CIRANO
Allier savoir et décision

Nudge vs. Financial Literacy in a Retirement Savings Laboratory Experiment

DINA **TASNEEM**
AUDREY **AZEROT**
MARINE **DE MONTAIGNAC**
JIM **ENGLE-WARNICK**

2018S-24
WORKING PAPER
CS

2018s-24

Nudge vs. Financial Literacy in a Retirement Savings Laboratory Experiment

Dina Tasneem, Audrey Azerot,
Marine de Montaignac, Jim Engle-Warnick

Série Scientifique
Scientific Series

Montréal
Juillet/July 2018

CIRANO
Allier savoir et décision

Centre interuniversitaire de recherche en analyse des organisations

ISSN 2292-0838 (en ligne)

Nudge vs. Financial Literacy in a Retirement Savings Laboratory Experiment[*]

Dina Tasneem[†], *Audrey Azerot*[‡], *Marine de Montaignac*[§], *Jim Engle-Warnick*[**]

Résumé/Abstract

We report results from an economics experiment that examines the role of financial literacy in retirement savings. In the experiments, participants make decisions in a retirement savings game, in which income during working years is uncertain. Participants are nudged to varying degrees with automatic savings in each period of the game. Some participants receive financial literacy training in the form of training to compute the expected savings needed at retirement to smooth consumption over the entire life cycle. We find evidence that literacy increases savings and improves efficiency. Our finding has implications for choice architecture for retirement savings.

Mots clés/Keywords: Precautionary Savings; Retirement Savings; Life-cycle Model; Dynamic Optimization; Nudge; Financial Literacy; Decision Heuristics

Codes JEL/JEL Codes: C90

---

[*] Acknowledgements: We acknowledge the Autorité des marchés financiers for funding. Corresponding author: Jim Warnick; jim.engle-warnick@mcgill.ca.
[†] American University of Sharjah.
[‡] McGill University.
[§] CIRANO.
[**] McGill University and CIRANO.

1 Introduction

The increased access to financial markets and conversion to defined contribution pension plans have put more responsibility for decision-making into the hands of savers. This responsibility can be addressed in at least three ways: getting advice, being nudged, or getting more financially literate.

The financial advice industry has grown substantially since 1980, with questions surrounding the performance of advice (Malkiel, 2013) and the extent to which advice is consistent with client heterogeneity (Foerster, Linnainmaa, Melzer, and Previtero, 2017). Nudge (Thaler and Sunstein, 2008; Benartzi and Thaler, 2013), or the attention to "choice architecture", has also attracted the attention of policy makers who use it to structure the decision-making environment so that increasing savings is made easy, given known and predictable biases of decision-makers. Financial literacy, on the other hand, is more related to autonomous or informed decision-making. Financial literacy is important not just for independent decision-making, but also for evaluating advice and being aware of and responding to nudges, thus in a sense it encompasses issues surrounding savings and investment decisions.

In this paper we address the effect of literacy training on savings decisions in a laboratory experiment. Studying the path from literacy training to decision-making in real life is difficult in the field for several reasons (Lusardi and Mitchell, 2014). For one, literacy is an investment in human capital, which means that not everyone should invest, and that not everyone who should invest should ultimately take an action based on the training. For another, it is unclear how financial literacy is developed. In addition, the quality of training can be difficult to assess, short term interventions must be targeted to specific needs, and it seems unlikely that short-term training would have identifiable long-term benefits.

We overcome several of these difficulties with a laboratory experiment, in which we present a game with only precautionary and retirement savings motives. We identify a salient and specific need for literacy: the computation of the expected amount of money needed at retirement age to smooth consumption over the life cycle. We nudge participants with and without financial literacy, and we test the effect of literacy training on their behavioral responses to the nudge.

Our experiment consists of two basic games. In the first game, participants make decisions in a twenty-period precautionary savings game. In each period there is a constant and independent 50/50 chance of earning either a high or a low income. Participants choose how much to consume and how much to spend. The consumption is turned into points through a CRRA utility function inducing the motive to smooth. The game is finite, borrowing is not permitted, and savings generates no interest income. The savings motive in this game is purely precautionary.

In the second game we add five periods to the end of the precautionary savings game, with no income, and call this the retirement period. Everything else is the same in this twenty-five period game. Thus, the savings motives in this case are both precautionary and retirement. Notice that the difference between the two games is retirement savings.

For the nudge treatments, we automatically place 0%, 20% or 30% of pay into savings (equivalently, cash-in-hand) at the beginning of each period of the retirement savings game. The 0% contribution is intended to be the treatment without nudge. The remaining treatments nudge participants about right and too much in expectation. Note that the constant savings rule that results from the nudge is not optimal, so there is scope for decision-making in all three nudge treatments.

For the literacy treatments, we run a financial literacy training module before the retirement savings game. The training module explains how to compute, in expectation, the amount of funds needed for retirement to smooth consumption across all twenty-five periods of the retirement savings game. The module takes the participant step-by-step through the process of computation, and requires correct answers from multiple choice questions before continuing to the next step.

Our precautionary savings game is identical to the on used in Tasneem and Warnick (2018). It is similar to Ballinger, Palumbo and Wilcox (2003) and Ballinger, Hudson, Karkovlata and Wilcox (2011) who reported results from social learning and cognitive ability on the precautionary game respectively. More broadly, Hey and Dardanoni (1988), Carbone and Hey (2004) and Carbone (2006) document heterogeneity in behavior in savings games. Taking a step back from the precautionary model, Carbone and Duffy (2014) report results from a deterministic life-cycle consumption optimization problem, and Zhikang, Chua and Camerer (2009) test for explanations for under-saving in life-cycle models. Our experiment adds the retirement savings motive to the precautionary game. We are not aware of a laboratory experiment that tests the effect of literacy training in the presence of nudging for retirement savings.

We find that only the large nudge increases savings compared with no nudge. Both the 0% and 20% nudge result in under-saving for retirement, and inefficient consumption smoothing. In fact, the 20% nudge results in higher consumption volatility than no nudge at all with the same savings level. The 30% nudge treatment shows significantly more saving. When financial training is included, the savings levels in the 0% and 20% nudges improve to a comparable level in the 30% nudge, and lower volatility in the 20% nudge is restored. On the other hand, literacy training decreases consumption volatility in the large nudge. We identify decision rules most likely responsible for these results.

In other words, we find that in the case of low nudge, literacy training can increase savings without improving consumption volatility, a result favorable to governments and financial institutions but neutral towards savers. In the case of high nudge, literacy training can reduce consumption volatility without affecting savings levels, which is a useful result for savers and neutral towards governments and financial institutions. It is thus important to take the interaction between nudge and literacy into account with choice architecture.

Overall literacy training, which was specifically targeted at a particular need, which was applied uniformly to all participants, and which was relevant to all participants, increased retirement savings and improved consumption efficiency. We conclude that choice architecture that consists of a combination of nudge and literacy can help to ensure that decision-making improves in our retirement savings game.

We introduce the experimental design in the next section, followed by experimental procedures, results, and the conclusion.

2 Experimental Design

2.1 The Model

The model is a finite time forward-looking intertemporal consumption problem, with an uncertain income in each period, and an incentive to smooth consumption (Ballinger et al., 2003). There is both a precautionary savings version of the game and a retirement savings version of the game.

In the twenty-period precautionary savings game, the income stream is given by $y = (y_1, y_2, ...y_{20})$, where each y_t takes on a low or a high value with equal probability at the beginning of each period. The decision in each period is simply how much money to save and how much to use for consumption, where the precautionary savings motive is induced by an incentive to smooth consumption over the lifespan. For simplicity the agent cannot borrow and does not earn interest on savings. In the retirement savings problem, all parameters are identical, but an additional five periods corresponding to retirement are added such that $y_{21} = y_{22} = ... = y_{25} = \$0$.

Let the instantaneous utility of consumption in period t be $u(c_t)$, the accumulated asset at the beginning of period t be A_t and the uncertain labour income realized at the beginning of each period be y_t. In general, utility is discounted at a constant rate β. During the T period life cycle the agent's objective is to choose c_s at each period $s = 1, 2, 3, ..., T$ to maximize the expected sum of discounted utility :

$$E_s \sum_{t=s}^{T} \beta^{(t-s)} u(c_t)$$

subject to the intertemporal budget constraint

$$A_{t+1} = A_t + y_t - c_t$$

where

$$A_t \geq 0 \ \forall t.$$

Utility in period t is given by a CRRA utility function, where the convex marginal utility along with a strict borrowing constraint creates a precautionary savings motive:

$$u(c_t) = k + \theta \frac{(c_t + \epsilon)^{(1-\sigma)}}{1-\sigma}.$$

As in Ballinger et al. (2003) the utility function has several parameters: ϵ is a flow of consumption that is independent of c_t, σ is the coefficient of relative risk aversion, and k and θ are scaling parameters needed to simulate the model in the laboratory. Note also that utility is scaled by an exchange rate of 0.16 to scale the experimental cash earnings in currency.

In this finite horizon model, the optimal consumption rule is a function of "cash-in-hand", $X_t = A_t + y_t$, and time, which can be denoted by $c^*(X_t, t, T)$.[1] In fact, the relationship between consumption and cash-in-hand is not a constant fraction in any certain period. Roughly speaking, the marginal propensity to save is increasing in cash in-hand (Deaton, 1992), and if the cash-in-hand goes below a critical value the consumer should spend everything. The optimal policy must be computed by solving the constrained maximization problem numerically using backward recursion that starts with finding c_T^*, given the terminal value function. Following that step, c_t^* for t = T-1: -1 :1 are derived successively in backward recursive steps (Miranda and Fackler (2002)).

[1] When the horizon is infinite the optimal consumption rule is a function of cash-in-hand and depends on the discount factor (Deaton,1992).

Table 1: Experimental Parameters

Treatment	k	θ	ϵ	σ	Income	Pr of low Income	Starting C-In-H	Retirement Period	T
Precautionary	10.105	476.19	2.7	3	3 or 9	0.5	6	0	20
Retirement	10.105	476.19	2.7	3	3 or 9	0.5	6	5	25

3 Experimental Procedures

In the experiments $y(t)$ took on a value of either \$9 or \$3 with equal probability, constant and independent of the previous period, for the first twenty periods of both the precautionary and the retirement savings games. Participants realized their income and then decided how much of their cash-in-and to spend and how much to save. The spending was transformed into consumption (i.e., cash payment for participation) by a CRRA utility function, and the savings was carried over (without interest) to the next period.[2] The utility function provides the motive to smooth consumption. The retirement savings game tacked on five extra periods of life with no income. Since the motive is to smooth consumption over the life cycle, this game adds a retirement savings motive to the precautionary savings game. Table 1 presents a summary of the design parameters of both games.

We drew the income streams before the experiment and presented identical draws to all subjects.[3] Figure 1 shows the income streams (black dot - identical for both precautionary and retirement savings), the optimal consumption policy for precautionary savings (green square), and the optimal consumption policy for retirement savings (red diamond). Spending is lower in the retirement treatment as agents are forced to save for retirement with the same total lifetime income. Notice that when retirement savings is an issue, optimal spending is

[2] We chose the same relative risk aversion parameter as in Ballinger et al. (2003), $\sigma = 3$. Note also that the there is no time discounting, i.e., $\beta = 1$.

[3] In Tasneem and Warnick (2018), in a more thorough study of the savings game, there were three distinct income sequences: one income sequence resulted in a majority of high income draws in the first ten periods, one resulted in a majority of low income draws in the first ten periods, and one resulted in a representative number of high and low draws in the first ten periods. We used the representative draw for this on-line study.

Figure 1: Income and Optimal Choice Histories for Precautionary and Retirement Savings Games

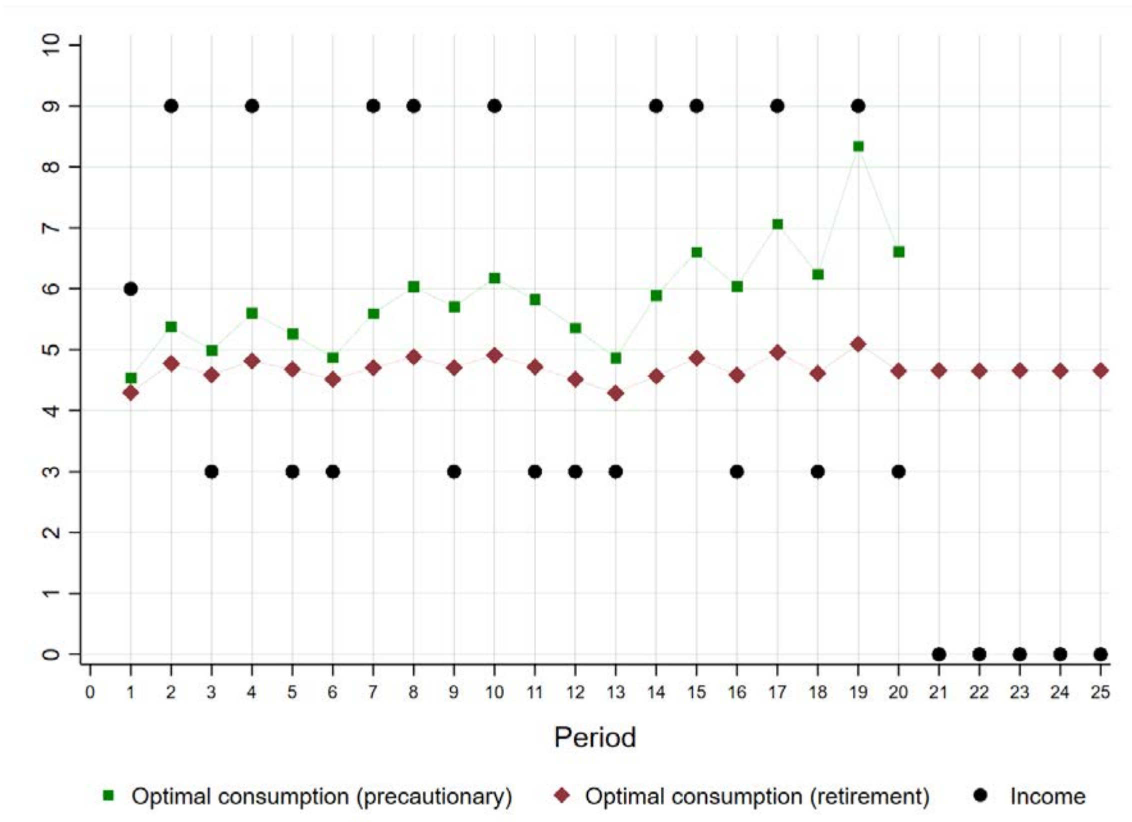

actually smoother across the total lifetime. Also notice that optimal retirement spending is mechanical at the end of a finite life, spending one-fifth of the cash-in-hand available at period twenty-one across the final five periods.

To test the effect of nudging retirement savings, we constructed three experimental treatments. The nudges, which automatically placed either 0%, 20% or 30% of period income into cash-in-hand, occurred at the beginning of each period of the retirement savings game, before the participant made her savings/consumption decision. A field on the screen showed how much money was currently in savings. Given the parameters of the experimental design, 20% is, on average, a close approximation to the the expected savings needed at retirement age to smooth consumption (it is not, however, an optimal strategy). Thus, every decision was made in deviation from one of the three automatic savings decisions, one of which is a baseline, one on average the right amount to save, and one too much.

To test the effect of financial literacy on the nudges, we added a training module before the retirement savings game that walked the participants step-by-step through the process of computing the expected amount of savings needed at the end of period twenty to be able to smooth consumption over the entire twenty-five periods of the game. The steps involved computing expected period income, then computing expected life-time income, then distributing the life-time income over the full twenty-five periods so that the amount needed for five periods would be known. At each step the participant was required to answer a multiple choice question (until they reported the correct answer) to ensure that they experienced the correct calculation.

The game was played in the experimental laboratory by a total of one-hundred eighteen participants, between the ages of nineteen and sixty-six. Approximately 51% of the participants were male and 49% were female, and the average age was thirty-three. After consenting to the study, participants were led through the experimental instructions. They were given a brief quiz to ensure that they understood the concept of income smoothing.

Savings (equivalently, consumption) decisions were made using a slider, which depicted consumption on the left and resulting cash-in-hand on the right. The experimental cash payoff to date was shown on the screen as the sum of spending decisions transformed into consumption through the CRRA utility function. There was a field that showed the amount of cash in savings. Participants played a minimum of three precautionary savings games for no pay for practice (they were permitted to play additional practice game if they wished), followed by an instruction screen that presented the retirement savings games as identical with the exception of the five extra periods. They then played a minimum of three retirement savings games with no pay for practice, followed by the retirement savings game for pay. Participants were paid in cash the sum of their consumption for the two games they played for pay.

3.1 Behavioral Hypotheses

Since the incentive is to smooth consumption, performance in these game is perhaps most appropriately measured in terms of consumption variance compared with the variance of the income stream. However, nudge is not oriented towards variance directly, as it is intended to induce more savings. Financial literacy training, on the other hand, should improve decision-making.

Conjecture 1: Nudge increases savings, which indirectly reduces consumption variance.

Conjecture 2: Financial literacy reduces consumption variance through the computation of expected savings needed for retirement.

4 Experimental Results

Figures 2 presents average consumption, aggregated over treatments, over the twenty-period precautionary savings game. The figure also includes the time path of optimal consumption (solid line). From the figure, participants were able to smooth their consumption to a significant extent: minimum consumption for a period was approximately $3.5 and maximum consumption was between $7 and $8 (recall that income varied between $3 and $9). There is room for improvement, as the optimal consumption path is smoother than the actual path. There is no nudge or financial literacy training in the precautionary savings game.

Figure 3 presents average consumption, broken out by treatment, over the twenty-five period retirement savings game. Both halves of the figure also include the time path of optimal consumption (solid line). Figure 3 shows that participants were able to improve smoothing in the retirement savings game, as theory predicts, in both treatments. As with the precautionary savings treatment, there is less smoothing than with the optimal consumption path.

Figure 2: Average consumption - precautionary savings game

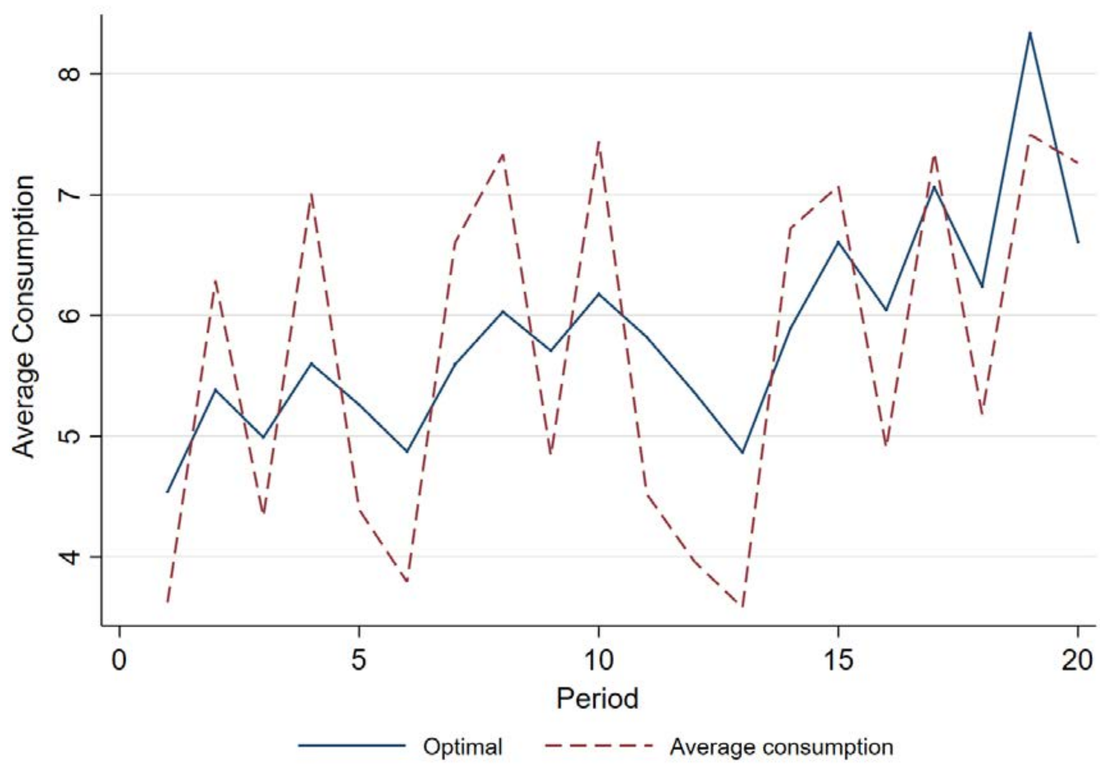

Figure 3: Average consumption by nudge and literacy training - retirement savings game

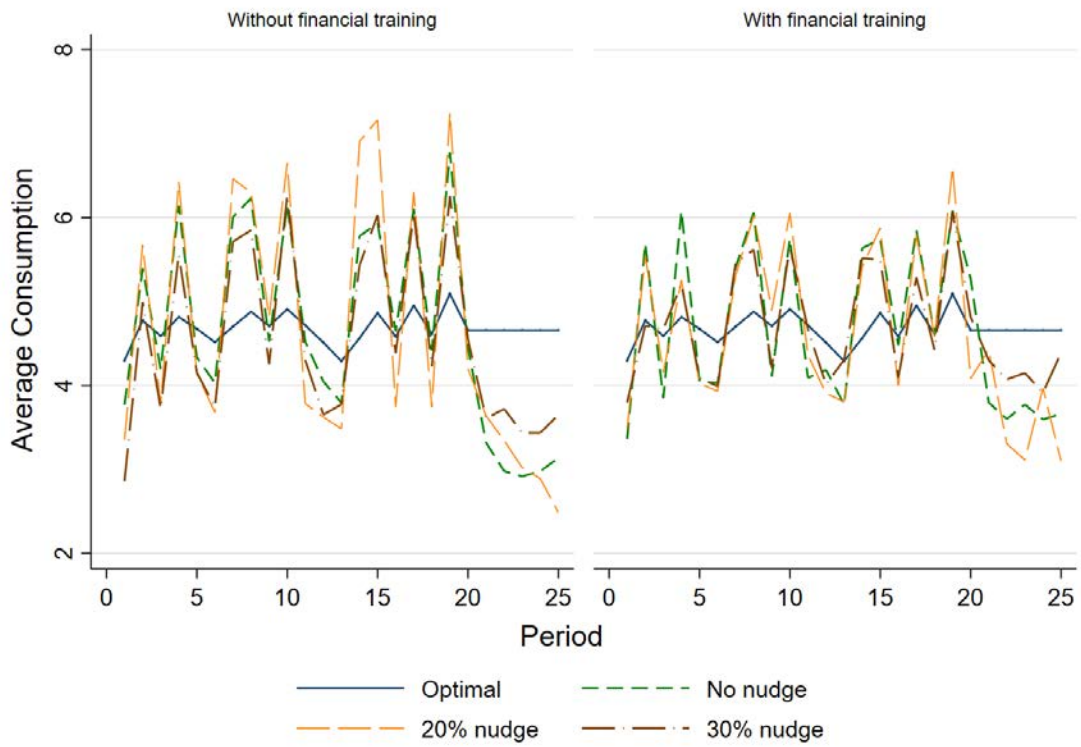

There appears to be an effect from financial training, with a reduced consumption variance evident in the right-hand panel of the figure.

To get an alternative look at the same behavior, Figure 4 presents average savings (cash-in-hand) over time, separated by treatment, and again including the optimal savings path, in the retirement savings game. The figure is divided into three panels: the left panel presents the results for no nudge, the middle for the 20% nudge, and the right panel for the 30% nudge.

These figures sum up nicely a central result of the paper: for both no nudge and 20% nudge, financial literacy appears to mitigate under-saving. By contrast, the 30% nudge results in a better level of saving than the other nudges, and the financial literacy training has no additional effect on the savings level.

Figure 4: Average savings by nudge with and without financial training - retirement savings game)

Table 2: Table 19 for the lab experiment with Ranksum test

Treatment	Variable	**Optimal**	Without financial literacy		With financial literacy		P-value from rank-sum test
			n	μ_A (s.d.)	n	μ_B (s.d.)	H_0: $A = B$
Zero nudge	Accumulated savings at round 20	**23.28**	23	15.75 (9.18)	18	18.91 (9.77)	0.24
	Efficiency	**1**		0.76 (0.32)		0.76 (0.39)	0.49
	Standard deviation of consumption decisions	**0.18**		1.57 (1.28)		1.43 (1.32)	0.60
20% nudge	Accumulated savings at round 20	**23.28**	22	15.51 (10.9)	14	19.9 (7.24)	0.19
	Efficiency	**1**		0.58 (0.58)		0.72 (0.35)	0.4
	Standard deviation of consumption decisions	**0.18**		2.12 (1.54)		1.63 (1.20)	0.44
30% nudge	Accumulated savings at round 20	**23.28**	22	21.22 (16.43)	19	20.88 (6.38)	0.53
	Efficiency	**1**		0.69 (0.39)		0.89 (0.15)	0.01**
	Standard deviation of consumption decisions	**0.18**		1.91 (0.85)		1.25 (1.04)	0.03**
Total			*67*		*51*		

***p<0.01, **p<0.05, *p<0.1.

Table 2 presents summary statistics for the behavior presented in the figures. In the table, the nudge treatment is listed in the left-most column, variables are located in the second column, optimal values are presented in the third column, treatment averages follow in the next two columns, and the p-value from the Wilcoxon-Mann-Whitney test of the null that the means are equivalent in the literacy training treatments is located in the right-most column. Multi-period averages are computed first by averaging across periods within participant, and then averaging those results across participants.

The first variable examined is average accumulated savings at the end of period twenty. Much of the literature on the inadequacy of retirement savings focuses on attempts to measure this variable; in our experiment we know it exactly. In the zero and 20% nudge treatments there is a substantial increase in retirement savings due to the literacy training: more than $3. However, a Wilcoxon-Mann-Whitney test (which we used due to the relatively small sample size) does not detect statistical significance. In the 30% nudge treatment, retirement savings in on the level of the other treatments with training, whether training was present or not.

We computed an efficiency statistic by dividing the sum of period consumption utility (presented to the participants as the cash they would earn) by the total utility that would have been obtained from optimal consumption. Recall that the incentive in the game was to maximize total consumption utility by smoothing consumption, thus, optimal efficiency is 1. Below efficiency in the table we also report the standard deviation of lifetime consumption. The table shows virtually no difference in efficiency due to the literacy training in the zero and 20% nudges, indicating that the increase in savings due to the training did not decrease the volatility of consumption. On the other hand, while literacy training did not increase aggregate savings in the 30% nudge, it is associated with a decrease in consumption volatility and (equivalently) an increase in efficiency, and this result is significant.

Recall that all subjects practiced the retirement savings game for no pay at least three times, and that they all experienced the same unique three income histories for each of those practices. While it is difficult to compare the time path of decision-making across those games due to the different optimal choice paths, we can look at the average savings participants had at retirement in all of the cases. We ran an OLS regression with savings at retirement as the dependent variable, and financial training as the primary dependent variable of interest on these data. Although three of these games were not incentivized, and participants were learning as they played, it seems unlikely to us that consistent results across all four games would be spurious.

The results are shown in Table 3. Note that parameter estimates for the practice session dummies, controlled for, are not shown. Table 3 makes clear that financial training is statistically associated with an increase in accumulated retirement savings as well as increase in efficiency. Notice also that the 20% nudge reduced efficiency compared with the zero nudge in a statistically significant sense, while not altering the level of savings. A look back at Figure 6 provides some visual evidence for this result: notice that the savings peaks in the 20% nudge are often higher than they are in the zero nudge.

Finally, we computed average savings at the time of retirement with and without financial literacy for all three nudges for the practice games and combined them with the paid games. A Wilcoxon-Mann-Whitney test confirms literacy training raised savings rates in these two cases. For zero and 20% nudges, the p-value is 0.083, and for the 30% nudge the p-value is 0.149. Note that the p-values generated by the test for the low nudges indicates that all four average retirement savings levels in the literacy treatment were higher than all four in the no-literacy treatment, in other words, this is as strong as the evidence can be with four data points that financial literacy increases savings levels.

The indications are the literacy training in these experiments affected behavior in a positive but complicated way. First, in the case of the smaller nudges, literacy is associated with higher savings levels, and a restoration of efficiency that was lost by the nudge. Second,

Table 3: Efficiency by nudge and financial training

	Accumulated sav.	**Efficiency**	**Consumption s.d.**
Financial training	4.31***	0.09**	-0.35***
	(1.04)	(0.04)	(0.11)
20% nudge	0.58	-0.15***	0.32**
	(1.27)	(0.05)	(0.14)
30% nudge	3.33***	-.004	0.17
	(1.23)	(0.05)	(0.13)
Constant	15.33***	0.73***	1.66***
	(1.33)	(0.05)	(0.14)
R2	0.083	0.054	0.062
N	472	472	472

in the case of the large nudge, literacy is associated with higher efficiency and unaffected savings levels. Overall the evidence is that literacy is capable of increasing savings and improving efficiency.

To better understand this behavior we inferred decision rules from their actions in the games. Several existing experimental studies provide evidence with regard to decision rules we might expect in our data, and we apply the heuristics in Tasneem and Warnick (2018).

Rule 1: Constant Consumption: This rule is is defined as $C_t = k_1 + \epsilon_{1t}\ \forall t \neq T$. It tries to maintain an approximately constant level of consumption except for the final period of the precautionary treatment, or the last income generating period and later in the retirement treatment.

Rule 2: Constant Savings: This rule is defined as $W_t - C_t = k_2 + \epsilon_{2t}\ \forall t \neq T$, where W_t is the cash-in-hand in period t. It tries to maintain an approximately constant level of savings except for the final period of the precautionary treatment, or the last income generating period and later in the retirement treatment.

Rule 3: Constant Propensity to Consume: This rule is defined as $\frac{C_t}{W_t} = k_3 + \epsilon_{3t}\ \forall t \neq T$. It tries to maintain an approximately constant propensity to consume from the cash-in-hand except for the final period of the precautionary treatment, or the last income generating

Figure 5: Decision rules by nudge and financial training - retirement savings game

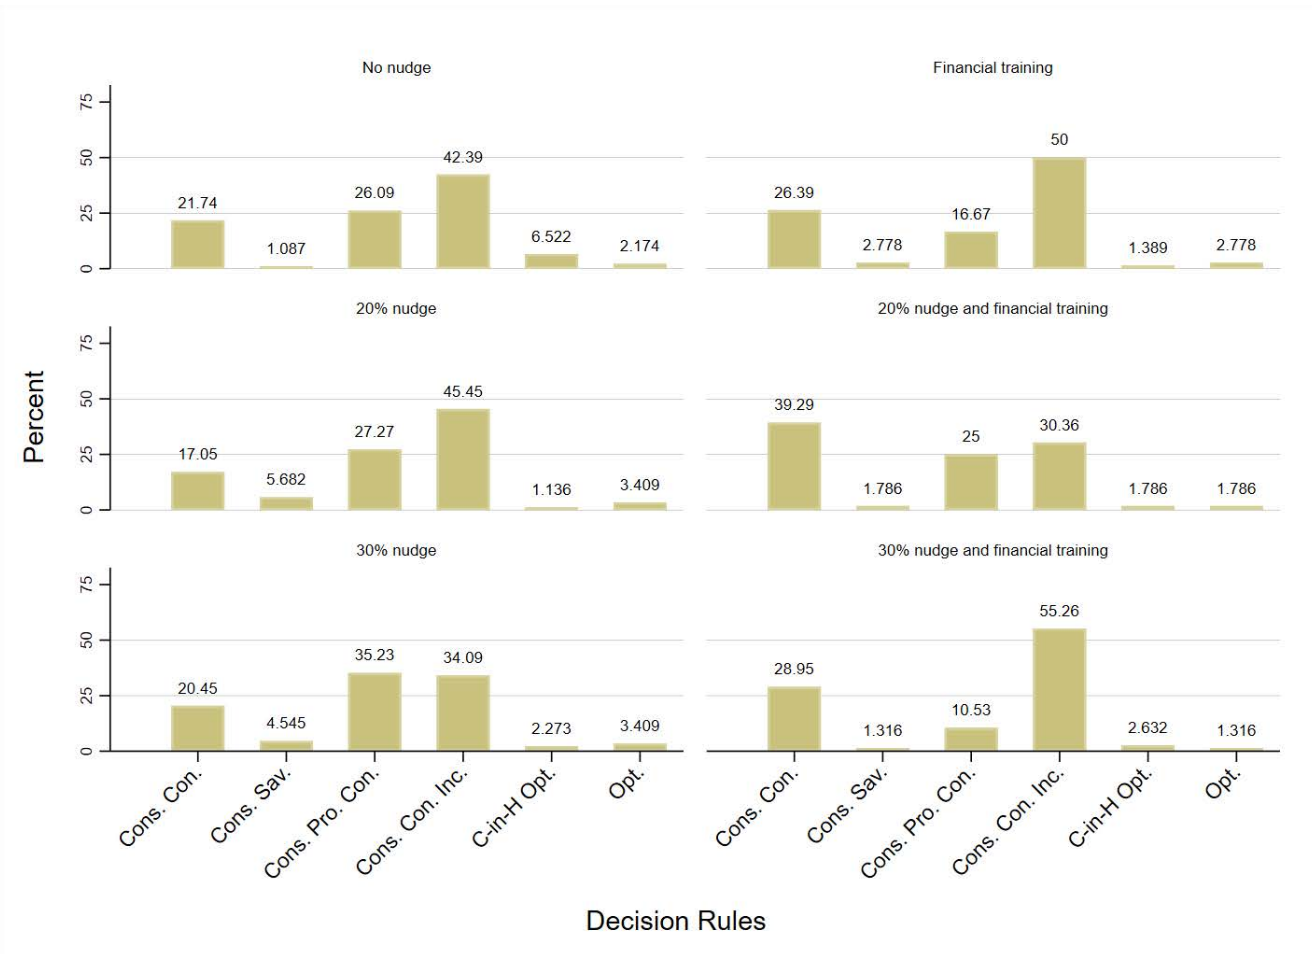

period and later in the retirement treatment. An important note here is that ϵ_{3t} is not an error in the consumption but in the propensity to consume which can be converted to error in consumption.

Rule 4: Constant Consumption Conditional on Income Level: This rule is defined as $C_t = k_{4,1} + k_{4,2}I_t + \epsilon_{4t}\ \forall t \neq T$, where I_t is an indicator variable assuming the value zero if the subject experiences low income in period t and one if the subject experiences high income in period t. It tries to spend a particular amount on consumption for each income level.

Rule 5: Cash-in-Hand Optimal Consumption: This rule is defined as $C_t = C_t^* + \epsilon_{5t}$. It follows the optimal consumption policy given the cash-in-hand at that period.

Rule 6: Optimal Consumption: This rule is defined as $C_t = C_t^o + \epsilon_{6t}$. It follows the optimal consumption policy.

The data from all four games, three practice and one for pay, are ideal for this inference procedure because they do not require the optimal choice path to be the same across observations. Indeed, the different realizations of the time paths of income can help to identify decision rules the participants appear to use. We infer a rule for each participant for each retirement savings game. We compute the likelihood score of each rule, and simply take the best one in each case.

The results of the estimation are presented in Figure 5. The figure is divided into three panels vertically: the top panel presents results from no nudge, the middle panel from the 20% nudge treatment, and the bottom panel from the 30% nudge treatment. Left-to-right presents without and with financial training.

In the top panel, with no nudge, the modal decision rule inferred is shown to be constant consumption conditional on income. This rule, and the other two rules that attempt to keep consumption constant in some way, account for nearly all of the organic behavior in the retirement savings game. There appears to be no change in inferred decision rules with literacy training, and a Fisher's exact test confirms with a p-value of 0.325.

When participants are nudged, there is an entirely different story. In the 20% nudge treatment, financial literacy appears to induce a movement out of the constant consumption conditional on income rule and into the constant consumption rule. This is all the more interesting because recall that the 20% nudge did not result in increased aggregate savings over no nudge, as was expected, however it did have a subtle effect on the apparent decision rules that were used. A Fisher's exact test rejects the null hypothesis that the decision rule distribution is the same for no literacy and literacy with a p-value of 0.048.

In the 30% nudge treatment we again see a difference induced by financial literacy training. Constant propensity to consume behavior is largely replaced by constant consumption conditional on income behavior, as participants correct the over-savings induced by the large nudge treatment. Recall that financial literacy was associated with a statistically significant reduction in consumption variance in this treatment. The difference in the decision rules explains how this occurred, and the p-value rejecting equivalent distributions in this case is 0.01.

5 Conclusion

We presented results from an economics experiment that examined the role of financial literacy in retirement savings. In the experiments, participants made decisions in a retirement savings game, in which income during working years was uncertain. Participants were nudged to varying degrees with automatic savings in each period of the game. Some participants received financial literacy training in the form of computing the expected savings needed at retirement to smooth consumption over the entire life cycle.

We found that the small nudge did not increase savings rates above no nudge, but it did decrease efficiency in the form of increased consumption variance. Literacy training moved savings to a level comparable to the large nudge, and restored the no-nudge level of efficiency. A decision rule inference procedure identified a constant consumption conditional on income rule as being a likely reason for the increase in savings, and that rule is expected to have a higher variance.

We also found that the large nudge increased savings rates, but not quite to the optimal level. Literacy training did not affect the savings rate, but it did tend to decrease consumption variance. This was likely due to the decline of a constant propensity to consume rule, that was induced by the high nudge, being swapped out for a constant consumption conditional on income rule. In this case, the switch to that rule decreased consumption variance.

Our literacy training design highlights the potential complexity of the effects of nudges. Apparently it is possible to positively affect savings rates without improving the underlying volatility of consumption. Such a result is good for companies and governments that wish to lower their pension liabilities, but does not improve the welfare of the savers.

Apparently it is also possible to improve outcomes for savers through financial literacy, by lowering consumption volatility, without increasing savings rates. Such a result is difficult to identify in the field, and to our knowledge, has not been at the center of discussion regarding choice architecture for savings decisions.

Overall literacy training, which was specifically targeted at a particular need, which was applied uniformly to all participants, and which was relevant to all participants, increased retirement savings and improved consumption efficiency.

References

Ballinger, T., E. Hudson, L. Karkoviata, and N. Wilcox (2011), "Saving Behavior and Cognitive Abilities", *Experimental Economics* 14:349-374.

Ballinger, T. Parker, M. Palumbo, and N. Wilcox (2003), "Precautionary saving and social learning across generations: an experiment", *Economic Journal* 113:920-947.

Benartzi, S. and R. Thaler (2007), "Heuristics and Biases in Retirement Savings Behavior", *Journal of Economic Perspectives* 21:81-104.

Benartzi, S. and R. Thaler (2013), "Behavioral Economics and the Savings Crisis", *Science* 339:1152-1153.

Brown, A., Z. Chua and C. Camerer (2009), "Learning and visceral temptation in dynamic saving experiments", *Quarterly Journal of Economics* 124: 197-231.

Carbone, E. and J. Duffy (2014), "LIfecycle Consumption Plans, Social Learning, and External Habits: Experimental Evidence", *Journal of Economic Behavior and Organization* 106:413-427.

Carbone, E. (2006), "Understanding Intertemporal Choices", *Applied Economics* 38: 889-898.

Carbone, E., and J. Hey (2004), "The Effect of Unemployment on Consumption: An Experimental Analysis", *The Economic Journal* 114:497:660-683.

Deaton, A. (1992), "Household Saving in LDCs: Credit Markets, Insurance and Welfare", *Scandinavian journal of Economics* 94(2): 253-273.

Foerster, S., J. Linnainmaa, B. Melzer, and A. Previtero (2017), "Retail Financial Advice: Does One Size Fit all?", *Journal of Finance* 72:4:1441-148.

Hey, J. and V. Dardanoni (1988), "Optimal Consumption Under Uncertainty: An Experimental Investigation", *Economic Journal* 98:105-116.

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

Mailkel, B. (2013), "Asset Management Fees and the Growth of Finance", *Journal of Economic Perspectives* 27:2:97-108.

Miranda, M.J. and P.L. Fackler (2002), *Applied Computational Economics and Finance*, MIT Press, Cambridge.

Tasneem, D, and J. Warnick (2018), "Decision Rules for Precautionary and Retirement Savings", CIRANO working paper.

Thaler, R. and C. Sunstein (2008), *Nudge: Improving Decisions about Health, Wealth and Happiness*, Yale University Press.