

Lifecycle Consumption Under Different Income Profiles: Evidence and Theory*

John Duffy[†] Yue Li[‡]

December 12, 2018

Abstract

We report on a series of economic decision-making experiments exploring how individuals make lifecycle consumption and saving plans when they face different income profiles, representing different pension replacement rates. We aim to assess whether variations in pension replacement rates might aid or hinder individuals' ability to make good lifecycle consumption and saving plans. We find that pension replacement rates matter for subjects' experimental payoffs and consumption behavior. In particular, our treatment with a 100% pension replacement rate yields the highest experimental payoff, and more subjects in this treatment choose the status quo strategy of consuming endowments in every period. We show that a model of rational inattention is useful for explaining subjects' responses to different pension replacement rates.

Keywords: Lifecycle model, consumption and savings, retirement planning, public pensions, behavioral and experimental economics, rational inattention.

JEL Codes: C91, D14, D91, H55.

*We thank three referees and editor Thomas A. Lubik for their valuable comments. We also thank seminar participants at SUNY Albany, the 2016 Barcelona GSE Summer Forum, and the 2017 AEA meeting for valuable comments and suggestions. We gratefully acknowledge funding for this project from the UC Irvine School of Social Sciences. The experimental protocol was approved by UC Irvine's Institutional Review Board, HS#2015-2082

[†]Department of Economics, University of California, Irvine. Email: duffy@uci.edu.

[‡]Department of Economics, University of Albany, SUNY, Email: yli49@albany.edu.

1 Introduction

Due to rising life expectancy and falling fertility, many developed countries with pay-as-you-go public pension programs are facing the possibility that promises made to future generations are un-sustainable, so that major changes to these programs will have to be made.¹ One much discussed change is to the pension replacement rate, i.e., the rate of pension benefits to pre-retirement after-tax average earnings. In this paper we use experimental methods to assess whether variations in pension replacement rates might aid or hinder individuals' ability to make good lifecycle consumption and saving plans. Our experiments with different pension replacement rates are timely from a policy perspective but would be difficult to conduct in the field and so we resort to laboratory experiments with paid human subjects.

The subjects in our experiment are induced to hold concave preferences over consumption in each of the 25 periods of their lifetimes so that lifecycle consumption smoothing is desirable according to the standard rational actor approach. Within this environment, we consider how subjects make consumption and saving decisions when facing several different lifecycle profiles for their income. These lifecycle income profiles are known to subjects with certainty, and some of these profiles involve explicit drops in income in the final periods of the lifecycle in order to mimic economies where pension replacement rates are less than 100 percent. To allow subjects an opportunity to learn, we have them participate in *two* 25-period lifetimes (sequences) under the same lifecycle income profile. Our assumption that the lifecycle income profile is deterministic and known in advance is clearly an idealized, best-case-scenario, but it provides us with an important benchmark for more rigorous analysis; if subjects are unable to intertemporally smooth their consumption and save for retirement in this very simple environment, then such difficulties are likely to be further compounded in settings with uncertain and highly variable income processes.

Our main innovation is to consider several different income profiles that have the same present value but which differ in the pension replacement rate. The optimal consumption path is the same across all four income profiles we consider, enabling us to easily evaluate the extent to which the pension replacement rate policy matters for lifecycle consumption and saving plans, which is our primary aim in designing these experiments. To preview our results, we find that under all lifecycle income profiles that we consider, subjects over-consume in the early periods of the lifecycle and under-consume in the later periods of the lifecycle relative to both the *unconditionally* and the *conditionally* optimal paths. On the

¹See, e.g., "Europe Faces Pension Predicament," *The Wall Street Journal*, March 6, 2016 (Samuel, 2016).

other hand, we also find evidence that subjects improve their performance in the direction of the optimal consumption path with repeated experience under the same endowment (income) profile, a finding that would be difficult to obtain in the field. This evidence that subjects can learn from experience suggests that proxies for experience, such as financial education/advice, may be a valuable. Most importantly, we find that the particular lifecycle endowment profile that subjects face matters for their experimental payoffs and consumption behavior. In particular, our treatment with a constant lifecycle endowment profile (100% pension replacement rate) yields the highest experimental payoff, and more subjects in that treatment choose the status quo strategy of consuming endowments in every period. While that strategy is *not* the optimal policy, it is not as far from the optimal policy as it would be a strategy of consuming endowments in the other treatments that we consider.

A second aim of our paper is to explore models that can rationalize the observed treatment difference in our experiment so that we can better understand agents' reactions to policy changes in pension replacement rates. The first model we consider is similar in spirit to the rational inattention model developed by Sims (2003); Matejka and McKay (2014); Carvalho and Silverman (2017), and Gabaix (2017). The version of the rational inattention model we consider allows subjects to choose whether to optimize or to follow a simple, consume endowments heuristic, and takes into account the cost of information processing, the ability to solve the intertemporal lifecycle problem and the loss from following the simple heuristic of consuming endowments. This rational inattention model yields results that help explain the treatment differences that we observe in our experiment. Specifically, we find that when the cost of using the consume endowments heuristic is low relative to the optimal policy (as in our constant endowment profile treatment), some subjects choose to use this heuristic. When the cost is higher, as when income drops to 0 in retirement, all subjects opt out of the status quo strategy of consuming endowments. We also consider several alternative explanations for the variations in consumption that are observed over the different lifecycle income profiles in our experiment but we do not find compelling evidence for these alternative explanations.

The remainder of the paper proceeds as follows. Section 2 reviews the literature. Section 3 describes the theoretical framework that motivates the design of our experiment. Section 4 explains the experimental environment, choice of model parameters and testable hypotheses. Section 5 reports our main experimental findings. Section 6 explores several explanations for the observed deviations from the rational choice theory. Finally, Section 7 provides a summary and suggestions for future research.

2 Related literature

This paper builds upon and complements three strands of existing research. The first strand, and the one closest to our approach, uses laboratory experiments to study consumption and saving decisions of individuals over a finite lifetime. Such papers are listed in Table 1. This table also provides a summary of some of the important design features of these

Table 1: Summary of Experimental Literature on Lifecycle Consumption/Saving Decisions

	Induced Utility (1)	Stochastic Income (2)	Retirement Phase (3)	Replacement Rate change (4)
Fehr and Zych (1998)	Y	N	N	N
Hey and Dardanoni (1988)	Y	N	N	N
Anderhub et al. (2000)	Y	N	N	N
Johnson et al. (2001)	N	N	Y	N
Ballinger et al. (2003)	Y	Y	N	N
Carbone and Hey (2004)	Y	Y	N	N
Carbone (2006)	Y	Y	N	N
Brown et al. (2009)	Y	Y	N	N
Ballinger et al. (2011)	Y	Y	N	N
Hey and Knoll (2011)	Y	Y	N	N
Carbone and Duffy (2014)	Y	N	N	N
Feltovich and Ejebu (2014)	Y	Y	N	N
Meissner (2014)	Y	Y	N	N
Koehler et al. (2015)	N	Y	Y	N
Levy and Tasoff (2015b)	Y	N	N	N
Meissner and Rostam-Afschar (2017)	Y	Y	N	N
Tasneem and Warnick (2018)	Y	Y	Y	N
Tasneem et al. (2018)	Y	Y	Y	N
This paper	Y	N	Y	Y

experimental studies, which we use to distinguish and differentiate the contribution of this paper. Specifically, Column 1 of Table 1 reports on whether or not experimental payments were determined by using an induced utility function. Inducing a utility (or payoff) function over consumption choices has several advantages. First, it incentivizes agents to exert effort to think about the tradeoffs they face from consumption versus savings. Second, it enables determination of the optimal consumption/saving path over the lifecycle and the extent of subjects' deviations from that path. As Table 1 reveals, most studies of lifecycle consumption choices induce a utility function. Column 2 reports on whether subjects faced a stochastic, as opposed to a deterministic income process over the life cycle. Column 3 reports on whether subjects faced a deterministic drop off in income in the later periods of the lifecycle mimicking a retirement phase. Finally, column 4 reports on whether the ratio of pension

income to wage income, i.e., the pension replacement rate, was varied across treatments.

As Table 1 makes clear, our experiment differs in design (as well as in focus) from this literature in several dimensions. First, most prior experimental studies do not consider an explicit retirement phase where non-capital income is subject to a large drop-off in the later periods of life as often occurs when agents stop working upon retirement. As a consequence, in most prior experimental studies of lifecycle consumption/saving decisions there are only *precautionary*, as opposed to *retirement* motivations for savings. Second, prior experimental studies have typically employed stochastic processes for lifecycle income in order to generate precautionary motives. While such processes are certainly realistic, it is not clear whether subjects deviate from the optimal consumption path because they misunderstand the probabilities of different income possibilities or because they do understand these probabilities but find it more difficult to calculate the optimal solution in the face of stochastic income. By contrast, we consider simpler settings with a deterministic lifecycle income process, which should work to minimize misunderstandings associated with the stochastic nature of income processes and which serves to focus attention on the intertemporal consumption smoothing objective, in particular, to contemplate a strategy that smooths consumption against the perfectly known drop-off in income in the retirement phase of life. Finally, and most importantly, none of these prior experimental studies consider how variations in the pension replacement rate affects lifecycle choices, since most of these studies do not have explicit retirement phases.² Indeed, as noted in the introduction, changes in such replacement rates are an active topic of discussion among policymakers and so our design varying replacement rates can contribute to that discussion.

A main finding from the experimental studies using induced utility (such as this one) is that subjects, on average, over-consume relative to the unconditional optimum in the early periods of life (For instance, see Fehr and Zych (1998) and Ballinger et al. (2003)). Consequently, they accumulate too little wealth and so they under-consume relative to the unconditional optimum in later periods of life. One concern with studying deviations from the *unconditionally* optimal path is that errors made in one period not only affect that period's deviation but also have a permanent effect on all future deviations, even if subjects

²Johnson et al. (2001) and Koehler et al. (2015) do include a retirement phase, but they do not design treatments with different replacement rates or induce subjects to hold a certain utility over differing consumption amounts. In concurrent research (that we were not aware of when we started this project), Tasneem and Warnick (2018) and Tasneem et al. (2018) use an inducted utility approach and consider a retirement phase that provides zero endowment to, separately, understand the effect of adding a retirement saving motive and of adding a default saving rate, but none of these papers consider the effect of changing pension replacement rates.

behave optimally from that period onwards. To address this concern, the previous literature using induced utilities also calculates the deviation of consumption from the conditional optimum. The evidence regarding the latter benchmark is mixed: Fehr and Zych (1998) and Ballinger et al. (2003) find that such deviations are almost always positive, while Brown et al. (2009) finds that such deviations are positive in early periods and become negative in the later periods of the lifecycle. One aim of this paper is to examine the patterns of life-cycle consumption relative to both unconditional and conditional optimal paths in a setting that has a retirement phase, which better represents the typical lifecycle income profile that many agents face.

A second strand of the related literature uses survey or field experiment methods to evaluate the efficacy of various interventions designed to aid in retirement planning (Brown et al., 2017, 2016; Liebman et al., 2015; Goda et al., 2014). Our research complements these studies by focusing on the effects of pension replacement rates on lifecycle consumption and saving decisions in the context of laboratory experiments.

A final strand of the literature uses behavioral models to explain various puzzles of consumption and saving decisions in the data. For instance, Laibson (1997) and O’Donoghue and Rabin (1999) propose that individuals may exhibit hyperbolic discounting leading them to consume more today than is optimal if no commitment device is available. Loewenstein and Prelec (1992) argue that the utility function involves a reference point such that the taste for gains and losses are asymmetric and framing and prior expectations affect intertemporal choices. Shefrin and Thaler (1988) introduce mental accounting to the lifecycle problem and argue that individuals treat different components of wealth to be non-fungible. Levy and Tasoff (2015a) argue that individuals may under-predict the compound interest earned on investments (exponential-growth bias), and hence save too little. In this paper, we apply a version of the rational inattention approach of Sims (2003), Matejka and McKay (2014), Carvalho and Silverman (2017) and Gabaix (2017) to study the relationship between consumption/saving decisions and pension replacement rates. Using this model, we demonstrate that given differences in information processing costs, some subjects rationally choose to behave like hand-to-mouth consumers. Further, the share of such subjects is non-decreasing in the pension replacement rate. To our knowledge, this is the first dynamic optimization experiment that uses a rational inattention model to describe subjects’ consumption behavior.³

³Carvalho and Silverman (2017) design an experiment to study portfolio choices in a *static* game and also find evidence that is in support of the rational attention model.

3 Theoretical Framework

Our theoretical framework is the standard, intertemporal model of lifecycle consumption and saving choices originating in the work of Modigliani and Brumberg (1954) and Friedman (1957). We adopt this framework as it remains the workhorse approach to the modeling of household consumption and saving behavior, and we wish to present our results using a framework that is familiar to this audience, as we hope to move the discussion of lifecycle planning in a direction that takes greater account of departures from the standard rational choice framework. We further restrict attention to the case of complete markets and no uncertainty as we wish to make the environment as simple as possible for our human subjects. As noted in the introduction, the world is far more complex—income is uncertain and markets are incomplete—but we wish to start with the simplest possible framework; if subjects are not able to optimally save for retirement in this simple setting, then it is unlikely that they will do better in more complicated settings.

The theoretical framework can be described as follows. Each household i makes consumption and saving decisions in each of periods $p = 1, 2, \dots, P$, where the final period, P , is perfectly known. Household i 's endowment process over the lifecycle is also perfectly known, and is denoted by $\{e_{ip}\}_{p=1}^P$. Household i 's objective is to:

$$\max_{\{c_{ip}, a_{i(p+1)}\}_{p=1}^P} \sum_{p=1}^P u(c_{ip})$$

subject to

$$c_{ip} + \frac{a_{i(p+1)}}{1+r} = e_{ip} + a_{ip}, \quad \forall p$$

$$a_{i(p+1)} \geq 0 \quad \forall p, \text{ and } a_{i1} = 0.$$

Here, c_{ip} denotes household i 's period p consumption, $a_{i(p+1)}$ denotes household i 's initial assets for period $p + 1$, and $r > 0$ is an exogenous fixed and known rate of interest; again, with the aim of simplicity, we are thus considering a partial equilibrium environment. We also impose a no-borrowing constraint, which should not be binding given a positive r and the considered income profiles; again the rationale is simplicity, though we recognize that borrowing constraints can be empirically important in lifecycle consumption and saving decisions.⁴ Notice further that we are ignoring any time discounting (i.e., we are setting

⁴Allowing for borrowing, we would have to specify borrowing constraints, which would further complicate

the discount factor, $\beta = 1$) as our experiment will be conducted over several hours in a laboratory setting.⁵

The utility function $u(\cdot)$ is assumed to be increasing and concave, and represents the payoff function by which subject's period-by-period consumption choices earn them monetary payments. Compared with an alternative setting without induced utility, the introduction of $u(\cdot)$ yields a unique, optimal consumption path against which we can compare the behavior of our subjects.⁶ We adopt a concave utility function specification as this implies that intertemporal smoothing of consumption is desirable (as in the lifecycle theory).

Given the concavity of the utility function and given a lifetime endowment *sequence*, $\{e_{ip}\}_{p=1}^P$, it is straightforward to calculate, by working backwards, the solution to the maximization problem stated above, which we refer to as the unconditionally optimal path and denote by $\{c_{ip}^{**}, a_{i(p+1)}^{**}\}_{p=1}^P$. Recognizing that subjects may make decision errors over their lifetimes, we will also consider their behavior relative to the *conditionally* optimal consumption path, which we denote by $\{c_{ip}^*\}_{p=1}^P$. Formally, c_{ip}^* involves re-optimization at each new period, p , of the lifecycle, based on *current* asset holdings, and is the solution to:

$$\max_{c_{ip}} \sum_{j=p}^P u(c_{ij})$$

subject to

$$c_{ij} + \frac{a_{i(j+1)}}{1+r} = e_{ij} + a_{ij}, \quad \forall j \geq p$$

$$a_{i(j+1)} \geq 0, \quad \forall j \geq p \text{ and } a_{ip} \text{ is given.}$$

the decision space of subjects. The optimal path for the model that we do implement does not require any borrowing.

⁵In general, the assumption of no discounting is consistent with an economy where there are two types of assets. In these models, the discount factor β is usually set to $\frac{1}{1+r^f}$, where r^f is the return on risk-free assets, which is close to zero. The interest rate considered in the model, r , resembles the return on risky assets (e.g., a stock market portfolio), though there is no risk associated with such an investment in our model.

⁶It would, of course, be interesting to study how individuals make consumption and saving decisions over their lifetimes given their own, "homegrown" preferences, but it might be difficult to determine the nature of those preferences and thus the optimal consumption path and hence we decided to induce preferences.

4 Experimental Design and Hypotheses

In our experiment, we set $P = 25$ and $u(c) = 0.2 \ln(0.01c + 1)$. Assuming that each period in the model represents 2.3 years, the choice of $P = 25$ periods corresponds to an economy where people begin life at age 23 and exit at age 79. More precisely, the first 17 periods correspond to ages 23-60 and represent the time of the lifecycle when people work and receive (after-tax) earnings; the last 8 periods correspond to ages 61-79 and represent the time when people are retired and receive pension benefit income or consume out of their accumulated asset positions. The implied coefficient of relative risk aversion is $\frac{0.01c}{1+0.01c}$. Over the range of optimal consumption amounts in our experiment (see Figure 1a below), this coefficient of relative risk aversion ranges from a value of 0.56 when $c = 126$, to a value of 0.95 when $c = 2121$. There is some support for this functional form in the empirical literature. For instance, using 33 sets of estimates of wage and income elasticities, Chetty (2006) reports that the implied value of relative risk aversion for an additive utility function ranges from 0.15 to 1.78 and has a mean of 0.71. The interest rate, r , was set at 10 percent per model period and remained constant across all of our treatments; considering that a model period consists of 2.3 years, this choice for the interest rate implies an average annual return of 4.5 percent on long term investments, which is consistent with evidence found in Munnell et al. (2013).

We consider four different lifecycle endowment profiles, which comprise our four main experimental treatments. All four endowment profiles have the *same* present value, that is, for any two treatments $i \neq j$, it was always the case that:

$$\sum_{p=1}^P \frac{e_{ip}}{(1+r)^p} = \sum_{p=1}^P \frac{e_{jp}}{(1+r)^p}.$$

Our four treatments differ according to the distribution of endowments over the lifecycle, or, equivalently, the pension tax and benefit scheme. For instance, the endowment flow in our first treatment is a constant 500 “tokens” for the first 17 (working) periods, but in the last 8 (retirement) periods, the endowment amount drops to 200 tokens. This first, baseline treatment thus involves a 40 percent pension replacement rate, which we chose to capture the replacement rate of Social Security benefits in the United States for a worker who earns the median U.S. income (Feldstein, 2005). For ease of reference we will henceforth refer to this treatment as “R40,” reflecting the 40 percent replacement rate. Treatment two, labeled “R0”, is designed to represent an economy with a 0 percent pension replacement rate, in

which subjects receive a higher endowment of 526 tokens during the first 17 working periods as compared to R40 since they pay no tax, but they now receive a zero endowment in each of the final 8 periods of retirement. Hence, subjects in R0 must actively save during their working years in order to have any consumption in the retirement phase of their lives. Treatment three, labeled as “LS”, is an extreme case where subjects receive a single lump-sum endowment of 4,644 tokens in the very first period of their lives and 0 in each of the remaining 24 periods, an alternative way of implementing a 0 pension replacement rate. Finally, treatment four, labeled as “R100”, represents the opposite extreme case where subjects receive a constant endowment of 465 tokens in each of the 25 periods of their lifetimes, which is lower than the endowment received during the working periods of R40 and R0. Treatment R100 can be viewed as a pension system that uses a greater payroll tax to fund a more generous retirement plan that has a 100 percent pension replacement rate. Figure A2 in Appendix B provides a visual illustration of these four different endowment processes, which were also shown and explained to subjects in our experiment.

The induced utility function serves as a mapping between subjects’ consumption choice (the number of tokens they consume) and their period money earnings. In each period, subjects are asked to decide how many tokens they wish to convert into money using the conversion rate implied by the utility function. The remaining amount of tokens are automatically saved for the future and earn the fixed interest rate of 10% per period. Following the completion of the final, 25th period, any tokens that were not converted into money were lost (had zero redemption value) and this fact was made known to subjects in the written instructions. This payment structure implies that maximization of monetary payoff in the experiment is equivalent to maximization of lifetime utility in the theoretical model. A sample screenshot for the decision process as well as the instructions distributed to subjects at the beginning of the experiment are provided in Appendix A.

In each session, subjects were first assigned to one endowment treatment condition for two, 25-period “sequences” (lifetimes), and then they are assigned to a second endowment treatment for an additional two, 25-period sequences. Thus, altogether, each subject makes choices in four, 25-period lifetimes under two different endowment treatments. We repeated each endowment condition twice to allow for some learning by subjects. In each experimental session, we implemented just one of four different treatment orders: R40-R0, R0-R40, LS-R100, or R100-LS. Thus, subjects who participated for example in a session with treatment order R40-R0 played two, 25-period sequences under the endowment profile of treatment 1 (R40) and then played 2 additional 25-period sequences under the endowment profile of

treatment 2 (R0). In all cases, subjects were instructed about the endowment profile of the first treatment, and following completion of the two lifecycle sequences under that endowment profile, they were instructed about the endowment profile of the second treatment. The change in endowment profile is the only change between treatments; all other factors, i.e., the utility function, the rate of interest on savings, and the number of periods of decision-making (25) were held constant.

Subjects' actual monetary payoff was the sum of their experimental lifetime utility from two, randomly selected sequences (lifetimes), one separately from each of the two treatments of a session, plus a show-up payment. Since subjects did not know which one of the two sequences they played of each treatment would be randomly chosen for payment, they were incentivized to perform their best in each sequence. Subjects were paid in cash at the end of a session with the average payment being around \$24 per subject for a two-hour experiment.

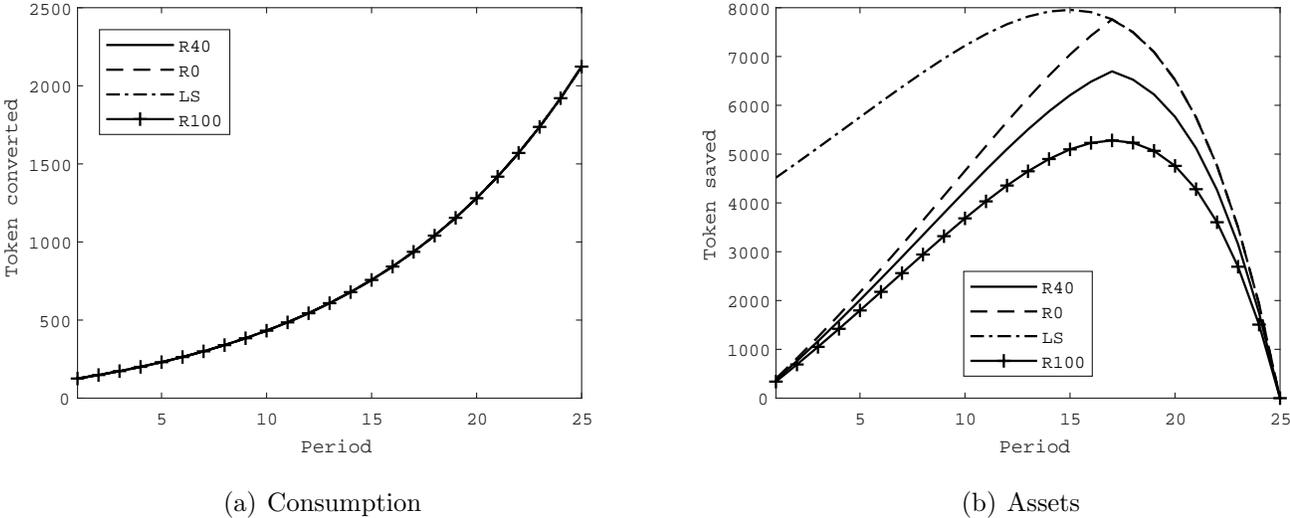


Figure 1: Optimal decisions implied by standard rational choice theory

Given the model parameters, Figure 1 plots the optimal lifecycle consumption and asset profiles for each treatment. Notice that the no borrowing constraint is never binding if subjects behave rationally. As shown in Figure 1(a), the optimal consumption path is upward sloping, which is due to the payment of interest on savings and the absence of discounting. Since the present value of the four different lifecycle income processes is, by design, exactly the same, the optimal consumption path shown in Figure 1(a) is also the same across all four treatments, enabling us to make comparisons of experimental decisions relative to this unique optimal path across our four treatments. However, since the distribution of endowments

differs by treatment, the optimal *asset* path differs across the four treatments as shown in Figure 1(b). In particular, to offset for the sharp drop-off in endowments in the later periods of life, subjects need to save more in treatments R0 and LS, as compared with treatments R40 and R100.

Our experiment was designed to test the following hypotheses:

Hypothesis 1 *Subjects' choices are no different between the first and second sequence of each treatment.*

This first hypothesis follows from the fact that the second sequence of each treatment is always a repeat of the lifecycle problem faced in the first sequence. Theoretically, repetition should not matter, and in reality agents live only a single lifetime. However, if repetition is found to improve subjects' choices, then such a result should be of interest to policy makers, as it may be possible to proxy for experience through better financial education or nudges to save more for retirement. The next hypothesis concerns predictions about subjects' consumption choices over the lifecycle that follow from our induced utility approach:

Hypothesis 2 *Subjects' consumption choices follow the unconditionally optimal path shown in Figure 1, with savings adjusting optimally as well.*

As an alternative to Hypothesis 2, we also consider:

Hypothesis 3 *Subjects' consumption choices are conditionally optimal, involving re-optimization at each new period, p , of the lifecycle given current wealth.*

Finally, given our parameterization of the model, the optimal consumption path, and thus, predicted earnings, are the same across all four treatments. Thus, even if subjects' behavior deviates from rational choice model predictions, we would expect such deviations to be similar across our four treatments. However, if the distribution of lifecycle income matters for subjects' choices, then that finding should also be of interest to policymakers in thinking about the consequences of implementing different pension replacement rates.

Hypothesis 4 *Given our parameterization of the model, there should be no difference in earnings or consumption behavior across our four treatments.*

The experimental subjects were all undergraduate students at the University of California, Irvine who had no prior experience participating in any of our experimental treatments.

No subject was allowed to participate in more than one treatment order (experimental session). The raw data are composed of 28, 32, 30, and 29 subjects for each of the four treatment orders R40-R0, R0-R40, LS-R100, and R100-LS, respectively (for a total of 119 subjects). As all of our treatments are individual-choice (there are no interactions between subjects), we treat each subject in a 25-period sequence as a single observation. Table 2 reports on some characteristics of our design.

Table 2: Characteristics of experimental sessions

	Lifetime endowment profile (1)	No. obs. (2)	Average earnings (3)	Potential earnings from converting endowments (4)
R40	500 for $p = 1, 2, \dots, 17$, 200 for $p = 18, 19, \dots, 25$	120	\$8.50	\$7.85
R0	526 for $p = 1, 2, \dots, 17$, 0 for $p = 18, 19, \dots, 25$	120	\$8.56	\$6.24
LS	4,644 for $p = 1$, 0 for $p = 2, 3, \dots, 25$.	118	\$8.31	\$0.77
R100	465 for $p = 1, 2, \dots, 25$.	118	\$8.70	\$8.66

Notes: The unit of observation in this table is one-subject-one-25-period sequence. The earnings from adopting the unconditionally optimal consumption path in a sequence was \$9.78 for all four treatments. Column (3) reports the average of final earnings from one sequence of the four different treatments. Subjects' total earnings from the experiment were the sum of their final earnings from one randomly chosen sequence in each of the *two* treatments they participated in plus a \$7 show-up payment, so average actual earnings were around \$24. Column (4) reports the amount that subjects would earn in a single sequence if they adopted the strategy of consuming endowments in all 25 periods.

5 Findings

We summarize our experimental results as a number of different findings. We begin with the finding that our experimental data run counter to Hypothesis 1:

Finding 1 *The quality of subjects' decisions improves with experience.*

We use three measures to assess and compare the quality of subject's decisions in order to evaluate Hypothesis 1. The first measure is subjects' final earnings, corresponding to lifetime utility in the model. The second measure is the Root-Mean-Square-Deviation (RMSD) of consumption choices from the unconditionally optimal level, which reflects the deviation of the main choice variable from the level that maximizes lifetime utility. The last measure is

the RMSD of assets from the unconditionally optimal level, which reflects actual deviations in the endogenous state variable, wealth, from the level that maximizes lifetime utility.⁷

Table 3 reports on the average final earnings and the average RMSD for each sequence, S1-S4, and whether the difference between the paired observation of two sequences follows a symmetric distribution around zero using a non-parametric Wilcoxon signed-rank test. Panel A reports results for final earnings, Panel B reports on the RMSD of consumption and panel C reports on the RMSD of assets. In each panel, Line 1 compares sequence 1 with each of sequences 2, 3, and 4; line 2 compares sequence 2 with that of sequences 3 and 4 and line 3 compares sequence 3 with sequence 4. It is clear that among all sequences, sequence 1 has the smallest final earnings (panel A), and the largest RMSD of consumption (panel B) or assets (panel C), indicating that the quality of individual decisions improves with experience. As shown in lines 2-4 however, the difference in means among the other three sequences is quite small, suggesting that subjects do not learn a lot beginning from the second lifecycle sequence onwards. Further, there is no evidence that decisions are converging to rational choice model predictions beyond the first repetition (S2); even in the final sequence, subjects' decisions deviate substantially from the unconditionally optimal levels, resulting in an average 11.9 percent reduction in lifetime utility compared with the unconditionally optimal path.

While in reality, individuals make lifecycle consumption/saving decisions just once, Finding 1 suggests that policies which approximate or substitute for greater experience, e.g., access to financial literacy programs or nudges to save more, might be effective in obtaining better outcomes.

Finding 2 *The quality of subjects' decisions differs with the different endowment profiles they faced, though these quality differences weaken with experience.*

To evaluate hypothesis 4, Table 4 reports on the average final earnings for each treatment in each sequence and whether the probability of an observation from treatment i exceeding an observation from treatment $j \neq i$ equals the probability of an observation from treatment j exceeding an observation from treatment i using a Wilcoxon rank-sum test.⁸ We focus

⁷The RMSD of the outcome variable $z \in \{c, a\}$ for subject i in sequence s is defined as follows:

$$RMSD_{is}^z = \sqrt{\frac{1}{P} \sum_{p=1}^P (z_{istp} - z_{tp}^{**})^2}.$$

where z_{istp} is the value of z for subject i in sequence s of treatment t for period p . z_{tp}^{**} is the associated unconditionally optimal level of treatment t for period p .

⁸Pairwise treatment differences in the RMSDs of consumption from the unconditionally optimal path show a similar pattern, and are therefore relegated to Table A1 in Appendix B.

Table 3: Comparison across sequences

	Mean	S2	S3	S4
	(1)	(2)	(3)	(4)
Panel A : Final earnings				
S1	8.10	<***	<***	<***
S2	8.65		<	<*
S3	8.61			<
S4	8.71			
Panel B: RMSD of consumption				
S1	865.96	>***	>**	>***
S2	704.48		>	>**
S3	730.69			>
S4	662.82			
Panel C: RMSD of assets				
S1	3787.49	>***	>***	>***
S2	3156.15		>	>
S3	3134.48			>
S4	3070.37			

Note: Column (1) reports the average value for each sequence. Columns (2)-(4) report the sign of the sum of the signed rank and its significance (***>0.01, **>0.05, *>0.1) using the Wilcoxon signed-rank test. < indicates by comparing the row sequence to the column sequence that the sum of the signed rank is negative, while > indicates the opposite.

first on sequence 1 (S1) of each treatment. Panel A shows that in this first sequence, subjects who are randomly assigned to R40 are significantly more likely to have smaller earnings than subjects randomly assigned to R0 (at the 10 percent level), or to R100 (at the 1 percent level). In addition, column 4 in the same Panel A indicates that subjects who are randomly assigned to R100 are significantly more likely to have greater earnings than those who are randomly assigned to any of the other three treatments in S1. Hence, given that subjects' decisions deviate from theoretical predictions, making the endowment flow perfectly constant over time, i.e., imposing a 100 percent pension replacement rate as is done in R100, induces subjects to make better intertemporal allocation decisions and to earn a higher lifetime utility. The lifetime utility in S1, however, is not a monotonic function of the pension replacement rate, since average earnings in R40 are smaller than those in R0 while the pension replacement rate in R40 is greater than that in R0. In summary, we find that deviations from rational choice model predictions *do* differ by treatment so that pension replacement rates really do matter for lifecycle decisions.

With experience, subjects' earnings tend to increase on average (see column 1 of Table 4), and the treatment differences found for S1 dissipate and become statistically insignificant in any pairwise comparisons of treatments by the final sequence, S4 (Panel D). The notable decrease in earnings for LS in S3 is likely because subjects' experience from the first two sequences of R100 (which preceded their play of LS in S3) provided little help to their solving the lifecycle problem faced in LS. However, we wish to stress that we view the differences observed across treatments in the very *first* sequence, S1, as a main finding, since the first sequence best resembles the decision process in the real world where every individual has only one life-time in which to make consumption and saving decisions. Note that our experimental design does allow individuals to adjust their consumption/saving decisions over the lifecycle, since in each sequence subjects are asked to make sequential choices for 25 different periods.

Finding 3 *Subjects depart from both the unconditionally and the conditionally optimal paths for consumption. In particular, they over-consume in early periods and under-consume in later periods of a lifecycle.*

To check on Hypothesis 2, Figure 2 presents the average period-by-period deviation from the *unconditionally* optimal path and the associated 95% confidence interval. To account for possible serial correlation of decision errors over time within subjects, standard errors are clustered at the subject level. This figure clearly indicates that decisions observed in the experiment depart from the unconditionally optimal decisions implied by rational choice

Table 4: Final earnings by treatment in each sequence

	Mean	T2	T3	T4
	(1)	(2)	(3)	(4)
Panel A: S1				
T1	7.79	<*	<	<***
T2	8.27		>	<*
T3	7.86			<**
T4	8.47			
Panel B: S2				
T1	8.45	<	<**	<***
T2	8.59		<	<*
T3	8.70			>
T4	8.84			
Panel C: S3				
T1	8.83	<	>***	<
T2	8.78		>**	<
T3	8.09			<***
T4	8.73			
Panel D: S4				
T1	8.83	>	>	<
T2	8.66		>	<
T3	8.59			<
T4	8.76			

Note: Column (1) reports the average value for each treatment. Columns (2)-(4) report the sign of the difference between treatments and its significance (**>0.01, **>0.05, *>0.1) using the Wilcoxon rank-sum test. < indicates the sum of ranks (beginning at 1 for the smallest value) of the row treatment is smaller than that of the column treatment, and > indicates the opposite.

theory. Subjects consume more than the optimal level in the early periods, and hence, relative to the optimal consumption level, subjects accumulate fewer assets and can afford less consumption in the later periods of their lives.

To further illustrate the magnitude of deviations from the unconditionally optimal path, Figure 3 plots the average *percentage* deviation of consumption and assets relative to the unconditionally optimal path. As shown in Figure 3(a), in the first three periods, consumption in our experiment is around 100 percent greater than that predicted by the rational choice model. Over time, the percentage deviation becomes smaller and negative, and reaches its lowest level at around period 18 (the retirement period).

As shown in Figure 3(b), actual asset positions are more than 40 percent lower than that predicted by the rational choice model from periods 7 through 22. The asset deviation becomes smaller toward the final periods, and the asset level is close to the unconditionally optimal level by the end of each 25-period sequence. Despite the significant difference in the average RMSD across treatments, we cannot identify any significant difference across treatments using a period-by-period comparison. As shown in Figure A3 in Appendix B, there is a very tiny difference in average consumption across treatments for each individual period in each sequence.

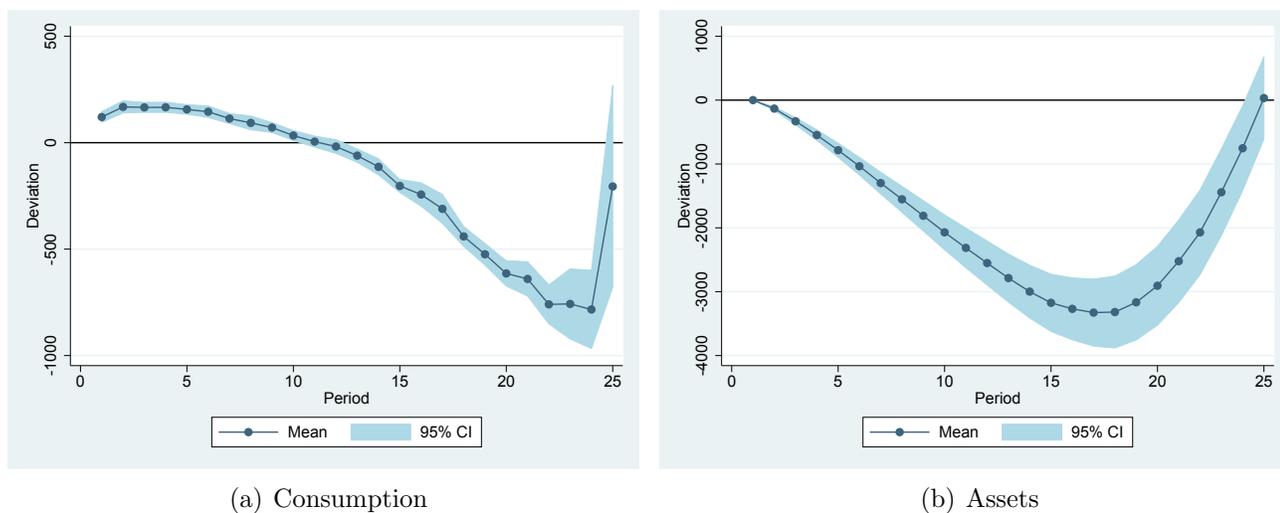


Figure 2: Average deviation from the unconditionally optimal path

Note: The sample includes observations from all treatments and all sequences. Robust standard errors are clustered at the subject level.

One concern with studying deviations from the *unconditionally* optimal path is that errors made in one period not only affect that period’s deviation but also have a permanent effect on

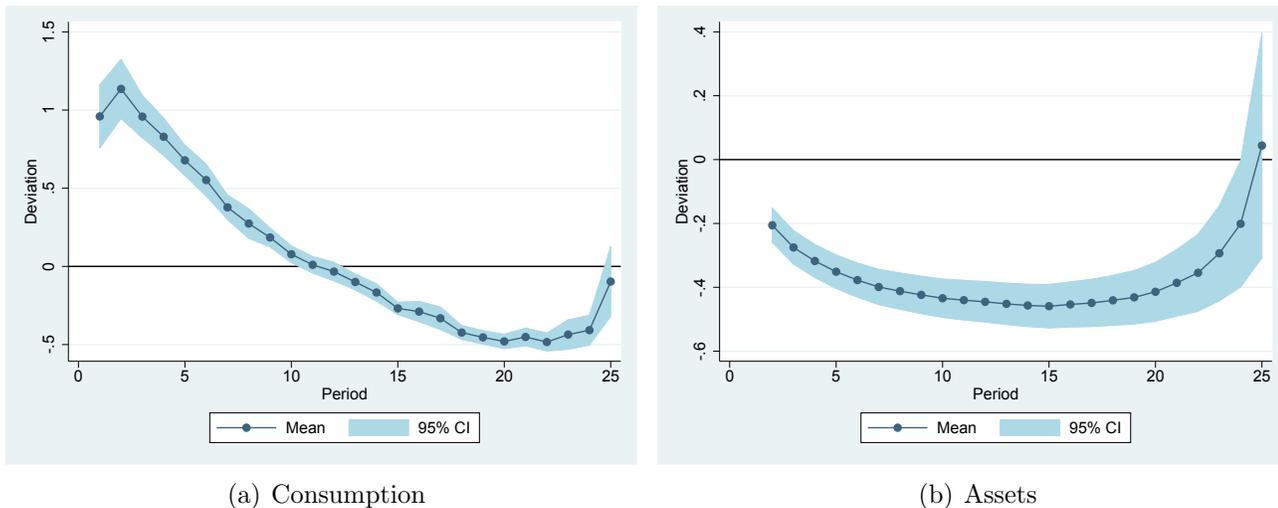


Figure 3: Average percentage deviation from the unconditionally optimal path

Note: The sample includes observations from all treatments and all sequences. Robust standard errors are clustered at the subject level. Figure 3(b) omits the first periods value, since the initial asset is 0.

all future deviations, even if subjects behave optimally from that period onwards. To address this concern and to test Hypothesis 3, Figure 4(a) plots the average consumption deviation from the *conditionally* optimal path, which eliminates the effect of past decision errors on current deviations, along with the associated 95 percent confidence interval. Notice that since the deviation in consumption exactly mirrors the deviation in assets when evaluated relative to the conditionally optimal level, there is no need to present a separate figure for the asset profile. It transpires that the pattern for consumption deviations found in Figure 2(a) is preserved in Figure 4(a): relative to the conditionally optimal path, subjects continue to over-consume in early periods and under-consume in later periods of the lifecycle. The only difference is that compared with Figure 2(a), in Figure 4(a), the deviation from the optimal path remains positive for an even *longer* period of time, and only around period 18 do subjects begin to consume less than the conditionally optimal level.⁹

Figure 4(a) reveals what appears to be substantial amounts of under-consumption even in the last period of a sequence, where it should be clear that consuming all remaining wealth is a dominant strategy. However, the number of such events is few. Out of the 476 subject-

⁹The percentage deviation from the conditionally optimal consumption path is relegated to Appendix Figure A4. This figure shows that the under-consumption pattern in later periods is not preserved in percentage terms, likely because as discussed later, less than 50% of subjects on average under consume in the last 8 periods, although the absolute level of their deviation from the conditionally optimal path is larger than those who over consume in the last 8 periods.

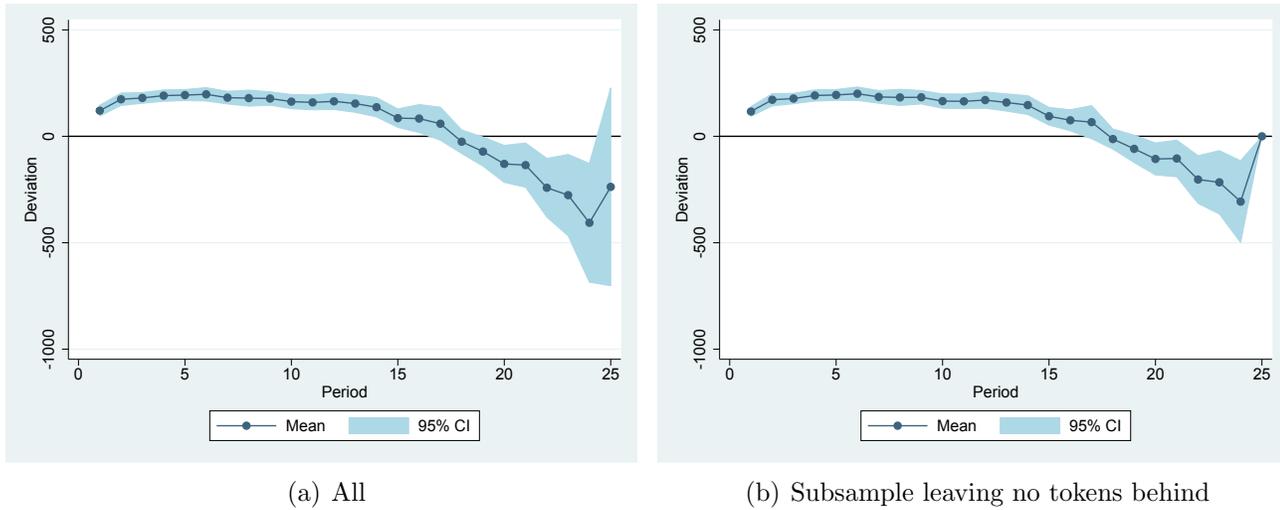


Figure 4: Average consumption deviation from the conditionally optimal path

Note: Robust standard errors are clustered at the subject level.

sequence pairs, there are just 11 instances (by six different subjects) where more than 1 token was unconsumed in the final period. To understand the behavior of the large subsample of subjects who understood the experimental incentives, in the analysis that follows, we restrict the sample to those subjects who leave no more than 1 token unconverted to cash earnings in the final period. A tolerance of 1 is chosen to accommodate the situation where subjects may round their consumption decisions to the largest integer that is affordable. Figure 4(b) displays the same statistics of Figure 4(a) after deleting these 11 subject-sequence pairs, and we observe that the pattern of over-and-under-consumption is robust to this change.

The switch from over- to under-consumption at around period 18 is coincident with the drop in endowment income that subjects received in treatments R40 and R0, the only two treatments with a reduction in endowments starting in period 18. To understand whether the endowment change in these two treatments is driving the switch from over- to under-consumption, Figure 5 plots the deviation of consumption from the conditionally optimal level for each of the four treatments *individually*. We find that for all four treatments, including treatments LS and R100 that have no change in endowment amounts at period 18, we continue to observe the same pattern of over-and-under consumption, and this pattern does not dissipate with experience (Figure 6). In sum, we identify a robust lifecycle pattern wherein subjects, on average, over-consume in early periods and under-consume in later periods of the lifecycle relative to the *conditionally* optimal path, which takes account of their asset position.

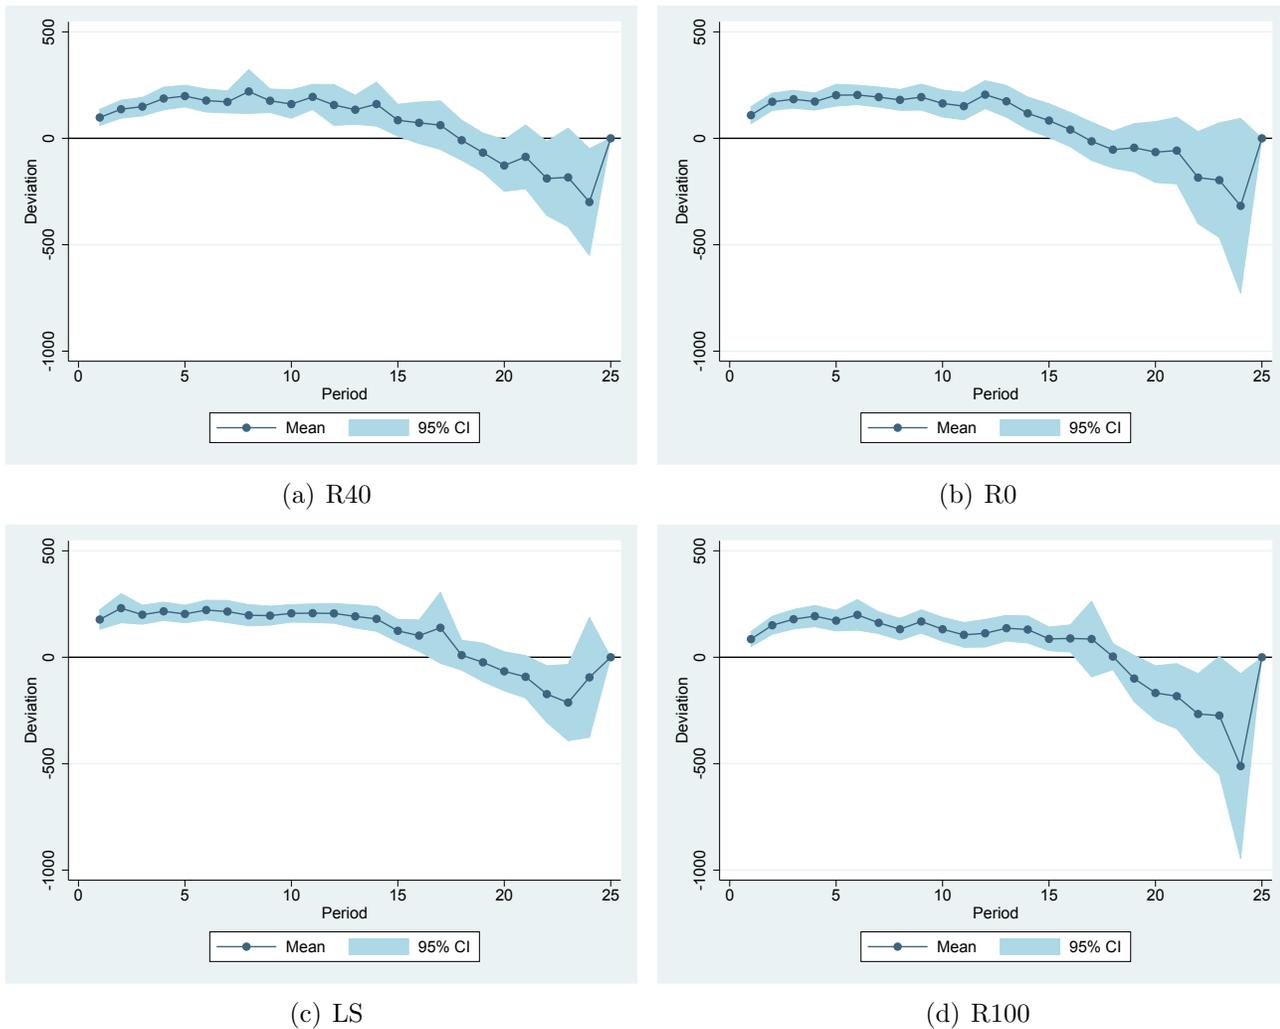


Figure 5: Average consumption deviation from the conditionally optimal path by treatment

Note: For each figure, the sample includes observations from all sequences. Robust standard errors are clustered at the subject level.

To further explore the under-consumption phenomenon, we divide subjects up according to whether they under-consumed, on average, relative to the conditionally optimal amount in the last 8 periods (the retirement period). We find that 41 percent of our subjects can be classified as “under-consumers” in these last 8 periods while the remaining 59 percent consumed more than the conditionally optimal amount in the last 8 periods. As revealed in Table 5, both types, labeled “Under-consumers” and “Others”, on average, over-consume in the first 17 periods, relative to their conditionally optimal consumption levels (compare columns 1 and 3), but the extent of over-consumption in the first 17 periods is less severe

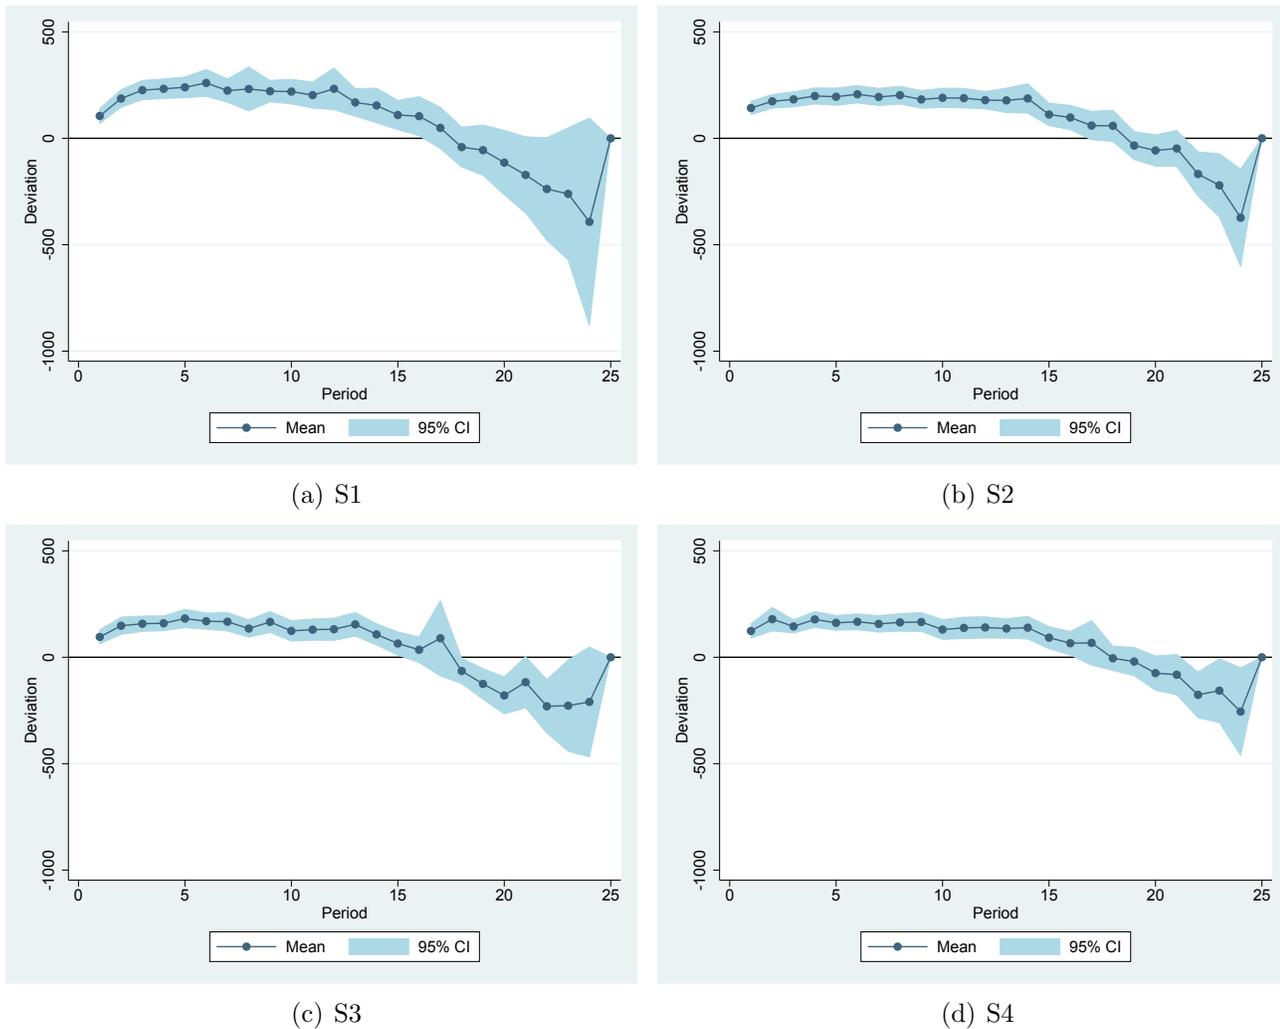


Figure 6: Average consumption deviation from the conditionally optimal path by sequence

Note: For each figure, the sample includes observations from all treatments. Robust standard errors are clustered at the subject level.

among those subjects classified as under-consumers (based on the last 8 periods) as compared with the other subjects who continued to over-consume in the last 8 periods. Those who under-consume in the last 8 periods bring more assets into the last 8 periods, and so they have a greater conditionally optimal consumption in the later periods. In the last 8 periods, these “under-consumers” consume less than the conditionally optimal level (compare columns 2 and 4) but they still consume more than the others, who bring fewer assets into the last 8 periods. Thus, it is possible that the over-and-under consumption pattern relative to conditional optimal levels discussed before may be a result of the experimental design which

Table 5: Comparison of average consumption between subjects who under-consumed in the last 8 periods versus the remaining subjects

	Avg. actual		Avg. con. opt.		Avg. uncon. opt.	
	(1)	(2)	(3)	(4)	(5)	(6)
Last 8 pd class.	P1-17	P18-25	P1-17	P18-25	P1-17	P18-25
Under-consumers	417.81	1367.26	354.09	1826.86	438.29	1530.76
Others	484.09	669.61	265.03	565.21	438.29	1530.76
Diff.	-66.28	697.65	89.06	1261.66		
P-value	0.00	0.00	0.00	0.00		

Note: Columns (1)-(2) report the average consumption of periods 1-17 (column 1) and of periods 18-25 (column 2) for those who, on average consume less than the conditional optimum in the last 8 periods (“Under-consume”) and for the remaining subjects (“Others”). Columns (3)-(4) report the corresponding statistics of columns (1)-(2) for the average, conditionally optimal consumption amounts. For benchmark purposes, the average unconditionally optimal amounts are given in Columns (5)-(6)

yields a much greater optimal consumption level for later periods relative to early periods.

Finding 4 *The high payoff from consuming endowments in each period in treatment 4 (R100) induces some subjects to follow the status quo strategy of consuming endowments.*

To better understand the differences across treatments, the four panels of Figure 7 show the distribution of (final) *earnings* across all subjects participating in each of our four treatments R40, R0, LS and R100, revealing a wide range of earnings (see Appendix Figure A5 for the CDFs of earnings). For reference purposes, each panel also shows, using vertical dashed lines with each of the four treatment labels, the final earnings that subjects would have earned had they followed the simple heuristic of consuming endowments in each of the 25 periods, which we henceforth denote by u_0 . A comparison across treatments reveals that there is no mass in the distribution of earnings around the values of u_0 for the R0 and LS treatments where consuming endowments would be a very poor strategy. By contrast, in R100, and to a smaller extent in R40, there is suggestive evidence for excess bunching around u_0 . As shown in Appendix Figure A6-A9, this excess bunching around u_0 in treatment R100 is persistent for all sequences.

Figure 8 plots the distribution of average *consumption* in relation to average endowment (the latter displayed with vertical lines). Panel A shows the distribution of average consumption over periods 1 to 17 and panel B does the same for periods 18 to 25. We report average consumption over several periods to reduce the impact of random decision errors. Figure 8 clearly reveals that, regardless of whether we consider periods 1–17 or 18–25, there is excess bunching of average consumption around the endowment level in R100, where u_0

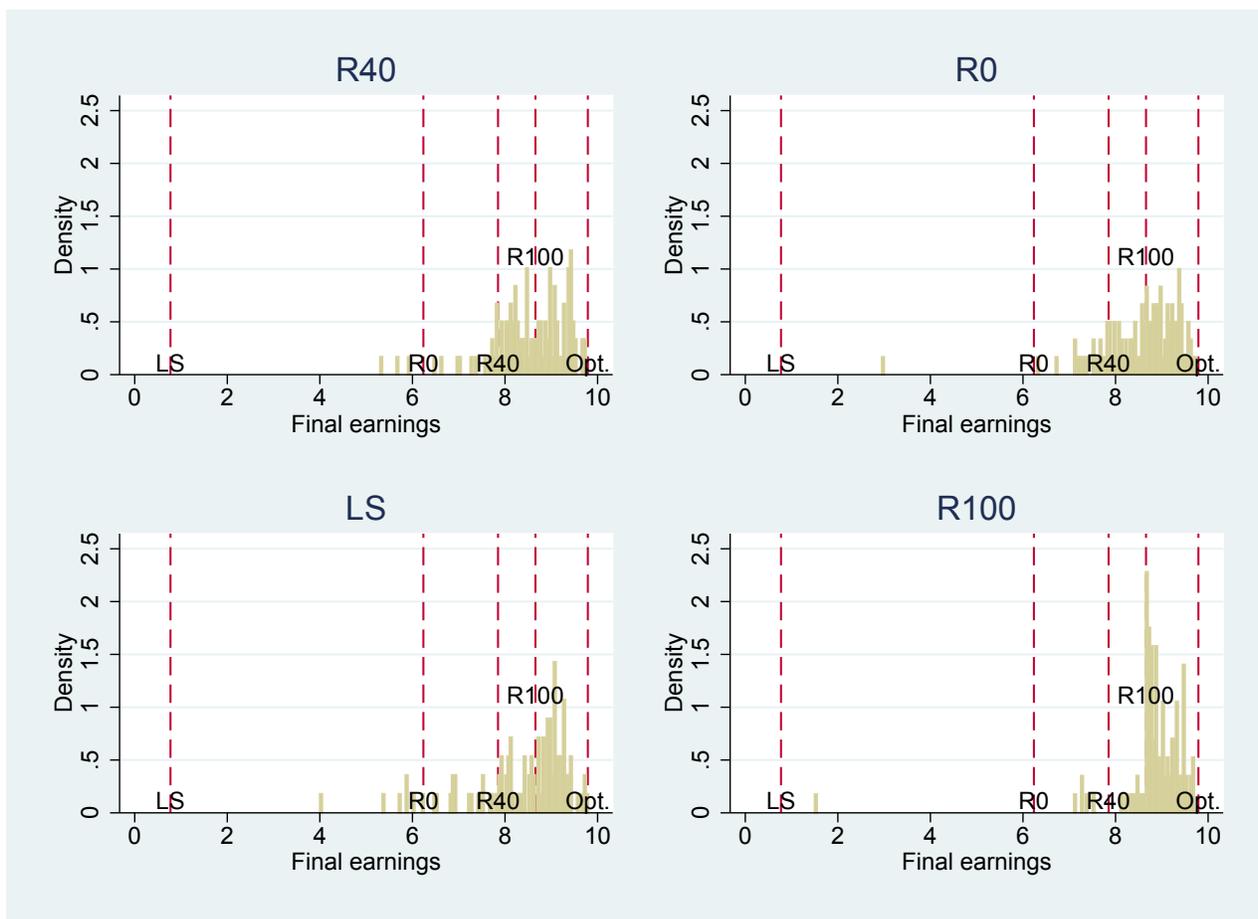


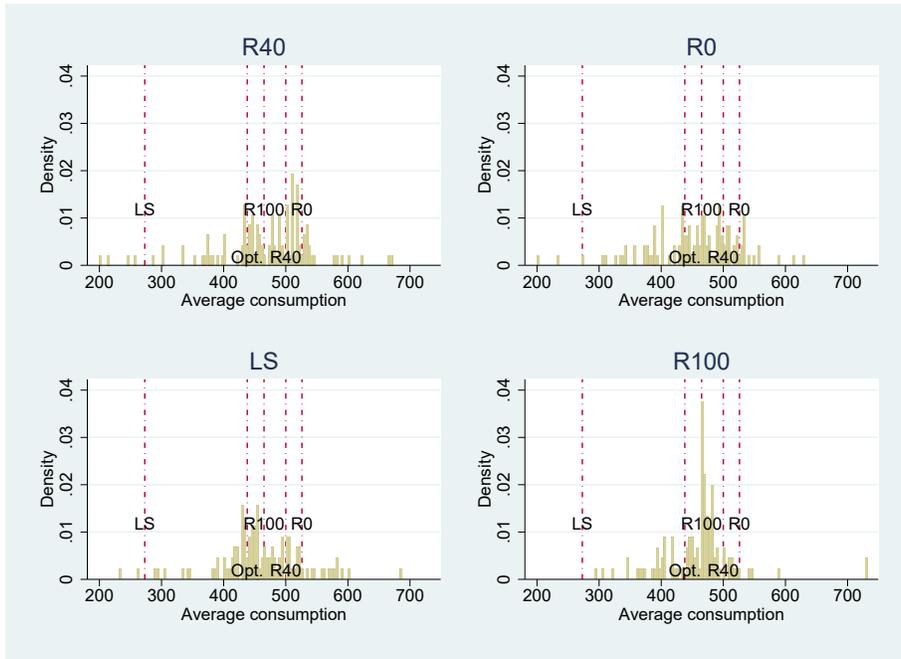
Figure 7: Distribution of final earnings across treatments

Note: The vertical lines indicate final earnings from a strategy of consuming endowments in each of the 25 periods of life in different treatments and from the optimal strategy (opt.).

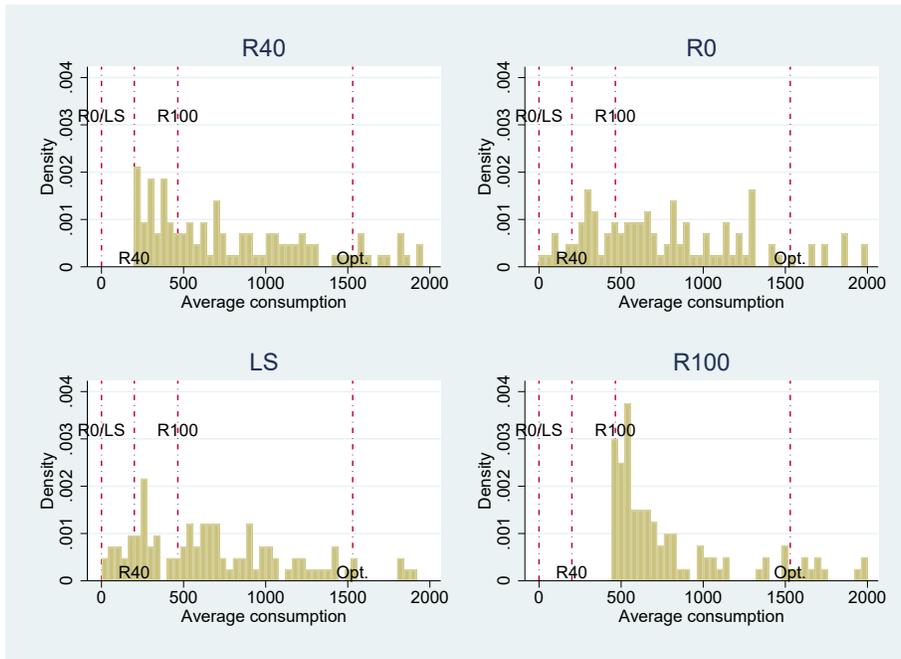
is largest, with considerably less bunching around the endowment level in R40, where u_0 is smaller, and no bunching around the endowment levels in treatments R0 and LS where u_0 is smaller and smallest.¹⁰

The presence of excess bunching in final earnings or in average consumption amounts for treatment R100 relative to treatments R0 and LS is formally tested by estimating the following equation using a linear OLS model. The sample excludes subjects assigned to R40,

¹⁰We tested whether the distribution of the three key outcome variables (final earnings, average consumption for periods 1-17, and average consumption for periods 18-25) are the same across treatments using a Kolmogorov-Smirnov test. As reported in Appendix Table A2, for all three variables, we cannot reject the null hypothesis that the distributions are the same in all pairwise comparisons between treatments R40, R0, and LS, but that there are significant differences in these distributions between R100 and the other three treatments.



(a) Periods 1-17



(b) Periods 18-25

Figure 8: Distribution of Average consumption across treatments

Note: The vertical lines indicate average endowments across periods for different treatments (R40-R100) and the average unconditionally optimal consumption across periods (opt.).

since as noted earlier, u_0 in R40 is moderately high and may affect some subjects' decisions.¹¹

$$I(y^* - \delta \leq y \leq y^* + \delta)_i = Cons. + \beta R100_i + \epsilon_i \quad (1)$$

where $I(y^* - \delta \leq y \leq y^* + \delta)_i$ is an indicator function that equals 1 if the distance between the outcome variable, y , (e.g., final earnings, average consumption for periods 1-17, and average consumption for periods 18-25) for subject i is no greater than some band width, δ , from some target value, y^* (e.g., R100's u_0 of 8.66 and R100's constant endowment of 465); $R100_i$ is an indicator function that equals 1 if subject i is assigned to treatment R100 and 0 otherwise; finally, ϵ_i is an error term. Standard errors are calculated using the bootstrap method with 1000 repetitions. There are three variables of interest: $Cons.$, measures the share of subjects around y^* in treatments R0 and LS; β measures the additional bunching around y^* in R100, and finally, the ratio $\beta/Cons.$ measures the size of excess bunching. Note that estimation of Equation (1) is similar in spirit to the approach used by Saez (2010) (among others) to estimate bunching around kink points. The difference here is that given the small number of observations in our experiment, we chose not to extrapolate the share of subjects in surrounding bands to construct the counterfactual share of subjects around the kink in the absence of the kink. Instead, we use the distribution of the benchmark treatments (R0 and LS) to approximate the counterfactual distribution.

When choosing the band width δ , we face a trade-off: if δ is too small, the estimation may be imprecise due to our small sample size, while if δ is too large, the amount of excess bunching will be attenuated by including subjects who do not react. We lack a good prior for the choice of δ and so we present results for six different δ values, ranging from approximately 5% to 30% of the standard deviation of each outcome variable. As reported in Table 6, for all three outcome variables and almost all δ s, there is statistically significant bunching around the status quo strategy of consuming endowments in each period of R100.

6 Behavioral Models

The previous section presents evidence that the different income profiles we consider with differing pension replacement rates matter for the quality of life-cycle consumption/saving

¹¹Using the same method, we also test for excess bunching for treatment R40 relative to treatments R0 and LS. As reported in Appendix Table A3, there is some suggestive but mostly insignificant evidence for excess bunching in treatment R40 relative to treatments R0 and LS.

Table 6: Excess bunching in R100 compared to R0 and LS

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Final Earnings						
δ	0.05	0.10	0.15	0.20	0.25	0.30
R100	0.09*** (0.03)	0.13*** (0.04)	0.16*** (0.05)	0.18*** (0.05)	0.21*** (0.06)	0.22*** (0.06)
Cons.	0.04*** (0.01)	0.10*** (0.02)	0.14*** (0.02)	0.18*** (0.03)	0.23*** (0.03)	0.28*** (0.03)
Ratio	2.05 (1.54)	1.30* (0.70)	1.16** (0.51)	0.99** (0.39)	0.92*** (0.32)	0.78*** (0.27)
Panel B: Average consumption, periods 1-17						
δ	4.00	8.00	12.00	16.00	20.00	24.00
R100	0.15*** (0.04)	0.16*** (0.05)	0.22*** (0.05)	0.24*** (0.05)	0.28*** (0.05)	0.26*** (0.05)
Cons.	0.06*** (0.01)	0.11*** (0.02)	0.16*** (0.02)	0.21*** (0.03)	0.26*** (0.03)	0.32*** (0.03)
Ratio	2.60* (1.56)	1.52** (0.69)	1.43*** (0.50)	1.16*** (0.38)	1.07*** (0.30)	0.82*** (0.23)
Panel C: Average consumption, periods 18-25						
δ	40.00	80.00	120.00	160.00	200.00	240.00
R100	0.12*** (0.04)	0.19*** (0.05)	0.20*** (0.05)	0.18*** (0.05)	0.12** (0.06)	0.09 (0.06)
Cons.	0.05*** (0.01)	0.09*** (0.02)	0.16*** (0.03)	0.23*** (0.03)	0.32*** (0.03)	0.41*** (0.03)
Ratio	2.39 (1.55)	2.26** (1.03)	1.32*** (0.51)	0.80*** (0.31)	0.38* (0.20)	0.22 (0.15)

Note: Table reports the estimated coefficients for equation (1). In each panel, the ratio is the coefficient for R100 divided by the constant. Bootstrapped standard errors in parentheses (1000 repetitions). *** indicates $p < 0.01$; ** indicates $p < 0.05$; * indicates $p < 0.1$.

decisions. We also found that there is some mass on the heuristic strategy of simply consuming endowments in treatment R100, while there is no mass on this heuristic strategy in treatments LS and R0, where agents receive 0 income in the retirement phase. In this section we consider whether some behavioral models might help us make better sense of these treatment differences (i.e., Findings 2 and 4).

6.1 Rational Inattention Model

One possible explanation for our finding of treatment effects is that subjects face some information processing costs to solving the intertemporal allocation problem, and their incentives to solve this problem depend on the earnings they could get from adopting the simpler strategy of just consuming their endowments in each period, or being a “hand-to-mouth” consumer. To formalize this behavioral explanation, we develop a *rational inattention* model, that is similar in spirit to the rational inattention framework developed by Sims (2003); Matejka and McKay (2014); Carvalho and Silverman (2017) and Gabaix (2017). Specifically, our application to the intertemporal allocation problem is similar to the one presented in Gabaix (2017) section 2.

Suppose that agent i , of ability level x_i , derives utility, w_i , from solving an intertemporal optimization problem of the type subjects faced in our experiment. Let $f(\cdot)$ denote the probability density function for x_i over its support $[0, 1]$, with $x_i = 1$ indicating the highest ability level. As before, let u_0 denote the utility from adopting the benchmark or “status quo” strategy of simply consuming endowments in each period or being a hand-to-mouth consumer. We choose this heuristic as our benchmark as it involves a minimal amount of decision-making; information on period endowments is provided to subjects in each period and consuming all of one’s period endowment maximizes the period utility value (i.e., it is the *statically* optimal choice). We assume that w_i has the following form:

$$w_i = \max\{u_0, \max_u u - g(u, x_i)\}$$

where $g(u, x_i)$ represents the information processing cost, or more precisely, the utility cost of contemplating a consumption plan that yields lifetime utility from consumption of u given ability x_i . We assume $\frac{\partial g}{\partial u} > 0$, meaning that the cost, g , is increasing in lifetime utility from consumption, and $\frac{\partial g}{\partial x_i} < 0$, meaning that the cost, g , is decreasing in agent i ’s ability. Note that agents do not need to pay the information processing cost if they adopt the status quo, hand-to-mouth, consume endowments strategy.

This problem can be solved in two steps. First, if an agent chooses to solve the problem, her choice of u needs to satisfy the following first order condition:

$$1 = \frac{\partial g(u, x_i)}{\partial u}$$

We use $z(x_i)$ to denote the level of lifetime utility from consumption that satisfies the above equation. Second, by comparing the utility from adopting the status quo (u_0), and the utility from solving the problem ($z(x_i) - g(z(x_i), x_i)$), agents choose the option that yields them the greater payoff.

Lemma 1 *Denote $t(u_0)$ as the ability level that satisfies $z(t(u_0)) - g(z(t(u_0)), t(u_0)) = u_0$. Agents with ability $x_i > t(u_0)$ choose to solve the intertemporal maximization problem and have their utility level $z(x_i) > u_0$; agents with ability $x_i \leq t(u_0)$ adopt the status quo strategy of consuming endowments each period and their utility is equal to u_0 .*

Proof. The benefit of solving the intertemporal problem rises with ability:

$$\begin{aligned} \frac{d(z(x_i) - g(z(x_i), x_i))}{dx_i} &= \left(1 - \frac{\partial g(z(x_i), x_i)}{\partial z(x_i)}\right) \frac{dz(x_i)}{dx_i} - \frac{\partial g(z(x_i), x_i)}{\partial x_i} \\ &= -\frac{\partial g(z(x_i), x_i)}{\partial x_i} > 0. \end{aligned}$$

Thus, $x_i > t(u_0)$ implies that $z(x_i) - g(z(x_i), x_i) > z(t(u_0)) - g(z(t(u_0)), t(u_0)) = u_0$; and $x_i \leq t(u_0)$ implies that $z(x_i) - g(z(x_i), x_i) \leq z(t(u_0)) - g(z(t(u_0)), t(u_0)) = u_0$. ■

In our experiment, the share of subjects whose consumption exactly equals their endowment income for the entire 25 periods is 3.8% for treatment R100, 2.5% for treatment R40, and 0% for treatments R0 and LS. The share of subjects whose consumption is within a small band (± 100 tokens) of the endowment level for the entire 25 periods is 8.5% for treatment R100, 2.5% for treatment R40, 0.1% for treatment R0, and 0% for treatment LS.

Proposition 1 *$t(u_0)$ increases with respect to u_0 .*

Proof. The derivative of $z(t(u_0)) - g(z(t(u_0)), t(u_0)) = u_0$ with respect to u_0 is:

$$\begin{aligned} \frac{dz(t(u_0))}{dt(u_0)} \frac{dt(u_0)}{du_0} - \frac{\partial g(z(t(u_0)), t(u_0))}{\partial z(t(u_0))} \frac{dz(t(u_0))}{dt(u_0)} \frac{dt(u_0)}{du_0} - \frac{\partial g(z(t(u_0)), t(u_0))}{\partial t(u_0)} \frac{dt(u_0)}{u_0} &= 1 \\ \left(1 - \frac{\partial g(z(t(u_0)), t(u_0))}{\partial z(t(u_0))}\right) \frac{dz(t(u_0))}{dt(u_0)} \frac{dt(u_0)}{du_0} - \frac{\partial g(z(t(u_0)), t(u_0))}{\partial t(u_0)} \frac{dt(u_0)}{du_0} &= 1 \\ -\frac{\partial g(z(t(u_0)), t(u_0))}{\partial t(u_0)} \frac{dt(u_0)}{u_0} &= 1 \end{aligned}$$

Thus, $\frac{dt(u_0)}{u_0} = \frac{1}{-\frac{\partial g(z(t(u_0)), t(u_0))}{\partial t(u_0)}} > 0$. ■

Proposition 1 explains finding 4: 1) for small u_0 that satisfies $\int_0^{t(u_0)} f(x) = 0$, measure zero of subjects choose the status quo strategy of consuming endowments in each period, similar to our findings for treatments R0 and LS; and 2) for large u_0 that satisfies $\int_0^{t(u_0)} f(x) > 0$, a positive measure of subjects choose the status quo strategy of consuming endowments each period, as we found in our treatment R100. This proposition predicts that the mass of subjects choosing the heuristic of consuming endowments is *non-decreasing* with respect to u_0 .

Proposition 2 *An increase in u_0 causes average utility from consumption to rise if $\int_0^{t(u_0)} f(x)dx > g(z(t(u_0)), t(u_0))f(t(u_0))t'(u_0)$ and to fall if $\int_0^{t(u_0)} f(x)dx < g(z(t(u_0)), t(u_0))f(t(u_0))t'(u_0)$.*

Proof. Let $U = \int_0^{t(u_0)} u_0 f(x)dx + \int_{t(u_0)}^1 z(x_i)f(x)dx$ denote the average utility from consumption. We have

$$\begin{aligned} \frac{\partial U}{\partial u_0} &= \int_0^{t(u_0)} f(x)dx + u_0 f(t(u_0))t'(u_0) - z(t(u_0))f(t(u_0))t'(u_0) \\ &= \underbrace{\int_0^{t(u_0)} f(x)dx}_{\text{gains of low ability agents who already adopted } u_0} - \underbrace{g(z(t(u_0)), t(u_0))f(t(u_0))t'(u_0)}_{\text{losses of high ability agents who switch to } u_0} \lesseqgtr 0 \quad (2) \end{aligned}$$

The second equality follows from the fact that $u_0 - z(t(u_0)) = -g(z(t(u_0)), t(u_0))$. ■

Proposition 2 provides some insight into Finding 2 regarding the non-monotonic relationship between u_0 and average earnings in the first sequence of our treatments. Depending on the value of u_0 , average earnings could be lower (higher) if some high ability subjects who could earn more than u_0 if they solved the optimization problem instead chose u_0 , and the lost earnings of those subjects (the second term in Equation (2)) are higher (lower) than the gain in earnings for some low ability subjects who would have earned less than u_0 if they had attempted to solve the problem but who chose u_0 instead (the first term). Note that it might be inappropriate to apply Proposition 2 to explain the differences across treatments beyond the first sequence, since the distribution of abilities in later sequences might change with the experience acquired in earlier sequences, as evidenced by the increase in earnings in S2-S4 (Finding 1).

Although the impact of increasing u_0 on observed average utility from consumption is ambiguous, after considering the cost of processing information and solving the model, a

better status quo choice (i.e., a higher u_0) always (weakly) improves welfare.¹²

Support for the predictions of the rational inattention model is also be found in data collected from other lifecycle consumption/saving experiments where endowment income followed a stochastic process. We focus on two predictions of the model. First we consider whether subjects ever adopt a status quo, consume endowments heuristic when they face a stochastic income profile. On the one hand, if income is subject to shocks, there should be a reduced tendency for subjects to adopt the hand-to-mouth heuristic of consuming endowments as such a strategy is more costly in the stochastic income case (as compared with the deterministic income case that we consider). On the other hand, stochastic income complicates the problem and raises the contemplation cost for finding the solution. This high contemplation cost may induce more subjects to choose the heuristic of consuming endowments. To address this question, we looked at data collected by Meissner (2014) where subjects faced either an upward or downward trending stochastic income process and (unlike in our experiment) were allowed to borrow. Similar to our experiment, we find that 7% of subjects in Meissner’s experiment have consumption that exactly equals their endowment for the entire 20 periods of a lifecycle. In a follow-up paper, Meissner and Rostam-Afschar (2017) test the Ricardian equivalence proposition in a setting with stochastic income over 25 periods, for which the experimental payoff from consuming endowments in every period is **zero**. In that experiment, there are still two subjects (out of 176) choosing the status quo strategy of consuming endowments.¹³ In summary, we do observe subjects in the lab to choose the status quo strategy of consuming endowments each period even if the associated punishment is high.

A second testable prediction of the rational inattention model is that the fraction of subjects who choose the status quo strategy of consuming endowments each period is *non-decreasing* with respect to the payoff from adopting the status quo strategy. The data from Meissner (2014) shows that in the first sequence of the experiment, the share of subjects who consume endowments in each period is 13.2% in the treatment with a high payment associated with the status quo strategy (i.e., the treatment with an upward trending income process) and is 0.0% in the treatment with a low payment associated with the status quo

¹²In recording subjects’ token conversion decisions, we also collected information on the total amount of time that subjects spent on the decision screen, and their usage of a utility calculator. We find these two variables are not good proxies for contemplation costs, likely because in our design, the experiment did not proceed to the next period until all subjects had submitted their decisions (there was no time pressure) and because there is not much need to use the utility calculator after learning the mapping between consumption and earnings from the instructions.

¹³We would like to thank Thomas Meissner for sharing the data.

strategy (i.e., the treatment with a downward trending income process).¹⁴

In addition to directly analyzing experimental data collected by others, we also use results reported in other papers to make inferences from treatment variations about the relationship between the incentives to adopt hand-to-mouth consumption strategies and the usage of this strategy. For instance, Carbone and Hey (2004) designs a life-cycle experiment with stochastic income to represent income changes associated with employment-unemployment dynamics, and uses that data to estimate the planning horizon of a boundedly rationality agent following the method developed by Ballinger et al. (2003). They report that in the treatments with a big difference between income when employed and income when unemployed, all subjects have a planning horizon that is greater than 1, meaning that they are forward looking. By contrast, in treatments with a small difference between employed and unemployed income, some subjects have a planning horizon of 1, i.e., they behave like hand-to-mouth consumers (see Table 10 of Carbone and Hey (2004)). Carbone (2006) compares the distribution of planning horizons for a typical subject pool (the sample used for Carbone and Hey (2004)) with that of a more representative sample of the population, and finds that the latter sample has a much greater share of subjects who have a planning horizon of 1 (see Figure 5 of Carbone (2006)). A recent paper by Tasneem and Warnick (2018) studies the effect of adding a retirement phase with 0 income to a setting with stochastic income. They report that when there is no need to save for retirement and the punishment for consuming endowments is low, about 12 percent of subjects can be classified as hand-to-mouth consumers. Increasing the punishment associated with consuming endowments by adding a retirement phase (of 0 income) reduces the share of hand-to-mouth consumers to just 2 percent (see Figure 4 of Tasneem and Warnick (2018)).

Moreover, we note that hand-to-mouth consumers have been included in economic models of lifecycle consumption in order to improve their fit to non-experimental field data. The influential work of Campbell and Mankiw (1989) shows that a two-type model with hand-to-mouth consumers explains well the volatility in aggregate quarterly consumption data. This type of model has been generalized by government agencies to develop economic forecasting models. For instance, the hand-to-mouth consumers have been incorporated in the Central Organising Model for Projection Analysis and Scenario Simulation (COMPASS) developed by the Bank of England (Burgess et al., 2013) and the Dynamic Stochastic General Equilibrium (DSGE) model developed by CPB Netherlands Bureau for Economic Policy Analysis (Elbourne et al., 2015).

¹⁴We cannot use the data from Meissner and Rostam-Afschar (2017) to test this prediction, since the payoff from adopting the status quo strategy is always 0 in that experiment.

Summarizing, there is evidence in support of the predictions of our rational inattention model, both in our own experimental data as well as the experimental and non-experimental data of other researchers working on understanding lifecycle consumption decisions.

6.2 Alternative explanations

The phenomenon of over-consumption and under-saving for retirement over the lifecycle is consistent with several alternative theories: 1) hyperbolic discounting, i.e., the tendency for people to choose a smaller-sooner reward over a larger-later reward (Laibson, 1997), 2) home production, the need to expend more on goods consumption when working (Aguiar and Hurst, 2005), and 3) exponential growth bias, the tendency to neglect the power of compound interest (Levy and Tasoff, 2015a). The former two explanations do not apply in the context of this experiment, since subjects are given one single reward at the end of the experiment and there is no home production. It is of interest to test whether the observed treatment difference in our experiment are consistent with the exponential-growth bias theory, though we did not design our experiment with this test in mind.¹⁵

One main prediction of the exponential-growth bias (EGB) theory that can be tested in our experimental setting is that budget-neutral delays in income will increase consumption (Levy and Tasoff, 2015b). This is because subjects with EGB use a distorted interest rate that causes them to both underestimate the price of current consumption and overestimate the present value of future income. To test this implication of EGB, following Levy and Tasoff (2015b), we examine the difference in log-normalized consumption, $\ln(c_{ip}/c_{ip}^*)$, by treatment, separately for sequence 1 (S1) and the other three sequences combined (S2-S4). Concerned that feedback received from early periods might impact on later period decisions, Levy and Tasoff (2015b) focus their analysis on just first period consumption. But subjects' decision errors might be greatest in the very first period. For this reason, we not only consider the first period, but also examine later periods and we normalize consumption by the conditionally optimal level of each period, c_{ip}^* . Since R100 has the greatest delay in lifetime income, followed by R40, R0, and LS, the EGB theory predicts that the normalized consumption should have the following rank order from lowest to highest: LS, R0, R40, and R100. We test this implication by estimating the following equation using a linear OLS

¹⁵To test for exponential growth bias, Levy and Tasoff (2015b) implement an experiment that is similar to ours where subjects live just six periods, and receive endowments only at the start, in the middle, or at the end of each sequence. Their experimental results are in support of the predictions of their exponential growth bias model, which posits that individuals use a hybrid mix of simple and compound interest to evaluate the present value of future income and consumption.

model:

$$\ln(c_{ip}/c_{ip}^*) = Cons. + \alpha_1 R0_{ip} + \alpha_2 R40_{ip} + \alpha_3 R100_{ip} + \epsilon_{ip} \quad (3)$$

where c_{ip} is the consumption for subject i in period p , and c_{ip}^* is the conditionally optimal consumption. $R0_{ip}$, $R40_{ip}$, and $R100_{ip}$, respectively, are indicators for treatments R0, R40, and R100. The EGB predicts that $\alpha_3 > \alpha_2 > \alpha_1 > 0$.

Table 7: Effect of income delays on normalized consumption

	$\ln(c/c^*)$			
	Period 1		Periods 2-24	
	S1	S2-S4	S1	S2-S4
	(1)	(2)	(3)	(4)
R0	-0.527 (0.384)	-0.024 (0.182)	-0.225 (0.193)	-0.066 (0.084)
R40	0.001 (0.212)	-0.020 (0.174)	-0.040 (0.124)	-0.122 (0.084)
R100	-0.549** (0.262)	-0.205 (0.163)	-0.105 (0.117)	-0.251** (0.115)
Cons	0.763*** (0.172)	0.604*** (0.140)	0.472*** (0.103)	0.367*** (0.058)
R ²	0.015	-0.001	0.008	0.013
N	99	312	2563	8016

Note: Table reports the estimated coefficients for equation (3). Robust standard errors are clustered at the subject level (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). The sample is observations of period 1 sequence 1 for column 1, of period 1 sequences 2-4 for column 2, of periods 2-24 sequence 1 for column 3, and of periods 2-24 sequences 2-4 for column 4.

As reported in Table 7, the data reject the EGB theory since subjects assigned to R100 consumed significantly less than subjects assigned to the baseline LS treatment. Further, the dummy variables for treatment R0 is also negative, although it is not statistically significant from 0. We conclude that the EGB theory, while very plausible, cannot account for the treatment differences we observe in our experiment, though we did not design our experiment to test this theory and we cannot exclude the possibility that exponential-growth bias may affect some subjects' decisions.

7 Conclusions

Understanding how individuals make consumption and saving decisions over the lifecycle in the face of different income profiles is important for understanding how individuals are likely to respond to future policy changes in pension replacement rates. In this paper we have gathered experimental evidence on lifecycle consumption and saving decisions over several different income profiles, representing different pension replacement rates. We have three main findings. First, the quality of consumption/saving decisions improves with experience. This finding suggests that policies aimed at improving financial education or nudging agents to save more might be effective proxies for experience, since, in reality, agents get only a single opportunity to tackle the lifecycle consumption/saving problem. Second, subjects deviate from the rational choice theory predictions, by over-consuming in the early periods of life and underconsuming in the later periods of life in all four of our treatments. Third, and most importantly, we find that pension replacement rates matter for subjects' experimental payoffs and consumption behavior. In particular, our treatment with a 100% pension replacement rate yields the highest experimental payoff, and more subjects in this treatment choose the status quo strategy of consuming endowments in each period. In addition to these three main experimental findings, we have shown how a model incorporating ability levels and contemplation costs that is inspired by the rational inattention literature can explain some of the variations that we observe in consumption choices and earnings across our different treatments.

An important issue is whether our experimental findings generalize outside of the laboratory to real, lifecycle decision-making. We think there are good reasons to believe that our laboratory results *would* generalize to the field and may even understate the actual extent of departures from the standard, rational choice model predictions. In particular, in our experiment, subjects' decision costs are incurred close to the time in which they also receive monetary payoffs. In reality, there will typically be a long delay between the time at which saving decisions are made, i.e., the time at which decision costs are realized, and the time at which the benefits from those savings (i.e., additional retirement income) are realized. If individuals were subject to present bias preferences (not really a factor in our design), then the long delay between decisions and rewards will make it even more costly to exert the mental effort needed to make good decisions today, making it even more likely that individuals choose the status quo option of consuming endowments. On the other hand, as we have noted, the real (and higher) costs associated with not saving enough for retirement may be a mitigating factor in extrapolating from our experimental findings to the real world.

While our experimental study relies on the “convenience sample” of university students, this subject population should not be dismissed as lacking in external relevance. Our university subjects are likely better educated than many households tasked with making lifecycle consumption-saving plans (so that again, our laboratory findings may *understate* the extent of suboptimal behavior) and our subjects are also at an age that is close to the first period of the lifecycle model, so that they provide a good approximation to “initial conditions” in terms of experience.

As for policy implications, our findings suggest that changes in pension replacement rates have to be considered with some care. As we have seen, and our rational inattention model formalizes, changes in replacement rates that work to further smooth income over the lifecycle may have the perverse effect of discouraging even high ability individuals from choosing an optimal plan in favor of a more heuristic hand-to-mouth strategy, and can possibly result in a welfare loss in terms of observed lifetime utility. On the other hand, an increase in replacement rates will definitely generate an increase in actual lifetime utility that includes the contemplation cost. In effect these results are reminiscent of the Lucas critique, that behavior is endogenous to policy, so that we cannot evaluate policy changes assuming that current behavior remains unchanged.

We think that an obvious next step in this research agenda is to introduce a contemplation cost to a rational choice model that is calibrated to match the data and use that model to examine the influence of different policy reforms while incorporating uncertainties people face over their life cycles. This extension is in a similar spirit to Kim et al. (2016) who consider the cost of managing portfolios and evaluating the welfare gain from investment management delegation and to Lusardi et al. (2017) who consider the accumulation of financial knowledge and examine the implications on wealth inequality. In addition, one could also introduce survival risks and health spending shocks to the retirement phase and examine the impact of these innovations on consumption and saving decisions over the lifecycle. We leave these extensions to future research.

References

- Aguiar, Mark and Erik Hurst**, “Consumption versus expenditure,” *Journal of Political Economy*, 2005, 113 (5).
- Anderhub, Vital, Werner Gäuth, Wieland Müller, and Martin Strobel**, “An experimental analysis of intertemporal allocation behavior,” *Experimental Economics*, 2000, 3 (2), 137–152.
- Ballinger, T Parker, Eric Hudson, Leonie Karkoviata, and Nathaniel T Wilcox**, “Saving behavior and cognitive abilities,” *Experimental Economics*, 2011, 14 (3), 349–374.

- , **Michael G Palumbo**, and **Nathaniel T Wilcox**, “Precautionary saving and social learning across generations: An experiment,” *The Economic Journal*, 2003, *113* (490), 920–947.
- Brown, Alexander L, Zhikang Eric Chua, and Colin Camerer**, “Learning and visceral temptation in dynamic savings experiments,” *Quarterly Journal of Economics*, 2009, *124* (1), 197–231.
- Brown, Jeffrey R, Arie Kapteyn, and Olivia S Mitchell**, “Framing and claiming: How information-framing affects expected Social Security claiming behavior,” *Journal of Risk and Insurance*, 2016, *83* (1), 139–162.
- , —, **Erzo FP Luttmer**, and **Olivia S Mitchell**, “Cognitive constraints on valuing annuities,” *Journal of the European Economic Association*, 2017, *15* (2), 429–462.
- Burgess, Stephen, Emilio Fernandez-Corugedo, Charlotta Groth, Richard Harrison, Francesca Monti, Konstantinos Theodoridis, Matt Waldron et al.**, “The Bank of England’s forecasting platform: COMPASS, MAPS, EASE and the suite of models,” Bank of England Working Paper 471 2013.
- Campbell, John Y and N Gregory Mankiw**, “Consumption, income and interest rates: Reinterpreting the time series evidence,” in “NBER Macroeconomics Annual 1989, Volume 4,” MIT Press, 1989, pp. 185–246.
- Carbone, Enrica**, “Understanding intertemporal choices,” *Applied Economics*, 2006, *38* (8), 889–898.
- and **John D Hey**, “The effect of unemployment on consumption: an experimental analysis,” *The Economic Journal*, 2004, *114* (497), 660–683.
- and **John Duffy**, “Lifecycle consumption plans, social learning and external habits: Experimental evidence,” *Journal of Economic Behavior & Organization*, 2014, *106*, 413–427.
- Carvalho, Leandro and Dan Silverman**, “Complexity and Sophistication,” Working Paper 2017.
- Chetty, Raj**, “A new method of estimating risk aversion,” *The American Economic Review*, 2006, pp. 1821–1834.
- Elbourne, Adam, Rob Luginbuhl, and Wind Joris de**, “An estimated DSGE model for the Netherlands with unemployment,” CPB Background Document 2015.
- Fehr, Ernst and Peter K Zych**, “Do addicts behave rationally?,” *The Scandinavian Journal of Economics*, 1998, *100* (3), 643–661.
- Feldstein, Martin**, “Structural reform of social security,” *Journal of Economic Perspectives*, 2005, *19* (2), 33–55.
- Feltovich, Nick and Ourega-Zoé Ejebu**, “Do positional goods inhibit saving? Evidence from a life-cycle experiment,” *Journal of Economic Behavior & Organization*, 2014, *107*, 440–454.
- Friedman, Milton**, *A theory of the consumption*, Princeton University Press Princeton, NJ, 1957.
- Gabaix, Xavier**, “Behavioral macroeconomics via sparse dynamic programming,” NBER Working Paper 2017.
- Goda, Gopi Shah, Colleen Flaherty Manchester, and Aaron J Sojourner**, “What will my account really be worth? Experimental evidence on how retirement income projections affect saving,” *Journal of Public Economics*, 2014, *119*, 80–92.

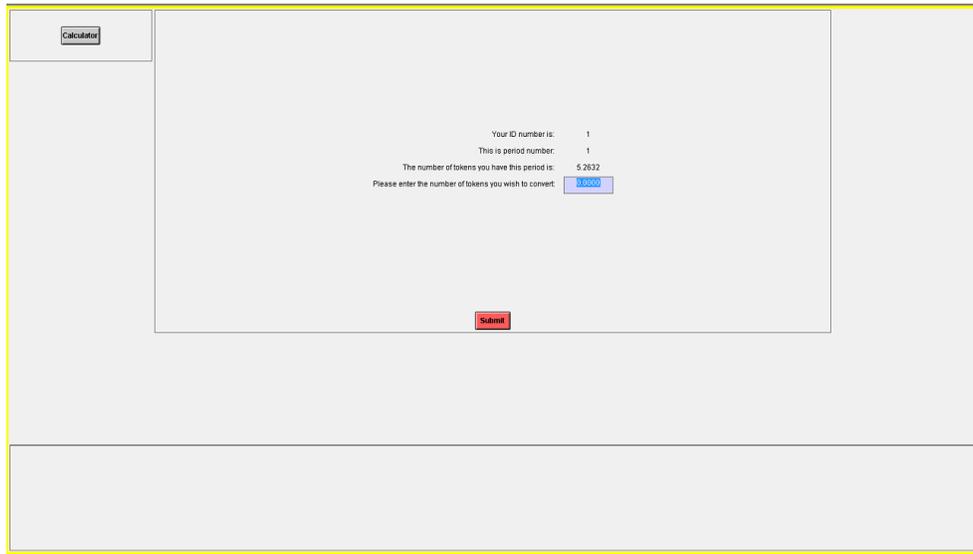
- Hey, John D and Julia A Knoll**, “Strategies in dynamic decision making—An experimental investigation of the rationality of decision behaviour,” *Journal of Economic Psychology*, 2011, *32* (3), 399–409.
- **and Valentino Dardanoni**, “Optimal consumption under uncertainty: An experimental investigation,” *The Economic Journal*, 1988, pp. 105–116.
- Johnson, Stephen, Laurence J Kotlikoff, and William Samuelson**, “Can people compute? An experimental test of the life cycle consumption model,” in “Essays on Saving, Bequests, Altruism, and Life-cycle Planning,” MIT Press, 2001, pp. 335–386.
- Kim, Hugh Hoikwang, Raimond Maurer, and Olivia S Mitchell**, “Time is money: Rational life cycle inertia and the delegation of investment management,” *Journal of Financial Economics*, 2016, *121* (2), 427–447.
- Koehler, Derek J, Jesse Langstaff, and Wu-Qi Liu**, “A simulated financial savings task for studying consumption and retirement decision making,” *Journal of Economic Psychology*, 2015, *46*, 89–97.
- Laibson, David**, “Golden eggs and hyperbolic discounting,” *The Quarterly Journal of Economics*, 1997, pp. 443–477.
- Levy, Matthew and Joshua Tasoff**, “Exponential-Growth bias and lifecycle consumption,” *Journal of the European Economic Association*, 2015.
- Levy, Matthew R and Joshua Tasoff**, “Exponential-Growth bias in experimental consumption decisions,” SSRN Working Paper 2015.
- Liebman, Jeffrey B, Erzo FP Luttmer et al.**, “Would people behave differently if they better understood social security? Evidence from a field experiment.,” *American Economic Journal: Economic Policy*, 2015, *7* (1), 275–99.
- Loewenstein, George and Drazen Prelec**, “Anomalies in intertemporal choice: Evidence and an interpretation,” *The Quarterly Journal of Economics*, 1992, pp. 573–597.
- Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia S Mitchell**, “Optimal financial knowledge and wealth inequality,” *Journal of Political Economy*, 2017, *125* (2), 431–477.
- Matejka, Filip and Alisdair McKay**, “Rational inattention to discrete choices: A new foundation for the multinomial logit model,” *The American Economic Review*, 2014, *105* (1), 272–298.
- Meissner, Thomas**, “Intertemporal consumption and debt aversion: an experimental study,” *Experimental Economics*, 2014, pp. 1–18.
- **and Davud Rostam-Afschar**, “Learning Ricardian Equivalence,” *Journal of Economic Dynamics and Control*, 2017, *82*, 273–288.
- Modigliani, Franco and Richard Brumberg**, “Utility analysis and the consumption function: An interpretation of cross-section data,” *Post-Keynesian Economics*, 1954, *1*.
- Munnell, Alicia H, Jean-Pierre Aubry, Josh Hurwitz et al.**, “How sensitive is public pension funding to investment returns,” CRR Working Paper 2013.
- O’Donoghue, Ted and Matthew Rabin**, “Doing it now or later,” *American Economic Review*, 1999, pp. 103–124.

- Saez, Emmanuel**, “Do taxpayers bunch at kink points?,” *American Economic Journal: Economic Policy*, 2010, 2 (3), 180–212.
- Samuel, J**, “Europe faces pension predicament: Mismatch of lifespans and birthrates means too few workers are paying into state pension plans,” *The Wall Street Journal Online*, 2016.
- Shefrin, Hersh M and Richard H Thaler**, “The behavioral life-cycle hypothesis,” *Economic Inquiry*, 1988, 26 (4), 609–643.
- Sims, Christopher A**, “Implications of rational inattention,” *Journal of monetary Economics*, 2003, 50 (3), 665–690.
- Tasneem, Dina and J Warnick**, “Decision Rules for Precautionary and Retirement Savings,” Technical Report, CIRANO 2018.
- , **Audrey Azerot, Marine de Montaignac, Jim Engle-Warnick et al.**, “A Laboratory Study of Nudge with Retirement Savings,” Working Paper, CIRANO 2018.

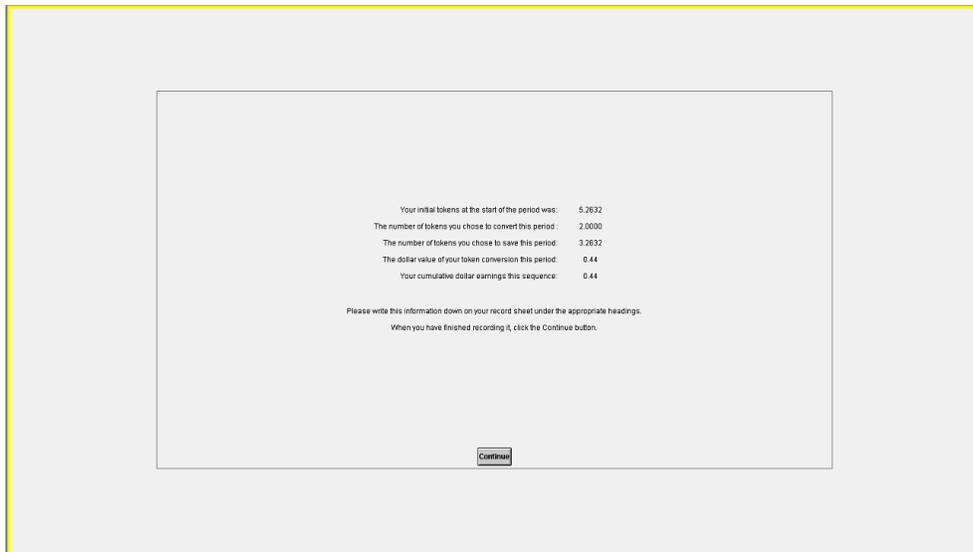
Appendix For Online Publication

A Sample Screenshots and Experimental Instructions

In this section, we include some sample screenshots from our experiment and the instructions used for the treatment order R40-R0. Instructions for the other treatment orders are similar to those reproduced here except that the distribution of endowments is different.



(a) Decision



(b) Results

Figure A1: Screenshots from the experiment

Instructions

Overview

Welcome to this experiment in the economics of decision-making. You are guaranteed \$7 for showing up and completing this experiment. These instructions explain how you can earn additional amounts of money from the decisions that you make in today's session. There is no talking for the duration of this session. Please silence all mobile devices. If you have a question, please raise your hand and your question will be answered in private.

Today's session consists of two parts. These instructions are for the first part. After completing the first part of the experiment, you will receive instructions for the second part. At the end of the second part you will be paid your earnings from both parts together with your \$7 show-up payment in cash and in private.

Part One Instructions

The first part of today's experiment involves two "sequences". Each sequence consists of 25 "periods" of decision-making. At the start of each period you are endowed with a certain number of "tokens." The exact number of tokens given to you in each of the 25 periods of a sequence is shown in Table 1 and is graphed in Figure 1. Please take a moment to look at this sequence of tokens that you will be given. Notice that in some periods, you are given a large number of tokens while in other periods you are given a small number of tokens. The number of tokens you are given each period will also be indicated on your computer screen. In addition to the tokens you are given each period you may have additional tokens that you have saved from prior periods which earn interest in terms of additional tokens as explained below. After viewing the *total* number of tokens you have available --the amount you are given for the period and your tokens from savings and interest-- you must decide how many of these tokens you wish to convert into money for the period. You can convert any number of tokens from 0 up to the maximum total tokens you have available in each period, and you can choose to convert fractions of tokens up to four decimal places. If the 25th period has not yet been reached, then any remaining tokens that you do not convert into money will be saved for your use in later periods, and these savings will earn interest in the form of additional tokens available to you in these later periods as explained in more detail below. In the 25th period, any tokens that you do not convert into money will become worthless.

Your earnings for each period depend on the number of tokens that you convert into money in that period and are shown in Table 2 and graphed in Figure 2. Notice several things. First, only some token amounts that you may wish to convert into money are shown, e.g., 0, 100,...,500,...2000,...10,000. The precise formula used to determine your earnings (in dollars) from converting tokens into money is given at the bottom of Table 2. Second, notice that the more tokens you convert, the greater are your money earnings for that period. Finally, notice also that the money you earn from converting tokens is proportionally *diminishing*; the difference in

your money earnings from converting 500 rather than 400 tokens is larger than the difference in your money earnings from converting 1500 rather than 1400 tokens.

At the start of all 25 periods in a sequence you receive some number of tokens as reported in Table 1 and Figure 1. In addition, in periods 2,...25, you may have *additional* tokens available to you depending on whether you have saved any tokens in prior periods; in that case, you will also receive interest on those savings paid to you in additional tokens. Specifically, you will earn an interest rate of 10 percent per period, paid to you in additional tokens, at the start of the next period. Thus, if in this period you saved $S > 0$ tokens, then at the start of the next period you would have $S + (S \times .10)$ (equivalently $(1.10) \times S$) tokens available to you in addition to the tokens you receive at the start of each new period as given in Table 1. Table 3 shows how various token amounts saved (S) in one period result in additional token amounts of $(1+.10) \times S$ in the following period.

Thus at the start of every period you may have some tokens available to you. Your decision screen will report this total available token number to you, breaking it down according to:

- 1) Token endowment this period: as given in Table 1.
- 2) Tokens saved from the last period: S
- 3) Interest earned on tokens saved from last period savings: $S \times .10$

The total tokens you have available to convert into money or to save in the current period will be the sum of these three numbers.

Your Decision

Type the number of tokens you wish to convert into money (up to four decimal places) in the input box on your decision screen. You may refer to Table 2 and Figure 2 to understand how your token conversion decision determines your earnings, but you can also use the calculator on the top left part of this decision screen to determine how your token conversion decision will translate into money this period. Once you have entered your choice click the Submit button to confirm your choice. You can change your mind anytime prior to clicking the Submit button.

Once the first 25-period sequence is completed, you will begin playing a second 25-period sequence. The second sequence will be just like the first sequence in that you will again receive the same endowment of tokens in each of the 25 periods as indicated in Table 1 and Figure 1 and you will again make token conversion decisions each period as before. Table 2, Figure 2 and Table 3 will continue to apply for determining your money earnings and how saving decisions determine additional tokens.

Information

Following the first period of a sequence, and after every period thereafter, you will be reminded of the total tokens you initially had available at the start of the period, your token conversion decision, your saved tokens, your money earnings for the period as well as your cumulative total money payoff for the current 25-period sequence. Please record this information on your Record Sheets under the appropriate headings. For your convenience, a complete history of this information will be provided at the bottom of your decision screen (following the first period of each sequence).

Earnings

After the second 25-period sequence has been completed, we will randomly select *one* of the two 25-period sequences you played. Both sequences have an equal chance of being chosen. Your cumulative money earnings from the one chosen sequence will comprise your earnings for this first part of today's experiment.

Questions? Now is the time for questions. If you have a question, please raise your hand and the experimenter will answer your question in private.

Quiz

Before continuing on to the experiment, we ask that you complete the following quiz for the part one instructions. In answering these questions, you may consult the instructions, tables and figures. Your performance on this quiz does not affect your payoff in any way. Write or circle your answers to the quiz questions as indicated. Do not put your name on this quiz. If any questions are answered incorrectly, we will go over the relevant part of the instructions again.

1. In part one you will participate in _____ sequences. Each sequence consists of _____ periods.
2. Suppose it is period 1. What is the maximum number of tokens that you can convert into money this period? _____. What is the minimum number of tokens you can convert into money this period? _____.
3. Suppose it is period 10. What is your endowment of tokens in this period? _____. If, in period 9 your savings was 1,000 tokens, how many *total* tokens, including savings, interest earnings and your endowment of tokens for period 10 will you be able to convert into money in period 10? _____.
4. Suppose it is period 20. What is your endowment of tokens in this period? _____. If, in period 19 your savings was 7,000 tokens, how many *total* tokens, including savings, interest earnings and your endowment of tokens for period 20 will you be able to convert into money in period 20? _____.
5. Suppose it is period 25. If you choose to save some of your tokens in period 25, will they have any future value to you? Circle one Yes No.
6. True or false: Your earnings will depend on your cumulative money earnings from one of the two 25-period sequences you play, but you will not know which sequence will be chosen until the end of the session. Circle one: True False

Period	Tokens You are Given
1	500
2	500
3	500
4	500
5	500
6	500
7	500
8	500
9	500
10	500
11	500
12	500
13	500
14	500
15	500
16	500
17	500
18	200
19	200
20	200
21	200
22	200
23	200
24	200
25	200

Tokens Converted	Money Earnings for the Period
0	0.00
100	0.14
200	0.22
300	0.28
400	0.32
500	0.36
600	0.39
700	0.42
800	0.44
900	0.46
1000	0.48
1100	0.50
1200	0.51
1300	0.53
1400	0.54
1500	0.55
1600	0.57
1700	0.58
1800	0.59
1900	0.60
2000	0.61
3000	0.69
4000	0.74
5000	0.79
6000	0.82
7000	0.85
8000	0.88
9000	0.90
10000	0.92

Tokens Saved	Interest Earned in Tokens	Savings+Interest in Tokens
0	0	0
100	10	110
200	20	220
300	30	330
400	40	440
500	50	550
600	60	660
700	70	770
800	80	880
900	90	990
1000	100	1100
1100	110	1210
1200	120	1320
1300	130	1430
1400	140	1540
1500	150	1650
1600	160	1760
1700	170	1870
1800	180	1980
1900	190	2090
2000	200	2200
3000	300	3300
4000	400	4400
5000	500	5500
6000	600	6600
7000	700	7700
8000	800	8800
9000	900	9900
10000	1000	11000

Money=\$ 0.2*ln(0.01*Tokens Converted+1)

Interest (in Tokens)=0.1*Tokens Saved

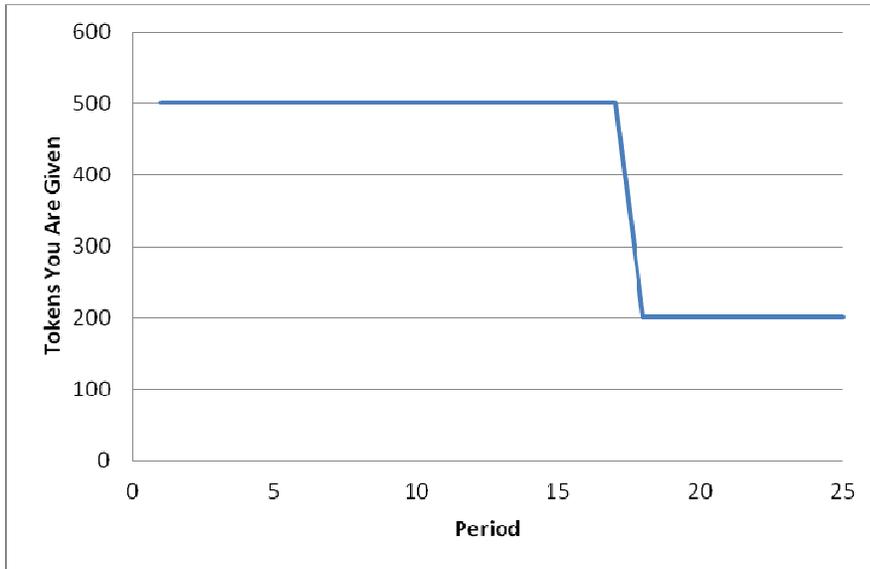


Figure 1: Token You are Given by Period

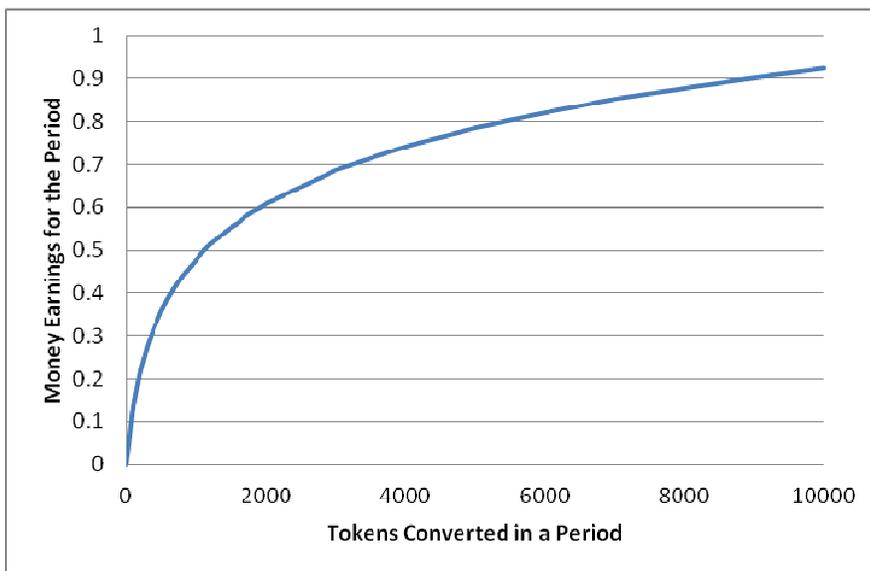


Figure 2: Token Conversions and Money Earned

Record Sheet Part 1 Player ID _____ Age _____ Sex(Circle) F M

Sequence	Period	Initial Total Tokens at the Start of this Period	Number of Tokens You Converted this Period	Number of Tokens You Saved this Period	Money Earned for this Period	Cumulative Money Earnings as of this Period
1	1					
1	2					
1	3					
1	4					
1	5					
1	6					
1	7					
1	8					
1	9					
1	10					
1	11					
1	12					
1	13					
1	14					
1	15					
1	16					
1	17					
1	18					
1	19					
1	20					
1	21					
1	22					
1	23					
1	24					
1	25					

Record Sheet Part 1 Player ID _____ Age _____ Sex(Circle) F M

Sequence	Period	Initial Total Tokens at the Start of this Period	Number of Tokens You Converted this Period	Number of Tokens You Saved this Period	Money Earned for this Period	Cumulative Money Earnings as of this Period
2	1					
2	2					
2	3					
2	4					
2	5					
2	6					
2	7					
2	8					
2	9					
2	10					
2	11					
2	12					
2	13					
2	14					
2	15					
2	16					
2	17					
2	18					
2	19					
2	20					
2	21					
2	22					
2	23					
2	24					
2	25					

Instructions, Continued

Part Two Instructions

The second part of today's experiment is similar to the first part and involves two more 25-period sequences of decision-making. At the start of each period you are again endowed with a certain number of tokens and must again decide how many of your total available tokens you wish to convert into money each period. The only difference from the first part is that in this second part of the experiment, the number of tokens given to you in each of the 25 periods of a sequence is *different* from before and is now shown in Table 4 and graphed in Figure 3. Please take a moment to look at this new sequence of token amounts. Notice that in some periods you are given a large number of tokens while in other periods you are given 0 tokens. The number of tokens you are given in each of the 25 periods will again be indicated on your computer screen. As in the first part, in addition to the tokens you are given each period you may have additional tokens that you have saved from prior periods which earn interest at the same rate of 10 percent as in the first part. After viewing the *total* number of tokens you have available -- the amount you are given for the period and your tokens from any prior period savings and interest-- you must decide how many of these tokens you wish to convert into money for the period. You can convert any number of tokens from 0 up to the maximum total tokens you have available for that period, and you can choose to convert fractions of tokens up to four decimal places. If the 25th period has not yet been reached, then the remaining tokens that you do not convert into money will be saved for your use in later periods, and these savings will earn 10 percent interest per period in the form of additional tokens available to you next period just as in the first part. In the 25th period, any tokens that you do not convert into money will become worthless.

The amounts of money you can earn from converting tokens each period is the *same* as in the first part and thus continues to be given by Table 2 for certain possible token conversion amounts and is graphed in Figure 2. (These are reprinted below). As before, a calculator is available on the top left side of your decision screen to help you determine how your token conversion decisions translate into money earnings each period. As noted above, the interest rate on savings remains the same at 10 percent per period, so that Table 3 (also reprinted below) continues to reveal how various token amounts saved this period earn interest for you in terms of additional tokens next period.

As in the first part you will complete two, 25-period sequences of decision-making. The second sequence will be just like the first sequence in that you will continue to receive the same endowment of tokens in each of the 25 periods as now indicated in the new Table 4 and you will make token conversion decisions each period just as before.

To reiterate, the *only* change from the first part is that the endowments of tokens that you are given in each of the 25 periods of each sequence in this second part of the experiment are *different* and are now given in the new Table 4.

Information

Following the first period of a sequence, and after every period thereafter, you will again be reminded of the total tokens you initially had available at the start of the period, your token conversion decision, your saved tokens, your money payoff for the period as well as your cumulative total money earnings for the current 25-period sequence. Please record this information on your Record Sheets under the appropriate headings. For your convenience, a complete history of this information will be provided at the bottom of your decision screen (following the first period of each sequence).

Earnings

After the second 25-period sequence has been completed, we will randomly select *one* of the two 25-period sequences you played. Both sequences have an equal chance of being chosen. Your cumulative money earnings from the one chosen sequence will comprise your earnings for this second part of today's experiment.

Following the completion of this second part, the experiment will be over. You will be paid your earnings from the first and second parts together with your \$7 show-up payment in cash and in private.

Questions?

Now is the time for questions. If you have a question, please raise your hand and the experimenter will answer your question in private.

Period	Tokens You are Given
1	526
2	526
3	526
4	526
5	526
6	526
7	526
8	526
9	526
10	526
11	526
12	526
13	526
14	526
15	526
16	526
17	526
18	0
19	0
20	0
21	0
22	0
23	0
24	0
25	0

Table 2: Token Conversions and Money Earned	
Tokens Converted	Money Earnings for the Period
0	0.00
100	0.14
200	0.22
300	0.28
400	0.32
500	0.36
600	0.39
700	0.42
800	0.44
900	0.46
1000	0.48
1100	0.50
1200	0.51
1300	0.53
1400	0.54
1500	0.55
1600	0.57
1700	0.58
1800	0.59
1900	0.60
2000	0.61
3000	0.69
4000	0.74
5000	0.79
6000	0.82
7000	0.85
8000	0.88
9000	0.90
10000	0.92

Table 3: Savings and Interest		
Tokens Saved	Interest Earned in Tokens	Savings+Interest in Tokens
0	0	0
100	10	110
200	20	220
300	30	330
400	40	440
500	50	550
600	60	660
700	70	770
800	80	880
900	90	990
1000	100	1100
1100	110	1210
1200	120	1320
1300	130	1430
1400	140	1540
1500	150	1650
1600	160	1760
1700	170	1870
1800	180	1980
1900	190	2090
2000	200	2200
3000	300	3300
4000	400	4400
5000	500	5500
6000	600	6600
7000	700	7700
8000	800	8800
9000	900	9900
10000	1000	11000

Money=\$ 0.2*ln(0.01*Tokens Converted+1)

Interest (in Tokens)=0.1*Tokens Saved

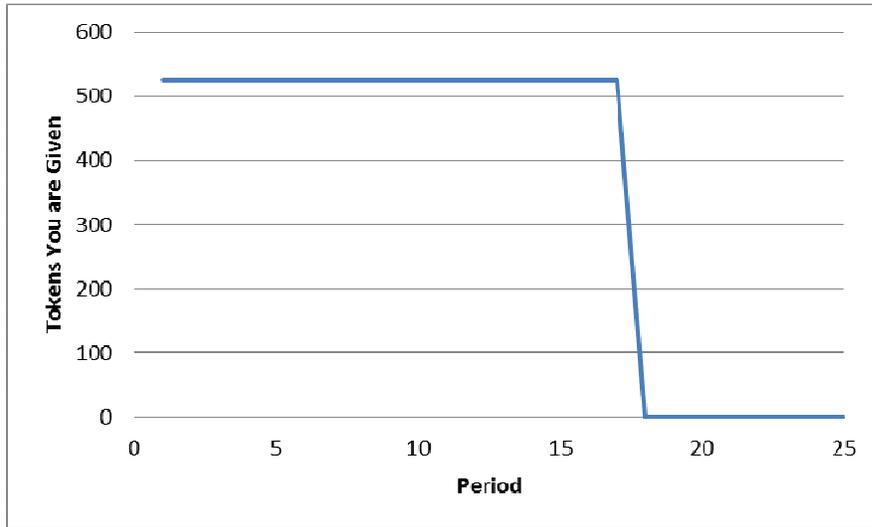


Figure 3: Token You are Given by Period

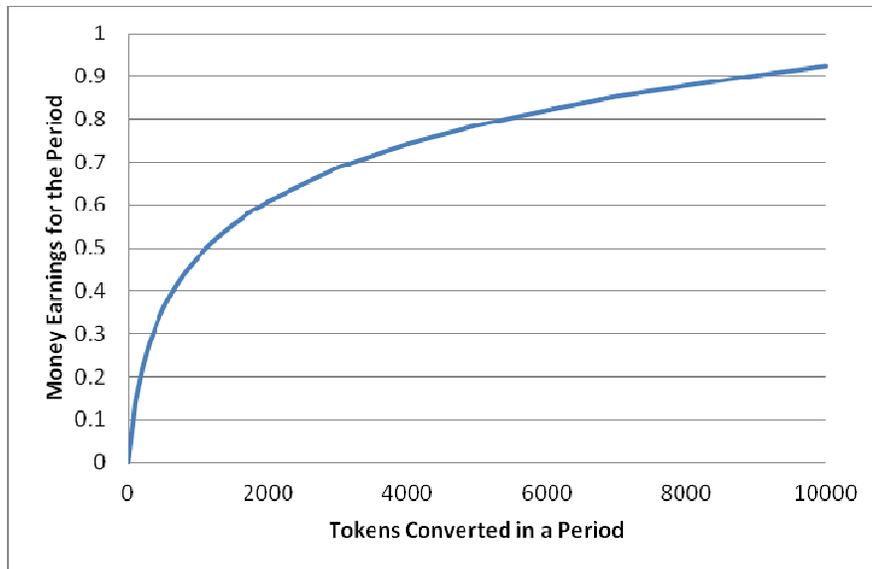


Figure 2: Token Conversions and Money Earned

Record Sheet Part 2 Player ID _____ Age _____ Sex(Circle) F M

Sequence	Period	Initial Total Tokens at the Start of this Period	Number of Tokens You Converted this Period	Number of Tokens You Saved this Period	Money Earned for this Period	Cumulative Money Earnings as of this Period
1	1					
1	2					
1	3					
1	4					
1	5					
1	6					
1	7					
1	8					
1	9					
1	10					
1	11					
1	12					
1	13					
1	14					
1	15					
1	16					
1	17					
1	18					
1	19					
1	20					
1	21					
1	22					
1	23					
1	24					
1	25					

Record Sheet Part 2 Player ID _____ Age _____ Sex(Circle) F M

Sequence	Period	Initial Total Tokens at the Start of this Period	Number of Tokens You Converted this Period	Number of Tokens You Saved this Period	Money Earned for this Period	Cumulative Money Earnings as of this Period
2	1					
2	2					
2	3					
2	4					
2	5					
2	6					
2	7					
2	8					
2	9					
2	10					
2	11					
2	12					
2	13					
2	14					
2	15					
2	16					
2	17					
2	18					
2	19					
2	20					
2	21					
2	22					
2	23					
2	24					
2	25					

B Additional Tables and Figures

Table A1: RMSD of consumption from the unconditional optimal path by treatment for each sequence

	Mean	R0	LS	R100
	(1)	(2)	(3)	(4)
Panel A: S1				
R40	924.37	>*	>	>***
R0	890.90		<	>
LS	779.94			>***
R100	871.03			
Panel B: S2				
R40	748.18	>	>	>**
R0	747.21		>	>
LS	702.37			<
R100	617.32			
Panel C: S3				
R40	677.27	>	<**	<
R0	608.13		<**	<
LS	852.39			>
R100	784.39			
Panel D: S4				
R40	675.03	<	<	>
R0	665.14		>	>
LS	632.12			>
R100	677.32			

Note: Column (1) reports the average value for each treatment. Columns (2)-(4) report the sign of the difference between treatments and its significance (**>0.01, *>0.05, >0.1) using a Wilcoxon rank-sum test. < indicates the sum of ranks (beginning at 1 for the smallest value) of the row treatment is smaller than that of the column treatment, and > indicates the opposite.

Table A2: p -values from Kolmogorov-Smirnov tests of the equality of distributions between pairs of treatments.

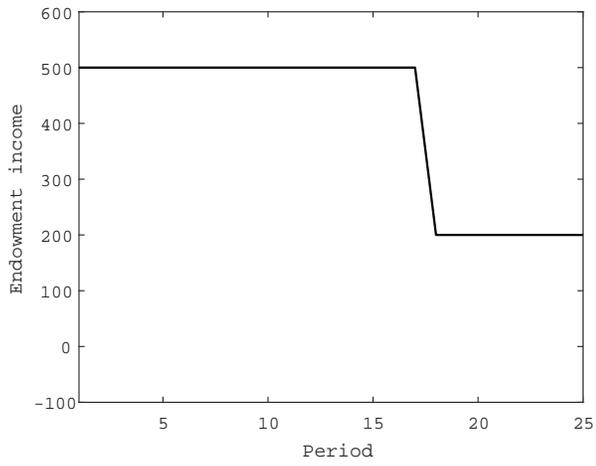
	R40 vs R0	R40 vs LS	R40 vs R100	R0 vs LS	R0 vs R100	LS vs R100
	(1)	(2)	(3)	(4)	(5)	(6)
Final earnings	0.54	0.56	0.00	0.61	0.00	0.00
Ave. con., p1-p17	0.21	0.17	0.00	0.25	0.02	0.07
Ave. con., p18-p25	0.94	0.20	0.00	0.32	0.00	0.00

Note: The table reports p -values from Kolmogorov-Smirnov tests of distributional differences between pairs of treatments for three outcome variables: final earnings, average consumption for periods 1-17, and average consumption for periods 18-25.

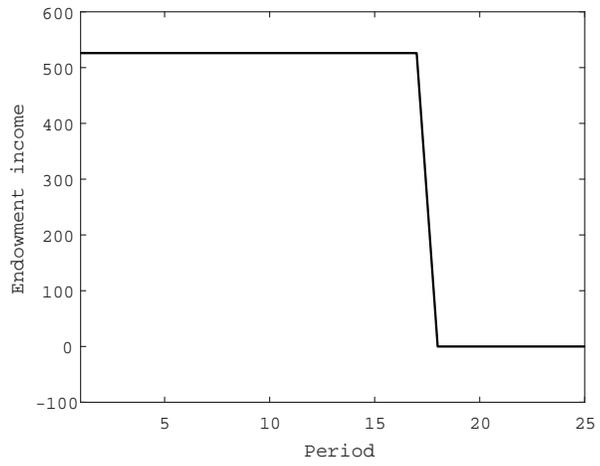
Table A3: Excess bunching in R40 compared to R0 and LS

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Final Earnings						
δ	0.05	0.10	0.15	0.20	0.25	0.30
R40	0.00	0.02	0.02	0.03	0.02	0.03
	(0.02)	(0.03)	(0.03)	(0.04)	(0.04)	(0.04)
Cons.	0.04***	0.06***	0.08***	0.11***	0.14***	0.17***
	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)
Ratio	0.08	0.35	0.23	0.25	0.16	0.20
	(0.87)	(0.69)	(0.49)	(0.41)	(0.34)	(0.29)
Panel B: Average consumption, periods 1-17						
δ	4.00	8.00	12.00	16.00	20.00	24.00
R40	0.02	-0.03	0.04	0.04	0.09*	0.10*
	(0.02)	(0.03)	(0.04)	(0.04)	(0.05)	(0.05)
Cons.	0.04***	0.12***	0.16***	0.18***	0.21***	0.26***
	(0.01)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)
Ratio	0.52	-0.28	0.23	0.24	0.43	0.37
	(1.00)	(0.27)	(0.29)	(0.28)	(0.27)	(0.23)
Panel C: Average consumption, periods 18-25						
δ	40.00	80.00	120.00	160.00	200.00	240.00
R40	0.02	-0.01	-0.01	-0.04	0.01	0.02
	(0.03)	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)
Cons.	0.05***	0.12***	0.19***	0.25***	0.26***	0.28***
	(0.01)	(0.02)	(0.03)	(0.03)	(0.03)	(0.03)
Ratio	0.46	-0.09	-0.07	-0.18	0.02	0.06
	(0.74)	(0.31)	(0.22)	(0.17)	(0.19)	(0.19)

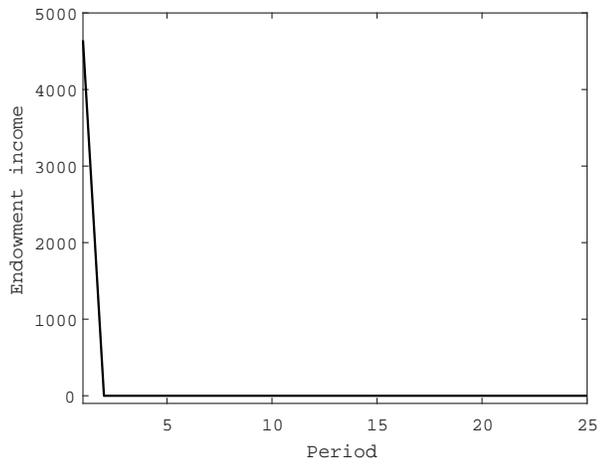
Note: Table reports the estimated coefficients for equation (1), where the indicator variable for treatment R100 is replaced by an indicator variable for treatment R40. In each panel, the ratio is the coefficient for R40 divided by the constant. Bootstrapped standard errors in parentheses (1000 repetitions). *** indicates $p < 0.01$; ** indicates $p < 0.05$; * indicates $p < 0.1$.



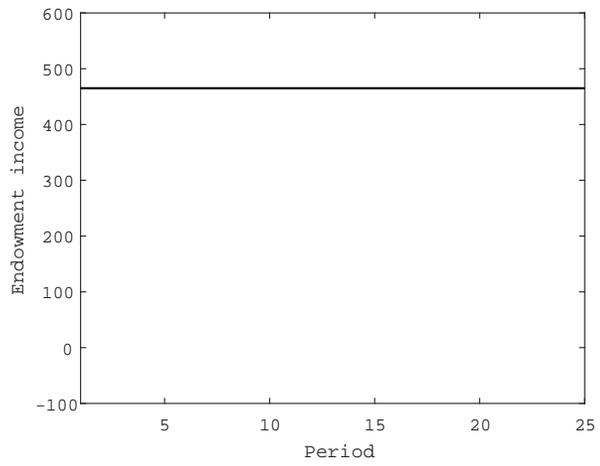
(a) R40



(b) R0

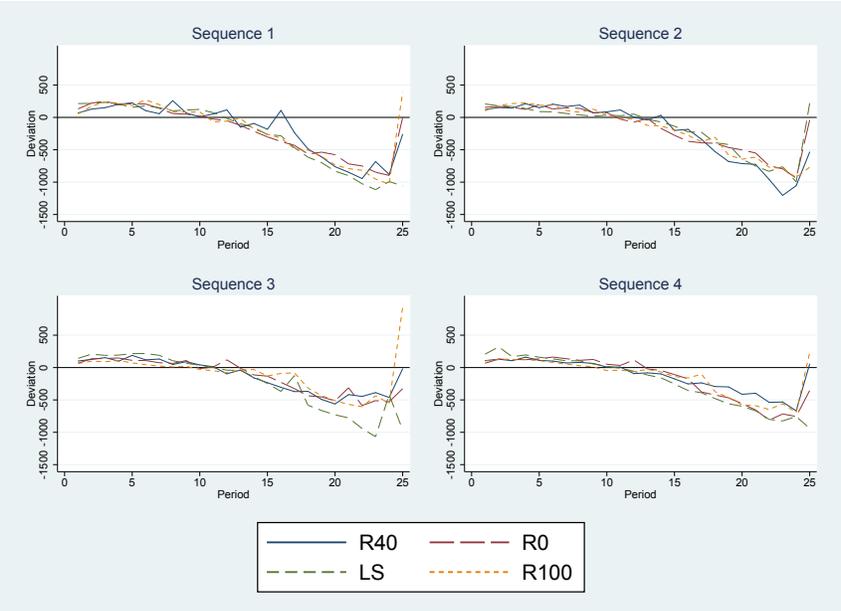


(c) LS

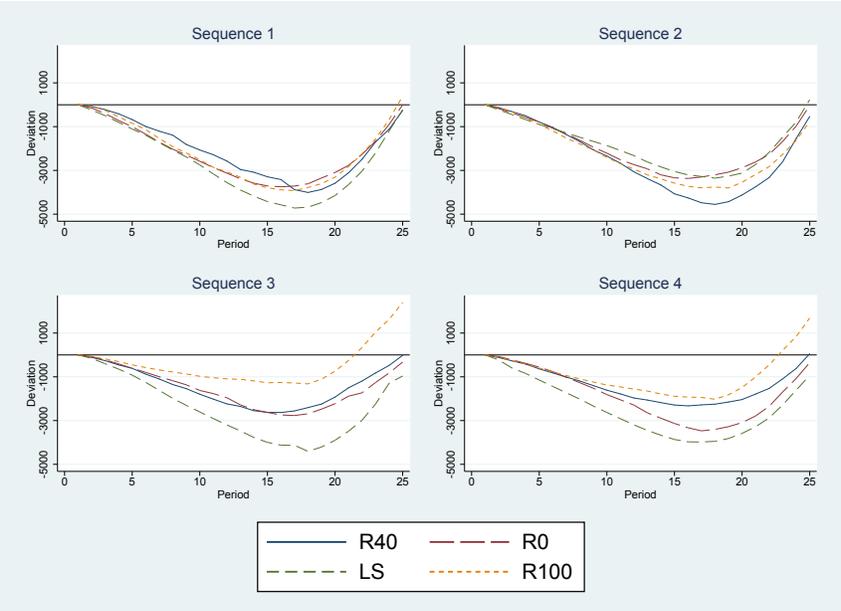


(d) R100

Figure A2: Endowment income by period

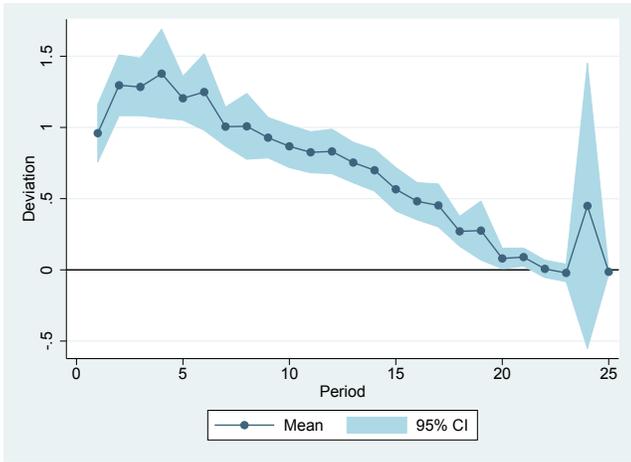


(a) Consumption

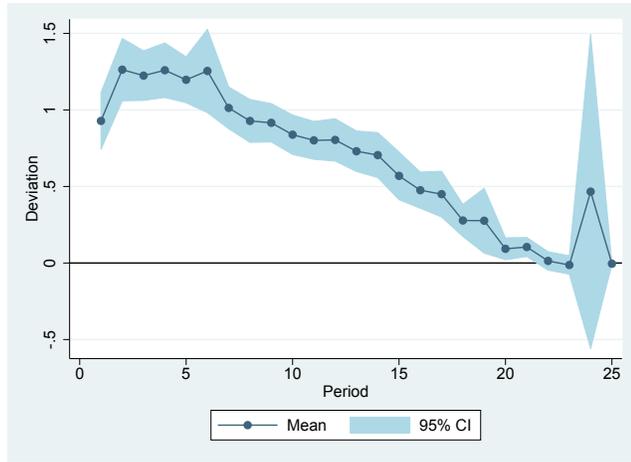


(b) Assets

Figure A3: Average deviation from the unconditionally optimal path by treatment and by sequence



(a) All



(b) Subsample leaving no tokens behind

Figure A4: Average percentage deviation from the conditionally optimal consumption path

Note: Robust standard errors are clustered at the subject level.

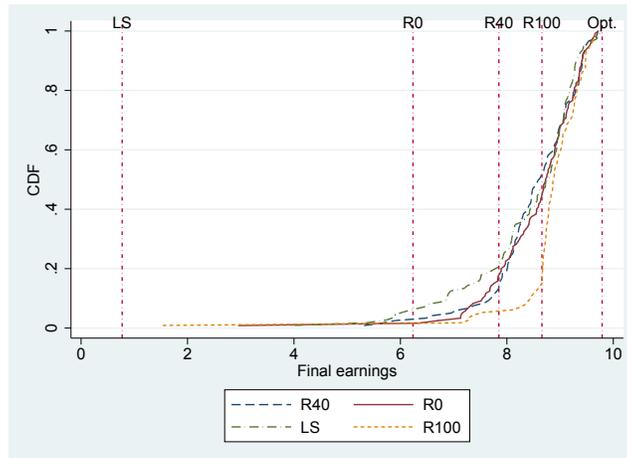


Figure A5: CDF of final earnings by treatment

Note: The vertical lines indicate final earnings from a strategy of consuming endowments in each of the 25 periods of life and from the optimal strategy (opt.).

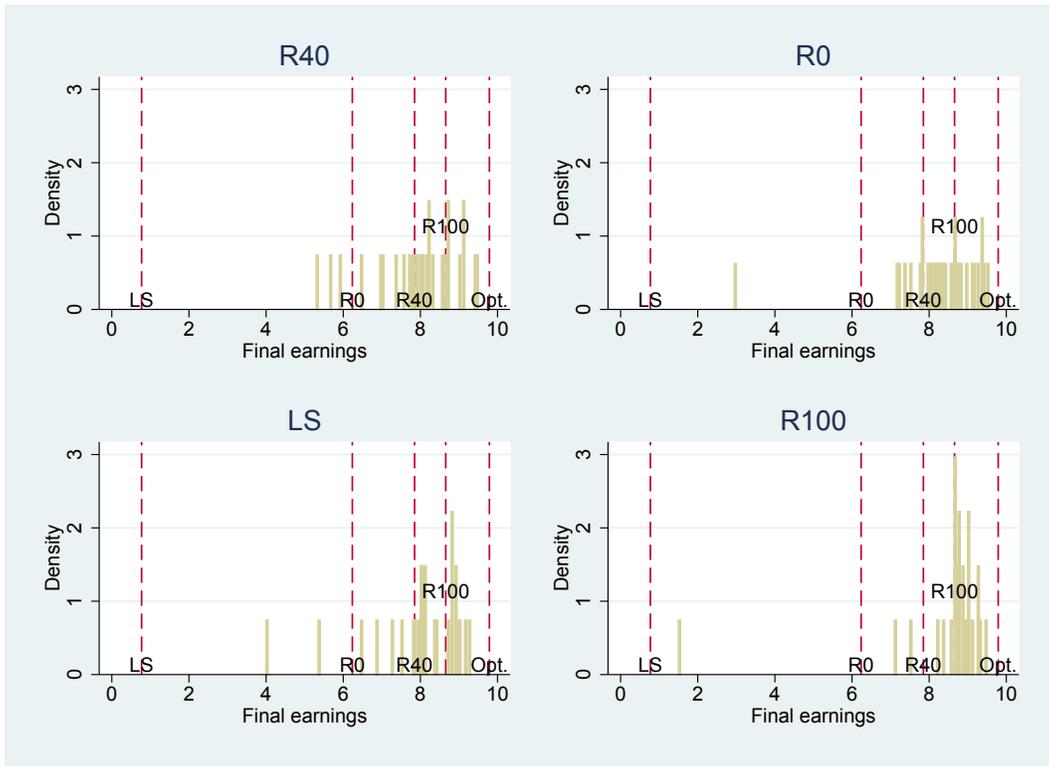


Figure A6: Distribution of final earnings across treatments in Sequence 1

Note: The vertical lines indicate final earnings from a strategy of consuming endowments in each of the 25 periods of life and from the optimal strategy (opt.).

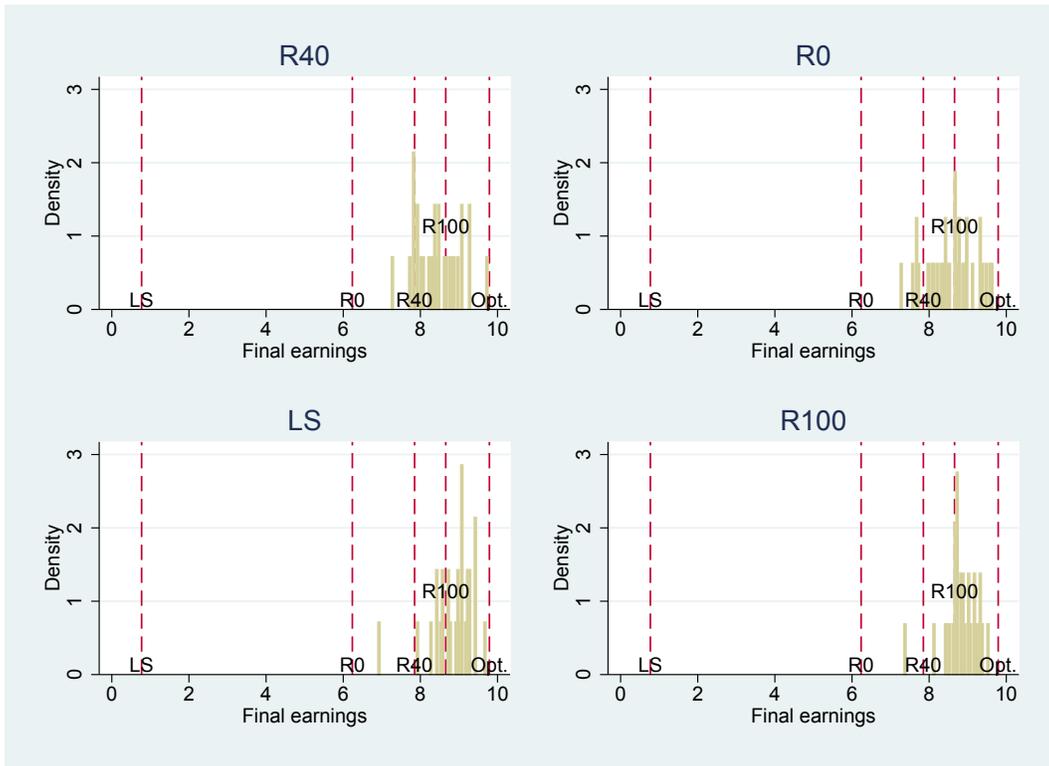


Figure A7: Distribution of final earnings across treatments in Sequence 2

Note: The vertical lines indicate final earnings from a strategy of consuming endowments in each of the 25 periods of life and from the optimal strategy (opt.).

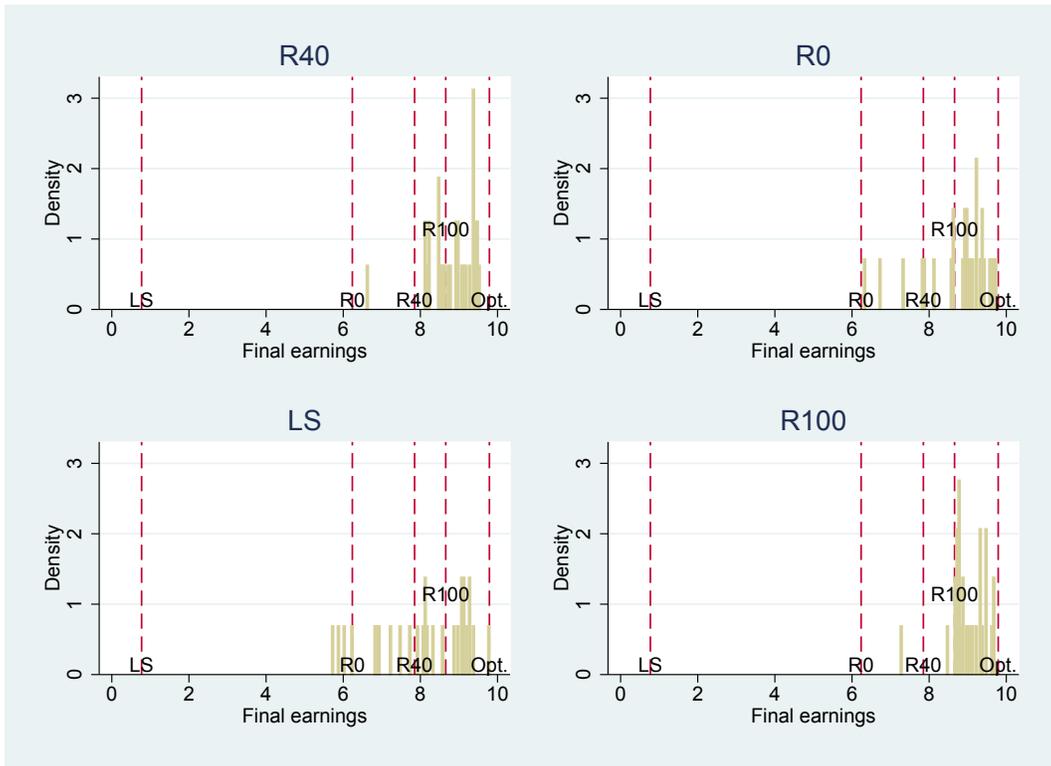


Figure A8: Distribution of final earnings across treatments in Sequence 3

Note: The vertical lines indicate final earnings from a strategy of consuming endowments in each of the 25 periods of life and from the optimal strategy (opt.).

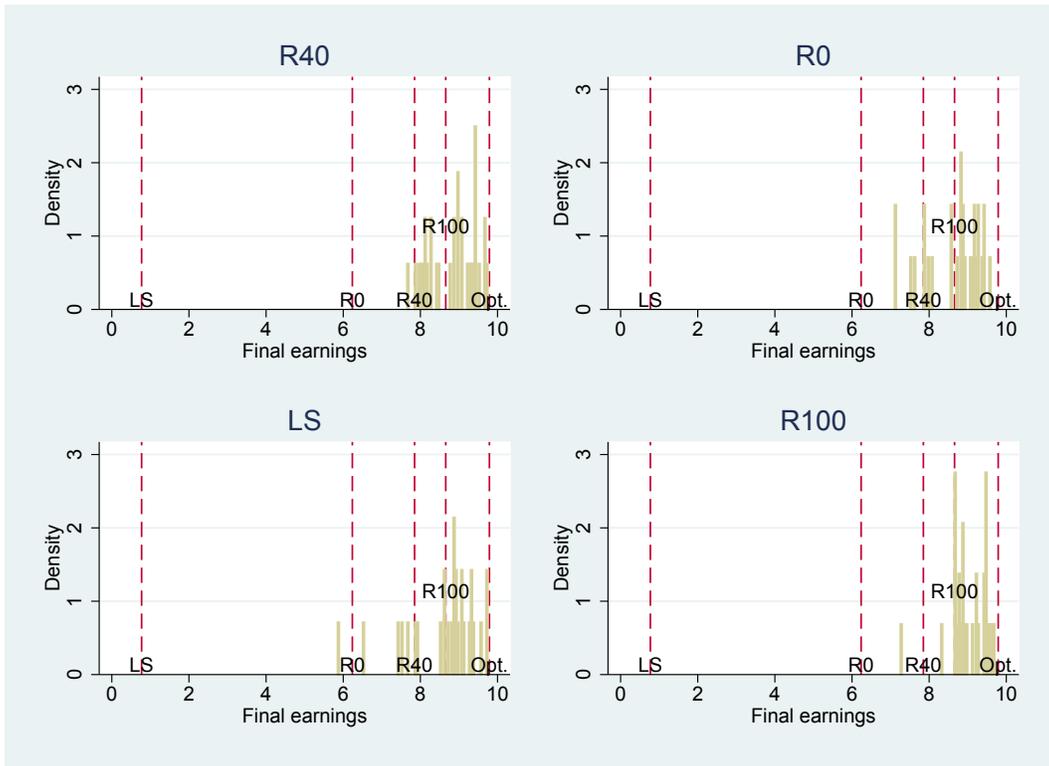


Figure A9: Distribution of final earnings across treatments in Sequence 4

Note: The vertical lines indicate final earnings from a strategy of consuming endowments in each of the 25 periods of life and from the optimal strategy (opt.).