CIRANO
Allier savoir et décision

A Laboratory Study of Nudge with Retirement Savings

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

2018S-23
WORKING PAPER
CS

2018s-23

A Laboratory Study of Nudge with Retirement Savings

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

CIRANO

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du gouvernement du Québec, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO is a private non-profit organization incorporated under the Quebec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the government of Quebec, and grants and research mandates obtained by its research teams.

Les partenaires du CIRANO

Partenaires corporatifs
Autorité des marchés financiers
Banque de développement du Canada
Banque du Canada
Banque Laurentienne
Banque Nationale du Canada
Bell Canada
BMO Groupe financier
Caisse de dépôt et placement du Québec
Énergir
Hydro-Québec
Innovation, Sciences et Développement économique Canada
Intact Assurance
Investissements PSP
Ministère de l'Économie, de la Science et de l'Innovation
Ministère des Finances du Québec
Mouvement Desjardins
Power Corporation du Canada
Rio Tinto
Ville de Montréal

Partenaires universitaires
École de technologie supérieure
HEC Montréal
Institut national de la recherche scientifique
Polytechnique Montréal
Université Concordia
Université de Montréal
Université de Sherbrooke
Université du Québec
Université du Québec à Montréal
Université Laval
Université McGill

Le CIRANO collabore avec de nombreux centres et chaires de recherche universitaires dont on peut consulter la liste sur son site web.

ISSN 2292-0838 (en ligne)

A Laboratory Study of Nudge with Retirement Savings[*]

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

Résumé/Abstract

We report results from an on-line economics experiment that examines the effect of nudging retirement savings decisions. In the experiments, participants make decisions in a finitely repeated retirement savings game, in which income during working years is uncertain, and retirement age is known. Participants, who are household financial decision-makers, are nudged with automatic savings in each period of the game. We find that that the nudge simply replaced natural decision-making observed in the absence of a nudge in this experiment, even to the extent that it resulted in nearly identical inferred decision rules. This surprising result highlights the unpredictability of the effect of nudging human behavior.

Mots clés/Keywords: Precautionary Savings; Retirement Savings; Life-cycle Models; Dynamic Optimization; Decision Heuristics; Nudge; Choice Architecture

Codes JEL/JEL Codes: C91; E21; C61

---

[*] 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

There is a common notion that many people save too little for retirement (Canada: Hamilton, 2015; United States: Chien and Morris, 2018). Fueling this perception is the fact that savings rates overall have sharply declined in the United States since the 1980's (Browning and Lusardi, 1996). Choice architects have stepped into the void to design institutional environments in such a way that saving more is made easier, given predictable biases that exist in decision-making (Thaler and Sunstein, 2008). The main principles behind nudging retirement savings essentially involve default options: the availability of a payroll deduction plan, automatic enrollment into the plan, default investment options, and automatic escalation of the amount saved over time (Benartzi and Thaler, 2013). Taken together, these principles increase saving.

However, it is not necessarily clear that people under-save for their retirement when life cycle consumption is carefully analyzed. For example, Banks, Blundell, and Tanner (1998) examine the "retirement savings puzzle" through the lens of the life cycle model in the UK. They find that accounting for labor market status and the components of consumption at retirement removes most of the difference between actual and predicted consumption over the life cycle, with the exception of a large dip in consumption in the early-mid sixties. Li, Shi and Wu (2015) find a similar and stronger result for China, as do Aguil, Attanasio, and Meghir (2011) and Scholz, Seshadri, and Khitatrakun (2006) for the United States.

We know that nudging causes people so save more, but we cannot be sure how much an individual should save in real life, thus the question we ask in this paper is whether nudge induces people to save better. This is a difficult question to answer in the field for at least two reasons. First, there is much relevant information available to the saver but not to the econometrician. Optimal savings depends on unobserved time and risk preferences, and many reasons to save can be complimentary (for example, retirement savings can be used for precautionary reasons). Second, measurement error makes observing both savings and (equivalently for this purpose) consumption difficult.

This paper reports results from an on-line economics decision-making experiment on retirement saving. The experiment eliminates much of the first difficulty by presenting an experimental design that induces consumption smoothing over the life-cycle with a utility function; it separates the precautionary savings motive from the retirement savings motive; it presents a finite game with no issue of time preference; and it makes the distribution of lifetime income clear and constant. It eliminates the second difficulty with field studies by presenting an environment where income and consumption are precisely known, and where their difference is exactly savings.

Our on-line experiment consists of two basic games. In the first game, participants, who are heads of households likely to be in the savings portion of their life-cycle, make decisions in a twenty-period game. In each period there is a constant and independent 50/50 chance of earning either a high or a low income. Participants realize their income and then choose how much to consume and how much to spend. The consumption is turned into cash earnings for experiment participation through a CRRA utility function inducing the motive to smooth. The savings motive in this game is precautionary.

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

To test the effect of nudging, we automatically place 0% or 20% of pay into savings at the beginning of each period. The 0% contribution is intended to be the treatment without nudge. The 20% treatment nudges on average about right, and follows the prescription for a nudge through automatic enrollment. Note that the constant savings rule that results from the nudge is not optimal, as the optimal rule is a function of cash in hand and the period. Thus, there is scope for improving the choice after experiencing the nudge.

Our precautionary savings game 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 Brown, 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 an experiment that tests nudge in this environment.

We find that participants are able to smooth consumption with savings, both for precautionary and retirement purposes, however, not to the degree predicted by theory. Surprisingly, we find no aggregate statistical difference in savings behavior when participants are nudged. Not only was aggregate behavior nearly indistinguishable between experimental treatments, but on a deeper level we inferred essentially identical distributions of decision rules in each treatment. This surprising result underscores the unpredictability of the effect of nudging human behavior.

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 basic 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 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 twenty-five period 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)).

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 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

[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)).

[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., $\beta = 1$.

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

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 history (black dots), the optimal consumption path in the precautionary savings game (green squares), and the optimal consumption path for the retirement savings game (red triangles). The figure makes apparent the degree to which savings reduces the variance induced by the random income stream. Notice also that the addition of the retirement phase makes it possible to smooth consumption better than under the precautionary regime, even during working life.

To test the effect of nudging retirement savings, we automatically placed either 0% or 20% of period income into a savings account 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 two automatic savings decisions, one of which is a baseline and one of which is approximately the right amount to save on average.

The game was played on-line by a total of one-hundred sixty nine participants, between the ages of twenty-seven and fifty, judged to likely be in the savings portion of their life cycle, in a large metropolitan area. Seventy-eight received the zero-nudge treatment and ninety-one received the nudge. 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.

[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

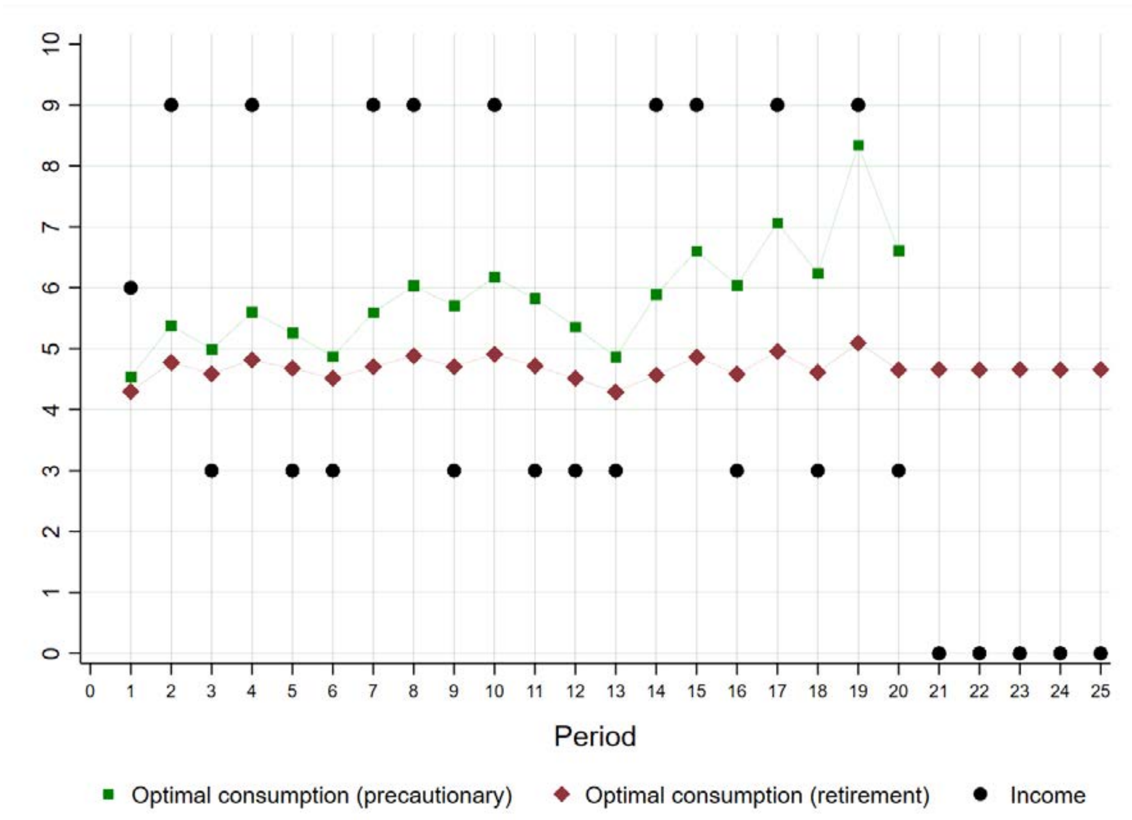

Savings (equivalently, consumption) decisions were made using a slider, which depicted consumption on the left and resulting cash-in-hand on the right. The experimental cash payoff to date was shown on the screen as the sum of spending decisions, transformed into consumption through the CRRA utility function. 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 games is probably most appropriately measured in terms of consumption variance compared with the variance of the income stream. However, nudge is not oriented directly towards efficiency, as it is intended to induce more savings.

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

4 Experimental Results

Table 2 presents participant socio-demographics of our specialized subject pool. The mean age of the participants was 40.6, and about 60% of participants were male. The modal income range was \$50,000-\$74,999 in this large metropolitan area. Nearly all of the participants were employed at the time of the experiment, and the reported education level was rather high: only about 20% reported not having completed college (note that college in this context refers to the last year of high school plus one additional year of education before three years of university). One-third reported having little savings and a fair amount of debt, while more than 60% of the participants reported having at least some savings.

Figures 2 and 3 present average consumption, broken out by treatment, over the twenty-period precautionary savings game and twenty-five period retirement savings game respectively. Both figures also include the time path of optimal consumption (solid line). Beginning with the pre-

Table 2: Socio-demographics

Age	
Mean	40.6
Standard deviation	5.9
Gender	
Male proportion	61.5%
Female proportion	38.5%
Annual income	
Less than $25,000	0.6%
$25,000 - $49,999	19.5%
$50,000 - $74,999	44.4%
$75,000 - $99,999	25%
$100,000 - $124,999	8.3%
$125,000 - $149,999	1.8%
Employment status	
Employed	95.9%
Unemployed	1.2%
Retired	0.6%
Student	0.6%
Homeworker	1.8%
Education level	
Some High-School	1.2%
Completed High-School	11.8%
Some college	6.5%
Completed college	24.9%
Some university	11.8%
Completed undergraduate	20.1%
Some graduate	4.7%
Master's	16.6%
PhD	2.4%
Financial situation	
No savings and significant debt	5.3%
Little savings and a fair amount of debt	32.5%
Some savings and some debt	26.0%
Some savings and little or no debt	26.6%
Significant savings and little or no debt	9.5%
N	169

Figure 2: Average consumption by nudge treatment - precautionary savings game)

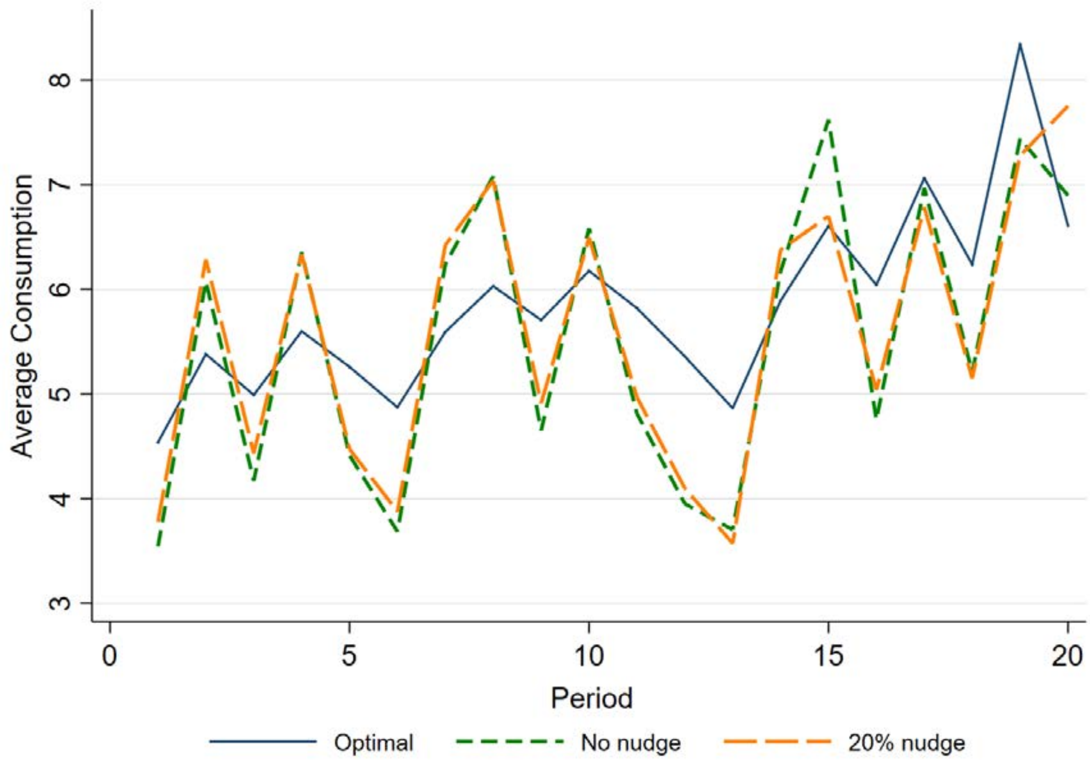

cautionary savings game (Figure 2), in both treatments, participants were able to smooth their consumption to a significant extent: minimum consumption for a period was approximately \$3.5 and maximum consumption was under \$8 (recall that income varied beteen \$3 and \$9). There is room for improvement, as the optimal consumption path is smoother than the actual paths. And there is little difference between experimental treatments, as is to be expected. There is no nudge in the precautionary savings game.

Figure 3 shows that participants were able to improve smoothing in the retirement savings game, as theory predicts. As with the precautionary savings treatment, there is less smoothing than with the optimal consumption path. Surprisingly, we see no effect from the nudge of the 20% savings account. It is almost as if participants are nudging themselves.

To get an alternative look at the same behavior, Figure 4 presents average savings (cash-in-hand) over time, again separated by treatment, and again including the optimal savings path, in the retirement savings game. In the early periods, average savings is a bit higher than optimal savings,

Figure 3: Average consumption by nudge treatment - retirement savings game

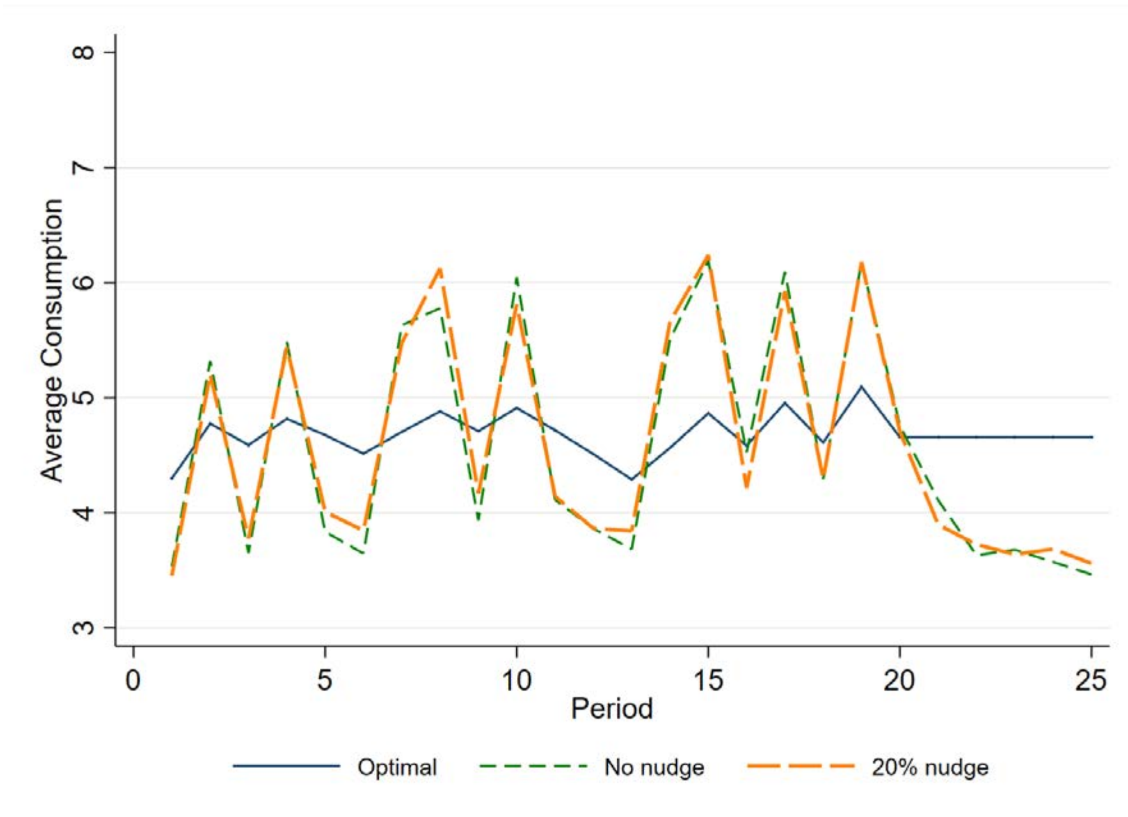

Figure 4: Average savings by nudge treatment - retirement savings game

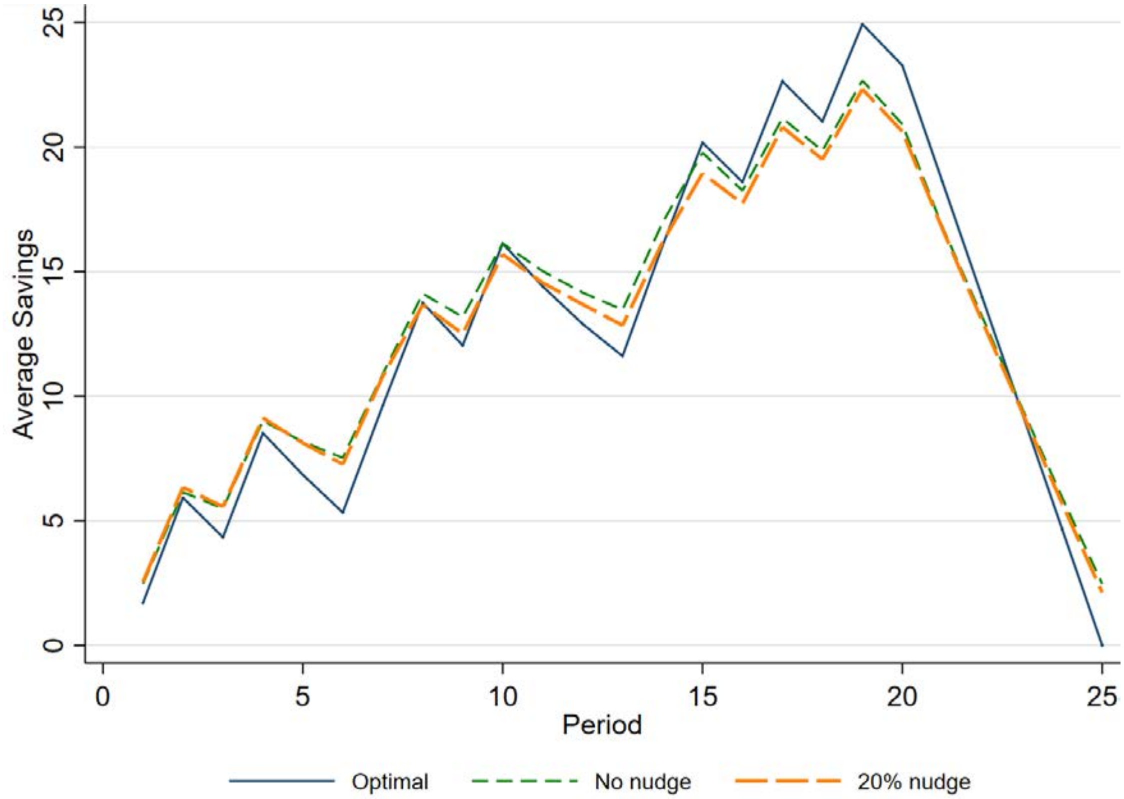

appearing as a risk averse strategy. After period 15, aggregate savings falls below optimal, and retirees are approximately \$2.50 poorer on average than they should be. That is roughly one-half optimal expected savings per period in the game. However one can see that savings closely tracks optimal, and nudge or not, participants do a rather good job managing their cash-in-hand.

Table 3 presents summary statistics for the behavior presented in the figures. In the table, variables are listed in the left-most column, optimal values are presented in the second column, then treatment averages follow in the next two columns, and the p-value from the t-test of the equivalence of means in the two nudge 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 both treatments, participants are more than \$2.00 short of what they would have if they behaved optimally. Since optimal average consumption is a bit more than \$4.50 per period, participants are lacking approximately half of a periods savings over the twenty periods of working life for optimal consumption at retirement. We fail to reject the null hypothesis of equivalent accumulated savings in the two treatments.

We computed an efficiency statistic by dividing the sum of period consumption utility (presented to the participants as points) 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. The table shows efficiency numbers of 0.72 with no nudge and 0.76 with the nudge, and no statistical difference between the two. Below this statistic we show that, equivalently, there is no significant difference in lifetime consumption variance.

Table 3: Summary Statistics

Variable		Optimal	No nudge μ_1 (s.d.)	20% nudge μ_2 (s.d.)	t-test p-value H_0: $\mu_1 = \mu_2$
Accumulated Savings at End of Round 20		**23.28**	20.92 (12.06)	20.63 (12.25)	0.88
Efficiency		**1**	0.72 (0.31)	0.76 (0.25)	0.35
Standard Deviation of Consumption Decisions		**0.18**	1.78 (1.02)	1.68 (1.07)	0.55
Average Consumption					
Periods 1-20 (β_1)		**4.69**	4.8 (0.6)	4.82 (0.61)	0.88
Periods 21-25 (β_2)		**4.66**	3.69 (2.19)	3.7 (2.22)	0.98
t-test p-value	H_0: $\beta_1 = \beta_2$		0.00***	0.00***	
	H_0: $\beta_1 \geq \beta_2$		0.00***	0.00***	
n			*78*	*91*	

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

We can also mimic a standard route of inquiry in field studies by quantifying the drop in living standard at retirement age. Unlike in field studies, we can see the precise effect of retirement on consumption in our experiment. The consumption portion of the table shows that average consumption during the working phase of the life cycle was about 4.8 per period in both treatments, which is about 0.1 above optimal. Consumption then dropped to 3.7 per period in retirement, which was about 1.0 below optimal. In our experiment, retirees experienced a 20% decline in lifestyle compared with their working lives.

It is surprising that the nudge did not have a discernible effect on aggregate behavior. To better understand this 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) to our data.

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 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.

We infer a rule for each participant on the their choice data in the paid retirement savings game. We compute the likelihood score of each rule, and simply take the best one in each case.

Figure 5: Decision rules by nudge

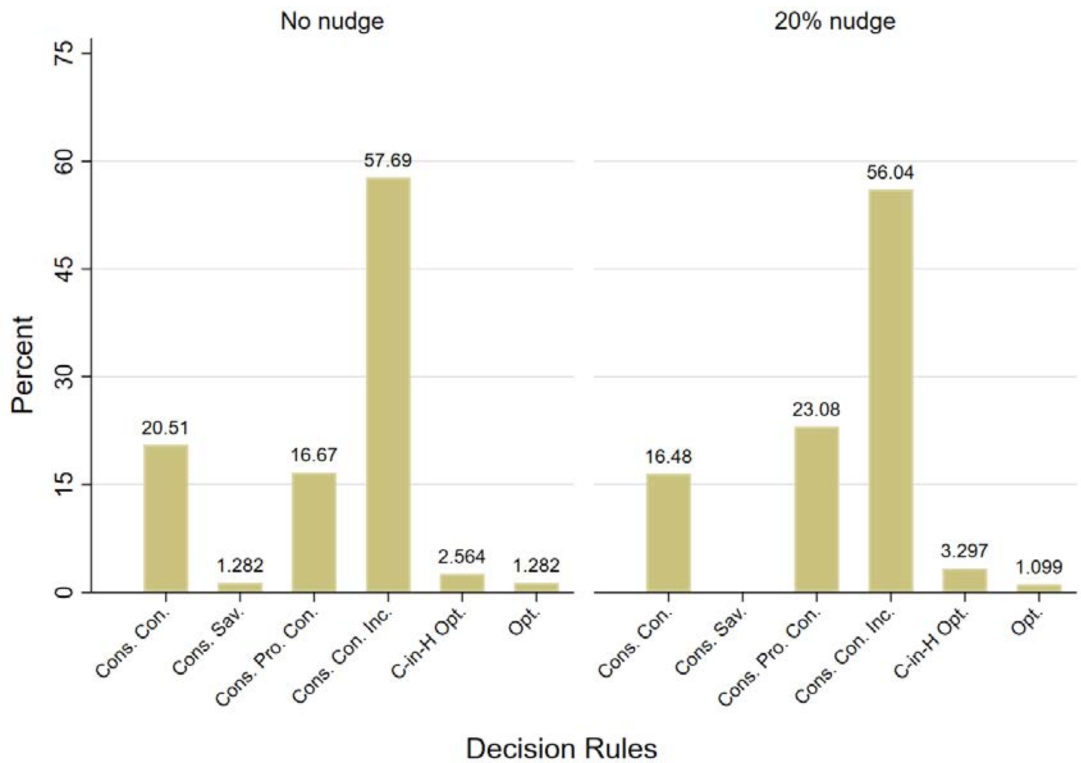

The results of the estimation are presented in Figure 5.

Figure 5 shows the distribution of the inferred rules for both experimental treatments. The left half of the figure presents the results for no nudge, and the right half presents the results for the 20% nudge. The result is rather strong: all three "constant consumption" rules are strongly inferred in the data, with modal constant consumption conditional on income rules in both cases. There is very little optimal behavior inferred from this exercise, and constant savings is also basically a no-show.

Further, a Fisher's exact test shows that there is no significant relationship between decision rules and presence of the nudge (p-value = 0.831). These results imply that the organic behavior in this game looks like a nudge on a fairly deep level. Nudging savers for retirement simply replaced their natural decision-making.

This study raises two questions. First, what level of a nudge would affect behavior? We think it is important to show that nudges may have no effect, and our result emphasizes the need for nudges that are strong enough to raise savings levels. However at some point, a nudge will kick in, and we address this question in subsequent work.

Second, why did participants in this study lose significant lifestyle to retirement, when at least four recent empirical studies cast doubt on the "retirement savings puzzle"? This also warrants further study. It is possible that the income distribution, while chosen for simplicity of exposition, may be more difficult to handle than real-life income distributions. It is also possible that people receive good advice, and that pension and benefits programs make the decision easier in real life. In any event, our study enables the exploration of this question.

5 Conclusion

We reported results from an economics experiment that examined the effect of nudging retirement savings decisions. In the experiments, participants made decisions in a retirement savings game, in which income during working years was uncertain. Participants were nudged with automatic savings in each period of the game. We found that that the nudge simply replaced natural decision-making in this experiment, even to the extent that it resulted in the nearly identical inferred decision rules.

This surprising result suggests that the effect of a nudge may be more difficult to measure than thought. We nudged participants roughly on average correctly for this game, though the nudge was far from the optimal strategy in the game. Not only was the resulting aggregate behavior nearly indistinguishable between experimental treatments, but on a deeper level we inferred essentially identical distributions of decision rules in each treatment. Our expectation was that a positive nudge of that magnitude would have a positive effect on savings.

Our study, while establishing how people who are in the savings portion of their lifecycle play the retirement savings game, raises at least two questions. First, what level of a nudge would have a positive effect on savings, or better, what is the optimal nudge? Second, why did we see a "retirement savings puzzle" in this study when at least four recent studies fail to do so with field data. There are numerous candidate answers, including the availability of financial advice, the possible simplification of the problem due to government pensions and health benefits, or possibly unintended difficulties with the 50/50 income distribution which is easy to understand but not a metaphor for real life expected income streams. It is also possible that there exists an undetected puzzle in the field data. With this foundation completed, we can go on to address this important question.

References

Aguila, E, O. Attanasio, and C. Meghir (2011), "Changes in Consumption at Retirement: Evidence from Panel Data", *Review of Economics and Statistics* 93:3:1094-1099.

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.

Banks, J. R. Blundell, and S. Tanner (1998), "Is There a Retirement Savings Puzzle?", *American Economic Review* 88:4:769-788.

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

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

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

Browning, M. and A. Lusardi (1996), "Household saving: Micro theories and micro facts" *Journal of Economic literature* 34(4): 1797-1855.

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

Chien, Y. and P. Morris (2018), "Many Americans Still Lack Retirement Savings", Federal Reserve Bank of St. Louis, First Quarter 2018.

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

Hamilton, M. (2015), "Do Canadians Save Too Little?", C.C. Howe Insitute Commentary no. 428.

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

Li, H., X. Shi, and B. Wu (2015), "The Retirement Consumption Puzzle in China", *American Economic Review* 105:5:437-441.

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

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

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.