016), 233285841667400. https://doi.org/10.1177/2332858416674007

[3] Joseph E. Beck and Yue Gong. 2013. Wheel-Spinning: Students Who Fail to Master a Skill. Springer, Berlin, Heidelberg, 431–440. https://doi.org/10.1007/978-3-642-39112-5{_}44

[4] Massimiliano Bratti and Stefano Staffolani. 2013. *ADRES GENES Student Time Allocation and Educational Production Functions.* Technical Report 111. 103–140 pages. https://www.jstor.org/stable/pdf/23646328.pdf?refreqid=excelsior%3A3d376429584992b41e02a2b3390afa78

[5] CG Brinton and S Buccapatnam. 2015. Mining MOOC Clickstreams: On the Relationship Between Learner Video-Watching Behavior and Performance. *ACD Conference on* (2015). https://pdfs.semanticscholar.org/26c4/5fac8de86ecde6eb6f1b99ee68df65babe3f.pdf

[6] G Christensen, A Steinmetz, and B Alcorn. 2013. The MOOC phenomenon: who takes massive open online courses and why? *Available at SSRN* (2013). http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2350964

[7] CW Cobb and PH Douglas. [n. d.]. A THEORY OF PRODUCTION1. *msuweb.montclair.edu* ([n. d.]). http://msuweb.montclair.edu/~lebelp/CobbDouglasProdAER1928.pdf

[8] Christophe Coupé, Yoon Oh, Dan Dediu, and François Pellegrino. 2019. Different languages, similar encoding efficiency: Comparable information rates across the human communicative niche. *Science Advances* 5, 9 (9 2019), eaaw2594. https://doi.org/10.1126/sciadv.aaw2594

[9] W Dobbie and RG Fryer Jr. 2013. Getting beneath the veil of effective schools: Evidence from New York City. *American Economic Journal: Applied* (2013). http://www.ingentaconnect.com/content/aea/aejae/2013/00000005/00000004/art00002

[10] BJ Evans, RB Baker, and TS Dee. 2016. Persistence patterns in massive open online courses (MOOCs). *The Journal of Higher Education* (2016). https://muse.jhu.edu/article/609047/summary

[11] Pedro Gomis-Porqueras, Jürgen Meinecke, and José A Rodrigues-Neto. 2011. *New Technologies in Higher Education: Lower Attendance and Worse Learning Outcomes?* Technical Report 1. 69–83 pages. https://www.jstor.org/stable/pdf/43200622.pdf?refreqid=excelsior%3Aa18f04c9e17fa69d41931a577e67627c

[12] Philip J. Guo, Juho Kim, and Rob Rubin. 2014. How video production affects student engagement. In *Proceedings of the first ACM conference on Learning @ scale conference - L@S '14*. ACM Press, New York, New York, USA, 41–50. https://doi.org/10.1145/2556325.2566239

[13] Julian Higgins, Simon Thompson, Jonathan Deeks, and Douglas Altman. 2003. Measuring inconsistency in meta-analyses [journal article as teaching resource, deposited by John Flynn]. *British Medical Journal* 327, 7414 (9 2003), 557–560. https://doi.org/10.1136/bmj.327.7414.557

[14] 360 JH Chen US Patent 8, 032 and undefined 2011. [n. d.]. System and method for high-quality variable speed playback of audio-visual media. *Google Patents* ([n. d.]). https://patents.google.com/patent/US8032360B2/en

[15] Mark R. Lepper, David Greene, and Richard E. Nisbett. 1973. Undermining children's intrinsic interest with extrinsic reward: A test of the "overjustification" hypothesis. *Journal of Personality and Social Psychology* 28, 1 (1973), 129–137. https://doi.org/10.1037/h0035519

[16] RE Mayer, R Moreno, and M Boire. 1999. Maximizing constructivist learning from multimedia communications by minimizing cognitive load. *Journal of educational* (1999). http://psycnet.apa.org/journals/edu/91/4/638/

[17] Toru Nagahama and Yusuke Morita. 2017. *Effect Analysis of Playback Speed for Lecture Video Including Instructor Images.* Technical Report 1. 50–58 pages. http://jaems.jp/contents/icomej/vol11/06_Nagahama.pdf

[18] I Nawrot, A Doucet Proceedings of the 23rd International Conference on, and undefined 2014. [n. d.]. Building engagement for MOOC students: introducing support for time management on online learning platforms. *dl.acm.org* ([n. d.]). https://dl.acm.org/citation.cfm?id=2580054

[19] P Oreopoulos and U Petronijevic. 2019. The remarkable unresponsiveness of college students to nudging and what we can learn from it. (2019). https://www.nber.org/papers/w26059

[20] E O'Rourke, E Peach, and CS Dweck. 2016. Brain Points: A Deeper Look at a Growth Mindset Incentive Structure for an Educational Game. *Proceedings of the Third (* (2016). http://dl.acm.org/citation.cfm?id=2876040

[21] R Pastore and AD Ritzhaupt. 2015. Using Time-Compression To Make Multimedia Learning More Efficient: Current Research and Practice. *TechTrends* (2015). http://link.springer.com/article/10.1007/s11528-015-0841-2

[22] Richard W. Patterson and Robert M. Patterson. 2017. Computers and productivity: Evidence from laptop use in the college classroom. *Economics of Education Review* 57 (4 2017), 66–79. https://doi.org/10.1016/J.ECONEDUREV.2017.02.004

[23] Jonathan E Peelle and Matthew H Davis. 2012. Neural Oscillations Carry Speech Rhythm through to Comprehension. *Frontiers in psychology* 3 (2012), 320. https://doi.org/10.3389/fpsyg.2012.00320

[24] Justin Reich and José A. Ruipérez-Valiente. 2019. The MOOC pivot. *Science* 363, 6423 (1 2019), 130–131. https://doi.org/10.1126/science.aav7958

[25] Sherry Ruan, Liwei Jiang, Justin Xu, Bryce Joe-Kun Tham, Zhengneng Qiu, Yeshuang Zhu, Elizabeth L Murnane, Emma Brunskill, James A Landay, and Zheng-neng Qiu. [n. d.]. QuizBot: A Dialogue-based Adaptive Learning System for Factual Knowledge. 13 ([n. d.]). https://doi.org/10.1145/3290605.3300587

[26] Graham P Shaw and David Molnar. [n. d.]. Article Non-native English Language Speakers Benefit Most from the Use of Lecture Capture in Medical School. ([n. d.]). https://doi.org/10.1002/bmb.20552

[27] Brent K. Simonds, Kevin R. Meyer, Margaret M. Quinlan, and Stephen K. Hunt. 2006. Effects of Instructor Speech Rate on Student Affective Learning, Recall, and Perceptions of Nonverbal Immediacy, Credibility, and Clarity. *Communication Research Reports* 23, 3 (9 2006), 187–197. https://doi.org/10.1080/08824090600796401

[28] Kristine Song, Amit Chakraborty, Matthew Dawson, Adam Dugan, Brian Adkins, and Christopher Doty. 2018. Does the Podcast Video Playback Speed Affect Comprehension for Novel Curriculum Delivery? A Randomized Trial. *The western journal of emergency medicine* 19, 1 (1 2018), 101–105. https://doi.org/10.5811/westjem.2017.10.36027

[29] Stefan Wager and Susan Athey. 2015. Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. (10 2015). http://arxiv.org/abs/1510.04342