Using a Life-Cycle Model to Evaluate Financial Literacy Program Effectiveness *

Annamaria Lusardi[†], Pierre-Carl Michaud[‡], and Olivia S. Mitchell[§]

December 7, 2015

Abstract

Prior studies disagree regarding the effectiveness of financial literacy programs, especially those offered in the workplace. To explain such measurement differences in evaluation and outcomes, we employ a stochastic life cycle model with endogenous financial knowledge accumulation to investigate how financial education programs optimally shape key economic outcomes. This approach permits us to measure how such programs shape wealth accumulation, financial knowledge, and participation in sophisticated assets (e.g. stocks) across heterogeneous consumers. We then apply conventional program evaluation econometric techniques to simulated data, distinguishing selection and treatment effects. We show that the more effective programs provide follow-up in order to sustain the knowledge acquired by employees via the program; in such an instance, financial education delivered to employees around the age of 40 can raise savings at retirement by close to 10%. By contrast, one-time education programs do produce short-term but few long-term effects. We also measure how accounting for selection affects estimates of program effectiveness on those who participate. Comparisons of participants and non participants can be misleading, even using a difference-in-difference strategy. Random program assignment is needed to evaluate program effects on those who participate.

[†]The George Washington University School of Business & NBER; alusardi@gwu.edu
[‡]Université du Québec à Montréal & RAND; michaud.pierre_carl@uqam.ca
[§]Wharton School & NBER; mitchelo@wharton.upenn.edu

1 Introduction

Employees and their families are increasingly responsible for securing their own financial well-being. Prior to the 1980s, U.S. workers relied mainly on Social Security and employer sponsored defined benefit (DB) pension plans for their retirement security. Today, by contrast, Baby Boomers are increasingly relying on defined contribution (DC) plans and Individual Retirement Accounts (IRAs) to finance their golden years. The transition to a DC retirement saving model has the advantage of permitting more worker flexibility and labor mobility than in the past, yet it imposes a greater responsibility on individuals to save, invest, and decumulate their retirement wealth sensibly. At the same time, financial markets have become more complex, offering products that are often difficult to understand. Whether individuals—in particular, older individuals—are equipped to deal with this new financial landscape is an important question that has implications for families, society, and policy makers.

Traditional economic models of saving and consumption decisions implicitly assume that people are able to formulate and execute saving and decumulation plans, all of which require expertise in dealing with financial markets, and that they have the capacity to undertake complex economic calculations. Yet, as Lusardi and Mitchell have reported (2008, 2009, 2011a,b), few people possess the financial knowledge adequate to make and execute complex financial plans. Moreover, acquiring such knowledge is likely to come at a cost. In our prior work (Lusardi, Michaud and Mitchell; hereafter LMM, forthcoming), we built and calibrated a stochastic life cycle model featuring uncertainty in income, longevity, capital market returns, and medical expenditures; that study also incorporated an endogenous knowledge accumulation process and a sophisticated saving technology. In the model, financial knowledge provided consumers with access to sophisticated financial products that boosted their expected return on financial assets. Naturally, those seeking to transfer resources over time by saving benefited most from financial knowledge.

The contribution of the present paper is to show how our stochastic life cycle model

incorporating endogenous knowledge accumulation may be used to help evaluate financial literacy programs. Specifically, since knowledge is at the core of the model, the approach permits us to evaluate how financial education policies can influence saving and investment decisions. Several prior studies have sought to measure how financial literacy programs changes behavior, but few have the kind of experimental data to capture precisely what the impact of the interventions actually is. Using our model, we evaluate the effectiveness of efforts to build workplace financial education using econometric methods commonly used to estimate the effect of such programs. Inasmuch as all counterfactuals are known in the context of our model, this allows us to compare "true" outcomes with estimates commonly generated by conventional program evaluation techniques. We show that it is frequently optimal for individuals to fail to invest in knowledge, as it is expensive to acquire and will not benefit everyone. Nevertheless, providing employees with financial knowledge can be valuable, depending on when it is offered and what reinforcement is provided. To this end, we use conventional program evaluation econometric techniques and simulated data to take into account selection and treatment effects: this allows us to measure how such programs shape wealth accumulation, financial knowledge, and participation in sophisticated assets (e.g. stocks) across heterogeneous consumers. Relatively more effective programs are those which embed follow-up or are continued over time, so as to help employees retain knowledge acquired via the program. In this case, financial education delivered to employees around the age of 40 will optimally enhance savings at retirement by close to 10%. By contrast, programs that provide one-time education can generate short-term but few long-term effects. Finally, we evaluate how important it is to account for selection in program participation. We conclude that comparing participants and non participants, even in a difference-in-difference framework, can deliver misleading estimates of program effectiveness.

The paper has several parts. First, we briefly summarize prior studies, and next, we describe our model and outline our calibration approach. We then present a series of scenarios where we evaluate the simulated impacts of alternative financial education programs. In turn, we use the resulting datasets to examine various econometric models conventionally

used to evaluate such programs. The paper concludes with a short discussion of the insights that policy and the finance and pension industry can gain from this work.

2 Prior Literature

In the wake of the financial crisis and ensuing Great Recession, interest has burgeoned in programs seeking to enhance financial literacy. For instance, the Organisation for Economic Co-operation and Development (OECD) has published a long list of reports on the importance of financial literacy and financial education programs. Several education programs in the U.S. focus on educational interventions for young people before they enter the labor market (Mandell, 2008; Walstad, Rebeck and MacDonald. 2010; Richardson and Seligman, 2014), while others examine programs offered to working-age adults, often by employers who seek to enhance employees' appreciation of and investment in their workplace-based financial literacy education (e.g., Bernheim and Garrett, 2003; Clark, d'Ambrosio, McDermed, and Sawant 2006; Lusardi, Keller, and Keller, 2008; Clark, Morrill, and Allen 2012).

Despite the widespread popularity of such programs in the U.S. and elsewhere, our recent literature review (Lusardi and Mitchell, 2014) as well as Collins and O'Rourke (2010) argued that relatively little could be learned from most of the existing evaluations to date. This is because analysts have typically not followed the protocol required by 'gold standard' randomized controlled trials, enabling researchers to extrapolate from observed results. More specifically, a good evaluation will compare outcomes for a randomly selected 'treatment' versus 'control' group, where the former will be exposed to a well-defined financial literacy program, while the latter will not (Imbens and Woolridge, 2009; Imbens, 2010). To this end, the modern program evaluation literature has identified three commonly used metrics for such comparisons: an Intent to Treat (ITT) measure, an Average Treatment Effect on the Treated (ATET) measure, and a Local Average Treatment (LATE) measure. In our context, the ITT compares outcomes of those who were versus were not offered the program, irrespective of whether and which people actually elected the program when offered. The

ATET measures the effect for the treated, not the average effect of moving someone into treatment, and hence it is often the only way to estimate program effects when selection is present; that is, one may not be able to evaluate a program's average treatment effect when those who do participate endogenously differ from those who do not.[1] Finally, the LATE measure, as defined by Angrist and Imbens (1994), captures the effect of the program for those who would participate in the program only if it was offered. .[2]Randomization of eligibility is a key ingredient for the recovery of LATE by instrumental variables regression.

In the context of financial education programs, some authors seeking to evaluate the impact of the programs have estimated ITT effects by comparing outcomes for people who were and were not exposed to the programs, given the option to undertake education programs. Good examples include studies of programs mandating high school financial literacy programs at different times across states (c.f., Bernheim, Garrett, and Maki 2001; Bayer, Bernheim and Sholz 2009). Yet other researchers have estimated the effect of participating in a program which may include both treatment and selection effects; numerous examples are cited in Lusardi and Mitchell (2014). And finally, several researchers have sought to estimate program effectiveness using instrumental variables estimation, seeking to control on potential unobserved factors driving program participation and thus recover the LATE (Lusardi and Mitchell, 2014). Our general conclusion, however, is that much remains to be learned about how financial education affects key outcomes of interest. Without a well-defined control group selected via randomized assignment, it is typically difficult to measure the effect of financial education programs, since assumptions needed to estimate what program adopters would have done in the absence of the program (the counterfactual) are probably too strong.

To remedy this problem, we show below how we can use our model (LMM forthcoming) to help clarify what can happen when a financial education program evaluation lacks a guiding theoretical framework. Most importantly, given individual heterogeneity and the

---

[1]In some cases, however, if a proper counterfactual can be identified, the average treatment effect can be estimated.

[2]In a randomized control trial with one-sided non-compliance (individuals not assigned to treatment cannot receive it), the LATE estimate may coincide with the ATET effect.

costs and benefits of financial literacy, not everyone will gain from financial education. Accordingly, one should not expect a 100% participation rate in every financial education program. Moreover, according to our model, financial education programs may not always boost savings, and in fact they may not increase savings at all for some. Therefore it is inaccurate to conclude that lack of saving means that financial education is ineffective. Instead, lack of saving can actually be optimal behavior for some, and financial education would not be expected to change that behavior. In this respect, our framework helps explain who is likely to participate in such programs, what behavioral outcomes can result, and whether lack of impact is proof of program ineffectiveness.

3   The Model and Calibration

3.1   Model

In what follows, we focus on workplace financial education programs of the sort most often offered by employers with defined contribution pensions.[3] We consider employees who can elect to take advantage of such programs, which for the present purposes can be conceptualized as financial education of one year's duration, delivered to employees who have not previously anticipated getting the offer.

We characterize each program in terms of three key parameters: an eligibility rule, a program cost, and the program's effectiveness. We assume eligibility is assigned randomly to all employees of a given age, which we vary across experimental settings (more on this below). The impact of the financial education program is to reduce the employee's cost of investing in knowledge. When a program is of high quality, it provides an incentive to acquire more knowledge, and individual employees will then decide whether to participate in the program. Costs matter as well: for instance, if the program were free, all workers will participate (or at best they will be indifferent). In order to capture the time/money costs of participating in the program, we model the participation cost for the program as a

---

[3]See for instance, Bernheim and Garrett (2003); Bayer, Bernheim, and Sholz (2009); Clark, d'Ambrosio, McDermed, and Sawant (2006), and Clark, Morrill, and Allen (2014).

fixed variable; a more general framework could depend on income or education, but for the present purposes we keep it fixed.

The remainder of the model follows our prior work (LMM forthcoming). Each individual is posited to select his consumption stream by maximizing expected discounted utility, where utility flows are discounted by β. Utility is assumed to be strictly concave in consumption and defined as $n_t u(c_t/n_t)$, where n_t is an equivalence scale capturing (known) differences in consumption patterns across demographic groups (Scholz, Seshadri, and Khitatrakun, hereafter SSK 2006). Each person's faces a stochastic mortality risk (in addition to income and medical expenditure risk), and decisions are made from time t=0 (age 25) to age T (or as long as the individual is still alive; T=100). We examine people of three different education profiles (High School dropouts or <HS; High school graduates or HS; and those with at least some college, whom we call the College+). It is important to allow for heterogeneity in earnings because different groups receive different rewards from the progressive social insurance system, as described in LMM (forthcoming), and they face differential patterns of income, mortality, demographics, and out-of-pocket medical expenditure risk.

We also posit that the individual can invest his resources using two different investment technologies. One is a basic technology (for example, a checking account) which yields a certain (low) return $\bar{r}$ ($\bar{R} = 1 + \bar{r}$). This represents the expected return to consumers without any financial know-how. The other is a more sophisticated technology which enables the consumer to receive a higher expected return which increases in financial knowledge f but comes at a cost. Specifically, the consumer must pay a direct cost (fee) to use the technology, c_d, and he must also invest time and money in acquiring the knowledge to generate a sufficiently high excess return. Obtaining knowledge in the form of investment i_t thus has a cost of $\pi_i(i_t)$; we assume that this cost function is convex, reflecting decreasing returns in the production of knowledge. We remain agnostic about whether the average cost of investing in additional knowledge is higher or lower for more educated households; rather, we assume initially that all households face the same cost function. The rate of return to the sophisticated technology is stochastic, with an expected return that depends

on the individual's level of financial knowledge at the end of t, $\widetilde{R}(f_{t+1})$. Thus the stochastic return function is log-normally distributed with $\log \widetilde{R}(f_{t+1}) = \bar{r} + r(f_t) + \sigma_\varepsilon \varepsilon_t$ where σ_ε is the standard deviation of a normally distributed shock ε_t. The function $r(f_{t+1})$ is increasing in f_{t+1} and it can be interpreted as an excess return function. Since the variance is assumed fixed, this also implies that individuals with higher financial knowledge obtain a higher Sharpe ratio (higher risk-adjusted returns) on their investments. We denote by κ_t the fraction of wealth that the consumer invests in the sophisticated technology in period t.

Financial knowledge evolves according to the following equation:

$$f_{t+1} = (1 - \delta)f_t + i_t$$

where δ is a depreciation rate and i_t is gross investment. Depreciation exists both because consumer financial knowledge may decay, and also because some knowledge may become obsolete as new financial products are developed. Alternatively, financial education can be modeled as a permanent boost to knowledge if the depreciation rate were to become smaller or even zero.

The consumer is also eligible for a government transfer tr_t which guarantees a minimum consumption floor of c_{min} (as in Hubbard, Skinner, and Zeldes, hereafter HSZ; 1995). This consumption floor can lower the expected variance of future consumption, which diminishes the precautionary motive for saving. Transfers are defined as $tr_t = \max(c_{min} - x_t, 0)$ where cash on hand is:

$$x_t = a_t + y_t - oop_t$$

where y_t is net household income and oop_t represents out-of-pocket medical expenditures. Both variables are stochastic over and above a deterministic trend. The sophisticated technology cannot be purchased if $x_t - c_d < c_{min}$ (that is, the government will not pay for costs

of obtaining the technology). End-of-period assets are given by:

$$a_{t+1} = \widetilde{R}_\kappa(f_{t+1})(x_t + tr_t - c_t - \pi(i_t) - c_d I(\kappa_t > 0))$$

where $\widetilde{R}_\kappa(f_{t+1}) = (1 - \kappa_t)\overline{R} + \kappa_t \widetilde{R}(f_t)$. We impose a borrowing constraint on the model such that assets a_{t+1} must be non-negative.

Following the literature, the individual's net income (in logs) during his worklife is given by a deterministic component which depends on education, age, and an AR(1) stochastic process; retirement occurs at age 65. After retirement, the individual receives retirement income which is a function of pre-retirement income and a similar stochastic AR(1) process is assumed for post-retirement out-of-pocket medical expenditures.[4]. Finally, we allow for mortality risk at all ages, denoting $p_{e,t}$ as the one-year survival probability. Mortality risk is allowed to differ across education groups, as in LMM (forthcoming).

The state-space in period t is defined as $s_t = (\eta_{y,t}, \eta_{o,t}, e, f_t, a_t)$ where $\eta_{y,t}$ and $\eta_{o,t}$ are shocks to income and medical spending. The consumer's decisions are given by (c_t, i_t, κ_t). Accordingly, there are three continuous control variables, consumption, investment, and the share of investment in the technology, and a discrete one, participation. There are five state variables. We represent the problem as a series of Bellman equations such that, at each age, the value function has the following form:

$$V(s_t) = \max_{c_t, i_t, \kappa_t} n_{e,t} u(c_t/n_{e,t}) + \beta p_{e,t} \int_\varepsilon \int_{\eta_y} \int_{\eta_o} V(s_{t+1}) dF_e(\eta_o) dF_e(\eta_y) dF(\varepsilon)$$

$$a_{t+1} = \widetilde{R}_\kappa(f_{t+1})(a_t + y_{e,t} + tr_t - c_t - \pi(i_t) - c_d I(\kappa_t > 0)), \ a_{t+1} \geq 0$$

$$f_{t+1} = (1 - \delta)f_t + i_t$$

$$\widetilde{R}_\kappa(f_{t+1}) = (1 - \kappa_t)\overline{R} + \kappa_t \widetilde{R}(f_t).$$

We index variables by e where education differences are assumed to be present. The model

---

[4]Because these expenditures are generally low prior to retirement (and to save on computation time), we allow only for medical expenditure risk after retirement (as in HSZ 1995)

is solved by backward recursion after discretizing the continuous state variables.[5]

3.2 Calibration

To explore the impact of financial education on employee behavior, we assume that $u(c_t/n_t)$ has a CRRA form with relative risk aversion σ for calibration purposes. Here we assume $\sigma = 1.6$, close to the value estimated by Attanasio, Banks, Meghir, and Weber (1999) using consumption data. Following SSK (2006), we define an equivalence scale that accounts for consumption differences in household size by education group and changes in demographics over the life cycle. Assuming that $z(j, k) = (j + 0.7k)^{0.75}$ where j is the number of adults in the household and k is the number of children under age 18, we then define $n_{e,t} = z(j_{e,t}, k_{e,t})/z(2, 1)$ where $j_{e,t}$ and $k_{e,t}$ are the average number of adults and children in the household by age and education group. We use data from the PSID to estimate the time series of average equivalence scales by education group. The age profile of those scales is hump-shaped and more amplified for less-educated households. For the base case, we use a discount factor of 0.96 (as in SSK, 2006, and Campbell and Viceira 2002). The annual minimum consumption floor is set at \$10,000 for a couple with one child.

Post-retirement income is defined to be a function of pre-retirement income, estimated from fixed-effect regressions analyzed separately by education level of net household income on age and a retirement dummy, as in LMM (forthcoming). This produces replacement rates of 0.75 for high school dropouts, 0.74 for high school graduates, and 0.63 for the College+, close to those based on total retirement income in the literature (e.g. Aon Consulting, 2008). Following retirement, we let income decline at the rate estimated in PSID data controlling for educational groups and cohort effects; that pattern is mostly due to changes in household composition (e.g. widowhood).

Turning to the financial market variables, we posit a safe asset return of $\bar{r} = 2\%$ (as in Campbell and Viceira 2002). As the excess return function has not been previously established, we note that the range of risk-adjusted excess portfolio returns reported by von

---

[5]For additional details on the solution method see LMM (forthcoming).

Gaudecker (2011), for example, is -0.017 (5th percentile) to 0.054 (95th percentile). Using Euler equations, Jappelli and Padula (2013) estimate that each point of financial literacy is associated with an expected increase in the return to saving from 0.2 to 1 percent. Clark, Lusardi and Mitchell (forthcoming) use administrative data on 401(k) participants and find that there is about a one percentage point difference in returns between those who have the lowest financial literacy score and those that have the highest. We therefore use a linear function by setting $r_{max} = r(f_{max}) = 0.04$ and $r_{min} = r(f_{min}) = 0$ where 0.04 is chosen to match the equity premium used in the portfolio literature. Below, we choose a convex cost function for investing in financial knowledge, which therefore embodies decreasing returns to producing knowledge. We set $\sigma_\varepsilon = 0.16$ in the simulations (Campbell and Viceira, 2002).[6]

Estimating the price of acquiring financial knowledge is difficult, as little is known regarding inputs to the production process (time and expenditures on financial services), along with investments in, as opposed to, the stock of financial knowledge. As in LMM (forthcoming), we model the process using the function $\pi(i_t) = 50i_t^{1.75}$, a form that posits that the first units of knowledge are inexpensive, while marginal costs rise thereafter. To parametrize the participation cost for the sophisticated technology (c_d), we use the median estimate of $750 (in $2004), following Vissing-Jorgensen (2003). We also require an estimate of the depreciation factor for financial knowledge, δ, though little is known on the size of this parameter. We use a value of 6 percent in our baseline calibration which is consistent with estimates of the depreciation of human capital.

Given this calibration, we can find optimal consumption, financial knowledge investment, and technology participation at each point in the state-space and at each age. Having done so, we then use our decision rules to simulate 2,500 individuals moving through their life cycles. We draw income, out-of-pocket medical expenditure, and rate of return shocks, and we use these to simulate the life cycle paths of all consumers. These consumers are given the initial conditions for education, earnings, and assets derived from the PSID for individuals age 25-30. We initialize financial knowledge at the lowest level (0). A list of the baseline

---

[6]For information on how we estimate income and medical expenditure processes as well as mortality risk by education, see LMM (forthcoming).

parameters and their values is provided in the Appendix.

4 Simulating the Impact of Financial Education Programs

4.1 The Programs

Given the model described above and the parameters of interest, we can evaluate the impact of employer-provided financial education programs on a variety of outcomes, including whether and which employees elect to participate, how much they invest in financial knowledge, their use of the sophisticated technology to invest, and how their lifetime consumption and utility levels change. We let eligibility for a particular financial education program offered in a given year be expressed using a binary variable d_{it}, and in what follows we assume eligibility is assigned randomly to all employees of that given age (which we vary across experimental settings). We model the financial education program as reducing the employee's cost of investing in knowledge. We expressed $\pi_p(i_t) = \vartheta \pi(i_t)$, where $\vartheta < 1$ captures the efficiency of the program. If the program is high quality, it provides an incentive to acquire more knowledge and more employees will then decide to participate in the program. Costs matter as well. For instance, if the program were free, all workers will participate (or be indifferent). In order to capture the fixed time and perhaps money costs of participating in the program, we define ψ as the participation cost for the program.

If the employee is eligible for a program, we define $V_p(s_t)$ $(p = 0, 1)$ as the value (indirect utility) of not participating versus participating, respectively. The individual participates if $v(s_t) = V_1(s_t) - V_0(s_t)$ is greater than zero. We add a zero mean disturbance to this difference, $\zeta_{it} \sim N(0, \sigma_v)$. Hence, participation is given by:

$$p_{it} = I(v(s_t) + \zeta_{it} > 0)$$

In order for ζ_{it} to have the correct scale, we fix σ_v to the standard deviation of the simulated utility differences (0.001).

The simulations to follow explore a number of different programs. First, program eligibility is a function of age, so we evaluate how results change depending on whether the program is provided to employees at age 30, 40, or 50. When a worker is of the targeted age, he is deemed to be eligible with probability 0.5. We also explore how program effectiveness affects outcomes, by varying $\vartheta \in [0.1, 0.5]$. Additionally, we vary the fixed cost of participating (i.e., $\psi \in [250, 500]$). And in a final and very important case, we also allow for the program to affect knowledge depreciation. That is, we posit that the financial education program provides knowledge that does not depreciate over time. This last experiment captures the possibility that a program could provide employees with financial advisers who can be accessed over time or that the program is continued over time. A total of six illustrative scenarios is considered below.

4.2   Who Participates in Financial Education Programs?

To understand who participates and who does not in a workplace financial education program of the sort described here, we first explore employees' participation patterns across various scenarios. Table 1 reports how participation rates in the program vary given (randomly assigned) employee eligibility, where it is clear that participation rates overall (last column) are generally below 100 percent. We emphasize that this is not a sign of program failure; rather, people must incur a cost when investing in knowledge, and knowledge depreciates with time. For both reasons, not everyone will partake of the opportunity to build knowledge. It is also worth noting that program participation rates rise depending when (at which age) program is offered. This is to be expected, since people tend to save most between the ages of 40 and 60; employees have little money to manage earlier in life. Furthermore, we find that program participation is higher for the better-educated, due to the larger gain from investing in knowledge for those individuals. Conversely, the least-educated are less likely to partake of the program offering. As we showed in LMM (forthcoming), the uneducated optimally save less, both as a result of their greater reliance on the social safety net, and their shorter life expectancies. The final two rows of the table indicate how participation rates

for a program offered at a given age, say age 40, vary depending on two factors: program efficiency, and the cost of participation. Logically enough, more efficient programs attract higher participation, whereas higher costs reduce participation.

[Insert Table 1]

In Table 2, we summarize the baseline characteristics of those who elect to participate in a financial education program when offered, versus those who do not (conditional on being eligible at a given age). Results indicate that program participants have higher earnings, more initial knowledge, and more wealth, while nonparticipants are poorer, earn less, and have little financial knowledge at baseline. This selectiveness occurs regardless of the age at which the program is offered. Importantly, it implies that an average program effectiveness measure which assumes that program and nonparticipants could benefit as much as participants will likely be biased.

[Insert Table 2]

The fact that those who optimally elect to undertake the financial education program differ systematically from those who do not underscores the fact that a careful program evaluation must take into account the process by which people endogenously elect into the program. That is, it would be misleading to compare outcomes for program participants versus nonparticipants, since each group has different reasons for their behavior. Moreover, any evaluation program that cannot carefully control the sample's baseline characteristics will be subject to such selection bias. Of course some of these characteristics – e.g., financial knowledge – may be difficult to measure precisely. Nevertheless, unless randomization is available, modeling the selection process is critical.

4.3 The Effect of Financial Education Programs over the Life Cycle

A useful aspect of our simulation approach is that the same simulated respondents are observed in different experimental settings, as they are, in turn, offered different financial education programs. Accordingly, we may compare life cycle investment, wealth, and saving profiles for the same individuals, along with information about whether they did or did not

optimally take part in each program.

Figures 1 to 6 report results, under six different financial education settings, of average profiles of investment in knowledge, stock of knowledge, changes in wealth (in percent), and the share of wealth invested in the sophisticated technology. Specifically, figures 1-3 analyze how results change when the program is offered to a worker at age 30, 40, or 50. Figure 4 reports results for a program offered to a 40-year old employee with an enhanced efficiency parameter, and in Figure 5 we lower the fixed cost of knowledge (shown in the same order as in Tables 1 and 2). Figure 6 illustrates how results change when financial knowledge depreciation is shut down, as for instance when an employer may maintain the employee's financial sophistication post-program via continued monitoring.

[Insert Figures 1 to 6]

A comparison of the first three Figures shows how results change when we vary the age at which the program is implemented. In each case, the upper left-hand panel depicts the impact on investment in financial knowledge, while the impact of the program on the stock of financial knowledge appears in the upper right-hand panel. In the lower left, we report the percentage change in wealth, and on the lower right, the share of the population using the sophisticated investment technology. Each panel includes three lines: the solid line refers to non-enrolled but eligible participants; the dashed line refers to enrolled participants; and the dotted line indicates how participants would have behaved without the program being introduced – a true counterfactual for those who did enroll when they could.

Figure 1 shows what happens with the program is made available to age 30 employees. Those who participate in the program do invest substantially in financial knowledge; this translates into a higher stock of financial knowledge compared to their own (no-program) counterfactual. We also see that those who participate in the program cut back on their investment after the program expires. Along with depreciation in financial knowledge, this leads to a dampening of the program's effect when it is over. Nevertheless, after the initial ramp-up in financial knowledge, the marginal effect on behavior compared to the proper counterfactual is quite small. Conversely, we see that those who do invest in the financial

knowledge program are markedly different from those who do not. In other words, both financial knowledge and sophisticated investment profiles are much higher compared to employees who optimally elect not to participate, underscoring the sample selection concern made earlier. In fact, if one were to compare program participants and nonparticipants, one would (erroneously) conclude that the program had an enormous impact on the stock of financial knowledge, producing a 20 percentage point advantage for participants. Yet the true counterfactual shows that the net effect of a one-year program offered at age 30 is quite small, particularly by the time the worker attains age 65. Results are similar across Figures 1-3, though when the program is implemented on older versus younger workers, the consequences appear slightly larger.

Somewhat larger program effects are evident in Figures 4 and 5. When the program offered becomes more efficacious for a 40-year old employee (Figures 2 versus 4), the employee experiences a much larger bump-up in knowledge which persists for some time, and savings rise detectably. Similar results obtain when the cost of knowledge is reduced (Figure 2 versus 5). Here again, investment in knowledge rises and some persistence in higher savings can be detected.

A much larger and longer-term impact results from shutting down the knowledge depreciation parameter, confirmed by a comparison of Figures 2 and 6. The 40-year old employee offered access to a financial education program whose effects do not decay will average three times more investment in knowledge, which in turn boosts his saving substantially. This effect persists until retirement, underscoring the long-term effect of not only building the knowledge, but also extending it throughout time. In other words, a one-time financial education program may have title effect, as expected, but the long term effects of a persistent financial education program can be sizable.

5 Evaluating Financial Education Programs

Next we use our simulated data to investigate the effect of the programs of interest using the different metrics employed in the financial education literature, as described above. We also evaluate program effectiveness on welfare, measured by changes in lifetime consumption and utility.

5.1 Long-Term Effects

Frequently, empirical researchers may not know when individuals in any given survey may have been exposed to or offered some sort of programs. In the present case, for instance, an employee may not recall whether his employer ever offered a financial education program and if so, when. Nevertheless, in some cases the econometrician may be able to observe wealth at some particular age (e.g., retirement), accompanied with an indicator of whether the person had ever been exposed to such a program earlier in life.[7] This can allow a determination of how offering an educational program affects outcomes of interest. In other cases, one might know which employees elected to take a program, permitting a comparison of outcomes between those who participated and those who did not. Rarely are both available, in practice, and the different outcomes are not directly comparable unless, as shown above, strong assumptions hold about the selection process into the program.

Results in Table 3 illustrate how results differ in our simulated setting where we can measure each of the key employee subgroups. For the six scenarios described earlier, we present four columns of retirement wealth values. The first column summarizes wealth levels for participants who elected to take the program when offered. The second column reports counterfactual wealth for the same people if the program had never been offered. The third column shows wealth levels for nonparticipants – those offered but who declined to participate – and the final column summarizes average wealth for those never offered the program. As before, each row represents a different policy experiment, with a program

---

[7]For instance the Health and Retirement Study has asked older individuals if their employers had offered them workplace-based financial education programs (Lusardi, 2004).

offered at age 30, 40, or 50 (first three rows), or at age 40 and three sets of other parameters comparable with those developed in Figures 1-5.

[Insert Table 3]

Turning to the first row of that Table, program participants held mean wealth at retirement of $524,271. Had they not participated in the program, the same people's mean wealth would have been about 1% lower (and the difference is statistically insignificant). This is the properly measured program effect on those who participated, consistent with Figure 1. In other words, the program did boost both financial knowledge and wealth at the time the employees were offered the program, but by retirement, the effect virtually disappeared.

In the real world, of course, we typically cannot observe the ideal counterfactual; instead, we must find ways to identify a counterfactual and therefore the average effect of the program on the treated. If one could reasonably assume that program participation were independent of wealth, then nonparticipants could be used to measure the counterfactual: the estimated program effect would be to raise retirement wealth by 75% ($225,292/$298,979).

These numbers would lead one to conclude that the program was extremely effective in boosting saving; however, as demonstrated earlier, this is a severely upward-biased metric because participation is correlated with wealth at baseline. Alternatively, we could investigate the effect of offering the program without conditioning on those who participated. Since program eligibility is random in our simulation, everyone who was eligible to elect the program comprises the ITT group. From Table 1, we know that 36% of those offered the program participated, which when combined with data in Table 3, yields an average wealth level of $381,480 for the eligible, versus $392,069 among the ineligible. Surprisingly, then, by this metric, offering the program decreases average retirement wealth by a statistically insignificant 3% (-$10,589/$381,480).

We can do better by recalling that program eligibility is random in our scenarios. Accordingly, we can recover the effect of the program on participants by comparing program participants and non participants. To do so, Imbens and Angris (1994) suggest using the

Wald estimator:

$$\Delta = \frac{E[w_{i,65}|d_i = 1] - E[w_{i,65}|d_i = 0]}{E[p_i|d_i = 1] - E[p_i|d_i = 0]}$$

where $w_{i,65}$ is wealth of respondent i at age 65, d_i denotes eligibility, and p_i participation. The expectation operator is $E[]$. Under certain assumptions, Imbens and Angrist (1994) show that this Local Average Treatment Effect (LATE) effectively captures the effect for a group of individuals who comply with the treatment being offered. Since the ineligible cannot participate, $E[p_i|d_i = 0] = 0$, we have one-sided non-compliance and therefore the effect becomes:

$$\Delta = \frac{E[w_{i,65}|d_i = 1] - E[w_{i,65}|d_i = 0]}{E[p_i|d_i = 1]}.$$

This delivers the average effect of the program on the treated, or the ATET (Imbens and Angrist, 1994). For the first scenario in Table 3, this yields a statistically insignificant change (-$10,589/0.36 = -$29,414), or a 7.4% drop, in percentage terms.

Continuing down the rows in Table 3, it is interesting to note that the largest bias generated by comparing participants and nonparticipants occurs when the program is offered to employees at age 40. At earlier ages, selection is less strong since participants and nonparticipants are more similar and wealth is lower. Later in life, however, the saving motive switches from precautionary to retirement preparation, and behavioral differences are exacerbated. After age 50, these differences again diminish. Since most financial education in the workplace occurs mid-career (around the age of 40), our model suggest that selection can be a major threat to the evaluation of such programs.

It is also of interest that the largest effects occur for most efficient programs provided at low cost. For example, the next to final row in Table 3 (where $\vartheta = 0.25, and \psi = 250$) shows that the true program effect slightly boosts retirement wealth by 1.3% ($6,491/ $472381 which is statistically insignificant). Comparing the ineligible with the eligible groups, we see an apparent negative impact of offering the program (by 3.1%, or -$12,361/395314). The Wald estimator of the effect for those who comply with the offer of the program yields an estimated 4.9 % effect of (-$19684/ $395314 and not statistically significant). In fact,

the only statistically significant effect across all program scenarios evaluated is found in the final row of Table 3, where depreciation has been shut down. This program does increase retirement wealth substantially, by 9% (($467371 - $428874)/$428874), yet that is much smaller than the 1.5 times wealth increment that would result from (incorrectly) using the nonparticipant pool as the comparator group.

To refine these estimates, next we implement these identification strategies in a regression framework which allows us to control for observable differences in outcomes.

5.1.1 Intent-to-Treat

As noted above, the intent-to-treat measure in our setting compares outcomes of those who were program-eligible to those who were not, assuming that program eligibility is exogenous to outcomes. To test this with our simulated data, we implement the following regression which controls for education and average lifetime income:

$$\log w_{i,65} = x_i\beta + \triangle d_i + \epsilon_i.$$

Under random assignment, we have $\epsilon_i \perp d_i$.

Table 4 reports for each of our six scenarios the point estimate of $\triangle$ along with its standard error. In five of the six cases, the program effects are small and statistically insignificant, ranging from -0.06 to 0.1236. This confirms the unconditional levels estimates we reported in Table 3. By contrast, the program effect is positive and statistically significant for the final experiment, where financial knowledge is preserved through time. The estimate suggest an effect of 30% with a standard error of 12%. Controlling for covariates yields an even larger ITT estimate.

[Insert Table 4]

5.1.2   OLS on Program Participation

Thus far, we have argued that, due to selection bias, comparisons of participants and non-participants do not identify the effect of the program on outcomes. But one might wonder whether this could be remedied by controlling for factors observed sometimes early in life, say at age 25. Since financial knowledge is zero at age 25, there are two exogenous outcomes on which we could condition: wealth at age 25, and average lifetime income (in addition to the education dummies). To evaluate this, we run the following OLS regression:

$$\log w_{i,65} = x_i\beta + \triangle p_i + \epsilon_i$$

This delivers the average effect of the program on the treated, if $\epsilon_i \perp p_i | x_i$. Table 5 reports the new point estimates of $\triangle$ along with their standard errors.

Results in Table 5 show that when a financial education program is offered early in life, such as at age 30, baseline controls can sufficiently correct for selection, since estimated effects are close to zero. After that, however, the controls and functional form are insufficient to control for biases imparted by endogenous selection. In other words, the estimated effect of participating in the program becomes large and statistically significant when the program is offered to older workers. This is mainly due to the fact that incentives to save and, thus, acquire knowledge are a function of the income, rather than simply its level.

[Insert Table 5]

5.1.3   Local Average Treatment Effects

The Wald estimator can also be implemented as an instrumental variables (IV) regression (Imbens and Angrist, 1994). In our case, the first-stage regression for participation is:

$$p_i = x_i\alpha + \eta d_i + \nu_i$$

assuming that eligibility is independent of ϵ_i. Results are reported in Table 6 along with standard errors. Our findings confirm that programs which do not affect depreciation have

little effect on retirement-age wealth levels. Although the point estimates are generally positive, the standard errors are often large. Only in the final scenario where the average effect on the treated (true) is positive does the LATE estimator pick up the effect and does the estimate become statistically significant. Accordingly, this IV estimator is a proper estimator of the average treatment effect on the treated (ATET) when eligibility or assignment to treatment is random.

[Insert Table 6]

5.2 Contemporaneous Effects

Several evaluations of financial education programs compare the same individuals prior to and after receiving the training. When the same is done for a control group, one can implement a difference-in-difference (DD) strategy of the following form:

$$\log w_{it} = \mu_i + \lambda_t + \triangle z_{it} + x_{it}\alpha + \epsilon_{it}$$

for $z_{it} = (p_{it}, d_{it})$. Identification of the average effect requires that $\epsilon_{it} \perp z_{it} | x_{it}, \lambda_t, \mu_i$. The common-trend assumption imposes that, in the absence of the program, the average change in wealth of those who participate ($z_{it} = 1$) would have been the same as for those not participating ($z_{it} = 0$). We can estimate this equation using fixed-effect regression using either p_{it} or d_{it}. As described above, estimates of $\triangle$ capture both the ATET and the ITT effects.

To implement this approach in our simulated data, we consider two periods: one year prior to the program, and five years after the program. Since we can directly compute the average effect of the program on those who participated (using the true counterfactual), we also report this estimate in column 4 of Table 7. We find that the true effect of the program on those who participate is generally small, except when the program is highly effective. Using non participants as the counterfactual (hence implementing DD with p_{it}) yields generally large and positive effects. The key explanation for why these estimates are

biased is that the common-trend assumption in fact does not hold for participants and non participants. That is, participants in financial education programs in our scenario would save more in the absence of the programs, compared to non participants. For this reason, using the trend on wealth of nonparticipants as a counterfactual grossly overestimates the effect of the programs. Implementing DD with eligibility yields relatively smaller biases, compared to using participation.

[Insert Table 7]

6 Discussion and Conclusions

In previous research we have demonstrated that important segments of the population are financially unsophisticated and do not understand simple interest, inflation, and risk diversification (Lusardi and Mitchell 2008, 2011a,b). We have also shown that it is actually optimal for many people to be unsophisticated, in that some people will rationally elect not to invest in knowledge as it is expensive to acquire and does not benefit everyone (LMM forthcoming). The present paper goes farther by using our theoretical model to evaluate the impacts of well-specified financial education programs that could be offered by employers to workers of different ages. In particular, we use our stochastic life cycle model incorporating endogenous knowledge accumulation to evaluate six different financial literacy program scenarios. This is useful since no empirical studies have the kind of information needed to capture precisely what the impact of the interventions will be. In our case, we know all relevant counterfactuals to compare "true" outcomes with program effectiveness estimates generated by conventional econometric techniques.

Our approach provides several important insights regarding financial education program evaluation. First, we show that low participation rates in such programs can be rational, once we recognize that improving financial literacy does not benefit everyone and acquiring knowledge is costly. In particular, the low-income and less-educated have less to gain from participating in such programs. For this reason, it is incorrect to conclude that financial

education programs are not valued and "preach only to the converted." Rather, the decision to invest in financial education depends on its costs and benefits, factors which differ across individuals. Second, our model emphasizes the role of self-selection in financial education, particularly at older ages. Accordingly, great care is required to rigorously evaluate the effectiveness of financial education in non-experimental settings, where self-selection tends to occur. Third, prior studies have taken too narrow a focus by overlooking the crucial role of knowledge retention, once the financial education is obtained. That is, financial education delivered to employees around the age of 40 can raise savings at retirement by close to 10%, if the knowledge gained can be maintained. Fourth, and relatedly, we show that short-term financial education programs are unlikely to dramatically alter saving, especially when offered to young people. They are more effective when targeted at peak saving years (e.g., post-age 40).

A final important lesson from our work is to point out how measures of financial education program effectiveness shape outcomes across heterogeneous individuals so that evaluators build several key elements into the study design. First, it is essential to have accurate measures are of what information the program delivers and what sort of follow-up is provided. Second, the researcher must measure baseline features of the eligible sample including wealth, income, and financial literacy. Third, it is necessary to randomize eligibility for the treatment. And fourth, longer-term follow-up is crucial.

References

Aon Consulting. 2008. *Replacement Ratio Study: A Measurement Tool for Retirement Planning.* Global Corporate Marketing and Communications (www.aon.com/retire).

Attanasio, O., J. Banks, C. Meghir, and G. Weber. 1999. "Humps and Bumps in Lifetime Consumption." *Journal of Economic and Business Statistics.* 17(1), 23-35.

Bernheim, D., and P. Bayer and J.K. Scholz. 2009. "The Effects of Financial Education in the Workplace: Evidence from a Survey of Employers." *Economic Inquiry.* 47(4): 605-24.

Bernheim, D., D. Garrett, and D. Maki. 2001. "Education and Saving: The Long-term Effects of High School Financial Curriculum Mandates." *Journal of Public Economics.* 80: 435–565.

Bernheim, D., and D. Garrett. 2003. "The Effects of Financial Education in the Workplace: Evidence from a Survey of Households." *Journal of Public Economics.* 87: 1487–1519

Campbell, J. Y. and L. Viceira. 2002. *Strategic Asset Allocation.* Oxford: Oxford University Press.

Clark, R., M.d'Ambrosio, A. McDermed, and K. Sawant. 2006. "Retirement Plans and Saving Decisions: The Role of Information and Education," *Journal of Pension Economics and Finance*: 45-67.

Clark, R., A. Lusardi, and O. S. Mitchell. Forthcoming. "Financial Knowledge and 401(k) Investment Performance." *Journal of Pension Economics and Finance.* [NBER Working Paper 20137].

Clark, R., M. Morrill, and S. Allen. 2012. "Effectiveness of Employer-Provided Financial Information: Hiring to Retiring." *American Economic Review,* May: 314-31

Collins, J. M. and C. M. O'Rourke. 2010. "Financial Education and Counseling: Still Holding Promise." *Journal of Consumer Affairs.* 44(3): 483-489.

Deuflhard, F., D. Georgarakos, and R. Inderst. 2013. "Financial Literacy and Savings Account Returns." Goethe University of Frankfurt Working Paper.

Hubbard, R. G., J. Skinner, and S. Zeldes. 1995. "Precautionary Savings and Social Insurance." *Journal of Political Economy.* 103(2): 360-399.

Imbens, G., and J. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects", *Econometrica,* 62, 467-475.

Imbens, G. and J. M. Wooldridge. 2009. "Recent Developments in the Econometrics of Program Evaluation." *Journal of Economic Literature.* 47(1): 5-86.

Imbens, G. 2010. "Better LATE Than Nothing: Some Comments on Deaton (2009) and Heckman and Urzua (2009)", *Journal of Economic Literature*, 48(2) , 399-423.

Jappelli, T., and M. Padula. 2013. "Investment in Financial Knowledge and Saving Decisions." *Journal of Banking and Finance*. 37: 2779-2792.

Lusardi, A. 2004. Saving and the Effectiveness of Financial Education," in O. S. Mitchell and S. Utkus (eds.). *Pension Design and Structure: New Lessons from Behavioral Finance*. Oxford: Oxford University Press: 157-184

Lusardi, A., P. Keller, and A. Keller. 2008. New Ways to Make People Save: A Social Marketing Approach, In *Overcoming the Saving Slump: How to Increase the Effectiveness of Financial Education and Saving Programs*. Ed. A. Lusardi. Chicago: University of Chicago Press: 209–236.

Lusardi, A., P-C. Michaud, and O. S. Mitchell. Forthcoming. "Optimal Financial Literacy and Saving for Retirement." *Journal of Political Economy*. [NBER Working Paper 18669]

Lusardi, A. and O. S. Mitchell. 2014. "The Economic Importance of Financial Literacy: Theory and Evidence." *Journal of Economic Literature*. 52(1): 5-44.

Lusardi, A. and O. S. Mitchell. 2011a. "Financial Literacy and Planning: Implications for Retirement Wellbeing." In *Financial Literacy: Implications for Retirement Security and the Financial Marketplace*. O. S. Mitchell and A. Lusardi, eds. Oxford: Oxford University Press: 17-39.

Lusardi, A. and O. S. Mitchell. 2011b. "Financial Literacy around the World: An Overview." *Journal of Pension Economics and Finance*. 10(4): 497-508.

Lusardi, A., and O. S. Mitchell. 2009. "How Ordinary Consumers Make Complex Economic Decisions: Financial Literacy and Retirement Readiness." NBER Working Paper 15350.

Lusardi, A., and O. S. Mitchell. 2008. "Planning and Financial Literacy: How Do Women Fare?" *American Economic Review* 98: 413–417.

Mandell, L. 2008. "Financial Education in High School." In *Overcoming the Saving Slump: How to Increase the Effectiveness of Financial Education and Saving Programs*. A. Lusardi, ed. Chicago, IL: University of Chicago Press: 257-279.

Organization for Economic Cooperation and Development (OECD). 2005. *Improving Financial Literacy. Analysis of Issues and Policies*. Geneva: OECD Publications.

Richardson, D. and J. Seligman. 2014. Financial Literacy Among Graduate Students. A Case Study from the Indiana University System. Unpublished paper, TIAA-CREF Institute.

Scholz, J., A. Seshadri, and S. Khitatrakun. 2006. "Are Americans Saving Optimally for Retirement?" *Journal of Political Economy*. 114(4): 607-643.

Vissing-Jorgensen, A. 2003. "Perspectives on Behavioral Finance: Does 'Irrationality' Disappear with Wealth? Evidence from Expectations and Actions." *NBER Macroeconomics Annual*.18(1): 139-194.

von Gaudecker, H. M. 2011. "How Does Household Portfolio Diversification Vary with Financial Sophistication and Advice?" Netspar Discussion Paper 2011-026.

Walstad, W, K. Rebeck and R. MacDonald. 2010. "The Effects of Financial Education on the Financial Knowledge of High School Students." *Journal of Consumer Affairs* 44(2): 336–357.

Figure 1: **Effects of the Financial Education Program over the Life-Cycle**: Intervention at age 30 with $\vartheta = 0.5$ and $\psi = 500$. We plot the average age profile of investment in knowledge, stock of knowledge, percent change in wealth, and the share of wealth invested in sophisticated products by participation status. For those who participated, we also plot the age profile had they not participated in the program.

Figure 2: **Effects of the Financial Education Programs over the Life-Cycle**: Intervention at age 40 with $\vartheta = 0.5$ and $\psi = 500$. We plot the average age profile of investment in knowledge, stock of knowledge, percent change in wealth, and the share of wealth invested in sophisticated products by participation status. For those who participated, we also plot the age profile had they not participated in the program.

Figure 3: **Effects of the Financial Education Programs over the Life-Cycle**: Intervention at age 50 with $\vartheta = 0.5$ and $\psi = 500$. We plot the average age profile of investment in knowledge, stock of knowledge, percent change in wealth, and the share of wealth invested in sophisticated products by participation status. For those who participated, we also plot the age profile had they not participated in the program.

Figure 4: **Effects of the Financial Education Programs over the Life-Cycle**: Intervention at age 40 with $\vartheta = 0.25$ and $\psi = 500$. We plot the average age profile of investment in knowledge, stock of knowledge, percent change in wealth, and the share of wealth invested in sophisticated products by participation status. For those who participated, we also plot the age profile had they not participated in the program.

Figure 5: **Effects of the Financial Education Program over the Life-Cycle**: Intervention at age 40 with $\vartheta = 0.25$ and $\psi = 250$. We plot the average age profile of investment in knowledge, stock of knowledge, percent change in wealth and the share of wealth invested in sophisticated products by participation status. For those who participated, we also plot the age profile had they not participated in the program.

Figure 6: **Effects of the Financial Education Program over the Life-Cycle**: Intervention at age 40 with $\vartheta = 0.1$ and $\psi = 100$ and no depreciation of knowledge among participants to the program. We plot the average age profile of investment in knowledge, stock of knowledge, percent change in wealth, and the share of wealth invested in sophisticated products by participation status. For those who participated, we also plot the age profile had they not participated in the program.

age	ϑ	ψ	less HS	HS	college	total
30	.5	500	.2349	.3302	.4453	.3662
40	.5	500	.3571	.4428	.5287	.4677
50	.5	500	.3438	.5109	.5792	.517
40	.25	500	.4048	.5084	.6169	.5405
40	.25	250	.4762	.6004	.6877	.6206
40	.1	100	.6131	.7899	.8429	.7882

Table 1: **Program Participation:** We report participation rates to the program among those eligible for a series of scenarios and for three education levels. Age refers to the time at which the program is implemented, ϑ is the relative marginal cost of investing in knowledge in the program, and ψ is the fixed cost of participating in the program.

age	ϑ	ψ	income (np)	fin (np)	wealth (np)	income (p)	fin (p)	wealth (p)
30	.5	500	34182	2.577	18742	55559	12.67	52396
40	.5	500	39939	13.86	35747	69325	43.46	111863
50	.5	500	49104	36.79	143971	66180	61.48	189600
40	.25	500	36277	10.31	27705	68482	42.49	108452
40	.25	250	36171	10.04	29141	64388	38.5	97148
40	.1	100	32171	9.837	34501	59463	32.5	81246

Table 2: **Characteristics of Participants and non participants:** We report means of baseline characteristics (income, financial knowledge, and wealth) for participants (p) and non participants (np). Age refers to the time at which the program is implemented, ϑ is the relative marginal cost of investing in knowledge in the program, and ψ is the fixed cost of participating in the program.

age	ϑ	ψ	participants	counterfactual	non-participant	non-eligible
30	.5	500	524271	522186	298979	392069
40	.5	500	519852	517146	253147	395314
50	.5	500	442676	440577	312325	394461
40	.25	500	515385	508452	220050	395314
40	.25	250	478872	472381	221023	395314
40	.1	100	467371	428874	187233	395314

Table 3: **Wealth at Retirement by Groups:** We report mean wealth at retirement (age 65) for those who participate in the program, mean wealth for those who participate had they not participated (counterfactual), non participants among those eligible and finally those not eligible. Age refers to the time at which the program is implemented, ϑ is the relative marginal cost of investing in knowledge in the program and ψ is the fixed cost of participating in the program.

age	ϑ	ψ	ITT	se
30	.5	500	.1141	.1129
40	.5	500	.1081	.1128
50	.5	500	-.02788	.113
40	.25	500	.1223	.1127
40	.25	250	.1236	.1127
40	.1	100	.2907	.1116

Table 4: **Effect of Offering Financial Education Program on Wealth at Retirement (Intent-to-Treat):** We report for each program the intent-to-treat estimate of the program along with standard error. This estimate is obtained by regressing log wealth at retirement on eligibility to the program and controls for education and average lifetime income.

age	ϑ	ψ	OLS	se
30	.5	500	-.005068	.1512
40	.5	500	.4073	.1367
50	.5	500	.672	.13
40	.25	500	.5023	.1302
40	.25	250	.4428	.1241
40	.1	100	.7365	.1146

Table 5: **Effect of Financial Education Program Participation on Wealth at Retirement (OLS)**: We report for each program the estimate of the effect of the program along with standard error. This estimate is obtained by regressing log wealth at retirement on participation to the program and controls for education, average lifetime income and initial wealth (at age 25).

age	ϑ	ψ	LATE	se
30	.5	500	.3102	.3069
40	.5	500	.2259	.235
50	.5	500	-.05387	.2183
40	.25	500	.2207	.2028
40	.25	250	.1964	.1786
40	.1	100	.3669	.1399

Table 6: **Effect of Financial Education Program Participation on Wealth at Retirement (LATE-IV)**: We report for each program the estimate of the local average treatment effect along with standard error. This estimate is obtained by instrumental variables regression of log wealth at retirement on participation to the program and controls for education and average lifetime income. The instrumental variable is eligibility to the program.

age	ϑ	ψ	counterfactual	non-participant	non-eligible
30	.5	500	-.01929	.7243	.4537
40	.5	500	.005797	.304	.1558
50	.5	500	.003612	-.01652	.009615
40	.25	500	.02136	.394	.1825
40	.25	250	.04464	.288	.1265
40	.1	100	.1142	.4037	.17

Table 7: **Difference-in-Difference Effect of Financial Education Program on Wealth:** We report estimates of the effect of the financial education program on wealth (in percent) 5 years after the program, relative to one year prior to the program. This is done using 3 potential counterfactuals. The first uses outcomes of those treated had they not participated (average effect on the treated). The second and third columns use different counterfactuals. The second uses non participants (but eligible). The last column uses those not eligible.

Appendix

Parameter	Value
σ	1.6
β	0.96
$\bar{r}$	0.02
$r(f_{max})$	0.04
π_0	50
π_1	1.75
c_d	750
δ	0.06
c_{min}	10,000

Table A.1: **Baseline Parameter Values.** Baseline values are as follows: relative risk aversion ($\sigma = 1.6$), financial knowledge depreciation rate ($\delta = 0.06$), investment production function ($\pi(i) = 50i^{1.75}$), participation cost ($c_d = 750$), discount factor ($\beta = 0.96$). The cost of investing in knowledge takes the form $\pi(i) = \pi_0 i^{\pi_1}$. See text.