CIRANO
Allier savoir et décision

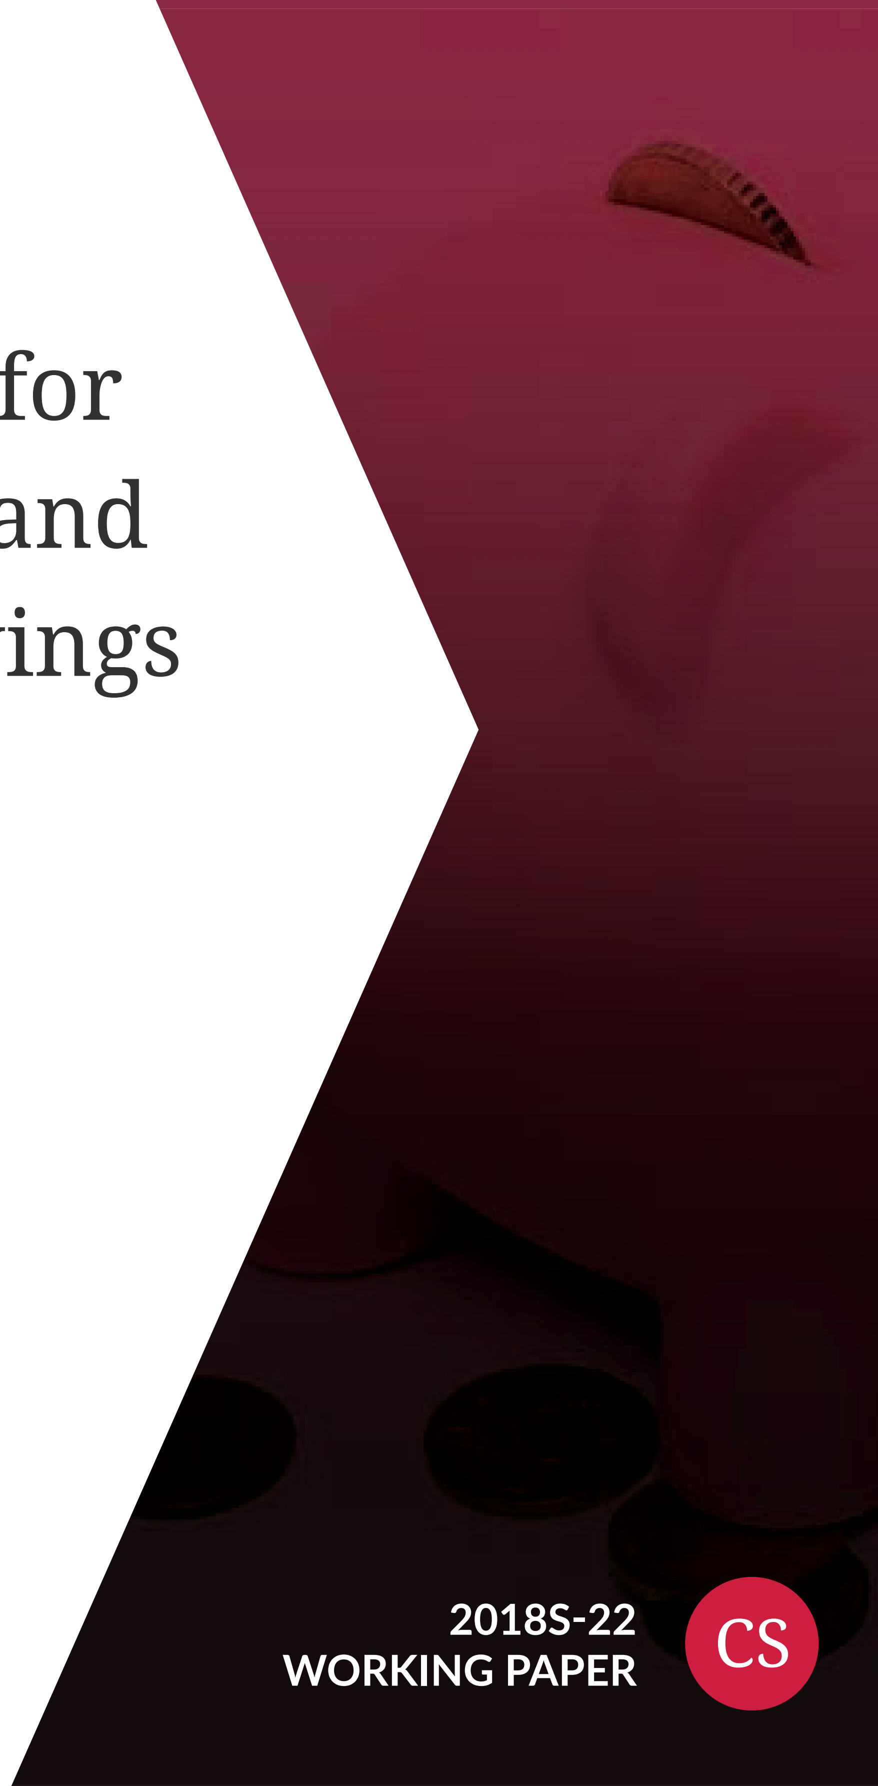

Decision Rules for Precautionary and Retirement Savings

DINA **TASNEEM**
JIM **ENGLE-WARNICK**

2018S-22
WORKING PAPER

CS

2018s-22

Decision Rules for Precautionary and Retirement Savings

Dina Tasneem, 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)

Decision Rules for Precautionary and Retirement Savings*

Dina Tasneem†, *Jim Engle-Warnick*‡

Résumé/Abstract

We report results from an experiment that compares precautionary savings behavior with retirement savings behavior. We find that more than 30% of precautionary savings behavior can be categorized as optimal or near optimal, while virtually all of this behavior disappears in favor of simple decision rules that specify constant consumption in each period when retirement savings is added as a motive. We discuss the the costs and benefits of these rules, which make a complex decision-making environment manageable. Our experiment is the first to identify how decision-making changes when agents are required to save for retirement.

Mots clés/Keywords: Precautionary Savings; Retirement Savings; Dynamic Optimization; Decision Heuristics

Codes JEL/JEL Codes: C91; E21; C61

* Acknowledgements: We acknowledge the Social Science and Humanities Research Council for funding. Corresponding author: Jim Engle-Warnick; jim.engle-warnick@mcgill.ca.

† American University of Sharjah.

‡ McGill University.

1 Introduction

Household saving in general in the United States has declined sharply since the 1980s (Browning and Lusardi, 1996), and retirement saving in particular has been identified as being inadequate (Benartzi and Thaler, 2013). Given the importance of retirement savings, it seems important to understand how people solve the problem. Life-cycle consumption models assume that agents solve a dynamic optimization problem, which often can only be solved numerically (Duffy, 2008; Lusardi and Mitchell, 2007). Identifying to what extent households save optimally is difficult to determine, because of information known to the saver but not the econometrician, measurement error, and complementarities between different types of savings.

Nevertheless there is empirical evidence that household and/or individual decisions show signs of optimization failure (Campbell, 2016). Choi, Laibson, and Madrian (2011) report the failure of the employees of seven firms, in different industries in US, to take full advantage of the 401(k) employers' matched contribution even though they could withdraw their contribution at any time without penalty. MacFarland, Marconi, and Utkus (2004) show that only half of their respondents possess adequate "planner" characteristics to optimally contribute to defined contribution retirement plans. Examining the impact of retirement savings policies on the population of Denmark, Chetty, et. al. (2014) categorize individuals as "active savers" vs "passive savers" based on their responses to different savings incentive policies, where active savers behave approximately optimally and passive savers employ a simple decision rule.

Empirical findings have led to the development of alternative models based on new behavioral hypotheses. Contrary to the permanent income hypothesis Gross and Souleles (2002) find that credit card holders increase spending as their credit limits increase, even though they did not reach their limit before the increase. They argue that this behavior can be better explained with buffer-stock models of precautionary saving. Campbell and Mankiw

(1989) use a two-type consumers model (optimal and rule-of-thumb) to explain empirical regularities in time series data on consumption, income and the interest rate. Mankiw (2000) proposes a similar model where a section of the population, called "savers", are like the economic agents in infinite horizon model of consumption, who smooth consumption and leave bequests. But the other type in the population, called "spenders", consumes all of their per-period disposable income. In similar spirit, Gali, Lopez-Salido, and Valles (2007), in an effort to explain consumption increases in response to increases in government spending, introduce an extension of the new Keynesian model with a fraction of rule-of-thumb consumers who do not borrow or save and who live on their current disposable income.

Experimental evidence tends to focus on the ability of participants to solve dynamic problems. In these studies, participants are presented with a decision problem of multiple periods using some version of the life-cycle model of consumption, where in each period there is an uncertain income realization, and the participants' task is to decide how much to consume and how much to save. The participants are incentivized to find the optimal solution to that problem. Examples include Hey and Dardanoni (1988), Ballinger et al. (2003, 2011), Carbone (2006), Brown et al. (2009), and Feltovich and Ejebu (2013). The experiments tend to find too little savings and too much consumption, however some have uncovered evidence for heterogeneity in sub-optimal behavior. For example, Ballinger et al. (2003, 2011) fit a model of bounded forward-looking to their precautionary savings data, and Carbone (2006) finds evidence for several decision rules that condition on savings or consumption in a life-cycle model.

In this paper we infer decision rules from choice data in precautionary and retirement savings games in an experiment. An experiment is the ideal tool with for this purpose because of our ability to isolate the savings decision. In our precautionary savings experiment, subjects know their income distribution (a high or low income with equal probability each period) and their lifespan (twenty periods). We induce preferences for smoothing consumption with a CRRA utility function. With the finite lifespan, we abstract from issues of self control and procrastination and focus directly on the ability to make an optimal savings choice. In each of twenty periods of life a known random process determines the period income, and the decision is simply how much to save and how much to consume.

In our retirement savings experiment, we simply add five periods to the end of the precautionary savings game with no income, while holding the actual income draws constant across the treatments. Since we know the optimal consumption (equivalently, savings) path (solved numerically), our experiment allows us to determine the quality of choices. Our experiment precisely determines how adding retirement savings onto precautionary savings changes decision-making.

We find that participants are able to smooth their consumption in both precautionary and retirement savings regimes, however less than theory predicts. We find evidence for five decision rules, two of which are either optimal or nearly optimal, for precautionary savings, while other rules include constant consumption and constant propensity to consume out of cash-in-hand are prevalent in retirement savings. In fact, more than 30% of the rules we infer in the precautionary savings game are optimal or near-optimal, while we find virtually none of these rules in the retirement savings game.

When retirement is added as a motive for saving, the optimal decision rules disappear from the data and constant consumption rules are substituted in their place. Subjects whose decision-making is consistent with these rules strive to hold their consumption constant, sometimes conditional on their current income level. When conditioning on current income level, participants over-save during periods of low income, and when consuming a proportion of their cash-in-hand, participants' consumption is upward sloping during their working lives. The use of these simple rules lowers efficiency by 7% compared with performance in the precautionary savings game, and the drop in efficiency is due to over-saving for retirement. Our experiment is the first to identify how decision-making changes when agents are required to save for retirement.

The next section of the paper details the experimental design. The results and conclusion follow.

2 Experimental Design

2.1 The Model

The model, similar to the model in the experimental study of Ballinger et al. (2003), is a finite time forward looking intertemporal consumption problem with an uncertain income in each period and an incentive to smooth consumption. There are no bequests, there is no investment motive, and the agents face a strict borrowing constraint. We induce preferences with a CRRA utility function, and assume that preferences are additively separable over time. Agents discount the future at a constant rate. We present two different decision problems: a precautionary savings problem and a retirement savings problem.

In the twenty-period precautionary savings problem, the income stream is $y = (y_1, y_2, ...y_{20})$, where each $y_t = \$3$ or \$9 with equal probability. 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 and an additional five periods corresponding to retirement are added such that $y_{21} = y_{22} = ... = y_{25} = \$0$. Note that both precautionary and retirement savings motives are present in this version of the game.

Following Ballinger et al. (2003) notation, denote the instantaneous utility of consumption in period t by $u(c_t)$, the accumulated asset at the beginning of period t by A_t and the uncertain labour income realized at the beginning of each period by y_t. 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. Note also that we multiplied the entire utility function 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, if the amount of cash-in-hand goes below a critical value the consumer should spend everything, and the marginal propensity to save is increasing in cash in-hand (Deaton (1992)). We cannot solve the model analytically, therefore we derive the optimal policy by solving the problem numerically using the backward recursion method 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 (Deaton (1992)).

2.2 Behavioral Hypotheses

While the dynamic stochastic optimization problem is solvable in our problems, as well as in typical models employed in empirical studies, the exact solutions are too difficult to achieve for subjects performing the tasks or for households making savings decisions. Nevertheless, empirical evidence suggests that households do a good job with regard to the optimal savings solutions.

One possible explanation is that rules of thumb exist that perform well compared with optimal behavior. For example, Winter et al. (2012) provide evidence for the fact that rules of thumb can result in relatively low utility losses with respect to the optimal choice path, where a key feature of a good rule is the presence of the anticipation of future income. Allen and Carroll (2001) investigate learning a "buffer stock savings rule", resulting from dynamic programming, and for which evidence has been found empirically. This rule requires consumers to achieve a buffer stock of liquid assets, and then consume their average income.

Several existing experimental studies provide evidence with regard to decision rules we might expect in our data. Hey and Dardanoni (1988) find that, in a consumption decision problem with an indefinite ending, optimal behavior did not depend on the period of the problem but actual behavior did, establishing existence of an apparent heuristic. Mueller (2001) presents a consumption decision problem that lasts at least three periods and at most six periods that admits heuristics that, for example, allocate consumption based on the expected number of periods in the problem or tries to achieve the same payoff for all possible histories. Carbone and Hey (2004), in an experiment where employment follows a Markov process, find that subjects overreact (i.e., their consumption is over-sensitive) to their current employment status. Carbone (2005) tests various heuristics in a savings problem against experimental choice data to best describe the behaviour of the subjects, which she identifies by inspection of her data, and some of which appear to be consistent with past experimental results.[2]

Using this literature as a guide, and adapting to our specific game, we hypothesize six decision rules to describe data in our experiments. The rules fall broadly into three categories. First, it seems sensible that participants would attempt to hold consumption constant in a game that incentivizes consumption smoothing, and we hypothesize three rules of this type: constant consumption, constant propensity to consume out of cash-in-hand (sensible for precautionary savings), and constant consumption conditional on current income level. Second, participants might flip the problem and attempt to hold savings constant, as many people do in real life: constant savings. Finally, we hold open the possibility of optimal behavior: purely optimal, and optimal conditional on current cash-in-hand.

The following 6 rules will thus form our behavioral hypotheses.

Rule 1: Constant Consumption: This rule 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. In this category we can include behavior where the subject tries to maintain an approximately constant per period payoff given that we do not have discounting over time. Included here is the extreme behavior where the subject consumes nothing, that is, saves everything over all the periods except for the last period (or last few periods), where he/she consumes everything (or not).

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. The policy of spending everything every period falls into this category implying zero savings always.

Rule 3: Constant Propensity to Consume: This rule is defined as $\frac{C_t}{W_t} = k_3 + \epsilon_{3t}\ \forall t \neq T$.

[2] Other experiments that study heuristics and decision rules in dynamic decision problems include Houser et al. (2004) and Houser and Winter (2004).

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. A similar type of policy can involve propensity to consume from income which is described bellow as our next possible policy.

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.[3] 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 $C_t = C_t^o + \epsilon_{6t}$. This rule follows the optimal consumption policy.

3 Experimental Procedures

In the experiments, at the beginning of each period, participants earned either a high income of \$9 or a low income of \$3 with equal probability, where this distribution is constant and independent of the previous period.[4] In the implementation of the decision problem called "precautionary treatment", an agent earns an income for exactly twenty periods. The income, which is realized at the beginning of each period, is either high or low with equal

[3] In this case we must ignore the first observation in the data. Also notice that the constant consumption rule is a special case of this rule where $k_{4,2} = 0$. Therefore first we categorize subjects' consumption decisions among different decision rules excluding Rule 1. Then for all the subjects who fall into the Rule 4 category we test if the mean consumptions for different incomes are significantly different or not. If the difference in the means is significant then we conclude in the favour of Rule 4 otherwise we select Rule 1.

[4] Note that $\$X$ denotes X experimental dollars.

Table 1: Experimental Parameters

Treatment	k	θ	ϵ	σ	**Income**	**Pr of low Income**	**Starting C-In-H**	**Retirement Period**	$\boldsymbol{T}$
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

probability. Subjects earn "experimental dollars" in income, which they either spend or save. The experimental dollars they spend are transformed into consumption by the utility function, which determines their cash payment at the end of the session. We chose the same relative risk aversion parameter as in Ballinger et al. (2003), $\sigma = 3$. We assume no discounting of utility over time and set $\beta = 1$.

In the second treatment called "retirement treatment", after twenty periods of the same uncertain income distribution, the agent then experiences exactly five periods with no income. These final five periods create a retirement savings motive. Table 1 summarizes the experimental treatments.

Each period of the decision task, participants realize their period income, which is immediately added to their cash-in-hand, and then they decide how much money to spend on consumption. Any cash left becomes cash-in-hand and is carried over to the next period. Participants spend "experimental cash", but are rewarded with real cash for their consumption after transforming their spending into a reward through the CRRA utility function. This transformation from spending to consumption induces the motive for consumption smoothing.[5]

We drew the income streams before the experiment and presented identical draws to all subjects. Since it is well known that subjects form non-standard beliefs about random process (Kahneman and Tversky, 1974), the particular draws we presented to the subjects

[5] A binary lottery mechanism exists to reward subjects to control for risk preferences, see Ballinger et al. (2003) for an example. For simplicity of understanding the experiment we did not employ this mechanism.

Figure 1: Income and Optimal Choice History by Income Draws

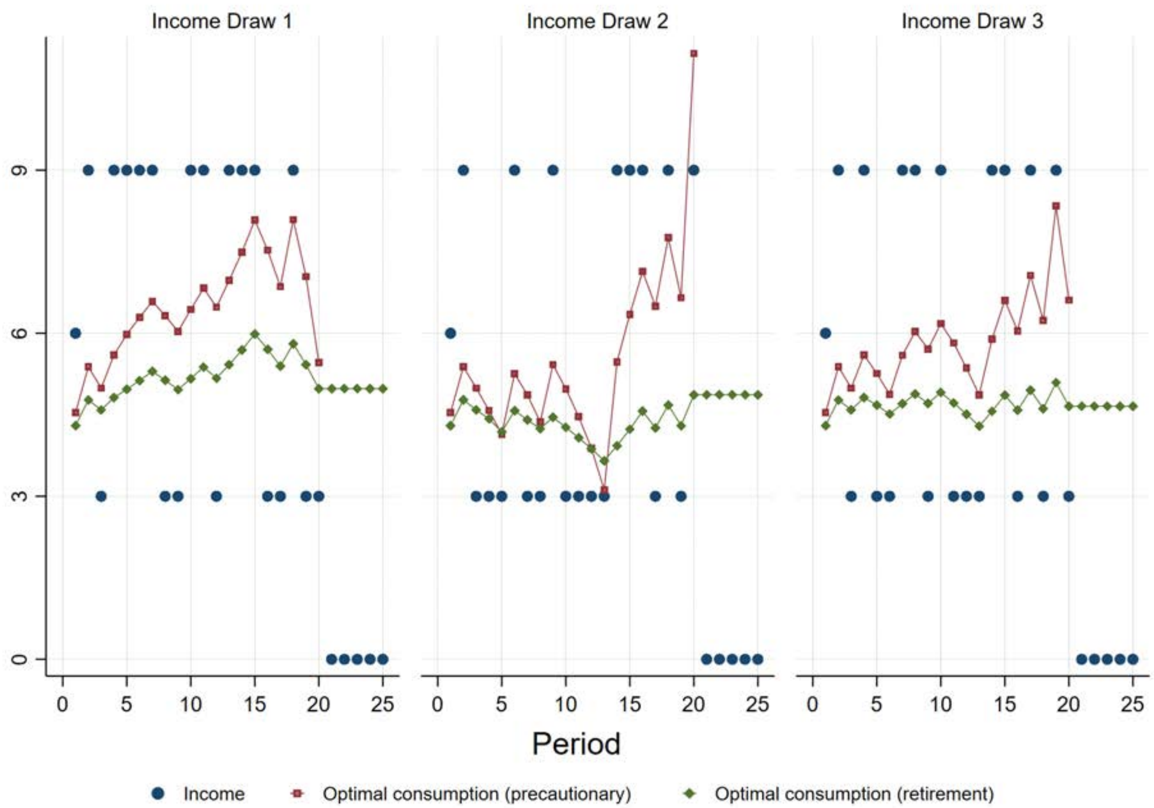

could influence decision-making in their own unique ways, thus we describe them here.

Figure 1 shows the three income draws in the experimental design. The left-hand panel presents the first draw, the middle panel presents the second draw, and the right-hand panel presents the third draw. In each panel, the horizontal axis represents the period in the game and the vertical axis shows income and spending in dollars. The circles represent income (either \$9 or \$3), the squares represent optimal spending on consumption in the precautionary savings game, and the diamonds represent optimal spending on consumption in the retirement savings game. Notice that the initial endowment is the average of the income draws, \$6, and that early-period income is relatively high compared with late-period income in the first draw, relatively low in the second draw, and approximately representative in the third draw.

Before playing the three savings tasks for pay, subjects made decisions in five tasks, each with different income draws, for no pay to maximize their opportunity to learn. At the end of each of the five tasks, subjects were told how much they would have earned had they played optimally, so that they had a relative measure for their performance. Subjects all received the same order of the three tasks for pay, and were paid for the sum of their earnings over all periods in one randomly chosen task of the three.

The experiments were conducted in a university experimental economics laboratory, and the subjects were mostly university undergraduates. Sixty-seven subjects participated in six sessions earning an average of $31.2 for their participation in the experiments including the show up fee of $10.

4 Experimental Results

Figure 2 presents average vs. optimal consumption, separated by income draw, for both treatments. The panels are arranged in the same manner as Figure 1, with precautionary and retirement savings treatments presented horizontally and income draws presented vertically. The statistic shown is the average amount spent on consumption by each participant, averaged across participant, in each period of the game. Recall that the income draws are the same in the precautionary and retirement treatments.

As predicted by theory, participants smoothed consumption better in the retirement savings treatment. In each treatment and for each draw, average consumption tracks optimal consumption, but is also more variable. However, recalling that income each period was either $3 or $9, smoothing was achieved in every treatment. Over-consumption during retirement indicates the presence of over-saving.

Figure 3 presents average vs. optimal savings, separated by income draw, for both treat-

Figure 2: Average vs. Optimal Consumption by Income Draw

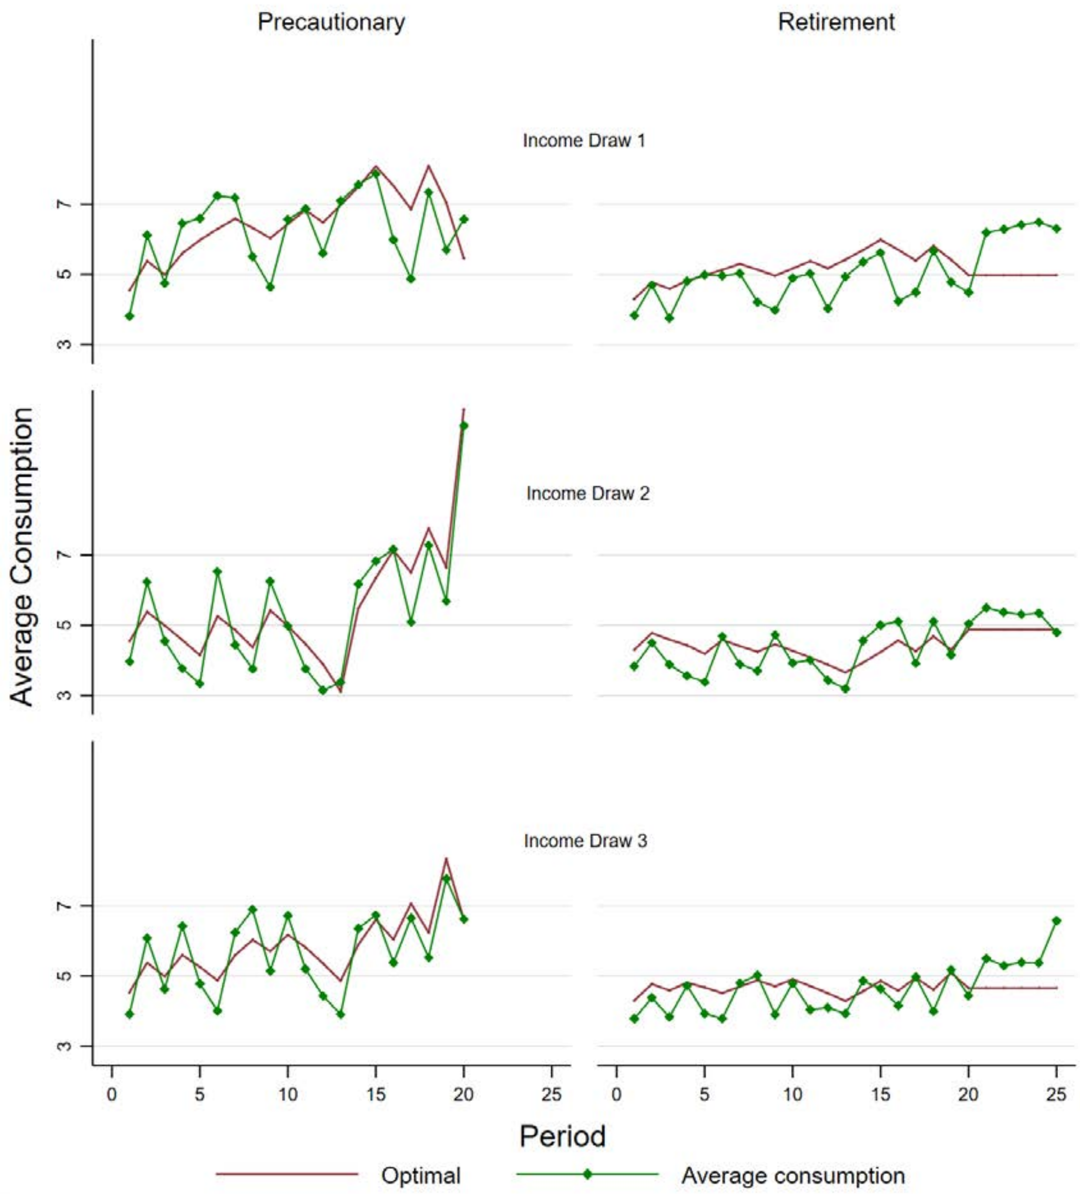

ments. The panels are arranged the in the same manner as Figure 1, with precautionary and retirement savings treatments presented horizontally and income draws presented vertically. The statistic shown is the average amount saved by each participant, averaged across participants, in each period of the game. Recall that the income draws are the same in the precautionary and retirement treatments.

As we suspected from Figure 2, participants on average over-saved during their working lives; this is a risk averse strategy. This is true across the board for all three treatments. Looking at the precautionary savings treatments, this over-saving appears to be due at least in part to over-saving for precautionary purposes.

To better understand behavior presented in Figures 2 and 3, we find the best-fitting decision rules on a game-by-game basis in our experimental data. For each subject in each incentivized game we directly estimate the parameters of decision rules 1-4 with a simple regression or, when relevant, simply by computing an arithmetic mean (Carbone, 2005). We then use these estimates to find the predicted values of consumption (note that decision rule 3 specifies a propensity to consume, rather than actual consumption), and use the predicted values of consumption to compute the sum of squared errors of the rules. For decision rules 5 and 6, optimal policies give the predicted consumption and we compute the sum of squared errors directly from the policies. We select the rule with the smallest sum of squared errors as the best fitting rule for an instance of savings task in the data.

Figure 4 presents the distribution of decision rules in the two experimental treatments, precautionary and retirement savings. The left-hand panel presents the results for precautionary savings. Notice that the constant savings rule (Rule 2), is not a factor in our data for either treatment. Of the five rules that occur with a significant frequency in precautionary savings, Rule 3 is the modal rule: constant propensity to consume. The remaining three rules best fit the data between approximately 12% and 20% of the time. It is interesting to note that the optimal Rules 5 and 6 (5: optimal given current cash-in-hand; 6: optimal

Figure 3: Average vs. Optimal Savings by Income Draw

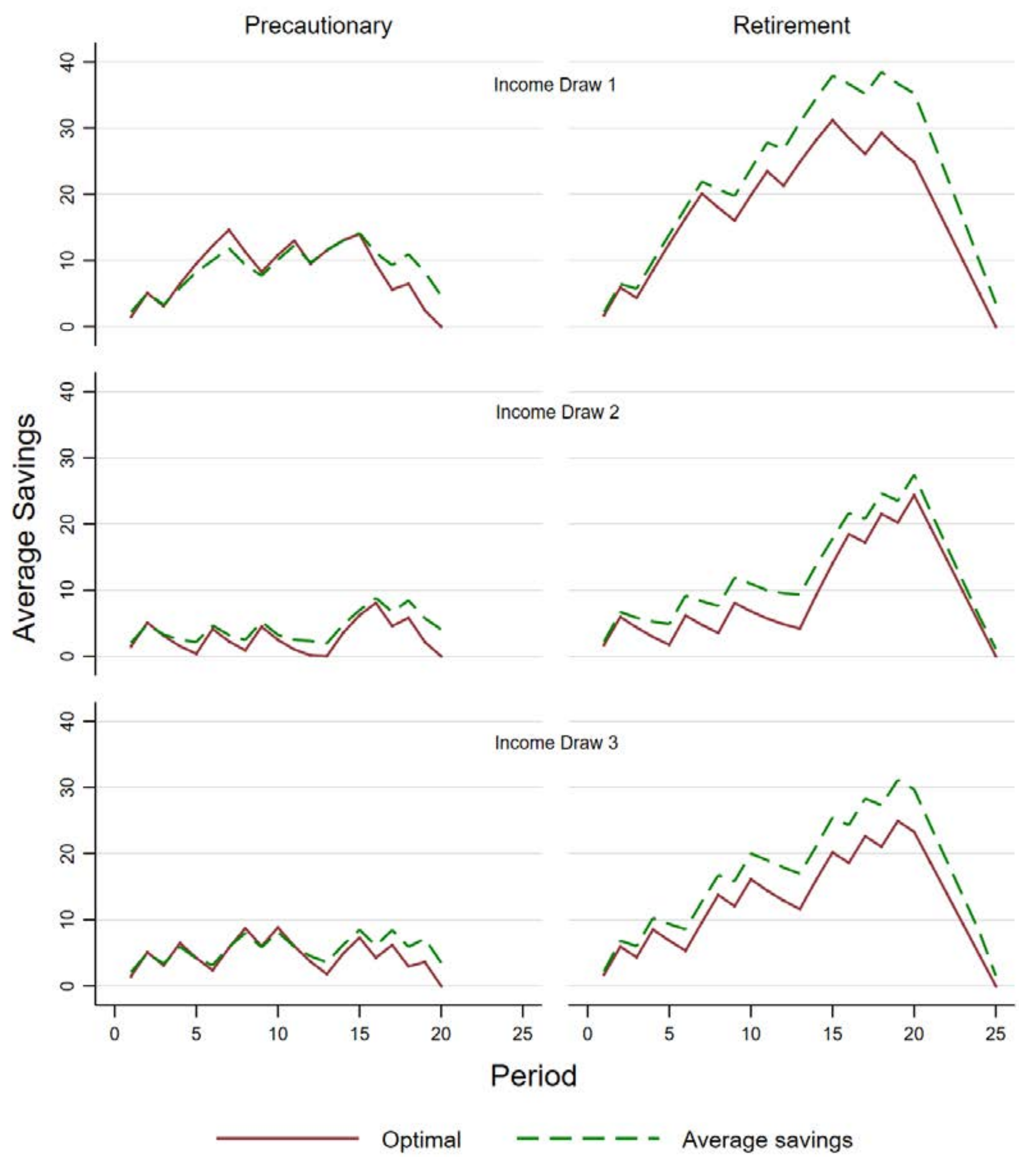

Figure 4: Best Fitting Decision Rules by Treatment

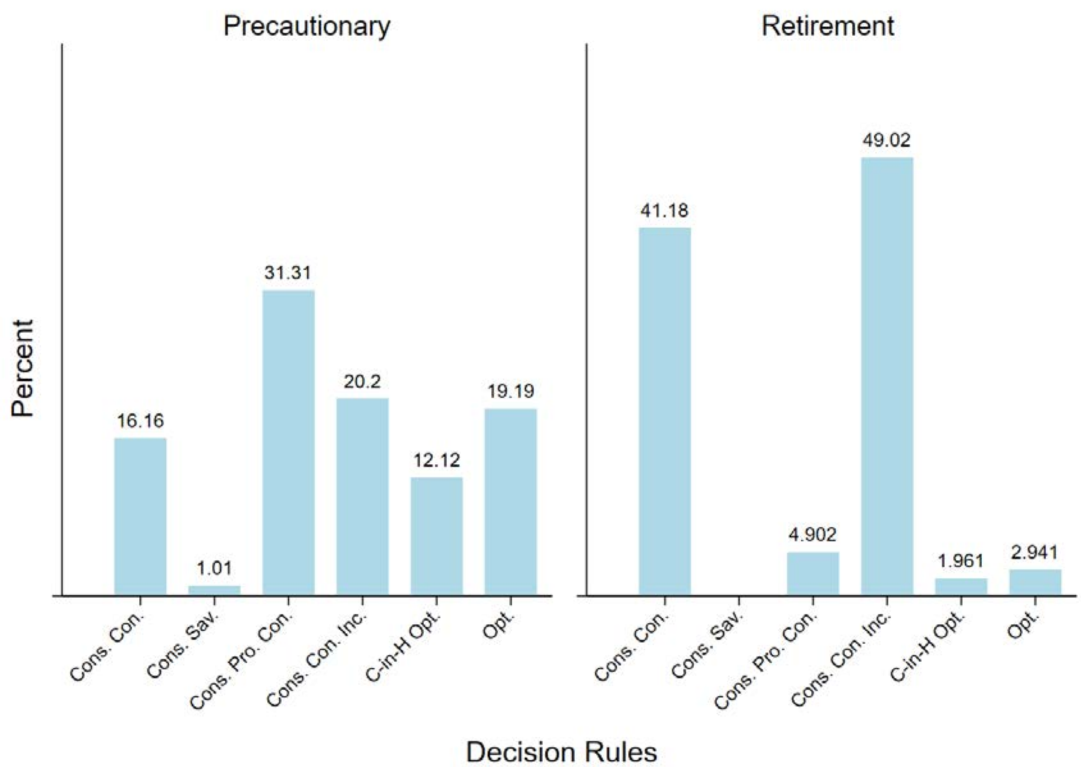

rule) together occur approximately as often as the modal rule. This suggests that subjects are rather good at smoothing for precautionary reasons.

The right-hand panel presents the results for retirement savings, and the treatment effect is stark. A chi-square test rejects the null hypothesis that the two distributions in Figure 1 are identical (Fisher's exact test p-value = 0.00), and the effect of the addition of retirement savings is clear from the figure. The best fitting rules become almost exclusively Rules 1 and 4: constant consumption and constant consumption conditional on current income level. When subjects are forced to consider retirement savings in addition to precautionary savings they revert to perhaps the simplest of rules. Why are these two rules simple? Because the choice each period is the level of consumption, so that a heuristic that keeps the actual choice constant requires little or no computation. Rule 4 is slightly more complex than Rule

Figure 5: Best Fitting Decision Rules by Treatment and Income Draw

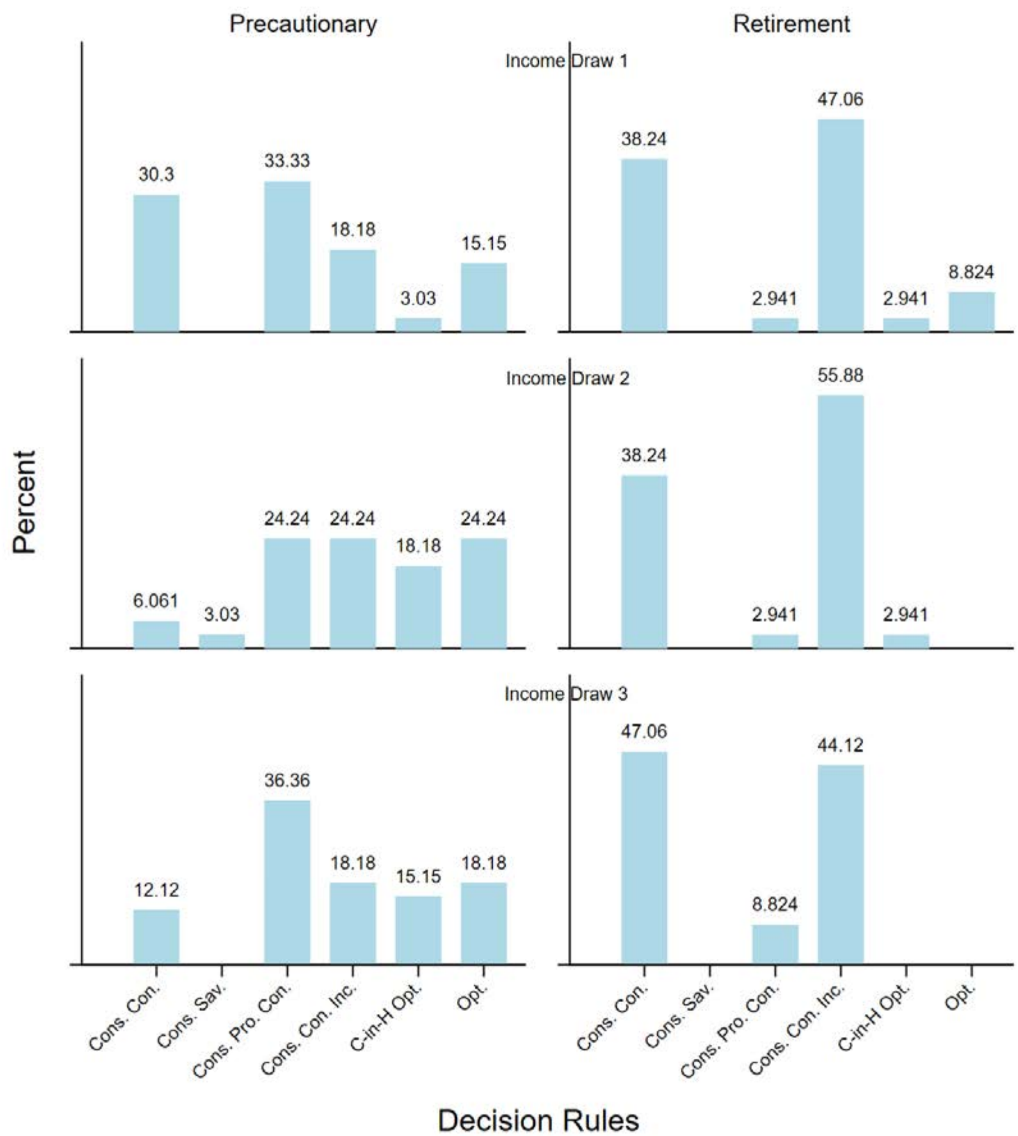

1 because it requires a different level of consumption for each income level. However, the fact remains that subjects switch out of all other rules and into constant consumption rules when retirement savings is required.

Since many of the rules that subjects abandon in favor of the rules that specify constant consumption are very good decision rules (recall that Rules 5 and 6 are optimal conditional on current cash-in-hand and optimal for the entire game), it seems likely that a loss of utility occurs when subjects must save for retirement. We compare the performance of all decision rules that best describe the data relative to optimal behavior. We use the measure of efficiency in Noussair and Matheny (2000), where we express the total actual payoff as a ratio of the optimal payoff for a subject in each individual savings task.[6] The efficiency measure shows what fraction of potential earnings was actually achieved by following a particular decision rule. An efficiency of one is equivalent to optimal consumption smoothing. Table 2 shows the average efficiency for each decision rule by treatment.

Table 2: Average Realized Efficiencies of Decision Rules by Treatment

Decision Rule	**Precautionary**	**Retirement**
Cons. Con	0.89 (0.18)	0.88 (0.24)
Cons. Sav.	0.63 (1 obs)	none
Cons. Pro. Con.	0.90 (0.11)	0.88 (0.05)
Cons. Con. Inc.	0.86 (0.12)	0.80 (0.19)
C-in-H Opt.	0.97 (0.03)	0.97 (0.01)
Opt.	0.96 (0.04)	0.97 (0.02)
All	0.91 (0.12)	0.84 (0.21)

Standard deviation in parentheses.

The table reveals a clear loss in efficiency under the retirement savings regime. To see

[6] Carbone and Duffy (2014) and others compute the mean squared deviation from the conditional optimal path to statistically test for treatment effects. We do not report this statistic because our two treatments result in different paths, one of which is smoother than the other.

this, note that Rules 3, 5 and 6 have realized efficiencies of 0.90, 0.97, and 0.96. These rules, which are present in the precautionary savings treatment, virtually disappear in the retirement savings treatment. The rules that increase in frequency, i.e., that take the place of these three rules in the retirement savings treatment, are Rules 1 and 4 with efficiencies of 0.88 and 0.80. The table also shows that average efficiency is lowered by 7% in the retirement savings treatment, a two-sided t-test rejects the null hypothesis that these mean efficiencies are identical (p-value of 0.0065). Apparently the added complexity of the task results in simple rules, and these rules come at the expense of the loss of better rules.

To ensure that we did not accidentally average these results over the three income draws, we broke out the decision rule analysis by income draw. The result is shown in Figure 5, and the characterization of the results remains the same: participants switched out of optimal rules and into constant consumption rules from precautionary to retirement savings treatments in each of the three draws. There is some evidence that the draws influenced precautionary savings, as constant consumption was far more prevalent with early income draws were relatively high. But otherwise, our conclusion that retirement savings induces constant consumption heuristics stands.

What are the consequences of using constant consumption rules for retirement savings? To answer this question, recall the histograms for retirement savings decision rule distributions in Figure 5. It is evident from the figure that three rules are prevalent in retirement savings: constant consumption, constant consumption conditional on income, and constant propensity to consume. The two optimal rules virtually disappeared under retirement savings, and the constant savings rule was not a factor in any treatment.

We went back to the data and we plotted the time series of the best performing constant consumption, constant propensity to consume, and constant consumption conditional on income rules from the data. The results are presented in Figure 6. The panels are arranged as before left to right. From top to bottom the panels present the result for the best constant consumption rule, constant propensity to consume rule, and constant consumption conditional on income rule for the third (representative) income draw.

Figure 6 shows that given the three rules available to switch into (since constant savings was not induced by the experiment), the two rules that dominated have the capability of doing very well, and visually obviously better than the constant propensity to consume rule. The best constant consumption rules result in extremely smooth consumption, while the best constant propensity to consume (which consumes a constant proportion of cash-in-hand) consumes more as cash-in-hand builds up over time, resulting in the positive slope, and over-consumption during retirement (i.e., under-consumption during working life).

We next plotted the average consumption path for each of these three rules in the final income draw and present the results in Figure 7. Note that these are average results from the rules with the best likelihood for all participants, regardless of how well the participants played the game, thus we cannot construct a hypothesis on which rule will do better on average. However, we can conclude something about the type of play they represent.

Beginning at the top of Figure 7, the constant consumption rule does extremely well for retirement savings. Note that spike in consumption in the last period is due to one participant who consumed virtually nothing the entire game until the last period. The middle panel shows the characteristic of the constant propensity to consume: it gradually builds up consumption as assets are built up. For a representative income draw, this is not the best rule to follow.

On average, participants whose behavior looked most like constant income conditional on income level did not consume enough during low income periods, resulting in over-savings overall. Thus in the aggregate, the simple constant consumption rule is associated with the smoothest consumption over the life cycle.

Figure 6: Best Decision Rule Performances in the Representative Income Draw

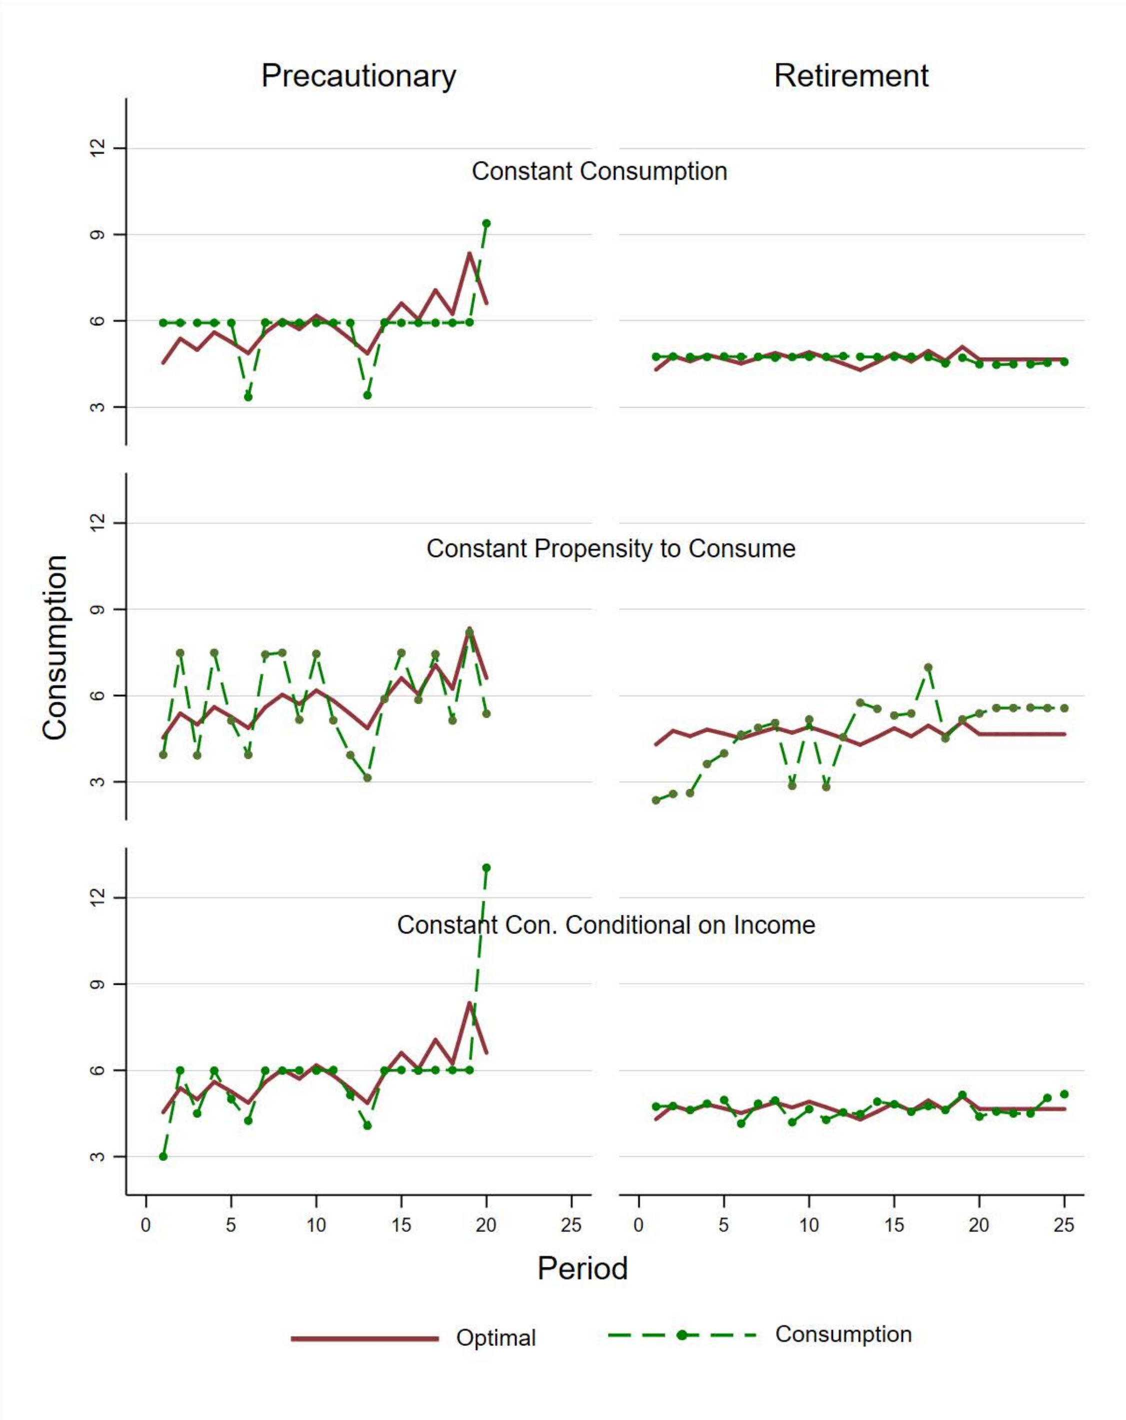

Figure 7: Average Decision Rule Performances in the Representative Income Draw

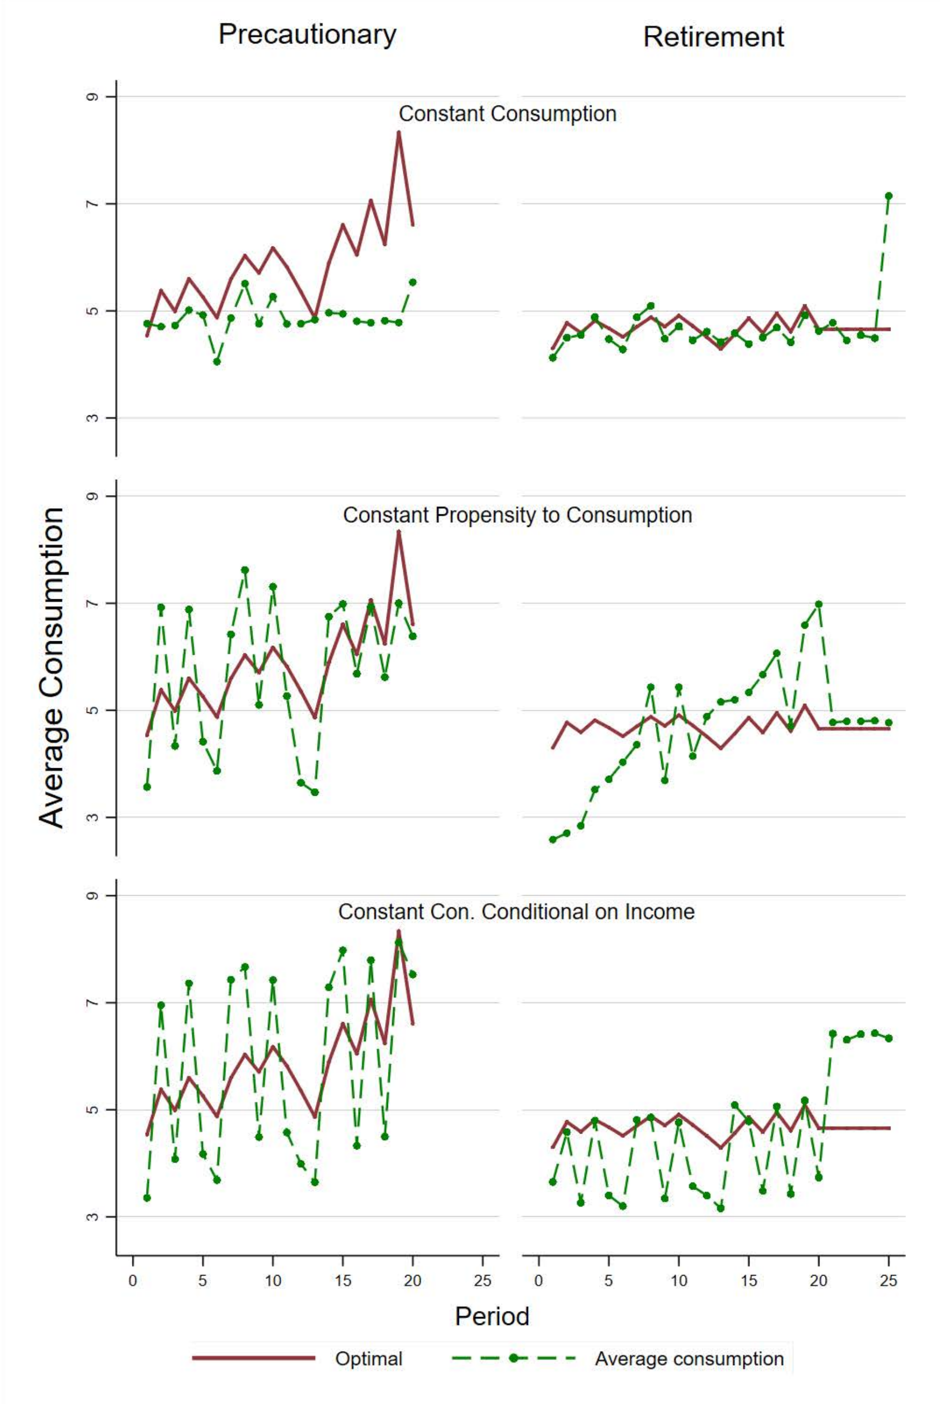

5 Conclusion

We presented an experiment to identify how behavior changes when retirement savings is added as a motive to precautionary savings. We inferred decision rules from choice data in an experiment under both precautionary and retirement savings motives. In our experiment, subjects were presented a random process for their income for a finite period of time and were given a financial incentive to smooth their consumption. The difference between precautionary savings and retirement savings treatments was simply and precisely the addition of five unemployed years at the end of life. Our experiment abstracted from self-control issues, investment decisions, and a host of other factors that affect savings decisions. Our experiment focused on the quality of the consumption smoothing decision.

We found that savings for precautionary purposes was of a high quality, compared to the optimal choice, with an average efficiency of 91%. More than 30% of the decision rules we inferred on the choice data were optimal or conditional-optimal. When subjects saved for retirement as well as for precautionary reasons, the optimal decision rules were replaced by constant consumption rules. The result was a statistically significant loss in efficiency of 7% due to over-saving for retirement.

When participants demonstrated a constant propensity to consume out of cash-in-hand, our experiment revealed that the consequences can be an upward-sloping consumption path during the life-cycle, which was contrary to the economic incentives of the experimental design. When participants conditioned their consumption on period income, they over-saved during periods of low income, resulting in over-saving during their working life.

Our results are the first to show that subjects adjust to the added complexity of saving for retirement with the simplest of decision rules, which do not require computation to implement, since they operate directly on the choice subjects are making. Our results suggest that there is room for improvement in the quality of the retirement savings decision, and our experimental design admits the study of institutional changes that have the potential to improve this choice.

References

Allen, T, and C. Carroll (2001). "Individual Learning about Consumption", *Macroeconomic Dynamics* 5:255-271.

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 (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.

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

Campbell, J. Y. (2016). "Restoring rational choice: The challenge of consumer financial regulation", *American Economic Review*, 106(5), 1-30.

Campbell, J., and N. Mankiw (1989). "Consumption, income, and interest rates: Reinterpreting the time series evidence", *NBER macroeconomics annual*, 4, 185-216.

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. (2005). "Demographics and behaviour", *Experimental Economics* 8:217-232.

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

Chetty, R., J. Friedman, S. Leth-Petersen, T. Nielsen, and T. Olsen (2014). "Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark", *The Quarterly Journal of Economics* 129(3), 1141-1219.

Choi, J., D. Laibson, and B. Madrian (2011). "$100 bills on the sidewalk: Suboptimal investment in 401 (k) plans', *Review of Economics and Statistics* 93(3), 748-763.

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

Duffy, J. (2008). “Macroeconomics: a survey of laboratory research”, *Handbook of experimental economics, 2*.

Feltovich, N., and O. Ejebu (2014), “Do positional goods inhibit saving? Evidence from a life-cycle experiment”, *Journal of Economic Behavior and Organization* 107:B:440-454.

Gali, J., J. Lopez-Salido, and J. Valles (2007). “Understanding the effects of government spending on consumption”, *Journal of the European Economic Association* 5(1), 227-270.

Gross, D., and N. Souleles (2002). “Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data”, *The Quarterly journal of economics* 117(1), 149-185.

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

Houser, D., and J. Winter (2004). “How do Behavioral Assumptions Affect Structural Inference? Evidence From a Laboratory Experiment”, *Journal of Business, Economics and Statistics* 22:64-79.

Houser, D., M. Keane, and K. McCabe (2004). “Behavior in a Dynamic Decision Problem: An Analysis of Experimental Evidence Using a Bayesian Type Classification Algorithm”, *Econometrica* 72:781-822.

Kahneman, D. and A. Tversky (1974). “Judgment under Uncertainty: Heuristics and Biases”, *Science* 185:1124-1131.

Lusardi, A., and O. Mitchell (2007). “Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth”, *Journal of monetary Economics* 54(1), 205-224.

MacFarland, D. C. Marconi, and S. Utkus (2004). “?Money Attitudes? and Retirement Plan Design: One Size Does Not Fit All”, *Pension design and structure: New lessons from behavioral finance* 97-120.

Mankiw, N. G. (2000). “The savers-spenders theory of fiscal policy”, *American Economic Review* 90(2), 120-125.

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

Müller, W. (2001). “Strategies, heuristics, and the relevance of risk-aversion in a dynamic decision problem”, *Journal of Economic Psychology* 22:493-522.

Noussair, C. and K. Matheny (2000). “An Experimental Study of Decisions in Dynamic Optimization Problems”, *Economic Theory* 15:389-419.

Winter, J.K., K. Schlafmann and R. Rodepeter, Ralf (2012). "Rules of Thumb in Life-cycle Saving Decisions", *Economic Journal* 122(560): 479-501.