



Knowing what it takes: The effect of information about returns to studying on study effort and achievement[☆]

Derek Rury^{a,1,2}, Scott E. Carrell^{b,c,*}

^a University of Chicago, United States of America

^b UC Davis, United States of America

^c NBER, United States of America

ARTICLE INFO

JEL classification:

I26

I29

D83

Keywords:

Beliefs

Study effort

Human capital

ABSTRACT

We study the effect of providing students with information on the returns to study effort in a large introductory microeconomics course. To do so, we use historical time-use data from the course's online homework module to estimate the association between study time and course performance. We measure the impact of providing students this information on subsequent study effort, class attendance, homework scores, and exam performance using a randomized research design. Results show that the information contained in our intervention increased time spent studying by approximately 7% throughout the entire course, though this effect is imprecisely measured. However, when examining shorter-run outcomes (prior to the next exam) we find larger and more precisely estimated treatment effects on time spent on homeworks (12%) and homework scores (14% of a standard deviation). Treatment effects on longer run outcomes in the course are negligible. We additionally estimate large, but somewhat imprecise, average treatment effects on class attendance and small positive and insignificant average treatment effects on exam performance throughout the course.

1. Introduction

A student's study effort is a critical part of their education production function (Bonesrønning & Opstad, 2015; Fraja et al., 2010; Gneezy et al., 2019; Stinebrickner & Stinebrickner, 2004, 2008). As such, understanding how students make decisions on how much time to allocate towards studying is of great importance for education policy-makers. Study effort decisions contain important trade-offs for students, as increased time towards studying implies less time for other activities such as leisure and work (Bound et al., 2010, 2012; Metcalf et al., 2019; Stinebrickner & Stinebrickner, 2003). However, in order to make these trade-offs efficiently, students must know the actual returns to effort. That is, how study effort maps onto academic outcomes such as performance on exams and course grades.

Despite the central importance of these study effort decisions, the research is decidedly thin regarding the returns to study effort and how students make these choices. This is likely due to the fact that valid measures of study effort are both difficult to obtain and are endogenous to other factors affecting achievement. As a consequence, it is unclear

how and whether students incorporate information about the returns to study effort into their beliefs and behavior, and whether changing those beliefs can lead to increases in academic achievement.

Previous work has shown that students often have incorrect beliefs about their own education production function (Ersoy, 2021; Fryer, 2016). Absent accurate information on their returns to study effort, students may over or under-invest in studying, leading to an inefficient allocation of time. While there is a growing literature examining information interventions in college classes, to our knowledge, no study has attempted to update students beliefs about the returns to study effort in a similar setting.

To fill this gap in the literature, we derive and administer an information intervention that both elicits and shocks students' beliefs about their returns to effort in an introductory microeconomics course. To obtain a valid measure of study effort to create the information treatment, we leverage historical, granular time-use data derived from the course's online homework application. After eliciting student's own beliefs about returns to study effort in a baseline survey, we randomly provide one-half of the students in the class information regarding the

[☆] This project was registered at the American Economic Associations RCT registry on June 4, 2019 under "Beliefs, Effort and College Success", with ID number AEARCTR-0004261.

* Correspondence to: Department of Economics, UC Davis, One Shields Avenue Davis, CA 95616, United States of America.

E-mail addresses: rury@uchicago.edu (D. Rury), secarrell@ucdavis.edu (S.E. Carrell).

¹ Addresses: Harris School of Public Policy at University of Chicago, Keller Center 1307 E. 60th st. Chicago, IL 60637.

² Both authors contributed equally to this paper.

average returns to effort using data from a previous version of the course taught by the same instructor. We then track subsequent study effort and course performance for all students.

Results show that the information contained in our intervention increased time spent studying by approximately 7% throughout the entire course, though this effect is imprecisely measured. However, when examining shorter-run outcomes (prior to the next exam) we find larger and more precisely estimated treatment effects on time spent on homeworks (12%) and homework scores (14% of a standard deviation). Treatment effects on longer run outcomes in the course are negligible. We additionally estimate large, but somewhat imprecise, average treatment effects on class attendance and small positive and insignificant average treatment effects on exam performance throughout the course. When controlling for multiple hypothesis testing across our study effort outcomes, we retain significance ($p < 0.05$) on short term study time.

Using responses from our baseline survey, we examine the role of initial beliefs about the returns to study effort when estimating our treatment effects. We fail to detect significant differences in treatment between students who have high or low estimates of their returns to study effort. As such, we posit that the short-run effects on study effort we observe are driven by one of two mechanisms: (1) an increase in the salience of study effort induced by the treatment; and/or, (2) a diminishing of the treatment effect in the long-run as students update their beliefs about their true returns to study effort.³ In our setting, students had multiple homework and exams post-intervention to learn about their own returns to study effort. Under these circumstances, student beliefs may naturally evolve and change, likely reducing the effectiveness of our information as the course progresses.

Finally, we use our experimentally induced short-term increase in study effort to estimate the causal impact of study time on homework scores and exam performance. We do so by using treatment status as an excluded instrument in an IV regression of course performance on study time. Results indicate ordinary least squares estimates tend to underestimate the returns to study effort, indicating study effort is negatively selected with performance.

This paper makes contributions to several related literatures. First, we contribute to the research examining student effort decisions and the effect of study effort on academic achievement (Ahn et al., 2019; Fraja et al., 2010; Metcalf et al., 2019; Stinebrickner & Stinebrickner, 2008). This literature has primarily focused on how changing incentives for achievement affect study effort and subsequent performance (Azmat & Iriberry, 2015; Golightly, 2020; Hishleifer, 2016; Stinebrickner & Stinebrickner, 2008), while our paper manipulates beliefs about effort while keeping incentives for achievement fixed. Additionally, in a pair of papers most related to ours, Ersoy (2019, 2021) uses data from an online language platform to demonstrate that students have incorrect beliefs about returns to study effort and shows that those beliefs become more accurate upon receiving information. In the spirit of Ersoy (2021), who studies how similar information impacts beliefs and performance for Duolingo users, our study focuses on updating students' beliefs in a classroom setting about the returns to effort.

Second, this paper contributes to the literature on performance feedback by examining how changes in beliefs translate into changes in achievement. Prior studies have found strong effects on achievement as a result of performance feedback (Azmat & Iriberry, 2010; Bandiera et al., 2015; Bobba & Frasinco, 2019a, 2019b; Brade et al., 2018; Gonzalez, 2017; Goulas & Megalokonomou, 2018; Li, 2018), although these effects are not always positive (Azmat et al., 2019). These papers rightly interpret these effects as a result of changing beliefs about students' own abilities. What is less clear, however, is the mechanisms by which changes in beliefs translate into changes in achievement. As students

learn about their ability, some inputs into the education production function must also change. The inputs most under the student's control, as well as those we believe most likely to be related to beliefs about ability, are those related to study effort. We aim to study this potential link and assess whether changes in beliefs about returns to study effort mimic findings from the research on beliefs about ability.

Third, we contribute to the literature examining beliefs, specifically in an education setting. A large literature has emerged over the past decade which demonstrates the importance of students' beliefs in decision making (Bobba & Frasinco, 2019a, 2019b; Conlon, 2020; Wiswall & Zafar, 2015a, 2015b; Zafar, 2011, 2013). To our knowledge, we are the first to document beliefs about returns to study effort in a common educational setting; a large introductory course at a selective public four-year university. We also examine the role *ex ante* beliefs about returns to study effort play in making study effort decisions.

Finally, we contribute to the literature studying "nudges" in education that attempt to alter student behavior via light-touch interventions, although with varying success (Carrell et al., 2020; Damgaard & Nielsen, 2018; Li, 2018; Oreopoulos et al., 2020). We show that our light-touch intervention changes short-run behavior in a way that is consistent with a common decision framework used in economics, further demonstrating that nudges may yet play an important role in the classroom.

The rest of the paper is organized as follows; Section 2 provides details on our field experiment; Section 3 describes our data; Section 3.3 presents our results; Section 4 concludes.

2. Experimental details

We implemented our effort experiment during the spring quarter of 2019 in two large introductory microeconomics courses with a total enrollment of over 700 students at a large selective public university on the west coast. We administered a *baseline survey* during the first week of class and short surveys at the beginning of each exam asking about time spent studying during the previous week. All surveys were completed by hand.

The baseline survey asked students questions regarding beliefs about their academic ability, preferences for majors, expected grade in the course, as well as beliefs about returns to study effort. Specifically, students were asked "how many hours do you think you would have to study per week to increase your grade by one letter?". The baseline survey also asked students to sign a Family Education Rights Protection Act (FERPA) to release their academic and demographic information from the university registrar.

The second survey was administered to all students in class prior to start of the first exam during the third week of the course. Two different surveys were randomly distributed to students: a *treatment* survey and a *control* survey. Both surveys asked questions regarding the amount of time spent studying for this class as well as other classes in the past week, followed by a paragraph of text and a short yes or no question, which was included to measure whether students actually read the paragraph.⁴ The paragraph in the *treatment* survey contained information on the returns to study effort. Specifically, students were told, "Using data from Prof. X's course last year, we found a significant relationship between the time students spent on homework and their course grade. Specifically, we found that for the average student, an additional three and a half hours of study time per week was associated with an improvement of a full letter grade in the course". The *control* survey contained a paragraph describing the benefits of participating in research on campus. The font and amount of text used in both the *treatment* and *control* surveys were designed so that the two surveys would appear identical at a quick glance.⁵

³ Recent research suggests informational interventions tend to have stronger effects through the salience of information rather than impacting beliefs (Bettinger et al., 2021).

⁴ We do not study students' self-assessments of study effort in this paper.

⁵ The *treatment* and *control* surveys can be found in the [Appendix](#).

The survey containing the information treatment appeared on the back of the first page of the exam. Exams were ordered such that treatment and control surveys alternated in their placement in order to provide an exogenous distribution of surveys. Teaching assistants handed out exams to students as they entered the classroom and took their seats. In section 3.2, we verify that assignment to treatment and control groups appears to be as good as random across student characteristics. Once the class began, students were given five minutes to turn over the first page of the exam and complete the survey. Students then turned the survey into the teaching assistants prior to the exam starting.

While we acknowledge that there are likely non-linearities in the returns in study effort (eg. very low returns for high levels of effort compared to high returns for low levels of effort), we see our information intervention targeting a linear approximation to a student's belief of returns to effort. In this setting, we assume students' hold imperfect information about returns to effort, which causes them to have uncertainty about how study time maps onto learning and achievement. Hence, we assume our information intervention will cause students to update their mean belief of returns to study effort closer to the information provided. Implicit in this argument is an assumption that students would not update in the opposite direction. For example, if a student estimates an increase in their letter grade after five hours of study time, we would expect the student to update their beliefs downward towards the 3.5 estimate provided in the intervention.⁶

3. Data and results

3.1. Data

A total of 456 students completed the baseline survey and 566 students completed the second survey. Our primary analytic sample consists of the 318 students who completed both surveys and signed the FERPA release, which represents just over two thirds of the baseline survey sample. Though we cannot rule out selection into the sample, we note that students in the analytic sample performed similar to those in the entire class.⁷

We then matched these survey data to course administrative data containing our primary outcomes of interest, including, time spent on homework assignments, class attendance, homework scores, and exam scores. Table 2 presents summary statistics for our pre-treatment classroom measures, survey results and background characteristics for the entire sample and separately for students by treatment status. Importantly, our data contains time-use data on each of the nine online homework assignments throughout the course.⁸ The time use data are measured in seconds, providing a granular measure of study effort. Specifically, these data measure the time spent between the moment when a student begins the homework assignment and either completes it or exits the homework module.⁹

⁶ We recognize that a highly sophisticated student could react in the opposite direction of the intended treatment if they understand the treatment is not a causal estimate.

⁷ Results available upon request.

⁸ These time-use data capture the exact amount of time a student has the homework open on their computer.

⁹ We used the same time use data from the previous time the instructor taught the course, spring 2018, to create our information treatment. To do so, we regress time spent on homework on overall percentage points in the course. We find that three and a half hours of study time is associated with a statistically significant ten percentage point increase in the course, which corresponds to an increase in one letter grade. Not surprisingly, we also find a strong positive relationship in our study sample between measures of time spent on homework throughout the course and overall course performance. Specifically, we find that an increase of one unit in log time spent on homework assignments is associated with a statistically significant increase

Homework scores are measured as the percentage of questions answered correctly. As the homework module allowed for multiple attempts, we view homework scores as a measure of both effort and learning. We also study the effect of our treatment on a rough measure of classroom attendance using an app called Pocket Points. Specifically, the instructor in the course incentivized students to come to class by providing extra credit. Students earned five points (one percentage point of the overall course grade) if they attended at least 16 h ($\approx 75\%$) of course lectures. To track attendance, Pocket Points requires students to be in class, open the app, and put their phone in sleep mode. The outcome we create to capture attendance is an indicator variable for whether the student met the 16 h threshold. These three measures (HW time, HW score, and class attendance) offer us the unique ability to study multiple dimensions of study effort.

We believe the information treatment to be most salient for study effort on the four homework assignments (homeworks two through five) that were due immediately following the information treatment and prior to students taking the second exam. That is, we believe time spent on these homework assignments to be our most valid measure of behavioral changes induced by the treatment, for a few reasons. First, save for exams, there were no other activities other than homework for which students received course credit. Second, students are most likely to remember the information and incorporate it into their studying decisions immediately after receiving the treatment. Third, students will likely (endogenously) update their own estimates of the returns to effort after receiving new signals (e.g., follow-on exam performance) about the returns to study effort. Lastly, time spent studying is the variable over which students have the most control, while other outcomes such as exam and homework scores are the result of a mapping of effort onto achievement.¹⁰

3.2. Balancing tests and descriptive results

First, we perform checks to examine whether our treatment was, in fact, assigned exogenously. As mentioned above, teaching assistants distributed exams by handing out the top exam from their pile to students as they entered the room. While we acknowledge that this mechanism is not truly "random", we see no reason, ex ante, that assignment to treatment would be significantly correlated with observable or unobservable student characteristics.

Table 1 shows results when regressing treatment status on pre-treatment time spent studying and academic achievement, demographic characteristics, and responses to questions in the baseline survey. For statistical inference, and to address for multiple hypothesis testing, in these regressions and our main results, we follow the method proposed by Athey and Imbens (2017) and Fischer (1925) and use a bootstrap-based procedure for testing the null hypotheses using random sampling to assign treatment status. Hence, in addition to reporting traditional standard errors in parentheses, square brackets contain empirical p-values from randomization-based inference using a counterfactual of randomly assigning treatment status 1000 times.¹¹

in course performance of 2.9 percentage points (p value = 0.000). Converting these results into time units, the amount of study time per week associated with a ten percentage point increase found in our study sample ranges from 3.8 to 4.5 h.

¹⁰ The experiment was registered on the American Economic Association trial website under the number AEARCTR-0004261. The trial was registered shortly after the intervention date, but prior to data collection, but did not contain a detailed pre-analysis plan. As such we implement tests to account for multiple hypothesis testing.

¹¹ Athey and Imbens (2017) recommend the use of randomization-based inference in lieu of sampling-based inference for experiments. Additionally, as discussed by List et al. (2019), "by incorporating information about dependence ignored in classical multiple testing procedures, such as the Bonferroni (1935) and Holm (1979) corrections randomization-based inference has much greater ability to detect truly false null hypotheses".

Table 1
Experimental balance.

	(1) TREAT b/se/pvalues	(2) TREAT b/se/pvalues	(3) TREAT b/se/pvalues	(4) TREAT b/se/pvalues	(5) TREAT b/se/pvalues	(6) TREAT b/se/pvalues	(7) TREAT b/se/pvalues	(8) TREAT b/se/pvalues
HW 1 time	0.030 (0.049) [0.620]			0.011 (0.055) [0.860]	0.017 (0.056) [0.820]			
HW 1 score	-0.090 (0.192) [0.620]			-0.112 (0.210) [0.640]	-0.153 (0.215) [0.460]			
Exam 1 score	-0.318 (0.261) [0.180]			-0.181 (0.307) [0.530]	-0.158 (0.309) [0.610]			
Female		-0.006 (0.060) [0.890]		-0.033 (0.063) [0.680]	-0.021 (0.064) [0.820]			
Low income		-0.024 (0.084) [0.790]		-0.044 (0.086) [0.530]	-0.039 (0.088) [0.560]			
African American		0.056 (0.185) [0.780]		0.063 (0.192) [0.720]	0.077 (0.193) [0.650]			
Hispanic		0.081 (0.102) [0.420]		0.087 (0.103) [0.360]	0.135 (0.105) [0.190]			
Asian		0.038 (0.075) [0.510]		0.053 (0.078) [0.410]	0.066 (0.078) [0.340]			
Other race		-0.224 (0.264) [0.430]		-0.295 (0.270) [0.250]	-0.250 (0.273) [0.370]			
SAT/ACT score		-0.000 (0.000) [0.490]		-0.000 (0.000) [0.950]	-0.000 (0.000) [0.970]			
First generation		0.025 (0.077) [0.740]		0.039 (0.078) [0.600]	-0.001 (0.081) [0.990]			
Study habits			-0.012 (0.009) [0.220]	-0.014 (0.010) [0.130]	-0.016 (0.010) [0.120]			
Returns to study effort			0.013* (0.007) [0.080]	0.014 (0.007) [0.100]	0.015* (0.007) [0.060]			
Econ top field			0.038 (0.064) [0.540]	0.016 (0.068) [0.790]	0.024 (0.068) [0.690]			
High control			0.334 (0.304) [0.340]	0.362 (0.313) [0.310]	0.296 (0.315) [0.370]			
Medium control			0.409 (0.306) [0.260]	0.444 (0.315) [0.240]	0.348 (0.317) [0.290]			
Expected A			-0.071 (0.065) [0.320]	-0.066 (0.070) [0.390]	-0.061 (0.071) [0.430]			
High Econ ability			0.005 (0.082) [0.930]	0.037 (0.087) [0.590]	0.020 (0.088) [0.820]			
Predicted HW 1 study time						0.132 (0.092) [0.170]		
Predicted HW 1 score							-0.053 (0.372) [0.870]	
Predicted exam 1 score								-0.698* (0.401) [0.090]
Observations	318	318	318	318	318	318	318	318

Note: Values are the result of regressing column variables on treatment status. Sample consists of all observations in our main analytical sample. Study habits is a variable that represents how much students study per week, Returns to study effort is a variable that represents answers to the survey question “How many hours do you need to study per week to increase your grade by one letter?”, econ top field is a variable that captures whether the student selected economics as their top choice of major, high control and medium control are variables that represent student’s responses to how much control they have other their ability, expected A is a variable that captures if a student expects to get an A in the class, high econ ability is an indicator that captures if a student selects that they have high ability in economics. Standard errors in parenthesis. Empirical P values in brackets.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2
Sample descriptives by treatment status.

	(1)	(2)	(3)
	Mean	Mean (TREAT = 1)	Mean (TREAT = 0)
HW 1 time (log)	5.49 (.623)	5.51 (.630)	5.46 (.617)
HW 1 score	.867 (.166)	.859 (.170)	.873 (.160)
Exam 1	.775 (.118)	.764 (.123)	.785 (.108)
Female	.588 (.493)	.589 (.494)	.587 (.494)
Low income	.255 (.436)	.270 (.445)	.239 (.428)
African American	.028 (.166)	.031 (.173)	.026 (.159)
Hispanic	.217 (.413)	.245 (.431)	.187 (.391)
Asian	.519 (.500)	.515 (.501)	.523 (.501)
Other	.013 (.111)	.001 (.078)	.019 (.138)
SAT/ACT	1240 (301)	1231 (299)	1250 (303)
First generation	.374 (.485)	.405 (.492)	.342 (.476)
Study habits	5.30 (3.72)	5.25 (3.52)	5.35 (3.92)
Returns to study effort	7.19 (5.00)	7.56 (5.14)	6.81 (4.82)
Econ top major choice	.286 (.453)	.307 (.463)	.265 (.443)
High control over ability	.774 (.419)	.748 (.435)	.800 (.401)
Medium control over ability	.217 (.413)	.245 (.431)	.187 (.391)
Expected grade A	.283 (.451)	.245 (.431)	.323 (.469)
High ability in econ (belief)	.154 (.362)	.153 (.361)	.155 (.363)
Observations	318	163	155

Sample consists of all observations in our main analytical sample. Study habits is a variable that represents how much students study per week, Returns to study effort is a variable that represents answers to the survey question “How many hours do you need to study per week to increase your grade by one letter?”, econ top field is a variable that captures whether the student selected economics as their top choice of major, high control and medium control are variables that represent student’s responses to how much control they have over their ability, expected A is a variable that captures if a student expects to get an A in the class, high econ ability is an indicator that captures if a student selects that they have high ability in economics. Standard errors in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Specifications 1–3 show results from separate regression for each of our three groups of pre-treatment variables (academic, demographics, and baseline survey), Specification 4 contains all of these variables in a single regression, and Specification 5 additionally adds teaching assistant fixed effects. In Specifications 6–8 we present results when regressing treatment status on predicted measures of time spent studying, homework score, and exam performance.¹² Results indicate that treatment assignment is largely uncorrelated with our pre-treatment measures, with only 3 of 60 coefficients significant at the ten percent level and none significant at the five or one percent levels.

As such, we see these associations as in line with what we expect to happen simply by chance. In light of this, we conclude that our randomization procedure worked as intended. We also note that in all of our upcoming results specifications we control for (pre-treatment) time spent on homework 1, homework 1 score, and exam 1 score.

¹² We predicted these measures by regressing time spent on homework 1, homework 1 score, and exam 1 score on all of the pre-treatment demographic and survey response variables.

Next, Fig. 1 shows the distribution of students’ estimates of returns to study effort by examining responses to the question asking “How many hours a week do you think you need to study to increase your grade by one letter?”. We see a wide dispersion across our sample, with a dramatic right skewness containing numerous high-value outliers. The median study hours students believe are required to increase one’s grade by a full letter is six, almost twice as large as the estimate provided in the information treatment. This implies that most students in our sample (80%) had beliefs about the number of hours required to increase their grade by one letter that were higher than the information we provided. We categorize these students as having low beliefs about their returns to study effort. Conversely, we have far fewer students (20%) who have (relative to our treatment) high beliefs of their returns to study effort.¹³

Finally, as a fidelity check on whether students read the information paragraph in our intervention, we examine the responses to the yes or no question asked to students below the information paragraph in the second survey. Specifically, the treatment survey asked students if they found the information on returns to study effort useful, while the control survey asked if students wanted to learn more about research on campus. For those who were given the treatment text, we find that 91.3% of students answered “yes” when asked if found the information useful. In contrast, for those students given the control text, when asked if they would like to learn about participating in research, only 49.6% of students answered “yes”. We see this as evidence that students not only read the treatment information provided to them carefully, but broadly speaking, they found it beneficial and were poised to incorporate the information into their beliefs.

3.3. Experimental results

To study the effects of the information treatment on our outcomes of interest, we estimate the following statistical model:

$$y_i = \alpha + \beta TREAT_i + PT_i\gamma + \epsilon_i$$

where y_i is an outcome of interest (eg. time spent on homework), PT_i (pre-treatment) is a vector containing exam one and homework one scores as well as time spent studying on the first homework and ϵ_i represents a random error term. $TREAT_i$ represents an indicator for assignment to the treatment group. Random assignment to the treatment group ensures that $corr(TREAT_i, \epsilon_i) = 0$, allowing us to estimate the causal effect of the information treatment on our outcomes of interest. For all results, traditional standard errors are in parentheses, while square brackets contain empirical p-values from randomization-based inference using a counterfactual of randomly assigning treatment status 1000 times.¹⁴

¹³ Following Ersoy (2021), we also asked students about their beliefs regarding how much “control” they have over their ability. This question was meant to elicit beliefs about whether the students had what is called a “growth mindset” or whether they believed their ability was “fixed”. Unfortunately, while we find this line of inquiry interesting, we found little variation in responses to this question, with only one student believing they had a below average amount of control over their ability.

¹⁴ In results not reported, estimated effects are largely unchanged when we also estimate models when controlling for a vector of background characteristics, including gender, race/ethnicity indicators, socioeconomically disadvantaged, first-generation status, and SAT score. We also include responses from the baseline survey, which capture important beliefs about choice of major, ability in economics, expected grade in the course, as well as beliefs about typical study habits and beliefs about the returns to study effort. We also control for teaching assistant fixed effects, to control for potential differences in the quality of teaching assistants.

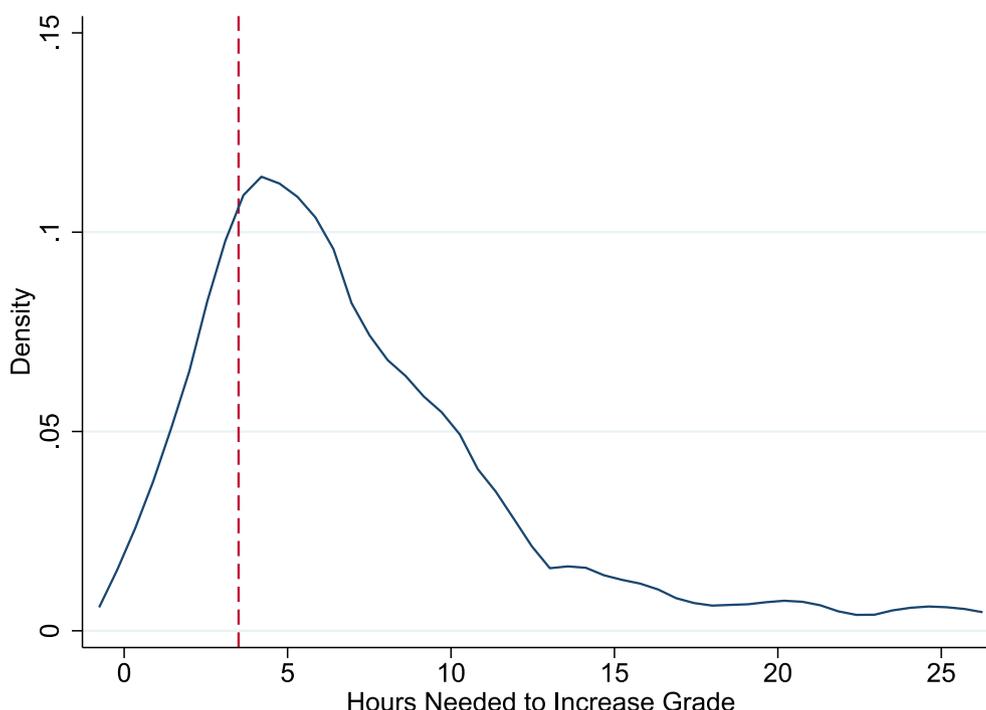


Fig. 1. Distribution of beliefs about returns to study effort. Note: This is a kernel density plot of the survey item “How many hours would you need to study to increase your grade by one letter?” Only primary analysis sample observations are included in this plot. The red dashed line presents the amount of time included in our information treatment.

3.3.1. Average treatment effects

Results showing the average treatment effects for the entire sample are presented in Table 3. Specifications include controls for (pre-treatment) time spent studying on homework 1, homework 1 performance, and exam 1 performance.¹⁵ Panel A presents results from time spent studying on the graded homework assignments. For these outcomes, we take the log of total time spent to approximate a percentage change in study time as a result of treatment. Panel B presents results for the percentage of points earned on the homework assignments. Panel C presents results for exam performance and an indicator measure of class attendance.

Specification 1 of Panel A show the treatment effect on time spent studying throughout the entire course. The coefficient of 0.070 indicates students in the treatment exhibited a 7% increase in study effort relative to the control, though this estimate is imprecisely measured ($p = 0.174$). Next, Specification 2 estimates the short-run treatment effects for time spent on the homework assignments assigned after treatment, but prior to the next exam. We examine effort during this period for two reasons. First, we would expect the treatment to be most salient for assignments that are due in the period immediately following the intervention. Second, as the course progresses and students received additional information signals (e.g., own exam performance), we would expect students to update their beliefs about their own “true” returns to study effort. The coefficient of 0.121 in Specification 2 indicates a 12% increase in study effort for students in the treatment, relative to students in the control, with the effect significant at the 5% level ($p = 0.026$). We interpret this result as strong evidence that the students who received the information treatment significantly increased their study effort in the period immediately following the intervention. Finally, when we look at time spent on homework towards the end of the course, however, we see imprecisely estimated negative treatment effects in Specification 3. We explore and discuss what may be driving

¹⁵ In results not shown, we also find nearly identical results when add controls for baseline survey responses, demographics, and TA fixed effects.

Table 3

The effect of treatment on homework time and scores, class attendance as well as exam and course performance.

Panel A: Effects on time spent on homeworks			
	(1) All homeworks time (2–9)	(2) Short-Run HW time (2–5)	(3) Long-Run HW time (6–9)
TREAT	0.070 (0.052) [0.174]	0.121** (0.054) [0.026]	-0.045 (0.075) [0.588]
Panel B: Effects on scores on homeworks			
	(4) Median HW score (All HW)	(5) Median HW score (Short-Run)	(6) Median HW score (Long-Run)
TREAT	0.024 (0.017) [0.150]	0.031* (0.018) [0.082]	0.010 (0.020) [0.646]
Panel C: Effects on exams and course performance			
	(7) Exam 2 score	(8) Max exam 3–4 score	(9) Attendance
TREAT	0.007 (0.015) [0.626]	0.008 (0.011) [0.478]	0.084 (0.056) [0.128]
Observations	318	318	318

Notes: Standard errors in parentheses; empirical p values in brackets. Statistical significance is based on empirical p values. All models include demographic controls and pre-treatment values of our outcome variables. Study time outcomes are represented in logs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

these effects when we present results based on beliefs about returns to study effort below.¹⁶

¹⁶ We also cannot rule out the possibility that there are spillover between the treatment and control group that would potentially bias our estimate

Next, in Panel B we examine treatment effects on homework scores. Similar to pattern of results in Panel A, we find positive, but statistically insignificant treatment effects throughout the entire course, but larger and more precisely estimated effects on short-run outcomes. Specifically, results in Specification 4 indicate that students in the treatment group showed a 3.1 percentage point (14 percent of a standard deviation) increase in homework score performance, relative to students in the control ($p = 0.082$).¹⁷ Likewise, these effects become smaller and less precise when examining performance on homework assignments later in the course as shown in Specification 6.

Finally, in Panel C, we examine measures of exam performance and class attendance. Though positive, the treatment effects on exam performance are economically small and indistinguishable from zero.¹⁸ These results are perhaps not too surprising, given the relatively modest positive treatment effects found in both study time and homework score. Finally, in Specification 9 we see large, but somewhat imprecise, effects on class attendance. Specifically, results indicate that students in the treatment group were 8.4 percentage points (16.4%) more likely to meet the class attendance incentive ($p = 0.128$).

Finally, as a robustness check, we correct for multiple hypothesis testing using the Romano–Wolf step-down adjusted method. When looking at our study effort measures (short and long term study time and scores), we retain significance only on short-term study time ($p < 0.05$).

3.3.2. Treatment effects by beliefs

As previously detailed, approximately 80% of students had estimated returns to study effort that were lower than our information treatment.

To explore treatment effects across beliefs about returns to study effort, we estimate separate effects depending on whether students had a higher or lower estimate of the returns to study effort compared to our information treatment. Specifically, we categorize students as having high (low) estimates of their own returns if their estimate was lower (higher) than the three and a half hours contained in our information treatment. We do so because these student types may react differently to the information treatment, depending on whether the income or substitution effect in study effort decisions dominates. A model of this decision about study effort can be found in [Appendix](#).¹⁹

Table 4 presents these results in the same format as **Table 3**. Panel A presents results from time spent studying on the graded homework

downward, particularly in the longer-run outcomes. That is, in addition to the potential that control students may be told the information by treated peers, we also recognize that there may be direct spillovers (e.g., peer effects) in behavior. That is, if a treatment student is induced to study more for the second exam, this may result in increased studying for non-treated peers through contemporaneous peer influence. We now discuss this in the text.

¹⁷ When considering scores across multiple homework assignments, our preferred measure is median homework score. We do so because as part of the course design, along with their lowest exam score, each student's lowest homework scores is automatically dropped from their final course point total. This creates an incentive for students to skip one homework without penalty. Taking the median homework score helps us avoid measurement issues that arise from this incentive.

¹⁸ For reasons similar to those regarding homework scores, our preferred measures of exam performance are percentage score on exam two (0–1 scale), as well as student's highest score on exams three and four. Along with dropping their lowest homework score, students' lowest exam score was also dropped from their total course points total.

¹⁹ The model presents a framework for interpreting the signs of the coefficients on homework time to allow us to determine whether the income or substitution effect dominates, on average, as students in the treatment group update their beliefs about study effort. More specifically, if the coefficient for those who had low estimates is positive (negative), and the coefficient for those who had low estimates is negative (positive), then the substitution (income) effect dominates.

Table 4
Heterogeneous effects by beliefs on homework time, scores and exam performance, and attendance.

<i>Panel A: Effects on time spent on homeworks</i>			
	(1) All homeworks time (2–9)	(2) Short-Run HW time (2–5)	(3) Long-Run HW time (6–9)
Low estimate of returns	0.088 (0.059) [0.142]	0.113* (0.062) [0.066]	0.028 (0.084) [0.766]
High estimate of returns	0.009 (0.117) [0.910]	0.158 (0.122) [0.144]	−0.323** (0.167) [0.028]
<i>Panel B: Effects on scores on homeworks</i>			
	(4) Median HW score (All HW)	(5) Median HW score (Short-Run)	(6) Median HW score (Long-Run)
Low estimate of returns	0.021 (0.019) [0.214]	0.033* (0.020) [0.100]	0.019 (0.023) [0.382]
High estimate of returns	0.033 (0.038) [0.436]	0.035 (0.040) [0.360]	−0.028 (0.045) [0.576]
<i>Panel C: Effects on exams and course performance</i>			
	(7) Exam 2 score	(8) Max exam 3-4 score	(9) Attendance
Low estimate of returns	0.010 (0.017) [0.538]	0.007 (0.013) [0.646]	0.148** (0.063) [0.014]
High estimate of returns	−0.017 (0.033) [0.674]	0.016 (0.025) [0.556]	−0.167 (0.124) [0.164]
Observations	318	318	318

Notes: Standard errors in parentheses; empirical p values in brackets. Statistical significance is based on empirical p values. All models include demographic controls and pre-treatment values of our outcome variables. Study time outcomes are represented in logs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

assignments, Panel B presents results for the percentage of points earned on the homework assignments, and Panel C presents results for exam and course performance and attendance.

Focusing on time spent on homework for those students who originally had high estimates of the returns to study effort, we see noticeable differences between the short and long-run effects of the intervention. Specifically, we find relatively large positive (16%), though imprecise, effects in the short run, followed by large negative (−32%) and significant effects in the long run. Conversely, the pattern of results is noticeably different for those students who originally had low estimates of the returns to study effort. We find a positive and marginally significant ($p = 0.066$) short run effects of 11%, though, this effect dissipates to zero when examining study time later in the course.

Next, results in Panel B and Panel C, examining homework and exam scores are less pronounced, but show some evidence of a positive treatment effect for student who had low estimates of the returns to study effort, with treatment effects on homework scores after the intervention significant at the 10% level. Likewise, we find a large, positive, and significant treatment effect of 14.8 percentage points on our measure of class attendance for students who originally had low estimates of the returns to study effort. On the contrary, we find a sizeable negative and insignificant effect on attendance for those students who had high estimates of returns.

Table 5
OLS and IV results of study effort on achievement.

	(1) (OLS) HW 2-5 score	(2) (IV) HW 2-5 score	(3) (OLS) Exam 2 median	(4) (IV) Exam 2 median
HW time	0.098*** (0.017)	0.320** (0.155)	0.017 (0.014)	0.055 (0.102)
Observations	318	318	318	318

Notes: Standard errors in parentheses. All models include demographic controls and pre-treatment values of our outcome variables. Models 1–2 and 5–6 use time spent on HW 2–5 as main explanatory variable while models 3–4 and 7–8 use time spent on homeworks 2–9. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In order for the model developed in [Appendix](#) to accurately describe student behavior, we would need to find opposing signs in our treatment effects between the groups who had low estimates versus those who had high estimates of returns. Interpreting these results through this framework, we see that the large increases in study time across both groups directly after the intervention cannot be fully explained by this framework. Rather, these short run effects are more consistent with our treatment increasing the *salience* of study effort across the distribution of beliefs about returns to studying ([Bettinger et al., 2021](#); [Sexton, 2015](#)). While we do see some evidence that is consistent with the model when looking at study time later in the course and through the effect on attendance, we cannot say for certain whether this model accurately describes students' study effort decisions. Another plausible explanation is that the effect treatment was effective in the short run, but was overpowered by relevant information students' learned about their own returns later in the course. Under this mechanism, our information treatment may have reduced some uncertainty about how effort maps onto achievement early in the course. As such, as students received more feedback from assignments and exams, information in our treatment became less valuable.

3.3.3. Effect of study effort on achievement

An important remaining question is whether or not increased study effort leads to a *causal* effect on academic achievement. This question is difficult to answer due to the likely endogeneity between study effort and academic performance. For example, students with higher ability may study less, on average, while still performing well on exams. Conversely, low-ability student may study more, on average, but perform worse on exams. In these scenarios, a simple regression of study effort on academic achievement will be biased downward, though we note that other scenarios may lead to bias in the opposite direction.

Fortunately, our experiment lends us the opportunity to examine this question. That is, because our experiment exogenously increased study effort (in the short run), we can causally estimate the returns to study effort by instrumenting for time spent studying with treatment status.

[Table 5](#) presents results from this analysis where we estimate two-staged least square models for our homework and exam score outcome variables. Given our treatment only increased study effort in the short run, we estimate these models for outcomes immediately following the intervention. Odd numbered specifications show OLS estimates, while even number specifications show IV estimates where the excluded instrument is treatment status. Though we are primarily interested in the exam score outcome, we also examine the effect of study time on homework scores. We do so knowing that there is likely a mechanical effect on homework scores since students are given multiple opportunities on each homework problem.

As shown in [Table 5](#), for both homework and exam score outcomes we find a consistent pattern of substantially larger estimates for the IV models compared to OLS, suggesting a negative bias in the OLS

estimates. For example, the 0.098 OLS coefficient in Specification 1 suggests that a 10 percent increase in time spent studying is associated with a nearly 1-percentage point increase in median homework score, whereas the IV estimate in Specification 2 shows this effect more than triples to 3.2 percentage points ($p < 0.05$). Turning to the effects on exam performance, we note the difference in magnitude between the OLS and IV are similar to the homeworks score results, though the effects are not precise.

4. Discussion and conclusion

In this paper, we study the effect of providing students with information about the returns to study effort. We measure the impact of receiving this information on several important outcomes such as study effort, class attendance, homework scores and exam performance. We are able to measure study effort using a granular measurement of effort based on time-use data from the course's online homework software. Importantly, we administer a survey that captures students beliefs about returns to study effort. In doing so, we show that around 80% of students had low estimates of their returns to study effort.

We find that our treatment led to a significant increase in time spent on homework assignments as well as homework performance before the subsequent midterm. When looking at how treatment effects evolve over time, we find that our estimates become noisy when looking at the entirety of the course. We also fail to detect significant effects on exam performance. Treatment effects in general do not appear to systematically depend on baseline beliefs about study effort, although class attendance does increase for treated students who had low estimates about their returns.

These results highlight several important findings for the literature examining the returns to study effort. First, our treatment appears to have increased the salience of the importance of studying for the period directly following our intervention for all students ([Bettinger et al., 2021](#)). However, as time progressed, the information we provided became less prevalent in the long-run as students update their beliefs about their true returns to study effort. In our setting, students had multiple homework and exams post-intervention to learn about their own returns to study effort. Under these circumstances, student beliefs may naturally evolve and change, likely reducing the value of our information as the course progresses.

We believe these results speak to findings from the performance feedback literature, which finds that student achievement can be manipulated by providing students with information about how their performance compares to a relevant standard or peer group ([Azmat & Iriberry, 2010](#); [Bandiera et al., 2015](#); [Bobba & Frasinco, 2019a, 2019b](#); [Brade et al., 2018](#); [Gonzalez, 2017](#); [Goulas & Megalokonomou, 2018](#); [Li, 2018](#)). As such, we demonstrate how similar results can be achieved through influencing beliefs about returns to study effort. Though, further research should attempt to measure beliefs about returns to study effort as a result of performance feedback to capture the relationship more formally.

Our results are also important for policy makers who wish to increase achievement and persistence in college. We document that the vast majority of students in our study had low estimates of the returns to study effort. Importantly, our results show that these students increase their study effort upon learning about the "true" returns to effort, albeit perhaps only in the short-run.

CRedit authorship contribution statement

Derek Rury: Conceptualization design, Methodology, Data collection, Econometric analysis, Writing. **Scott E. Carrell:** Conceptualization design, Methodology, Data collection, Econometric analysis, Writing.

Data availability

The authors do not have permission to share data.

Appendix

A.1. Model of beliefs about returns to effort

Below we discuss how information about returns to study effort could influence subsequent study effort through income and substitution effects. To motivate our discussion of the role of beliefs in study effort decisions, consider the following utility function for a representative student. Following Ersoy (2021), the student receives utility from both academic achievement, A , and leisure, L , and allocates their time, \bar{T} , over both study effort, e , and leisure, l

$$\max_{A,L} U(A, L) \tag{1}$$

$$s.t. e + l = \bar{T} \tag{2}$$

We also assume that study effort maps onto academic achievement linearly. Specifically, we assume that $A = f(e) = \alpha e$. We refer to this rate, α , as the “returns to study effort”. For simplicity we also assume that $L = l$. After substituting into Eq. (1), the student’s optimization problem becomes:

$$\max_{e,l} U(\alpha e, l) \tag{3}$$

$$s.t. e + l = \bar{T} \tag{4}$$

where e^* and l^* are the solution to the above problem. Lastly, we assume the utility function is strictly concave so that a unique solution exists and that marginal utility is decreasing for both arguments. Under this framework, the student faces a linear budget constraint in time with which they are endowed \bar{T} and over which they allocate leisure and study effort. We represent the student’s problem in Fig. 2. using the familiar graph used in studying consumption decisions between two goods.

Because effort maps onto achievement at a rate of α , the slope of the budget line is $-\alpha$. The linear axis represents both time spent on leisure, l , as well as time devoted to study effort, e , as $e = \bar{T} - l$. As in consumer theory, optimal effort and leisure are found where the student’s indifference curve is tangent to the budget line, or more formally where $MU_A = MU_L$. Assume also that students do not know the value of α , but have beliefs, $\hat{\alpha}$, about its value.

One way to frame how student behavior changes as beliefs about returns effort change is in terms of the “income” and “substitution” effects. Assume that a student holds beliefs about returns to study effort such that $\hat{\alpha} < \alpha$. We assume here that α is the same for all students but understand that, in reality, returns to study effort are likely to be heterogeneous. This would imply that $\alpha = \alpha_i$ for each student i . In our information intervention, we provide students with the average returns across a large sample of students from a previous course. In doing so, we make a trade off between offering specific information to students with offering feasible information in the form of average returns. In the end, our aim is less to provide individualized information to students but rather shock their beliefs about returns to study effort.

When the student is provided information on the true value of α , we assume the student fully updates their beliefs about their returns to study effort. This leads to a rotation of the budget constraint up

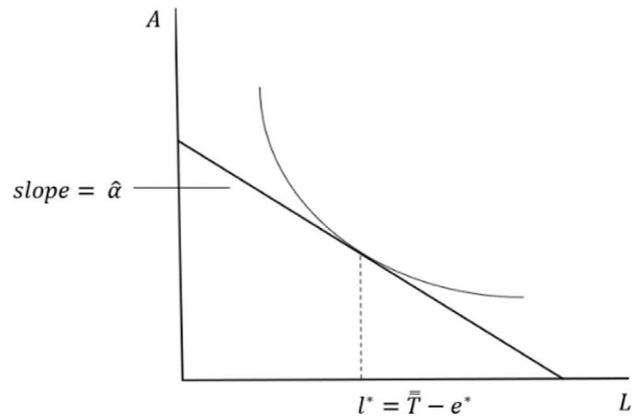


Fig. 2. Utility maximization over study effort and leisure.

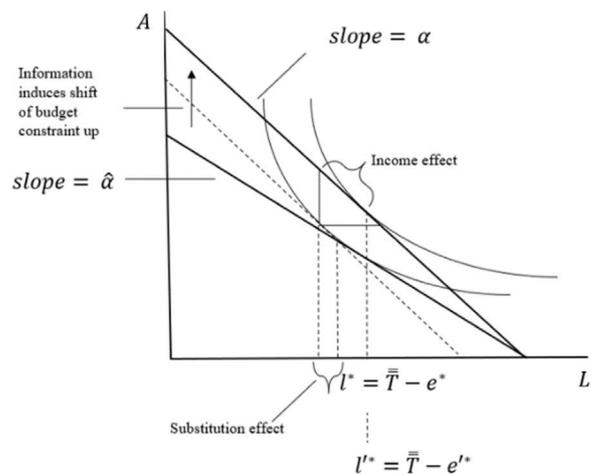


Fig. 3. Substitution and income effect from a shock to beliefs about the returns to effort.

the vertical axis, due to this relative “price” decrease. Similar to a relative price change in consumer theory, this rotation leads to a new equilibrium e^* and l^* resulting from a combination of the income and substitution effects as depicted in Fig. 3.

The substitution effect in this case will lead the student to study more, as academic achievement is now “cheaper”. Likewise, the income effect makes the student study less, as they are now “richer” in time available and leisure is a normal good. As a result of these two opposing effects, the students overall change in study effort, e^* , is ambiguous. The same is true in cases where $\hat{\alpha} > \alpha$, although the income and substitution effects moves e^* in the opposite direction (see Fig. 3).

A.2. Treatment and control information

Treatment Text: Benefits of Study Effort

Recent research has demonstrated that students underestimate the benefits of study effort

Using data from Prof. X’s course last year, we found a significant relationship between the time students spent on homework and their course grade.

Specifically, we found that for the average student **an additional three and a half hours of study time per week** was associated with an improvement of a **full letter grade for the course**.

Control Text:**Benefits of Research Participation**

Student participation in research is an integral part of the research process here at UC Davis

As social science research continues to study how people make important decisions, students have been asked to dedicate more of their time to research participation. There are many benefits of research participation in general, including learning how research is conducted.

We hope you have found your participation in this project to be interesting.

A.3. Baseline survey**Consent to Release Student Information Under FERPA**

Please release the following information to Derek Rury for use in research Mr. Rury's research into understanding student study patterns

- Course enrollment and grades
- Information from my application form about my background including SAT/ACT score, high school GPA, gender, ethnicity, and broad indicators of parent's education and income level. This does NOT include FAFSA or other financial aid information.

Before signing this consent to release your information, please read the following:

- You do not have to agree to release your student information.
- You can withdraw your consent at any time by contacting Derek Rury at drury@ucdavis.edu
- The consent applies only to the categories of data for the specific purpose listed above;
- Only the Individual(s) listed above will receive these data;
- Manuscripts and reports resulting from data analysis will include only aggregated data so you will not be identifiable.

Name:

Student ID:

Signature:

Date:

Survey Questions

- (1) What is your first choice of major? (if you already have a major, please write that major)
- (2) What is your second choice of major?
- (3) How would you describe your ability in your top choice of major? (please circle) a. High b. Medium c. Low
- (4) How would you describe your ability in your second choice of major? (please circle) a. High b. Medium c. Low
- (5) How would you describe your ability in economics? (please circle) a. High b. Medium c. Low
- (6) What grade do you anticipate you will receive in this class? (please only put letter grades with marks, eg. "B+")
- (7) How much control do you think you have over your ability to do well in this class? a. A lot b. Above average amount c. Average amount d. Below average amount e. None
- (8) How much time (in hours) do you think you'd need to study each week to increase your grade in this course by one letter grade?
- (9) How much time (in hours) do you study each week for a typical class?

References

- Ahn, T., Arcidiacono, P., Hopson, A., & Thomas, J. (2019). *Equilibrium grade inflation with implications for female interest in STEM majors: Working Paper*.
- Athey, S., & Imbens, G. (2017). The econometrics of randomized experiments. In *Handbook of economic field experiments, Vol. 1* (pp. 73–140).
- Azmat, G., Bagues, M., Cabrales, A., & Iriberry, N. (2019). What you don't know...can't hurt you? A natural field experiment on relative performance feedback in higher education. *Management Science*, 65(8), 3714–3736.
- Azmat, G., & Iriberry, N. (2010). The importance of relative performance feedback information: Evidence from a natural experiment using high school students. *Journal of Public Economics*, 94(7–8), 435–452.
- Azmat, G., & Iriberry, N. (2015). The provision of relative performance feedback: An analysis of performance and satisfaction. *Journal of Economics and Management Strategy*, 25(1), 77–110.
- Bandiera, O., Larcinese, V., & Rasul, I. (2015). Blissful ignorance? A natural experiment on the effect of feedback on students' performance. *Journal of Labour Economics*, 34.
- Bettinger, E., Cunha, N., Lichand, G., & Madeira, R. (2021). *Are the effects of informational interventions driven by salience? Working Paper No. 350*, University of Zurich, Department of Economics.
- Bobba, M., & Frasinco, V. (2019a). *Learning about oneself: The effect of signaling academic ability on school choice: Working Paper*.
- Bobba, M., & Frasinco, V. (2019b). *Self-perceptions about academic achievement: Evidence from Mexico city: Working Paper*.
- Bonesrønning, H., & Opstad, L. (2015). Can student effort be manipulated? *Applied Economics*.
- Bonferroni, C. (1935). *Il calcolo delle assicurazioni su gruppi di teste*. Rome: Tipografia del Senato.
- Bound, J., Lovenheim, M., & Turner, S. E. (2010). Why have college completion rates declined?: An analysis of changing student preparation and college resources. *American Economic Journal: Applied Economics*, 7(3), 1–31.
- Bound, J., Lovenheim, M., & Turner, S. E. (2012). Increasing time to baccalaureate degree in the United States. *Education Finance and Policy*, 7(4), 365–425.
- Brade, R., Himmler, O., & Jäckle, R. (2018). *Normatively framed relative performance feedback - Field experiment and replication: MRPA Working Paper*.
- Carrell, S., Lusher, L., & Rury, D. (2020). *Major disappointment: A large scale experiment on (non-)pecuniary information and college major: Working Paper*.
- Conlon, J. (2020). Major malfunction: A field experiment correcting undergraduates' beliefs about salaries. *Journal of Human Resources*, Forthcoming.
- Damgaard, M. T., & Nielsen, H. S. (2018). Nudging in economics. *Economics of Education Review*, 64, 313–342.
- Ersoy, F. (2019). *Returns to effort: Experimental evidence from an online language platform: Working Paper*.
- Ersoy, F. (2021). Effects of perceived productivity on study effort: Evidence from a field experiment. *Experimental Economics*.
- Fischer, R. A. (1925). *Statistical methods for research workers*. London: Oliver and Boyd.
- Fraja, G. D., Oliveira, T., & Zanchi, L. (2010). Must try harder: Evaluating the role of effort in educational attainment. *The Review of Economics and Statistics*, 92(3), 577–597.
- Fryer, R. (2016). Information, non-financial incentives, and student achievement: Evidence from a text messaging experiment. *Journal of Public Economics*, 144, 109–121.
- Gneezy, U., List, J., Jin, X., & Sadoff, S. (2019). Measuring student success: The role of effort on the test itself. *American Economic Review: Insights*, Forthcoming.
- Golightly, E. K. (2020). *Does college access increase high school effort? Evaluating the impact of the texas top 10% rule: Working Paper*.
- Gonzalez, N. (2017). *How learning about one's ability affects educational investments: Evidence From the advanced placement program: Mathematica Policy Research Working Paper*.
- Goulas, S., & Megalokonomou, R. (2018). *Knowing who you are; The effect of feedback on short and long-term outcomes: Working Paper*.
- Hishleifer, S. R. (2016). *Incentives for inputs or outputs? A field experiment to improve student performance: Working Paper*.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70.
- Li, H.-H. (2018). Do mentoring, information and nudge reduce the gender gap in economics. *Economics of Education Review*, 64, 165–183.
- List, J., Saik, A. M., & Xu, Y. (2019). Multiple hypothesis testing in experimental economics. *Experimental Economics*, 22, 773–793.
- Metcalfe, R., Burgess, S., & Proud, S. (2019). Students' effort and educational achievement: Using the timing of the world cup to vary the value of leisure. *Journal of Public Economics*, 172(4), 111–126.
- Oreopoulos, P., Patterson, R. W., Petronijevic, U., & Pope, N. G. (2020). Lack of study time is the problem, but what is the solution? Unsuccessful attempts to help traditional and online college students. *Journal of Human Resources*, Forthcoming.
- Sexton, S. (2015). Automatic bill payment and salience effects: Evidence from electricity consumption. *The Review of Economics and Statistics*, 97(2), 229–241.
- Stinebrickner, R., & Stinebrickner, T. (2003). Working during school and academic performance. *Journal of Labor Economics*, 93, 473–491.
- Stinebrickner, R., & Stinebrickner, T. (2004). Time use and college outcomes. *Journal of Econometrics*, 21(2), 243–269.
- Stinebrickner, R., & Stinebrickner, T. (2008). The causal effect of studying on academic performance. *The B.E. Journal of Economic Analysis and Policy*, 14(8), 1–53.
- Wiswall, M., & Zafar, B. (2015a). Determinants of college major choices: Identification from an information experiment. *Review of Economic Studies*, 82(2), 791–824.
- Wiswall, M., & Zafar, B. (2015b). How do college students respond to public information about earnings. *Journal of Human Capital*, 9(2), 117–169.
- Zafar, B. (2011). How do students form expectations? *Journal of Labor Economics*, 29(2), 301–348.
- Zafar, B. (2013). College major choice and the gender gap. *Journal of Human Resources*, 48(3), 545–595.