When the Student Becomes the Master: A Field Experiment on Learning by Teaching*

Rohen Shah†

September 29, 2025

Abstract

As AI tools become commonplace, there is an increasing concern that students may put less effort into their homework. To mitigate this potential learning loss, it is important for teachers to assign tasks that actively engage students at home. This study presents evidence from a large-scale field experiment that randomly assigned students to experimental conditions in which they are either (1) assigned to make TikTok-style "explanation" videos for homework, (2) assigned additional practice problems for homework, or (3) placed in a pure control condition. The explanation treatment improved short-run scores by 0.17 standard deviations and long-run grades by 0.07 standard deviations relative to the practice-problem group. Notably, while both treatment groups improved relative to control, only the explanation treatment improved performance on novel problems, suggesting that explaining concepts enhances one's ability to understand deeply and generalize concepts.

JEL Classifications: D91, I21, J24, 015

Keywords: Field Experiments, Economics of Education, Behavioral Economics

*This study was pre-registered on the AEA RCT Registry (AEARCTR-0011884) and approved by the Social and Behavioral Sciences IRB from the University of Chicago (IRB23-0041). I would like to thank Ariel Kalil, John List, Jens Ludwig, Lesley Turner, Andrew Johnston, Noah Liu, Emileigh Harrison, Phil Oreopoulos, Goya Razavi, Derek Rury, Rebecca Shah, and Haruka Uchida for valuable feedback. I also thank the school administrators and teachers at participating districts for being generous with their time and willingness to make this study happen.

This project was supported by the Administration for Children and Families (ACF) of the United States (U.S.) Department of Health and Human Services (HHS) as part of a financial assistance award (Award#: 90PD0322-01-01) totaling $24,948 with 30 percent funded by ACF/HHS and $57,960 amount and 70 percent funded by the following non-government sources: J-PAL North America and the Horowitz Foundation for Social Policy. The contents are those of the author(s) and do not necessarily represent the official views of, nor an endorsement, by ACF/HHS, or the U.S. Government. For more information, please visit the ACF website, Administrative and National Policy Requirements.

†Yale University, Department of Economics. Email: rohen.shah@yale.edu

1 Introduction

Suppose that Jenny is a typical student who puts a modest amount of effort in school. Her friend, Jack, asks Jenny for help with the most recent homework assignment. As Jenny attempts to help, she realizes that she did not understand the material as well as she thought. Jenny engages more with the material, and through the process of generating explanations for Jack, Jenny comes to understand the concepts more deeply.

As modern economies increasingly rely on human capital for growth, economists have sought ways to develop human capital efficiently (Fryer, 2017). While economists have widely documented "Learning by Doing" (Arrow, 1962) in the context of firm production, it is not clear how this can be leveraged to improve human capital development for students. In this study, I propose that *teaching* is the "doing" that can lead to deep and efficient student learning. The problem is that the Jennys of the world typically do not have a friend asking them for help for every homework assignment. Is there a scalable way to provide students like Jenny more opportunities to engage in "doing" math?

In this paper, I test the impact of creating weekly math explanation videos on students' math skills. I conduct a field experiment in partnership with 20 public middle and high schools in the Midwestern United States during the spring 2024 semester. Math teachers at these schools who taught multiple periods (sections) of the same math class were invited to enroll their classrooms in the study. For each teacher, their class periods were randomly assigned to 1) Control, 2) "Passive Task", and 3) "Video Creation" conditions. Upon enrollment, all three class periods were given a baseline math assessment and survey, and then a post-test after three months. The control classroom had business-as-usual homework assignments between the baseline and post-test. The "Passive Task" classroom was assigned one weekly PSAT[1] math question (hereafter, "Task") on Google Form, in addition to their regular homework assignments, for eight weeks. The "Video Creation" classroom was assigned the exact same weekly PSAT math question on the same schedule as the "Passive Task" classroom, with a sole difference: the Video Creation

[1] in some cases, questions were taken from the ACT, SAT, or middle school state tests depending on the grade level and topic covered in that teacher's class

classroom's tasks asked them to create a video explaining their solution to the task to a hypothetical student. Students submitted their videos directly to their teachers via Flipgrid, Google Drive, Google Form, etc. based on the teacher's preference. This study enrolled 128 classrooms, for a total of 2,523 students.

The post-test at the end of the intervention contained 15 PSAT questions on topics that were relevant to the intervention "Tasks". A subset of these (6 questions) were *exact* Task questions that students in the Passive Task and Video Creation groups were assigned (hereafter, "Task Questions"). The rest (9 questions) were PSAT questions not previously assigned, but which still assessed the topics covered by the intervention Tasks (hereafter, "Novel Questions").

I find that the overall average completion rate across all students and tasks in the Passive Task classrooms was 84%, whereas that of the Video Creation classrooms was 62%. This difference might be expected given that the effort required to complete a task is higher for the Video Creation classrooms than the Passive Task classrooms. I estimate an Intent-to-Treat (ITT) effect of the treatment on overall math class grades and the post-test score, including a breakdown of performance on the "Task Questions" and "Novel Questions" subsets of the post-test. I find that relative to the Passive Task treatment, the Video Creation treatment led to a statistically significant improvement in both overall math class grade (0.07σ) and post-test score (0.22σ). There was no significant difference in either outcome between the Passive Task and the control groups. Further, when considering the "Task Questions" and "Novel Questions" subsets of the post-test, I find that the Video Creation group outperforms the control and Passive Task groups for both types of questions, whereas the Passive Task group only outperforms the control group for "Task Questions" but not for "Novel Questions." This implies that the Video Creation treatment was particularly effective at helping students generalize what they learned.

One potential mechanism driving the overall treatment effect might be the total amount of effort that a student would choose to spend on math as a result of being assigned the Video Creation task. On the one hand, the video creation might sufficiently

motivate students to put in additional study effort beyond the time it takes to record themselves solving the problem. On the other hand, students might have an inelastic total amount of time that they would spend on math (Cotton et al., 2020). In this case, the time it would take to record themselves solving the math problem would take away from the time they spend preparing. To disentangle these potential mechanisms, I estimate a proxy for student effort for each task. Below the text of the Task problems, I provided a link to a "help" video on YouTube for each task that explains the solution to a similar problem. While the help video for each task was identical for the Passive Task and Video Creation classrooms, I could track the total number of clicks on the help video separately for each classroom. I estimate the probability of clicking on a help video for any given task to be 11% for the Passive Task classrooms, and 25% for the Video Creation classrooms. This difference was statistically significant, and highlights that asking students to create videos could be way to induce student effort.

Unlike many cluster RCTs in education that randomly assign a treatment at the teacher or school level, the design of this experiment allows for the estimation of a "teacher-level" treatment effect. This is because the random assignment of classrooms was stratified by teacher, implying that every teacher has both a "Passive Task" and "Video Creation" classroom. This yields a distribution of treatment effects, and I find that the treatment effect of Video Creation relative to the Passive Task was positive for approximately 80% of teachers, and negative for 20%.

It is not obvious that the Video Creation task should lead to a positive impact on math skills relative to the Passive Task because of the lower likelihood of students completing the Video Creation task. If the completion rate for the Passive Task is sufficiently higher than the Video Creation task, we might expect a negative Intent-to-Treat effect on math skills. For instance, if most students in a teacher's Passive Task classroom complete their tasks, but virtually no students in that teacher's Video Creation classroom complete their tasks, then we would expect students in the Passive Task classroom to have a greater improvement in their math skills than those in the Video Creation classroom. I leverage the aforementioned distribution of "teacher-level" treatment effects to explore

the relationship between relative task completion and treatment effect. I measure the "relative compliance rate" as the task completion rate in a teacher's Video Creation classroom minus that of their Passive Task classroom. I find a statistically significant, positive correlation between the relative compliance rate and treatment effect. This suggests that the effectiveness of the Video Creation treatment depends on a teacher's ability to induce their students to create videos. It might also suggest that for the sake of maximizing student welfare, teachers who are unable to induce students to create videos might be better off switching to the passive task.

This work relates to recent laboratory experiments in behavioral economics that show that people are substantially more likely to learn things they discover on their own as opposed to hearing from others (Conlon et al., 2022), as well as work showing that students benefited from being asked to "give advice" via a short survey (Eskreis-Winkler et al., 2019). This also contributes to our understanding of how humans "transfer" knowledge to new situations. The cognitive science literature denotes different stages of "generalizing" knowledge, where the initial step is to recognize a "new problem" as one similar to a familiar scenario, followed by creating a mental map between the familiar and new scenario, and finally using that map to solve the new problem (Samat et al., 2019). A part of what might hinder this knowledge transfer is the ability to retrieve memory of the familiar scenario (Bordalo et al., 2020, 2021). Another hindrance might be not understanding the structure of the original problem well enough to be able to create the "map" between the old and new scenario (Ahn et al., 1992). In this experiment, I showed that students who did the passive task were able to retrieve the memory of those tasks because they outperformed the control group, but only the students who created videos were able to outperform the control group on "new scenario" questions. This might suggest that memory retrieval is not the issue, but rather the video creation process helps students understand the structure of the problem well enough to be able to make the requisite "mental map" relating their knowledge of the topic to the "new problem," thereby generalizing what they had learned.

This also contributes to the literature on scaling in economics. Many interventions

that have large positive impacts in the lab often fail to replicate when brought to the field (List, 2022). While lab studies indicate the potential for improving outcomes through "learning by teaching", this study presents a way to successfully implement this idea in a large-scale field setting to improve non-lab outcomes. Similarly, qualitative and lab studies in education research highlight the importance of engaging students in "active learning" to improve outcomes (Bonwell and Eison, 1991; Freeman et al., 2014; Markant et al., 2016). However, there has been mixed success with implementing active learning strategies in classroom settings, sometimes resulting in a negative effect relative to the status quo (Berlinski and Busso, 2017). Given the challenge in asking teachers to change the way they teach, this study shows that a lighter-touch intervention that simply adds a few assignments might be a more reliable way to engage students in active learning such that the effect holds at scale.

This also relates to the work by labor economists on motivating students. Provision of high-quality resources alone does not guarantee utilization by students who would benefit (Robinson et al., 2022). It is even difficult to motivate students when providing financial incentives (Brownback and Sadoff, 2020; Burgess et al., 2021; Fryer, 2017; Sadoff, 2014). One reason why financial incentives on "outputs" such as test scores might not be effective is because students might not know how to convert inputs (i.e. studying) into outputs (Cotton et al., 2020). Nevertheless, student motivation to put forth effort can also play a role in test performance (Gneezy et al., 2019). In general, incentivizing inputs is more reliable than incentivizing outputs (Gneezy and List, 2013). This study contributes to this literature by highlighting a way to incentivize student effort without the use of financial incentives. Randomly assigning students to create videos doubled their likelihood of using the "help video" resource, which indicates a substantial difference in student effort. This treatment might therefore be a more scalable way to induce student effort as it does not require schools to budget for prizes or financial incentives for students.

Finally, in a world where Generative Artificial Intelligence (GenAI) such as Large Language Models (LLM) like ChatGPT are more commonplace, a common concern is that students can now easily put in less cognitive effort in tasks completed at home,

where LLM usage cannot be monitored. This study shows that video creation at home is one type of homework assignment that could mitigate this concern, given the finding that students are more likely to put in effort in this task than the passive-task counterpart. Another contribution of this study is to highlight the potential downside in plans to scale up tutoring provision by replacing human peer tutors with LLMs. Doing so would lead to a loss in potential human capital gains by the would-be peer tutors, and quantifying this is necessary for a full cost-benefit accounting of such plans.

This paper is organized as follows. Section 2 discusses the prior work on learning by teaching and the rationale for this intervention. Section 3 presents a theoretical framework of student effort choice. Section 4 describes the experimental design, and Section 5 shows the results. I discuss mechanisms and the distribution of treatment effects in Sections 6 and 7 respectively, and conclude in section 8.

2 Background

2.1 Prior Evidence for Learning by Teaching

Conceptually, teaching may lead to more learning than traditional techniques because teaching may be more engaging in two different senses. First, teaching may lead people to spend more time and effort with academic content. This might happen if the teacher altruistically cares about the learner's outcome, or if they socially care about the learner's perception of the teacher's abilities. Second, teaching may force the teacher to engage more actively with the content, including thinking about the content in different ways as part of thinking through how to communicate that idea to someone else.

Evidence for "learning by teaching" primarily comes from lab settings, but the extent to which the treatment effect generalizes to the field is unclear. Lab studies have shown that students score better on a quiz when they are preparing to tutor someone, as opposed to preparing for a quiz (Fiorella and Mayer, 2013; Guerrero and Wiley, 2021). Students are randomly assigned to either a control condition or a "tutoring expectation" condition. Control participants are told that they have 10 minutes to study for a quiz

on a physics topic (the Doppler Effect) and are provided with study materials. The treatment participants are told that they have 10 minutes to prepare to tutor someone on the Doppler Effect, and are given the same study materials that the control group is provided. However, the treatment group is then given a quiz rather than actually tutoring someone (they are debriefed about this deception afterward), and the authors find a significant, positive effect on the quiz score. This sheds light on a mechanism through which tutors might learn from tutoring: they prepare more deeply before their tutoring sessions. While this evidence is suggestive that students might learn by tutoring, it is not clear whether the results from this lab setting would generalize to the field. In particular, the impact measured in the lab is from a single preparation session, and the effect might not be sustained over an entire semester. Additionally, if students in a field setting choose not to review content before tutoring, then the results of the study would not apply.

There have not been many field experimental studies that have identified the effect of tutoring on the tutor's knowledge. Those that do often suffer from substantial identification threats. For instance, some studies on peer tutoring measure what students learn after both receiving and providing tutoring to a peer (De Backer et al., 2012; King et al., 1998). These designs measure the improvement in skills over time (pre-post design), but do not have a "control group" of students that do not engage in the activity. 5th grade students in Greene et al. (2018) were assigned to tutor 3rd and 4th graders. The 5th grade students were randomly assigned to either receive training or not receive training prior to this tutoring. However, there was no random assignment to a control group (5th graders who did not tutor), and therefore the study is not designed to identify the effect of tutoring on the tutor's knowledge. AbdulRaheem et al. (2017) randomly assigns a single classroom from one school to control and a single classroom from another school to a peer-tutoring treatment. This design makes it impossible to disentangle the treatment effect of peer tutoring from the teacher-, classroom-, or school- fixed effects.

Mitchell et al. (2016) does randomly assign 4th grade students to either a control condition or a condition where they tutor a 2nd grade student (with or without training). The teacher decided the student-tutor pairs based on personality matches, which in and

of itself is not a threat to internal validity. The researchers found that while the 2nd grade students benefitted from receiving tutoring, the 4th grade students did not benefit from providing tutoring. The biggest issue with this study is the sample size. With less than 15 students per treatment arm (43 4th grade students for three experimental conditions), the study is underpowered to detect modest treatment effects.

Two studies attempt to identify the impact of tutoring on learning by randomizing which subject a student tutors. Romero et al. (2022) studied cross-age tutoring within a primary school in Kenya. The sample size is large, and the outcome is the tutor's own grades. They find that the tutors had little impact on their own test scores from tutoring math (as opposed to tutoring English). Noteworthy here is that there is a 5-year gap between the tutor and student. This might be too large of a gap to expect the tutor's own math performance (5th grade) to improve from tutoring the student (1st grade), especially because the tutors are self-selected high performers. Fuchs and Malone (2021) assigned master's students in education to tutor either Math ($n = 25$) or English ($n = 17$) to elementary school students. However, the assignment was based on scheduling constraints rather than random assignment. They find large effects (0.5 to 0.75 SD) for an assessment on the math topic that was tutored (fractions).

2.2 Rationale for Video Creation as an Intervention

Given the effectiveness of learning by teaching in lab settings, we might wonder why this is not already universally used as a tool for human capital development in schools. While there have been calls for universal school-wide peer tutoring programs, few exist (Kraft and Falken, 2021).

One reason might be that only the type of person who chooses to teach would benefit, or what economists call "selection on gains." However, it might be the case that the benefits from teaching are universal, yet only a select few engage in the activity. For instance, it might be the case that while everyone would benefit from providing peer tutoring, only students near the top of the achievement distribution are asked to provide peer tutoring. If this is the case, then making the opportunity to provide peer tutoring

universal would help mitigate the achievement gap.

There are at least three barriers that could explain why peer tutoring programs are not widespread. 1) It is logistically difficult to coordinate peer tutoring sessions. Scheduling synchronous sessions in a way where most students have an opportunity to provide tutoring takes considerable effort, and this high transaction cost may not be worth the perceived benefit. 2) Programs designed to give peer tutors the opportunity to learn might not be ideal for the students receiving tutoring. The tutoring programs with the highest impact on students are ones where screened professionals provide the tutoring (Nickow et al., 2020). This implies that the opportunity cost for students receiving tutoring from an untrained peer might be too high to justify their participation. 3) Students might be unwilling to engage in an activity that reveals their academic ability to their peers. Prior work has shown that making academic effort publicly visible to students' peers can have an adverse impact if students highly value their social status (Bursztyn et al., 2019).

An intervention that overcomes these concerns is one where students create video explanations to a hypothetical peer as homework assignments. 1) The asynchronous nature of this task mitigates the need to coordinate schedules, making participation more feasible. 2) The risk that the recipient of tutoring might receive a lower quality experience than alternative uses of their time is no longer a concern. 3) Adverse peer effects are avoided because the audience is the teacher rather than peers. This intervention might not capture all of the benefits that synchronous peer tutoring might entail. For instance, if the back-and-forth interaction in peer tutoring leads to a substantial impact on the tutor, then this video creation intervention misses out on this benefit by only focusing on the "initial explanation" phase (Kobayashi, 2022). Additionally, if the tutor's motivation to put effort into their tutoring stems from caring about the student's outcome, then this intervention would not capture this benefit because the audience of the video is a hypothetical rather than an actual peer. Despite this lower potential benefit, the video creation intervention might still have a high benefit-cost ratio because of its low cost and potential for scalability.

3 Theoretical Framework

This section describes a model of "learning by teaching" that highlights mechanisms as well as potentially adverse affects. I also describe a model of peer tutoring, highlighting the tradeoffs from introducing AI tutoring, in part because of the learning loss from the tutors themselves if they are replaced with AI tutors.

3.1 A Model of Student Effort Choice

I adapt the worker effort framework in DellaVigna et al. (2022) to describe student i's optimal effort as a function of grade incentives, the cost of effort, and a social preference parameter A. The utility function to maximize is:

$$u(e_i) = p \cdot e_i - c_i(e_i) + A \cdot e_i$$

Where e_i is the amount of effort, the piece rate $p > 0$ is the incentive in the form of class grades, $c_i(\cdot)$ is the cost of effort for student i, and A is a social parameter that reflects how much you care about others (i.e. altruism) or their perception of you (i.e. social image). I assume that effort maps one-to-one to observable grade outcomes. A unique solution is guaranteed by assuming $c'() > 0$, $c''() > 0$, and $\lim_{e\to\infty} c'(e) = \infty$. The first order condition to this maximization problem is:

$$\begin{aligned} 0 &= p - c'(e_i) + A \\ c'(e_i) &= p + A \\ e_i &= (c^{-1})'(p + A) \end{aligned}$$

Now, suppose there are two types of effort that a student can choose to engage in. I denote e_1 as *receptive effort* and e_2 as *generative effort*, which signifies the type of learning the student is engaged in (passive versus active). Generative effort is effort that results in generating an output, such as an explanation, as opposed to receptive effort, which might involve reading an explanation. The piece-rates associated with receptive

and generative effort are p_1 and p_2 respectively, where $p_2 > p_1$. The cost functions associated with receptive effort and generative effort are $c_1(\cdot)$ and $c_2(\cdot)$ respectively, where $c_2(e) > c_1(e)$, $\forall e$. Students are aware that generative effort has a higher benefit, but also that generative effort has a higher cost. The cost functions for these efforts can vary for each student, which can result in different levels of investment of each type of effort among students that face the same grade incentives. For now, I assume that the costs and benefits of each type of effort type are independent from each other.

Suppose that a teacher is considering two alternative assignments where the student's work will be visible to the teacher. One assignment ($T = 0$) is a direct function of receptive effort e_1, and the other assignment ($T = 1$) is a direct function of generative effort e_2. This results in the following utility maximization problem:

$$\max_{e_1,e_2\geq 0} u(e_1, e_2) = p_1e_1 + p_2e_2 - c_1(e_1) - c_2(e_2) + A(e_1 \cdot \mathbb{1}_{T=0} + e_2 \cdot \mathbb{1}_{T=1})$$

For a student who is assigned task $T = 0$, where receptive effort is socially incentivized, the optimal level of e_1 effort is $e_1^* = (c_1^{-1})'(p_1 + A)$. For a student who is assigned $T = 1$, where generative effort is socially incentivized, the optimal level of e_2 effort is $e_2^* = (c_2^{-1})'(p_2 + A)$. Note that if no socially visible task is assigned, then the optimal effort values are $e_1^* = (c_1^{-1})'(p_1)$ and $e_2^* = (c_2^{-1})'(p_2)$.

For students who are assigned $T = 0$, there is no change in the optimal value of e_2 relative to receiving no socially incentivized task. The change in optimal effort e_1 relative to no socially incentivized task is $\Delta e_1 = (c_1^{-1})'(p_1 + A) - (c_1^{-1})'(p_1) > 0$.

For students who are assigned $T = 1$, there is no change in the optimal value of e_1 relative to receiving no socially incentivized task. The change in optimal effort e_2 relative to no socially incentivized task is $\Delta e_2 = (c_2^{-1})'(p_2 + A) - (c_2^{-1})'(p_2) > 0$.

In other words, relative to being assigned $T = 0$, a student being assigned $T = 1$ results in them exerting a higher level of e_2 effort, but a lower level of e_1 effort. This implies that a treatment that socially incentivizes generative effort relative to a counterfactual of incentivizing receptive effort could result in a *negative treatment effect* if Δe_2 is sufficiently smaller than Δe_1. That is, a student might be better off being assigned a receptive task

rather than a generative one if the receptive task results in a substantial increase in receptive effort, whereas the generative task results in a small increase in generative effort. This would depend both on the student's relative costs of generative and receptive effort, and also on the mapping between each type of effort and math skills (or utility) outcomes.

3.2 Effort Complementarity

The previous section assumes additive separability in the cost and benefit of the two effort types. Dropping this assumption implies that the two components of effort could be either substitutes or complements in production. Students might view the total amount of time spent on math as relatively inelastic (Cotton et al., 2020). If that is the case, then the two types of effort might be treated as substitutes. This would imply that being assigned a generative task such as creating videos might make you invest less effort in receptive tasks such as watching videos. On the other hand, a student may view these two types of effort as complementary. In this case, being assigned a task requiring generative effort would make them more likely to also engage in receptive effort. That is, being assigned to create a video would make them *more* likely to watch a math video as well.

3.3 Skill Transfer

Suppose the skill variable y has two components, y_1 and y_2 corresponding to *local* and *generalizable* skills respectively. y_1 refers to knowing how to do something exactly a certain way, whereas y_2 refers to one's ability to generalize how to do tasks similar to ones they know. The work in cognitive psychology implies that active learning effort (e_2) is a pre-requisite to develop generalizable (y_2) skills (Ahn et al., 1992; Boaler et al., 2022). This would imply that $\frac{\partial y_2}{\partial e_1} = 0$. This implies that the likelihood of the aforementioned negative treatment effect of being assigned to teach someone is much lower for generalizable skills, and also that the counterfactual passive task treatment would not have a positive treatment effect on the y_2 component of skill.

3.4 A Model of Peer Tutoring

Suppose a school has a population of N students, and of these τ have the ability to serve as peer tutors. Each tutor has the capacity to tutor g students. Suppose that $\tau + g\tau < N$, implying some students do not receive tutoring in equilibrium.

There is a distribution of "benefits from tutoring" amongst the $N - \tau$ students who are not peer tutors, and the way to select the $g\tau$ students who receive tutoring is simply by selecting those with the highest benefit.

Now, suppose that *AI Tutoring* is introduced, which allows for all remaining $(N - \tau - g\tau)$ students to receive tutoring. Assume this benefit is quantifiable. While this benefit may seem to come at no cost, there might be a substitution where some human tutoring amongst the $g\tau$ students gets switched to AI tutoring. Even if the quality of AI tutoring is the same from the students' perspective, the *tutors would lose out* on skill development. This negative effect should be factored in when estimating the total impact of AI tutoring.

Additionally, suppose that γ percent of the students who don't otherwise choose to get tutoring are "cheaters", where they would use the AI tutoring as a way to get answers to homework questions and reduce their cognitive load of studying at home as a result of having access to this AI tutoring. This decrease in their own learning would impact performance on in-school administered tests, and thus be quantifiable.

This implies that the introduction of AI tutoring would need to weigh the benefit of the students who now get tutoring that previously could not with the costs of 1) decreased learning opportunities for the peer tutors themselves, and 2) decreased skills for those who would now lessen their own at-home effort as a result of having access to AI tutoring.

This project conducts a "learning by teaching" intervention in a way that provides an estimate for cost (1) describe above, but also explores whether cost (2) could be mitigated by assigning an at-home task that is more difficult to "cheat" on, thus inducing more student effort relative to a passive task where AI could do all of the work.

4 Experimental Design

4.1 Sample and Recruitment

4.1.1 District and School Recruitment

During December 2023 and January 2024, I recruited school districts via emails to the superintendents. Emails were sent to all districts with at least 5,000 students in Illinois, Wisconsin, Iowa, Michigan, and Ohio. I sent 371 emails and received 94 replies, of which 29 indicated interest in learning more. Of these, 13 agreed to participate and shared information about the study with middle and high school math teachers. The most common reasons for districts not being able to participate were scheduling conflicts with other school improvement initiatives and extensive research review processes that would not fit in our timeline for the current school year.

The characteristics of participating school districts varied widely. The average percentage of students on Free or Reduced Lunch was 41%, and ranged from 8% to 93%. The high school graduation rates varied from 72% to 98%, with an average of 88%.

4.1.2 Teacher Recruitment

Next, math teachers at participating school districts were sent an interest form with information about the study. Teachers who taught at least two periods (sections) of the same math class were eligible. Teachers received $500 for their effort in helping implement the study. Recruitment took place on a rolling basis between February 5th and March 8th, 2024. In total, 47 teachers filled out the interest form, of which 41 teachers chose to proceed with the study after learning more. Of these 41, 28 were high school teachers and 13 were middle school teachers.

Finally, teachers distributed parent permission forms to all students in their classes. The permission form indicated whether the students were allowed to take a short survey with demographic information and their math background, as well as whether the teacher could share identifiable data with the researchers (including student-generated videos). Students whose parents did not give permission still participated in the tasks according

to their class period, but they did not take the survey and only de-identified data was shared for these students.

4.2 Randomization

Random assignment was done at the class period level, stratified by teacher. Individual-level randomization is difficult to implement because the treatment involves assignments given at the classroom level. SUTVA violations (List, 2024a) might especially be a concern for individual-level randomization in this context.

Teachers were eligible to enroll their classes in the study if they taught at least two periods (sections) of the same math class. If they taught exactly two periods, then one was randomized to the "Video Creation" treatment and the other to the "Passive Task" treatment (see Section 4.3 for a description of these treatments). If they taught three or more periods, then they were randomized to 1) Video Creation, 2) Passive Task, and 3) Control as long as each period had at least 15 students. If there were fewer than 15 students in a period, then the smaller two periods were treated as a single unit for the purpose of randomization (See the Pre-Analysis Plan[2] for the full description of the relative class size cutoff algorithm used to determine random assignment). This was done to ensure that there was sufficient power to detect differences between the Passive Task and Video Creation treatments, with the pure control classrooms being an additional comparison arm in case enough teachers with more than two periods were recruited.

Randomization occurred after all students in a teacher's class periods took a pre-test consisting of consisting of 15 multiple-choice questions taken from either the PSAT, ACT, SAT, or grade-level state standardized math test. The questions were chosen based on the topics the teachers indicated they planned to cover between March and May 2024. This in-class pre-test was 25 minutes, and teachers administered these on a rolling basis between March 4th and March 29th. Randomization occurred on a rolling basis after all the pre-tests were completed for each teacher.

[2] https://www.socialscienceregistry.org/trials/11884

4.3 Treatment Description

Each week, teachers assigned students in their Passive Task classroom a PSAT (or in some cases ACT, SAT, or grade-level state test) question to complete for homework via Google Form. Teachers had some discretion on whether to skip weeks or whether to have two tasks in the same week depending on their preference and alignment with their curriculum. Teachers gave input for the topics that they preferred that the tasks covered, and the researcher selected questions from a pool of aforementioned standardized test questions. Teachers also had discretion on how they incentivized the task, with the recommendation being that it be for completion credit. While most teachers followed this, some were unable to do so because of school-wide policies that disallowed the use of homework for credit.

Below the Task question on each weekly Google Form, students were provided with a link to a "help" video on YouTube that explained the solution to a similar problem. These videos were selected by the researcher. In some cases where no videos were easily available, a video was created for the purpose of this study. The videos were shared on the Google Form in a bitly[3] link, which allowed the researcher to measure the total number of clicks on that link. While it was not possible to measure whether an individual student clicked on the link, the bitly link was unique for each classroom for each task. As pre-registered, I use this information to determine the probability of clicking on the link for each student. An example of this is shown in Figure 1.

In the Video Creation classrooms, the teachers assigned the same weekly task on the same schedule as their Passive Task classrooms. The only difference was that their Google Form asked them to submit a video explaining the solution. Teachers chose the exact mechanism by which students submitted their video depending on the teacher's preference. Most teachers used FlipGrid to collect student videos. Some teachers used Google Form directly, or asked students to send the video via email or uploading to a Google Drive folder. The Google Form for the video creation classrooms also had the same YouTube help video, but with a different bitly link that allowed for a separate

[3] https://bitly.com/

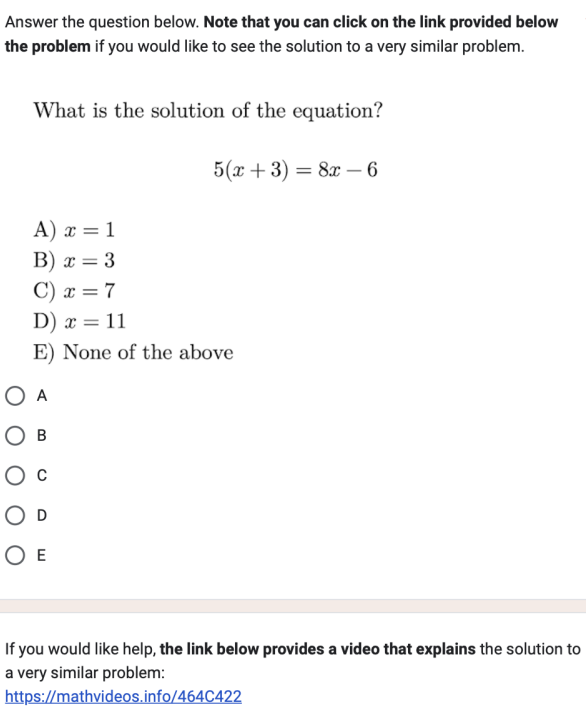


Figure 1: Example of Passive Task

identification of the number of total clicks in the Passive Task versus Video Creation classrooms. An example of this is shown in Figure 2.

Three teachers assigned this as an in-class activity because their school did not allow for assigning homework. As pre-registered, these three teachers were dropped after it was determined that an exception could not be made for the purpose of this study.

4.4 Outcomes

Two outcomes I use to estimate math skills are: 1) Semester 2 math class grade, and 2) A 15-question post-test, consisting primarily of PSAT and ACT questions, on the topics related to the Tasks for that teacher's classes. This 15-question test included 6 questions that were taken from the exact assigned Task questions, and 9 additional questions that were not exact task questions but covered the same topics.

The two secondary outcomes that shed light on mechanisms are 1) the students' self-reported math confidence, and 2) the students' likelihood of viewing the "help" video. The students' likelihood of viewing the "help" video is a proxy for the amount of effort they put into the task, and was measured as the number of total clicks on their class

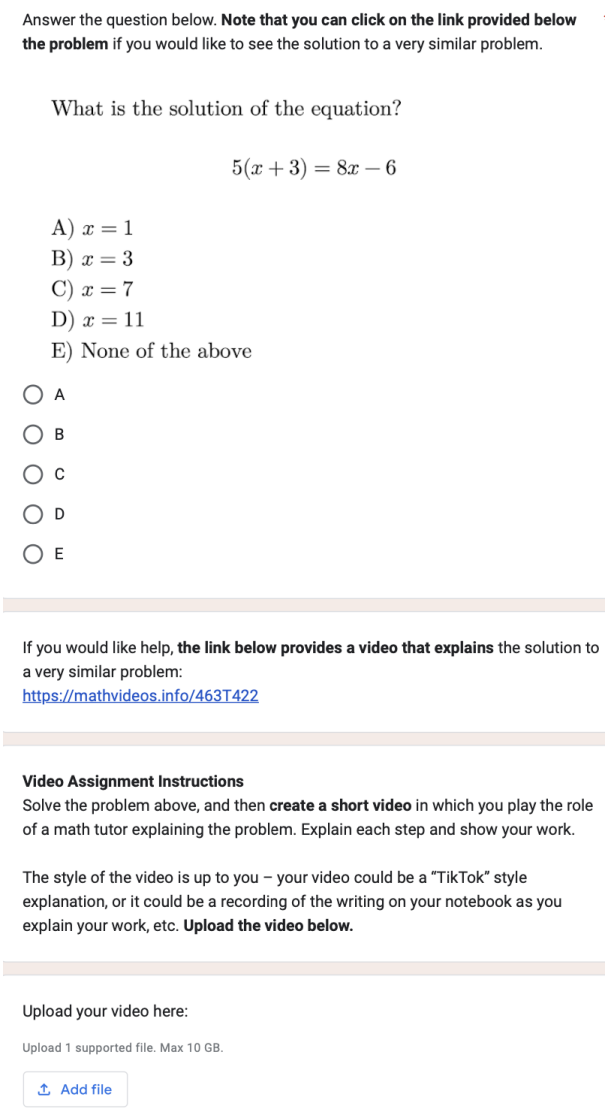


Figure 2: Example of Video Creation Task

period's link divided by the total number of students in that period.

4.5 Covariates

The teachers administered the pre-test along with a survey that asked about students' age, gender, race, math confidence, growth mindset, as well as past experience with tutoring and video creation. Only students whose parents gave active consent and identifiable data-sharing permission took the survey.

5 Results

5.1 Descriptive Results

Table 1 shows the descriptive statistics of baseline characteristics along with a balance test of the difference in the mean values for each treatment group. The first column represents the mean value of the control classroom students for each variable, with the standard deviation of each variable in parenthesis below it. The second column shows the regression coefficient of the difference between the value of the Passive Task classroom students and that of the control classroom students, with clustered standard errors in parenthesis below along with the p-value of this coefficient. The third column is the same as the second, but for the Video Creation classrooms. Finally, the fourth column is the difference in the mean value of the students in the Passive Task and Video Creation classrooms, with the standard error and p-value below each mean difference.

The table shows us that there is no significant difference (pairwise) between any of the experimental arms. About half of the students were women, 10% were black, and 23% Hispanic.

5.1.1 Attrition

Some students did not take the post-test at the end of the year, primarily due to absences on the day the post-test was administered. While some teachers were able to administer make-up tests for students who were absent on the post-test day, others were constrained by final exam scheduling and chronic absenteeism from students who missed both days. Given that the post-test is the primary outcome variable, I consider a student to have dropped out of the study if they did not take the post-test.

Using this definition of "drop out", I measure the attrition rate as the fraction of students who dropped out. The attrition rate was 13.6%, 13.0%, and 13.9% for the control, Passive Task, and Video Creation groups respectively. The first row of Table 1 shows that these rates are not (pairwise) significantly different. I use baseline covariates to conduct a selective attrition test and determinants of attrition test as per List (2024a).

I further use baseline outcomes to conduct the attrition test as per Ghanem et al. (2023) in the appendix.

5.1.2 Compliance

Given the additional effort required to create a video (Video Creation treatment) relative to simply completing the google form (Passive Task treatment), we would expect the task completion rate for the Video Creation students to be lower than that of the Passive Task students. I define the task "completion rate" for each student as a continuous variable between 0 and 1 that indicates the proportion of all 8 tasks that student completed[4]. I find that the overall task completion rate is 83.7% for the Passive Task students, and 62.4% for the Video Creation students. When subsetting on only the students who did not attrite from the study, these numbers are 87.8% for the Passive Task students and 66.5% for the Video Creation students.

I additionally pre-registerd a binary "compliance" variable for each student that has a value of 1 if the student completed at least half of the assigned tasks (i.e. if the student completed 4 or more tasks if their teacher assigned all 8 tasks), and 0 otherwise. With this definition, the compliance rate for the Passive Task students is 90.1% and that of the Video Creation students is 69.8%. When subsetting to the students who did not drop out of the study, the compliance rate is 94.7% for the Passive Task students, and 74.6% for the Video Creation students. While I use these compliance rates to estimate a Local Average Treatment Effect (LATE) in the appendix, the primary results I report in this study are Intent-to-Treat (ITT) estimates because they are more policy-relevant.

5.2 Impact on Math Skills

The following regression was used to estimate the treatment effect on math skills. This regression was pre-registered in the Pre-Analysis Plan[5]:

[4] a few teachers had to drop tasks and had fewer than 8. In these instances, the completion rate is the total number of tasks that student completed divided by the total number assigned by the teacher.

[5] https://www.socialscienceregistry.org/trials/11884

Table 1: Descriptive Statistics and Balance Tests

	Control (N=523)	Passive Task (N=1,031)	Video Creation (N=969)	Difference (Video-Passive)
Attrition Rate	0.136	−0.0058	0.0036	0.0093
	(SD: 0.343)	(0.041)	(0.043)	(0.035)
		$p = 0.889$	$p = 0.934$	$p = 0.787$
Female	0.533	−0.0111	0.0015	0.0127
	(SD: 0.499)	(0.032)	(0.036)	(0.029)
		$p = 0.728$	$p = 0.966$	$p = 0.660$
Age	14.4	0.180	0.234	0.054
	(SD: 1.74)	(0.441)	(0.442)	(0.387)
		$p = 0.684$	$p = 0.597$	$p = 0.889$
Black	0.103	−0.040	−0.038	0.002
	(SD: 0.305)	(0.025)	(0.025)	(0.016)
		$p = 0.109$	$p = 0.135$	$p = 0.901$
Hispanic	0.231	−0.004	0.003	0.007
	(SD: 0.422)	(0.056)	(0.060)	(0.052)
		$p = 0.938$	$p = 0.962$	$p = 0.888$
White	0.369	0.084	0.071	−0.013
	(SD: 0.483)	(0.058)	(0.058)	(0.050)
		$p = 0.149$	$p = 0.224$	$p = 0.791$
Class Size	22.1	−0.263	−0.356	−0.094
	(SD: 4.36)	(1.07)	(1.09)	(1.00)
		$p = 0.806$	$p = 0.745$	$p = 0.926$
Baseline Score	6.16	0.427	0.475	0.048
(out of 15)	(SD: 3.34)	(0.512)	(0.521)	(0.443)
		$p = 0.406$	$p = 0.364$	$p = 0.914$
Baseline Grade	0.105	−0.156	−0.105	0.050
(Standardized)	(SD: 0.941)	(0.274)	(0.261)	(0.235)
		$p = 0.572$	$p = 0.688$	$p = 0.832$
Baseline Confidence	6.63	0.012	0.127	0.115
(10-point scale)	(SD: 1.932)	(0.130)	(0.134)	(0.135)
		$p = 0.925$	$p = 0.342$	$p = 0.394$

Note: The first column represents the mean value of the control classroom students for each variable, with the standard deviation of each variable in parenthesis below it. The second column shows the regression coefficient of the difference between the value of the Passive Task classroom students and that of the control classroom students, with clustered standard errors in parenthesis below along with the p-value of this coefficient. The third column is the same as the second, but for the Video Creation classrooms. Finally, the fourth column is the difference in the mean value of the students in the Passive Task and Video Creation classroom, with the standard error and p-value below each mean difference. *** p<0.01, ** p<0.05, * p<0.1.

$$y_{i,t} = \beta_0 + \beta_1 T_i + \beta_2 C_i + \beta_3 y_{i,t-1} + \tau_j + \varepsilon_i$$

Here, $y_{i,t}$ is the value of outcome at the end of the experiment, $y_{i,t-1}$ is the baseline level of outcome, T_i is a binary indicator where 1 = either Video Creation or Passive Task classroom and 0 otherwise, C_i is a binary indicator where 1 = Video Creation task classroom and 0 otherwise, τ_j is the teacher fixed effect for teacher j, and ε_i is the error term. Standard errors of coefficients are clustered at the classroom level, because that was the level at which the treatment was assigned.

5.2.1 Math Grades

Table 2 shows the treatment effect on the overall class grade, normalized for each teacher and measured in standard deviation units. Class grades were measured either on a percentage scale (0-100) or on a 4-point scale. Class grades were unavailable for some teachers depending on the Data Use Agreement with that district. Additionally, a few middle school teachers from one district had no overall numerical course grades available due to district policy. With these restrictions, normalized course grades were available for 1,603 students. Overall, I find that the video creation treatment led to a statistically significant 0.068 Standard Deviation increase in math class grade relative to the Passive Task treatment.

5.2.2 Post-Test Scores

Table 3 shows the main treatment effect on the PSAT post-test. The first column shows the results for the full sample, indicating that the Passive Task did not lead to a significant improvement in the Post Test score overall, whereas the Video Creation treatment led to a significant impact compared to the Passive Task.

For the first four teachers that participated, the post-test questions were chosen by the teachers. These four teachers graded the post-test themselves and sent final scores to the researcher, and a breakdown of how each student performed on each question is unavailable. The second column excludes these four teachers, showing the treatment effect for only the subset of students where the post-test grading was automated on Google Form (Quiz Mode) and student performance on each question is available.

Table 2: Treatment Effect on Math Grades

	(1) Math Class Grade
Assigned Either Task (Passive or Video Creation)	0.00288 (0.0553)
Video Creation	0.0680** (0.0315)
Baseline Grade Controlled	Yes
Teacher Fixed Effects	Yes
Observations	1,603

Note: The Assigned Either Task variable in the first row is binary with a value of 0 for students in the control group, and a value of 1 for students in either the Passive Task or Video Creation group. The Video Creation variable in the second row is binary with a value of 0 for students in either the control or Passive Task group, and a value of 1 for students in the Video Creation group. The outcome is the student's overall math course grade for semester 2, measured in standard deviation units. Clustered standard errors (at the classroom-level) are in parentheses.
*** p<0.01, ** p<0.05, * p<0.1.

The 15-question post-test included 6 questions that were *exact* intervention questions completed by the Passive Task and Video Creation classrooms (Task Qs), and 9 questions that covered the same topics, but had not been assigned as tasks (Novel Qs). The treatment effect on the subset of Task and Novel post-test questions is shown in columns (3) and (4) respectively, using the same sample of students as in column (2). The third column shows that there was a significant impact of the Passive Task on the exact Task questions, and the fourth column shows that there was no significant impact of the Passive Task on the Novel questions. The Video Creation treatment has a significant impact relative to the Passive Task for both types of questions.

5.2.3 Knowledge Generalization

For instruction to be broadly useful, it is important for students to be able to generalize what they learn and apply it to other contexts, rather than learning in a way that is akin to rote memorization (Bonwell and Eison, 1991). Cognitive psychologists refer to this type of generalization as *knowledge transfer* (Samat et al., 2019). One aspect of this transfer process depends on rote memory of the initially learned concept, and another

Table 3: Treatment Effect on Post-Test

	(1) Post-Test	(2) Post-Test
Assigned Either Task (Passive or Video Creation)	0.113*	0.124*
	(0.0581)	(0.0625)
Video Creation	0.216***	0.167***
	(0.0414)	(0.0429)
Baseline Score Controlled	Yes	Yes
Teacher Fixed Effects	Yes	Yes
Observations	2,130	1,869

Note: The Assigned Either Task variable in the first row is binary with a value of 0 for students in the control group, and a value of 1 for students in either the Passive Task or Video Creation group. The Video Creations variable in the second row is binary with a value of 0 for students in either the control or Passive Task group, and a value of 1 for students in the Video Creation group. The outcome in Columns 1 and 2 is the 15-question post-test consisting of standardized test questions that were relevant to the teacher's curriculum. Column 1 is the full sample, whereas Column 2 only contains teachers for whom a breakdown of the type of post-test question is available. All outcomes are in standard deviation units. Clustered standard errors (at the classroom-level) are in parentheses.
*** p<0.01, ** p<0.05, * p<0.1.

depends on understanding the concept well enough to be able to create a map between the previously learned concept and a new problem (Ahn et al., 1992).

Table 4 shows the impact of the Passive Task and Video Creation treatments. These regressions are the same as those in columns (2)-(4) of Table 3, except that Table 4 shows the treatment effect for each treatment relative to the control group. Note that the treatment effect coefficients in column (1) are a weighted average of those in columns (2) and (3). We see that the Passive Task leads to a significant impact on the exact Task Questions relative to the control group, but no significant impact on Novel Questions. On the other hand, the Video Creation treatment leads to a significant impact on both exact Task Questions and on Novel Questions. This highlights that while Passive Tasks might be effective at helping students learn in a way that is akin to rote memorization, the Video Creation treatment is effective at helping students learn in a way that allows for knowledge transfer.

Table 4: Treatment Effect on Post-Test relative to Control

	(1) Post-Test	(2) Task Qs	(3) Novel Qs
Passive Task	0.124**	0.183***	0.0563
	(0.0625)	(0.0664)	(0.0579)
Video Creation	0.291***	0.354***	0.187***
	(0.0625)	(0.0670)	(0.0574)
Baseline Score Controlled	Yes	Yes	Yes
Teacher Fixed Effects	Yes	Yes	Yes
Observations	1,869	1,869	1,869

Note: The Passive Tasks variable in the first row is binary with a value of 1 for students in Passive Tasks group and 0 otherwise. The Video Creations variable in the second row is binary with a value of 1 for students in the Video Creations group and 0 otherwise. The outcome in Column 1 is the 15-question post-test consisting of standardized test questions that were relevant to the teacher's curriculum. The outcome in Column 2 is the student's score on the subset of the post-test questions that were exact treatment tasks, while Column 3's outcome is the student's score on the remaining questions on the post-test. All outcomes are in standard deviation units. Clustered standard errors (at the classroom-level) are in parentheses.
*** p<0.01, ** p<0.05, * p<0.1.

6 Mechanisms

Two potential mechanisms that could explain this treatment effect are 1) student effort and 2) student confidence. I pre-register that student effort would be proxied by the likelihood that students click on the "Help Video" link provided to them for each task, and that student confidence is estimated by a self-reported survey question administered right before the pre-test and post-test.

6.1 Student Effort

In each task, students in both the Passive Task and Video Creation classrooms were provided a link to a "Help" video on YouTube. While the video was the same for the Passive Task and Video Creation classrooms for any given teacher-task dyad, the links in the Google Forms were embedded in separate bitly links for each classroom. The bitly link allows for tracking of the total number of times that a link has been clicked. While it cannot be ascertained whether a given student clicked on a link, or even whether the

same link was clicked multiple times by the same person, this total provides an estimate of the likelihood that a student in a given classroom clicked on the link.

I estimate V_k as the Viewing likelihood for any given classroom. This is computed by taking the total number of clicks for a given classroom across all tasks, and then dividing it by the total number of students in the classroom, and then again dividing by the number of tasks. I estimate the following regression, where the unit of observation is a classroom:

$$V_k = \alpha_0 + \alpha_1 C_k + u_k$$

Here, V_k is the Viewing likelihood, C_k is a binary indicator where 1 = Video Creation class period and 0 otherwise, and u_k = error term. Note that control classrooms are excluded from this by construction because not having any tasks means that a Viewing likelihood cannot be estimated. The results for this regression are shown in Table 5.

Table 5: Impact of Video Creation Treatment on Click Likelihood

	(1) Click Likelihood
Video Creation	0.143*** (0.0341)
Constant	0.110*** (0.0239)
Observations	100
R-squared	0.153

Note: Each unit of observation in this regression is a classroom. Video Creation is a binary variable with a value of 1 if that classroom was assigned to the Video Creation treatment, and 0 if assigned to the Passive Task treatment. The outcome variable is the probability of a student in that classroom clicking on the help video resource for any given task. The standard error of the coefficients are in parenthesis.
*** p<0.01, ** p<0.05, * p<0.1.

The average likelihood of clicking on a video is 11% for students in Passive Task classrooms, and is over twice that (25%) for studentss in Video Creation classrooms. This result contributes to disentangling the opposing forces on student effort described in section 3. Here, a student who is inelastic with regards to the total amount of time invested in math would be *less* likely to click on this help video link when assigned to

the Video Creation treatment relative to the Passive Task treatment. On the other hand, a student who puts a high enough value on the video viewer's perception of them and treats passive and active effort as complementary would be more likely to click on the link. These results highlight that the video creation task might be sufficient enough to induce student motivation to apply more effort, which is generally hard to move (Burgess et al., 2021; Cotton et al., 2020). Additionally, this shows that a mechanism through which teaching can help people learn is that teaching induces people to put in more time preparing, implying that lab studies that hold the total amount of preparation time constant are underestimating the true general equilibrium impact of "learning by teaching" interventions.

6.2 Student Confidence

One possible mechanism through which this intervention would improve student skills is through improving their confidence in their math skills, and this confidence would lead to a skill-improving change in "mindset" (Yeager and Dweck, 2012). If this were the case, we would expect the treatment students to end the semester with higher changes in confidence than the control group. Another possibility is that as the students in the Video Creation classrooms explained math in their videos, their lack of initial understanding became clearer to them. This is what psychologists call "Illusion of Explanatory Depth" (Rozenblit and Keil, 2002). If this is prevalent, then we would expect a zero or negative change in math confidence among the treatment group. Table 6 shows the regression results of the treatment on math confidence at the end of the semester, normalized for each teacher and measured in standard deviation units.

I find that the video creation treatment led to an insignificant change in math confidence. One possibility is that both mechanisms are at play, and canceling each other out, where math confidence is initially lowered when students explain math and get a realistic sense of their skills, but then sustained explanations over time makes them improve their overall math confidence as they gain experience with explaining math. Future iterations of this study could have a weekly confidence check-in question to explore the dynamics

Table 6: Impact of Video Creation Treatment on Math Confidence

	(1) Math Confidence
Assigned Either Task (Passive or Video Creation)	0.0152 (0.0430)
Video Creation	-0.0290 (0.0313)
Baseline Confidence	0.665*** (0.0182)
Teacher Fixed Effects	Yes
Observations	1,779

Note: The Assigned Either Task variable in the first row is binary with a value of 0 for students in the control group, and a value of 1 for students in either the Passive Task or Video Creation group. The Video Creation variable in the second row is binary with a value of 0 for students in either the control or Passive Task group, and a value of 1 for students in the Video Creation group. The outcome is the student's overall self-reported level of confidence in their math skills (on a scale from 1 to 10), measured in standard deviation units. Clustered standard errors (at the classroom-level) are in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

of confidence as students continue to produce videos.

7 Distribution of Treatment Effects

One feature of this experimental design is that it allows for the estimation of an individual treatment effect for each teacher. Class sections (e.g. 3rd versus 4th period) are typically randomly assigned in schools, and this study ensures that each teacher has at least one class period that creates videos and one period that does the passive task. While there might still be classroom peer effects at play, and while the sample size for any given teacher is typically only between 30-90 students, we can think of this study as having conducted a set of 41 small RCTs, one for each teacher.

I estimate the treatment effect of the Video Creation treatment relative to the Passive Task treatment on the PSAT post-test for each teacher individually, controlling for the baseline PSAT pre-test. This allows for a distribution of treatment effects which is shown in Figure 3. From this, we can see that approximately 80% of teachers had a positive

treatment effect, and 20% had a negative treatment effect.

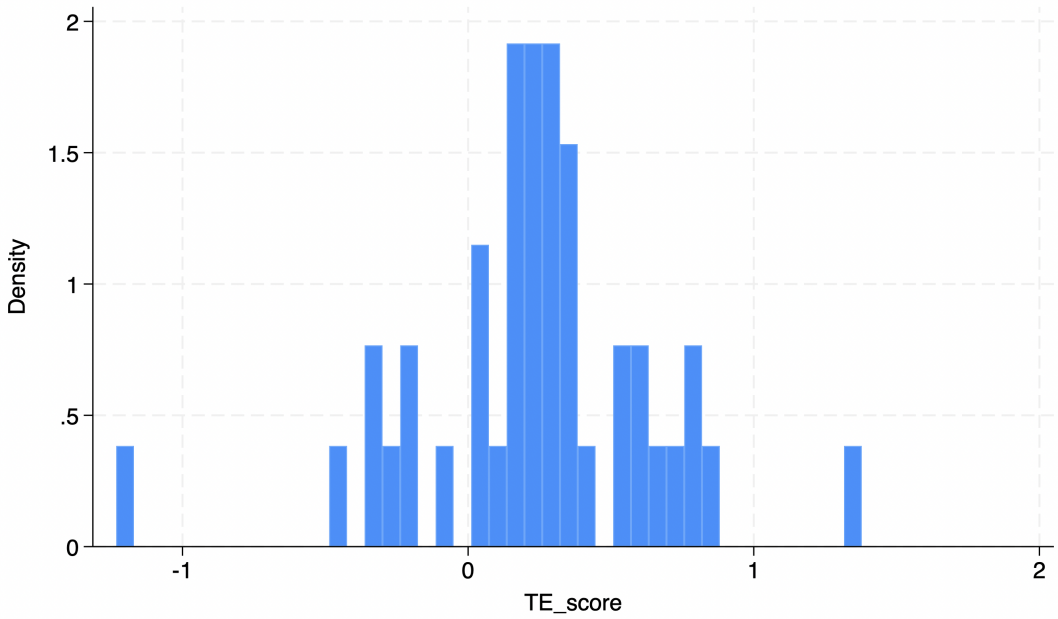


Figure 3: Distribution of Treatment Effects

As described in Section 3, a negative treatment effect is possible in this design for students who would choose to complete their assigned task if they are in the Passive Task classroom, but not if they are in the Video Creation classroom. The distribution in Figure 3 shows that some teachers had negative treatment effects, including one that had a large (greater than 1 SD) negative treatment effect. I assess whether the difference in task completion rate between a teacher's Passive Task and Video Creation classrooms predicts a teacher's treatment effect size.

Figure 4 shows the relationship between compliance and treatment effect. The x-axis shows a teacher's task completion rate for their Video Creation classroom minus that of their Passive Task classroom, and the y-axis shows the treatment effect size. Note that most teachers have a negative value for the x-axis variable because the task completion rate for their Passive Task classrooms was greater than that of their Video Creation classroom. We visually see a positive relationship, which is confirmed by a statistically significant positive slope as shown in Table 7.

This implies that teachers who were able to induce high compliance in their video creation classroom were able to achieve a high treatment effect. This relationship is not

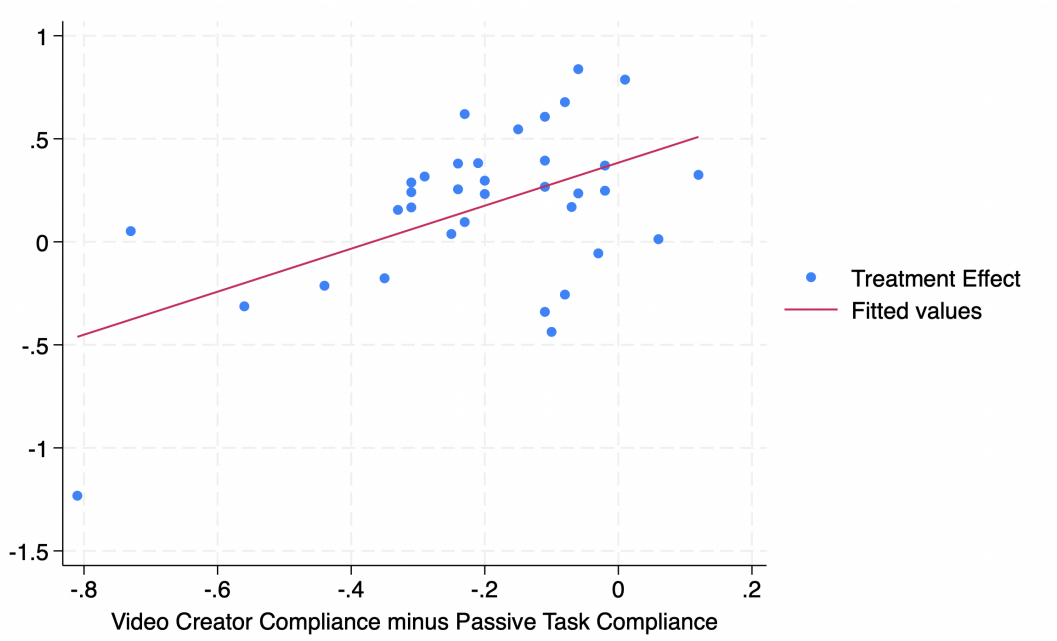


Figure 4: Relationship Between Compliance and Treatment Effect

Table 7: OLS regression of Compliance on Treatment Effect

	(1) Treatment Effect
Difference in Compliance Rate (Video Creation minus Passive Task)	1.044*** (0.290)
Observations	35
R-squared	0.281

Note: The independent variable in this regression is the compliance rate of that teacher's Video Creation class section minus the compliance rate of that teacher's Passive Task class section. The outcome variable is the treatment effect, measured in standard deviation units, for each teacher. The standard error of the coefficient is in parenthesis.
*** p<0.01, ** p<0.05, * p<0.1.

causal, so it is unclear whether this relationship is because of the high compliance rate or because of some other characteristic about the teachers who managed to induce a high compliance rate. Future work could see whether experimentally inducing higher compliance, for example by having these assignments count for a high fraction of the overall course grade, leads to a higher treatment effect.

8 Conclusion

One goal of social science research is to estimate the impact of interventions that have the potential to be implemented widely. A pilot version of this intervention involved randomly assigning some students to a treatment where they tutored a student in a grade below them. Even when measures were taken to ensure that the tutor had higher baseline skills than the student, the compliance rate was exceedingly low. High financial incentives were not enough to induce a sufficiently large compliance rate to detect an Intent-to-Treat effect. There were additional costs involved in logistically ensuring that students had a space (even if it was virtual, a shared link) to meet. The nature of the intervention tested in this project was a result of attempting to ensure that this paper tested a version of a program that we may expect when implemented on a larger scale, or "policy-based" evidence as per List (2024b). The direct implementation cost of this project was approximately $10 per student.[6] Given the treatment effects in this study, the benefit-cost ratio is estimated to be on the high end of other successfully implemented educational interventions (Guryan et al., 2023; Kline and Walters, 2016).

However, as this intervention and others like it are scaled, we should be mindful of the distribution of treatment effects. In particular, if some students have a high probability of a negative treatment effect, then scaling should be done with caution. The treatment effect in this study was consistently positive for teachers who had a high compliance rate of video creation relative to the passive task. If a teacher notices that their compliance rate is low and cannot do much to change it, then having the non-compliers engage in the passive task might be a way to maximize societal welfare. Of course, there is a moral hazard if students know that they will be assigned a less challenging task if they simply refuse to participate for the first few assignments. Future work in this area can analyze the student-generated videos to see if there are characteristics of created videos that are predicted to have higher treatment effects[7]

[6] $25,000 cost, primarily on teacher incentives, for approximately 2,500 students

[7] The data agreements made with districts in this study did not allow for matching videos with test scores and did not allow researcher access to videos from enough classrooms.

References

AbdulRaheem, Y., Yusuf, H. T., and Odutayo, A. O. (2017). Effect of peer tutoring on students' academic performance in economics in ilorin south, nigeria. *Journal of Peer Learning*, 10(1):95–102.

Ahn, W.-k., Brewer, W. F., and Mooney, R. J. (1992). Schema acquisition from a single example. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(2):391.

Arrow, K. J. (1962). The economic implications of learning by doing. *The review of economic studies*, 29(3):155–173.

Berlinski, S. and Busso, M. (2017). Challenges in educational reform: An experiment on active learning in mathematics. *Economics Letters*, 156:172–175.

Boaler, J., Brown, K., LaMar, T., Leshin, M., and Selbach-Allen, M. (2022). Infusing mindset through mathematical problem solving and collaboration: Studying the impact of a short college intervention. *Education Sciences*, 12(10):694.

Bonwell, C. C. and Eison, J. A. (1991). *Active learning: Creating excitement in the classroom. 1991 ASHE-ERIC higher education reports.* ERIC.

Bordalo, P., Coffman, K., Gennaioli, N., Schwerter, F., and Shleifer, A. (2021). Memory and representativeness. *Psychological Review*, 128(1):71.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2020). Memory, attention, and choice. *The Quarterly journal of economics*, 135(3):1399–1442.

Brownback, A. and Sadoff, S. (2020). Improving college instruction through incentives. *Journal of Political Economy*, 128(8):2925–2972.

Burgess, S., Metcalfe, R., and Sadoff, S. (2021). Understanding the response to financial and non-financial incentives in education: Field experimental evidence using high-stakes assessments. *Economics of Education Review*, 85:102195.

Bursztyn, L., Egorov, G., and Jensen, R. (2019). Cool to be smart or smart to be cool? understanding peer pressure in education. *The Review of Economic Studies*, 86(4):1487–1526.

Conlon, J. J., Mani, M., Rao, G., Ridley, M. W., and Schilbach, F. (2022). Not learning from others. Technical report, National Bureau of Economic Research.

Cotton, C., Hickman, B. R., List, J. A., Price, J., and Roy, S. (2020). Disentangling motivation and study productivity as drivers of adolescent human capital formation: Evidence from a field experiment and structural analysis. Technical report, National Bureau of Economic Research.

De Backer, L., Van Keer, H., and Valcke, M. (2012). Exploring the potential impact of reciprocal peer tutoring on higher education students' metacognitive knowledge and regulation. *Instructional science*, 40(3):559–588.

DellaVigna, S., List, J. A., Malmendier, U., and Rao, G. (2022). Estimating social preferences and gift exchange at work. *American Economic Review*, 112(3):1038–1074.

Eskreis-Winkler, L., Milkman, K. L., Gromet, D. M., and Duckworth, A. L. (2019). A large-scale field experiment shows giving advice improves academic outcomes for the advisor. *Proceedings of the national academy of sciences*, 116(30):14808–14810.

Fiorella, L. and Mayer, R. E. (2013). The relative benefits of learning by teaching and teaching expectancy. *Contemporary Educational Psychology*, 38(4):281–288.

Freeman, S., Eddy, S. L., McDonough, M., Smith, M. K., Okoroafor, N., Jordt, H., and Wenderoth, M. P. (2014). Active learning increases student performance in science, engineering, and mathematics. *Proceedings of the national academy of sciences*, 111(23):8410–8415.

Fryer, R. G. (2017). The production of human capital in developed countries: Evidence from 196 randomized field experiments. In *Handbook of economic field experiments*, volume 2, pages 95–322. Elsevier.

Fuchs, L. S. and Malone, A. S. (2021). Can teaching fractions improve teachers' fraction understanding? insights from a causal-comparative study. *the elementary school journal*, 121(4):656–673.

Ghanem, D., Hirshleifer, S., and Ortiz-Becerra, K. (2023). Testing attrition bias in field experiments. *Journal of Human Resources*.

Gneezy, U. and List, J. A. (2013). *The why axis: Hidden motives and the undiscovered economics of everyday life*. Random House.

Gneezy, U., List, J. A., Livingston, J. A., Qin, X., Sadoff, S., and Xu, Y. (2019). Measuring success in education: The role of effort on the test itself. *American Economic Review: Insights*, 1(3):291–308.

Greene, I., Tiernan, A. M., and Holloway, J. (2018). Cross-age peer tutoring and fluency-based instruction to achieve fluency with mathematics computation skills: A randomized controlled trial. *Journal of Behavioral Education*, 27(2):145–171.

Guerrero, T. A. and Wiley, J. (2021). Expecting to teach affects learning during study of expository texts. *Journal of Educational Psychology*.

Guryan, J., Ludwig, J., Bhatt, M. P., Cook, P. J., Davis, J. M., Dodge, K., Farkas, G., Fryer Jr, R. G., Mayer, S., Pollack, H., et al. (2023). Not too late: Improving academic outcomes among adolescents. *American Economic Review*, 113(3):738–765.

King, A., Staffieri, A., and Adelgais, A. (1998). Mutual peer tutoring: Effects of structuring tutorial interaction to scaffold peer learning. *Journal of Educational Psychology*, 90(1):134.

Kline, P. and Walters, C. R. (2016). Evaluating public programs with close substitutes: The case of head start. *The Quarterly Journal of Economics*, 131(4):1795–1848.

Kobayashi, K. (2022). Learning by teaching face-to-face: the contributions of preparing-to-teach, initial-explanation, and interaction phases. *European Journal of Psychology of Education*, 37(2):551–566.

Kraft, M. A. and Falken, G. T. (2021). A blueprint for scaling tutoring and mentoring across public schools. *AERA Open*, 7(1):1–21.

List, J. A. (2022). *The voltage effect: How to make good ideas great and great ideas scale.* Crown Currency.

List, J. A. (2024a). *Experimental Economics: Theory and Practice.* University of Chicago Press.

List, J. A. (2024b). Optimally generate policy-based evidence before scaling. *Nature*, 626(7999):491–499.

Markant, D., Ruggeri, A., Gureckis, T., and Xu, F. (2016). Enhanced memory as a common effect of active learning. mind, brain, and education, 10 (3), 142-152.

Mitchell, R. J., Morrison, T. G., Feinauer, E., Wilcox, B., and Black, S. (2016). Effects of fourth and second graders' cross-age tutoring on students' spelling. *Reading Psychology*, 37(1):147–166.

Nickow, A., Oreopoulos, P., and Quan, V. (2020). The impressive effects of tutoring on prek-12 learning: A systematic review and meta-analysis of the experimental evidence. working paper 27476. *National Bureau of Economic Research.*

Robinson, C. D., Bisht, B., and Loeb, S. (2022). The inequity of opt-in educational resources and an intervention to increase equitable access. *Annenberg Institute at Brown University EdWorkingPaper*, (22-654).

Romero, M., Chen, L., and Magari, N. (2022). Cross-age tutoring: Experimental evidence from kenya. *Economic Development and Cultural Change*, 70(3):000–000.

Rozenblit, L. and Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive science*, 26(5):521–562.

Sadoff, S. (2014). The role of experimentation in education policy. *Oxford Review of Economic Policy*, 30(4):597–620.

Samat, C., Chaijaroen, S., and Wattanachai, S. (2019). The designing of constructivist web-based learning environment to enhance problem solving process and transfer of learning for computer education student. In *Innovative Technologies and Learning: Second International Conference, ICITL 2019, Tromsø, Norway, December 2–5, 2019, Proceedings 2*, pages 117–126. Springer.

Yeager, D. S. and Dweck, C. S. (2012). Mindsets that promote resilience: When students believe that personal characteristics can be developed. *Educational psychologist*, 47(4):302–314.

A Appendix A

district_id	Grades	Size	%FRL	%ELL	%White	%Blk	%Hsp	%Asian
1	9-12	8,000	88	N/A	5	3	90	1
2	9-12	8,000	15	2	77	5	5	7
3	9-12	6,000	30	N/A	53	3	27	15
4	6-8	12,000	24	12	73	9	7	7
5	9-12	7,000	62	N/A	21	19	56	1
6	6-8	7,000	26	16	68	2	24	3
7	9-12	9,000	43	4	42	34	7	5
8	6-8	7,000	28	17	50	3	42	1
9	6-12	5,000	11	0.7	80	2	8	5
10	6-12	7,000	36	14	48	11	30	3
11	6-12	9,000	31	6	70	7	7	8

Figure A.1: District Demographics

State	district_id	school_id	Grades	Size	% Math Proficiency	Graduation Rate (HS)	%FRL	%ELL	%White	%Blk	%Hsp	%Asian
IL	1	1	9-12	3,620	10	72	82		6	5	86	1
IL	1	2	10-12	3,360	7	73	93		4	2	94	0
IA	2	3	9-12	2,270	74	96	21		78	6	5	6
IL	3	4	9-12	2,450	52	98	8		82	0	11	4
MI	4	5	6-9	570	46		41		55	16	9	13
IL	5	6	9-12	3,400	11	72	72		12	15	70	0
IL	5	7	9-12	3,550	19	81	50		31	23	42	1
IL	6	8	6-8	790	24		28		63	1	30	3
IL	6	9	6-8	850	26		28		68	2	25	3
OH	7	10	9-12	1,770	25	91	39		49	32	6	4
IL	8	11	6-8	910	9		39		60	3	33	1
OH	9	12	5-6	800	76		12	1	80	2	8	5
OH	9	13	9-12	1,600	74	98	11	1	82	2	7	6
OH	9	14	PK-4	600	86		9	1	78	1	9	7
IL	10	15	9-12	2,170	16	90	42	14	53	9	30	2
IL	10	16	6-8	760	8		57	17	47	7	39	1
WI	11	17	6-8	850								
WI	11	18	6-8	440	49		15		85	3	3	6
WI	11	19	6-8	340	26		47		70	5	7	9

Figure A.2: School Demographics

Figure A.3: Study Information Sheet



Research Study Information

Study Number: IRB23-0041

Study Title: Learning By Creating Math Videos

Researcher: Rohen Shah (shahr@uchicago.edu)

Study Summary

We are conducting a research study on boosting middle and high school students' math skills. Anecdotally, teaching is a good way to learn something yourself. However, students are not often put in a position where they have to explain math to someone else. This project aims to create opportunities for students to learn by explaining math concepts to a hypothetical peer. To make this scalable, students will be assigned to create short video explanations of math concepts. We are seeking to partner with math teachers who teach at least two sections/periods of the same math class. We will randomly assign one section to receive "business as usual" practice problems, and another section to receive a video-creation assignments for the same math problems. Teachers will have autonomy in choosing the specific math topics so that the problems align with their planned curriculum. We also ask teachers to administer a 20-minute pre-test and post-test. Teachers will be compensated $500 for their time. The study will begin in Spring 2024.

Who are you looking for as participants?

Middle and high school math teachers who teach at least two sections of the same math class.

Will participants be compensated?

Yes, participating teachers will receive $500.

Who is funding this project?

This project is funded by J-PAL (https://www.povertyactionlab.org/sector/education).

1

THE UNIVERSITY OF CHICAGO

What will students get out of the study?

Students will participate in math tasks which may improve their math skills and confidence.

What is the study design?

For each participating teacher, we will choose one class section at random to receive "business as usual" homework problems. The other section will receive a modified version of the same task where students will also have to create a video explaining the solution. The randomization is important so that any difference is due to the task itself and not other confounding variables. We will administer a pre- and post-test (based on topics the teacher plans to cover that semester) to measure differences in student learning between the two groups. The study will take place between February and May 2024.

What would teachers be asked to do?

Teachers will be asked to do the following:

(1) Meet with researchers to go over the requirements of the study.

(2) Administer brief (20 min) pre- & post-tests to students at the beginning and end of the study.

(3) Assign aforementioned weekly tasks to students in the respective sections for a grade.

(4) Share final math course grades with researchers at the end of the semester.

What data will you need from administrators?

We will need student math grades from the 2023-24 Fall and Spring semesters, as well as past GPA information. Our university will have a formal data sharing agreement which will outline all of the steps we take to ensure data security.

Figure A.4: Parent Permission Form

THE UNIVERSITY OF CHICAGO

Version: 01/02/2024

Parental Permission Form for Participation in Research

Study Number: IRB23-0041
Study Title: Learning Math through Videos
Researcher(s): Rohen Shah and John List

This form is to be completed by parents of students under the age of 18. Your child is being asked to participate in a research study about improving mathematics skills, conducted by the University of Chicago. This form has important information about the reason for doing this study, what we will ask your child to do, and the way we would like to use information about your child if you choose to allow your child to be in the study. The purpose of the study is to learn about the effectiveness of various math assignments on student learning.

What will my child be asked to do if my child is in this study?
Your child will be asked to fill out two short surveys as well as complete normal classroom tasks. The surveys will each take approximately five minutes and will ask your child about their experiences and attitudes about learning math.

What are the possible benefits for my child or others?
In addition to helping us learn about various techniques to help students learn, your child may also directly benefit in terms of mathematics skills.

What are the possible risks or discomforts to my child?
To the best of our knowledge, the things your child would be doing in this study have no more risk of harm than the risks of everyday life. As with all research, there is a chance that confidentiality of the information we collect about your child could be breached – we will take steps to minimize this risk. Please be aware that under the Protection of Pupils Rights Act (20 U.S.C. Section 1232(c)(1)(A)), you have the right to review a copy of the questions asked of or materials that will be used with students. If you would like to do so, you should contact the research study team to obtain a copy of the questions or materials.

Data and Confidentiality
Your child's educational records, particularly class grades and test scores, will be shared with the research team at the University of Chicago. Your child's homework assignment data might also be shared with the research team. These records will only be used for research purposes, and the researchers will de-identify the data once it is received. Identifiable data will never be shared outside the research team. De-identified information from this study may be used for future research studies or shared with other researchers for future research without seeking your additional permission.

What are my child's rights as a research participant?
Participation in this study is voluntary. Your child may withdraw from this study at any time -- you and your child will not be penalized in any way or lose any sort of benefits for deciding to stop participation. Participation or lack of participation will not affect the child's grades or their relationship with the school/teachers. If you or your child decide to withdraw from this study, the researchers will ask if the information already collected from your child can be used.

Financial Information
Participation in this study will involve no cost to you or your child. Your child will not be paid for participating in this study.

Contact:
If you have questions or concerns about the study or your child's participation, you can reach out to the Co-Principal Investigator, Rohen Shah at shahr@uchicago.edu or 734-578-5684, or the Principal Investigator, John List, at jlist@uchicago.edu. This study is approved by the University of Chicago's Social & Behavioral Sciences Institutional Review Board (IRB23-0041). If you have any questions about your child's rights as a participant in this research, or to discuss other study-related concerns with someone who is not part of the research team, you can contact the University of Chicago Social & Behavioral Sciences Institutional Review Board (IRB) Office by phone at (773) 702-2915, or by email at sbs-irb@uchicago.edu.

Parental Permission for Participation in Research

Please indicate below whether you agree with the following:

> **I have read this form; the research study has been explained, and I know whom to contact if I have any questions. I give permission for my child to participate in the research study described above and will receive a copy of this form. I also give permission to my school to share educational records with the researchers.**

Check one of the options below:

_____ *Yes, I agree with the statement above and give permission for my child to participate.*

_____ *No, I disagree with the statement above and do not permission for my child to participate.*

_____ _____
Parent/Legal Guardian's Signature Date

_____
Child's Name (printed)

_____
Parent/Legal Guardian's Name (printed)

Figure A.5: Student Baseline Survey

Math Research Survey

You are being asked to participate in a research study about improving mathematics skills, conducted by the University of Chicago.

You will be asked to fill out two short (under 5 minutes) surveys about your experiences and attitudes about learning math. All of your responses will be strictly confidential and any identifiable data will not be shared with any third party.

What are my rights as a research participant?
Participation in this study is voluntary. You may withdraw from this study at any time -- you will not be penalized in any way or lose any sort of benefits for deciding to stop participation. Participation or lack of participation will not affect your grades or your relationship with the school/instructors. If you decide to withdraw from this study, the researchers will ask if the information already collected from you can be used.

Financial Information
Participation in this study will involve no cost to you. You will not be paid for participating.

What are the possible benefits?
In addition to helping us learn about various techniques to help students learn, you may also directly benefit in terms of mathematics skills.

What are the possible risks or discomforts?
To the best of our knowledge, the things you would be doing in this study have no more risk of harm than the risks of everyday life. As with all research, there is a chance that confidentiality of the information we collect about you could be breached – we will take steps to minimize this risk. Please be aware that under the Protection of Pupils Rights Act (20 U.S.C. Section 1232(c)(1)(A)), you have the right to review a copy of the questions asked of or materials that will be used with students. If you would like to do so, you should contact the research study team to obtain a copy of the questions or materials.

Contact:
If you have questions or concerns about the study or your participation, you can reach out to the Co-Principal Investigator, Rohen Shah at shahr@uchicago.edu or 734-578-5684, or the Principal Investigator, John List, at jlist@uchicago.edu. This study is approved by the University of Chicago's Social & Behavioral Sciences Institutional Review Board (IRB23-0041). If you have any questions about your rights as a participant in this research, or to discuss other study-related concerns with someone who is not part of the research team, you can contact the University of Chicago Social & Behavioral Sciences Institutional Review Board (IRB) Office by phone at (773) 702-2915, or by email at sbs-irb@uchicago.edu.

* Indicates required question

1. Do you consent to participate in this study? *

   Mark only one oval.

   ◯ Yes *Skip to question 2*
   ◯ No

Survey Questions

2. Name: *

_____

3. Age (years) *

Mark only one oval.

- 10
- 11
- 12
- 13
- 14
- 15
- 16
- 17
- 18+

4. Gender

Mark only one oval.

- Female
- Male
- Other

5. Race

Mark only one oval.

- Black
- Hispanic
- White
- Asian
- Native American
- Other

6. Which grade are you currently in? *

Mark only one oval.

- 6th
- 7th
- 8th
- 9th
- 10th
- 11th
- 12th

7. What class period do you have math class? *

_____

8. How confident are you in your current math abilities? *

Mark only one oval.

	1	2	3	4	5	6	7	8	9	10	
Not											Very confident

9. How much do you agree with the following statement: "My math skills would improve if I spend more time on it" *

Mark only one oval.

- Strongly Disagree
- Disagree
- Neutral
- Agree
- Strongly Agree

10. Which of the following best describes how often you've **received** math tutoring? *

Mark only one oval.

- ◯ Not at all
- ◯ Every now and then
- ◯ About 1 hour per week, most weeks
- ◯ 2 or more hours per week, most weeks

11. Which of the following best describes how often you've tutored someone else? *

Mark only one oval.

- ◯ I have never helped anyone with (or tutored) math
- ◯ I've helped some friends with math occasionally
- ◯ I regularly tutor math to students in classes below me
- ◯ I regularly tutor math to peers in my class

12. How much experience do you have making videos (TikTok, YouTube, etc.)? *

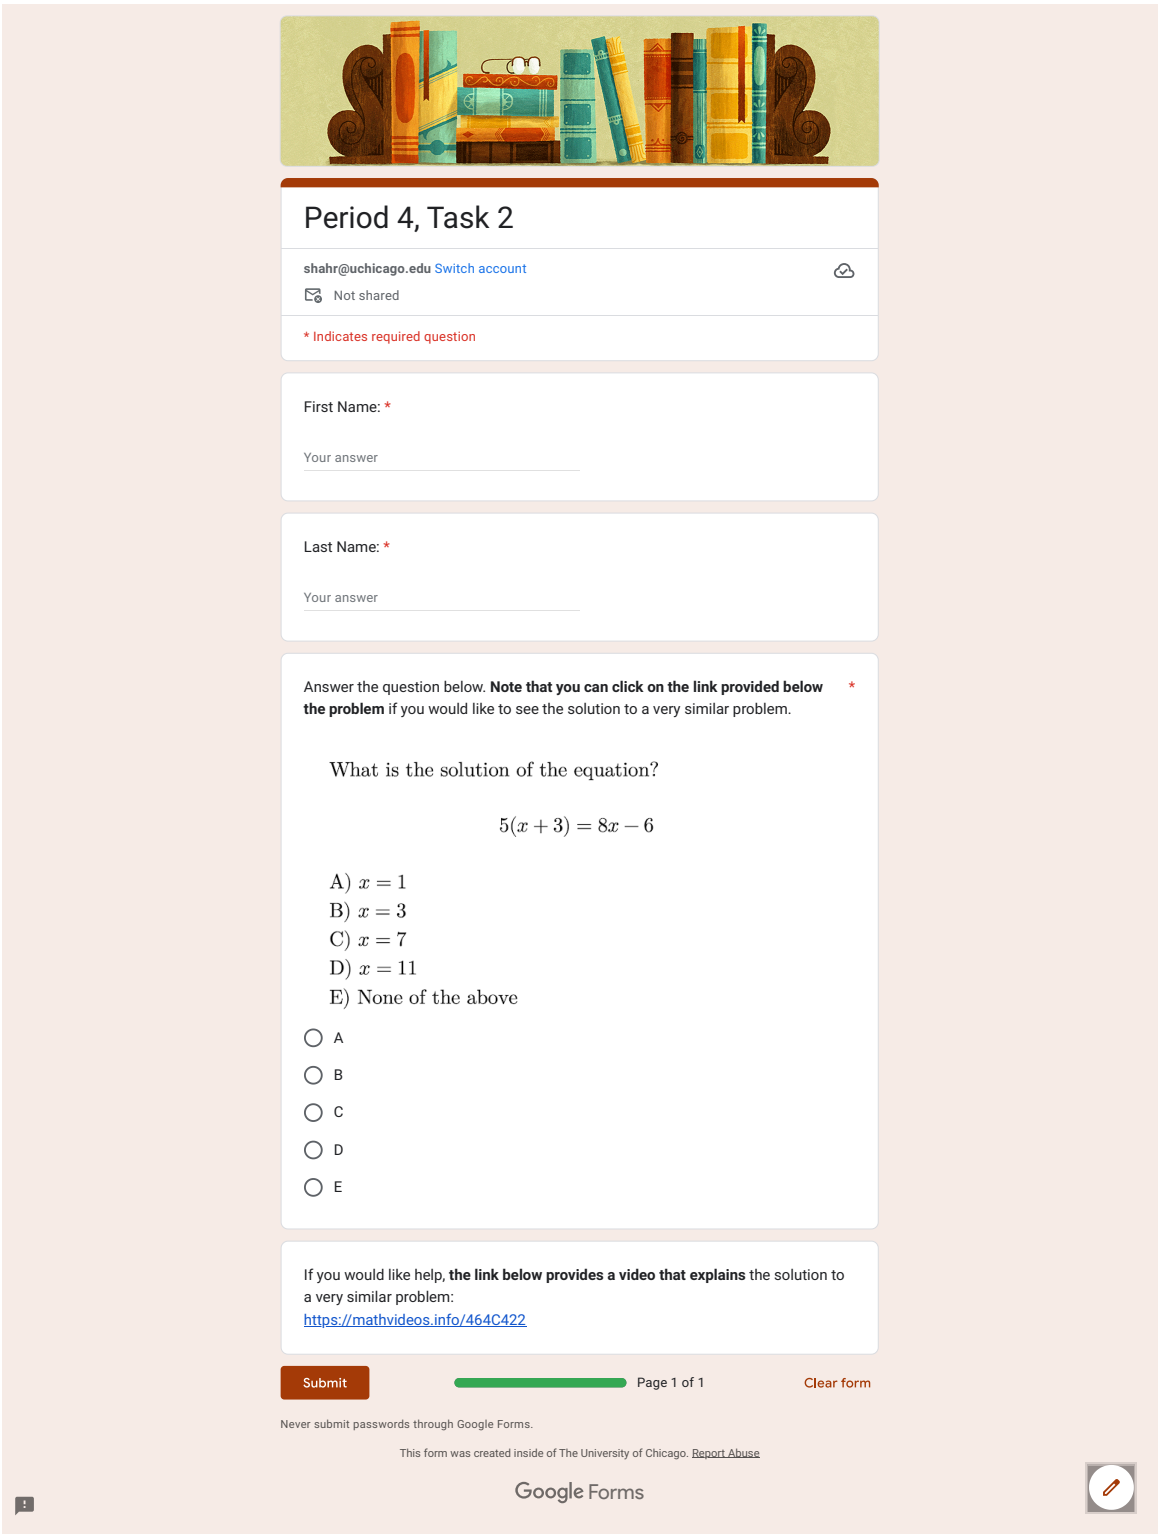
Mark only one oval.

- ◯ None
- ◯ A little
- ◯ Moderate amount
- ◯ A lot

This content is neither created nor endorsed by Google.

Google Forms

Figure A.6: Example of Passive Task



Period 4, Task 2

shahr@uchicago.edu Switch account

Not shared

* Indicates required question

First Name: *

Your answer

Last Name: *

Your answer

Answer the question below. **Note that you can click on the link provided below the problem** if you would like to see the solution to a very similar problem. *

What is the solution of the equation?

$$5(x+3) = 8x - 6$$

A) $x = 1$
B) $x = 3$
C) $x = 7$
D) $x = 11$
E) None of the above

- A
- B
- C
- D
- E

If you would like help, **the link below provides a video that explains** the solution to a very similar problem:
https://mathvideos.info/464C422

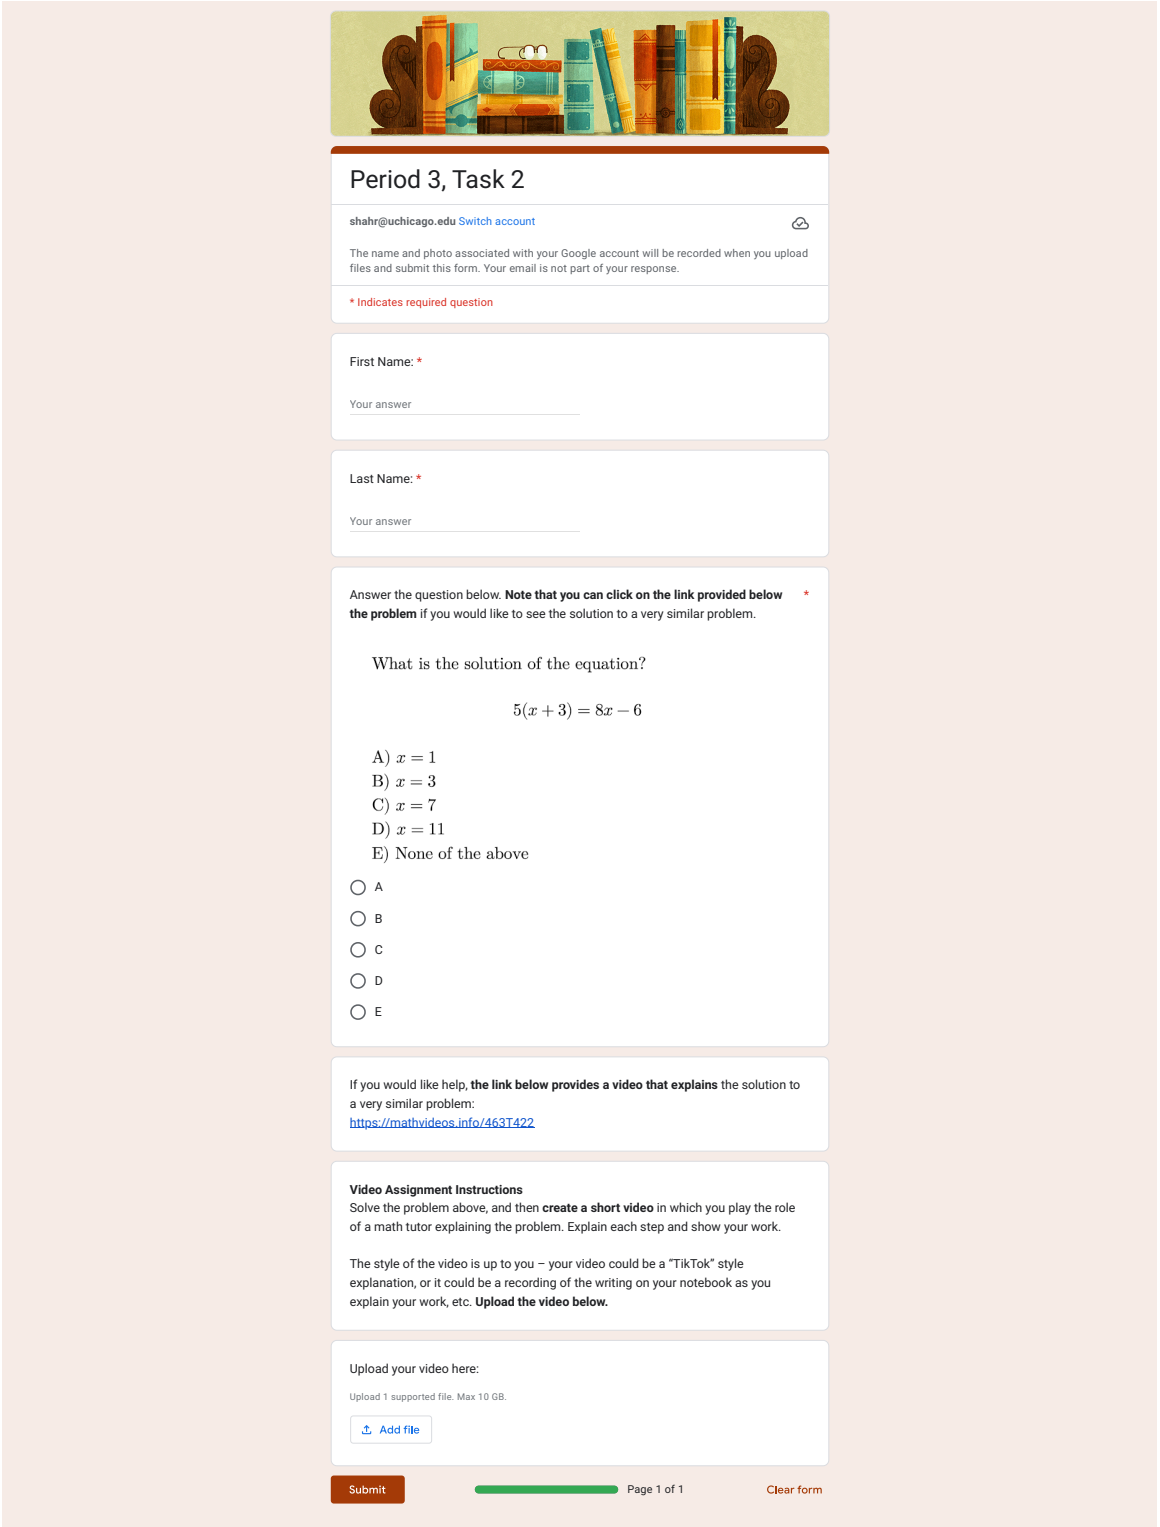
Submit Page 1 of 1 Clear form

Never submit passwords through Google Forms.

This form was created inside of The University of Chicago. Report Abuse

Google Forms

Figure A.7: Example of Video Creation Task



Period 3, Task 2

shahr@uchicago.edu Switch account

The name and photo associated with your Google account will be recorded when you upload files and submit this form. Your email is not part of your response.

* Indicates required question

First Name: *

Your answer

Last Name: *

Your answer

Answer the question below. **Note that you can click on the link provided below the problem** if you would like to see the solution to a very similar problem. *

What is the solution of the equation?

$$5(x+3) = 8x - 6$$

A) $x = 1$
B) $x = 3$
C) $x = 7$
D) $x = 11$
E) None of the above

- A
- B
- C
- D
- E

If you would like help, **the link below provides a video that explains** the solution to a very similar problem:
https://mathvideos.info/463T422

Video Assignment Instructions
Solve the problem above, and then **create a short video** in which you play the role of a math tutor explaining the problem. Explain each step and show your work.

The style of the video is up to you – your video could be a "TikTok" style explanation, or it could be a recording of the writing on your notebook as you explain your work, etc. **Upload the video below.**

Upload your video here:

Upload 1 supported file. Max 10 GB.

Add file

Submit Page 1 of 1 Clear form