

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

The causal mechanism of financial education: Evidence from mediation analysis [☆]

Fenella Carpena ^{a,*}, Bilal Zia ^b^aOslo Business School, Oslo Metropolitan University^bThe World Bank

ARTICLE INFO

Article history:

Received 4 March 2019

Revised 15 April 2020

Accepted 1 May 2020

JEL Codes:

C93

D14

G21

O12

Keywords:

Causal Mediation Analysis

Mechanism of Impact

Financial Education

Financial Literacy

Financial Knowledge

Impact Evaluation

Randomized Control Trial

ABSTRACT

This paper uses a field experiment in India and mediation analysis to investigate the causal mechanisms between financial education and financial behavior. Focusing on the mediating role of financial literacy, we propose a broader definition of financial knowledge that includes three dimensions: numeracy skills, financial awareness, and attitudes towards personal finance. We then employ causal mediation analysis to investigate the proportion of the treatment effect that can be attributed to these three channels. Strikingly, we find that numeracy does not mediate any effects of financial education on financial outcomes. For simple financial actions such as budgeting, both awareness and attitudes serve as pathways, while for more complex financial activities such as opening a savings account, attitudes play a more prominent role—though these patterns appear to be sensitive to confounding. We also compare our mediation analysis results to other empirical techniques that have been typically used to study mechanisms, and we discuss how mediation analysis differs from these approaches.

© 2020 Published by Elsevier B.V.

1. Introduction

Over the last decade, financial education programs have become an increasingly popular tool for promoting financial inclusion, consumer welfare, and stable financial systems (e.g., [Lusardi and Mitchell, 2013](#); [World Bank, 2017](#)). Nevertheless, the merits of such programs remain a hotly contested policy issue. On the one hand, critics maintain that financial education is a fallacy because it is neither economical nor effective, arguing instead for other forms of financial regulation such as retirement savings defaults and pro bono financial advisory services (e.g., [Willis, 2011](#)). On the other hand, many govern-

[☆] We are very grateful to Saath Microfinance for their constant support. We thank Stuti Tripathi, Bhakti Shah, and the Center for Microfinance at the Institute for Financial Management and Research for excellent field work and research assistance. This study is part of a broader research project in India that was undertaken in collaboration with Shawn Cole and Jeremy Shapiro, to whom we are grateful. We are also grateful to Simon Galle, Xavier Giné, Arie Kapteyn, Annamaria Lusardi, and David McKenzie, as well as participants at the Oslo Business School Research Group in Economics and Finance Seminar and the 2nd CEAR-RSI Household Finance Workshop for many helpful comments and suggestions. Earlier versions of this paper benefitted from coverage in the World Bank's *Development Impact* and *All About Finance* blogs. Research funding from the World Bank Research Department is gratefully acknowledged. The study is registered in the AEA RCT Registry, ID # AEARCTR-0000173.

* Corresponding author.

E-mail addresses: fenella.carpena@oslomet.no (F. Carpena), bzia@worldbank.org (B. Zia).

ments and organizations worldwide continue to champion financial education, as evidenced in the membership of hundreds of countries and public institutions in the *International Gateway for Financial Education*.¹

At the center of this policy discourse lies a growing—yet incomplete—body of evidence on the causal effect of financial education. Although most early studies demonstrate only modest effects (e.g., [Fernandes et al., 2014](#); [Hastings et al., 2013](#); [Miller et al., 2015](#)), recent research indicates a more promising role for financial education. For example, financial training delivered through popular media ([Berg and Zia, 2017](#)), complemented with goal setting and counseling ([Carpena et al., 2017](#)), or targeted for specific groups ([Doi et al., 2014](#); [Bruhn et al., 2016](#)) all result in significant improvements in household financial outcomes. At the same time, much of this literature thus far focuses on establishing causality, placing little weight on the pathways of the effects. Hence, the mechanisms through which financial education operates are not well-understood.

This paper carefully unpacks the causal mechanisms of financial education. We examine the gains in financial literacy as an intermediate channel by investigating two questions. First, fundamentally, how should financial literacy be measured, especially in developing countries? Existing studies typically use the “Big Three” questions covering interest rates, inflation, and diversification from [Lusardi and Mitchell \(2009\)](#), but they may not be suitable in emerging markets where consumers have poor education and lack the most basic skills. And second, how do different aspects of financial literacy (e.g., numeric calculations, familiarity with bank accounts, perceptions of financial products) mediate the impact of financial education on financial outcomes?

We examine the above questions using a field experiment among the urban poor in India, which we previously studied in [Carpena et al. \(2017\)](#). As part of the experiment, a randomly selected two-thirds were invited to a video-based financial education program. A subset of those assigned to financial education were likewise provided with one of three add-on treatments: (1) concrete financial goal setting; (2) individualized financial counseling; or (3) both goal setting and counseling. In [Carpena et al. \(2017\)](#), we focus on the Average Treatment Effect (ATE) and find that financial education alone did not bring about large changes in financial behavior. But financial education with goal setting encouraged relatively *simple* follow-up activities (e.g., writing a budget), while financial education with counseling fostered more *complex* financial actions (e.g., opening a savings account).

In contrast to [Carpena et al. \(2017\)](#), this paper moves beyond the ATE and explores the underlying channels of financial education, goal setting, and counseling. We employ causal mediation analysis (e.g., [Baron and Kenny, 1986](#); [Imai et al., 2011](#); [VanderWeele, 2015](#); [Acharya et al., 2016](#)), an approach that quantifies the extent to which the treatments influence outcomes through a specific mediating variable. To this end, our empirical method consists of two components. The first concerns measuring financial literacy, our proposed mediator. Compared to the specific “Big Three” questions by [Lusardi and Mitchell \(2009\)](#), we take a more comprehensive approach, conceptualizing financial literacy into broader categories: financial numeracy, financial awareness, and financial attitudes.

Through a series of endline survey questions, we designed each of these dimensions to capture a different aspect of financial literacy. Financial *numeracy* deals with calculating interest rates, summing expenses, and similar computations. These skills may facilitate more effective fiscal management and comparisons of financial products. Next, financial *awareness* emphasizes fundamental financial concepts (e.g., household budgeting) and basic information about financial products (e.g., deposit insurance, loan fees). This type of knowledge may promote wider adoption of financial instruments and simple financial habits. Finally, financial *attitudes* encompass individual perspectives on the benefits of financial services. We view attitudes as an essential element of financial literacy because they may have critical consequences for financial behavior: for example, if one sees no advantage of savings, one may be less inclined to set money aside for the future.

To understand how numeracy, awareness, and attitudes function as mediators, the second component of our study implements causal mediation analysis. We follow the framework outlined in [Imai et al. \(2011\)](#) to decompose the ATE into an Average Causal Mediation Effect (ACME), representing a particular mechanism, and an Average Direct Effect (ADE), representing all other pathways. We estimate the ACME empirically using coefficients from two regressions: one for the effect of the treatment on the mediator and another for the effect of the mediator on the outcome conditional on the treatment. Intuitively, the product of these two coefficients—the ACME—captures the portion of the ATE that can be attributed to the mediator. We then test the sensitivity of the ACME to the key identifying assumption that the mediator is statistically independent of the outcome, given treatment and baseline characteristics. We also show that our results are robust to using an alternative estimation strategy—specifically, sequential *g*-estimation as proposed in [Acharya et al. \(2016\)](#)—which requires weaker identification assumptions.

Our analysis reveals four critical insights on mechanisms. First, while conventional intuition suggests that quantitative skills may be an important ingredient for financial capability, we find that financial numeracy does not play a significant intermediary role in any of our study’s financial education interventions. We obtain an ACME of precisely nil for numeracy across all treatments.

Second, we demonstrate that both financial awareness and attitudes appear to be meaningful channels from financial education to household budgeting. For all treatment combinations of financial education, goal setting, and counseling that we study, we detect an increase in awareness and attitudes scores of up to 15 percentage points. Correspondingly, our results indicate that across all treatments, up to 20 and 21 percent of the ATE for budgeting operates through awareness

¹ See <http://www.financial-education.org/about.html>, accessed June 5, 2018.

and attitudes, respectively. Moreover, these estimates are statistically indistinguishable, suggesting that both channels are equally important for changes in household budgeting behavior.

Third, while awareness and attitudes play analogous roles as channels for budgeting, we uncover key differences in how they mediate treatment impacts on household savings. Specifically, our analysis shows that for all treatments, financial attitudes but not awareness likely fosters adoption of formal savings accounts. Indeed, the ACME of awareness is very close to zero regardless of whether the treatment consists of financial education alone or augmented with goal setting and/or counseling, though we cannot rule out that these null effects are due to low statistical power. On the other hand, the ACME of financial attitudes across all types of treatments is between 2 and 3 percentage points, amounting to as much as 31 percent of the ATE for the treatment where all three of financial education, goal setting, and counseling are offered. These findings provide suggestive evidence that changing perceptions about financial products is an important mechanism for financial education.

And fourth, we show that although the most intensive treatments (i.e., financial education with counseling or financial education with both counseling and goal setting) have significant positive effects on borrowing and insurance outcomes, none of the three financial literacy dimensions appear to act as mediators. Our estimates of the ACME of awareness and attitudes on these behaviors are all close to zero and not statistically significant. Furthermore, the treatment effects of financial education on borrowing and insurance are much smaller in absolute terms relative to budgeting and savings outcomes. Overall, these patterns indicate that changing borrowing and insurance behavior through financial education is a difficult task, and the impacts of the treatments on these outcomes are likely mediated through channels other than improved financial knowledge.²

Our study contributes to both research and policy on financial education. While many studies have considered the causal effects of financial education on financial behavior, few have examined the *mechanisms* through which such education is effective. Addressing this gap in the literature, this paper is, to our knowledge, the first to implement causal mediation analysis to explore the channels of financial education. From a methodological perspective, our study illustrates how field experiments can be structured to identify causal mechanisms. Indeed, if mediators are incorporated in the study design, one can apply causal mediation analysis as we have done here to investigate mediation effects. Importantly, we compare our mediation analysis results with other techniques that are typically used to study mechanisms (e.g., instrumental variables), and we discuss in detail how mediation analysis differs from these approaches. In so doing, this paper adds to the methodological discourse by illustrating how mediation analysis can be carried out in economics research.

2. Experiment design and summary statistics

Our study takes place in Ahmedabad, a large metropolis in Gujarat, India. We exploit a field experiment with three types of interventions: financial education, concrete financial goal setting, and individualized financial counseling, all of which were randomly assigned at the respondent level. In what follows, we provide a brief summary of the treatments, and we refer the reader to [Carpena et al. \(2017\)](#) for more details on the experiment design.

2.1. Financial education

Our main intervention is a video-based financial education program. This program was offered to a randomly selected two-thirds of participants, while the other one-third formed the control group. The control group was offered health education (instead of no training at all) to account for Hawthorne effects, which may be especially important in our setting due to the program's intensity: whereas many studies have examined only short, one-off financial trainings,³ the financial program in our experiment consisted of five weekly meetings, each lasting 2 to 3 hours.

The financial education program screened videos on five topics: budgeting, savings, loans, insurance, and financial management. For consistency, the health program also used videos and discussed five topics unrelated to financial knowledge: cleanliness and hygiene; midwifery; maternal and child health; condoms, AIDS, and syphilis; and night-blindness.⁴ To further ensure comparability across subjects, both financial and health training sessions were implemented with the same logistics. For example, both were carried out in a classroom environment, where each class met at the same time every week and consisted of about 20 respondents assigned to the same type of program.⁵ All participants received a show-up

² An important caveat to all our results is that the ACME estimates appear to be sensitive to violations of the key identifying assumption that the mediator is as-if random given treatment and baseline controls. See [Section 5.4](#).

³ For instance, Miller et al. (2014) find that more than one-third of financial education programs are delivered within one week or less. Since our study involves an in-depth curriculum, providing the control group with health education was necessary to maintain the same level of everyday disruption across all respondents throughout the study.

⁴ The financial education videos were based on standard materials previously used in the literature, but they were adapted to our context by a local media company together with input from the research team and our local research and implementation partners. The health education videos were produced by the United Nations in India.

⁵ The study was carried out over several waves. Each wave comprised of about 15 classes. Of these, 10 were for financial education, and the remaining 5 were for health training. All financial and health classes were facilitated by a trained instructor, who answered outstanding questions as well as promoted discussion on the course topics.

Table 1
Experimental Design.

Financial Education Videos	Counseling	Goal Setting	N	% of Sample
No	No	No	316	33
Yes	No	No	171	18
Yes	No	Yes	149	16
Yes	Yes	No	152	16
Yes	Yes	Yes	160	17

Notes: This table describes the randomization across the various treatments. The total sample consists of 948 respondents to whom knowledge questions on financial literacy (i.e., financial numeracy, financial awareness, and financial attitudes) were administered at endline and who have non-missing baseline controls.

fee of Rs. 50 (approximately US\$ 1) for every session they attended as well as free transportation to and from the training center.

2.2. Concrete financial goal setting

Our second intervention, concrete financial goal setting, encouraged participants to set short-term achievable but non-compulsory financial goals. This treatment involved a household visit where: (1) respondents were interviewed about their use of financial services; (2) respondents were asked to voluntarily choose a target date for completing one or more financial goals (i.e., opening a savings account, increasing savings, reducing expenditure, and/or purchasing insurance); and (3) enumerators listed the respondents' target dates on a calendar provided at no cost by the study, and subsequently posted the calendar in the respondent's home.

We administered the goal setting treatment by design to a randomly selected half of all respondents assigned to financial education. In so doing, we are able to estimate the marginal effect of goal setting beyond financial education alone. The other half of participants then served to separate the effects of goal setting versus a household visit. In particular, this group received the same household interview about financial services over the same fieldwork period, but they were not asked to set financial goals nor given calendars. Thus, the impact of goal setting that we measure represents the combined effect of both the target dates and free calendars provided to respondents.

2.3. Financial counseling

Our third treatment, financial counseling, consisted of one-on-one instruction and individualized advice, carried out for free at the participant's home. Skilled financial counselors—trained rigorously by the Center for Microfinance, our Indian research partner—guided participants on money management depending on their specific needs. For example, counselors assisted respondents in preparing a budget, gathering documents for opening a bank account, or contacting an insurance provider. This counseling treatment was offered to a randomly selected half of those in the financial education group, orthogonally to goal setting. Shortly after goal setting fieldwork was completed, those assigned to counseling received one household visit per month from the counselor for four months, with more frequent meetings available at the respondent's request.

2.4. Summary statistics

The data in this paper come from a baseline survey before the start of the interventions and an endline survey almost ten months after the final session of the video-based education program. To accommodate the large sample, the sample was split into separate waves. Respondents were drawn from various neighborhoods (*chalis*) that were mutually exclusive across waves, and all treatments were stratified based on the respondent's gender, neighborhood, and status as a microfinance client. Attrition between baseline and endline is only 6% and is uncorrelated with treatment.

Table 1 outlines our study sample ($N = 948$) as well as the proportion of respondents assigned to each treatment combination.⁶ As can be seen in the table, 18% were allocated to receive only financial education; 16% to financial education and goal setting (but not counseling); 16% to financial education and counseling (but not goal setting); and 17% to all three treatments. The remaining 33% serve as control and were assigned to watch health education videos. Across both treatment and control, take-up of the education program was quite high; both financial and health training sessions received nearly 100% attendance for the duration of the five-week program.

⁶ The sample in this paper contains fewer observations than Carpena et al. (2017) because we focus on respondents with endline measures of numeracy, awareness, and attitudes. These questions were added only in Waves 2 to 4. Hence, the sample in this paper consists of subjects from Waves 2 to 4, while Carpena et al. (2017) uses all four waves.

Table 2
Baseline Summary Statistics.

	Median	Mean	Standard Deviation	Coefficient of Variation	Correlation with Baseline Financial Knowledge Score	Test of Joint Equality of Means Across All Treatments (F-test p-value)
	(1)	(2)	(3)	(4)	(5)	(6)
Household size	5.00	5.71	2.42	0.42	-0.01	0.27
Household monthly income (Rs.)	5500.00	6892.01	5878.08	0.85	0.15	0.39
Household monthly income per capita (Rs.)	1026.79	1280.61	980.71	0.77	0.14	0.41
Household has phone		0.84			0.10	0.89
Household has water connection		0.76			0.04	0.76
Household has non-farm enterprise		0.25			0.10	0.96
Respondent is Female		0.60			-0.05	
Respondent is Hindu		0.79			0.00	0.92
Respondent has completed secondary school		0.04			0.13	0.59
Respondent is microfinance client		0.49			0.03	
Respondent has hard time saving		0.95			-0.03	0.72
Respondent is interested in financial matters		0.87			0.09	0.95
Respondent has inconsistent time preferences		0.45			0.02	0.08*
Respondent monthly discount rate	0.14	1.40	4.45	3.18	0.03	0.67
Respondent is risk averse		0.16			-0.02	0.82
Respondent math score (out of 8)	5.00	4.73	2.05	0.43	0.27	0.56
Respondent financial knowledge score (out of 3)	2.00	1.58	0.62	0.39		0.18

Notes: This table provides baseline summary statistics for our sample. Column (4) reports the coefficient of variation, defined as the ratio of the standard deviation to the mean. Column (5) reports the correlation between the given variable and the baseline financial knowledge score (i.e., the last variable in the table). Column (6) is the p-value of the F-test of joint significance of all treatment coefficients in regressions of the baseline characteristics on treatment dummies. The four treatments consist of: (1) financial education only; (2) financial education and goal setting; (3) financial education and counseling; and (4) financial education, goal setting, and counseling. The regression specification that was used for the F-test in the last column also controls for strata dummies, where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client. Standard errors are clustered at the wave-class level. *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Next, [Table 2](#) presents baseline summary statistics for our sample. Households have about six members on average, with a mean monthly income of Rs. 6892 (US\$ 118). Sixty percent of the respondents were female, and very few (i.e., 4%) completed secondary school. About half were members of a microfinance organization, yet, almost everyone in our sample (i.e., 95%) reported having difficulty saving. Additionally, we evaluated computational skills using eight simple math questions as well as financial knowledge skills using questions similar to the “Big Three” by [Lusardi and Mitchell \(2009\)](#). The average score was just over 50% on both of these measures.

To better understand the preferences of our study subjects, we also measured discount rates and risk aversion. Discount rates were assessed in a standard manner, by asking respondents to provide the minimum amount they would be willing to hypothetically accept in one month in lieu of a hypothetical payment of Rs. 350 today. Respondents in our sample reported relatively high monthly discount rates: the median was 0.14, while the average was 1.40. We then measured risk aversion by allowing respondents to choose between a payment of Rs. 10 with certainty or playing a lottery that pays out Rs. 25 or Rs. 0 with equal probability. Sixteen percent of our sample chose the safe payment, and these respondents were coded as risk averse.

Finally, the p-values in [Table 2](#), Column 6 report the statistical significance of a joint test for the difference between the means across all treatments including the control group. The p-values are fairly large, suggesting no significant difference across treatment in baseline measures. Only one baseline characteristic—namely, the indicator for having inconsistent time preferences—exhibits imbalance across treatments. We control for this variable in all regression specifications.

3. Measuring financial knowledge

Because our goal is to study financial knowledge as a mediator, a critical ingredient is measurement. Financial knowledge is typically assessed through survey questions designed to capture respondents' basic financial understanding and capability in applying financial concepts to financial decisions. Currently, the standard approach is based on three questions developed

by Lusardi and Mitchell (2009), which cover interest rates, inflation, and diversification.⁷ These “Big Three” questions have been implemented in many developed countries (e.g., Italy, Germany, Netherlands, USA) and have been shown to be strong predictors of financial outcomes. They have likewise been used in many emerging markets, such as Indonesia, India, Sri Lanka, and Mexico.

Despite the ubiquity of the “Big Three,” they are not necessarily comprehensive nor appropriate in many settings. Indeed, Lusardi and Mitchell (2009) themselves write that it is “imperative to expand the range of measures of financial literacy, so as to better evaluate the types of problems that people find most difficult” (p. 6). Developing countries in particular possess many features—among others, high poverty, low access to finance, and lack of consumer finance protections—that are necessary to consider when measuring financial knowledge (Holzmann, 2010). For instance, if most households are uneducated and hold informal savings, it may be more important to evaluate understanding of bank account opening requirements than savings returns computations. Thus, the notion of financial literacy should be extended, especially in the developing country context.

To this end, we propose a broader approach to assessing financial literacy. We conceptualize financial knowledge not as specific topics as the “Big Three” does, but rather as three general and distinct dimensions.⁸ The first dimension is financial *numeracy*, which concerns the ability to add income and expenses, determine interest rates, and similar calculations. In particular, as part of our endline survey, we asked respondents the following two questions.

- (1) Let’s assume that you deposited Rs. 10,000 in a bank account at an 8% monthly interest rate. How much money will you have in your account in one year if you do not withdraw from or add to this account any money? (a) More than 10,800; (b) Less than 10,800; (c) Exactly 10,800.
- (2) Suppose you had Rs. 50 to save. You could either save this for 1 month in an account which earns 14% interest per month or save it for 1 month in an account that earns 2% interest per week. Which would you choose? (a) 14% per month; (b) 2% per week.

Such quantitative skills are important not only in selecting optimal financial products, but also in navigating one’s day-to-day personal finances. As shown in the above two questions, financial *numeracy* encompasses whether respondents can reason with numbers and whether they can use elementary arithmetic to solve financial problems.

The second dimension of financial knowledge we introduce is financial *awareness*. In comparison to numeracy which deals with mathematics-related questions, awareness emphasizes knowledge about fundamental financial planning tools as well as the details of basic financial products and services. We capture this type of knowledge through the following four questions in our endline survey.

- (1) Shantiben is preparing a budget for her household. Which of the following needs to be included in the budget? (a) Income only; (b) Expenses only; (c) Income and expenses.
- (2) Do you think you can open a savings account in a bank with amount as low as Rs. 50? (a) Can open an account; (b) Cannot open an account.
- (3) If I have a savings account in a bank and the bank closes down for some reason, will I get my money back? (a) Will get my money back; (b) Will not get my money back.
- (4) Manojbhai recently borrowed some money from a local moneylender. He wanted to buy some clothes for his children for Diwali (festival). What do you think about Manojbhai’s loan? (a) It is a productive loan; (b) It is an unproductive loan.

Notably, financial awareness does not involve any calculations. Instead, it relates to whether individuals are familiar with the different parts of a household budget, the bank account opening requirements or deposit insurance in their local context, and the use of loans for productive purposes. Financial awareness may therefore be regarded as a simpler and more rudimentary form of financial knowledge than financial numeracy.

The third dimension we propose is financial *attitudes*, which involve individual perspectives about the benefits of financial products. We consider attitudes as an important element of financial literacy because it has serious implications for financial outcomes: for instance, if one holds negative attitudes towards saving, then one will be less inclined to set money aside for the future. To assess financial attitudes, our endline presented households with the following three questions.

- (1) Rameshbhai does plastering on tall buildings. It is a dangerous job and he is worried that if he gets injured, his family’s income will become inadequate to meet their needs. If Rameshbhai comes to you for advice what would you suggest? (a) Take up some other (different) work; (b) Purchase health/life/accident insurance; (c) Increase savings.
- (2) Vimlaben has a very bright child who is currently in secondary school but will probably do well in university. She is worried how her family will pay for the child’s education. If Vimlaben comes to you for advice what would you

⁷ The “Big Three” financial literacy questions by Lusardi and Mitchell (2009) are: (1) Suppose you had \$100 in a savings account and the interest rate was 2 percent per year. After 5 years, how much do you think you would have in the account if you left the money to grow: more than \$102, exactly \$102, or less than \$102?; (2) Imagine that the interest rate on your savings account was 1 percent per year and inflation was 2 percent per year. After 1 year, would you be able to buy more than, exactly the same as, or less than today with the money in this account?; (3) Do you think that the following statement is true or false? “Buying a single company stock usually provides a safer return than a stock mutual fund.”

⁸ Appendix A provides additional information on the development of the numeracy, awareness, and attitudes measures.

suggest? (a) Buy child life insurance policy; (b) Borrow money from a moneylender; (c) Open a savings account in a bank; (d) Save at home; (e) Discontinue education.

(3) Do you think making a budget is helpful? (a) Yes; (b) No.

4. Empirical method

Our empirical approach consists of two components. In the first part, we measure average treatment effects (ATE) on financial outcomes as in Carpena et al. (2017). Specifically, we employ the intent-to-treat estimator and obtain the ATE using separate regressions comparing each treatment arm against the control arm. In the second step, we implement causal mediation analysis.

4.1. Estimating average treatment effects

Our experimental design with random assignment enables us to isolate the ATE in a straightforward manner. That is, we estimate the linear regression

$$Y_i = \alpha_1 + \beta_1 T_i + \xi_1^T \mathbf{X}_i + \epsilon_{1i} \quad (1)$$

where T_i is a dummy variable equal to 1 if individual i was randomly assigned to the treatment group, and 0 if s/he was assigned to the control group. We implement Eq. (1) separately for four different types of treatment variables T_i : (i) financial education only; (ii) financial education with goal setting; (iii) financial education with counseling; and (iv) all three of financial education, goal setting, and counseling. \mathbf{X}_i represents our baseline controls, which consists of the respondent's monthly discount rate, together with indicators for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters.⁹ Further, we include stratification controls in all specifications; strata are defined by gender, whether the respondent is currently a microfinance client, and neighborhood. Note that since neighborhoods were mutually exclusive across waves, we do not add wave fixed effects. Standard errors are clustered at the wave-class level.

For outcomes Y_i , we use endline data on financial knowledge and financial behavior. The financial behaviors we examine are budgeting (i.e., whether the household tried making a budget in the last six months); savings (i.e., whether the household has a savings account); borrowing (i.e., whether, conditional on borrowing in the last six months, the loan was for productive purposes); and insurance (i.e., whether the household purchased life insurance in the last six months). We chose these outcomes because as shown in Carpena et al. (2017), they exhibit statistically significant ATEs. We then examine the mechanisms of these treatment effects using mediation analysis.

4.2. Causal mediation analysis

The principal aim of this paper is to go beyond the ATE and to quantify the effect of the treatment that operates through a particular channel. That is, we focus on the causal mechanism by which the treatment T causally affects outcomes Y through a mediator M . In our analysis, we consider three mediator variables corresponding to our three dimensions of financial knowledge, namely, numeracy, awareness, and attitudes. Our goal then is to decompose the ATE into an indirect effect (representing a given mechanism) and a direct effect (representing all other channels).

Decomposing the ATE to study mechanisms is not a trivial exercise. As Green, Ha, and Bullock (2010) point out, conventional regression approaches to isolate causal mechanisms that have been used in different fields—such as political science, psychology, and public health—all rely on strong and often implausible identification assumptions. Moreover, in the presence of confounding variables that are influenced by the treatment and affect both the mediator and the outcome, conditioning on the mediator can lead to spurious effects (Rosenbaum, 1984). As a result, such an analysis suffers from bias and loses causal inference, even in experimental studies.

Recent developments in causal mediation analysis take these bias concerns seriously, and several papers have presented new estimation methods to isolate the influence of a mediating variable (e.g., Imai et al., 2011; Acharya et al., 2016). The common thread among these methods is that they focus on issues of estimation as well as clarify the underlying identification assumptions. For our study, we apply two leading approaches to mediation analysis, which we describe below.

4.2.1. Estimating average causal mediation effects

While the experimental design allows us to obtain the ATE, it does not by itself offer possibilities for capturing the mechanisms underlying the change in financial outcomes. To this end, our first approach adopts the causal mediation analysis proposed in Imai et al. (2011) and Imai, Keele, and Yamamoto (2010). We start by defining the *indirect effect* or the *causal mediation effect* as

$$\delta_i(t) \equiv Y_i(t, M_i(1)) - Y_i(t, M_i(0))$$

⁹ These variables were chosen as controls since they all relate to financial behavior. Further, the variable for inconsistent time preferences showed imbalance at baseline.

where $M_i(1)$ and $M_i(0)$ are the potential value of the mediator when individual i is assigned to treatment and control, respectively. $\delta_i(t)$ is the change in Y_i due to the change in M_i from control to treatment, holding the person's treatment status constant at t . This quantity captures the effect of the treatment on the outcome via the mediating variable: since we fix the treatment and change only the mediator, $\delta_i(t)$ isolates the impact of the variable M from all other channels. Further, we see that if $M_i(1) = M_i(0)$ —so that the treatment has no effect on the mediator—then $\delta_i(t)$ is zero.

Next, the *direct effect* of the treatment encompasses all other mechanisms and is given by

$$\zeta_i(t) \equiv Y_i(1, M_i(t)) - Y_i(0, M_i(t)).$$

$\zeta_i(t)$ is the impact of the treatment that is not from the mediator. In this sense, $\zeta_i(t)$ is the portion of the treatment effect that remains after the indirect effect $\delta_i(t)$ is accounted for, and the treatment effect is the sum of the direct and indirect effects. For our analysis, we are primarily interested in the *Average Causal Mediation Effect (ACME)*, $\bar{\delta}(t)$, and the *Average Direct Effect (ADE)*, $\bar{\zeta}(t)$. As in the ATE, these averages are obtained by taking the expected value over all individuals i .¹⁰

Given the formal definitions of the ACME and the ADE, an outstanding question is how these parameters are identified empirically. Our randomized experiment allows us to estimate the ATE, but it is not sufficient to estimate the ACME and the ADE because the potential outcomes required for $\delta_i(t)$ and $\zeta_i(t)$ are never observable. Although we know the values of $Y_i(1, M_i(1))$ for the treatment and $Y_i(0, M_i(0))$ for the control, we observe neither $Y_i(1, M_i(0))$ nor $Y_i(0, M_i(1))$.

To identify the ACME and ADE, we follow [Imai, Keele, and Yamamoto \(2010\)](#) and impose a set of assumptions collectively referred to as *Sequential Ignorability (SI)*.¹¹ SI consists of two parts.

SI Assumption 1. Given baseline characteristics, the treatment must be ignorable or as-if random, meaning it is independent of potential outcomes and potential mediators. Mathematically, this assumption is expressed as $\{Y_i(t', m), M_i(t)\} \perp T_i \mid X_i = x$. Moreover, in our study, this assumption holds because all treatments are randomly assigned.

SI Assumption 2. Conditional on the actual treatment status and baseline characteristics, the observed mediator is ignorable or as-if random. Mathematically, this statement is expressed as

$$Y_i(t', m) \perp M_i(t) \mid T_i = t, X_i = x.$$

This means that once T and X are taken into account, there are no omitted variables that influence both M and Y . In other words, this statement says that the observed mediator is statistically independent of the potential outcome, given the individual's treatment assignment and pre-treatment covariates. This assumption will be violated if there are any unobservable pre-treatment confounders that affect both the mediator and the outcome variable. The assumption is also violated if there are any unobserved or observed post-treatment confounders. The assumption embodied in SI Assumption 2 is therefore quite strong, and we return to this point below.

As [Imai, Keele, and Yamamoto \(2010\)](#) show, the upside for making this strong assumption is that it becomes possible to consistently estimate the ACME and ADE without any additional distributional or functional form assumptions regarding the mediator or outcome variables. Under linearity, the ACME and ADE can be estimated using the following system of two regressions:

$$M_i = \alpha_2 + \beta_2 T_i + \xi_2^T X_i + \epsilon_{2i} \quad (2)$$

$$Y_i = \alpha_3 + \beta_3 T_i + \gamma_3 M_i + \xi_3^T X_i + \epsilon_{3i}. \quad (3)$$

The ACME is given by the product of the ordinary least squares (OLS) estimates $\hat{\beta}_2 \cdot \hat{\gamma}_3$ from [Eqs. \(2\) and \(3\)](#) above,¹² while the ADE is $\hat{\beta}_3$ from [Eq. \(3\)](#). Importantly, [Imai, Keele, and Yamamoto \(2010\)](#) prove that $\hat{\beta}_2 \cdot \hat{\gamma}_3$ is a valid estimate of the ACME under SI.¹³ Standard errors and confidence intervals for the ACME are obtained using a quasi-Bayesian Monte Carlo approximation ([King et al., 2000](#)) based on the implementation by [Hicks and Tingley \(2011\)](#).

In practice, we estimate [Eqs. \(2\) and \(3\)](#) using four different types of treatment variables T_i as described earlier in [Section 4.1](#) (i.e., financial education only; financial education and goal setting; financial education and counseling; and all three treatments). We also consider three different types of mediating variables M_i (i.e., financial numeracy, financial awareness, and financial attitudes). Furthermore, the regression in [Eqs. \(2\) and \(3\)](#) are estimated using the same sample, which as in the estimation of ATEs, consist of those respondents belonging to the control group and the particular treatment group represented by T_i .

SI Assumption 1 allows us to identify the coefficients in [Eq. \(2\)](#), and as already mentioned, it is satisfied in our context because the treatments were randomized. Meanwhile, SI Assumption 2 allows us to identify the coefficients in [Eq. \(3\)](#). Because this second part of SI is a strong assumption, we now turn to understanding the confounding variables that result in its violation.

¹⁰ The ACME is $\bar{\delta}(t) \equiv E[Y_i(t, M_i(1)) - Y_i(t, M_i(0))]$, the ADE is $\bar{\zeta}(t) \equiv E[Y_i(1, M_i(t)) - Y_i(0, M_i(t))]$, and the ATE is $\bar{\tau} \equiv E[Y_i(1, M_i(1)) - Y_i(0, M_i(0))]$.

¹¹ [Imai, Keele, and Yamamoto \(2010\)](#) call these two assumptions sequential ignorability because two ignorability assumptions are made sequentially: first, it is assumed that treatment is ignorable; second, it is assumed that the mediator is ignorable given the actual treatment status and baseline characteristics.

¹² An equivalent method is to calculate $ACME = \hat{\beta}_1 - \hat{\beta}_3$, where the $\hat{\beta}_1$ is from the OLS regression of [Eq. \(1\)](#).

¹³ In addition, we assume that the ACME does not depend on the treatment (i.e., $\bar{\delta}(1) = \bar{\delta}(0)$). This assumption can be relaxed, and we explain this point further in [Section 6.3](#).




Fig. 1. Unobserved Pre-Treatment Confounder

Notes: In this diagram, sequential ignorability is not satisfied because U is unobserved and influences both M and Y .

There are two potential confounders in the $M \rightarrow Y$ relationship: (1) unobserved (or omitted) pre-treatment confounders, and (2) post-treatment confounders (whether observed or unobserved). We now explain each confounder in more detail. Throughout, we use the example where T is the financial education treatment, M is financial awareness, and Y is opening a bank savings account.

4.2.2. Unobserved pre-treatment confounders

The first potential confounder for SI Assumption 2 is an *unobserved pre-treatment* variable, which we denote as U . Fig. 1 shows a diagram for this confounder. In this figure, the direct effect of T is captured by the $T \rightarrow Y$ path, while the indirect effect is $T \rightarrow M \rightarrow Y$.

To illustrate how a pre-treatment confounder violates SI Assumption 2 in our study, suppose that among those receiving the financial education treatment, subjects who are *ex ante* more interested in finance (U) have higher *ex post* financial awareness (M). This may be because those who are more enthusiastic about finance put more effort into learning from the education program. As a result, U and M are correlated. At the same time, people who are more interested in finance might also be more likely to open a bank account, because they may feel more confident approaching formal institutions. Thus, U impacts Y as well. The joint effects of U on both M and Y imply that U is an omitted variable in Eq. (3), particularly if it is unobserved and cannot be included in the baseline controls X . SI Assumption 2 is not satisfied, and as seen in Eq. (3), we obtain a biased estimate of γ_3 , one of the parameters needed to calculate the ACME.

Another example of a pre-treatment variable that may confound the $M \rightarrow Y$ link is baseline risk aversion. More risk averse respondents may have different financial behaviors than their less risk averse counterparts. As such, there may be a path $U \rightarrow Y$, where U is baseline risk aversion. If, in addition to affecting outcomes, baseline risk aversion also influences financial awareness M , then SI Assumption 2 is violated. For instance, among those receiving financial education, individuals who are more risk averse at baseline may become more aware about deposit insurance due to the financial education program. This means that $U \rightarrow M$. Because of these simultaneous impacts of baseline risk aversion U on both M and Y , SI Assumption 2 fails, and the ACME is not identified.

In practice, we observe the pre-treatment variables for both interest in finance and risk aversion in our study since our baseline questionnaire covered these topics,¹⁴ and we include these variables in our baseline controls. Because we have controlled for baseline interest in finance and risk aversion in regression Eqs. (2) and (3), these variables no longer induce violations of SI Assumption 2. Still, we acknowledge that even with these controls, we cannot guarantee that there are no unobserved pre-treatment variables contaminating the $M \rightarrow Y$ relationship. From a practical perspective, it is impossible to think that all baseline confounders U are measurable. And putting measurement issues aside, it is perhaps unlikely that all possible confounders would have been included in our baseline survey. For this reason, SI is indeed a strong assumption.

4.2.3. Unobserved or observed post-treatment confounders

Apart from unobserved pre-treatment covariates, the second potential confounder in the $M \rightarrow Y$ link is a *post-treatment* variable (whether observed or unobserved), which we denote as N . By post-treatment variables, we mean those variables that are themselves a result of the treatment.

Fig. 2 illustrates the case of post-treatment confounding. In this figure, the variable N is a post-treatment confounder because it is impacted by T . Additionally, N affects both M and Y . Under this scenario, SI Assumption 2 does not hold. Recall that SI Assumption 2 states that M is ignorable given the treatment T and baseline variables X . There is nothing in this assumption about post-treatment variables N , so the ignorability of M must be satisfied even without conditioning on post-treatment variables. The presence of any post-treatment variable N that confounds the $M \rightarrow Y$ relationship thus violates Assumption 2, regardless of whether N is observed or unobserved.

What are examples of post-treatment confounding in our study? We return to the variable risk aversion. Assuming that the financial education treatment T impacts risk preferences, *ex post* risk aversion (N) could violate SI Assumption 2. For instance, suppose that financial education makes individuals less risk averse ($T \rightarrow N$). Further, suppose that as a result of

¹⁴ We measure interest in finance by asking “Generally, how interested are you in financial matters?” We measure risk aversion as specified in Section 2.4.




Fig. 2. Post-Treatment Confounder

Notes: In this diagram, sequential ignorability is not satisfied because the variable N , which is itself a result of the treatment, influences both M and Y .

this lower risk aversion, individuals are more likely to be aware about financial products ($N \rightarrow M$) and to open a bank savings account ($N \rightarrow Y$). If so, then *ex post* risk aversion N is a post-treatment confounder for the link between M and Y . Even if we observe post-treatment risk aversion and control for it in Eq. (3), doing so does not overcome the problem of post-treatment confounding due to N . This is because SI Assumption 2 states that M is ignorable given only the treatment and baseline covariates. Hence, the ignorability of M should hold without controlling for *ex post* risk aversion.

In summary, the second part of the SI assumption—that the mediator M is as-if random conditional on the treatment T and baseline controls X —embeds two requirements about the $M \rightarrow Y$ relationship: first, there are no unobserved pre-treatment confounders, and second, there are no observed or unobserved post-treatment confounders. For example, in our study, this means that all baseline variables that jointly cause financial awareness (M) and opening a bank savings account (Y) are observable and included in our controls. But in addition, none of these variables should be affected by financial education (T). If either of these conditions fail, then the direct and indirect effects of the treatment—namely, the ACME and the ADE—will not be identified.

4.2.4. Sequential g-Estimation

Because the SI assumption in Imai et al. (2011) is quite stringent, a number of papers in the methodological literature have explored the identification of mediation effects under weaker assumptions. One related study in this vein is Acharya et al. (2016), who propose sequential-g estimation, the second approach that we employ in this paper.

Sequential-g estimation is concerned with isolating the *Average Controlled Direct Effect (ACDE)*, defined as

$$ACDE = E[Y_i(1, m) - Y_i(0, m)].$$

We note that the ACDE is different from the ADE, $E[Y_i(1, M_i(t)) - Y_i(0, M_i(t))]$. The difference is that the ACDE captures the effect of the treatment when the mediator is held fixed at the same value m for all persons i . In contrast, with the ADE, the mediator is held constant at $M_i(t)$, the person-specific potential value of the mediator under the treatment status t . In this sense, the ACDE does not capture the “natural” direct effect of the treatment, since the value m is not the “natural” level of the mediator that would arise under a given treatment assignment. Intuitively, the ACDE can then be thought of as the direct effect when manipulating the mediator to be m .

To obtain the ACDE, we follow the procedure explained in Acharya et al. (2016), which consists of three steps. The first step is to regress the outcome on the mediator, treatment, and controls. Specifically, we estimate a regression that is similar to Eq. (3), but note that the controls include both pre-treatment and post-treatment variables, as will be explained below.

In the second step, we “demediate” the outcome Y using the estimate of γ_3 (i.e., the coefficient on M_i) from the previous regression. In other words, we obtain \tilde{Y}_i which is defined as

$$\tilde{Y}_i = Y_i - \hat{\gamma}_3 M_i.$$

Intuitively, by “demediating” the outcome and removing the effect of the mediator M , the variation that is left over in Y is due to the direct effect of the treatment.

In the third and final step, we regress \tilde{Y}_i on the treatment and baseline controls X_i as follows:

$$\tilde{Y}_i = \alpha_4 + \beta_4 T_i + \xi_4^T X_i + \epsilon_{4i} \quad (4)$$

where the coefficient β_4 gives us an estimate of the ACDE. Standard errors for this estimate are obtained using non-parametric bootstrapping.

As described previously, the ACDE differs from the ADE in that the ACDE is not the “natural” direct effect. The upside of focusing on the ACDE, however, is that it can be identified with weaker assumptions than SI, which is what we need to identify the ADE and the ACME. In particular, to identify the ACDE, we need a set of assumptions that Acharya et al. (2016) refer to as *Sequential Unconfoundedness (SU)*. SU consists of two parts.

SU Assumption 1. $\{Y_i(t, m), M_i(t)\} \perp T_i \mid X_i = x$. This means there are no omitted variables for the effect of T on Y , conditional on pre-treatment confounders X .

SU Assumption 2. $Y_i(t, m) \perp M_i \mid T_i = t, X_i = x, N_i = n$. This means there are no omitted variables for the effect of M on Y , conditional on the treatment T , pre-treatment confounders X , and post-treatment confounders N .

The first part of SU is analogous to the first part of SI, which is satisfied when the treatments are randomly assigned. The second part of SU, however, relaxes the second part of SI by allowing for conditioning on post-treatment confounders N .

While the advantage of sequential-g estimation is that it relies on less stringent identification assumptions, the drawback is that we cannot recover the ADE (and consequently, the ACME, since $ATE = ACME + ADE$) without additional assumptions. To obtain the ADE using sequential-g estimation, we need to impose a no-interactions assumption (Robins, 2003), which holds that

$$Y_i(1, m) - Y_i(0, m) = Y_i(1, m') - Y_i(0, m').$$

This no-interactions assumption means that for any person i , the controlled direct effect is the same regardless of the value m at which the mediator is fixed. In other words, the controlled direct effect for any person i does not depend on m . In our study, this implies the effect of the financial education intervention (T) when we force individuals to have a high level of financial awareness (M) is equivalent to when we force them to have a low level of financial awareness, and this must be the case for all individuals. If this assumption holds, then the ACDE is equal to the ADE, and we can obtain the ACME through sequential-g estimation (Acharya et al., 2016).

We highlight that this no-interactions assumption is also quite strong, as it must be true at the individual level (i.e., for each person i). Hence, there is an important trade-off between the causal mediation approach in Imai et al. (2011) and the sequential-g estimation in Acharya et al. (2016). To identify the ACME and ADE as in Imai et al. (2011), we need that M is ignorable conditional on treatment and only pre-treatment variables. Acharya et al. (2016) has less restrictive ignorability assumptions as it allows for conditioning on post-treatment variables, but we require an additional no-interactions assumption to identify the ACME. Given this trade-off, we show estimates using both approaches. As we will describe in the results section, we obtain very similar estimates whether we use causal mediation analysis, sequential-g estimation with only pre-treatment controls, and sequential-g estimation with both pre-treatment and post-treatment controls.¹⁵

4.2.5. Sensitivity analysis

Finally, because identification of the ACME requires a strong assumption with SI, it is important to understand how our results change when SI is violated. To this end, we employ the sensitivity analysis proposed by Imai, Keele, and Yamamoto (2010) and Imai, Keele and Tingley (2010). Although the SI assumption cannot be tested directly, sensitivity analysis allows us to understand how the ACME would change for different degrees of violation of SI Assumption 2. The SI assumption implies that the correlation between the error terms ϵ_{2i} from Eq. (2) and ϵ_{3i} from Eq. (3) would be zero. Conversely, non-zero values of this correlation, which we denote as ρ , would imply that sequential ignorability has been violated.

To illustrate the sensitivity analysis, consider a setting where an individual's unobserved ability positively correlates with financial awareness (M) and opening a bank savings account (Y). This means $\rho > 0$, and because of the non-zero correlation between ϵ_{2i} and ϵ_{3i} , the ACME estimate is biased. Thus, ρ serves as the sensitivity parameter, where the larger values of ρ in absolute terms result in larger bias in the ACME. The sensitivity analysis relaxes the condition that $\rho = 0$ and then estimates Eqs. (2) and (3) for different values of ρ . With these estimates, we then show the graphical plot of a given value of ρ against the true ACME. Doing so allows us to quantify the degree of sensitivity by looking at how large ρ must be for the mediation effect to be insignificant.

While sensitivity analysis can be informative about whether the estimates obtained under the SI assumption are valid, there are two critical limitations to consider (Imai et al., 2011). First, this sensitivity analysis evaluates the sensitivity of the estimates to only unobserved *pre-treatment* confounding. Hence, it is not able to say anything about post-treatment confounding. Nevertheless, as we will show later, we obtain very similar estimates when we implement the causal mediation analysis framework of Imai et al. (2011)—controlling for only pre-treatment variables—and sequential-g estimation proposed by Acharya et al. (2016)—controlling for both pre-treatment and post-treatment variables. These patterns provide suggestive evidence that relative to pre-treatment confounding, post-treatment confounding may be less of a concern in our study, though we acknowledge that there may be other post-treatment variables that we have not accounted for.

Second, the SI assumption cannot be proven or disproven with observable data. There may be unobserved pre-treatment confounders that violate SI, but we will never know whether this is the case because the only information we have is what we can observe in our dataset. For this reason, when we implement sensitivity analysis, we do not have an objective criterion for determining whether the SI assumption holds (Imai et al., 2011). We can only investigate how our estimates of the ACME would change for different degrees of violation of SI.

¹⁵ The pre-treatment controls we use in sequential-g estimation are the same as when we implement Imai et al. (2011). These controls are the respondent's monthly discount rate and dummy variables for whether the respondent is risk averse, has consistent time preferences, self-reports having a hard time saving, and self-reports being interested in financial matters. The post-treatment controls are the same set as the pre-treatment variables, with the exception of interest in financial matters. We are not able to include this variable as a post-treatment control as it was not collected in our endline survey for logistical reasons, particularly to manage the length of the questionnaire.

Table 3
Average Treatment Effects.

	Budgeting	Savings	Borrowing	Insurance
	Has tried making a budget in the last 6 months (1)	Has a savings account (2)	Loan purpose: Business, education, or purchase of durable goods (3)	Bought life insurance in the last 6 months (4)
Panel A. Financial Education Only				
Financial Education Only	0.138*** (0.038)	-0.005 (0.039)	0.059 (0.087)	0.001 (0.017)
Adj. R-squared	0.113	0.107	0.100	-0.008
Number of Observations	487	487	235	487
Panel B. Financial Education and Goal Setting				
Financial Education and Goal Setting	0.169*** (0.045)	0.081* (0.045)	-0.008 (0.087)	-0.001 (0.016)
Adj. R-squared	0.099	0.126	0.038	-0.024
Number of Observations	465	465	223	465
Panel C. Financial Education and Counseling				
Financial Education and Counseling	0.399*** (0.049)	0.168*** (0.043)	0.176** (0.077)	0.045* (0.025)
Adj. R-squared	0.226	0.151	0.102	0.016
Number of Observations	468	468	227	468
Panel D. All Three Treatments				
All Three Treatments	0.475*** (0.041)	0.107** (0.045)	-0.057 (0.071)	0.042* (0.021)
Adj. R-squared	0.247	0.119	0.046	0.013
Number of Observations	476	476	228	476
Control Group Mean	0.155	0.310	0.333	0.035

Notes: This table presents regressions estimating the Average Treatment Effects (ATE) on financial outcomes. Each panel refers to a different treatment. Financial Education Only is a dummy equal to 1 for an individual who was assigned to the financial education program but not financial counseling nor goal setting. Financial Education and Goal Setting is a dummy equal to 1 for an individual who was assigned to the financial education and goal setting treatments, but not the financial counseling treatment. Financial Education and Counseling is a dummy equal to 1 for an individual who was assigned to the financial education and financial counseling treatments, but not the goal setting treatment. All Three Treatments is a dummy equal to 1 for an individual who was assigned to all three (i.e., financial education, counseling, and goal setting). In each panel, the regression sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to a particular treatment. For example, in Panel A, the regression sample consists of those in the Financial Education Only treatment and those assigned to control. All regressions include baseline control variables (i.e., monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters). In addition, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). Standard errors are clustered at the wave-class level. *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

5. Empirical results

5.1. Average treatment effects

Table 3 presents ATE estimates using the specification described in Eq. (1). The outcomes in this table consist of dummy variables for the four main financial behaviors targeted by the financial education curriculum: (1) budgeting (whether the respondent tried to write a budget in the past 6 months); (2) savings (whether the household has a savings account); (3) loans (whether, conditional on borrowing in the past 6 months, the loan was obtained for productive purposes such as business, education, or durable goods); and (4) insurance (whether the household purchased life insurance in the last 6 months). Estimating the ATEs for these behaviors is a necessary first step in our study, because in the mediation analysis below, we focus only on those outcomes for which the interventions had a statistically significant effect.

The results in Table 3 indicate that financial education alone is not a panacea for changing financial behavior. We find that the effects of the financial education only intervention on savings, borrowing, and insurance are all statistically insignificant, with magnitudes that are very close to zero (Table 3, Panel A, Columns 2 to 4). Nevertheless, the financial education only treatment does foster better budgeting: individuals assigned to financial education are 13.8 percentage points more likely than control to have written a budget in the past 6 months (Table 3, Panel A, Column 1). Since the control group average is only 15.5 percent, the treatment coefficient represents a sizable improvement in respondents' propensity to start writing a household budget.

Despite the limited scope of financial education alone in shaping financial outcomes, complementing it with individualized add-ons results in substantial positive effects. Adding goal setting to financial education leads to a 16.9 and 8.1 percentage point increase in the likelihood of making a budget and opening a savings account, respectively (Table 3, Panel

Table 4
Effects on Financial Knowledge.

	Financial Numeracy (1)	Financial Awareness (2)	Financial Attitudes (3)
Panel A. Financial Education Only			
Financial Education Only	-0.037 (0.032)	0.133*** (0.023)	0.071*** (0.023)
Adj. R-squared	0.080	0.098	0.124
Number of Observations	487	487	487
Panel B. Financial Education and Goal Setting			
Financial Education and Goal Setting	-0.022 (0.034)	0.129*** (0.026)	0.121*** (0.022)
Adj. R-squared	0.100	0.136	0.147
Number of Observations	465	465	465
Panel C. Financial Education and Counseling			
Financial Education and Counseling	0.015 (0.029)	0.111*** (0.023)	0.102*** (0.029)
Adj. R-squared	0.059	0.172	0.135
Number of Observations	468	468	468
Panel D. All Three Treatments			
All Three Treatments	0.006 (0.038)	0.152*** (0.021)	0.147*** (0.021)
Adj. R-squared	0.022	0.163	0.124
Number of Observations	476	476	476
Control Group Mean	0.703	0.672	0.727

Notes: This table presents regressions estimating the effect of the various treatments on the mediator variables, which capture three different facets of financial knowledge. Each panel refers to a different treatment. Financial Education Only is a dummy equal to 1 for an individual who was assigned to the financial education program but not financial counseling nor goal setting. Financial Education and Goal Setting is a dummy equal to 1 for an individual who was assigned to the financial education and goal setting treatments, but not the financial counseling treatment. Financial Education and Counseling is a dummy equal to 1 for an individual who was assigned to the financial education and financial counseling treatments, but not the goal setting treatment. All Three Treatments is a dummy equal to 1 for an individual who was assigned to all three (i.e., financial education, counseling, and goal setting). In each panel, the regression sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to a particular treatment. For example, in Panel A, the regression sample consists of those in the Financial Education Only treatment and those assigned to control. All regressions include baseline control variables (i.e., monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters). In addition, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). Standard errors are clustered at the wave-class level. *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

B, Columns 1 and 2). Augmenting financial education with counseling shows even stronger effects, with a 39.9 percentage point gain in the probability of making a budget, 16.8 percentage points of having a savings account, 17.6 percentage points of borrowing for productive purposes, and 4.5 percentage points of buying life insurance (Table 3, Panel C, Columns 1 to 4). Further, the combination of all three treatments exhibit positive effects of 47.5 percentage points for budgeting, 10.7 percentage points for savings, and 4.2 percentage points for insurance (Table 3, Panel D, Columns 1, 2, and 4).

Since this paper aims to investigate financial knowledge as a mechanism, we likewise examine the ATEs of the interventions on financial knowledge scores measured through a series of questions at endline. This analysis is presented in Table 4, where the outcome variables are the scores for financial numeracy (column 1), awareness (column 2), and attitudes (column 3). These scores are the proportion of correct answers in each of the financial knowledge dimensions. Our estimates show that none of the treatments have any significant effect on numeracy skills. However, all four treatment combinations have statistically significant positive impacts on both awareness and attitudes. Additionally, the magnitude of these effects is non-trivial: compared to the control group average of 67 and 73 percent, financial awareness and attitudes each improve by as much as 15 percentage points. Importantly, the null effects of the treatments on financial numeracy allow us to rule out numeracy as a potential mechanism. We discuss this point in more detail in Section 6.1.

5.2. Mediating effects of financial numeracy, awareness, and attitudes

We now consider the principal question of interest in this paper: how do the different dimensions of financial knowledge mediate the impact of financial education? We examine this question in Tables 5 to 8, where we present causal mediation analysis results using the approach outlined in Section 4.2.1. In each table, we focus on one experimental intervention as well as the financial outcomes for which that intervention had statistically significant ATEs. Panel A reports the coefficients from estimating Eq. (3), while Panel B presents the ACME and ADE estimates.

Our results reveal four notable patterns that shed light on the role of numeracy, awareness, and attitudes as mechanisms. First, consistent with the nil ATEs on numeracy scores, we find that numeracy does not serve as a mediator for financial

Table 5
Causal Mediation: Financial Education Treatment.

	Budgeting		
	Has tried making a budget in the last 6 months		
	(1)	(2)	(3)
Panel A. Coefficient Estimates			
Financial Education Only	0.140*** (0.034)	0.110*** (0.035)	0.113*** (0.032)
Endline Numeracy Score	0.065 (0.064)		
Endline Awareness Score		0.206*** (0.065)	
Endline Attitudes Score			0.348*** (0.061)
R-Squared	0.035	0.046	0.087
Number of Observations	487	487	487
Control Group Mean		0.155	
Panel B. Estimates of ACME, ADE, and ATE			
ACME	-0.002 (0.003)	0.027*** (0.010)	0.025*** (0.008)
ADE	0.140*** (0.034)	0.110*** (0.035)	0.113*** (0.032)
ATE		0.138*** (0.038)	
% of ATE Mediated	-1%	20%	18%

Notes: This table presents estimates of the Average Causal Mediation Effect (ACME) and the Average Direct Effect (ADE) of the Financial Education Only treatment. The mediator variables considered are financial numeracy (column 1), financial awareness (column 2), and financial attitudes (column 3). Financial Education Only is a dummy equal to 1 for an individual who was assigned to the financial education program but not financial counseling nor goal setting. In Panel B, the ADE is replicated from the coefficient estimate of the treatment variable in Panel A, while the ATE is replicated from the coefficient estimate of the treatment variable from Table 3. In all regressions, the sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to the Financial Education Only treatment. All regressions include baseline control variables (i.e., monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters). In addition, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

education in any way. For instance, the ACME of numeracy in Table 5, Column 1 is very close to zero, indicating that none of the positive impacts of the financial education only treatment on budgeting is channeled through numeracy. This pattern of zero ACME for numeracy persists in all financial education interventions and all financial outcomes that we study—that is, whether we consider financial education augmented with goal setting vis-à-vis budgeting or savings (Table 6, Columns 1 and 4); financial education with counseling vis-à-vis budgeting, savings, borrowing, or insurance (Table 7, Columns 1, 4, 7, and 10); and all three treatments vis-à-vis budgeting, savings, or insurance (Table 8, Columns 1, 4, and 7). Importantly, the null ACME of numeracy are estimated relatively precisely, as the confidence intervals are all tightly centered around zero.

Second, notwithstanding the zero ACMEs of numeracy, we observe that across all treatment combinations, the mediating effects of both attitudes and awareness on budgeting appear to be positive and statistically significant. Specifically, Table 5 shows an ACME of 2.7 percentage points for awareness and 2.5 percentage points for attitudes, both statistically significant at the 1% level. These ACMEs amount to 20% and 18% of the ATE of the financial education only intervention on budgeting, respectively. Moreover, we find similar results when we examine the ACMEs for the rest of the treatments. For example, in the intervention involving financial education with counseling, 8% of the ATE on budgeting operates through awareness, while 10% is through attitudes (Table 7, Columns 2 and 3). Notably, for all variants of financial education treatments in our study, we cannot statistically distinguish the ACME of awareness on budgeting from that attitudes.¹⁶ Thus, these estimates provide suggestive evidence that both awareness and attitudes may be equally important for improvements in budgeting behavior.

Third, whereas awareness and attitudes function similarly in fostering household *budgeting*, there seems to be important differences in how they mediate household *savings*. Our results show that only attitudes (and not awareness) mediates the

¹⁶ In particular, across the estimates in Tables 5 to 8, the differences in the ACME of awareness and attitudes all have 95% confidence intervals that contain zero.

Table 6
Causal Mediation: Financial Education and Goal Setting Treatment.

	Budgeting			Savings		
	Has tried making a budget in the last 6 months			Has a savings account		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Coefficient Estimates						
Financial Education and Goal Setting	0.169*** (0.039)	0.149*** (0.043)	0.132*** (0.042)	0.082** (0.040)	0.080** (0.039)	0.061 (0.042)
Endline Numeracy Score	0.000 (0.064)			0.035 (0.071)		
Endline Awareness Score		0.155* (0.078)			0.005 (0.067)	
Endline Attitudes Score			0.301*** (0.073)			0.160* (0.082)
Adj. R-Squared	0.035	0.042	0.072	0.010	0.010	0.018
Number of Observations	465	465	465	465	465	465
Control Group Mean		0.155			0.310	
Panel B. Estimates of ACME, ADE, and ATE						
ACME	-0.000 (0.002)	0.020** (0.011)	0.036*** (0.011)	-0.001 (0.003)	0.001 (0.009)	0.019** (0.011)
ADE	0.169*** (0.039)	0.149*** (0.043)	0.132*** (0.042)	0.082** (0.040)	0.080** (0.039)	0.061 (0.042)
ATE		0.169*** (0.045)			0.081* (0.045)	
% of ATE Mediated	0%	12%	21%	-1%	1%	23%

Notes: This table presents estimates of the Average Causal Mediation Effect (ACME) and the Average Direct Effect (ADE) of the Financial Education and Goal Setting treatment. The mediator variables considered are financial numeracy (columns 1 and 4), financial awareness (columns 2 and 5), and financial attitudes (columns 3 and 6). Financial Education and Goal Setting is a dummy equal to 1 for an individual who was assigned to the financial education and goal setting treatments, but not the financial counseling treatment. In Panel B, the ADE is replicated from the coefficient estimate of the treatment variable in Panel A, while the ATE is replicated from the coefficient estimate of the treatment variable from Table 3. In all regressions, the sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to the Financial Education and Goal Setting treatment. All regressions include baseline control variables (i.e., monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters). In addition, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

impact of financial education on savings. This is true for all treatment combinations. Depending on the intervention, we estimate an ACME of attitudes on savings ranging from 2 to 3 percentage points, corresponding to 11 to 31% of the ATE (Table 6 to 8, Column 6). In contrast, the ACME of awareness on savings is very close to zero and is statistically insignificant in all treatments (Tables 6 to 8, Column 5)—although as we will explain in the robustness section, we cannot rule out that these null effects are due to lack of statistical power. These results appear to suggest that simply increasing awareness about account opening requirements or deposit insurance may not be sufficient to enable individuals to undertake the complex financial action of opening a savings account. Changing attitudes and perceptions about financial products may be more important for financial education to more effectively induce poor households to start a bank account.

Finally, we find that none of our three dimensions of financial knowledge causally mediate effects on borrowing and insurance. On the one hand, the two most intensive interventions—financial education with counseling (which provides one-on-one financial advice) and the combination of all three treatments (financial education, counseling, and goal setting)—had significant positive ATEs on whether the respondent's household took out a loan for productive purchases or purchased life insurance (Table 4). On the other hand, the ACME for numeracy, awareness, and attitudes are all close to zero and statistically insignificant (Table 7, Columns 7 to 12; Table 8, Columns 7 to 9). Consequently, the treatment impacts on borrowing and insurance are likely mediated through pathways other than financial knowledge. Indeed, our results suggest that borrowing and insurance behavior are quite difficult to influence using financial education, as the magnitude of the ATEs for such outcomes are much smaller in absolute terms relative to budgeting and savings.

While the above patterns in our data shed light on the causal mechanisms, it is important to emphasize that these results rely on the strong assumption that the mediators are ignorable conditional on the treatment and baseline controls. As we will explain in Section 5.4, our findings are very sensitive to even small violations of this assumption, so it is important to keep SI in mind when interpreting our results. Further, statistical power may also be a concern, a point we return to in the robustness section.

5.3. Sequential-g estimation

We now investigate how our causal mediating effect estimates from the previous subsection differ when we employ sequential-g estimation, which relaxes the SI assumption by allowing for post-treatment controls. The estimates are

Table 7
Causal Mediation: Financial Education and Counseling Treatment.

	Budgeting			Savings			Borrowing			Insurance		
	Has tried making a budget in the last 6 months			Has a savings account			Loan purpose: Business, education, or purchase of durable goods			Bought life insurance in the last 6 months		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A. Coefficient Estimates												
Financial Education and Counseling	0.398*** (0.043)	0.367*** (0.046)	0.360*** (0.047)	0.167*** (0.038)	0.163*** (0.040)	0.149*** (0.037)	0.141*** (0.047)	0.123*** (0.044)	0.141*** (0.049)	0.045** (0.022)	0.048* (0.024)	0.042* (0.023)
Endline Numeracy Score	0.054 (0.069)			0.046 (0.074)			0.009 (0.097)			0.022 (0.021)		
Endline Awareness Score		0.295*** (0.098)			0.047 (0.089)			0.166 (0.133)			-0.025 (0.054)	
Endline Attitudes Score			0.388*** (0.081)			0.184** (0.086)			0.004 (0.118)			0.030 (0.029)
Adj. R-Squared	0.196	0.215	0.245	0.035	0.035	0.045	0.013	0.020	0.013	0.004	0.004	0.005
Number of Observations	468	468	468	468	468	468	227	227	227	468	468	468
Control Group Mean		0.155			0.310			0.333			0.035	
Panel B. Estimates of ACME, ADE, and ATE												
ACME	0.001 (0.002)	0.033*** (0.013)	0.040*** (0.013)	0.001 (0.002)	0.005 (0.010)	0.019** (0.010)	0.000 (0.005)	0.018 (0.016)	0.000 (0.016)	0.000 (0.001)	-0.003 (0.006)	0.003 (0.003)
ADE	0.398*** (0.043)	0.367*** (0.046)	0.360*** (0.047)	0.167*** (0.038)	0.163*** (0.040)	0.149*** (0.037)	0.141*** (0.047)	0.123*** (0.044)	0.141*** (0.049)	0.045** (0.022)	0.048* (0.024)	0.042* (0.023)
ATE		0.399*** (0.049)			0.168*** (0.043)			0.176** (0.077)			0.045* (0.025)	
% of ATE Mediated	0%	8%	10%	1%	3%	11%	0%	10%	0%	0%	-7%	7%

Notes: This table presents estimates of the Average Causal Mediation Effect (ACME) and the Average Direct Effect (ADE) of the Financial Education and Counseling treatment. The mediator variables considered are financial numeracy (columns 1, 4, 7, and 10), financial awareness (columns 2, 5, 8, and 11), and financial attitudes (columns 3, 6, 9, and 12). Financial Education and Counseling is a dummy equal to 1 for an individual who was assigned to the financial education and financial counseling treatments, but not the goal setting treatment. In Panel B, the ADE is replicated from the coefficient estimate of the treatment variable in Panel A, and the ATE is replicated from the coefficient estimate of the treatment variable from Table 3. In all regressions, the sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to the Financial Education and Counseling treatment. All regressions include baseline control variables (i.e., monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters). In addition, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table 8
Causal Mediation: All Three Treatments.

	Budgeting			Savings			Insurance		
	Has tried making a budget in the last 6 months			Has a savings account			Bought life insurance in the last 6 months		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A. Coefficient Estimates									
All Three Treatments	0.474*** (0.036)	0.449*** (0.037)	0.417*** (0.035)	0.106** (0.040)	0.093** (0.037)	0.074* (0.041)	0.042** (0.019)	0.044* (0.023)	0.035* (0.019)
Endline Numeracy Score	0.084 (0.065)			0.055 (0.062)			0.005 (0.025)		
Endline Awareness Score		0.169* (0.089)			0.093 (0.086)			-0.014 (0.048)	
Endline Attitudes Score			0.393*** (0.072)			0.225*** (0.076)			0.047 (0.035)
Adj. R-Squared	0.261	0.264	0.303	0.019	0.019	0.033	0.010	0.010	0.013
Number of Observations	476	476	476	476	476	476	476	476	476
Control Group Mean		0.155			0.310			0.035	
Panel B. Estimates of ACME, ADE, and ATE									
ACME	0.001 (0.003)	0.026* (0.014)	0.058*** (0.013)	0.000 (0.003)	0.014 (0.013)	0.033*** (0.012)	0.000 (0.001)	-0.002 (0.008)	0.007 (0.005)
ADE	0.474*** (0.036)	0.449*** (0.037)	0.417*** (0.035)	0.106** (0.040)	0.093** (0.037)	0.074* (0.041)	0.042** (0.019)	0.044* (0.023)	0.035* (0.019)
ATE		0.475*** (0.041)			0.107** (0.045)			0.042* (0.021)	
% of ATE Mediated	0%	5%	12%	0%	13%	31%	0%	-5%	17%

Notes: This table presents estimates of the Average Causal Mediation Effect (ACME) and the Average Direct Effect (ADE) of All Three Treatments. The mediator variables considered are financial numeracy (columns 1, 4 and 7), financial awareness (columns 2, 5, and 8), and financial attitudes (columns 3, 6, and 9). All Three Treatments is a dummy equal to 1 for an individual who was assigned to all three of financial education, counseling, and goal setting. In Panel B, the ADE is replicated from the coefficient estimate of the treatment variable in Panel A, and the ATE is replicated from the coefficient estimate of the treatment variable from Table 3. In all regressions, the sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to All Three Treatments. All regressions include baseline control variables (i.e., monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters). In addition, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

shown in Appendix Tables A1 to A4. Panel A of these tables reports the results of sequential-g estimation controlling for only pre-treatment variables in the first step of the estimation process, while Panel B reports results when controlling for both pre-treatment and post-treatment variables in the first step of the estimation process. As the tables show, the sequential-g estimates remain very similar in both panels. Moreover, all estimates using the sequential-g estimation in Acharya et al. (2016) are very similar to the estimates from the causal mediation analysis method in Imai et al. (2011). Hence, our findings in the previous subsection appear to be robust to both approaches and their different identifying assumptions.

5.4. Sensitivity analysis

Our empirical results demonstrate that both financial awareness and attitudes may be important mediators for the impact of financial education, goal setting, and counseling on financial behavior. Nevertheless, the validity of these findings depends on the assumption that the mediators are ignorable, conditional on treatment and baseline characteristics. This assumption is quite strong because it is unlikely that respondents' awareness of and attitudes towards personal finance are random: for instance, it may be that those who have better awareness and attitudes also have higher unobserved ability. If unobserved ability influences financial outcomes through channels other than awareness and attitudes, then the assumption on the ignorability of the mediator is violated; the estimated ACME will be confounded with the impacts of unobservable characteristics.

To understand the robustness of our results to such biases, we conduct sensitivity analysis. Under SI, the correlation between the error terms in Eqs. (2) and (3), denoted $\rho \equiv \text{corr}(\epsilon_{2i}, \epsilon_{3i})$, is zero. We relax the condition that $\rho = 0$ by specifying hypothetical values of ρ . Then, we estimate the ACME under these non-zero correlations. The results from this exercise are shown in Figs. 3 to 6, where we plot the ACME vs. ρ for different combinations of the treatment, awareness and attitudes mediators, and outcomes from the main mediation analysis results. Here, the dashed line represents the ACME, while the shaded area shows the bootstrapped 95% confidence interval. Additionally, the ACME at $\rho = 0$ corresponds to the ACME estimates in Tables 5 to 8.

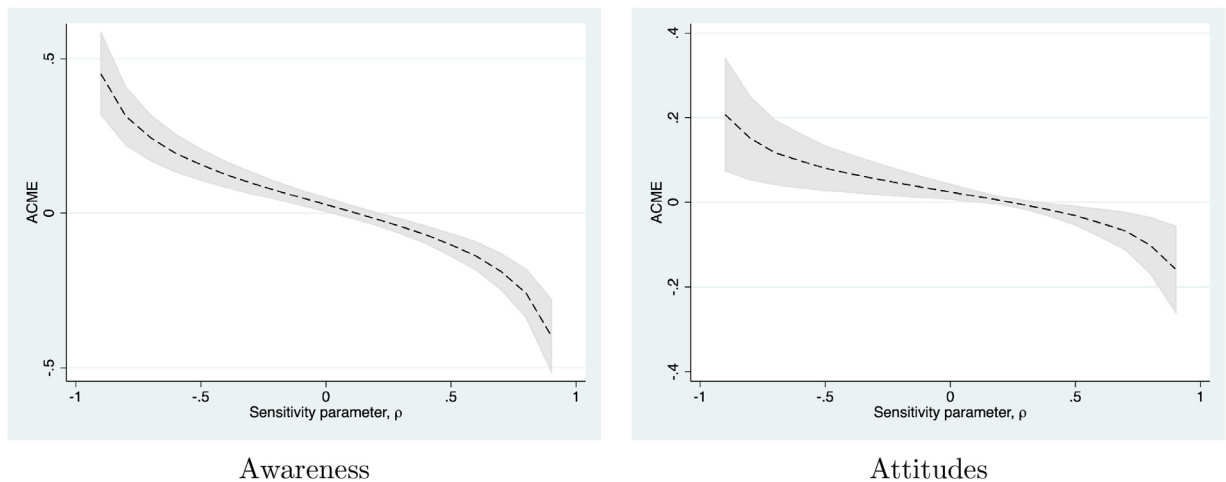


Fig. 3. Sensitivity Analysis (Treatment: Financial Education Only) Has tried making a budget in the last 6 months

Notes: This figure shows how the estimates of the Average Causal Mediation Effect (ACME) change with different values of ρ , defined as the correlation between the error terms in the regression of the mediator on the treatment (ϵ_{2i}) and the regression of financial outcomes on the treatment and the mediator (ϵ_{3i}). The dashed line represents the estimated ACME for the given mediator and for different values of ρ , while the shaded area represents the 95% confidence interval. The figure on the left considers financial awareness as the mediator, and that on the right uses financial attitudes as the mediator. The treatment variable is *Financial Education Only*, defined as a dummy equal to 1 for an individual who was assigned to financial education but not goal setting nor counseling. Sequential ignorability implies that ρ is equal to zero, so the ACME for $\rho = 0$ in the above figures corresponds to the ACME estimate in Table 5.

The results of the sensitivity analysis tell us how large ρ must be to drive the mediation effect to zero. If this value of ρ is small and our ACME estimate is statistically significant, then even small violations of SI will change our conclusions. Conversely, the larger the value of this ρ , the more robust the results are to unobserved baseline confounders. In Figs. 3 to 6, we see that our estimates of the ACME appear to be quite sensitive to changes in ρ ; if the results were not sensitive, the dashed line would be relatively flat, but this is not what we observe. This point regarding sensitivity should therefore be kept in mind when interpreting our ACME results. We also note that generally speaking, the results for attitudes appear to be much less sensitive than the results for awareness. For example, when we consider budgeting and the financial education only treatment, the value of ρ for which the ACME of awareness is zero is 0.112, while that for attitudes is 0.231 (Fig. 3).¹⁷

6. Related empirical approaches for studying causal mechanisms

In this section, we discuss how the results from our mediation analysis relate to—and differ from—the following alternative approaches: (1) testing the effects of the treatment on the mediator; (2) regressing outcomes on the treatment conditional on the mediator; (3) examining heterogeneity in the effects of the mediator by treatment and control groups; (4) instrumental variables (IV); (5) the front-door criterion (Pearl, 1995; Pearl 2000); and (6) randomizing the mediator.


6.1. Testing the effects of the treatment on the mediator

A simple alternative to the mediation analysis that we have implemented in the paper is to examine only the effects of the treatment T on the mediator M , without conducting further investigation of causal mechanisms (Keele, 2015). If T does not influence M , then the impact of T on outcomes Y through M —that is, the ACME—is zero. As explained in the empirical methods section, the ACME is given by $\hat{\delta}(t) \equiv E[Y_i(t, M_i(1)) - Y_i(t, M_i(0))]$, where $Y_i(t, m)$ as the potential outcome of person i if the treatment and mediator equal t and m , respectively. $M_i(1)$ is the potential value of the mediator under the treatment and $M_i(0)$ under control. Null effects of the treatment on the mediator imply that $M_i(1) = M_i(0)$, so we have that $\hat{\delta}(t) = 0$. Consequently, we can eliminate M as a potential mechanism for the treatment (Imai et al., 2011).

To illustrate how testing the effects of T on M works in our study, let us revisit the regression results Table 4. We see that the treatments increased the awareness and attitudes scores, with statistically significant effects at the 1% level. In contrast, for numeracy, we see no impact from any of the interventions. These results support the view that awareness and attitudes are potential mechanisms, but numeracy is not. Moreover, the null effects on numeracy appear to be precisely estimated as the confidence intervals are relatively tight around zero. We find 95% confidence intervals of [-0.10, 0.03] for the financial

¹⁷ Figs. 3 to 6 also indicate the direction of the bias in the ACME: if $\rho < 0$, we would underestimate the ACME, but if $\rho > 0$, we would overestimate the ACME. We explore the potential direction of the bias in Appendix B. This appendix also examines sensitivity to selection on unobservables, following Altonji et al. (2005) and Oster (2019).

Panel A: Has tried making a budget in the last 6 months



Panel B: Has a savings account




Fig. 4. Sensitivity Analysis (Treatment: Financial Education and Goal Setting) Panel A: Has tried making a budget in the last 6 months


Note: This figure shows how the estimates of the Average Causal Mediation Effect (ACME) change with different values of ρ , defined as the correlation between the error terms in the regression of the mediator on the treatment (ϵ_{2i}) and the regression of financial outcomes on the treatment and the mediator (ϵ_{3i}). The dashed line represents the estimated ACME for the given mediator and for different values of ρ , while the shaded area represents the 95% confidence interval. The figures on the left consider financial awareness as the mediator, and those on the right use financial attitudes as the mediator. The treatment variable is Financial Education and Goal Setting, defined as a dummy equal to 1 for an individual who was assigned to financial education and goal setting, but not counseling. Sequential ignorability implies that ρ is equal to zero, so the ACME for $\rho = 0$ in the above figures corresponds to the ACME estimate in Table 6.

education only treatment; [-0.09, 0.05] for financial education and goal setting; [-0.04, 0.07] for financial education with counseling; and [-0.07, 0.08] for financial education with both goal setting and counseling. The tightness of the confidence intervals around zero is a critical point: to conclude that numeracy is not a mediator, we must have precise estimates of null effects, rather than just noisy estimates that are statistically insignificant.


Although we observe relatively precise zeros for the average effect on numeracy, it may be that the treatments had heterogeneous impacts, so positive effects among some participants are cancelled out by negative effects among others. Hence, to more convincingly eliminate numeracy as a mediator, we must also rule out heterogeneity in the effects. Adopting the techniques outlined in Gerber and Green (2012), we provide two pieces of evidence demonstrating that the variance of the treatment effect on financial numeracy is likely to be zero. First, we consider the lower bound for the variance of the treatment effect on numeracy, $\text{var}(\tau_i)$, where τ_i is the difference between the treated and untreated potential outcomes. We find a value of around 0.01 for all treatments. Second, we tested the null hypothesis $\text{var}(\tau_i) = 0$ (i.e., homogenous treatment effects on numeracy) versus the alternative $\text{var}(\tau_i) > 0$. We find large p-values for this test for all treatments.¹⁸

¹⁸ Appendix C describes how the lower bound for $\text{var}(\tau_i)$ and the p-values were calculated.

Panel A: Has tried making a budget in the last 6 months




(A1) Awareness




(A2) Attitudes

Panel B: Has a savings account




(B1) Awareness




(B2) Attitudes

Panel C: Loan purpose: Business, education, or purchase of durable goods




(C1) Awareness




(C2) Attitudes

Panel D: Bought life insurance in the past six months



(D1) Awareness




(D2) Attitudes


Fig. 5. Sensitivity Analysis (Treatment: Financial Education and Counseling) Panel A: Has tried making a budget in the last 6 months

This figure shows how the estimates of the Average Causal Mediation Effect (ACME) change with different values of ρ , defined as the correlation between the error terms in the regression of the mediator on the treatment (ϵ_{2i}) and the regression of the financial outcomes on the treatment and the mediator (ϵ_{3i}). The dashed line represents the estimated ACME for the given mediator and for different values of ρ , while the shaded area represents the 95% confidence interval. The figures on the left consider financial awareness as the mediator, and those on the right use financial attitudes as the mediator. The treatment variable is *Financial Education and Counseling*, defined as a dummy equal to 1 for an individual who was assigned to financial education and counseling, but not goal setting. Sequential ignorability implies that ρ is equal to zero, so the ACME for $\rho = 0$ in the above figures corresponds to the ACME estimate in Table 7.

Panel A: Has tried making a budget in the last 6 months




(A1) Awareness




(A2) Attitudes

Panel B: Has a savings account




(B1) Awareness

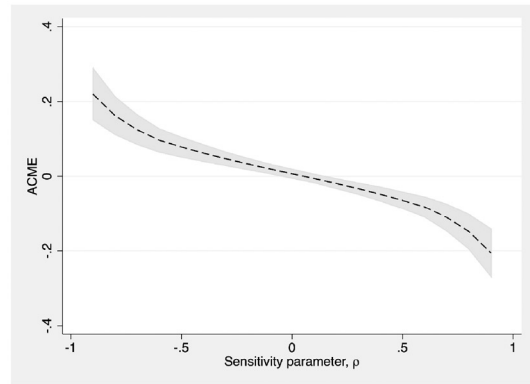


(B2) Attitudes

Panel C: Bought life insurance in the past six months



(C1) Awareness



(C2) Attitudes

Fig. 6. Sensitivity Analysis (Treatment: All Three Treatments) Panel A: Has tried making a budget in the last 6 months
 This figure shows how the estimates of the Average Causal Mediation Effect (ACME) change with different values of ρ , defined as the correlation between the error terms in the regression of the mediator on the treatment (ϵ_{2i}) and the regression of the financial outcomes on the treatment and the mediator (ϵ_{3i}). The dashed line represents the estimated ACME for the given mediator and for different values of ρ , while the shaded area represents the 95% confidence interval. The figures on the left consider financial awareness as the mediator, and those on the right use financial attitudes as the mediator. The treatment variable is *All Three Treatments*, defined as a dummy equal to 1 for an individual who was assigned to all three treatments of financial education, goal setting, and counseling. Sequential ignorability implies that ρ is equal to zero, so the ACME for $\rho = 0$ in the above figures corresponds to the ACME estimate in [Table 8](#).

In summary, our treatments influenced only two out of the three mediators. We find positive and statistically significant effects of the treatments on awareness and attitudes, coupled with precisely estimated average null effects on numeracy. Additionally, the null effects on numeracy seem to apply at the individual level, as we do not find evidence of heterogeneous effects. These patterns make us confident that we can dismiss numeracy as a mediator. So, throughout this paper, we concentrate on the potential mediating roles of awareness and attitudes, but not numeracy.

The above approach of studying the effects of the treatment T on the mediator M has several advantages and disadvantages. The main advantage is that it is straightforward to estimate and understand. It also requires the usual OLS identification assumption that T is orthogonal to the error term in Eq. (2); this is satisfied in an experimental context like ours with randomized treatments, allowing us to identify the causal effect of T on M . Unlike the mediation framework in Imai et al. (2011) and Acharya et al. (2016), we would not need the strong assumption necessary for causal mediation analysis, namely, that M is as-if random conditional on treatment and controls.

Still, in isolation, exploring whether and how the treatment T affects the mediator M provides limited information about mechanisms. If we had restricted our analysis to only this method, we would have been able to eliminate financial numeracy as a channel. Yet, finding significant positive treatment effects on awareness and attitudes does not say anything about their relationship to the outcome Y : if we had examined the relationship between T and M alone, we would not have been able to demonstrate that awareness and attitudes are indeed mechanisms for the effects of T on Y . Likewise, we would not have been able to measure the magnitudes of their potential mediating effects. Testing for the impacts of T on M is therefore useful primarily for excluding particular mediators, while mediation analysis using the methods in Imai et al. (2011) and Acharya et al. (2016) allows for a deeper study of those mechanisms that cannot be ruled out.

6.2. Regressing outcomes on the treatment conditional on the mediator

Apart from analyzing the effects of T on M , existing studies have also used regressions of Y on T conditional on M as evidence for mechanisms (e.g., Fearon and Laitin, 2003; Alesina, Nunn, and Giuliano, 2013). This approach for investigating mediation effects is very common in the social sciences. For instance, Acharya et al. (2016) write that of all papers published in three of the top journals in political science from 2010 to 2015, about two-thirds regress Y on T and M . Of these papers, around one-quarter do so to explicitly test or adjudicate between potential mechanisms.

The central regression in this approach is Eq. (3), which we reproduce below:

$$Y_i = \alpha_3 + \beta_3 T_i + \gamma_3 M_i + \xi_3^T X_i + \epsilon_{3i}.$$

The idea then is to examine whether the effect of the treatment goes away (or becomes smaller and less significant) after including the mediator in the regression. If so, this is taken as evidence that the variable M is indeed a channel for the treatment. How does this line of reasoning compare to the mediation methods in Imai et al. (2011) and Acharya et al. (2016)?

Estimating Eq. (3) is, as a matter of fact, similar to causal mediation analysis. This is because Eq. (3) is necessary to calculate the ACME, along with Eq. (2). Specifically, the product of γ_3 from Eq. (3) and β_2 from Eq. (2) is equivalent to the ACME, representing the indirect effect of T through M (i.e., $\bar{\delta}(t) = \gamma_3 * \beta_2$). The parameter β_3 of Eq. (3) on the other hand, which is the coefficient of interest in many previous studies that estimate this regression, captures the Average Direct Effect (ADE), representing all other mechanisms through which T impacts Y (i.e., $\bar{\zeta}(t) \equiv E[Y_i(1, M_i(t)) - Y_i(0, M_i(t))] = \beta_3$). If we find a precisely estimated null coefficient β_3 , this implies that the ADE is zero: the average treatment effect is due entirely to M . In this way, focusing on the effect of T on Y after controlling for M can provide evidence for mediating effects.

Nevertheless, it is important to note that to identify the coefficients in Eq. (3), we need the SI assumption (Imai, Keele, and Yamamoto, 2010). Since M is a post-treatment variable—meaning that it is itself an outcome of the treatment—adding M in Eq. (3) can result in selection bias in β_3 . If T impacts M , then comparing outcomes across treatment and control groups while holding M constant is no longer an apples-to-apples comparison, even in an experiment where T is randomized (Angrist and Pischke, 2008). Moreover, if there are any variables (e.g., interest in finance) that correlate with both the mediator (e.g., financial awareness) and the outcome (e.g., having a bank account), then γ_3 in Eq. (3) is not identified either. Hence, for estimates of Eq. (3) to be valid, M must be as-if random conditional on T and X . This assumption corresponds to the second part of SI, which is necessary for the mediation methods discussed in Imai et al. (2011).

Given that Eq. (3) calls for the same identification assumption as causal mediation analysis, we argue that estimating Eq. (3) alone and concentrating on β_3 is not ideal for studying mechanisms. One reason is that it cannot give a complete picture of mediating pathways, particularly when the ADE is nonzero. As an example, consider the results of estimating Eq. (3) in our study, shown in Tables 5 to 8, Panel A. We find that β_3 is positive, and in almost all cases, statistically significant. Here, we cannot say that the treatment effect is due solely to the mediator; we can only conclude that there are mechanisms other than numeracy, awareness, and attitudes. In addition, we cannot assess the strength of our proposed mediators since to do so, we need to estimate Eq. (2).

Because of these limitations, our view is that there is little to gain in using Eq. (3) on its own rather than combining Eqs. (2) and (3) as part of causal mediation analysis. Both approaches require the assumption that M is ignorable conditional on treatment and controls. However, the former approach provides much less information than the latter. Further, we note that in many previous papers that estimate a specification similar to Eq. (3), the identifying assumption regarding the ignorability of M remains obscure. By providing an example of causal mediation analysis in our study, we hope that we can contribute to clarifying the necessary conditions to interpret the estimates of Eq. (3) as direct and indirect treatment effects.

6.3. Examining heterogeneity in the effects of the mediator by treatment and control groups

Heterogeneity analysis is another popular method that has been frequently employed to explore causal mechanisms. One example of an economics paper that has used this approach is [Ahamed and Mallick \(2019\)](#), published in the *Journal of Economic Behavior & Organization*, which finds that higher financial inclusion (T) contributes to a bank's financial stability (Y). The authors then use heterogeneous effects to study the mechanisms, arguing that by reaching out to a larger pool of customers, banks garner more retail deposits (M), resulting in greater financial stability.

The objective in this type of analysis is to investigate whether the effect of the mechanism M on the outcome Y varies between the treatment and control groups. In practice, this analysis can be implemented by augmenting [Eq. \(3\)](#) with T^*M and using the full sample to estimate

$$Y_i = \alpha_5 + \beta_5 T_i + \gamma_5 M_i + \kappa_5 T_i * M_i + \xi_5^T X_i + \epsilon_{5i}. \quad (5)$$

Equivalently, one can estimate the $M \rightarrow Y$ effect separately for the treatment and control subsamples.

The coefficient of interest in [Eq. \(5\)](#) is κ_5 , which measures the difference in the effect of M on Y in the treatment relative to the control group. In the economics literature, researchers often use κ_5 (particularly its sign and whether it is different from zero) to test for causal mechanisms. To see how this might work, let us imagine that T positively impacts both M and Y . The logic goes that if M is a mediator, we should see a larger effect of M on Y among the treatment group, by virtue of the gains in M that this group experienced due to the intervention. Thus, a positive estimate of κ_5 (i.e., a larger average effect of M on Y in the treatment group) would be used to support the conclusion that M is an important treatment mechanism.

Are such claims based on heterogeneous effects of M on Y well-founded for mediation analysis? We now delve deeper into this issue. A statistically significant interaction κ_5 can certainly be informative about the existence of causal mechanisms. The rationale is similar to the regression of Y on T and M in [Eq. \(3\)](#). As explained earlier, the system of [Eqs. \(2\)](#) and [\(3\)](#) gives us the ACME $\bar{\delta}(t) = \gamma_3 * \beta_2$. If we replace [Eq. \(3\)](#) with [Eq. \(5\)](#) in this system, we obtain the ACME $\bar{\delta}(t) = \beta_2(\gamma_5 + t * \kappa_5)$. In other words, [Eq. \(5\)](#) sheds light on mediating effects because it can be used as an alternative to [Eq. \(3\)](#) when estimating the ACME. The main difference is that with [Eq. \(3\)](#), the ACME is posited to be the same for both treatment ($t = 1$) and control ($t = 0$), whereas with [Eq. \(5\)](#), we allow the ACME to vary by treatment status.

An estimate of κ_5 that is statistically different from zero implies that the mediating effect of M differs between treatment and control. For example, suppose that financial education (T) increases financial awareness (M) (i.e., $\beta_2 > 0$) and the effect of financial awareness (M) on opening a bank account (Y) is larger in the treatment than control group (i.e., $\kappa_5 > 0$). Correspondingly, the ACME is also larger in the treatment: $\bar{\delta}(1) > \bar{\delta}(0)$. Intuitively, this means that the proportion who open a bank account in the control group would have increased by a relatively small amount, even if they had the treated level of financial awareness, $M_i(1)$. But for the treatment group, the proportion who open a bank account would have been substantially lower if they had the untreated level of financial awareness, $M_i(0)$. Together, these patterns indicate that the awareness mechanism is present. Moreover, because $\kappa_5 > 0$, awareness is a stronger channel in the treatment than the control group.

Similar to the regression of Y on T and M in [Eq. \(3\)](#), a central issue to take into account when estimating [Eq. \(5\)](#) is the identification of the parameter κ_5 . And, as with [Eq. \(3\)](#), given that the mediator M (e.g., financial awareness) in [Eq. \(5\)](#) is a post-treatment variable, we require that M is ignorable conditional on T and X to identify the coefficients. Otherwise, [Eq. \(5\)](#) suffers from selection and intermediate confounder bias, as discussed in the previous subsection. Therefore, while it is possible to show evidence of causal mechanisms using the coefficient on the interaction of T and M , it is important to be cautious about the key underlying assumption: regression [Eq. \(5\)](#) is valid only under SI.

Assuming that SI is satisfied, we present the results of estimating [Eq. \(5\)](#) in our study in [Appendix Table A5](#). Our purpose here is to demonstrate how analyzing the heterogeneous effects of M by treatment status would be implemented in our context. To facilitate comparison, [Appendix Table A5](#) shows the interaction results for the same combination of treatment and outcomes that we consider in our causal mediation analysis in [Tables 5](#) to [8](#). For brevity, we report only the interaction coefficient κ_5 of [Eq. \(5\)](#), and we do not consider financial numeracy as a mediator M because it was not impacted by our financial education treatments.

For the most part, the interaction effects we find in [Appendix Table A5](#) are not statistically different from zero: we fail to reject the null hypothesis that the ACME is equal between treatment and control groups. However, these null effects neither support nor contradict the existence of potential mechanisms. They suggest that the mediating effects of awareness and attitudes is the same for both treatment and control, but we do not know whether these effects (i.e., the ACME) are zero or non-zero. On a separate but related point, the lack of interaction effects between T and M also help us choose between [Eq. \(5\)](#) or [Eq. \(3\)](#) in our system of equations for estimating the ACME. Our results in [Appendix Table A5](#) indicate that [Eq. \(3\)](#)—which assumes that the ACME is the same across treatment and control—appears to be a more appropriate specification than [Eq. \(5\)](#), since we do not find any indication that the ACME differs by treatment status.

Throughout the above discussion, we have examined the case with interactions of T and a *post-treatment* variable M . Aside from this approach, researchers have also used interactions of T with a *pre-treatment* (i.e., baseline) variable W to examine causal channels. A recent example is [Dalton et al. \(2019\)](#), who conduct a field experiment to study how e-payments (T) promote the financial inclusion of small- and medium-sized enterprises (SMEs). The authors find that e-payments allow SMEs to better access loans (Y). They hypothesize that this is because e-payments foster transparency in financial records (M), thereby allowing lenders to observe the creditworthiness and track the business transactions of SMEs. To corroborate

this treatment pathway, the authors investigate heterogeneity in the effect of e-payments by baseline establishment size (W), measured as the number of employees. The idea is that smaller firms are more likely to have opaque financial records and to lack hard information to attract external financing. Therefore, if their hypothesis on the mechanism is correct, the effects of e-payments should be more pronounced for smaller firms.

We note that regressing outcomes on T , the baseline variable W , and the interaction $T*W$ does not necessarily provide evidence for causal mediation. As [Baron and Kenny \(1986\)](#) write, there is a distinction between a mediator and a moderator. A *moderator* affects the strength and/or direction of the relationship between T and Y . A *mediator*, on the other hand, is an intervening variable or mechanism through which T influences Y . ATE heterogeneity obtained through interacting T and W suggests that W moderates the relationship between T and Y . Yet, this does not directly imply that W mediates the treatment. For heterogeneity in ATE by baseline characteristics to provide information on mediation, we need an additional assumption—that the magnitude of the ADE does not depend on W ([Imai et al., 2011](#)). Because the ATE equals the ADE plus the ACME, if this assumption holds, then heterogeneity in the ATE is driven entirely by heterogeneity in the ACME. A statistically significant coefficient on $T*W$ would thus imply that the ACME varies by treatment, and in turn, that the ACME for the mediator represented by W does exist.

In summary, regressions with interactions of the treatment T and the mediator M can inform us about the presence of mediating effects. However, this approach calls for the same SI assumption that we impose to estimate the ACME. Moreover, this approach has important drawbacks: it can only tell us whether mediating effects are present, but not how large the effects are. Consequently, causal mediation analysis is still more informative than examining the heterogenous effects of M . Both methods require SI, but causal mediation analysis enables us to estimate the size of the mediating effect as well as to build confidence intervals for these effect sizes.

As with interactions of T and M , interactions of T and a baseline variable W can also provide evidence for causal mechanisms. The upside of this approach is that it allows us to study mechanisms even if we do not or are unable to measure the post-treatment variable M . Nevertheless, the downside is that to interpret the interaction of T and W in terms of causal mechanisms, we need to impose the additional assumption that there is no heterogeneity in the ADE by W ([Imai et al., 2011](#)). Further, similar to the interaction of T and M , the interaction of T and W can only tell us whether the ATE and the ACME vary by W , but not the sign, magnitude, nor confidence interval of the mediating channels. Hence, this approach is unable to provide a full picture of causal mediation.

6.4. Instrumental variables

Many previous studies have studied causal mechanisms by employing an IV framework. Here, the basic idea is to use the treatment T —which is often experimentally randomly assigned—as an instrument in estimating the effects of the mediator M on outcomes Y . An example of an economics paper that has used this method is [Sayinzoga et al. \(2016\)](#), who conduct a field experiment in Rwanda to investigate the effects of financial education on financial knowledge and behavior. The authors study the causal chain that financial literacy training (T) builds financial knowledge (M), resulting in improved financial decision-making and economic well-being (Y). To probe into this theory of change, the authors estimate an IV regression. Specifically, they employ the random assignment of the financial literacy training as an instrument for financial knowledge.

The IV estimate can be obtained from Two-Stage Least Squares. In the first stage, the randomized treatment T is used to explain M ; this regression corresponds to [Eq. \(2\)](#). In the second stage, we use the predicted level of the mediator from the first stage, denoted \hat{M} , to explain outcomes:

$$Y_i = \alpha_6 + \beta_6 \hat{M}_i + \xi_6^T X_i + \epsilon_{6i}. \quad (6)$$

The coefficient of interest is then β_6 .

IV requires two key assumptions. The first is relevance, meaning that the instrument T must have an effect on M in the first stage. In our context, the first stage corresponds to [Table 4](#), which reports the impact of the different financial education treatments T on our proposed mediators M (i.e., numeracy, awareness, and attitudes). [Table 4](#) shows that the effects of T on M are statistically significant at the 1 percent level for awareness and attitudes, but not for numeracy. Thus, the relevance condition is satisfied only when awareness and attitudes are the proposed mechanisms.

The second IV assumption is the exclusion restriction: the instrument T is uncorrelated with all other determinants of Y . This means that the only reason for the relationship between Y and the instrument T is the first stage equation ([Angrist and Pischke, 2008](#)). The exclusion restriction requires that T must be as good as randomly assigned, so that it is uncorrelated with potential outcomes, conditional on covariates. Our experiment supports this condition. Nevertheless, randomization of T is not sufficient to satisfy the exclusion restriction, as we also need that T impacts Y only through M . As we explain below, this restriction makes the IV technique unsuitable for analyzing causal mediation effects, given that it does not allow for any direct effects of T on Y .

[Appendix Table A6](#) presents IV regressions using the data from our study. The instrument is the treatment T , an indicator for whether the respondent was randomly assigned to a particular treatment group. For comparability with our ACME estimates, we focus on the same combination of treatment, mediators, and outcomes as in our main mediation analysis, and we implement the IV separately for four types of treatments T : (i) financial education only (Panel A); (ii) financial education with goal setting (Panel B); (iii) financial education with counseling (Panel C); and (iv) all three of financial education, goal setting, and counseling (Panel D). In all panels, we report IV estimates considering the mediator to be either financial aware-




Fig. 7. The Front-Door Criterion

Notes: This diagram illustrates the front-door criterion, following Pearl (2000). Here, U is an unobservable confounder that causes both T and Y , and all of the effect of T on Y is mediated through M , which is not influenced by U .

ness (odd-numbered columns) or financial attitudes (even-numbered columns). We do not study the IV effects of financial numeracy because as mentioned above, none of the treatments are relevant instruments for this variable.

The IV results in Appendix Table A6 show that for any given outcome Y and treatment T , the coefficients on financial awareness and financial attitudes are all very similar. In other words, when we instrument either financial awareness or financial attitudes with the randomly assigned treatment, the magnitude for the effect of financial awareness on outcomes is comparable to that for the effect of financial attitudes. If we view these IV estimates as mediating effects, we would have concluded that awareness and attitudes are equally important mechanisms for financial education. This result stands in contrast to our mediation analysis using the framework in Imai et al. (2011) and Acharya et al. (2016), where we find that attitudes play a more prominent role in linking financial education with complex financial behaviors, such as opening a bank savings account.

To understand why IV and mediation analysis reach different conclusions, it is important to clarify the purpose of IV and what it estimates. Generally speaking, IV is used when we are interested in the causal effect of an endogenous variable. In our context, the IV allows us to overcome the potential endogeneity of the mediator M , so that we can consistently estimate the causal effect of M on Y . Importantly, as Imbens and Angrist (1994) show, the IV yields a Local Average Treatment Effect (LATE). The LATE corresponds to the ATE of M among individuals for whom T influences M , known as compliers. For example, when we instrument financial awareness with the financial education only treatment, we estimate the average effect of financial awareness among a particular subset of respondents: those whose level of awareness can be changed by financial education.

But the LATE obtained through IV is not ideal for causal mediation analysis because it differs fundamentally from mediation effects (Keele, 2015). The IV is used to measure the causal effect of the endogenous variable M , and this is not the same as estimating the mediating effect. Moreover, with causal mediation analysis, the goal is to decompose the ATE into an indirect effect, which is the effect of T on Y due to the mechanism M , and the direct effect, representing all other channels. The IV design does not allow for this decomposition because the exclusion restriction requires that T impacts Y only through M . In other words, the validity of the IV rests on M being the sole mechanism of impact. The IV method therefore assumes, *ex ante*, that direct effects are zero and that there are no mechanisms other than M . Assuming no direct effects is not well-suited for a study like ours because the treatment likely impacts outcomes through mechanisms other than our proposed mediator. Thus, using IV is not ideal for causal mediation analysis.

6.5. The front-door criterion

Pearl's front-door criterion (e.g., Pearl, 1995; Pearl, 2000)—shown in Fig. 7—is also related in some ways to causal mediation analysis. However, there is an important conceptual difference between their objectives. With the front-door criterion, the goal is to use a separate variable (e.g., the mediator M) to identify and estimate the average causal effect of the treatment T , in the presence of an unobserved confounder U that influences T and the outcome Y (i.e., T is endogenous). In contrast, with causal mediation analysis, we assume that the ATE is already identified (i.e., T is exogenous); the goal is to decompose the ATE into the direct effect of T on Y and the indirect effect of T via M .

Nevertheless, there are similarities between the front-door criterion and causal mediation analysis. Consider again the diagram in Fig. 7. In this figure, we are interested in the effect of an *endogenous* treatment T on Y , and we assume that all effects of T flow through M , so T has no direct effect on Y . Further, the unobserved confounder U does not influence M . Then, we obtain the causal effect of T on Y through a front-door adjustment: we multiply the effect of T on M with the effect of M on Y . By taking this product of coefficients, we circumvent the confounding by U , since U does not confound the $T \rightarrow M$ and $M \rightarrow Y$ relationships.

Notably, this product of coefficients is analogous to what we use to estimate the ACME, following the method described in Imai et al. (2011). In causal mediation analysis, the ATE is equivalent to the sum of the ACME and the ADE. With Pearl's front-door criterion, we assume that the ADE is zero, and as a result, the ATE is equal to the ACME. In other words, using Pearl's front-door criterion to estimate the ATE will give us the same estimate as the causal mediating effect of M (i.e., the ACME), under the assumption that the ADE is zero.

Although the ACME equals the ATE in this case, using the front-door criterion differs fundamentally from causal mediation analysis. The reasoning runs parallel to using a randomized treatment T as an instrument for the mediator M , which

we discuss in the previous subsection. Such an IV framework requires the exclusion restriction (i.e., T has no direct effects on Y), so we assume *ex ante* that the ADE is zero. In a similar way, Pearl's front-door criterion assumes *a priori* that the ADE is zero, and doing so does not allow for decomposing the ATE into direct and indirect effects, which is the goal of causal mediation analysis. As mentioned in the case of IV, assuming no direct effects *ex ante* may not be appropriate in many settings like ours, where there are likely to be direct effects which are themselves of particular interest.

6.6. Randomizing the mediator

As explained in the empirical methods section, causal mediation analysis is ambitious because it entails counterfactual outcomes that can never be observed, even in experimental designs. For this reason, we must rely on a selection-on-observables argument—as embodied in the second part of the SI assumption—to estimate causal mediating effects.

As an alternative to relying on selection-on-observables, one might perhaps envision an empirical design where the mediator M is randomly assigned. Unfortunately, randomizing M is not possible in many settings (Keele, 2015), including ours. In our study, the financial knowledge mediators we examine are three-fold: numeracy (e.g., interest calculations), awareness (e.g., basic information about financial products), and attitudes (e.g., views about the benefits of financial services)—all of which are difficult to directly manipulate as they pertain to individual perception and capabilities. Though our study shows that these mediators can be influenced by financial education, it is hard to imagine practicable approaches to engineer these mediators to take on specific levels. Real-world situations that allow experimenters to randomly assign a person's financial numeracy, awareness, and attitudes would likely be farfetched.

Even if it were feasible to randomly assign M in our context, doing so does not guarantee that we can estimate causal mediation effects (Imai et al., 2011). Experiments where M is randomized may be instructive and interesting in their own right, but they cannot provide empirical estimates of causal mediation effects without additional assumptions. The key idea is that the causal mediation effect captures the effect of T on Y that is transmitted through M . When we randomize M , we are externally (and artificially) assigning M to take on a certain value, and this value may not be the same as what would naturally arise because of the treatment. Consequently, changes in Y that result from randomly assigning M are not necessarily equivalent to the treatment-induced changes in Y due to M .

To illustrate this point further, consider an experiment among seafarers during the 18th century, which showed that eating limes reduces scurvy because of the vitamin C content. In this study, discussed in Gerber and Green (2012), T is an indicator for receiving a daily dietary supplement of one lime, M represents vitamin C, and Y is an indicator for scurvy. Now, suppose that we hold T constant and randomize M . For example, among those who do not receive limes (i.e., $T = 0$), a randomly selected subset receives a pill containing a lime's worth of vitamin C. Comparing the incidence of scurvy between those who did and did not receive the vitamin C pill, will we get an estimate of the ACME?

Note that the ACME of vitamin C, fixing T at zero, is $\bar{\delta}(0) = E[Y_i(0, M_i(1))] - E[Y_i(0, M_i(0))]$. While the comparison proposed above resembles $\bar{\delta}(0)$, it does not, in fact, give us an empirical estimate of $\bar{\delta}(0)$. The difference is that $M_i(1)$ is the value of the mediator when person i receives the treatment of limes. Because of differences between limes and pills in how they are eaten and how nutrients are absorbed, a person's vitamin C status from consuming a lime is not necessarily the same as when taking a vitamin C pill. As a result, a lime's worth of vitamin C in pill form may not have the same indirect or mediating effect on scurvy as a lime itself. In addition, taking vitamin C may bring about other dietary consequences that, in turn, affect scurvy. In this state of the world, the outcomes we observe when randomizing vitamin C pills would be different from the potential outcomes that would have occurred when M takes its "natural" value under the lime treatment.

For such an experiment to provide estimates of the ACME, we need a new assumption: we must assume that the effect of M on Y is the same irrespective of how changes in M are induced—whether through vitamin C pills, limes, oranges, or other means (Gerber and Green, 2012). If this condition is satisfied, then the average of potential outcomes such as $E[Y_i(0, M_i(1))]$ and $E[Y_i(0, M_i(0))]$ can be approximated by the averages of observed outcomes in our experiment with randomly assigned M . Randomizing M (assuming it was possible to begin with) would therefore have allowed us to get around selection-on-observables arguments. Nevertheless, this benefit comes at a cost, as we would not have been able to obtain causal mediating effects without imposing an additional assumption.

Lastly, we note that in the existing methodological literature on causal mediation analysis, several papers have studied how the ACME can be identified without assuming the mediator is ignorable (e.g., Albert, 2008; Jo, 2008; Sobel, 2008; Imai, Tingley, and Yamamoto, 2013). But similar to the case of randomizing M , additional assumptions to identify the ACME will be necessary. For example, Imai, Tingley, and Yamamoto (2013) discuss how encouragement designs (where randomly selected subjects are encouraged to take on a high or low value of the mediator) can be used to study mediation effects, under weaker assumptions than the ignorability of M . However, this encouragement design again requires an extra assumption (i.e., outcomes would have been the same if the mediator level was spontaneously chosen by the subjects themselves).

7. Robustness

In this section, we investigate the robustness of our results to two issues related to estimation. First, statistical power could be an important concern, given that we are looking at multiple treatment combinations with a sample size of around 1,000 respondents. Since statistical power tends to be lower with a more diverse sample, in Table 2, we assess the amount of

baseline heterogeneity among our respondents using the Coefficient of Variation (CV)—the ratio of the standard deviation to the mean. We find CVs well below one for almost all continuous variables. Further, we estimate *ex post* Minimum Detectable Effect Sizes (MDEs) for the ACMEs. We obtain MDEs for the ACME of numeracy that are all close to zero. However, for awareness and attitudes, we find relatively large MDEs, so we may be underpowered in estimating the mediating effects in these cases.¹⁹

Second, because we conduct many different hypothesis tests in this paper, we examine robustness to multiple inference. We find that most of our results survive multiple inference correction.²⁰

8. Discussion and conclusion

This paper applies causal mediation analysis to a field experiment on financial education in India, to understand the causal pathways from financial knowledge to financial behavior. As causal mediation analysis is a challenging exercise (Green, Ha, and Bullock, 2010), there are important caveats to our results. The ACME estimates depend on the strong SI assumption that the mediators are as-if random, conditional on the treatment and pre-treatment variables; our sensitivity analysis shows that the estimates are sensitive to even small violations of this assumption. Moreover, we acknowledge that the validity of SI cannot be empirically tested: we can never be certain whether we have controlled for all pre-treatment confounders or whether post-treatment confounders are present. In this sense, causal mediation analysis is similar to observational studies (Keele, 2015), and we recognize that there may be uncertainty in the validity of our estimates of mediating effects. Apart from SI, an additional caveat is statistical power. We may be underpowered in estimating the ACMEs of awareness and attitudes, and this is an important consideration in interpreting our results.

All this being said, we do highlight that one of our objectives in this paper is to demonstrate how causal mediation analysis can be carried out by social science researchers. Drawing on methodological studies from fields such as political science and biostatistics, we aim to provide an empirical example of the ingredients and assumptions that are required for such an analysis. In particular, we do so in experimental study where treatments were randomized so that the identifying assumption regarding the ignorability of the treatment T is satisfied. But critically, we also have a set-up where mediators were explicitly and carefully measured. Ultimately then, having shown a practical application of causal mediation analysis, we leave it up to readers to decide for themselves whether they believe the identifying assumptions are plausible in our study.

Despite the necessary caveats, our paper speaks to the ongoing research and policy debate on the merits of financial education. Our analysis suggests that for financial education to effectively impact financial outcomes through knowledge gains, two links in the causal chain must be operative.

The first link is that financial education must increase financial knowledge. Whereas financial knowledge is commonly measured using specific survey questions on interest rates, inflation, and diversification, our study points to a broader notion. We highlight financial literacy not as a singular concept, but rather one that involves multiple dimensions: numeracy, awareness, and attitudes. We find that financial education affects financial knowledge through avenues other than numeracy, by improving understanding of basic financial concepts and beliefs in the value of financial planning. Notably, these seemingly small advances appear to be important mediators that encourage the poor to improve their finances. Hence, our results call for measuring financial knowledge using a combination of awareness and attitudinal assessments, particularly in development settings where financial education beneficiaries may have low education. Our study then provides a field-tested list of such questions to complement existing measures.

The second step in the causal chain is that financial literacy in turn must translate into improved financial behavior. Our causal mediation analysis shows that numeracy does not mediate the treatment effects on any financial behaviors that we study. However, we find suggestive evidence that awareness and attitudes may be important channels for financial education. Specifically, we find that for *simple* financial actions such as attempting to write a household budget, awareness and attitudes seem to be equally critical, but for more *complex* financial actions, such as opening a savings account, attitudes may play a more meaningful mediating role. Responses from the open-ended questions in our endline are likewise consistent with our finding that the main hindrance subjects faced in opening a bank account is not awareness, but rather attitudes. For instance, respondents stated that one of the reasons they are trying to save more money now is because they realized the value and importance of savings. Thus, the financial education treatments likely operate by improving perceptions about personal finance, consistent with our finding that the ACME for attitudes is much larger than that of awareness for savings outcomes.

The substantial role of attitudes for simple financial actions and awareness for complex ones indicate that higher and cumulative levels of learning may be necessary to encourage sophisticated financial behaviors. This may be the case because building financial awareness is a much less cumulative exercise than building financial attitudes. Increasing financial awareness may be achieved by supplying respondents with information (e.g., about bank account opening requirements) in

¹⁹ Appendix D discusses the robustness of our results to statistical power concerns in more detail.

²⁰ See results in Appendix Table 9. This table presents the adjusted and non-adjusted p-values for the ATE, ACME, and ADE. The adjusted p-values are the Benjamini and Hochberg (1995) q-values as described in Anderson (2008).

a one-off event. In contrast, improving financial attitudes involves continuously internalizing that knowledge into one's everyday life. Our ACME estimates appear to support this view: we find a cumulative mediation channel for financial attitudes in that a more intense treatment results in higher mediation effects. When financial education is combined with either counseling or goal setting, the ACME of attitudes on savings outcomes is similar at 1.9 percentage points each. In contrast, when providing financial education with both counseling and goal setting, the ACME of attitudes on savings increases to 3.3 percentage points, a 70% change. Interestingly, this pattern indicates potential complementarities between counseling and goal setting, such that a more comprehensive treatment increases both the ATEs and mediating efficacy of financial attitudes.

More generally, our results reveal that broad pessimism on the value of financial literacy may be unwarranted. We show that financial education can be effective in increasing different facets of financial knowledge, and in turn, better financial attitudes and awareness appear to be critical pathways for improved financial behavior. If the objective of a financial education initiative is to encourage complex financial behaviors, our results indicate that placing more emphasis on changing beliefs about the benefits of financial products may be a productive approach. We believe these insights are key, particularly for designing successful financial literacy programs that deliver meaningful impacts. Likewise, understanding how the patterns of causal mediation change over time remains an important avenue for future research, especially if individuals' financial attitudes evolve further as they gain more experience with financial products.

Appendix A. Further details on financial numeracy, awareness, and attitudes

We developed the measures for numeracy, awareness, and attitudes iteratively, in close collaboration with our local Indian partners: the Center for Microfinance (a non-profit research institution) and Saath Microfinance (an organization

Table A1
Sequential g-Estimation: Financial Education Treatment.

	Budgeting		
	Financial Numeracy (1)	Financial Awareness (2)	Financial Attitudes (3)
Panel A. Controlling for Pre-Treatment Variables Only			
ATE-ACDE \approx ACME	-0.002 (0.004)	0.027** (0.013)	0.025** (0.011)
ACDE \approx ADE	0.140*** (0.042)	0.110** (0.044)	0.113*** (0.041)
ATE		0.138*** (0.038)	
% of ATE Mediated	-1	20	18
Panel B. Controlling for Pre-Treatment and Post-Treatment Variables			
ATE-ACDE \approx ACME	-0.003 (0.004)	0.028** (0.013)	0.026** (0.012)
ACDE \approx ADE	0.140*** (0.042)	0.109** (0.044)	0.112*** (0.041)
ATE		0.138*** (0.038)	
% of ATE Mediated	-2%	20%	19%

Notes: This table presents estimates of the Average Controlled Direct Effect (ACDE), obtained using sequential g-estimation as outlined in Acharya, Blackwell, and Sen (2016), for the Financial Education Only treatment. The ACDE is equivalent to the Average Direct Effect (ADE) under the no-interactions assumption (Robins, 2003). With this assumption, the Average Causal Mediating Effect (ACME) can be obtained by subtracting the ACDE from the Average Treatment Effect (ATE). The ATE given above is replicated from the coefficient estimate of the treatment variable from Table 3. The mediator variables considered are financial numeracy (column 1), financial awareness (column 2), and financial attitudes (column 3). In all regressions, the sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to the Financial Education Only treatment. In Panel A, the regressions control for only pre-treatment variables in the first step of sequential-g estimation. The pre-treatment variables are the monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters. In Panel B, the regressions control for both pre-treatment and post-treatment variables in the first step of sequential-g estimation. The post-treatment controls consist of the same set as the pre-treatment variables, with the exception of interest in financial matters. Further, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A2
Sequential g-Estimation: Financial Education and Goal Setting Treatment.

	Budgeting			Savings		
	Has tried making a budget in the last 6 months			Has a savings account		
	Financial Numeracy (1)	Financial Awareness (2)	Financial Attitudes (3)	Financial Numeracy (4)	Financial Awareness (5)	Financial Attitudes (6)
Panel A. Controlling for Pre-Treatment Variables Only						
ATE-ACDE \approx ACME	-0.000 (0.003)	0.020* (0.011)	0.036*** (0.011)	-0.001 (0.004)	0.001 (0.015)	0.019* (0.011)
ACDE \approx ADE	0.169*** (0.046)	0.149*** (0.048)	0.132*** (0.047)	0.082 (0.050)	0.080 (0.051)	0.061 (0.051)
ATE		0.169*** (0.045)			0.081* (0.045)	
% of ATE Mediated	0%	12%	21%	-1%	1%	23%
Panel B. Controlling for Pre-Treatment and Post-Treatment Variables						
ATE-ACDE \approx ACME	0.000 (0.003)	0.019 (0.012)	0.038*** (0.012)	-0.001 (0.005)	-0.004 (0.015)	0.018 (0.011)
ACDE \approx ADE	0.169*** (0.046)	0.150*** (0.049)	0.131*** (0.047)	0.082 (0.050)	0.085* (0.051)	0.063 (0.051)
ATE		0.169*** (0.045)			0.081* (0.045)	
% of ATE Mediated	0%	11%	22%	-1%	-5%	22%

Notes: This table presents estimates of the Average Controlled Direct Effect (ACDE), obtained using sequential g-estimation as outlined in Acharya, Blackwell, and Sen (2016), for the Financial Education and Goal Setting treatment. The ACDE is equivalent to the Average Direct Effect (ADE) under the no-interactions assumption (Robins, 2003). With this assumption, the Average Causal Mediating Effect (ACME) can be obtained by subtracting the ACDE from the Average Treatment Effect (ATE). The ATE given above is replicated from the coefficient estimate of the treatment variable from Table 3. The mediator variables considered are financial numeracy (columns 1 and 4), financial awareness (columns 2 and 5), and financial attitudes (columns 3 and 6). In all regressions, the sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to the Financial Education and Goal Setting treatment. In Panel A, the regressions control for only pre-treatment variables in the first step of sequential-g estimation. The pre-treatment variables are the monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters. In Panel B, the regressions control for both pre-treatment and post-treatment variables in the first step of sequential-g estimation. The post-treatment controls consist of the same set as the pre-treatment variables, with the exception of interest in financial matters. Further, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

providing microfinance services). These measures were designed to be closely related to the curriculum presented in the financial education program. The content of the education program, in turn, was based on materials that were used in previous research. In particular, we developed the curriculum based on standard materials produced by Freedom from Hunger, Microfinance Opportunities, and Citi Foundation that have been employed in other studies. We then iterated on and adapted these materials to our study context, urban India, together with our local partners. Importantly, we note that all materials in our study—including all questionnaires as well as the content of the financial education program—were carefully piloted prior to implementation.

Next, we examine to what extent financial knowledge correlates with baseline measures in our sample. To avoid Hawthorne-type effects (i.e., to avert priming respondents about the content of the financial education course), our end-line measures on financial numeracy, awareness, and attitudes were not collected at baseline. However, we do have some pre-treatment measures related to financial knowledge. More specifically, the baseline includes the following three questions: (1) If you borrowed Rs. 5,500 and were charged 12% interest per month, how much interest would you pay in the first month? (2) Suppose you had Rs. 100 in a savings account and the same amount saved at home, which of the two will yield returns at the end of one year? (3) Suppose your friend inherits Rs. 10,000 today and his brother inherits Rs. 10,000 three years from now. Who is richer because of the inheritance?

The proportion of correct answers to these three questions is the variable that corresponds to the financial knowledge score in Table 2. Column 5 of this table reports the correlation between this financial knowledge score and the different baseline variables. We find low correlation between the financial knowledge score and household background (e.g., household income), respondent demographics (e.g., gender and education), and other respondent characteristics (e.g., risk aversion)—with correlations ranging from -0.05 to 0.27. Hence, financial knowledge appears to be capturing an aspect that is different from socio-economic status and mathematical ability.

Table A3
Sequential g-Estimation: Financial Education and Counseling Treatment.

	Budgeting			Savings			Borrowing			Insurance		
	Has tried making a budget in the last 6 months			Has a savings account			Loan purpose: Business, education, or purchase of durable goods			Bought life insurance in the last 6 months		
	Financial Numeracy (1)	Financial Awareness (2)	Financial Attitudes (3)	Financial Numeracy (4)	Financial Awareness (5)	Financial Attitudes (6)	Financial Numeracy (7)	Financial Awareness (8)	Financial Attitudes (9)	Financial Numeracy (10)	Financial Awareness (11)	Financial Attitudes (12)
Panel A. Controlling for Pre-Treatment Variables Only												
ATE-ACDE \approx ACME	0.001 (0.003)	0.033*** (0.012)	0.040*** (0.014)	0.001 (0.003)	0.005 (0.013)	0.019* (0.011)	0.000 (0.012)	0.021 (0.024)	0.003 (0.022)	0.000 (0.001)	-0.003 (0.007)	0.003 (0.004)
ACDE \approx ADE	0.398*** (0.051)	0.367*** (0.052)	0.360*** (0.051)	0.167*** (0.051)	0.163*** (0.052)	0.149*** (0.051)	0.175* (0.096)	0.155 (0.098)	0.172* (0.099)	0.045* (0.027)	0.048* (0.029)	0.042 (0.028)
ATE		0.399*** (0.049)			0.168*** (0.043)			0.176** (0.077)			0.045* (0.025)	
% of ATE Mediated	0%	8%	10%	1%	3%	11%	0%	12%	2%	0%	-7%	7%
Panel B. Controlling for Pre-Treatment and Post-Treatment Variables												
ATE-ACDE \approx ACME	0.001 (0.003)	0.032*** (0.012)	0.038*** (0.014)	0.001 (0.004)	0.001 (0.013)	0.014 (0.011)	0.002 (0.013)	0.015 (0.023)	0.010 (0.023)	0.000 (0.001)	-0.003 (0.007)	0.002 (0.004)
ACDE \approx ADE	0.398*** (0.051)	0.367*** (0.052)	0.361*** (0.051)	0.167*** (0.051)	0.167*** (0.052)	0.154*** (0.051)	0.173* (0.097)	0.161 (0.098)	0.166* (0.098)	0.045* (0.027)	0.048* (0.029)	0.044 (0.028)
ATE		0.399*** (0.049)			0.168*** (0.043)			0.176** (0.077)			0.045* (0.025)	
% of ATE Mediated	0%	8%	10%	1%	1%	8%	1%	9%	6%	0%	-7%	4%

Notes: This table presents estimates of the Average Controlled Direct Effect (ACDE), obtained using sequential g-estimation as outlined in Acharya, Blackwell, and Sen (2016), for the Financial Education and Counseling treatment. The ACDE is equivalent to the Average Direct Effect (ADE) under the no-interactions assumption (Robins, 2003). With this assumption, the Average Causal Mediating Effect (ACME) can be obtained by subtracting the ACDE from the Average Treatment Effect (ATE). The ATE given above is replicated from the coefficient estimate of the treatment variable from Table 3. The mediator variables considered are financial numeracy (columns 1, 4, 7, and 10), financial awareness (columns 2, 5, 8, and 11), and financial attitudes (columns 3, 6, 9, and 12). In all regressions, the sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to the Financial Education and Counseling treatment. In Panel A, the regressions control for only pre-treatment variables in the first step of sequential-g estimation. The pre-treatment variables are the monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters. In Panel B, the regressions control for both pre-treatment and post-treatment variables in the first step of sequential-g estimation. The post-treatment controls consist of the same set as the pre-treatment variables, with the exception of interest in financial matters. Further, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A4
Sequential g-Estimation: All Three Treatments.

	Budgeting			Savings			Insurance		
	Has tried making a budget in the last 6 months			Has a savings account			Bought life insurance in the last 6 months		
	Financial Numeracy (1)	Financial Awareness (2)	Financial Attitudes (3)	Financial Numeracy (4)	Financial Awareness (5)	Financial Attitudes (6)	Financial Numeracy (7)	Financial Awareness (8)	Financial Attitudes (9)
Panel A. Controlling for Pre-Treatment Variables Only									
ATE-ACDE \approx ACME	0.001 (0.004)	0.026* (0.015)	0.058*** (0.014)	0.000 (0.003)	0.014 (0.017)	0.033** (0.014)	0.000 (0.001)	-0.002 (0.008)	0.007 (0.007)
ACDE \approx ADE	0.474*** (0.046)	0.449*** (0.050)	0.417*** (0.048)	0.106** (0.050)	0.093* (0.050)	0.074 (0.052)	0.042 (0.027)	0.044 (0.029)	0.035 (0.028)
ATE		0.475*** (0.041)			0.107** (0.045)			0.042* (0.021)	
% of ATE Mediated	0%	5%	12%	0%	13%	31%	0%	-5%	17%
Panel B. Controlling for Pre-Treatment and Post-Treatment Variables									
ATE-ACDE \approx ACME	0.000 (0.004)	0.028* (0.016)	0.058*** (0.014)	0.000 (0.004)	0.015 (0.017)	0.029** (0.014)	0.000 (0.001)	-0.003 (0.008)	0.005 (0.008)
ACDE \approx ADE	0.474*** (0.046)	0.447*** (0.050)	0.417*** (0.048)	0.106** (0.050)	0.092* (0.051)	0.077 (0.052)	0.042 (0.027)	0.045 (0.030)	0.037 (0.028)
ATE		0.475*** (0.041)			0.107** (0.045)			0.042* (0.021)	
% of ATE Mediated	0%	6%	12%	0%	14%	27%	0%	-%	12%

Notes: This table presents estimates of the Average Controlled Direct Effect (ACDE), obtained using sequential g-estimation as outlined in Acharya, Blackwell, and Sen (2016), for All Three Treatments of financial education, goal setting, and counseling. The ACDE is equivalent to the Average Direct Effect (ADE) under the no-interactions assumption (Robins, 2003). With this assumption, the Average Causal Mediating Effect (ACME) can be obtained by subtracting the ACDE from the Average Treatment Effect (ATE). The ATE given above is replicated from the coefficient estimate of the treatment variable from Table 3. The mediator variables considered are financial numeracy (columns 1, 4, and 7), financial awareness (columns 2, 5, and 8), and financial attitudes (columns 3, 6, and 9). In all regressions, the sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to All Three Treatments. In Panel A, the regressions control for only pre-treatment variables in the first step of sequential-g estimation. The pre-treatment variables are the monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters. In Panel B, the regressions control for both pre-treatment and post-treatment variables in the first step of sequential-g estimation. The post-treatment controls consist of the same set as the pre-treatment variables, with the exception of interest in financial matters. Further, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A5
Interaction of Mediator and Treatment.

	Budgeting		Savings		Borrowing		Insurance	
	Has tried making a budget in the last 6 months		Has a savings account		Loan purpose: Business, education, or purchase of durable goods		Bought life insurance in the last 6 months	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Treatment: Financial Education Only								
Treatment * Endline Awareness Score	0.071 (0.139)							
Treatment * Endline Attitudes Score	0.310* (0.184)							
Number of Observations	487		487					
Panel B. Treatment: Financial Education and Goal Setting								
Treatment * Endline Awareness Score	0.194 (0.121)		0.118 (0.209)					
Treatment * Endline Attitudes Score	-0.006 (0.207)		0.386 (0.238)					
Number of Observations	465		465		465		465	
Panel C. Treatment: Financial Education and Counseling								
Treatment * Endline Awareness Score	0.370** (0.175)		0.120 (0.189)		-0.007 (0.463)		0.086 (0.075)	
Treatment * Endline Attitudes Score	0.194 (0.219)		0.159 (0.192)		-0.638 (0.429)		0.035 (0.096)	
Number of Observations	468		468		227		227	
Panel D. Treatment: All Three Treatments								
Treatment * Endline Awareness Score	0.303 (0.195)		0.256 (0.155)				-0.049 (0.110)	
Treatment * Endline Attitudes Score	0.051 (0.241)		0.238 (0.194)				-0.036 (0.105)	
Number of Observations	476		476		476		476	

Notes: This table presents regressions of the outcomes on the treatment, mediator, and their interaction. The mediators considered are awareness (odd columns) or attitudes (even columns). Each panel refers to a different treatment. Financial Education Only is a dummy equal to 1 for an individual who was assigned to the financial education program but not financial counseling nor goal setting. Financial Education and Goal Setting is a dummy equal to 1 for an individual who was assigned to the financial education and goal setting treatments, but not the financial counseling treatment. Financial Education and Counseling is a dummy equal to 1 for an individual who was assigned to the financial education and financial counseling treatments, but not the goal setting treatment. All Three Treatments is a dummy equal to 1 for an individual who was assigned to all three (i.e., financial education, counseling, and goal setting). In each panel, the regression sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to a particular treatment. For example, in Panel A, the regression sample consists of those in the Financial Education Only treatment and those assigned to control. All regressions include baseline control variables (i.e., monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters). In addition, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). Standard errors are clustered at the wave-class level. *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Appendix B. Exploring the direction of bias and selection on unobservables in ACME

Figs. 3 to 6 not only quantify the degree of sensitivity, but likewise tell us the direction of the potential bias in our estimate of the ACME if we know the correlation ρ . Indeed, across all figures, we see that if the true ρ is less than 0, we would underestimate the ACME, but if the true ρ is greater than 0, then we would overestimate the ACME. Further, recall that the ACME is the product of β_2 from Eq. (2) and γ_3 from Eq. (3). Since the former parameter is identified with randomization and we find that the treatments had a positive impact on the mediators (i.e., $\beta_2 > 0$), the sign of the bias in the ACME will be determined by the sign of the bias in γ_3 .

To understand the direction of the bias in the $\hat{\gamma}_3$, we note that as the sample size increases, the estimator $\hat{\gamma}_3$ has the limit (Gerber and Green, 2012)

$$\hat{\gamma}_{3_{N \rightarrow \infty}} = \gamma_3 + \frac{\text{cov}(\epsilon_{2i}, \epsilon_{3i})}{\text{var}(\epsilon_{2i})}.$$

This indicates that the sign of the bias in γ_3 is governed by $\text{cov}(\epsilon_{2i}, \epsilon_{3i})$, the covariance of the error terms in Eqs. (2) and (3). Note that the population regression error ϵ_{2i} in Eq. (2) captures factors impacting the financial literacy mediator other than the financial education treatment, while ϵ_{3i} in Eq. (3) represents all other variables impacting financial outcomes that are not contained in the financial literacy mediator, the financial education treatment, or the controls. This means that intuitively, the covariance $\text{cov}(\epsilon_{2i}, \epsilon_{3i})$ could be nonzero because of unobserved or omitted pre-treatment variables U that influence both the mediator M and the outcome Y . One example that was mentioned in the empirical methods section is

Table A6
IV Regressions, Instrumenting Mediator with Treatment.

	Budgeting		Savings		Borrowing		Insurance	
	Has tried making a budget in the last 6 months		Has a savings account		Loan purpose: Business, education, or purchase of durable goods		Bought life insurance in the last 6 months	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Instrument: Financial Education Only Treatment								
Endline Awareness Score	1.036*** (0.339)							
Endline Attitudes Score	1.949*** (0.665)							
Number of Observations	478	478						
Panel B. Instrument: Financial Education and Goal Setting Treatment								
Endline Awareness Score	1.308*** (0.452)			0.626* (0.371)				
Endline Attitudes Score	1.396*** (0.416)		0.668* (0.380)					
Number of Observations	452	452	452	452				
Panel C. Instrument: Financial Education and Counseling Treatment								
Endline Awareness Score	3.610*** (0.903)			1.519*** (0.495)	1.953** (0.963)			0.409 (0.245)
Endline Attitudes Score	3.919*** (1.194)		1.649*** (0.522)		1.333** (0.552)		0.444* (0.234)	
Number of Observations	453	453	453	453	204	204	453	453
Panel D. Instrument: All Three Treatments								
Endline Awareness Score	3.116*** (0.441)			0.701** (0.289)			0.275* (0.142)	
Endline Attitudes Score	3.239*** (0.443)		0.728** (0.315)				0.286* (0.145)	
Number of Observations	457	457	457	457			457	457

Notes: This table presents instrumental variable regressions where the mediator is instrumented with the treatment assignment. The mediators considered are awareness (odd columns) or attitudes (even columns). In each panel, a different treatment variable is used as an instrument. Financial Education Only is a dummy equal to 1 for an individual who was assigned to the financial education program but not financial counseling nor goal setting. Financial Education and Goal Setting is a dummy equal to 1 for an individual who was assigned to the financial education and goal setting treatments, but not the financial counseling treatment. Financial Education and Counseling is a dummy equal to 1 for an individual who was assigned to the financial education and financial counseling treatments, but not the goal setting treatment. All Three Treatments is a dummy equal to 1 for an individual who was assigned to all three (i.e., financial education, counseling, and goal setting). In each panel, the regression sample consists of the control group (which was not assigned to any intervention) and those individuals assigned to a particular treatment. For example, in Panel A, the regression sample consists of those in the Financial Education Only treatment and those assigned to control. All regressions include baseline control variables (i.e., monthly discount rate and dummy variables for whether the respondent has inconsistent time preferences, is risk averse, self-reported having difficulty in saving, and self-reported being interested in financial matters). In addition, all regressions include strata dummies (where strata are defined by gender, neighborhood, and whether the respondent is a microfinance client). Standard errors are clustered at the wave-class level. *** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

interest in financial matters. Specifically, suppose that individuals who are more interested in finance at baseline (U) have higher post-treatment financial awareness (M) and are also more likely to open a bank savings account (Y). Then, U results in a positive correlation between ϵ_{2i} and ϵ_{3i} . In this case, if we are not able to control for U in our specifications, the estimate of γ_3 (and correspondingly, the ACME) will be biased upward.

Since the error terms ϵ_{2i} and ϵ_{3i} are inherently unobservable, we will never know the true sign of $cov(\epsilon_{2i}, \epsilon_{3i})$ and ρ . Nevertheless, we can use the residuals— $\hat{\epsilon}_{2i}$ (i.e., the sample residual from the regression of the mediator on the treatment) and $\hat{\epsilon}_{3i}$ (i.e., the sample residual from the regression of the financial outcome on the treatment and the mediator)—as our estimates of the population errors. We can then examine the relationship between these residuals to understand the potential sign of $cov(\epsilon_{2i}, \epsilon_{3i})$ and ρ . This is precisely our goal in Appendix Figs. A1 to A4, which provide the scatter plots of $\hat{\epsilon}_{2i}$ in the y -axis vs. $\hat{\epsilon}_{3i}$ in the x -axis. In these figures, the red line represents the fitted line from the bivariate linear regression of $\hat{\epsilon}_{2i}$ on $\hat{\epsilon}_{3i}$, thus indicating the correlation between the two. The figures show that the residuals appear to have a random relationship and a correlation of zero. We recognize that these patterns are not proof that our estimates of the ACME are unbiased, nor does it directly show that the SI assumption for identifying the ACME holds. Indeed, SI cannot be confirmed or refuted with our observable information. Even so, we find it encouraging that the figures show a random pattern—if the pattern were non-random, this would suggest that there are important confounders that are not captured in our regression specification.

Finally, throughout the discussion of sensitivity analysis in Section 5.4, we have focused on understanding how the ACME estimates change with varying levels of ρ . A similar approach to assessing sensitivity is provided by Altonji et al. (2005), which studies the effect of Catholic school attendance on educational outcomes. In such a study, selection on unobservables

Table A7
Relative Degree of Selection on Unobservables to Observables.

	Budgeting	Savings	Borrowing	Insurance
	Has tried making a budget in the last 6 months	Has a savings account	Loan purpose: Business, education, or purchase of durable goods	Bought life insurance in the last 6 months
	(1)	(2)	(3)	(4)
Panel A. Treatment: Financial Education Only				
Mediator: Endline Numeracy Score	0.118			
Mediator: Endline Awareness Score	1.331			
Mediator: Endline Attitudes Score	0.612			
Panel B. Treatment: Financial Education and Goal Setting				
Mediator: Endline Numeracy Score	0.127	0.252		
Mediator: Endline Awareness Score	1.179	0.402		
Mediator: Endline Attitudes Score	0.629	0.344		
Panel C. Treatment: Financial Education and Counseling				
Mediator: Endline Numeracy Score	0.138	0.241	-0.150	0.234
Mediator: Endline Awareness Score	1.034	0.367	1.502	0.044
Mediator: Endline Attitudes Score	0.674	0.333	2.305	0.342
Panel D. Treatment: All Three Treatments				
Mediator: Endline Numeracy Score	0.163	0.245		0.234
Mediator: Endline Awareness Score	0.600	0.321		-0.001
Mediator: Endline Attitudes Score	0.571	0.310		0.203

Notes: This table examines selection on unobservables in the regression of outcomes on the treatment, mediator, and baseline controls (i.e., the regressions in Panel A of Tables 5 to 8). The estimates in this table represent the relative degree of selection on unobservables that would be required to drive the coefficient on the mediator to zero, as proposed by Altonji et al. (2005). A high value implies that the results are robust to selection on observables. For example, the value 2.305 in column 3 indicates that in the regression of the borrowing outcome on the attitudes score, the financial education and counseling treatment dummy, and baseline controls, selection on unobservables must be 2.305 times more important than selection on observables to explain away the effect of attitudes. The estimates in this table were obtained using the implementation by Oster (2019).

is an important issue, as there may be many unobservable characteristics that determine school choice. The authors then provide a method for assessing the robustness of the results to selection on unobservables. In particular, assuming that selection on unobservables can be fully recovered from selection on observables, Altonji et al. (2005) estimate the ratio of two covariances—the covariance of Catholic school attendance and unobservables, and the covariance of Catholic school attendance and observables—which would result in a Catholic school effect of zero. This ratio tells us the degree of selection on unobservables relative to observables that would be required to explain away the effect of Catholic schools. This ratio is valid under the null hypothesis of no Catholic school effect, and it also requires the assumption that if we can observe the complete set of unobservables and include them in the regression, the regression R-squared would equal 1 (Oster, 2019).

The above ratio can then be used to understand the robustness of the results to selection on unobservables. If this ratio is high, then the results are robust to selection on unobservables. This is because attendance in Catholic schools would have to covary with unobservables a lot more than observables to drive the effect of Catholic schools to zero. Conversely, if the ratio is small, then the results are not robust to selection on unobservables: even a relatively weak relationship between unobservables and Catholic schools can explain away the result. Although there is no strict rule for how large this ratio must be to conclude robustness, Altonji et al. (2005) propose a cutoff of one. In other words, when the ratio is above one—that is, unobservables are more important than observables for explaining away the effects—Altonji et al. (2005) suggest that the results can be viewed as robust to selection on unobservables.

In our study context, selection on unobservables is an important issue when estimating the parameter γ_3 in Eq. (3) (i.e., the effect of the mediator on outcomes). We now calculate the ratio proposed by Altonji et al. (2005) for this parameter. Doing so answers the following question: how important would unobservables have to be, relative to observables, to drive the effect of the mediator (e.g., financial awareness) on outcomes to zero? We note that this question is conceptually very similar to the goal of the sensitivity analysis in Figs. 3 to 6.

We report estimates of this ratio in Appendix Table A7. To obtain these estimates, we use the implementation of Oster (2019), which extends the analysis in Altonji et al. (2005). We find that some of the estimated ratios are greater than one, but most are below one. These results indicate that our estimates of γ_3 in Eq. (3) are sensitive to selection on unobservables: a small amount of selection on unobservables (relative to observables) can drive the estimate of γ_3 to zero. Since γ_3 is also used to calculate the ACME, our ACME estimates would be sensitive to selection on unobservables as well. These results are consistent with our sensitivity analysis in Figs. 3 to 6, where we find that our ACME estimates are sensitive to unobserved pre-treatment confounding.

Table A8
Minimum Detectable Effects for the ACME.

	Budgeting	Savings	Borrowing	Insurance
	Has tried making a budget in the last 6 months (1)	Has a savings account (2)	Loan purpose: Business, education, or purchase of durable goods (3)	Bought life insurance in the last 6 months (4)
Panel A. Treatment: Financial Education Only				
Mediator: Endline Numeracy Score	0.010			
Mediator: Endline Awareness Score	0.027			
Mediator: Endline Attitudes Score	0.023			
Panel B. Treatment: Financial Education and Goal Setting				
Mediator: Endline Numeracy Score	0.006	0.008		
Mediator: Endline Awareness Score	0.031	0.024		
Mediator: Endline Attitudes Score	0.030	0.030		
Panel C. Treatment: Financial Education and Counseling				
Mediator: Endline Numeracy Score	0.007	0.007	0.013	0.002
Mediator: Endline Awareness Score	0.035	0.028	0.045	0.017
Mediator: Endline Attitudes Score	0.037	0.028	0.044	0.009
Panel D. Treatment: All Three Treatments				
Mediator: Endline Numeracy Score	0.010	0.008		0.002
Mediator: Endline Awareness Score	0.039	0.036		0.021
Mediator: Endline Attitudes Score	0.036	0.033		0.014

Notes: This table shows estimates of the ex-post Minimum Detectable Effect (MDE) sizes for the ACME, under 80% power and 5% significance level. The MDEs were obtained as 2.8 times the standard error (SE) of the ACME. The SE of the ACME are from Tables 5 to 8.

Appendix C. Variance of the treatment effect on financial numeracy

This appendix describes, in greater detail, two pieces of evidence demonstrating that the variance of the treatment effect on financial numeracy is likely to be zero. The discussion below draws on the methods outlined in Gerber and Green (2012).

First, we calculate the lower bound for the variance of the treatment effect, $var(\tau_i)$, on financial numeracy. The treatment effect τ_i is defined as the difference between the treated and untreated potential outcomes, $Y_i(1) - Y_i(0)$. To investigate the presence of heterogenous effects, we are interested in whether $var(\tau_i) = var[Y_i(1)] + var[Y_i(0)] - 2cov[Y_i(1), Y_i(0)]$ is greater than zero. Nevertheless, we cannot directly estimate $var(\tau_i)$ using our data: although our randomized experiment provides us with sample estimates of the variances of $Y_i(1)$ and $Y_i(0)$, we lack information on their covariance, $cov[Y_i(1), Y_i(0)]$.

Fortunately, as discussed in Gerber and Green (2012), we can obtain a lower bound for $var(\tau_i)$ as follows. We begin by sorting the financial numeracy scores (Y) in increasing order within treatment and control. Next, we pair each percentile of Y within the treatment distribution to its corresponding percentile in the control distribution. This pairing implies that the sample analog of $cov[Y_i(1), Y_i(0)]$ is as large as possible, and it allows us to obtain $\hat{\tau}_i$ (i.e., the sample estimates of τ_i). The summary statistic we are interested in is then the variance of these $\hat{\tau}_i$, which provides a lower bound for $var(\tau_i)$. If this lower bound is markedly greater than zero, this would indicate heterogeneity in effects of the treatment on financial numeracy. In fact, the variance of $\hat{\tau}_i$ that we estimate in the data is very small. We find a variance of around 0.01 for all treatments in our study, consistent with homogeneous treatment effects.

Second, we tested the null hypothesis that $var(\tau_i) = 0$ (i.e., homogenous treatment effects on numeracy) versus the alternative that $var(\tau_i) > 0$ (i.e., heterogenous treatment effects on numeracy). Note that because $Y_i(1) = Y_i(0) + \tau_i$, if $\tau_i = \tau$ for all persons i , then $var[Y_i(1)] = var[Y_i(0)]$. Conversely, if $var[Y_i(1)] \neq var[Y_i(0)]$ then we have heterogenous treatment effects. Therefore, rejecting $var[Y_i(1)] = var[Y_i(0)]$ amounts to rejecting $var(\tau_i) = 0$, and the above hypothesis test is equivalent to the two-sided test that the variance of numeracy scores is equal across treatment and control: $var[Y_i(1)] = var[Y_i(0)]$.

Using randomization inference, we implemented the test of $var[Y_i(1)] = var[Y_i(0)]$ based on the approach in Gerber and Green (2012). This exercise proceeds as follows. We first assume that the null hypothesis is true: we have a constant treatment effect on numeracy that is equal to the estimate $\hat{\beta}_2$ from Eq. (2). These estimates of $\hat{\beta}_2$ correspond to column 1 of Table 4. Next, for each person in the control group, we calculate the potential treated numeracy score as $Y_i(1) = Y_i(0) + \hat{\beta}_2$. Similarly, for each individual in the treatment group, we calculate the potential control numeracy score as $Y_i(0) = Y_i(1) - \hat{\beta}_2$. Afterwards, we randomly assign all subjects to a “fake” treatment, and we do so 100,000 times. In each iteration, we record the estimated difference $var[Y_i(1)] - var[Y_i(0)]$ under the “fake” random assignment. The proportion of the 100,000 estimated differences that is greater than the actual difference in variances we observe in the data (all in absolute value terms) is then the p-value of our hypothesis test.

The results are shown graphically in Appendix Figure A5. Each sub-figure plots the distribution of the 100,000 estimated differences in the variance of numeracy scores between the “fake” treatment and control groups. The red line is the difference in the variance of numeracy scores that we observe in the actual data. We find that under the null hypothesis

Table A9
Multiple Inference Adjustment.

	Budgeting			Savings			Borrowing			Insurance		
	Has tried making a budget in the last 6 months			Has a savings account			Loan purpose: Business, education, or purchase of durable goods			Bought life insurance in the last 6 months		
	Financial Numeracy (1)	Financial Awareness (2)	Financial Attitudes (3)	Financial Numeracy (4)	Financial Awareness (5)	Financial Attitudes (6)	Financial Numeracy (7)	Financial Awareness (8)	Financial Attitudes (9)	Financial Numeracy (10)	Financial Awareness (11)	Financial Attitudes (12)
Panel A. Treatment: Financial Education Only												
ATE p-value		0.001			0.899			0.502			0.937	
ATE q-value		0.004			0.993			0.658			0.993	
ACME p-value	0.489	0.005	0.003									
ACME q-value	0.652	0.014	0.011									
ADE p-value	0.000	0.003	0.001									
ADE q-value	0.001	0.010	0.004									
Panel B. Treatment: Financial Education and Goal Setting												
ATE p-value		0.000			0.082			0.924			0.956	
ATE q-value		0.003			0.125			0.993			0.993	
ACME p-value	0.997	0.069	0.001	0.777	0.937	0.068						
ACME q-value	0.997	0.116	0.003	0.920	0.993	0.116						
ADE p-value	0.000	0.001	0.003	0.047	0.048	0.150						
ADE q-value	0.001	0.005	0.010	0.096	0.096	0.224						
Panel C. Treatment: Financial Education and Counseling												
ATE p-value		0.000			0.000			0.027			0.076	
ATE q-value		0.001			0.002			0.061			0.123	
ACME p-value	0.748	0.009	0.003	0.786	0.611	0.065	0.943	0.246	0.98	0.699	0.645	0.340
ACME q-value	0.917	0.024	0.010	0.920	0.787	0.116	0.993	0.354	0.993	0.872	0.817	0.470
ADE p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.005	0.007	0.007	0.045	0.053	0.070
ADE q-value	0.001	0.001	0.001	0.001	0.002	0.002	0.014	0.020	0.019	0.096	0.105	0.116
Panel D. Treatment: All Three Treatments												
ATE p-value		0.000			0.023			0.426			0.056	
ATE q-value		0.001			0.054			0.579			0.107	
ACME p-value	0.884	0.067	0.000	0.903	0.272	0.006				0.974	0.774	0.181
ACME q-value	0.993	0.116	0.001	0.993	0.383	0.016				0.993	0.920	0.265
ADE p-value	0.000	0.000	0.000	0.012	0.016	0.081				0.032	0.058	0.077
ADE q-value	0.001	0.001	0.001	0.029	0.039	0.125				0.070	0.108	0.123

Notes: This table reports results of multiple inference adjustments. The table presents the non-adjusted p-values of the Average Treatment Effect (ATE), the Average Causal Mediating Effect (ACME), and the Average Direct Effect (ADE). Note that the ACME p-values shown here are normal-based. The adjusted p-values are the [Benjamini and Hochberg \(1995\)](#) q-values as described in [Anderson \(2008\)](#).




Fig. A1. Scatterplot of $\hat{\epsilon}_{2j}$, $\hat{\epsilon}_{3j}$ (Treatment: Financial Education) Has tried making a budget in the last 6 months. This figure shows a scatter plot where the y-axis is the sample residuals from the regression of the mediator on the treatment ($\hat{\epsilon}_{2j}$), and the x-axis is the sample residuals from the regression of financial outcomes on the treatment and the mediator ($\hat{\epsilon}_{3j}$). The red line represents the fitted line from the bivariate linear regression.

of no heterogeneous effects, the difference in variance that we see in our data is not surprising. This holds true for all treatments: the p-values are 0.76, 0.29, 0.66, and 0.22 for the financial education only, the financial education and goal setting, financial education and counseling, and all three interventions, respectively. As a result, we fail to reject the null hypothesis that the effects of the treatment on numeracy—which we estimate to be zero on average—is the same for all subjects.

Appendix D. Robustness to statistical power concerns

In this appendix, we discuss two points related to statistical power: (1) heterogeneity in the sample, and (2) Minimum Detectable Effect (MDE) sizes of the ACME.

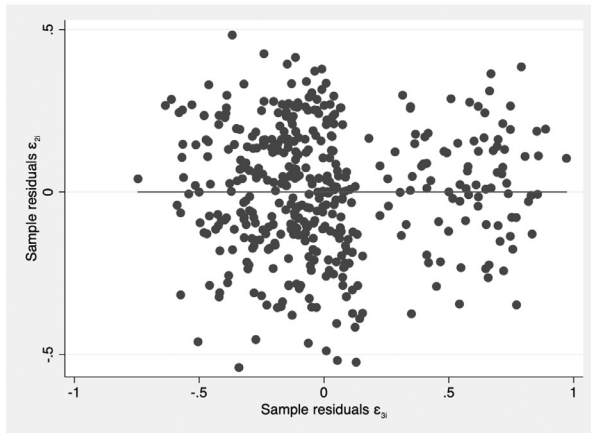
First, in any study, statistical power is influenced not only by the magnitude of the treatment effect and the sample size, but also by the amount of heterogeneity in the sample. In particular, with a more diverse mix of individuals, it is more difficult to measure average differences in outcomes due to the treatment (e.g., [Glennerster and Takavarasha, 2013](#)). Therefore, to better understand whether statistical power may be a potential issue in our setting, it is important to examine the diversity in the characteristics of our study sample.

One way we can formally assess this heterogeneity is to look at the Coefficient of Variation (CV)—i.e., the ratio of the standard deviation to the mean—of baseline variables, shown in column 4 of [Table 2](#). As the table shows, the CV for almost all continuous variables in the table are well below one. The only exception is the respondent's baseline monthly discount rate, which has a CV of 3.18. The high CV for this variable may be because time preferences are difficult to measure in the field, given that other factors such as the survey context may influence responses (e.g., [Cohen et al., 2016](#)). Still, the CVs in [Table 2](#) show that there appears to be little variability across subjects in their household background as well as their baseline math and financial literacy scores.


This lack of diversity in baseline characteristics is also echoed in our endline results for the treatment and mediating effects, as we discussed in [Sections 6.1 and 6.3](#). In particular, we do not find heterogeneity in the impact of our treatments on financial numeracy, suggesting that the null effects of numeracy we observe may hold at the individual level and not just on average. We also do not find evidence of heterogeneity in the ACME by treatment group. Further, we would like to highlight that in our experimental design, we stratified randomization by baseline characteristics (i.e., the respondent's gender, neighborhood, and current status as a microfinance client), which helps to increase statistical power ([Bruhn and McKenzie, 2009](#)).

Second, although it is important to assess power, *ex post* power calculations can be problematic. Using the observed estimates of the effect size to calculate power retrospectively can provide misleading results since there is noise in the observed effect size ([Gelman, 2019; McKenzie and Ozier, 2019](#)). Hence, in lieu of *ex post* power calculations, we consider the MDEs of the mediating effects. We note that the standard errors of the ACME reported throughout the paper are based on either non-parametric bootstrapping (meaning that we do not assume an underlying distribution or rely on asymptotic arguments) or quasi-Bayesian Monte Carlo approximation. Nevertheless, to obtain back-of-the-envelope, rough estimates of the MDEs, we assume normality as a simplification.

Panel A: Has tried making a budget in the last 6 months




(A1) Awareness




(A2) Attitudes

Panel B: Has savings account



(B1) Awareness



(B2) Attitudes

Fig. A2. Scatterplot of $\hat{\epsilon}_{2i}$, $\hat{\epsilon}_{3i}$ (Treatment: Financial Education and Goal Setting)

Notes: This figure shows a scatter plot where the y-axis is the sample residuals from the regression of the mediator on the treatment ($\hat{\epsilon}_{2i}$), and the x-axis is the sample residuals from the regression of financial outcomes on the treatment and the mediator ($\hat{\epsilon}_{3i}$). The red line represents the fitted line from the bivariate linear regression.


With 80% power and 5% significance level, we obtain the *ex post* MDE as 2.8 times the observed standard error. This follows from the formula for calculating the MDE (e.g., [Glennester and Takavarasha, 2013](#)), given by

$$MDE = (t_{(1-k)} + t_{\alpha/2}) * \sqrt{\text{Var}(\hat{E})}.$$

Here, $\text{Var}(\hat{E})$ denotes the variance of the estimator (denoted \hat{E}), $1 - k$ is the power, and α is the significance level. The parameter t_α is the critical value such that $\Phi(-t_\alpha) = \alpha$, where Φ is the standard normal cumulative distribution function. Since we assume 80% power and 5% significance, we have that $t_{0.8} + t_{0.025} = 0.84 + 1.96 = 2.8$. We then use the observed standard error as the value for $\text{Var}(\hat{E})$.

Importantly, unlike calculating *ex post* power using observed effect sizes, calculating the *ex post* MDEs using observed standard errors is a valid approach. Although there will still be sampling variability in the *ex post* MDEs, this variability is much lower than with *ex post* power calculations because the MDEs do not rely on potentially noisy estimates of the effect size. Our calculation of the MDE therefore answers the following question: if we assume that the amount of uncertainty in the estimates is the same as what we observe in our data, what is the smallest effect that we can pick up given 80% power and 5% significance?


Panel A: Has tried making a budget in the last 6 months



(A1) Awareness

(A2) Attitudes


Panel B: Has savings account



(B1) Awareness

(B2) Attitudes


Panel C: Loan purpose: Business, education, or purchase of durable goods



(C1) Awareness

(C2) Attitudes

Panel D: Bought life insurance in the last 6 months



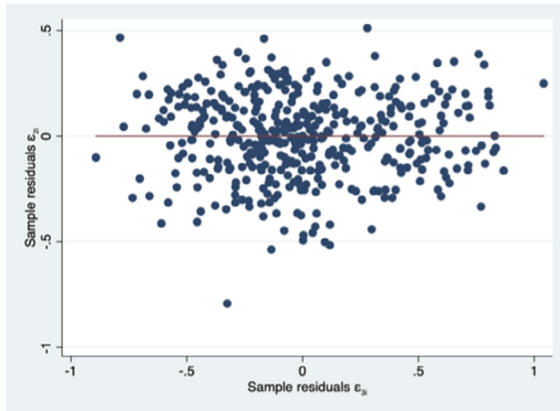
(D1) Awareness

(D2) Attitudes


Fig. A3. Scatterplot of $\hat{\epsilon}_{2i}$, $\hat{\epsilon}_{3i}$ (Treatment: Financial Education and Counseling)

Notes: This figure shows a scatter plot where the y-axis is the sample residuals from the regression of the mediator on the treatment ($\hat{\epsilon}_{2i}$), and the x-axis is the sample residuals from the regression of financial outcomes on the treatment and the mediator ($\hat{\epsilon}_{3i}$). The red line represents the fitted line from the bivariate linear regression.

Panel A: Has tried making a budget in the last 6 months




(A1) Awareness




(A2) Attitudes

Panel B: Has savings account




(B1) Awareness




(B2) Attitudes

Panel C: Bought life insurance in the last 6 months



(D1) Awareness



(D2) Attitudes

Fig. A4. Scatterplot of $\hat{\epsilon}_{2i}$, $\hat{\epsilon}_{3i}$ (Treatment: All Three Treatments)

Notes: This figure shows a scatter plot where the y-axis is the sample residuals from the regression of the mediator on the treatment ($\hat{\epsilon}_{2i}$), and the x-axis is the sample residuals from the regression of financial outcomes on the treatment and the mediator ($\hat{\epsilon}_{3i}$). The red line represents the fitted line from the bivariate linear regression.




Fig. A5. Test for the Variance of the Treatment Effect on Numeracy

Note: This figure presents results of randomization inference to test the null hypothesis that the variance of the treatment effect on numeracy scores is zero. Each sub-figure above considers a different treatment variable. The approach here follows Gerber and Green (2012), and the exercise proceeds as follows. We first assume that the null hypothesis is true: we have a constant treatment effect that is equal to the estimates from column 1 of Table 4, denoted $\hat{\beta}_2$ from Eq. (2). Next, for each person in the control group, we calculate the potential treated numeracy score as $Y_i(1) = Y_i(0) + \hat{\beta}_2$. Similarly, for each individual in the treatment group, we calculate the potential control numeracy score as $Y_i(0) = Y_i(1) + \hat{\beta}_2$. Afterwards, we randomly assign all subjects to a “fake” treatment, and we do so 100,000 times. In each iteration, we record the estimated difference $var[Y_i(1)] - var[Y_i(0)]$ under the “fake” random assignment. The proportion of the 100,000 estimated differences that is greater than the actual difference in variances we observe in the data (all in absolute value terms) is the p -value of our hypothesis test. Each sub-figure above plots the distribution of the 100,000 estimated differences in the variance of numeracy scores between the “fake” treatment and control groups. The red line is the difference in the variance of numeracy scores that is observed in the actual data. The corresponding p -values are 0.76, 0.29, 0.66, and 0.22 for the financial education only, the financial education and goal setting, financial education and counseling, and all three interventions, respectively.

The results of the MDE calculation are reported in Appendix Table A8. We find that the MDEs for the ACME for financial numeracy are all very close to zero. Because we would have been able to detect even small ACMEs for numeracy, this suggests that the null estimates of the ACME that we find for numeracy represent true null effects. Furthermore, it is consistent with our earlier findings that the effect of the treatments on financial numeracy is zero, both on average and across all individuals. With these patterns, we can reject financial numeracy as a mechanism for the treatment effects. Meanwhile, for awareness and attitudes, we acknowledge that we may be underpowered in estimating the mediating effects in some cases. For example, the ACME estimate in column 5 of Table 6 suggests that awareness does not play a mediating role for the effect of the financial education with goal setting treatment on savings: the ACME estimate of 0.001 is close to zero and not statistically significant. However, looking at the MDE for this effect in Appendix Table A8, Panel B (column 2), we find a value of 0.024. Therefore, it is possible that the ACME of awareness for the financial education with goal setting treatment on savings outcomes is some positive number less than 0.024, but we are not able to detect this effect with our data.

References

- Acharya, A., Blackwell, M., Sen, M., 2016. Explaining causal findings without bias: Detecting and assessing direct effects. *American Political Science Review* 110 (3), 512–529.
- Albert, J.M., 2008. Mediation analysis via potential outcomes models. *Statistics in Medicine* 27 (8), 1282–1304.
- Alesina, A., Giuliano, P., Nunn, N., 2013. On the origins of gender roles: Women and the plough. *The Quarterly Journal of Economics* 128 (2), 469–530.
- Altonji, J.G., Elder, T.E., Taber, C.R., 2005. Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools. *Journal of Political Economy* 113 (1), 151–184.
- Ahamed, M.M., Mallick, S.K., 2019. Is financial inclusion good for bank stability? International evidence. *Journal of Economic Behavior & Organization* 157, 403–427.
- Anderson, M.L., 2008. Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American Statistical Association* 103 (484), 1481–1495.
- Angrist, J.D., Pischke, J.S., 2008. *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press, Princeton, NJ.
- Baron, R.M., Kenny, D.A., 1986. The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology* 51 (6), 1173–1182.
- Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: series B (Methodological)* 57 (1), 289–300.
- Berg, G., Zia, B., 2017. Harnessing emotional connections to improve financial decisions: Evaluating the impact of financial education in mainstream media. *Journal of the European Economic Association* 15 (5), 1025–1055.
- Bruhn, M., Leão, L.D.S., Legovini, A., Marchetti, R., Zia, B., 2016. The impact of high school financial education: Evidence from a large-scale evaluation in Brazil. *American Economic Journal: Applied Economics* 8 (4), 256–295.
- Bruhn, M., McKenzie, D., 2009. In pursuit of balance: Randomization in practice in development field experiments. *American Economic Journal: Applied Economics* 1 (4), 200–232.
- Carpena, F., Cole, S., Shapiro, J., Zia, B., 2017. The ABCs of financial education: Experimental evidence on attitudes, behavior, and cognitive biases. *Management Science* 65 (1), 346–369.
- Cohen, J. D., Ericson, K. M., Laibson, D., & White, J. M. (2016). Measuring time preferences. *NBER Working Paper No. 22455*.
- Dalton, P. S., Pamuk, H., Ramrattan, R., van Soest, D., & Uras, B. (2019). Transparency and Financial Inclusion: Experimental Evidence from Mobile Money. *Center Discussion Paper No. 2018-042*.
- Doi, Y., McKenzie, D., Zia, B., 2014. Who you train matters: Identifying combined effects of financial education on migrant households. *Journal of Development Economics* 109, 39–55.
- Fearon, J.D., Laitin, D.D., 2003. Ethnicity, Insurgency, and Civil War. *American Political Science Review* 97 (01), 75–90.
- Fernandes, D., Lynch Jr, J.G., Netemeyer, R.G., 2014. Financial literacy, financial education, and downstream financial behaviors. *Management Science* 60 (8), 1861–1883.
- Gelman, A., 2019. Don't calculate post-hoc power using observed estimate of effect size. *Annals of Surgery* 269 (1), e9–e10.
- Gerber, A.S., Green, D.P., 2012. *Field experiments: Design, analysis, and interpretation*. W.W. Norton & Company, New York, NY.
- Glennerster, R., Takavarasha, K., 2013. *Running randomized evaluations: A practical guide*. Princeton University Press, Princeton, NJ.
- Green, D.P., Ha, S.E., Bullock, J.G., 2010. Enough already about “black box” experiments: Studying mediation is more difficult than most scholars suppose. *The Annals of the American Academy of Political and Social Science* 628 (1), 200–208.
- Hastings, J.S., Madrian, B.C., Skimmyhorn, W.L., 2013. Financial literacy, financial education, and economic outcomes. *Annual Review of Economics* 5 (1), 347–373.
- Hicks, R., Tingley, D., 2011. Causal mediation analysis. *The Stata Journal* 11 (4), 1–15.
- Holzmann, R., 2010. Bringing Financial Literacy and Education to Low and Middle Income Countries: The Need to Review, Adjust, and Extend Current Wisdom. *World Bank Social Protection & Labor Discussion Paper No. 1007*.
- Imai, K., Keele, L., Tingley, D., 2010. A general approach to causal mediation analysis. *Psychological Methods* 15 (4), 309–334.
- Imai, K., Keele, L., Tingley, D., Yamamoto, T., 2011. Unpacking the black box of causality: Learning about causal mechanisms from experimental and observational studies. *American Political Science Review* 105 (4), 765–789.
- Imai, K., Keele, L., Yamamoto, T., 2010. Identification, inference and sensitivity analysis for causal mediation effects. *Statistical Science* 51–71.
- Imai, K., Tingley, D., Yamamoto, T., 2013. Experimental designs for identifying causal mechanisms. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 176 (1), 5–51.
- Imbens, G.W., Angrist, J.D., 1994. Identification and estimation of Local Average Treatment Effects. *Econometrica* 62 (2), 467–475.
- Jo, B., 2008. Causal inference in randomized experiments with mediational processes. *Psychological Methods* 13 (4), 314.
- Keele, L., 2015. Causal mediation analysis: Warning! Assumptions ahead. *American Journal of Evaluation* 36 (4), 500–513.
- King, G., Tomz, M., Wittenberg, J., 2000. Making the most of statistical analyses: Improving interpretation and presentation. *American Political Science Review* 44 (2), 341–355.
- Lusardi, A., Mitchell, O.S., 2009. How Ordinary Consumers Make Complex Economic Decisions: Financial Literacy and Retirement Readiness. *NBER Working Paper No. 15350*.
- Lusardi, A., & Mitchell, O.S. (2013). The Economic Importance of Financial Literacy: Theory and Evidence. *NBER Working Paper No. 18952*.
- McKenzie, D., Ozier, O., 2019. Why ex-post power using estimated effect sizes is bad, but an ex-post MDE is not Retrieved from <https://blogs.worldbank.org/impactevaluations/why-ex-post-power-using-estimated-effect-sizes-bad-ex-post-mde-not>.
- Miller, M., Reichelstein, J., Salas, C., Zia, B., 2015. Can You Help Someone Become Financially Capable? A Meta-Analysis of the Literature. *The World Bank Research Observer* 30 (2), 220–246.
- Oster, E., 2019. Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics* 37 (2), 187–204.
- Pearl, J., 1995. Causal diagrams for empirical research. *Biometrika* 82 (4), 669–688.
- Pearl, J., 2000. *Causality: Models, Reasoning, Inference*. Cambridge University Press, Cambridge, UK.
- Robins, J.M., 2003. Semantics of causal DAG models and the identification of direct and indirect effects. *Oxford Statistical Science Series* 70–82.
- Rosenbaum, P., 1984. The Consequences of Adjustment for a Concomitant Variable That Has Been Affected by the Treatment. *Journal of the Royal Statistical Society* 147 (5), 656–666 Series A (General).
- Sayinzoga, A., Bulte, E.H., Lensink, R., 2016. Financial literacy and financial behaviour: Experimental evidence from rural Rwanda. *The Economic Journal* 126 (594), 1571–1599.
- Sobel, M.E., 2008. Identification of causal parameters in randomized studies with mediating variables. *Journal of Educational and Behavioral Statistics* 33 (2), 230–251.
- VanderWeele, Tyler J., 2015. *Explanation in causal inference: Methods for mediation and interaction*. Oxford University Press, Oxford, England.
- Willis, L.E., 2011. The Financial Education Fallacy. *American Economic Review, Papers & Proceedings* 101 (3), 429–434.
- World Bank. (2017). *Global Financial Inclusion and Consumer Protection Survey, 2017 Report*. World Bank, Washington, DC.