Dollars and Sense: Self-Selection in Financial Literacy Education*

Elaine Shen† Junru Lyu Laila Voss

March 26, 2026

Preliminary: Do not cite or distribute

Abstract

A central prediction of human capital theory is that individuals with the highest marginal returns to education should be the most likely to invest in it. We document a stark violation of this prediction in financial literacy education: individuals with the lowest baseline knowledge – and thus the highest potential returns – have the lowest demand for education. Using motivating evidence from a large university personal finance course and a controlled online experiment (n=1,047), we document that demand for financial education is positively correlated with baseline financial literacy, potentially exacerbating pre-existing inequalities. In the lab setting, we identify the mechanism driving this inefficiency. Low-ability individuals systematically miscalibrate their beliefs, overestimating their baseline knowledge while underestimating the returns to education. To test whether reducing perceived cognitive costs can correct this distortion, we vary the format of the education material, comparing a traditional lecture-style education to video- and AI-based alternatives. While the AI and video versions significantly increase aggregate demand, they fail to correct the adverse selection on ability. Together, the results suggest that supply-side interventions alone may be insufficient to close financial literacy gaps without addressing underlying belief distortions.

*IRB of record: UC Berkeley Protocol #2024-04-17372 (Observational Data), 2024-03-17249 (Lab Experiment); AEA Trial Registration #AEARCTR-0016395.
We thank our advisors Stefano DellaVigna, Ned Augenblick, Dmitry Taubinsky, Ulrike Malmendier, and Chris Walters for their immensely valuable feedback. We further thank Luisa Cefala, Terry Odean, Gautam Rao, Matteo Saccarola, Afras Sial, Sharath Sonti, Karthik Tadepalli, and Carly Urban for helpful comments. We are also grateful to Andrew Eppig and Jennifer Lynn Hogg for their assistance with the data linkage and for providing valuable context on the Registrar and UC-CCP data. Elaine is especially grateful to the Institute for Business Innovation (IBSI), the UC Berkeley Experimental Social Science Laboratory, and the National Endowment for Financial Education (NEFE) for generous financial support. All errors are our own.

†Shen: UC Berkeley (elaineshen@berkeley.edu), Lyu: UC Berkeley (jrlyu823@berkeley.edu), Voss: UC Berkeley (lailavoss@berkeley.edu).

1 Introduction

Standard economic models often assume that individuals make investment decisions with complete information and rational expectations. As a result, models of educational attainment predict that individuals invest until marginal benefit equals marginal cost; namely, that individuals with the greatest potential benefits are most likely to invest in education (Willis and Rosen, 1979; Card, 1999). Existing literature suggests this is also true with respect to financial literacy education; better educated, higher-income individuals may have the most to gain and therefore may be most likely to invest in acquiring financial knowledge (Lusardi et al., 2017).

However, several frictions may prevent individuals who *would* benefit from financial education from investing optimally. For example, it may be difficult for individuals to correctly value the benefits of financial literacy education or predict the costs of acquiring financial knowledge when their starting level of financial knowledge is low. Individuals with limited financial knowledge may also misjudge their starting level of financial literacy. This could lead to underinvestment in financial education due to overconfidence if they believe they already know everything they need to, or underconfidence if they incorrectly assume they cannot learn effectively. Finally, individuals with low starting levels of financial literacy may also lack access to education resources.

This paper studies self-selection in financial literacy education – specifically, whether individuals correctly value financial education and choose to invest in it. First, we document selection on baseline ability in both observational data and in a lab setting. We show that individuals with higher levels of financial knowledge are more likely to invest in financial literacy education. Next, we use a lab experiment to measure demand for education and incentivize individuals to complete education within the experiment. This allows us to identify whether there is selection *into* education on learning gains or treatment effects. To evaluate whether new technologies can lower barriers to participating in financial literacy education, we vary the format and perceived cognitive difficulty of the education material offered. Experimentally varying the perceived cognitive cost of education allows us to better understand the mechanisms driving current selection patterns while exploring policy-relevant ways to change existing selection patterns.

We study self-selection into financial literacy education using two complementary approaches. The first uses rich observational data on UC Berkeley students to document selection into a large, highly rated introductory personal finance class. We combine survey measures of financial literacy with student records from 2010-2025 from the UC Berkeley Registrar's Office with individual-level credit panel data from the California Policy Lab (CPL). To estimate treatment effects for the class, we use inverse propensity score weighting to create a synthetic control group and estimate the treatment effect of the class using academic and credit outcomes. Finally, we control for the intensive margin or the time and effort put into the class using the time students spend on the course website. This allows us to control for another previously unobserved aspect of selection:

selection on attention.

To explore the mechanisms behind self-selection into financial education and to test for selection on treatment effects on a separate, more representative population, we move to a controlled lab setting. We conduct an online experiment with n=1,047 US participants on Prolific, a large online survey platform. In the experiment, each participant completes two sets of incentivized financial decision-making exercises. Between the two exercises, participants are given an opportunity to purchase different versions of financial education material and a portion of participants are randomly assigned to complete the education material. This allows us to precisely measure beliefs about ability and expected benefits from education. We then examine whether the true benefits from education are positively correlated with WTP, as standard models of investment in education would predict, or if individuals *mispredict* the benefits of education. Finally, we experimentally vary the perceived difficulty, or cognitive cost, of the education material to test whether new technology can lower barriers to entry and change existing selection patterns.

We document selection on observables into financial literacy education in both settings, among two distinct populations. We find that students who enroll in a large financial literacy class (UGBA 135) tend to be positively selected on family income, quantitative skills, and financial decision-making ability. Similarly, in the US Prolific population, individuals with the highest latent return to financial education have the lowest willingness to pay for education. Using measurements from the lab setting, we find that this is driven by systematic belief errors: low-ability individuals tend to underestimate the benefits from education and overestimate their starting level of knowledge. While technology-backed education solutions that lower the cost of education do increase demand compared to the traditional "classroom style" lecture education material, the AI and video options fail to alter selection patterns, indicating that additional behavioral barriers persist even when cognitive costs are reduced.

This paper contributes to a large literature on mistakes in financial decision-making and financial literacy education. Given the increasingly complex financial landscape, there is growing concern about the ability of households to make sound financial decisions. Studies have shown that individuals with lower levels of financial literacy have lower rates of retirement savings (Mitchell and Lusardi, 2015), higher costs of borrowing (Huston, 2012; Lusardi and de Bassa Scheresberg, 2013), lower rates of asset accumulation, and lower rates of investment in stocks and other assets with high rates of return (Lusardi and Mitchell, 2023; van Rooij et al., 2011). These studies also find that financial mistakes are more prevalent among low-income individuals and minorities, suggesting that disparities in financial literacy may exacerbate existing income and wealth gaps.

Despite the seemingly large benefits of financial literacy, several studies have pointed out difficulties encouraging participation in financial education (Bruhn et al., 2013; Lara Ibarra et al., 2017). Furthermore, the individuals who take up financial literacy education are often *not* those who

would ex ante have been expected to have the largest benefits. In fact, existing studies suggest that individuals who choose to invest in financial education are often *positively selected* along some dimension. For example, Meier and Sprenger (2013) find that individuals who choose to participate in a free credit counseling program tend to be more patient (less present-focused) than individuals who do not participate. Similarly, Bhattacharya et al. (2012) find that investors who opt into free investment advice tend to be more financially sophisticated—older, wealthier, and investors who hold more diversified, less home-biased portfolios—than those who do not. This paper advances our understanding of selection by documenting both selection on baseline ability and selection on treatment effects.

Selection may limit the efficacy of financial literacy education even when education is mandatory. For example, financial literacy education mandates may be less effective if the students who pay attention or exert effort in the class are the ones who have a higher starting level of financial knowledge and therefore recognize the importance of acquiring additional knowledge. These selection patterns may also have distributional consequences, especially as states increasingly adopt high school financial literacy mandates and as financial literacy education becomes more prevalent on college campuses (Luedtke and Urban, 2023; Urban et al., 2020). If high-income, high-literacy students disproportionately invest time and effort into improving their personal finances then state-mandated publicly funded financial literacy programs may actually contribute to increasing income inequality by teaching wealthy individuals how to grow their wealth.[1]

This paper proceeds as follows. Section 2 describes the observational data used to identify selection on observables and presents preliminary results. Section 3 describes the lab experiment used to decompose mechanisms. We then test for selection on treatment effects and decompose the mechanisms driving selection. Finally, Section 4 concludes and discusses areas for future research.

2 Observational Data

2.1 Data and Empirical Strategy

We use two complementary approaches to understand the mechanisms behind self-selection in financial education. The first combines observational data from three sources to document selection and estimate heterogeneous treatment effects for a large personal finance class at UC Berkeley: Intro to Personal Finance (UGBA 135). The class is taught by Professor Terrance Odean and currently enrolls approximately 600 students per semester. The class is highly rated, available to any UC Berkeley student, and has no finance or math prerequisites. UGBA 135 has been offered since 2008 (and in its current form since 2013) and is often fully enrolled.

[1]Over 29 states in the US either currently have or plan to adopt a financial literacy mandate that requires students to take a financial literacy class before graduating from high school.

To build the dataset, we combine observational data from three sources: 1) course records from the UC Berkeley Registrar's Office, 2) credit bureau data from the University of California Consumer Credit Panel (UC-CCP), and 3) survey data from a personal finance survey fielded through the Associated Students of the University of California (ASUC).

Data from the Registrar's office includes anonymized individual-level degree, admissions, and course data for all undergraduate entry cohorts from 2007-2025.[2] This includes college application data (parental education level, the applicant's self-reported family annual income, University of California Application Fee Waiver status, and intended major), degree date (major/GPA at graduation, graduation date), and demographic data (gender, race/ethnicity, and age), transfer student status, course enrollment data, and student's GPA for each semester.

We are also interested in whether attention or effort (although endogenously allocated) impacts treatment effects from the class. To control for attention, we include data on the total amount of time (in minutes) that each student spends on the course website as a proxy for attention and effort. Notably, time spent on the course website is an especially relevant measure of attention for this class: the weekly homework assignment for the class involves watching pre-recorded videos and completing online multiple choice quizzes on the course website. The instructors also regularly post recordings of the in-class lectures, lecture notes, and budgeting and other financial resources on the course website.

To measure financial literacy knowledge using similar measures to our lab setting, we partner with the ASUC Financial Literacy Task Force to administer a survey to students at UC Berkeley. In the survey we measure financial knowledge and stress and attitudes towards personal finances. To measure financial knowledge, we ask the Big Three financial literacy questions and have students answer three hypothetical decision-making questions. Finally, we ask students to provide their hypothetical willingness to pay or accept for UGBA 135.

The hypothetical decision-making questions ask students to make decisions about investing (index funds), debt, and compound interest in scenarios where there is a financially dominant choice. Students are told to maximize the amount of money they have in the scenario. While unincentivized, these scenarios give us a measure of financial decision-making ability against an objective benchmark and allow us to measure financial mistakes as the amount of money the participants lose by choosing a dominated option. While both the willingness to pay elicitation and the financial decision-making questions are unincentivized, these questions are similar to those used in our experiment, and give us another measure of selection patterns.

Finally, to study the long-term financial impact of financial literacy education, we plan to link course information from the Registrar's office with the University of California Consumer Credit Panel (UC-CCP). The UC-CCP is a longitudinal panel of approximately 40 million consumers

[2]Transfer student data is from 2016-2025.

starting in 2004 and continuing quarterly. The sample comprises of anatomized credit records of a nationally representative 2% sample of U.S. adult consumers with credit records along with a full sample of 100% of Californians with credit histories, allowing us to observe outcomes for the majority of students at UC Berkeley.

2.2 Observational Data Results – Selection on Baseline Ability

Standard human capital accumulation theories suggest a negative correlation between baseline ability and demand for education, as those with the lowest starting knowledge typically have the lowest marginal cost (i.e. opportunity cost) and highest marginal benefits from education. However, both our results from the observational data and lab setting contradict this prediction.

We first document selection into the personal finance class at UC Berkeley. Using individual-level data, we compare the demographic and major characteristics of UGBA 135, the large personal finance class at UC Berkeley, against the broader undergraduate population. We focus on demographic variables and proxies for income such as Pell Grant status and family education background.

We find strong evidence of positive selection on observables. As shown in Table 1, students who enroll in the financial literacy course are significantly different from non-enrollees across several dimensions. Specifically, enrollees are 12 percentage points more likely to come from quantitative majors (e.g., math, economics, and data science) (78% vs. 66%, $p < 0.01$), tend to have higher cumulative GPAs, and are significantly less likely to be underrepresented minorities (12% vs. 21%, $p < 0.01$). Furthermore, students from higher socioeconomic backgrounds appear more likely to invest in financial education; enrollees are less likely to be low-income or Pell Grant Eligible (proxied by eligibility for UC Fee Waivers) and are less likely to come from families without a college degree when compared to the broader population of Berkeley students.

	All Mean (SE)	Fin Lit Class Mean (SE)	No Fin Lit Class Mean (SE)	Difference Δ	Reg P-value
Male	0.45	0.48	0.45	0.03	0.04
	(0.01)	(0.01)	(0.01)		
White	0.26	0.22	0.27	-0.05	0.00
	(0.01)	(0.01)	(0.01)		
Asian	0.41	0.54	0.38	0.15	0.00
	(0.01)	(0.01)	(0.01)		
Underrepresented Minority	0.20	0.12	0.21	-0.10	0.00
	(0.00)	(0.01)	(0.01)		
Quantitative Majors	0.68	0.78	0.66	0.12	0.00
	(0.01)	(0.01)	(0.01)		
Cumulative GPA	3.54	3.60	3.53	0.07	0.00
	(0.00)	(0.01)	(0.01)		
Parents - No College	0.16	0.13	0.17	-0.03	0.01
	(0.00)	(0.01)	(0.01)		
Parents - No 4Yr College	0.24	0.20	0.25	-0.05	0.00
	(0.01)	(0.01)	(0.01)		
UC Fee Waiver	0.22	0.18	0.23	-0.05	0.00
	(0.01)	(0.01)	(0.01)		
N	6,639	1,234	5,405		

Table 1: UC Berkeley Student Characteristics by Financial Literacy Class - Class of 2023

Notes: Table 1 shows demographics of students from the 2023 graduating class who did and did not take a large personal finance class at UC Berkeley. Each row shows the results from a one-sample t-test comparing students who take a large personal finance class at UC Berkeley (UGBA 135) to students who do not take the class from the same graduating year. Negative numbers and p-values less than 0.10 are highlighted in red. Students who self-selected into the personal finance class tend to have higher cumulative GPAs and are more likely to have quantitative majors, less likely to be underrepresented minorities, less likely to be first-generation college students, and less likely to qualify for a Pell Grant or UC Fee Waiver (a proxy for low-income status).

Using data from the time spent on the course website, we also find evidence of selection on the intensive margin, or selection on attention. Again, we measure attention using the time (in minutes) spent on the course website. Unlike other courses that do not actively use their website, the majority of homework assignments for this class are completed on the course website. The instructors also regularly post recordings of in-class lectures, lecture notes, and other financial resources on the course website.

We find that even among the positively selected students who choose to enroll in UGBA 135, there are large differences in the amount of time students spend on the course website. These differences also tend to be correlated with education outcomes – students who are above the median in time spent on the class website tend to earn a higher grade in the class (UGBA 135 GPA) and have higher GPAs in general (Term GPA). These differences highlight selection on attention as another important dimension and may help explain the limited efficacy of financial literacy mandates.

	All Mean (SE)	High Attention Mean (SE)	Low Attention Mean (SE)	Difference Δ	Reg P-value
UGBA 135 GPA	3.63	3.72	3.53	0.18	0.00
	(0.02)	(0.02)	(0.03)		
Term GPA	3.46	3.54	3.37	0.17	0.00
	(0.02)	(0.03)	(0.03)		
Male	0.51	0.48	0.55	-0.07	0.02
	(0.02)	(0.02)	(0.02)		
White	0.19	0.16	0.22	-0.06	0.03
	(0.01)	(0.02)	(0.02)		
Asian	0.52	0.56	0.48	0.08	0.01
	(0.02)	(0.02)	(0.02)		
Underrepresented Minority	0.17	0.16	0.18	-0.02	0.38
	(0.01)	(0.02)	(0.02)		
Quantitative Majors	1.00	1.00	1.00	0.00	.
	(0.00)	(0.00)	(0.00)		
Parents - No College	0.19	0.20	0.17	0.03	0.27
	(0.01)	(0.02)	(0.02)		
Parents - No 4Yr College	0.27	0.30	0.25	0.04	0.15
	(0.01)	(0.02)	(0.02)		
UC Fee Waiver	0.22	0.25	0.20	0.04	0.13
	(0.01)	(0.02)	(0.02)		
N	977	489	488		

Table 2: Student Characteristics by Attention in UGBA 135 - 2023-2024 Academic Year

Notes: Table 2 shows student characteristics by attention for students who took UGBA 135 during the 2023-2024 academic year. Attention is measured as the amount of time (in minutes) the student spent on the UGBA 135 course website, which includes time spent watching videos and completing comprehension check quizzes, watching lecture videos, and reviewing course notes and other materials. Students are in the high attention category if they are above the median in time spent on the course website. Students in the low attention category are below the median. Note that time spent on the course website may be missing for a small number of students. Negative numbers and p-values less than 0.10 are highlighted in red.

We also provide suggestive evidence of selection on financial decision-making ability. Using the ASUC survey data, we examine the relationship between the hypothetical willingness to pay (WTP) for the course and performance on hypothetical financial decision-making questions (e.g., maximizing returns in compound interest questions and debt repayment tasks). Figure 1 plots these results. We find that WTP is positively correlated with financial literacy scores: students with the highest demand for the course are arguably students who already possess high levels of financial knowledge. While the decisions in the survey are unincentivized, the correlation raises further concerns that students who stand to benefit the most from financial literacy education – students starting with lower levels of financial knowledge or from less privileged backgrounds – may not be the ones self-selecting into the class.

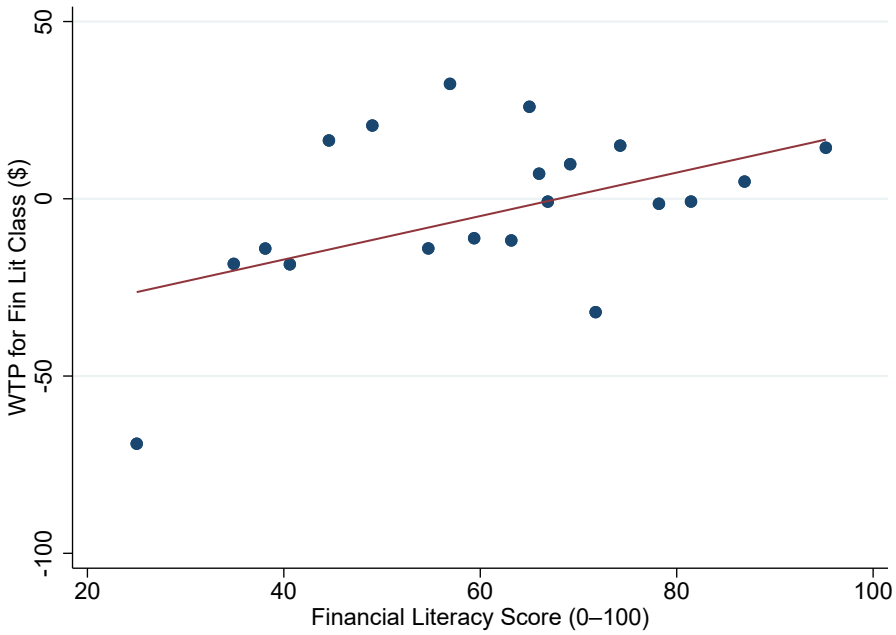


Figure 1: WTP for UGBA 135 versus Financial Decision-Making Score

Notes: Figure 1 shows the results from the ASUC survey conducted at UC Berkeley. We plot the relationship between the students' score on a set of financial decision-making questions and the hypothetical demand for UGBA 135. Both the decision-making questions and the questions are hypothetical and unincentivized, however the results provide additional suggestive evidence of the positive selection trend we observe in the observational data.

Taken together, the observational evidence suggests strong positive selection into financial education along observable characteristics and financial decision-making ability. However, these patterns are inherently descriptive. In the administrative data, demand for education is jointly determined by expected returns, perceived costs, institutional constraints, peer effects, and information frictions, making it difficult to isolate the underlying mechanism driving selection. In particular, we cannot observe individuals' beliefs about their own baseline ability or the returns to education, nor can we measure how these beliefs relate to actual learning outcomes. Our controlled lab experiment, described below in Section 3, helps us address these limitations and isolate the relationship between baseline ability, expected returns, and demand for education.

3 Lab Experiment

3.1 Data and Empirical Strategy

To demonstrate that the selection patterns we observe are not limited to classroom settings and to better understand mechanisms, we turn to our lab experiment. We start with a sample of $1,159$ Prolific participants who are 18 years or older and located in the United States. After excluding participants with crossover switching in the willingness-to-pay multiple price lists, who are flagged

by our invisible text as having used AI to complete their written responses, and participants who have their education material choices implemented, we are left with a final sample of n=1,047 observations. Participants in our experiment have an average age of about 43 years old, are 66% white, and have an average income of about $65,000. The key demographic characteristics of our final sample are presented in Table 3.

	Obs.	Mean	SD	P(25)	P(75)
Age (yrs)	1,047	42.79	12.48	33.00	51.00
Male	1,047	0.43	0.50	0.00	1.00
Time (in minutes)	1,047	37.32	17.42	25.65	44.82
Race - White	1,047	0.66	0.47	0.00	1.00
Race - Hispanic	1,047	0.07	0.25	0.00	0.00
Race - Asian	1,047	0.17	0.37	0.00	0.00
Race - Black	1,047	0.06	0.24	0.00	0.00
Income ($ est)	1,047	65031.04	49604.36	30000.00	87500.00
College Degree	1,047	0.68	0.46	0.00	1.00

Table 3: Summary Statistics for Experiment Participants

Notes: Table 3 presents the summary statistics of the demographics of the participants in the lab experiment. Participants are recruited through Prolific, a large online survey platform commonly used by researchers. All participants must be 18 years or older and located in the United States. We exclude participants with crossover switching in the willingness-to-pay multiple price lists, participants who are flagged by our invisible text as having used AI to complete their written responses, and participants who have their education material choices implemented.

In the experiment, each participant completes two sets of incentivized financial decision-making exercises. In each financial decision-making exercise, participants choose between two financial options: a fixed dollar amount, or an investment in which the starting amount and growth rate are varied. We then elicit demand for each version of the education material. 10% of participants have their choices implemented and 90% of participants are randomly assigned to one of the four experiment arms: 1) AI education, 2) Video education, 3) University Lecture education, or 4) control. Finally, all participants complete an endline set of financial decisions and are asked multiple choice questions about their experience with financial education and their use of AI in financial decisions. We also collect information on age, gender, race, and income. Figure 2 provides an overview of the experiment.

Demand for education, or willingness to pay (WTP) or willingness to accept (WTA), is measured using three Becker, Degroot, and Marschak (BDM) Mechanism (Becker et al., 1964) multiple price lists. Participants are given the choice between the following options:

A) Dollar amounts ranging from $10 to +$10 + no education material, or

B) Dollar amounts ranging from -$10 to +$10 + complete the education material

If participants receive the education material, they must fully complete the material by passing attention checks and correctly answering three multiple choice comprehension check questions

before they are able to move on in the experiment.

We then test whether these selection patterns are in part driven by the perception that traditional versions of financial literacy education may be unengaging or cognitively difficult. We offer three different formats of the education material: 1) a classroom style university lecture format, 2) a video education format, and 3) an interactive AI-powered learning assistant. The video and AI based learning formats allow us to test whether supply-side policies making education more accessible and less difficult can change selection patterns. All versions of the education material teach the exact same material; the Rule of 72, a simple mental shortcut for calculating the value of the investment options. We intentionally try to hold all versions of the education material constant by keeping the definitions, comprehension check questions, and explanations to the multiple choice questions exactly the same. We also emphasize to participants that all of the materials are based on the same content and vary only in presentation format. The university lecture format is framed as the most difficult version while the AI-powered learning assistant is framed as the most beginner friendly, interactive, and easy to understand version. We describe the video education version using similar language and framing as popular online learning platforms like Khan Academy and Coursera.[3] Screenshots of the different versions of the education material are included in Appendix A.1.

We also measure each participant's beliefs about their starting or "pre" score (measured as their expected bonus), their expected bonus if they receive each version of the education material (their perceived benefit), and the reasons they prefer each version. Eliciting expectations also allows us to measure how overconfident participants are in their own knowledge. Finally, we ask the participants a set of five unincentivized multiple choice questions on compound interest from Ambuehl et al. (2022), as well as questions about their experience with financial literacy education, personal finances, and AI usage.

The treatment condition is assigned based on the actual WTP/WTA for 10% of participants. The remaining 90% are randomly assigned to one of the four treatment arms; 1) AI education material, 2) video education material, 3) university lecture education material, or 4) no education material. Randomly incentivizing participants with both positive and negative WTPs to receive financial literacy education allows us to identify whether the true benefits from education are positively correlated with demand, as standard models of investment in education would predict, or if individuals actually *mispredict* the benefits of education. Payment for the education bundle is implemented only if they receive a bonus for their endline decision. This ensures that the costs or benefits education are realized *only if* the participant receives the bonus. Table 4 shows the balance table for the four treatment arms.

[3]https://www.khanacademy.org/college-careers-more/personal-finance/x12915ee2df79104f:welcome-to-personal-finance/x12915ee2df79104f:untitled-171/a/welcome-to-personal-finance-v2

	Control		AI Edu		Video Edu		Lecture Edu	
	Mean	(SD)	Mean	(SD)	Mean	(SD)	Mean	(SD)
Male	0.46	(0.50)	0.44	(0.50)	0.42	(0.49)	0.39	(0.49)
Age	42.34	(12.45)	41.93	(12.23)	43.60	(12.41)	43.22	(12.82)
Time (mins)	28.44	(15.01)	41.46	(18.83)	39.58	(17.57)	39.62	(15.03)
Race - White	0.68	(0.47)	0.65	(0.48)	0.61	(0.49)	0.70	(0.46)
Race - Black	0.05	(0.23)	0.07	(0.25)	0.08	(0.27)	0.05	(0.22)
Income ($ est)	65830.08	(49237.02)	62263.78	(46519.61)	69606.74	(55007.13)	62351.85	(46974.71)
College	0.66	(0.47)	0.69	(0.47)	0.73	(0.45)	0.66	(0.47)
Observations	256		254		267		270	

Table 4: Balance Table by Treatment Arm

Notes: Table 4 presents the balance table of the key participant demographics by treatment arm. Treatment status is randomly assigned by the computer at the individual level.

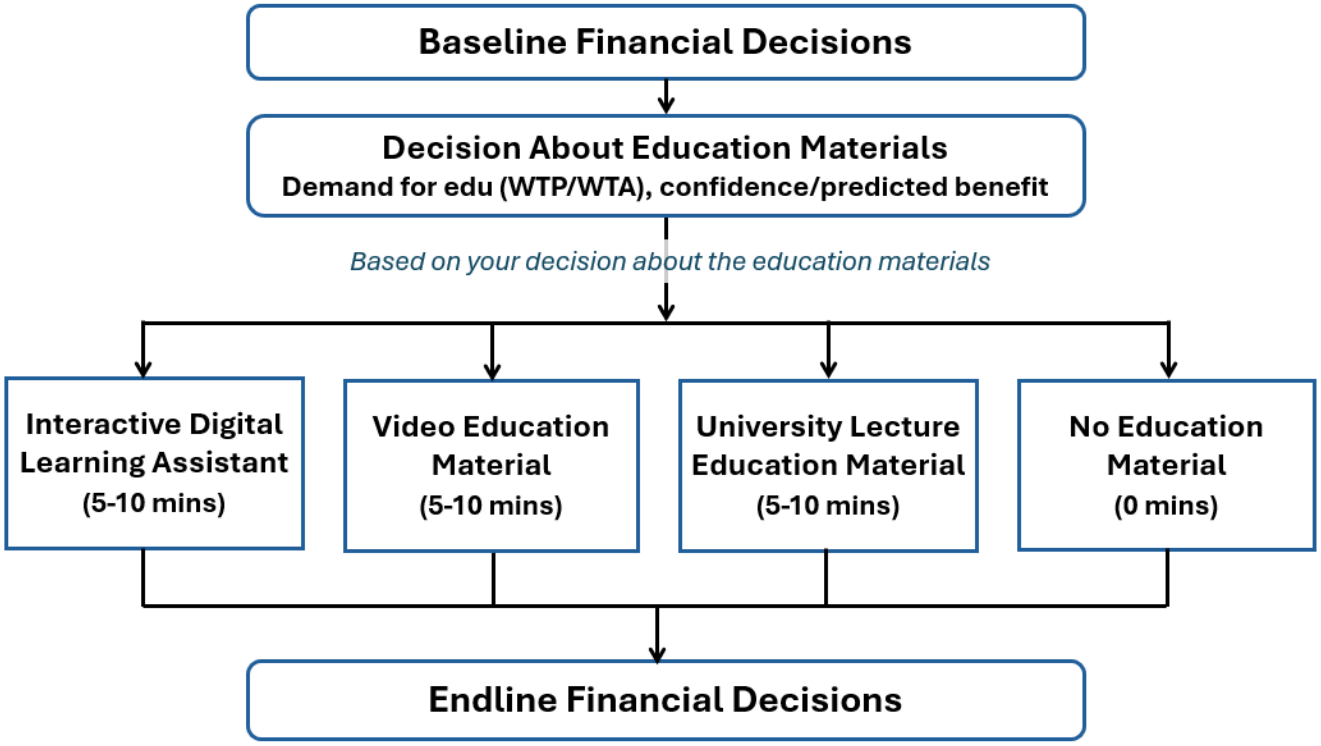


Figure 2: Lab Experiment Overview

Notes: The diagram in Figure 2 provides an overview of the experimental design for the lab experiment. Participants were first asked to answer a set of 10 baseline financial decision-making questions that test the participant's knowledge of the compound interest. We then ask participants to predict how they performed on the 10 baseline questions, elicit a willingness to pay for each version of the education material, and ask participants to predict how they would do on the endline questions *if they were to receive the specified version of the education material.* 90% of participants are then randomized into one of four experimental arms: the AI education material, video education material, lecture education material, or control (no education material). Finally, all participants answer a set of 10 endline financial decision-making questions, which test the same concepts and are very similar in structure and difficulty to the baseline questions.

One challenge the financial education literature has often faced is the question of how to measure financial literacy in objective, welfare-relevant terms. This stems from the fact that financial literacy encompasses a wide variety of skills, including numeracy, statistics and probability, attention to detail, knowledge of financial terms and definitions, and the ability to read and process complex financial contracts. Additionally, in many cases these shortcomings can be overcome through exerting effort (for example, looking up the concepts online, or following the advice of knowledgeable friends or experts). This raises the question of whether teaching financial literacy is simply increasing motivation or teaching individuals to exert effort. Finally, there is the question of how knowledge of financial concepts are translated into real-world behavior that improves welfare, as measured against an objective benchmark.[4] In fact, there is evidence that being able to *recall* financial concepts is not the same as being able to *apply* these concepts and translate them into improvements in financial decision-making (Ambuehl et al., 2022).

Acknowledging this difficulty, we choose to focus on what we believe is one of the core components of financial literacy: numeracy. Specifically, our questions test for an understanding of interest rates and compound interest. We focus on these concepts for two main reasons. The first is that we believe that quantitative principles are a core component of making good financial decisions. In fact, "The Big Three," the widely-used, standard measures of financial literacy developed by Lusardi and Mitchell (2011), are at their core questions about numeracy.[5] The second benefit is that focusing on numeracy allows us to measure financial knowledge in a way that directly maps to monetary gains and losses. This allows us to measure and quantify the benefits of our financial education intervention using an objective, welfare-relevant metric.

Our question is based on the methodology developed in Ambuehl et al. (2022). We ask participants to choose between two financial instruments: one that is a complexly-framed dollar amount and one that is a simply-framed dollar amount. Participants are told that there is a chance $1 will be added to their bonus for each decision where they choose the financial instrument that is worth more money in 72 days. Since there are ten questions in total, the potential bonus ranges from $0 to $10. We also tell participants that there is no interest rate or investment risk and no uncertainty. Figure 3 shows an example of the question.

Of the two financial instruments, there is one option that maximizes the amount of money the participant will receive at the end of the study. Since one option is framed as a compound interest calculations, participants should calculate the dollar value of the complexly framed option, thus testing their knowledge of compound interest. As in Ambuehl et al. (2022), the complexly framed amount can be calculated quickly using the Rule of 72, a mental shortcut that allows you to quickly

[4]Early papers focus on outcomes such as saving which may not be universally welfare improving. See Bernheim and Taubinsky (2018) for a discussion.

[5]Two of the Big Three questions test for an understanding of interest rates (one directly and one indirectly by asking about inflation). The third question asks about diversification, a statistical concept that is closely related to numeracy.

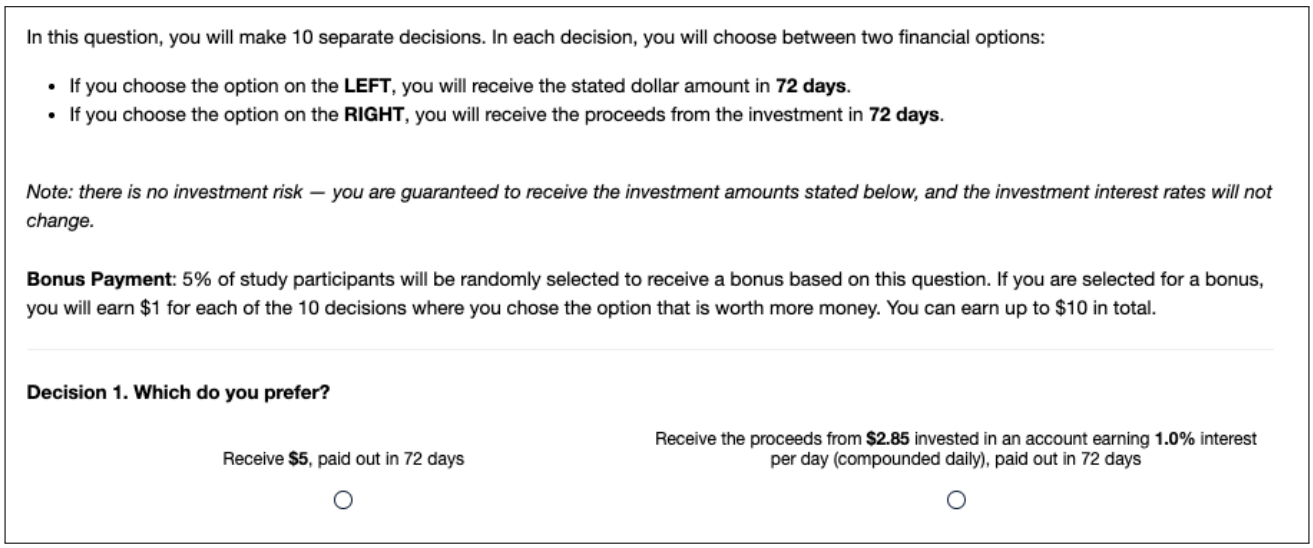


Figure 3: Financial Decision-Making Question

Notes: Figure 3 shows the financial decision-making question used in the lab experiment. The question is based on a core concept in personal finance; understanding compound interest and exponential growth and is based on a similar question in Ambuehl et al. (2022).

estimate how long it would take a dollar amount to double at a fixed interest rate. We then define the change in score (and educational benefit) as the endline score (0-10) minus the baseline score (0-10).

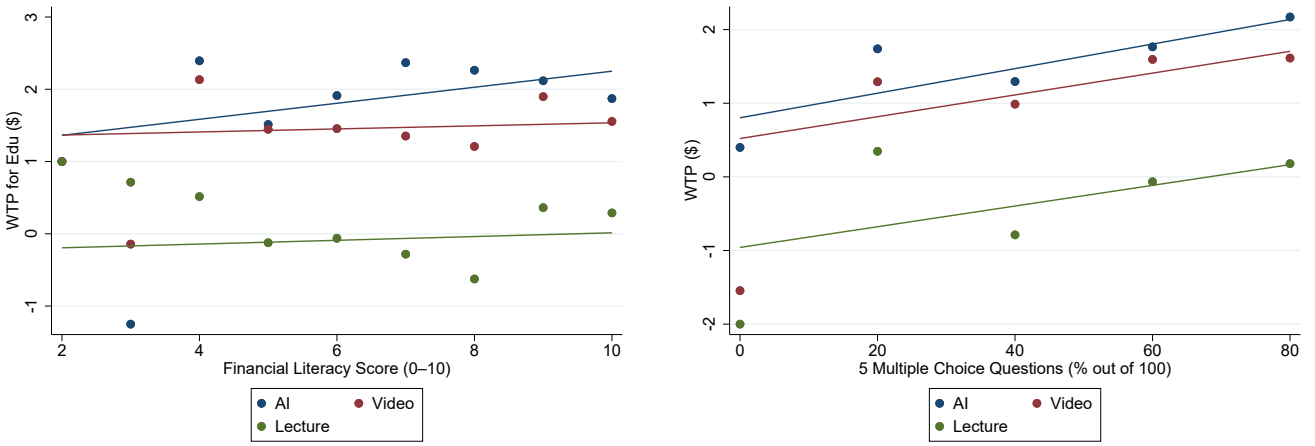
3.2 Results – Lab Experiment

3.2.1 Selection on Treatment Effects

We first seek to replicate the selection patterns we see in the observational data in our lab setting. One potential concern is that a controlled setting may attenuate the selection mechanisms we observe in the real world. Beyond the ability to correctly predict the costs and benefits of education, real-world selection patterns may encompass factors such as present focus, uncertainty about the future, and concerns about the decay of knowledge that we control for in the lab. For example, we remove any uncertainty about the potential benefits of education by telling participants they will use the material immediately after (in the endline questions). We may also reduce uncertainty about the quality of the education material by providing it as a part of the experiment, essentially signaling its usefulness. Finally, we remove any concerns about knowledge decays over time, thereby increasing the potential benefits of the education material, by asking the endline questions immediately after the education material. Thus, it would be unsurprising if any selection patterns are more muted, or even non-existent, in a lab setting.

Despite these limitations, we still find evidence that demand for education has a slight positive correlation with baseline financial knowledge in our lab setting. As illustrated in Figures 4(a)

and 4(b), participants with higher scores on the baseline "Rule of 72" decision-making questions and five financial literacy multiple choice questions about compound interest exhibit higher WTPs for the education modules. Conversely, individuals with the lowest baseline performance—who theoretically stand to gain the most—report the lowest WTP.



(a) Rule of 72 Questions

(b) Compound Interest Multiple Choice Questions

Figure 4: Selection on Baseline Ability

Notes: Figures 4(a) and 4(b) show selection on baseline financial decision-making ability. Figure 4(a) shows the relationship between WTP for each version of the education material and baseline financial decision-making ability, measured as the baseline or pre score on the Rule of 72 questions. The Financial Literacy Scores can range from zero to ten (one point for each question correct). Figure 4(b) shows the relationship between demand for education and the participant's score on five multiple choice questions about compound interest. These five questions are taken from Ambuehl et al. (2022) and are unincentivized. We elicit demand for each of the 3 different versions of the education material: 1) the AI version, 2) the video version, and 3) the university lecture version.

To estimate the treatment effect of the education material, we run the following specification, which corresponds to Column (1) of Table 5:

$$Y_{i,1} = \alpha + \beta_1 T_i + \beta_2 Y_{i,0} + \mathbf{X_i} + \varepsilon_i \tag{1}$$

where $Y_{i,1}$ is endline score, $Y_{i,0}$ is baseline score, T_i is an indicator for treatment status, and $\mathbf{X_i}$ is a vector of baseline demographic controls for participant i. Baseline demographic controls include include college, gender, race, income, and age. Column (2) of Table 5 runs the same specification with an indicator for each education material type.

To examine treatment effect heterogeneity, we augment Equation 1 and estimate the following specification, which corresponds to Column (3) of Table 5:

$$Y_{i,1} = \alpha + \beta_1 T_i + + \beta_2 Y_{i,0} + \beta_3 H_i + \beta_4 (T_i \times H_i) + \varepsilon_i \tag{2}$$

where $Y_{i,1}$ is endline score, $Y_{i,0}$ is baseline score, T_i is an indicator for treatment status, H_i is the

dimension of heterogeneity we are interested in, and $\mathbf{X_i}$ is a vector of baseline demographic controls for participant i. We look at treatment effect heterogeneity by education material and by baseline ability.

Table 5 reports treatment effects on endline test performance using an ANCOVA specification that controls for baseline ability. Across all specifications, assignment to AI, video, or lecture education leads to large and statistically significant improvements in test scores. Pooling across formats yields an average treatment effect of approximately 1.3 points (s.e. = 0.117). The coefficient on baseline performance indicates substantial persistence in test scores, while remaining well below one. Columns (5) and (6) show significant heterogeneity in treatment effects: the interaction between treatment assignment and baseline score is negative, implying that educational interventions generate larger gains for participants with lower initial ability. These results establish that the interventions are effective and that their benefits are concentrated among participants who start from a lower baseline.

	(1) Post Score	(2) Post Score	(3) Post Score
Pre Score	0.431***	0.429***	0.619***
	(0.028)	(0.028)	(0.050)
Any treatment = 1	1.311***		2.944***
	(0.117)		(0.400)
AI treatment = 1		1.454***	
		(0.150)	
Video treatment = 1		1.318***	
		(0.145)	
Lecture treatment = 1		1.171***	
		(0.153)	
Pre Score × Any treatment			-0.251***
			(0.059)
Constant	3.932***	3.939***	2.718***
	(0.301)	(0.301)	(0.396)
Observations	1047.000	1047.000	1047.000
R^2	0.237	0.239	0.246
Pre score control	X	X	X
Demographics controls	X	X	X

Table 5: Treatment Effects of Education Material

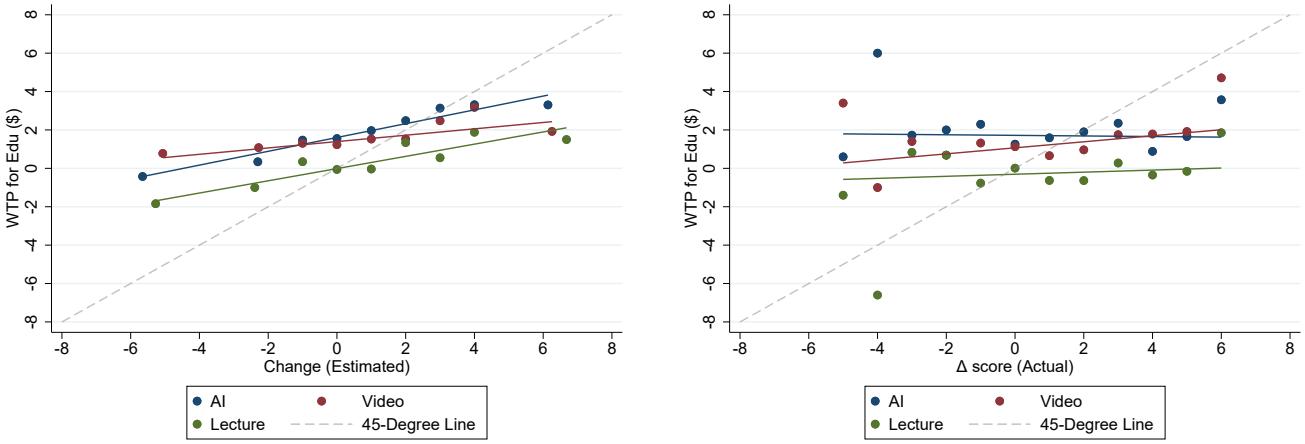
Notes: This table reports treatment effects on endline test performance using an ANCOVA specification. All columns regress the post-treatment test score on indicators for assignment to educational interventions, with the control group omitted. Column (1) pools all interventions into a single "Any treatment" indicator. Column (2) includes separate indicators for AI, video, and lecture treatments. Column (3) includes the specification in Column (1) but additionally allows treatment effects to vary with baseline test performance by interacting treatment assignment with the pre-treatment score. Baseline test scores are included in all specifications as "Pre Score." All columns additionally control for demographics (college education, gender, race, income, and age). Robust standard errors are reported in parentheses.

3.2.2 Mechanisms: Belief Errors and Misprediction

Another central prediction of human capital theory is that individuals have rational expectations and continue to invest in education while the marginal benefits (or the gains from education) equal the expected marginal costs. To test these patterns, we turn to a laboratory setting where our experimental design allows us to test the accuracy of these predictions by measuring key variables that we cannot observe in administrative datasets: financial ability, beliefs, and treatment effects.

We first compare participants' willingness to pay (WTP) against both their predicted score improvement (perceived return) and their actual score improvement (realized return). Figure 5 shows the relationship between demand for financial education and the predicted benefits of education. The contrast between these two relationships is striking. As shown in Figure 5(a), demand for financial education is strongly and positively correlated with *predicted* benefits; participants who believe they will learn more are willing to pay more.

However, Figure 5(b) highlights the disconnect between beliefs and reality. While there is a positive relationship between expected benefit and demand, there is essentially zero correlation between actual gains from education and WTP for education. This implies that individuals are essentially "optimizing on noise"; given incorrect beliefs, individuals are not able to direct resources in a way that lines up with true benefits.



(a) WTP versus Predicted Change

(b) WTP versus Actual Change

Figure 5: WTP versus Predicted and Actual Change

Notes: Figures 5(a) and 5(b) show the relationship between participant demand (WTP) for the different types of education material versus their predicted and actual change in score. Figure 5(a) shows that demand is increasing in estimated or predicted benefit, confirming that the WTP elicitation is positively correlated with the perceived benefits from education, however Figure 5(b) shows that WTP is much less correlated with the *actual* change in score. This suggests that individuals have difficulty predicting the *true* benefits from education.

To test whether prediction errors vary by baseline ability, Figure 6 plots the perceived change in score from the education material and actual change in score versus baseline ability. The figure highlights that the prediction error, or the gap between actual change in score and expected change in score, is larger for individuals with lower starting scores and decreases as baseline score increases. Although this pattern is partly driven by the natural cap on the scores at 10 points (given the 0-10 range), this is consistent with the idea that individuals with higher levels of financial knowledge have a better understanding of their abilities.

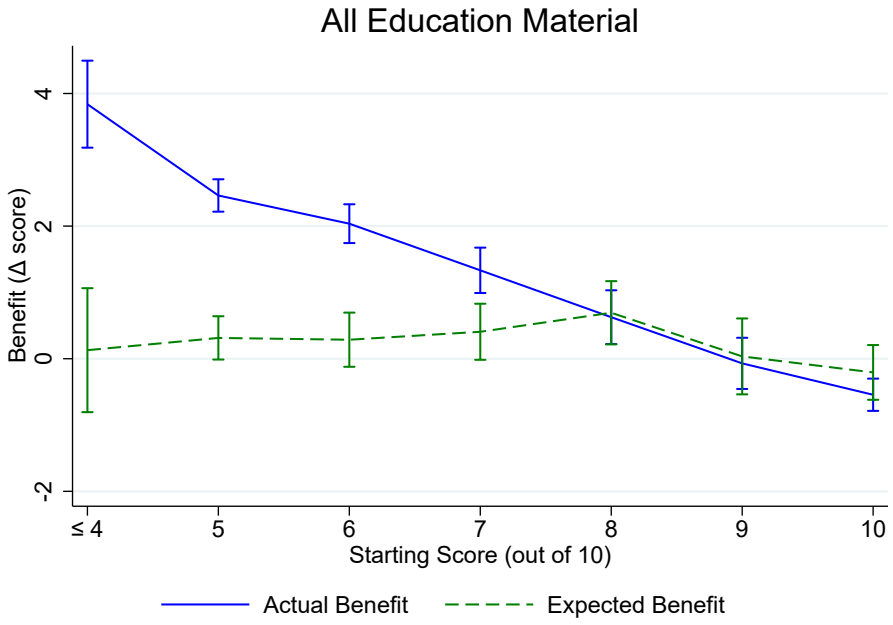


Figure 6: Actual versus Expected Benefit by Ability

Notes: Figure 6 plots expected or perceived benefits from all versions of the education material (in change in score) against the actual benefits from education by baseline ability. The figure shows that forecast errors, or the gap between the actual and perceived benefit, are smaller for higher ability individuals. This is consistent with the idea that individuals with higher levels of financial knowledge have a better understanding of their abilities. This figure pools all the versions of the education material together and only shows the treated individuals. Figures for each version of the education material are included in Appendix A.2. Error bars show 95% confidence intervals.

Table 6 examines how actual learning gains (improvements in score between the pre and post tasks, subtracting out the average score of control participants with the same starting score) vary with baseline ability, willingness to pay, and expected benefits across the three educational formats. Across all materials, learning gains are highest among participants with the lowest baseline scores. In contrast, willingness to pay for a given material is only weakly related to realized learning gains. The estimated coefficient on willingness to pay is small and statistically insignificant for AI and lecture, and only marginally positive for video. Similarly, expected learning gains relative to the control group are not predictive of actual learning outcomes. Together, these results show that treatment effects are strongly heterogeneous and concentrated among lower-ability participants,

while willingness to pay is only weakly related to realized gains, consistent with demand being driven by perceived rather than actual returns.

	(1) T-C Change (AI)	(2) T-C Change (Video)	(3) T-C Change (Lecture)
Baseline number correct (pre)	-0.304***	-0.195***	-0.157***
	(0.054)	(0.055)	(0.060)
WTP for This Material	0.019	0.048*	-0.030
	(0.030)	(0.026)	(0.025)
Est. Change (rel. to control)	0.001	-0.036	0.026
	(0.039)	(0.049)	(0.055)
Constant	3.513***	2.638***	2.187***
	(0.422)	(0.410)	(0.451)
R^2	0.093	0.050	0.027
Observations	248	259	263

Table 6: Selection on Treatment Effects

Notes: Table 6 examines selection on treatment effects by relating actual learning gains to willingness to pay (WTP) and beliefs. The dependent variable in each column is actual learning, defined as the participant's endline score minus the average endline score of control-group participants with the same baseline score (i.e., baseline-score–adjusted treatment–control differences). Each column restricts the sample to participants who received the indicated education format (AI, video, or lecture). "Baseline number correct (pre)" measures initial performance on the financial decision-making task. "WTP for This Material" is the participant's elicited willingness to pay for the specific education format received. "Est. Change (rel. to control)" is the participant's expected learning gain relative to control, measured prior to treatment. Standard errors are reported in parentheses.

We further dissect this relationship in Table 7, where we regress actual and estimated baseline scores and benefits on demand (WTP) for different versions of the education material. Across all three formats, perceived benefit is the clearest and most consistent predictor of WTP. Participants who believed they would learn more were willing to pay more, with especially large effects for the AI-based lesson and smaller but still meaningful effects for video and lecture. Taking the average across all versions of the education material, a one-point increase in expected learning gains is associated with an increase in average willingness to pay of 29.2 cents (s.e. $=$ 0.076). This validates our willingness to pay measure and suggests that demand for education is "rational" conditional on beliefs.

Table 7 also highlights that perceived difficulty moves in the opposite direction. Participants who thought the education material would be more difficult had lower demand, particularly for the video and lecture formats. The estimate for AI is negative as well, although smaller and less precisely estimated. These patterns indicate that the cognitive costs of learning do influence demand for education, but to a lesser extent than the predicted benefits of education given the smaller coefficients.

Finally, Table 7 highlights that objective measures of learning (or the change in score) show much weaker relationships with WTP. Pre score and actual improvement have small and imprecise

coefficients, which is consistent with participants basing their WTP primarily on their beliefs about how much they would benefit, even when these beliefs did not reflect actual performance gains.

	(1) WTP for AI	(2) WTP for Video	(3) WTP for Lecture	(4) Average WTP
Baseline number correct (pre)	0.074	0.038	-0.063	0.007
	(0.095)	(0.086)	(0.092)	(0.080)
Estimated pre score	0.228***	0.183***	0.250***	0.200***
	(0.067)	(0.064)	(0.069)	(0.061)
Change in number correct	0.043	0.133*	0.040	0.066
	(0.082)	(0.078)	(0.082)	(0.071)
Estimated change	0.380***	0.246***	0.402***	0.292***
	(0.076)	(0.076)	(0.073)	(0.076)
Perceived difficulty (0 to 10)	-0.080	-0.133**	-0.276***	-0.071
	(0.055)	(0.056)	(0.062)	(0.066)
AI Belief Index (z)	1.855***	0.303	0.374*	0.894***
	(0.240)	(0.199)	(0.216)	(0.186)
Constant	-0.493	-0.012	1.895*	0.230
	(0.889)	(0.867)	(0.984)	(0.812)
R^2	0.127	0.035	0.080	0.053
Observations	1014	1020	1026	966

Table 7: Determinants of Willingness to Pay for Educational Materials (OLS)

Notes: Table 7 reports OLS regressions of willingness to pay (WTP) for each educational format (AI, video, lecture) and the average WTP across all formats per participant. Columns (1)–(3) correspond to WTP for AI, video, and lecture, respectively; Column (4) reports results for the average WTP across all three materials. *Baseline number correct (pre)* is the number of questions the participant got right on the first set of decisions. *Change in number correct* is the post–pre difference in actual performance. *Perceived difficulty* and *Estimated change* capture participants' beliefs about how hard the material is and how much they expected to improve. The *AI Belief Index* is the standardized average of participants' attitude and trust in AI. All regressions include demographic controls (college, gender, race, age, and income), which are omitted from the table for brevity. Standard errors are reported in parentheses.

Table 8 helps us understand what is driving the prediction errors. Column 1 of Table 8 shows the extent to which participants mispredict their starting score; for each additional point (or question) answered correctly, participants increase their estimate of their baseline score by only 0.32 points. Column 2 highlights an interesting heterogeneity: lower-ability participants systematically *overestimate* their baseline knowledge. We refer to this as overconfidence, following Definition One in Moore and Healy (2008). Individuals with lower starting ability are more likely to be overconfident – each additional point on the baseline quiz is correlated with a 0.68 point reduction in overestimation of the baseline score. The percent of overconfident individuals by baseline score is also shown in Figure 7.

Column 3 of Table 8 shows that for 1 additional point of actual improvement, predicted improvement only increases by 0.13 points. Comparing the coefficients in Column 1 and 3 suggests that participants have more accurate estimates of their starting score than of their change in score with

education. For every 1 point increase in baseline score, participants only increase their estimated baseline score by 0.321, but for every 1 point increase in their actual change in score, participants only increase their estimated change by 0.164 points. Finally, Column 4 shows that the participants who actually see the most improvement in score are more likely to underestimate their improvement: each additional point of actual improvement is correlated with an underestimation of 0.836 points.

Taken together, these results imply a misalignment between demand and returns: individuals who would benefit the most from education are not the ones with the highest willingness to pay. In particular, lower-ability participants tend to overestimate their starting level of knowledge and underestimate the potential gains from education. Combined, this leads to underinvestment in education. Higher-ability participants, whose beliefs are more accurately calibrated, express higher willingness to pay despite experiencing smaller realized gains.

	(1) Estimated pre score	(2) Overestimate pre (est-actual)	(3) Estimated change	(4) Underestimate change (actual-est)
Pre score	0.321***	-0.679***	0.043	-0.043
	(0.045)	(0.045)	(0.055)	(0.055)
Actual change			0.164***	0.836***
			(0.049)	(0.049)
Constant	3.302***	3.302***	-0.988**	0.988**
	(0.459)	(0.459)	(0.501)	(0.501)
R^2	0.054	0.178	0.016	0.301
Observations	1047	1047	1047	1047

Table 8: Actual and Perceived Ability and Learning Gains

Notes: Table 8 examines the relationship between actual and perceived ability and learning gains. Column (1) regresses participants' estimated baseline quiz score on their actual baseline score. Column (2) regresses the signed estimation error in baseline ability—defined as estimated pre score minus actual pre score—on actual baseline ability; positive values indicate overestimation. Column (3) regresses participants' estimated learning gains on their actual learning gains (change in number correct from pre to post). Column (4) regresses the signed estimation error in learning gains—defined as actual change minus estimated change—on actual learning gains; positive values indicate underestimation of learning gains. All regressions include controls for college education, gender, race, income, and age. Standard errors are reported in parentheses. Table A1 shows the same table split by high and low starting ability.

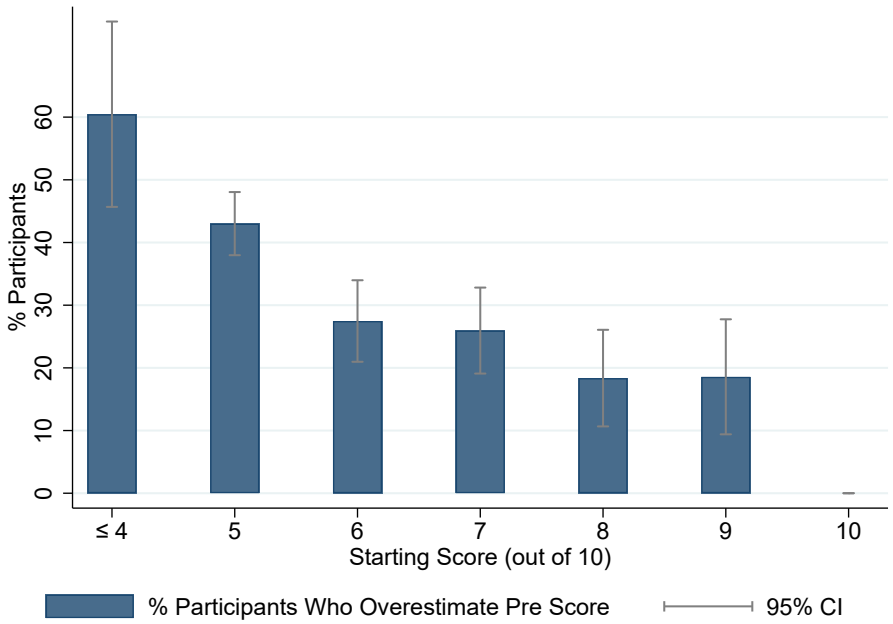


Figure 7: Overconfidence by Pre Score

Notes: Figure 7 shows the portion of participants who overestimate their starting score by pre score. The figure highlights that participants with lower pre scores are more likely to overestimate their starting score than participants with higher pre scores. Error bars show 95% confidence intervals.

3.2.3 Reducing the Costs of Education

We test whether lowering the perceived cognitive cost of education can correct these selection distortions by offering participants different versions of the education material. We then randomized participants into one of the three educational formats: a standard text-based university lecture (high cognitive cost), a video-based module, and an interactive AI-based learning assistant (low cognitive cost).

Figure 8 shows that we are successfully able to vary perceived difficulty of the education material. We elicit perceived difficulty, or perceived cognitive cost, using two measures: 1) how difficult participants perceive the education material to be, and 2) how much time they expect it will take them to complete and fully understand that version of the education material, both before completing the education material). Participants all find the AI material to be the least difficult version, the video education material to be the second most difficult, and the university lecture version to be the most difficult. Participant free-response answers about why they preferred specific materials shows that the primary reason for choosing the AI and video materials was because they were easier to understand and clearer. The share of participants who chose each reason can be found in Appendix Figures A.4 and A.5.

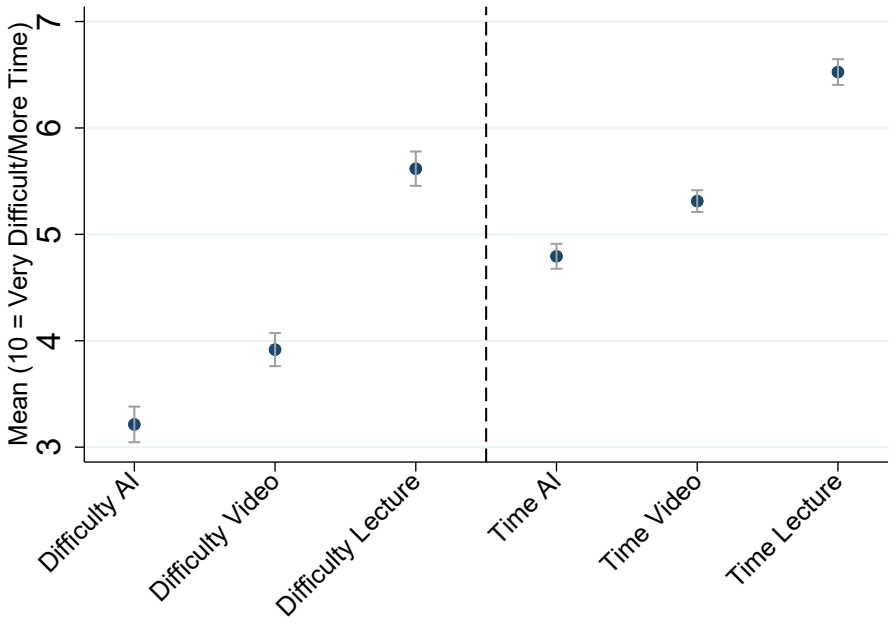


Figure 8: Perceived Cost of Education Material

Notes: Figure 8 shows the perceived difficulty, or cognitive cost, of the different versions of the education material. Participants perceive the AI education material to be the least cognitively difficult, the video education material to be slightly more difficult, and the lecture education material to be the most cognitively difficult. This is in line with the way we intentionally described the material.

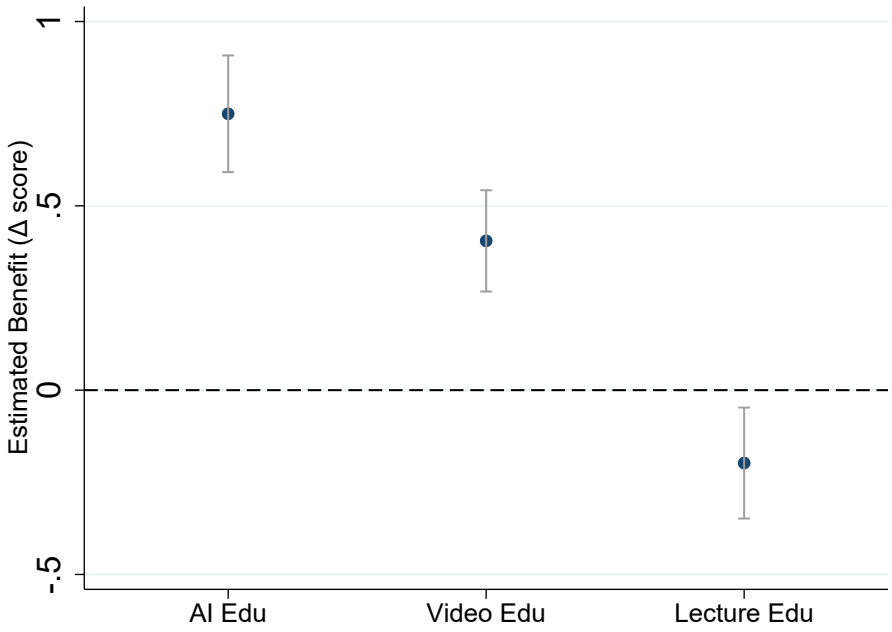


Figure 9: Estimated Benefit of Education Material

Notes: Figure 9 shows the estimated benefits from education for each version of the education material. Participants expect to see the largest changes in score from the AI education material, followed by the video education material and the lecture education material. This mirrors the perceived cognitive difficulty of the material. Notably, participants expect a positive change in score from the AI and video education materials and zero (or slightly negative) change in score from the lecture education material.

Figure 9 shows the estimated benefits from education for each version of the material and Figure 10 shows how demand for education varies with each version. We find that the video and AI-versions of the education material (the technology-backed versions of the education material designed to reduce cognitive difficulty) increase the perceived benefit of education and increase demand across all ability levels. This suggests that reducing the cognitive cost of education does increase aggregate demand, however, lowering the cost of education fails to correct the selection pattern. While demand increased across the board, the positive correlation between baseline ability and WTP persisted in the AI and video treatment arms. Low-ability individuals remained less willing to pay than high-ability individuals. The persistence of the selection pattern raises concerns that high ability individuals with higher demand will continue to crowd-out lower ability individuals, even in settings where financial literacy education is taught in an easy, beginner-friendly way. Our results suggest that cognitive cost is not the primary barrier preventing low-ability individuals from investing in education. Instead, the binding constraint appears to be the belief frictions identified above – specifically, mispredicting the benefits of education and an overconfidence in one's baseline ability.

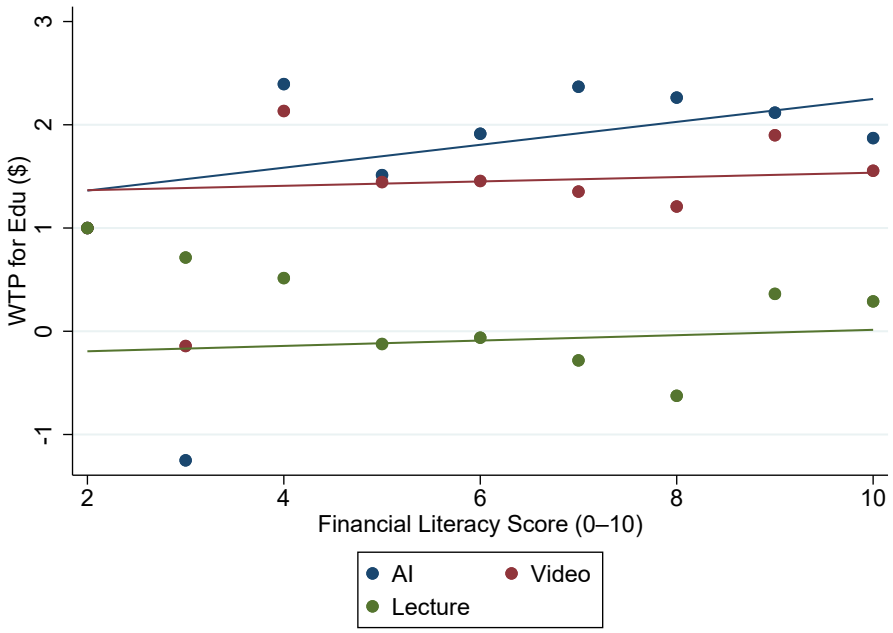


Figure 10: WTP by Baseline Ability for each version of the Education Material

4 Conclusion

This paper investigates a critical friction in the accumulation of financial knowledge: the self-selection mechanisms that deter high-return individuals from enrolling in financial literacy education. Standard human capital theory predicts that individuals will invest in education when the expected marginal returns exceed the marginal costs. Under this framework, individuals with the lowest baseline level of knowledge – and thus the highest potential returns to learning – should be the most eager to invest in education. Our findings challenge this assumption in the context of financial literacy education, documenting a troubling . pattern where the individuals who stand to benefit the most from education are the least likely to demand it.

Using administrative data from a large university, we first show that enrollment in financial education is characterized by strong positive selection. Students from higher socioeconomic backgrounds and those with stronger quantitative skills effectively select into the course, while underrepresented minorities and those from less privileged backgrounds select out. Our laboratory experiment isolates the mechanism driving this inefficiency. We find that low demand among individuals with lower starting levels of knowledge is not driven by rational cost-benefit analysis, but by systematic belief errors. Low-ability individuals are more likely to underestimate their returns to education. This is in part driven by an overconfidence in their baseline level of knowledge, which leads low-ability individuals to underestimate their potential for improvement.

Additionally, we show that supply-side interventions designed to reduce cognitive costs are insuffi-

cient to correct this market failure. While lowering barriers via AI and video-based learning tools does increase aggregate demand compared to traditional lecture-based learning, it fails to alter the selection patterns. The composition of the market remains skewed toward higher ability individuals, suggesting that the binding constraint is not on the difficulty of investing in education, but the misperception of its returns.

These results have important policy implications. Governments and institutions often rely on individuals to "opt-in" to financial education programs to improve decision-making and reduce inequality. Furthermore, even in cases where financial education is required, individuals have to "opt-in" on the intensive margin and choose to allocate attention and effort. Our findings suggest that self-selection patterns may render these education resources less effective, and may even exacerbate inequality by subsidizing the human capital accumulation of individuals who are already financially sophisticated. Because the friction is driven by systematic belief errors about potential benefits rather than access, standard subsidies or even technologies that lower the real and cognitive costs of participation are unlikely to solve the problem.

Ultimately, this paper highlights a critical friction in human capital accumulation. When returns to education are difficult to assess and beliefs are miscalibrated, individuals who are most in need of education may have the most difficulty evaluating the returns to education. While this paper focuses on financial literacy education, the implications are likely to be true for other domains such as mental health take-up, preventative health, or job retraining. Acknowledging those most in need of resources often have the most difficulty evaluating the benefits of said resources is an important step towards effective policy design and opens important areas for further research.

References

Ambuehl, Sandro, B. Douglas Bernheim, and Annamaria Lusardi, "Evaluating Deliberative Competence: A Simple Method with an Application to Financial Choice," *American Economic Review*, November 2022, *112* (11), 3584–3626.

Becker, Gordon M., Morris H. Degroot, and Jacob Marschak, "Measuring utility by a single-response sequential method," *Behavioral Science*, 1964, *9* (3), 226–232. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/bs.3830090304.

Bernheim, B. Douglas and Dmitry Taubinsky, "Behavioral Public Economics," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 1, Elsevier, 2018, pp. 381–516.

Bhattacharya, Utpal, Andreas Hackethal, Simon Kaesler, Benjamin Loos, and Steffen Meyer, "Is Unbiased Financial Advice to Retail Investors Sufficient? Answers from a Large Field Study," *The Review of Financial Studies*, April 2012, *25* (4), 975–1032.

Bruhn, Miriam, Gabriel Lara Ibarra, and David McKenzie, "Why is Voluntary Financial Education so Unpopular? Experimental Evidence from Mexico," May 2013.

Card, David, "The Causal Effect of Education on Earnings," in "Handbook of Labor Economics," Vol. 3, Elsevier, 1999, pp. 1801–1863.

Huston, Sandra J., "Financial literacy and the cost of borrowing," *International Journal of Consumer Studies*, 2012, *36* (5), 566–572.

Ibarra, Gabriel Lara, David McKenzie, and Claudia Ruiz Ortega, "Learning the Impact of Financial Education When Take-Up Is Low," November 2017.

Luedtke, Allison Oldham and Carly Urban, "High school financial education courses in the United States.: What is the importance of setting state policies?," *Journal of Financial Literacy and Wellbeing*, 2023, *1* (3), 431–449. Edition: 2024/02/15.

Lusardi, Annamaria and Carlo de Bassa Scheresberg, "Financial literacy and high-cost borrowing in the United States," Technical Report, National Bureau of Economic Research 2013.

_ **and Olivia S. Mitchell**, "Financial Literacy and Planning: Implications for Retirement Wellbeing," May 2011.

_ **and** _ , "The Importance of Financial Literacy: Opening a New Field," *Journal of Economic Perspectives*, November 2023, *37* (4), 137–154.

_ , **Pierre-Carl Michaud, and Olivia S. Mitchell**, "Optimal Financial Knowledge and Wealth Inequality," *Journal of Political Economy*, April 2017, *125* (2), 431–477.

Meier, Stephan and Charles D. Sprenger, "Discounting financial literacy: Time preferences and participation in financial education programs," *Journal of Economic Behavior & Organization*, November 2013, *95*, 159–174.

Mitchell, Olivia S. and Annamaria Lusardi, "Financial Literacy and Economic Outcomes: Evidence and Policy Implications," *The Journal of Retirement*, June 2015, *3* (1), 107–114.

Moore, Don A. and Paul J. Healy, "The trouble with overconfidence.," *Psychological Review*, April 2008, *115* (2), 502–517.

Urban, Carly, Maximilian Schmeiser, J. Michael Collins, and Alexandra Brown, "The effects of high school personal financial education policies on financial behavior," *Economics of Education Review*, October 2020, *78*, 101786.

van Rooij, Maarten C. J., Annamaria Lusardi, and Rob J. M. Alessie, "Financial literacy and retirement planning in the Netherlands," *Journal of Economic Psychology*, August 2011, *32* (4), 593–608.

Willis, Robert J. and Sherwin Rosen, "Education and Self-Selection," *Journal of Political Economy*, October 1979, *87* (5, Part 2), S7–S36.

A Appendix

A.1 Experimental Design Details

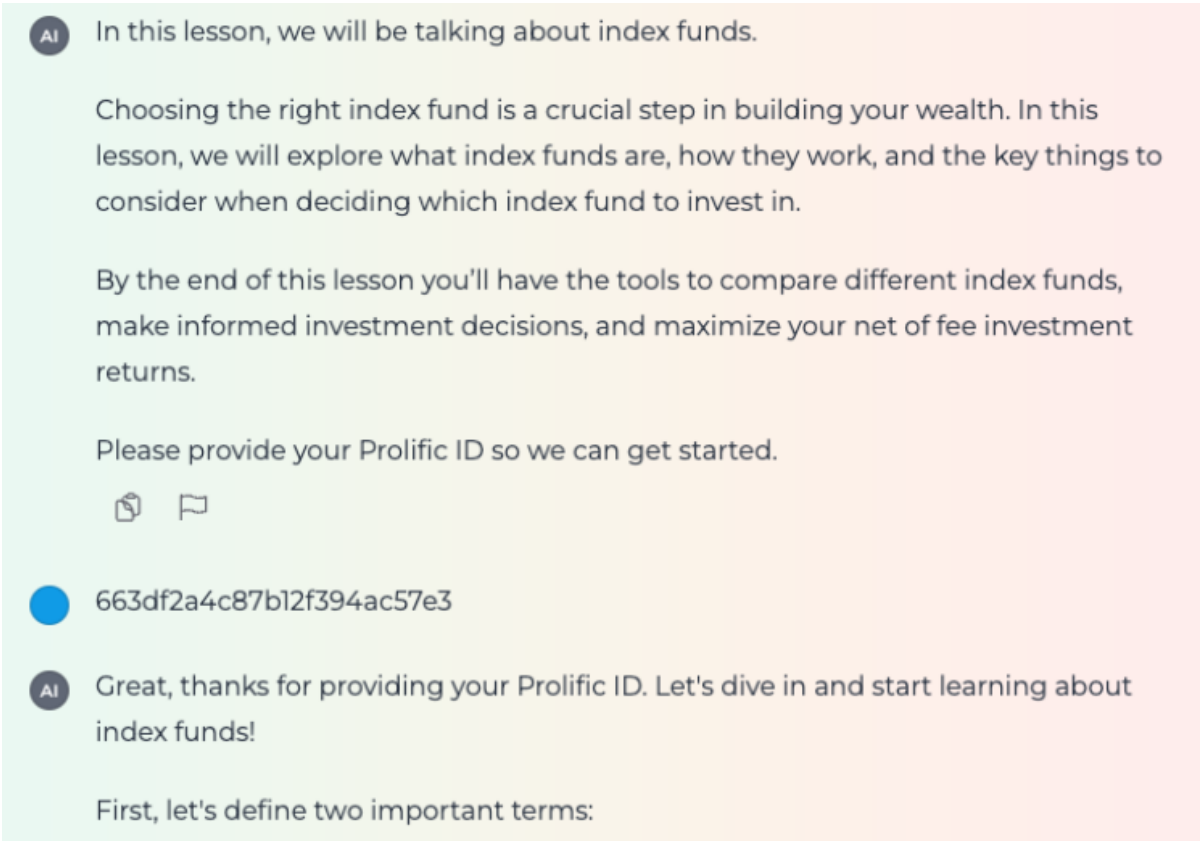


Figure A.1: Example AI Education Material – Interactive digital learning assistant that provides real-time personalized feedback



Question 1. True or False: If you invest in any S&P 500 index fund, you essentially own shares of all 500 stocks that make up the S&P 500 index.

- ○ True
- ○ False

Question 2. You are comparing two S&P 500 Passively Managed Index Funds, Fund A and Fund B. Fund A has a return since inception of 12.8%. Fund B has a return since inception of 7.3%. Both Index Funds track the same index (the S&P 500). How is this possible?

- ○ Index Fund A is trading with illegal insider information
- ○ Index Fund A is managed by better fund managers
- ○ This is impossible since both index funds are passive and track the same index (the S&P 500)
- ○ Index Fund A and B may have different start dates which lead to different returns since inception

Figure A.2: Video Education Material – a pre-recorded video, multiple choice questions with correct answers and explanations, and a PDF summary of key definitions

Index Funds Lecture Notes Outline

MBA-101A - Lecture 4

- Key Takeaways
- Definitions
  - S&P 500
  - Index Fund
  - Expense Ratio
  - Return Since Inception
- Main Formulas



Figure A.3: Lecture Education Material – Lecture notes, comprehension check questions, and a lecture video recorded from a personal finance course at a top university

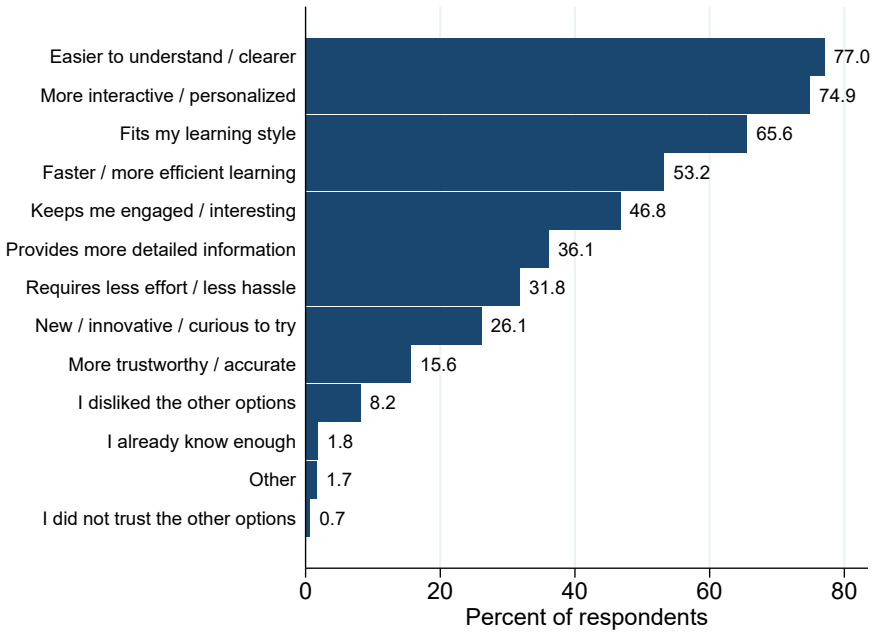


Figure A.4: Reasons for choosing the AI Education format

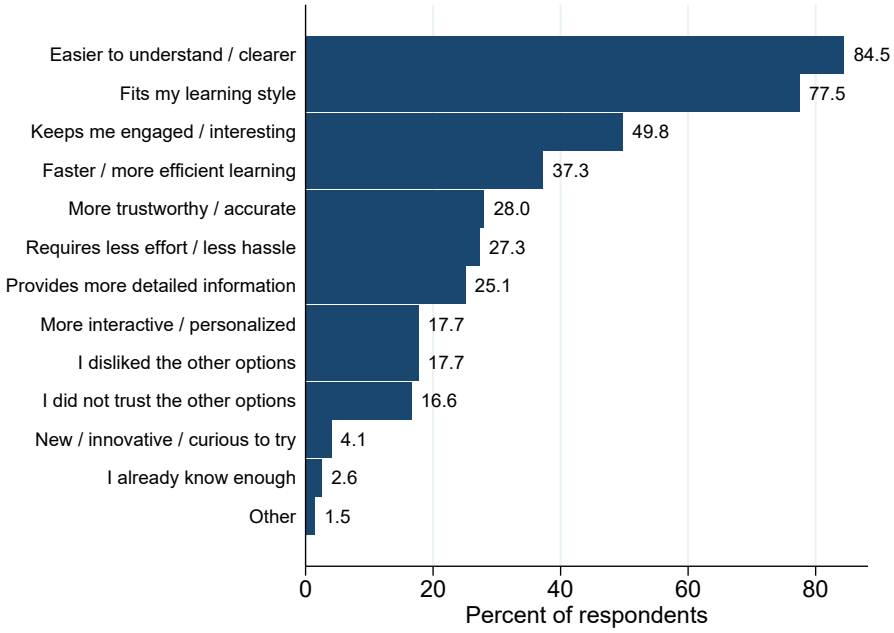


Figure A.5: Reasons for choosing the Video Education format

A.2 Additional Results

Table A1: Actual and Perceived Ability and Learning Gains

	(1) Estimated pre score	(2) Overestimate pre (est-actual)	(3) Estimated change	(4) Underestimate change (actual-est)
Low ability	-0.417**	2.393***	-0.289	0.289
	(0.176)	(0.176)	(0.249)	(0.249)
Change in number correct			0.180***	0.820***
			(0.056)	(0.056)
Low ability=1 × Change in number correct			-0.018	0.018
			(0.091)	(0.091)
Constant	5.585***	-2.130***	-0.592	0.592
	(0.352)	(0.355)	(0.361)	(0.361)
R^2	0.015	0.160	0.018	0.302
Observations	1047	1047	1047	1047

Notes: This table examines the relationship between actual and perceived ability and learning gains. Column (1) regresses participants' estimated baseline quiz score on whether they are low ability. Column (2) regresses the signed estimation error in baseline ability—defined as estimated pre score minus actual pre score—on actual baseline ability; positive values indicate overestimation. Column (3) regresses participants' estimated learning gains on their actual learning gains (change in number correct from pre to post). Column (4) regresses the signed estimation error in learning gains—defined as actual change minus estimated change—on actual learning gains; positive values indicate underestimation of learning gains. All regressions include controls for college education, gender, race, income, and age. Standard errors are reported in parentheses.

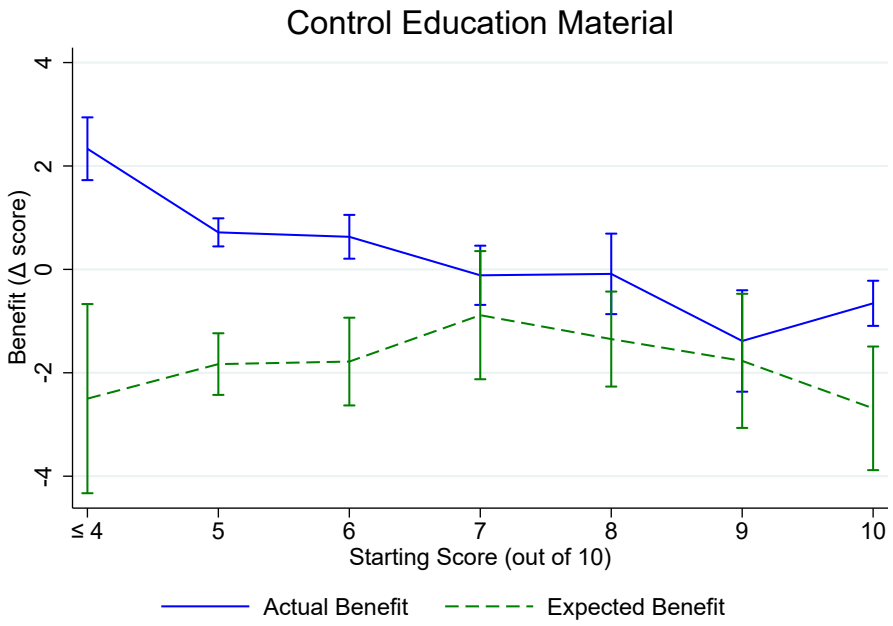


Figure A.6: Actual versus Expected Benefit Across Control Participants (Expected Benefit Control Only)

Notes: The actual benefit line is for participants who were in the control, and the expected benefit line is for only those in the control – i.e., the expected benefit of the control for participants who actually were assigned to the control.

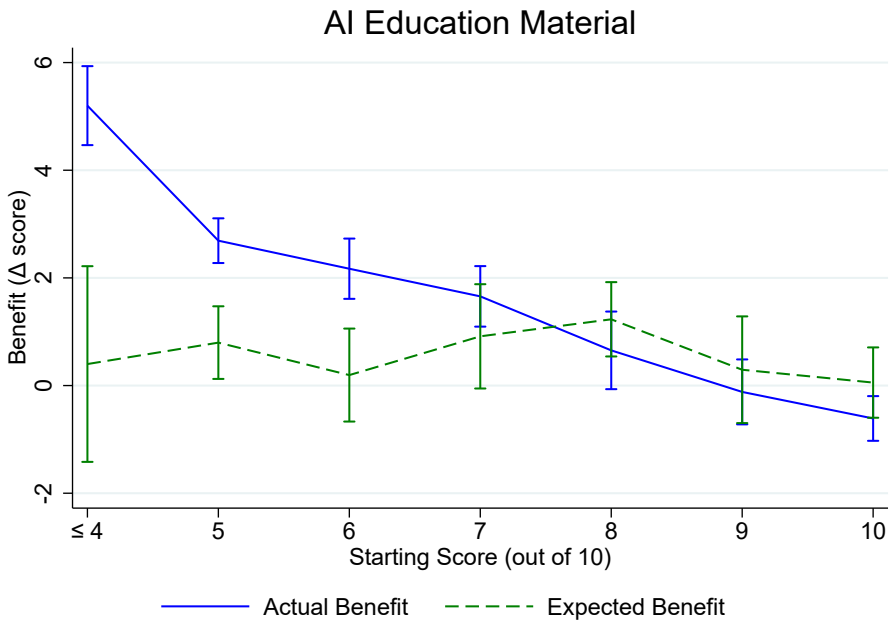


Figure A.7: Actual versus Expected Benefit Across AI Participants (Expected Benefit AI Only)

Notes: The actual benefit line is for participants who were in the AI treatment, and the expected benefit line is for only those in the AI treatment – i.e., the expected benefit of the AI treatment for participants who actually were assigned to the AI treatment.

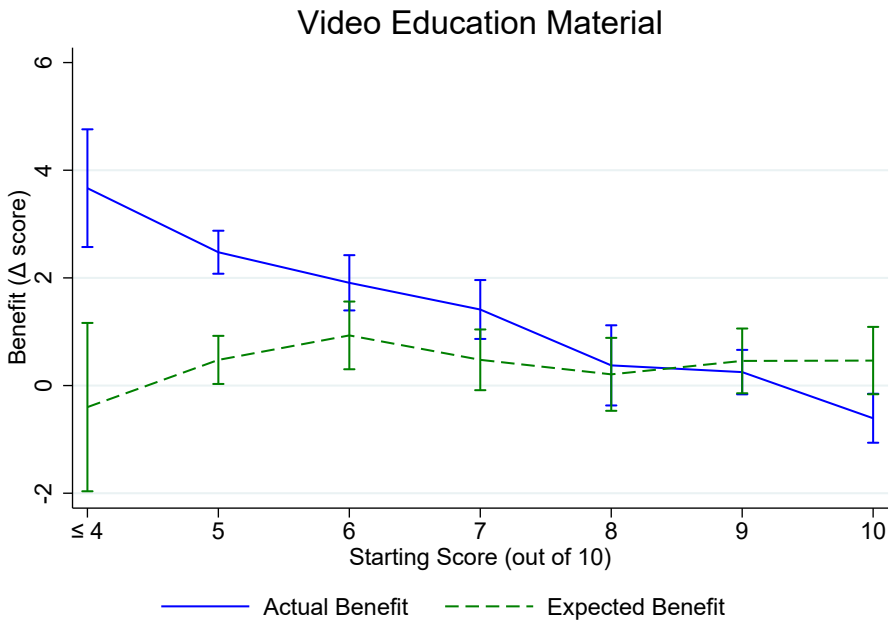


Figure A.8: Actual versus Expected Benefit Across Video Participants (Expected Benefit Video Only)

Notes: The actual benefit line is for participants who were in the Video treatment, and the expected benefit line is for only those in the Video treatment – i.e., the expected benefit of the Video treatment for participants who actually were assigned to the Video treatment.

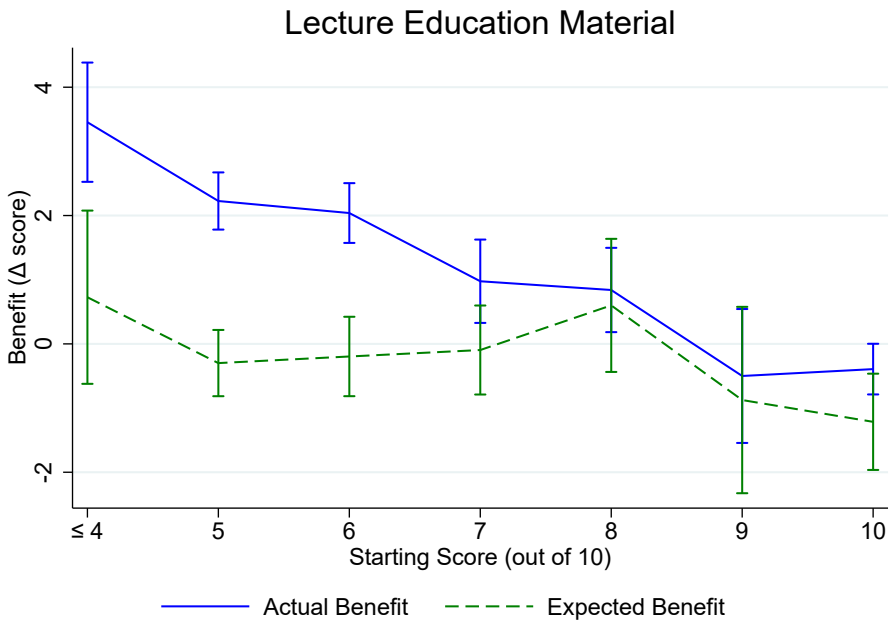


Figure A.9: Actual versus Expected Benefit Across Lecture Participants (Expected Benefit Lecture Only)

Notes: The actual benefit line is for participants who were in the Lecture treatment, and the expected benefit line is for only those in the Video treatment – i.e., the expected benefit of the Lecture treatment for participants who actually were assigned to the Lecture treatment.

Table A2: Mechanisms Behind Willingness to Pay for AI (OLS)

Panel A: Raw Specification

	(1)	(2)	(3)	(4)	(5)
Baseline number correct (pre)	0.115		0.132	0.268	0.268
	(0.083)		(0.095)	(0.203)	(0.203)
Change in number correct		-0.034	0.011	0.000	0.000
		(0.089)	(0.097)	(.)	(.)
Received this material		-0.942**	-0.837*		
		(0.447)	(0.448)		
Received material X Change in number correct		0.060	0.054		
		(0.171)	(0.170)		
Perceived difficulty (0 to 10)			-0.152***	-0.036	-0.036
			(0.057)	(0.136)	(0.136)
Estimated change			0.330***	0.529**	0.529**
			(0.070)	(0.217)	(0.217)
Change in perceived difficulty				-0.222*	-0.222*
				(0.126)	(0.126)
Actual minus perceived benefit				-0.184	-0.184
				(0.186)	(0.186)
Baseline number correct (pre)					0.000
					(.)
Constant	1.111*	2.118***	1.412*	-0.903	-0.903
	(0.580)	(0.210)	(0.765)	(1.815)	(1.815)
R^2	0.002	0.006	0.049	0.081	0.081
Observations	1014	1014	1014	248	248

Panel B: Control-Adjusted Specification

	(1)	(2)	(3)	(4)	(5)
Baseline number correct (pre)	0.115		0.132	0.081	0.081
	(0.083)		(0.095)	(0.163)	(0.163)
Change in number correct (minus control)		0.012	0.017	0.419*	0.419*
		(0.116)	(0.112)	(0.216)	(0.216)
Received this material		-0.818*	-0.658		
		(0.496)	(0.486)		
Received material X T-C improvement		0.078	0.094		
		(0.208)	(0.203)		
Perceived difficulty (0 to 10)			-0.186***	-0.032	-0.032
			(0.061)	(0.135)	(0.135)
Estimated change (relative to control)			0.366***	0.239**	0.239**
			(0.071)	(0.099)	(0.099)
Change in perceived difficulty				-0.222*	-0.222*
				(0.125)	(0.125)
Actual minus perceived benefit				0.226**	0.226**
				(0.113)	(0.113)
Baseline number correct (pre)					0.000
					(.)
Constant	1.111*	1.902***	0.720	-0.035	-0.035
	(0.580)	(0.271)	(0.786)	(1.470)	(1.470)
R^2	0.002	0.005	0.073	0.102	0.102
Observations	1014	767	767	248	248

Notes: This table reports regressions examining the mechanisms underlying willingness to pay (WTP) for the AI educational format. **Panel A** presents raw specifications using actual improvement and estimated change. **Panel B** uses control-adjusted learning measures: T–C improvement and estimated change relative to the control mean. Both panels sequentially add baseline accuracy, perceived difficulty, expected benefit, and changes in perceptions. Standard errors in parentheses.

Table A3: Mechanisms Behind Willingness to Pay for Video (OLS)

Panel A: Raw Specification

	(1)	(2)	(3)	(4)	(5)
Baseline number correct (pre)	0.026		0.093	0.309*	0.000
	(0.071)		(0.083)	(0.185)	(.)
Change in number correct		0.101	0.121	0.000	0.000
		(0.077)	(0.085)	(.)	(.)
Received this material		-0.268	-0.374		
		(0.403)	(0.404)		
Received material X Change in number correct		0.043	0.050		
		(0.166)	(0.167)		
Perceived difficulty (0 to 10)			-0.156***	-0.085	-0.085
			(0.056)	(0.162)	(0.162)
Estimated change			0.136**	0.397*	0.397*
			(0.065)	(0.225)	(0.225)
Change in perceived difficulty				-0.040	-0.040
				(0.113)	(0.113)
Actual minus perceived benefit				-0.288	-0.288
				(0.177)	(0.177)
Baseline number correct (pre)					0.309*
					(0.185)
Constant	1.292**	1.384***	1.327*	-0.897	-0.897
	(0.503)	(0.196)	(0.701)	(1.604)	(1.604)
R^2	0.000	0.003	0.019	0.022	0.022
Observations	1020	1020	1020	259	259

Panel B: Control-Adjusted Specification

	(1)	(2)	(3)	(4)	(5)
Baseline number correct (pre)	0.026		0.045	0.215	0.000
	(0.071)		(0.085)	(0.165)	(.)
Change in number correct (minus control)		0.181*	0.154	0.151	0.151
		(0.104)	(0.103)	(0.202)	(0.202)
Received this material		-0.029	-0.157		
		(0.439)	(0.428)		
Received material X T-C improvement		0.055	0.092		
		(0.194)	(0.189)		
Perceived difficulty (0 to 10)			-0.180***	-0.191	-0.191
			(0.063)	(0.156)	(0.156)
Estimated change (relative to control)			0.281***	0.479***	0.479***
			(0.066)	(0.128)	(0.128)
Change in perceived difficulty				0.017	0.017
				(0.111)	(0.111)
Actual minus perceived benefit				-0.139	-0.139
				(0.154)	(0.154)
Baseline number correct (pre)					0.215
					(0.165)
Constant	1.292**	1.052***	0.930	-0.740	-0.740
	(0.503)	(0.260)	(0.736)	(1.330)	(1.330)
R^2	0.000	0.007	0.046	0.079	0.079
Observations	1020	771	771	259	259

Notes: This table reports regressions examining the mechanisms underlying WTP for the video educational format. Panel A uses raw improvement and expected benefit; Panel B replaces these with control-adjusted counterparts. Specifications parallel those in Table A2. Standard errors in parentheses.

Table A4: Mechanisms Behind Willingness to Pay for Lecture (OLS)

Panel A: Raw Specification

	(1)	(2)	(3)	(4)	(5)
Baseline number correct (pre)	0.004		0.033	0.099	0.099
	(0.081)		(0.090)	(0.197)	(0.197)
Change in number correct		0.110	0.110	0.000	0.000
		(0.087)	(0.095)	(.)	(.)
Received this material		-0.576	-0.463		
		(0.429)	(0.421)		
Received material X Change in number correct		-0.221	-0.211		
		(0.173)	(0.166)		
Perceived difficulty (0 to 10)			-0.299***	-0.022	-0.022
			(0.062)	(0.162)	(0.162)
Estimated change			0.270***	0.209	0.209
			(0.062)	(0.220)	(0.220)
Change in perceived difficulty				-0.229*	-0.229*
				(0.120)	(0.120)
Actual minus perceived benefit				0.009	0.009
				(0.176)	(0.176)
Baseline number correct (pre)					0.000
					(.)
Constant	-0.092	0.030	1.512*	-0.705	-0.705
	(0.573)	(0.206)	(0.805)	(1.917)	(1.917)
R^2	0.000	0.008	0.060	0.046	0.046
Observations	1026	1026	1026	263	263

Panel B: Control-Adjusted Specification

	(1)	(2)	(3)	(4)	(5)
Baseline number correct (pre)	0.004		-0.066	0.037	0.037
	(0.081)		(0.095)	(0.167)	(0.167)
Change in number correct (minus control)		0.173	0.166	0.051	0.051
		(0.118)	(0.115)	(0.213)	(0.213)
Received this material		-0.293	-0.209		
		(0.469)	(0.453)		
Received material X T-C improvement		-0.357*	-0.335*		
		(0.203)	(0.195)		
Perceived difficulty (0 to 10)			-0.331***	-0.049	-0.049
			(0.070)	(0.159)	(0.159)
Estimated change (relative to control)			0.286***	0.227	0.227
			(0.080)	(0.147)	(0.147)
Change in perceived difficulty				-0.207*	-0.207*
				(0.119)	(0.119)
Actual minus perceived benefit				0.121	0.121
				(0.121)	(0.121)
Baseline number correct (pre)					0.000
					(.)
Constant	-0.092	-0.192	1.672*	-0.433	-0.433
	(0.573)	(0.271)	(0.863)	(1.621)	(1.621)
R^2	0.000	0.010	0.066	0.055	0.055
Observations	1026	773	773	263	263

Notes: This table reports regressions examining the mechanisms underlying WTP for the lecture educational format. Raw specifications appear in Panel A, while control-adjusted versions using T–C improvement and control-adjusted estimated change appear in Panel B. Standard errors in parentheses.

Table A5: Relationship Between AI Beliefs and WTP for AI

	WTP for AI								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
General AI Use (z)	0.802***				0.303*				
	(0.167)				(0.165)				
Uses AI for Finances (z)		0.897***			0.345**				
		(0.161)			(0.170)				
Believes AI Outperforms Me (z)			0.524***		0.297*				
			(0.174)		(0.163)				
Trusts AI Information (z)				1.344***	1.021***				
				(0.168)	(0.188)				
AI Usage Index (Use + Fin Use, z-mean)						0.881***			
						(0.205)			
AI Attitude Index (Superiority + Trust, z-mean)						1.250***			
						(0.224)			
AI Belief Index (All Four, z-mean)							2.116***	2.184**	2.039***
							(0.256)	(0.929)	(0.308)
Pre Score								0.124	
								(0.080)	
AI Belief Index X Pre Score								-0.010	
								(0.132)	
Score Change (post - pre)									-0.004
									(0.071)
AI Belief Index X Score Change									0.059
									(0.117)
Constant	1.894***	1.904***	1.856***	1.891***	1.903***	1.897***	1.904***	1.086*	1.912***
	(0.153)	(0.153)	(0.155)	(0.150)	(0.150)	(0.150)	(0.150)	(0.556)	(0.179)
R^2	0.026	0.032	0.011	0.073	0.083	0.076	0.075	0.077	0.075
Observations	1014	1014	1014	1014	1014	1014	1014	1014	1014

Notes: This table reports regressions of willingness to pay for AI on specific belief components (AI usage, financial AI use, perceived AI superiority, trust in AI), on composite belief indices, and on interactions between beliefs and either baseline performance or actual improvement. Standard errors in parentheses.

Table A6: AI Belief Index and Measures of Ability

	(1) AI Belief Index (All Four, z-mean)	(2) AI Belief Index (All Four, z-mean)
Pre Score	-0.006	
	(0.012)	
Change in number correct		-0.015
		(0.009)
Constant	0.030	0.009
	(0.078)	(0.023)
R^2	0.000	0.003
Observations	1047	1047

Standard errors in parentheses

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

Notes: This table reports regressions of the AI Belief Index on participants' baseline quiz performance (Column 1) and actual improvement (Column 2). Standard errors in parentheses.

Table A7: Demographic Correlates of AI Beliefs

	(1) AI Belief Index (All Four, z-mean)	(2) AI Belief Index (All Four, z-mean)
College Educated=1	0.132***	
	(0.044)	
Income (Bin Midpoint)		0.000**
		(0.000)
Constant	-0.101***	-0.068*
	(0.038)	(0.035)
R^2	0.009	0.005
Observations	1047	1047

Standard errors in parentheses

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

Notes: Table A7 reports regressions of the *AI Belief Index*—a standardized index combining participants' AI usage, attitudes, trust, and perceived comparative advantage—on demographic variables. Column (1) reports a specification with a college indicator; Column (2) uses income (midpoint of self-reported income bin). Higher values of the AI Belief Index indicate more positive attitudes toward using AI tools. Standard errors are reported in parentheses.

Table A8: Prediction Error versus Financial Behaviors and Attitudes

	Error (Predicted - Actual Change)			
	(1)	(2)	(3)	(4)
Talk Finance = Rarely	0.128			
	(0.334)			
Talk Finance = Sometimes	0.176			
	(0.413)			
Talk Finance = Often	0.323			
	(0.341)			
Talk Finance = Always	2.260**			
	(0.980)			
Taken Class = Yes		0.477*		
		(0.246)		
Owns Stocks=1			0.253	
			(0.230)	
Decision Diff = Doable				0.318
				(0.303)
Decision Diff = Hard				0.766*
				(0.426)
Constant	-1.491***	-1.436***	-1.430***	-1.589***
	(0.278)	(0.135)	(0.177)	(0.268)
R^2	0.008	0.005	0.002	0.004
Observations	791	791	791	729

Standard errors in parentheses

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

Notes: Prediction error is defined as perceived benefit minus actual benefit. This table regresses prediction error on talking about finance, prior course experience, stock ownership, and perceived decision difficulty. The base category for Talk Finance is Never, the base category for Taken Class and Owns Stocks is No, and the base category for Decision Diff is Easy. Standard errors in parentheses.

Table A9: Relationship Between Baseline Ability and Education Beliefs

	(1) Pre Score	(2) Pre Score	(3) Pre Score
Decision Diff = Doable	-0.339**		
	(0.159)		
Decision Diff = Hard	-0.716***		
	(0.228)		
Edu enjoyable = Somewhat unenjoyable		0.002	
		(0.294)	
Edu enjoyable = Neither enjoyable nor unenjoyable		0.040	
		(0.198)	
Edu enjoyable = Somewhat enjoyable		0.125	
		(0.160)	
Edu enjoyable = Very unenjoyable		-0.879***	
		(0.275)	
Edu learning = Learned very little			-0.003
			(0.797)
Edu learning = Learned a moderate amount			-0.300
			(0.770)
Edu learning = Learned quite a bit			-0.195
			(0.771)
Edu learning = Learned a great deal			-0.244
			(0.773)
Constant	6.864***	6.545***	6.800***
	(0.144)	(0.123)	(0.762)
R^2	0.011	0.011	0.002
Observations	972	729	729

Standard errors in parentheses

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

Notes: This table regresses baseline quiz performance on perceived decision difficulty, enjoyment, and perceived learning potential. The base category for Decision Diff is Easy, the base category for Edu enjoyable is very enjoyable, and the base category for Edu learning is learning nothing. Standard errors in parentheses.

Table A10: Treatment Effect versus Perceived Effort and Engagement

	Treatment Effect			
	(1)	(2)	(3)	(4)
Edu effort = Low	1.500**			
	(0.632)			
Edu effort = Moderate	0.871**			
	(0.403)			
Edu effort = High	1.054***			
	(0.378)			
Edu effort = Very high	1.093***			
	(0.376)			
Edu learning = Learned very little		0.131		
		(0.576)		
Edu learning = Learned a moderate amount		0.652		
		(0.549)		
Edu learning = Learned quite a bit		1.046*		
		(0.550)		
Edu learning = Learned a great deal		1.158**		
		(0.560)		
Edu enjoyable = Somewhat unenjoyable			-0.950***	
			(0.308)	
Edu enjoyable = Neither enjoyable nor unenjoyable			-0.386	
			(0.237)	
Edu enjoyable = Somewhat enjoyable			-0.418**	
			(0.201)	
Edu enjoyable = Very unenjoyable			-0.486	
			(0.418)	
Talk Finance = Rarely				-0.498**
				(0.211)
Talk Finance = Sometimes				-1.271***
				(0.259)
Talk Finance = Often				-0.929***
				(0.206)
Talk Finance = Always				-1.437***
				(0.538)
Constant	0.500	0.700	1.893***	2.206***
	(0.355)	(0.532)	(0.157)	(0.159)
R^2	0.002	0.021	0.013	0.039
Observations	729	729	729	791

Standard errors in parentheses

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

Notes: This table relates treatment effect (post–pre change among treated participants) to perceived effort, perceived learning, enjoyment, and financial discussion frequency. The base category for effort is Very low effort, the base category for Decision Diff is Easy, the base category for Edu enjoyable is very enjoyable, and the base category for Edu learning is learning nothing. Standard errors in parentheses.

Table A11: Determinants of WTP for Educational Materials (Tobit)

	(1) WTP for AI	(2) WTP for Video	(3) WTP for Lecture	(4) Average WTP
Baseline number correct (pre)	0.165	0.094	0.025	0.071
	(0.108)	(0.093)	(0.102)	(0.081)
Change in number correct	0.050	0.146*	0.071	0.075
	(0.095)	(0.087)	(0.094)	(0.074)
Perceived difficulty (0 to 10)	-0.111*	-0.163***	-0.340***	-0.100
	(0.064)	(0.062)	(0.072)	(0.070)
Estimated change	0.303***	0.142*	0.319***	0.160**
	(0.081)	(0.073)	(0.072)	(0.066)
AI Belief Index (z)	2.281***	0.358	0.443*	0.955***
	(0.298)	(0.221)	(0.245)	(0.197)
Constant	0.168	0.719	3.075***	1.004
	(1.023)	(0.925)	(1.131)	(0.838)
Pseudo-R^2	0.020	0.004	0.011	0.007
Observations	1014	1020	1026	966

Notes: Tobit regressions censor WTP at $-11 and $11. Columns correspond to WTP for AI, video, lecture, and average WTP. Coefficients represent latent WTP. Standard errors in parentheses.

Table A12: Determinants of WTP Using Control-Adjusted Learning Measures (Tobit)

	(1) WTP for AI	(2) WTP for Video	(3) WTP for Lecture	(4) Average WTP
Baseline number correct (pre)	0.145	0.048	-0.097	0.017
	(0.106)	(0.094)	(0.107)	(0.082)
T-C improvement	0.082	0.203**	0.083	0.118
	(0.105)	(0.096)	(0.106)	(0.082)
Perceived difficulty (0 to 10)	-0.150**	-0.185***	-0.386***	-0.156**
	(0.069)	(0.068)	(0.083)	(0.077)
Estimated change (relative to control)	0.356***	0.304***	0.335***	0.283***
	(0.083)	(0.076)	(0.094)	(0.077)
AI Belief Index (z)	2.003***	0.167	0.170	0.782***
	(0.335)	(0.255)	(0.271)	(0.219)
Constant	-0.620	0.112	3.340***	0.773
	(1.037)	(0.947)	(1.212)	(0.859)
Pseudo-R^2	0.023	0.009	0.012	0.010
Observations	767	771	773	729

Notes: Tobit regressions parallel the OLS control-adjusted models, substituting T–C improvement and control-adjusted expected benefit. Columns correspond to AI, video, lecture, and average WTP. Standard errors in parentheses.

Table A13: All Characteristics - Willingness to Pay for Educational Materials (OLS)

	(1) WTP for AI	(2) WTP for Video	(3) WTP for Lecture	(4) Average WTP
Estimated change	0.232***	0.141*	0.289***	0.141*
	(0.086)	(0.082)	(0.077)	(0.078)
Perceived difficulty (0 to 10)	-0.152**	-0.127*	-0.300***	-0.158*
	(0.068)	(0.073)	(0.087)	(0.087)
Perceived time (0 to 10)	0.128	-0.040	-0.036	0.139
	(0.106)	(0.125)	(0.115)	(0.116)
Edu effort = Low	1.202	-2.547	-1.092	-0.338
	(2.334)	(6.391)	(6.941)	(3.959)
Edu effort = Moderate	3.134*	-1.422	-0.506	0.451
	(1.660)	(6.278)	(6.834)	(3.775)
Edu effort = High	2.879*	-1.637	-0.988	0.146
	(1.635)	(6.273)	(6.816)	(3.765)
Edu effort = Very high	3.365**	-0.953	-0.478	0.791
	(1.673)	(6.277)	(6.826)	(3.777)
Edu enjoyable = Somewhat unenjoyable	-1.519**	-2.187***	-1.355*	-1.786***
	(0.741)	(0.692)	(0.704)	(0.598)
Edu enjoyable = Neither enjoyable nor unenjoyable	-1.252**	-1.520***	-0.293	-1.094**
	(0.586)	(0.524)	(0.581)	(0.501)
Edu enjoyable = Somewhat enjoyable	-0.449	-0.335	0.054	-0.358
	(0.485)	(0.462)	(0.501)	(0.426)
Edu enjoyable = Very unenjoyable	-2.256*	-1.374	-0.345	-1.688
	(1.210)	(1.094)	(1.148)	(1.067)
Previously took finance course	0.411	0.033	-0.107	0.186
	(0.429)	(0.382)	(0.425)	(0.337)
Talk Finance = Rarely	0.091	0.010	-0.187	-0.081
	(0.486)	(0.481)	(0.518)	(0.436)
Talk Finance = Sometimes	0.012	0.333	0.288	-0.134
	(0.633)	(0.566)	(0.672)	(0.547)
Talk Finance = Often	-0.221	-0.218	-0.302	-0.410
	(0.491)	(0.460)	(0.540)	(0.434)
Talk Finance = Always	1.856	0.270	-0.499	1.546
	(1.652)	(1.784)	(1.715)	(1.412)
AI Belief Index (z)	1.533***	-0.136	0.047	0.519**
	(0.308)	(0.259)	(0.270)	(0.245)
College Educated=1	-0.464	-0.302	-0.561	-0.372
	(0.416)	(0.380)	(0.411)	(0.355)
Male=1	-0.923***	-0.634*	-0.548	-0.532*
	(0.356)	(0.335)	(0.378)	(0.310)
White=1	0.717*	0.152	0.023	0.400
	(0.371)	(0.350)	(0.399)	(0.325)
Income (Bin Midpoint)	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Constant	-1.106	4.117	3.230	0.947
	(1.735)	(6.352)	(6.890)	(3.835)
R^2	0.137	0.057	0.068	0.074
Observations	706	709	711	668

Notes: This table reports OLS regressions of willingness to pay (WTP) for educational materials on perceived benefit, perceived cognitive costs, engagement, enjoyment, prior exposure to financial education, background characteristics, and beliefs about AI. The base category for Edu effort is Very low effort. The base category for Edu enjoyable is Very enjoyable. The base category for Talk Finance is Never. The base category for College Educated, Male, and White indicators is 0. Perceived benefit, perceived difficulty, and perceived time are measured separately for each educational format in Columns (1)–(3); in Column (4), these variables are averaged across formats. Standard errors are reported in parentheses.