# Liquidity constraints and buffer stock savings: Theory and experimental evidence[*]

John Duffy[†] Andreas Orland[‡]

September 25, 2020

**Abstract**

We provide a direct, experimental test of the buffer stock model of savings behavior. We use a three-period intertemporal model of consumption/savings decisions where liquidity in the second period is constrained (and, thus, borrowing is not possible). We contrast behavior in this constrained version of the model with an unconstrained version where there is no liquidity constraint. A second treatment variable is the variance of the stochastic income process, resulting in a $2 \times 2$ experimental design. We test the comparative statics predictions of the model and find that, in contrast to these predictions, the liquidity constraint does not increase savings in the first period of the constrained model relative to the first period of the unconstrained model. However, we find strong support for all the other comparative statics predictions of the model, e.g., the impact of a higher variance of income on savings behavior and differences between period 1 and period 2 savings. In further analyses, we find that we can rationalize the departures we observe from model predictions by some combination of debt aversion, heterogeneity in cognitive abilities and/or learning.

**Keywords:** Experimental economics, intertemporal optimization, liquidity constraints, consumption, saving, debt aversion.

**JEL codes:** C91, D92, E52

---

[*]This experimental study has been approved by UCI's Institutional Review Board. A pre-registration of this study can be found here: `https://osf.io/hxfrj`.

[†]Department of Economics, University of California, Irvine, United States. E-mail: duffy@uci.edu.

[‡]Department of Economics, University of Potsdam, Potsdam, Germany. E-mail: aorland@uni-potsdam.de.

# 1 Introduction

A workhorse model in theoretical and empirical analyses of consumption is the buffer stock model of precautionary savings (Deaton, 1991; Besley, 1995; Carroll, 1997). This model posits that agents, facing an uncertain income, target a certain ratio of cash-on-hand (i.e., wealth plus disposable income) to permanent income and use savings as a buffer to smooth out fluctuations in income relative to this target.[1] When individuals are below their target, they save more (consume less) while when they are above their target they save less (consume more). In a world with perfect credit markets, the only explanation for such buffer stock savings is that individuals have "precautionary" motives in the face of fluctuating income. Such precautionary motives arise naturally if agents have certain preferences (e.g., constant relative risk aversion; CRRA) where the marginal utility of consumption is convex so that agents are "prudent" (Kimball, 1990). The buffer stock target arises as a compromise between prudent savings behavior, when cash-on-hand is below target, and impatience with regard to delayed consumption, when cash-on-hand is above target. However, buffer stock savings motives do not *require* prudent preferences. If credit markets are imperfect, so that agents face liquidity or borrowing constraints, then buffer stock savings behavior can also arise, even if agents do not have prudent preferences. As it is difficult to know the functional form that best characterizes individual preferences, the liquidity constraint explanation for buffer stock savings seems more generalizable for testing purposes. Our aim in this paper is to directly examine this liquidity constraint rationale for buffer stock savings behavior, and more generally to provide an empirical evaluation of the buffer stock model of precautionary savings.

The role played by liquidity constraints in buffer stock savings behavior is difficult to evaluate in the field. For instance, we rarely observe the precise constraints that agents face or the shocks

[1] This buffer stock model is at odds with the lifecycle/permanent income hypothesis of Modigliani and Brumberg (1954) and Friedman (1957), which posits that agents' primary motivation for savings are lifecycle concerns such as purchasing a home or having income in retirement.

to their income. For this reason, we test the role of liquidity constraints for buffer stock savings in a controlled laboratory setting, using experiments with paid human subjects.[2] While several experiments have been conducted testing the lifecycle/permanent income theory of savings (see Ballinger et al 2003, Brown et al 2009, and Ballinger et al 2011 for prominent previous studies, and Duffy 2016 for a survey), we are not aware of any experimental research testing how the presence or absence of liquidity constraints affects savings behavior when agent face uncertain income. An experiment conducted by Meissner (2016) does not involve any liquidity constraints, but does find evidence for *debt aversion*, a result that we also model and use to explain our own findings.

The empirical relevance of liquidity constraints is well established. Jappelli (1990, p. 220) estimates that 19% of U.S. households are liquidity constrained and discusses the characteristics of these households. Gross and Souleles (2002, p. 153) estimate the share of potentially liquidity constrained households in the US to be over 66%. Using data from credit card accounts, they show that an increase in credit card borrowing limits generates an immediate rise in debt and is strongest for people close to their credit limit prior to the increase. Liquidity constraints may not be restricted to those with little or no wealth; Boar et al (2017) suggest that 82% of all U.S. *homeowners* may be considered liquidity constrained in the sense that they would benefit from an increase in their liquid assets and mortgage debt by the same amount, keeping their real wealth unchanged. Thus, we view liquidity constraints as a common and empirically relevant phenomenon and we ask how savings behavior responds to the presence or absence of such constraints.

[2]There are other advantages of laboratory experiments over work with field data: In a laboratory experiment, the experimenter has control over all important variables (including subjects' information about income, the planning horizon, etc.) and can induce a specific utility function (and thus eliminate questions about which function best reflects subjects' preferences). Most importantly, the experimenter randomly assigns subjects to treatments and avoids problems with self-selection. For example, people may choose or avoid jobs with a high income uncertainty (e.g., German civil servants enjoy a very certain income, which may be correlated with their risk preferences, see Fuchs-Schündeln and Schündeln 2005), or they are liquidity constrained (which correlates with many consumer characteristics, see Jappelli 1990). Thus, we consider laboratory experiments to be an important complement to empirical work on savings behavior using field data.

Our experiment implements a simple, three-period model of intertemporal consumption and savings choices due to Besley (1995, pp. 2141–2144) that is further detailed in Jappelli and Pistaferri (2017, pp. 115–118). In this framework, individuals face uncertain income realizations in each of the three periods, and the real possibility that liquidity constraints may be binding for how much they can borrow in order to intertemporally smooth their consumption. Besley considers the case where agents face liquidity (or no borrowing) constraints in the first two periods of the three-period model; in the final third period, there is also no borrowing as agents simply consume all remaining resources. Here, we consider a different version of the three-period model described in Jappelli and Pistaferri where individuals face a liquidity constraint only in period 2, as this setup allows agents to *anticipate* the possibility, in period 1, that the liquidity constraint will be binding on their period 2 consumption choice so that they may rationally adjust their period 1 savings behavior accordingly. We contrast the case where agents face a known, period 2 liquidity constraint with the case where they are unconstrained in their period 2 borrowing. In both cases, the model predicts that agents target a critical level of cash-on-hand, defined by parameters of the model, in order to determine their savings/borrowing decisions. If realized income is above this target level, agents save a fraction of that income while if realized income is below the target level they borrow a fraction of the shortfall. In the case where a liquidity constraint is known to be binding in period 2, agents should anticipate the impact of that constraint and adjust their savings and borrowing decisions in period 1. Thus, we can examine the threshold predictions in both settings and whether liquidity constraints affect saving and borrowing behavior in the manner prescribed by the theory. In addition, we also consider two different values for the variance of the uncertain income realizations, as theory predicts that savings are positively associated with the volatility of uncertain income realizations. Specifically, we consider cases where income realizations are drawn uniformly over a large or small interval so that there is high or low variance in these income realizations. In the high variance case, the liquidity constraint is an even more binding concern than in the low variance case, so this dimension of the choice problem also has implications for savings and borrowing behavior. Thus, we have a $2 \times 2$ experimental design where the treatment variables are: 1) liquidity constraints or no liquidity constraints, and 2) the variance of income realizations is high or low. We test the model's comparative statics predictions concerning savings behavior in all four cells. To our knowledge, this is the first experimental test of the importance of liquidity constraints for buffer stock savings behavior.

The solution to the optimization problem requires that agents proceed recursively, applying backward induction. We explore whether individual subjects differ in their predispositions to apply backward induction in our intertemporal savings experiment by collecting data on whether and how quickly they can solve a single-player Tower of Hanoi game (introduced in 1883 by Édouard Lucas). This game provides a measure of subjects' ability to apply recursive or backward induction reasoning. We also elicit measures of subjects' cognitive abilities and risk attitudes using standard measures (a cognitive reflection test and self-reported grade point average for cognitive abilities and paired lottery choice lists for risk aversion). In this dimension of our experiment, we contribute to the literature on the heterogeneity in savings behavior (see, e.g., Meghir and Pistaferri 2004 and Schaner 2015).

There are a number of empirical and simulation studies evaluating the predictions of the buffer stock savings model. Most studies consider the buffer stock model without liquidity constraints, as such constraints are difficult to observe. On the one hand, there is research supportive of the buffer stock saving model's predictions. For instance, Carroll et al (1992) find that expectations about future unemployment rates are closely associated with the level of saving; Carroll (1997) reports that buffer stock savings explain three empirical puzzles of the lifecycle/permanent income theory; Love (2006) presents results from a model where both unemployment benefits and a tax-sheltered retirement account are included; Jappelli and Pistaferri (2020) find that consumers adjust their *target* wealth approximately one-to-one with their permanent income. On the other hand, there is also research that challenges the predictions of the buffer stock savings model. For example, Ludvigson and Michaelides (2001) find that it cannot (fully) explain excess smoothness and excess sensitivity of consumption; Jappelli et al (2008) test simulated predictions of the model with Italian data and reject the predictions of buffer stock model; Fulford (2015) finds that income uncertainty does not affect liquid savings and that households would rather save to smooth expenditure shocks. Fewer papers tackle liquidity constraints in the context of buffer stock savings models. Zeldes (1989) compares models with and without liquidity constraints and carefully concludes that liquidity constraints affect the consumption of a significant portion of the population; Campbell and Hercowitz (2018) set up a model where the household consumes both a standard good and a special good that is consumed on a regular basis but not often (think of college education or housing) and report a good performance of their model.

To preview our results, we find mixed support for the predictions of the buffer stock savings model. While the data show that subjects behave in line with most of the comparative statics predictions of the model, we reject the main hypothesis that a known liquidity constraint in the second period of our three-period model increases savings in the first period relative to the model without this liquidity constraint. We identify three behavioral explanations for this observed departure. First, we find strong evidence for aversion to borrowing or "debt aversion" both at the extensive margin (of whether or not to borrow or to save given a particular income realization) and at the intensive margin of how much to borrow or to save. We derive the savings function for a debt-averse individual and test the performance of this debt-aversion model relative to the constrained or unconstrained models. We find that when the variance in income is high, the debt-averse model yields the best fit to our experimental data in terms of minimizing the root-mean-square error. Second, we find that cognitive abilities also play a role;

subjects who score better in a cognitive reflection test generally behave closer to theoretical predictions. Finally, there is evidence for learning over time; subjects reduce their deviations from theoretical predictions by about 15 to 25% over the course of the experiment.

The remainder of the paper is structured as follows. Section 2 introduces the theory, and Section 3 explains the experimental design and procedures. Section 4 presents the model's comparative statics which serve as our testable hypotheses. Section 5 reports on the main experimental findings and Section 6 puts forward behavioral explanations for the observed deviations from the theory. Finally, Section 7 summarizes and concludes.

## 2 Theory

In this section, we describe the theoretical model we test in our experiment. Specifically, we derive the savings functions for the unconstrained model and the liquidity constrained model.[3] In all settings, an individual lives for three equidistant periods, $t = 1$, $t = 2$ and $t = 3$. We use a three-period model as it is the simplest framework in which to characterize the role of liquidity constraints on buffer stock savings behavior.[4] In all three periods, the individual receives an uncertain income that is known to be independently and identically distributed according to a discrete uniform distribution with support $[y_{min}, y_{max}]$ and mean income, $\mu$. At the start of both the first and the second periods, individuals first learn their income realization for that period and then make a consumption decision for the period which implies also a certain saving/borrowing decision. Savings/borrowings are then automatically paid back in the following period without bearing interest, that is, all loans are one-period. Since there is no discounting between periods,

[3]The derivation of similar models is also shown in Besley (1995, pp. 2141–2144), Carroll and Kimball (2001, pp. 36–38), and Jappelli and Pistaferri (2017, pp. 115–118). More detailed derivations can be found in Section A.1 of the Appendix.

[4]Further, with a three-period model (as opposed to a many-period lifecycle model), subjects can be repeatedly confronted with making consumption and savings decisions in that three-period model, so that we can also consider the role of learning or experience when evaluating model predictions. If subjects cannot achieve the optimum in a three-period model, it seems unlikely they would fare better in a many-period lifecycle model.

the risk-free rate of interest is set to zero.[5] In the third period, the individual learns her income for that period but simply consumes all remaining resources and thus makes no decisions in that period. We assume that individuals evaluate per-period consumption using the period utility function $u(c) = \ln(c)$.[6] Figure 1 illustrates the timing of the model.

**Figure 1:** Timing in the three-period lifecycle model

| $t = 1$ | $t = 2$ | $t = 3$ |
|---|---|---|
| The individual<br>• receives income $y_1$,<br>• saves/borrows $s_1$,<br>• consumes $c_1 = y_1 - s_1$,<br>• receives utility $u(c_1)$. | The individual<br>• receives income $y_2$,<br>• saves/borrows $s_2$,<br>• consumes $c_2 = y_2 + s_1 - s_2$,<br>• receives utility $u(c_2)$. | The individual<br>• receives income $y_3$,<br>• takes no savings decision,<br>• consumes $c_3 = y_3 + s_2$,<br>• receives utility $u(c_3)$. |

Thus, the individual's optimization problem in period 1 is given by:

$$\max_{s_1, s_2} u(c_1) + \mathbb{E}_{t=1}\, u(c_2) + \mathbb{E}_{t=1}\, u(c_3) \tag{1}$$

subject to

$$c_1 + s_1 = y_1, \tag{2}$$

$$c_2 + s_2 = s_1 + y_2, \tag{3}$$

$$c_3 = s_2 + y_3. \tag{4}$$

Given our assumption of log preferences and substituting the budget constraints (Equations 2-4) into the objective function (Equation 1), the optimization problem can be rewritten as:

$$\max_{s_1, s_2} \ln(y_1 - s_1) + \mathbb{E}_{t=1} \ln(y_2 + s_1 - s_2) + \mathbb{E}_{t=1} \ln(y_3 + s_2). \tag{5}$$

First, we derive the savings functions in periods 1 and 2 for the *unconstrained* model where

[5] One reason for this choice is that the periods in our experiment are not that far apart in time. The main advantage of having a zero interest rate is that it simplifies the task for subjects in the experiment. Note that our design does not imply impatience, which is, besides liquidity constraints, another way to induce buffer stock savings (Deaton, 1991; Carroll, 2004).

[6] Note that with log preferences, $u'''(c) = c^{-4} > 0$ so the individual should prudently respond to risk, i.e., there is an operative precautionary savings motive.

the individual can save *and* borrow in periods 1 and 2. We use backward induction and solve the individual's two-period problem in period 2 for $s_2$. Applying the expectation operator to the uncertain income in period 3 leaves us with a savings function that depends only on the deviation of total wealth (current income plus last period's savings), $w_2$, from mean income:

$$s_2^{\text{Unc}} = s_2^*(y_2, s_1) = \frac{1}{2}(y_2 + s_1 - \mu) = \frac{1}{2}(w_2 - \mu). \tag{6}$$

Thus, in period 2, the individual should *save* if her wealth is above mean income and *borrow* if her wealth is below mean income. We use this solution and substitute it into the initial three-period optimization problem and use the procedure described above to solve for $s_1$:

$$s_1^{\text{Unc}} = s_1^*(y_1) = \frac{2}{3}(y_1 - \mu). \tag{7}$$

Similarly to period 2, we find that the individual should save if wealth (here, only current income) is above mean income and borrow if it is below. Both solutions are linear functions with different slopes.

Second, we derive the savings functions for the *constrained* model. The only difference between the two models is the known liquidity constraint that an individual faces in period 2: without constraints, $s_2$ can take negative values (i.e., borrowing is possible); with the liquidity constraint, however, borrowing is not possible so that $s_2 \geq 0$. Again, we start in period 2 and solve for $s_2$. Solving the inequality leaves us with a savings function that is defined piecewise:

$$s_2^{\text{Con}} = s_2^*(y_2, s_1) = \begin{cases} \frac{1}{2}\left[y_2 + s_1 - \mu\right] & \text{if } \ y_2 + s_1 \geq \mu \\ 0 & \text{otherwise} \end{cases} = \begin{cases} \frac{1}{2}\left[w_2 - \mu\right] & \text{if } \ w_2 \geq \mu \\ 0 & \text{otherwise} \end{cases}. \tag{8}$$

This convex function has a kink at the point where total wealth is equal to mean income. When wealth is below mean income, savings are zero; above this point the savings function is identical to the one in the unconstrained model in Equation 6. We substitute this solution for period 2 savings into the three-period problem and we solve for period 1 savings. This involves solving a quadratic equation. Thus, the savings function in period 1 has two solutions, of which only one makes economic sense.[7] As in Equation 8, it is convex:

$$s_1^{\textsc{Con}} = s_1^*(y_1) = \pm 2\sqrt{\mu}\sqrt{7\mu - 2y_{max} - y_1} + 5\mu - y_{max}. \tag{9}$$

We use the unconstrained and constrained solutions for the design of our experimental treatments and, later, for the analyses of our experimental data. In the next section, we introduce the parameterization of our experimental treatments and we display the unconstrained and constrained savings functions for those parameterizations.

# 3 Experimental design and procedures

## 3.1 Experimental design

In our experiment, we vary both the existence of a liquidity constraint (as shown in Section 2) and the variance of the income distribution (by manipulating the mean-preserving spread) in a $2 \times 2$ experimental design shown in Table 1. The rationale for the latter treatment variable is that differences in the variance of income affect period 1 and 2 savings functions differently, as detailed below.

**Table 1:** Experimental design

| | | Variance | |
|---|---|---|---|
| | | High | Low |
| Liquidity | Unconstrained | UncHigh | UncLow |
| | Constrained | ConHigh | ConLow |

We label the treatments UncHigh (unconstrained savings decision in a high variance environment), UncLow (unconstrained savings decision in a low variance environment), ConHigh (constrained savings decision in a high variance environment), and ConLow (constrained

[7] In both parameterizations we introduce in the next section, the economically sensible solutions are the negative solutions.

savings decision in a low variance environment). In the HIGH treatments, we implement the following parameters for the discrete uniform distribution: $y_{min} = 35$ and $y_{max} = 105$ (with $\mu = 70$ and $\sigma = 20.5$). In the LOW treatments, we instead set $y_{min} = 60$ and $y_{max} = 80$ (so that $\mu = 70$ and $\sigma = 6.1$). The standard deviation in the HIGH treatments is thus more than three times as high as in the LOW treatments (while keeping the mean constant).

Using our parameterization of the model, Figure 2 presents the two models' savings functions (and the difference between the two) for different possible income levels in period 1 and for different possible wealth levels in period 2 in both the HIGH and LOW treatments. In Figure 2a, one can see that for the HIGH treatment, optimal savings, in the constrained case, $s_1^{\text{CON}}$, is always greater than optimal savings in the unconstrained case, $s_1^{\text{UNC}}$ for all possible period 1 income levels.[8] By contrast for the LOW treatment, one can see in Figure 2b that (i) for low income levels, the constrained optimal savings, $s_1^{\text{CON}}$ is greater than the unconstrained optimal savings, $s_1^{\text{UNC}}$; (ii) the two savings functions intersect close to the mean income; and (iii) for high income levels, the unconstrained optimal savings, $s_1^{\text{UNC}}$ is greater than the constrained optimal savings, $s_1^{\text{CON}}$. Figures 2c and 2d display the optimal savings functions for period 2. Here, the functions are defined for the range of possible wealth levels (defined from zero up to twice the maximum income level, $y_{max}$; see also the description of the natural savings and borrowings constraints in our experiment in the next paragraph). Both figures are very similar with a positive difference between the constrained optimal period 2 savings, $s_2^{\text{CON}}$ and the unconstrained optimal period 2 savings, $s_2^{\text{UNC}}$, for wealth levels below the mean income (70) and no difference above it. Later, we use the functions as they are shown here to derive predictions for the different treatments which we use to analyze the experimental data both at the aggregate and individual levels.

[8]Note that we need a model with *at least* three periods in order to generate this anticipation effect of the liquidity constraint on period 1 savings. As noted earlier, our experiment uses the simplest possible model to examine our research questions.

**Figure 2:** The savings functions in all experimental treatments

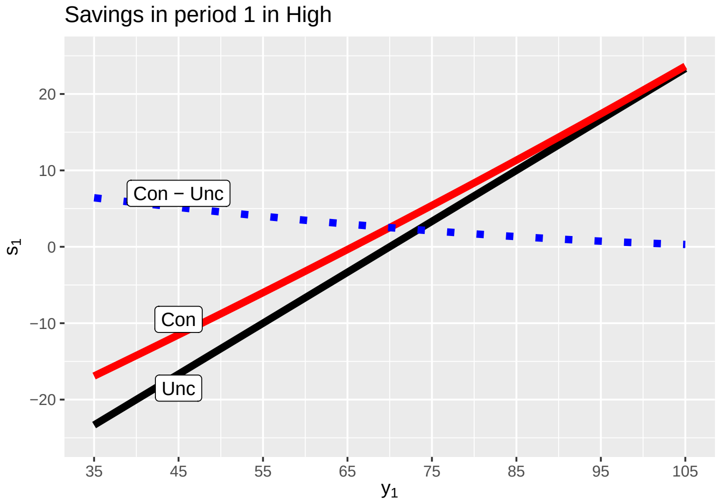


**(a)** Savings in period 1 in HIGH

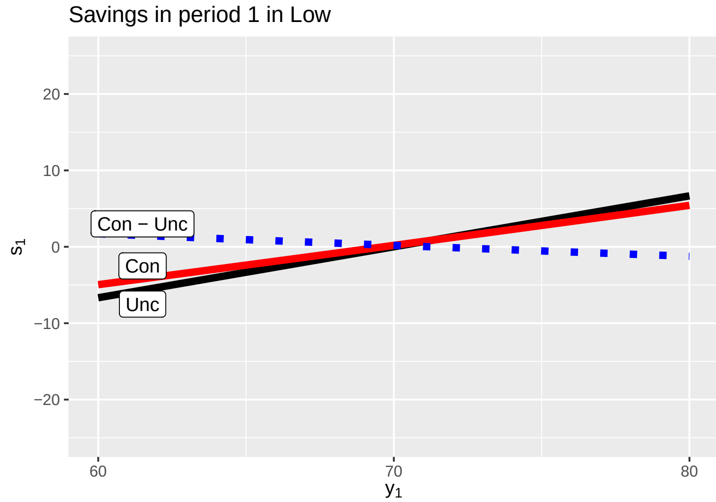


**(b)** Savings in period 1 in LOW

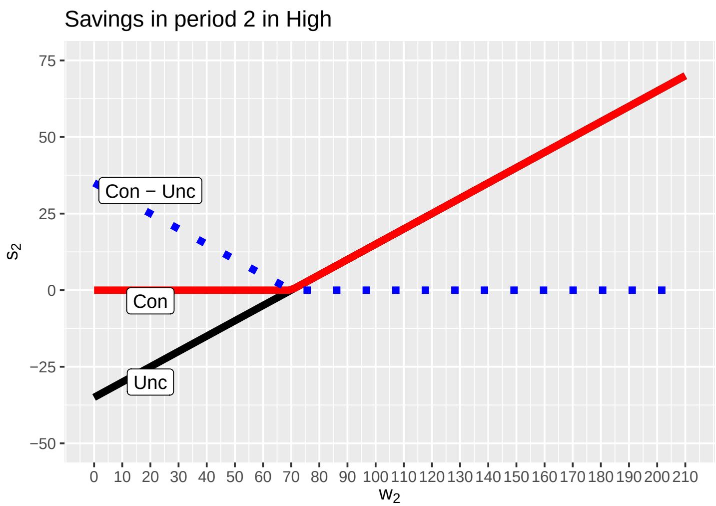


**(c)** Savings in period 2 in HIGH

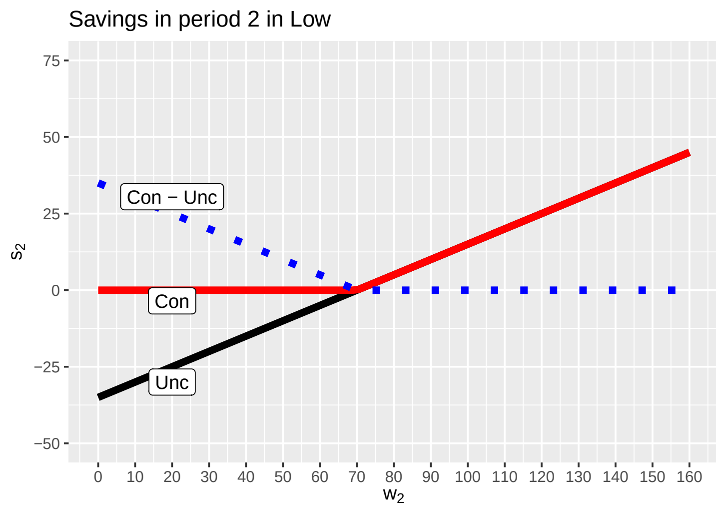


**(d)** Savings in period 2 in LOW

In all four treatments, we induce the same, per-period utility function $u(c) = 0.77 \ln(c)$; this function transforms experimental currency consumption in each period into dollar payoffs. This utility function requires strictly positive consumption. Thus, we need to introduce the following natural saving and borrowing constraints (that are in addition to the liquidity constraint explicitly spelled out in the model). The natural saving constraints ensure that current consumption is strictly non-negative: $s_1 < y_1$ and $s_2 < w_2$. The natural borrowing constraints ensure positive future consumption (in case the worst-possible outcome occurs): $s_1 > -y_{min}$ and $s_2 > -y_{min}$. In the experiment, we further introduce a minimum consumption constraint, $c_1 \geq 1$ and $c_2 \geq 1$, and maximum consumption constraints of $c_1 \leq y_1 + y_{min} - 1$ and $c_2 \leq w_2 + y_{min} - 1$. These constraints prevent subjects from losing money in a period of the experiment and do not affect the theoretical predictions in either the HIGH or LOW treatment parameterizations.

### 3.2 Experimental procedures

We combine a within-subject design approach (where each subject went through 15, three-period lifecycles of both CON and UNC) with a between-subject approach (each subject was randomly assigned to either the HIGH or to the LOW treatment). One half of subjects in both the HIGH and the LOW treatments first made decisions in 15 lifecycles of the CON setting, followed by 15 lifecycles of the UNC setting (the CON-UNC order); the other half of subjects made decisions the other way around (the UNC-CON order). This design allows us to control for learning and order effects. At the end of the experiment, we randomly chose one three-period lifecycle (or "round", as lifecycles are called in the instructions) in the CON treatment and we randomly chose another one in the UNC treatment for payoff and this fact was known to subjects.

Given our experimental design, our model calibration, and our testable hypotheses, we determined the number of subjects that would be required to properly evaluate our hypotheses using a power calculation; the details are provided in the next section and in Section A.2 of the

Appendix). In line with the power calculations, 100 subjects took part in the HIGH treatment (50 in the CON-UNC order, 50 in the UNC-CON order) and 50 subjects took part in the LOW treatment (25 in the CON-UNC order, 25 in the UNC-CON order). Thus, we report data from a total of 150 subjects with no prior experience with the environment of the experiment.

At the start of each session, all subjects were told that the experiment consisted of three independent parts and that they would receive instructions for the next part after finishing the previous part. Part 1 consisted of the first 15 rounds and Part 2, the second 15 rounds. The experiment was framed using a savings/consumption context. In all rounds, subjects could enter their consumption decisions with up to one decimal place. At the end of each round, subjects were provided with feedback about all past rounds of the relevant part.

Part 3 consisted of two tasks: a risk preference elicitation task (as more risk-averse individuals should save more, relative to the predictions of the theory) and a Tower of Hanoi (ToH) game (which requires the ability to think recursively, a skill that is closely related to backward induction). Both tasks in Part 3 were monetarily incentivized; we randomly chose one of the two tasks (with equal probability) for payment, but subjects did not know which of the two tasks would be chosen. To elicit subjects' risk preferences, we used the procedure proposed by Drichoutis and Lusk (2016) which consists of multiple choices between pairs of lotteries. Table A.2 in the Appendix shows the ten pairs of lotteries each subject had to choose between. (If the risk elicitation task was chosen for payoff, then one of the ten decisions was chosen and the subject's payoff was determined randomly according to their preferred lottery for that decision.) The ToH game is a one-player mathematical puzzle (first described by the French mathematician Édouard Lucas). Subjects have to move different-sized disks from the leftmost peg to the rightmost peg according to some rules and using a minimum number of moves. (If the ToH task was chosen for payoff, then the more moves a subject needed in order to solve the game, the lower was that subject's payoff.) Details on the ToH game can be found in Section A.5 in the Appendix.

After all payoff-relevant decisions have been taken by the subjects (but before feedback about their total dollar earnings was shown), all subjects filled out a post-experimental questionnaire asking for their age, gender, field of study, and grade point average (GPA). They then completed an unincentivized cognitive reflection test (CRT) by Toplak et al (2014). The CRT score (the sum of correct answers) provides a measure of subjects' abilities to override reactive "system 1" thinking and instead employ more reflective, "system 2" thinking to solve problems; it asks four questions (shown in Table A.3 in the Appendix) similar to the three original CRT questions by Frederick (2005). The four questions we used are not as widely known as the original three CRT questions, and that is why we chose to use these four questions instead.

Upon arrival at the laboratory, all subjects were seated at computer workstations with privacy walls. Communication between subjects was prohibited. The subjects all received printed instructions which included a graph and a table showing monetary payoffs (utility) for all integer values of the per-period utility function. These instructions were read aloud and thereafter, subjects had to correctly answer a set of control questions in order to proceed.[9] The experiment was computerized and programmed using oTree (Chen et al, 2016). Each session lasted about 90 minutes. After the experiment, all subjects were paid in cash and in private. We conducted twelve sessions with a total of 150 subjects in the Economic Social Science Laboratory (ESSL) at the University of California, Irvine between October 2019 and February 2020. Subjects were undergraduate students from various fields of study. We did not apply any exclusion criteria to the registered subjects in the database used to recruit subjects. Every subject took part in one session only. Subjects earned on average \$24.08 (minimum \$18.30, maximum \$26.80), plus a show-up fee of \$7.

[9] The appendix contains the experimental instructions (Section B.1), the graph and table (Section B.2), the control questions (Section B.3), and example screenshots (Section B.4).

# 4 Aggregate predictions and hypotheses

In this section we present aggregate predictions for all treatments and use these to form testable hypotheses. Table 2 displays means and standard deviations for optimal savings in periods 1 and 2 of all four treatments. These numerical solutions are based on the savings functions and income and wealth ranges as shown in Figure 2. In Section A.2 in the Appendix, we use these predictions to calculate the number of required subjects.[10]

**Table 2:** Aggregate predictions and hypotheses

| | | High | | | | | Low | | |
|---|---|---|---|---|---|---|---|---|---|
| | | Unc | Con | | | | Unc | Con | |
| | | | | | H4 | | | | |
| | | H1a | | | | | H1b | | |
| $s_1$ | | 0.0 | 2.8 | | $s_1$ | | 0.0 | 0.2 | |
| | | (13.7 ) | (11.9 ) | | | | (4.0) | (3.1) | |
| | H3a | | | H3b | | H3c | | | H3d |
| $s_2$ | | 17.5 | 23.4 | | $s_2$ | | 5.0 | 12.7 | |
| | | (30.5 ) | (23.4 ) | | | | (23.2 ) | (14.9 ) | |
| | | H2a | | | | | H2b | | |
| | | | | | H5 | | | | |
| | | | | | H6 | | | | |

Note: The cells in the middle present mean predictions and their standard deviations (in parentheses). The curly brackets show which variables are compared to test our hypotheses (marked with H).

The first two sets of hypotheses (H1a-H1b and H2a-H2b) test the consequences of the liquidity constraint. These two, between-treatment, within-subject hypotheses are the main focus of our study. Namely, we expect that the liquidity constraint makes subjects save *more* in both periods 1 and 2 of the Con treatment as compared with the Unc treatment of the same income variance. In period 2, subjects should save more because of the liquidity constraint, while in period 1 it is the anticipation of the liquidity constraint that should increase their savings.

**Hypothesis 1a:** In the ConHigh treatment, $s_1$ is higher than in the UncHigh treatment.

[10] Means and standard deviations change if we restrict the range in period 2 to the income range. All hypotheses but H3a and H3c, however, do not change (H3a and H3c depend on an asymmetric wealth distribution around the mean income with more higher wealth levels than lower wealth levels). As we did not know how subjects would behave in period 1 (which influences total wealth in period 2), we chose the complete possible wealth range in period 2 (from 1 to $2y_{max} - 1$) for the predictions and power calculations.

**Hypothesis 1b:** In the ConLow treatment, $s_1$ is higher than in the UncLow treatment.

**Hypothesis 2a:** In the ConHigh treatment, $s_2$ is higher than in the UncHigh treatment.

**Hypothesis 2b:** In the ConLow treatment, $s_2$ is higher than in the UncLow treatment.

The next set of hypotheses examine savings behavior within each treatment:

**Hypothesis 3a:** In the UncHigh treatment, $s_2$ is higher than $s_1$.

**Hypothesis 3b:** In the ConHigh treatment, $s_2$ is higher than $s_1$.

**Hypothesis 3c:** In the UncLow treatment, $s_2$ is higher than $s_1$.

**Hypothesis 3d:** In the ConLow treatment, $s_2$ is higher than $s_1$.

Finally, the last three between-treatment between-subject hypotheses examine how the variance of income affects savings behavior: higher income uncertainty should cause higher savings.

**Hypothesis 4:** In the ConHigh treatment, $s_1$ is higher than in the ConLow treatment.

**Hypothesis 5:** In the UncHigh treatment, $s_2$ is higher than in the UncLow treatment.

**Hypothesis 6:** In the ConHigh treatment, $s_2$ is higher than in the ConLow treatment.

# 5 Results

## 5.1 Comparative statics

Table 3 summarizes all of our main, aggregate-level findings. This table reports mean savings (and standard errors) for the two periods ($s_1$, $s_2$) of all four treatments (note that we do not exclude any observations in the following analyses unless explicitly stated). In a first step, we test the hypotheses formulated in Section 4 using two-sided t-tests (if the difference has the predicted sign) and one-sided t-tests (if the difference does *not* have the predicted sign).[11] While we reject the main hypothesis set H1a-H1b—that the anticipation of a liquidity constraint in the second period increases savings in the first period compared to a situation without a liquidity

[11] We use the two-sided (non-directional) t-test if the difference between the compared means has the predicted sign (this, in comparison to the one-sided t-test, more conservative test is also pre-registered). If the sign is contrary to prediction, we report the one-sided t-test (which is adequate, given our directional hypotheses).

constraint—we cannot reject hypotheses H2-H6 concerning various other comparative statics predictions of the model. Table 4 shows the detailed findings.[12]

**Table 3:** Aggregate results

| | | High | | | | | Low | | |
|---|---|---|---|---|---|---|---|---|---|
| | | Unc | Con | | | | Unc | Con | |
| | | | | | H4: ✓✓✓ | | | | |
| | | H1a: ✗ | | | | | H1b: ✗ | | |
| $s_1$ | H3a: ✓✓ | 8.62 | 6.97 | H3b: ✓✓✓ | $s_1$ | H3c: ✓ | 1.42 | −0.55 | H3d: ✓✓✓ |
| | | ( 1.66) | ( 1.76) | | | | ( 1.64) | ( 1.51) | |
| $s_2$ | | 12.34 | 19.69 | | $s_2$ | | 4.53 | 8.99 | |
| | | ( 2.61) | ( 2.52) | | | | ( 2.75) | ( 1.91) | |
| | | H2a: ✓✓✓ | | | | | H2b: ✓✓ | | |
| | | | | H5: ✓✓ | | | | | |
| | | | | | H6: ✓✓✓ | | | | |

Note: Means with standard errors clustered at the subject-level in parentheses. ✓✓✓, ✓✓, and ✓ mark hypotheses that cannot be rejected at the 1%-, 5%-, and 10%-level, respectively. ✗ marks hypotheses that are rejected. Significance tests based on t-tests clustered at the subject-level (details in Table 4).

**Table 4:** Tests of the hypotheses

| | Hypothesis | | | Difference | | | |
|---|---|---|---|---|---|---|---|
| Result 1a | $s_1^{\text{ConHigh}}$ | $>$ | $s_1^{\text{UncHigh}}$ | $s_1^{\text{ConHigh}} - s_1^{\text{UncHigh}}$ | $=$ | $-1.65$ | $(p = 0.835)$ |
| Result 1b | $s_1^{\text{ConLow}}$ | $>$ | $s_1^{\text{UncLow}}$ | $s_1^{\text{ConLow}} - s_1^{\text{UncLow}}$ | $=$ | $-1.97$ | $(p = 0.9575)$ |
| Result 2a | $s_2^{\text{ConHigh}}$ | $>$ | $s_2^{\text{UncHigh}}$ | $s_2^{\text{ConHigh}} - s_2^{\text{UncHigh}}$ | $=$ | 7.35 | $(p = 0.001)$ |
| Result 2b | $s_2^{\text{ConLow}}$ | $>$ | $s_2^{\text{UncLow}}$ | $s_2^{\text{ConLow}} - s_2^{\text{UncLow}}$ | $=$ | 4.46 | $(p = 0.013)$ |
| Result 3a | $s_2^{\text{UncHigh}}$ | $>$ | $s_1^{\text{UncHigh}}$ | $s_2^{\text{UncHigh}} - s_1^{\text{UncHigh}}$ | $=$ | 3.72 | $(p = 0.011)$ |
| Result 3b | $s_2^{\text{ConHigh}}$ | $>$ | $s_1^{\text{ConHigh}}$ | $s_2^{\text{ConHigh}} - s_1^{\text{ConHigh}}$ | $=$ | 12.72 | $(p < 0.001)$ |
| Result 3c | $s_2^{\text{UncLow}}$ | $>$ | $s_1^{\text{UncLow}}$ | $s_2^{\text{UncLow}} - s_1^{\text{UncLow}}$ | $=$ | 3.11 | $(p = 0.083)$ |
| Result 3d | $s_2^{\text{ConLow}}$ | $>$ | $s_1^{\text{ConLow}}$ | $s_2^{\text{ConLow}} - s_1^{\text{ConLow}}$ | $=$ | 9.54 | $(p < 0.001)$ |
| Result 4 | $s_1^{\text{ConHigh}}$ | $>$ | $s_1^{\text{ConLow}}$ | $s_1^{\text{ConHigh}} - s_1^{\text{ConLow}}$ | $=$ | 7.52 | $(p = 0.001)$ |
| Result 5 | $s_2^{\text{UncHigh}}$ | $>$ | $s_2^{\text{UncLow}}$ | $s_2^{\text{UncHigh}} - s_2^{\text{UncLow}}$ | $=$ | 7.81 | $(p = 0.040)$ |
| Result 6 | $s_2^{\text{ConHigh}}$ | $>$ | $s_2^{\text{ConLow}}$ | $s_2^{\text{ConHigh}} - s_2^{\text{ConLow}}$ | $=$ | 10.70 | $(p = 0.001)$ |

Note: $p$-levels of Results 1a and 1b are based on one-sided t-tests clustered at the subject-level, all other $p$-levels on two-sided t-tests clustered at the subject-level.

In a second step, we compare observed savings with theoretical predictions (this cannot be done by just comparing predicted savings in Table 2 and observed mean savings in Table 3 because savings in the second period depend on both the income realizations and the savings from the previous period, in contrast to savings in the first period, which only depend on the

[12] We also compare $s_1^{\text{UncHigh}}$ with $s_1^{\text{UncLow}}$. This comparison is not among the hypotheses in the previous section or in our pre-registration since both savings predictions are zero. However, the observed difference, $s_1^{\text{UncHigh}} - s_1^{\text{UncLow}} = 7.21$, is significantly different from zero ($p = 0.002$). Hence, the higher variance in income also increases savings in the first period of the unconstrained treatment.

income realizations in that period). Table 5 shows the results of our analysis of *conditional* savings behavior. In all periods of all treatments but period 1 in the Low treatments (where we have both fewer observations and a lower expected effect size than in the High treatments), we observe significant *oversaving* (with savings up to 25 times the prediction).[13] In the first period, oversaving in the unconstrained treatment is higher than in the constrained treatment, whereas in the second period it is the other way around.

**Table 5:** Observed savings, predictions and deviations from predictions

| | High | | Low | |
|---|---|---|---|---|
| | Unc | Con | Unc | Con |
| Period 1 | | | | |
| Observed savings $s_1$ | 8.62*** | 6.97*** | 1.42 | −0.55 |
| | ( 1.66) | ( 1.76) | ( 1.64) | ( 1.51) |
| Optimal savings $s_1^*(y_1)$ | 0.34 | 2.97*** | 0.08 | 0.21* |
| | ( 0.39) | ( 0.29) | ( 0.15) | ( 0.11) |
| Deviation $s_1 - s_1^*(y_1)$ | + 8.29*** | + 4.00** | + 1.34 | - 0.76 |
| | ( 1.68) | ( 1.76) | ( 1.62) | ( 1.52) |
| | difference=-4.28** ($p = 0.010$) | | difference=-2.10* ($p = 0.069$) | |
| Period 2 | | | | |
| Observed savings $s_2$ | 12.34*** | 19.69*** | 4.53 | 8.99*** |
| | ( 2.61) | ( 2.52) | ( 2.75) | ( 1.91) |
| Conditionally optimal savings $s_2^*(y_2, s_1)$ | 4.49*** | 8.17*** | 0.72 | 2.68*** |
| | ( 0.81) | ( 0.65) | ( 0.82) | ( 0.50) |
| Deviation $s_2 - s_2^*(y_2, s_1)$ | + 7.85*** | +11.52*** | + 3.81* | + 6.31*** |
| | ( 1.97) | ( 1.96) | ( 2.17) | ( 1.61) |
| | difference=3.67** ($p = 0.034$) | | difference=2.50, $p = 0.128$ | |

Note: Optimal savings (conditioned on income realization) and conditionally optimal savings (conditioned on income realization and previous period's savings decision) are based on the solutions derived in Section 2. Standard errors clustered at the subject-level in parentheses. ***, **, and * show difference from zero (based on two-sided t-tests clustered at the subject-level) at the 1%-, 5%-, and 10%-level, respectively.

## 5.2 Estimated savings functions

In this section, we estimate savings *functions* based on subjects' decisions and we compare those estimated savings functions with theoretical predictions. We first consider the savings functions for period 1. We estimate a period 1 savings function by regressing savings decisions $s$ by each subject $i$ on period 1 income $y$ using a panel regression estimator (with $u$ being the individual effect and $e$ the disturbance term):

[13]Oversaving, especially for potentially liquidity constrained individuals, might not seem plausible at first sight. However, according to Sahadi (2015), many U.S. taxpayers overpay taxes throughout the year (thus, giving the government an interest-free loan). This results in about 80% of taxpayers receiving refunds (with a mean refund of about $2800 in 2020, see Internal Revenue Service 2020).

$$s_{1,i,\text{round}} = \text{constant} + \beta y_{1,i,\text{round}} + u_i + e_{i,\text{round}} \tag{10}$$

Figure 3 shows the fitted value of this estimated period 1 saving function for all four treatments, along with 95% confidence intervals. These estimated savings functions are shown together with scatterplots of actual (jittered) period 1 savings decisions, $s_1$, against actual period 1 incomes, $y_1$, as well as the predicted period 1 savings function of the theory. We observe the following: (i) there is considerable heterogeneity in observed savings decisions in all treatments as the scatterplots make clear;[14] (ii) in the HIGH treatments (the two top panels) we see that for low period 1 incomes the estimated period 1 savings function is significantly higher than predicted but becomes indistinguishable from the theoretical prediction for higher period 1 incomes; (iii) in the LOW treatments (the two bottom panels) the estimated and theoretical period 1 savings functions overlap except for very high period 1 incomes in UNC, where estimated savings are slightly below predictions (the difference is small and can be observed in tables with exact values).[15]

Next, we consider the savings behavior for period 2, $s_2$, in a manner similar to period 1, but where $s_2$ is linear function of actual period 2 wealth, $w_2$. We again plot actual (jittered) period 2 savings decisions against actual period 2 wealth levels for all four treatments in Figure 4. In those same figures, we report linear estimated savings functions for the UNC treatments that are similar in specification to Equation 10 except that $w_2$ replaces $y_1$, and linear *threshold* estimated savings functions for the CON treatments.[16] Since the theoretical constrained period 2 savings

[14]Such heterogeneity can also be found outside the lab. An enormous dispersion of wealth and savings even among U.S. households with similar socioeconomic characteristics was already reported by Venti and Wise (1998) with data from the Health and Retirement Survey.

[15]Figure A.1 in the Appendix shows the estimated period 1 savings functions together with locally weighted scatterplot smoother filters with a bandwidth of 0.8 (LOWESS filters; Cleveland 1979). In all four treatments, the non-parametric LOWESS filter lies inside the 95% confidence interval. Thus, our linear regression estimations fit the data well.

[16]The idea for threshold estimations was introduced by Hansen (1999). We use the Stata package `xthreg` for panel regressions (due to Wang 2015; the article also gives an introduction to threshold panel regressions).

function in the CON treatments is piecewise linear and the slope of the savings function depends on the wealth realization, $w_2$ the estimated threshold savings function is able to determine the threshold level of wealth, $\gamma$, at which period 2 savings behavior changes:

$$s_{2,i,\text{round}} = \text{constant} + w_{2,i,\text{round}}\left(w_{2,i,\text{round}} < \gamma\right)\beta_1 + w_{2,i,\text{round}}\left(w_{2,i,\text{round}} \geq \gamma\right)\beta_2 + u_i + e_{i,\text{round}} \tag{11}$$

Thus, to estimate the period 2 savings function for the CON treatments, we use Equation 11 to estimate the slope, $\beta_1$ before the estimated threshold, $\gamma$, and the slope, $\beta_2$, after that estimated threshold. In the threshold regression approach, the threshold $\gamma$ is *endogenously* determined by minimizing the residual sum of squares; $u$ is the individual effect and $e$ is the disturbance term.

In Figure 4, we observe the following: (i) in UNCHIGH, period 2 savings is higher than predicted for low period 2 wealth levels ($w_2$) but for higher wealth levels, period 2 savings are indistinguishable from predictions; (ii) in UNCLOW, period 2 savings are indistinguishable from theoretical predictions for all period 2 wealth levels; (iii) in both CON treatments, the endogenously determined thresholds, $\gamma = 115.9$ in CONHIGH and $\gamma = 84.4$ in CONLOW) are well above the predicted thresholds of 70; the estimated slope before the endogenously determined threshold is greater than the theoretical prediction of zero, while after the threshold the estimated slope is even greater than before the threshold, but less steep than the theoretical prediction for these period 2 wealth levels.[17]

Table 6 shows the estimated coefficients from the savings functions displayed in Figures 3 and 4. As the optimal period 1 savings functions in the CON treatments are convex functions (see

[17] Figure A.2 in the Appendix again replaces the scatterplot with LOWESS filters. Here, we observe that the filters mostly lie inside the confidence interval before they leave it for high wealth levels. In all four treatments, this happens at wealth levels to the right of the income range (of $[60, 80]$ in LOW and of $[35, 105]$ in HIGH). This might be explained by subjects with a preference for high savings: Those who saved a high amount in period 1 (and received a high income in period 2), subsequently also save a high amount in period 2. (This might anticipate results from a stimulus-response model presented in Section A.6 in the Appendix.)

Equation 9 on p. 10), the predictions shown here are also results from linear panel regressions of optimal savings decisions on income.[18] We observe that, in period 1 of all treatments, the estimated intercept is significantly greater than predicted and the estimated savings function slope coefficient is significantly smaller than predicted. In period 2, for the Unc treatments, the estimates are similar: estimated intercepts are significantly greater than predicted and estimated savings slope coefficients are smaller than predicted (though only the UncHigh slope coefficient is significantly less than predicted, and only at the 10%-level). For the period 2 savings coefficients in the Con treatments, we only have theoretical predictions for the coefficients *above* the threshold (they lie completely in the region where positive savings are predicted; the coefficient of savings before the threshold consists of a region where no savings *or* some positive savings are predicted). We observe that in the estimated post-threshold region, these savings slope coefficients are significantly smaller than predicted.[19,20] Summarizing, we generally find that, relative to theoretical predictions, the estimated savings functions have a higher intercept and a smaller slope, resulting in greater than predicted savings, particularly at low income or wealth levels. In the constrained case, we find that the endogenously determined threshold for a break in the period 2 savings function occurs at a much higher level of period 2 wealth than is predicted by the model, which indicates a greater precautionary motive than the theory predicts.

[18]The convexity of the theoretical functions is small and the linear approximations provide a very good fit with $R^2 > 0.999$ in both the High and Low treatment.

[19]We also compare the coefficients before and after the threshold in the Con treatments. The coefficients after the threshold are significantly larger than the ones before the threshold in both treatments ($p = 0.001$ in ConHigh; $p = 0.0018$ in ConLow).

[20]Table A.4 in the Appendix shows estimated savings functions for period 2 savings in the Con treatments, separately for the region before and after the theoretical threshold. We can observe the following: before the threshold, the coefficients on wealth are smaller than in the threshold panel regressions in Table 6 (though significantly different from zero). The intercepts before the threshold are not significantly different from zero. After the threshold, the coefficients on wealth are larger than the ones in Table 6. They are significantly larger than the predictions. The intercepts are larger than predicted.

**Figure 3:** Savings decisions (jittered), estimated linear savings functions, and savings predictions in period 1

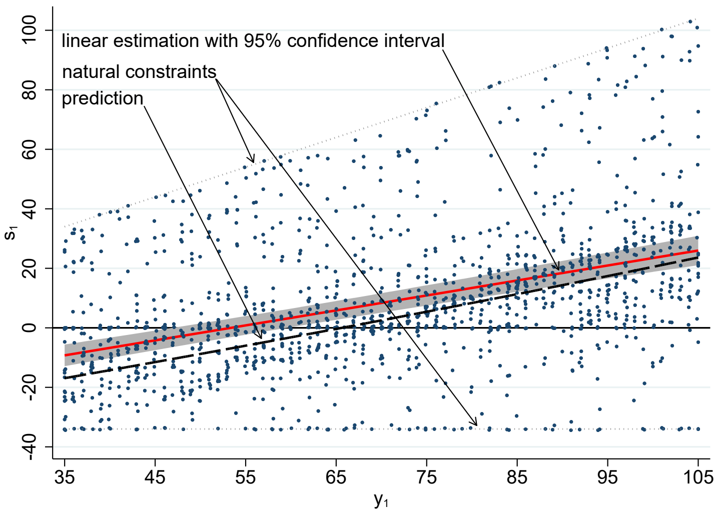


**(a)** Savings in period 1 in UncHigh

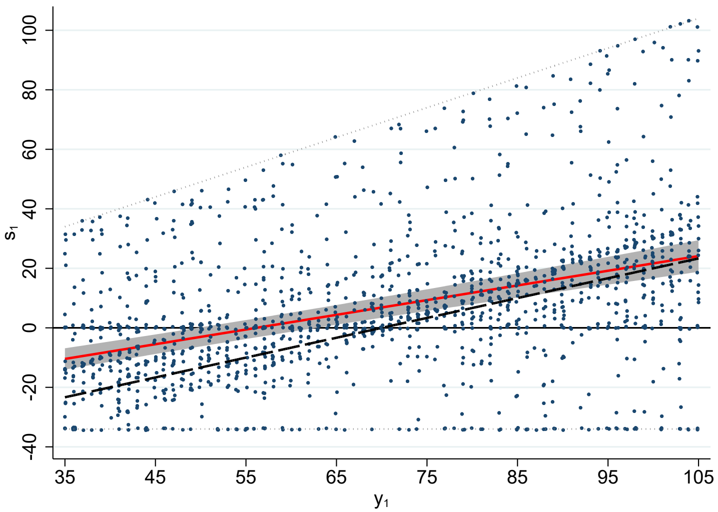


**(b)** Savings in period 1 in ConHigh

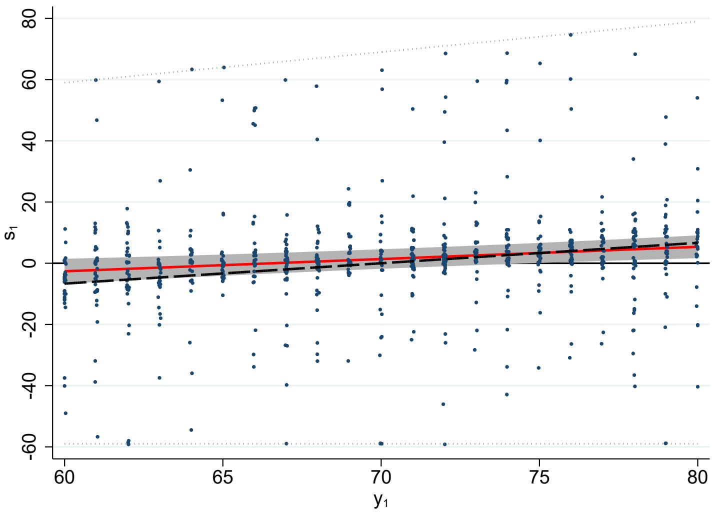


**(c)** Savings in period 1 in UncLow

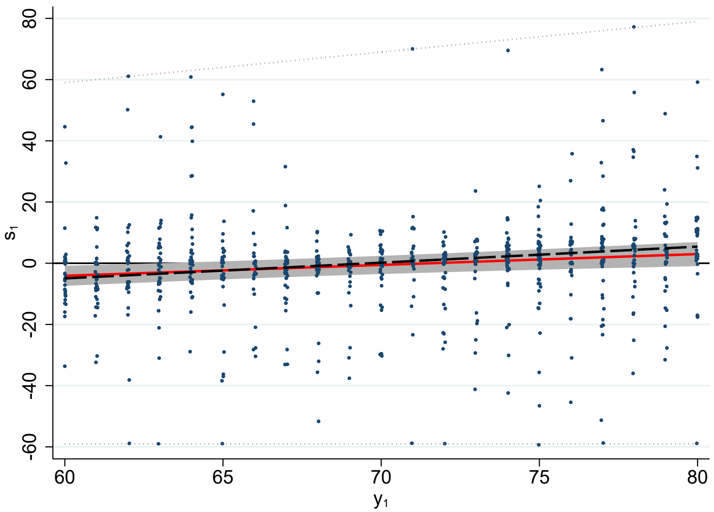


**(d)** Savings in period 1 in ConLow

**Figure 4:** Savings decisions (jittered), linear (threshold) savings functions, and savings predictions in period 2

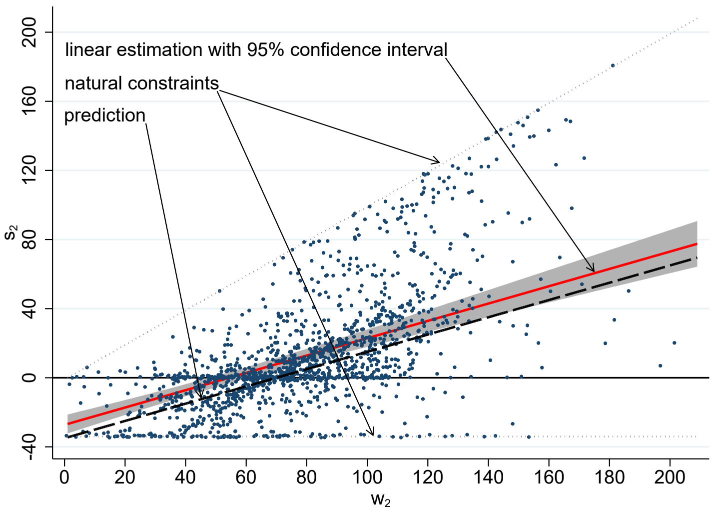


**(a)** Savings in period 2 in UncHigh

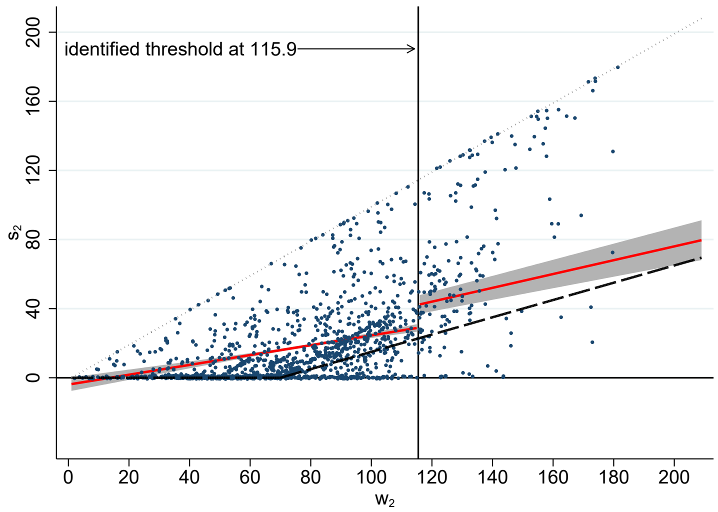


**(b)** Savings in period 2 in ConHigh

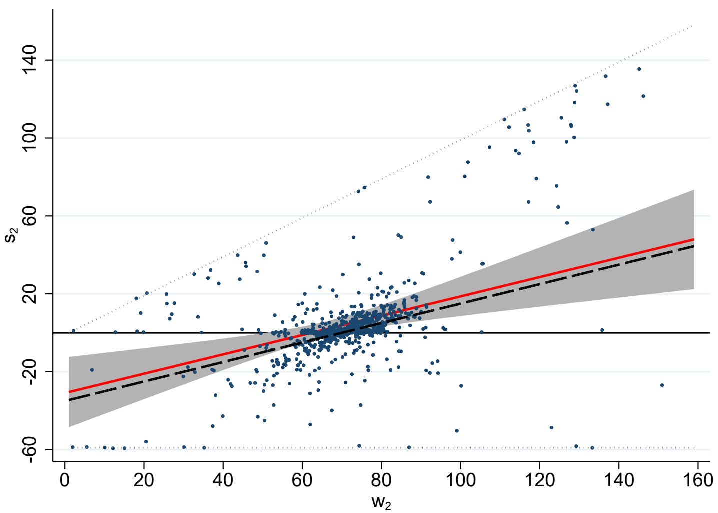


**(c)** Savings in period 2 in UncLow

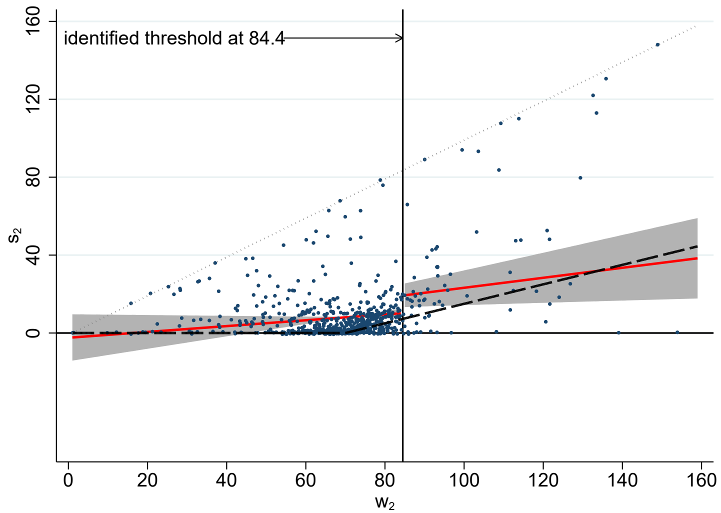


**(d)** Savings in period 2 in ConLow

**Table 6:** Estimated linear (threshold) savings functions

| | UncHigh | | | | ConHigh | | | | UncLow | | | | ConLow | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Prediction | | Estimation | | Prediction | | Estimation | | Prediction | | Estimation | | Prediction | | Estimation |
| $s_1$ | $y_1$ | 0.667 | >>> | 0.505*** | $y_1$ | 0.579 | >>> | 0.495*** | $y_1$ | 0.667 | >>> | 0.392*** | $y_1$ | 0.519 | >> | 0.358*** |
| | | | | (0.021) | | | | (0.019) | | | | (0.084) | | | | (0.074) |
| | constant | -46.667 | <<< | -26.974*** | constant | -37.685 | <<< | -27.793*** | constant | -46.667 | <<< | -26.083*** | constant | -36.118 | << | -25.633*** |
| | | | | (1.534) | | | | (1.385) | | | | (5.943) | | | | (5.230) |
| | #obs. | | | 1,500 | #obs. | | | 1,500 | #obs. | | | 750 | #obs. | | | 750 |
| | #clusters | | | 100 | #clusters | | | 100 | #clusters | | | 50 | #clusters | | | 50 |
| | $R^2$ (within) | | | 0.293 | $R^2$ (within) | | | 0.328 | $R^2$ (within) | | | 0.030 | $R^2$ (within) | | | 0.032 |
| | $R^2$ (between) | | | 0.005 | $R^2$ (between) | | | 0.003 | $R^2$ (between) | | | 0.043 | $R^2$ (between) | | | 0.002 |
| | $R^2$ (overall) | | | 0.158 | $R^2$ (overall) | | | 0.162 | $R^2$ (overall) | | | 0.026 | $R^2$ (overall) | | | 0.015 |
| $s_2$ | $w_2$ | 0.500 | > | 0.465*** | $w_2 < 115.9$ | | | 0.285*** | $w_2$ | 0.500 | | 0.446*** | $w_2 < 84.4$ | | | 0.150*** |
| | | | | (0.018) | | | | (0.017) | | | | (0.038) | | | | (0.033) |
| | | | | | $w_2 \geq 115.9$ | 0.500 | >>> | 0.400*** | | | | | $w_2 \geq 84.4$ | 0.500 | >> | 0.257*** |
| | | | | | | | | (0.014) | | | | | | | | (0.026) |
| | constant | -35.000 | <<< | -24.385*** | constant | | | -3.948*** | constant | -35.000 | <<< | -27.350*** | constant | | | -2.434 |
| | | | | (1.495) | | | | (1.242) | | | | (2.734) | | | | (2.210) |
| | #obs. | | | 1,500 | #obs. | | | 1,500 | #obs. | | | 750 | #obs. | | | 750 |
| | #clusters | | | 100 | #clusters | | | 100 | #clusters | | | 50 | #clusters | | | 50 |
| | $R^2$ (within) | | | 0.324 | $R^2$ (within) | | | 0.378 | $R^2$ (within) | | | 0.168 | $R^2$ (within) | | | 0.178 |
| | $R^2$ (between) | | | 0.713 | $R^2$ (between) | | | 0.767 | $R^2$ (between) | | | 0.608 | $R^2$ (between) | | | 0.388 |
| | $R^2$ (overall) | | | 0.421 | $R^2$ (overall) | | | 0.497 | $R^2$ (overall) | | | 0.337 | $R^2$ (overall) | | | 0.259 |

Note: Estimations are based on fixed-effect panel regressions; in the Con treatments in period 2, threshold fixed-effect panel regressions using the Stata package `xthreg` with 5% trimming (Wang, 2015). Standard errors in parentheses. ***, **, and * show differences from zero at the 1%-, 5%-, and 10%-level, respectively. >>>, >>, and > (<<<, <<, <) display the results from two-sided Wald tests where the estimated coefficient is significantly smaller (larger) than the prediction (significant at the 1%-, 5%-, and 10%-level, respectively). Predictions in period 1 of the Con treatments are derived from linear panel regressions of optimal savings on income (all $R^2 > 0.999$ of these regressions).

# 6 Behavioral explanations

In this section, we consider three behavioral explanations as to why we reject Hypotheses H1a and H1b, that first-period savings are not higher in the constrained treatment relative to the unconstrained treatment. The first explanation we consider is that subjects are *debt-averse*; they are more likely to avoid borrowing than they are to avoid saving when borrowing/saving is optimal. Consequently, the presence or absence of a constraint on borrowing, as is varied the Con and Unc treatments, does not matter as much for their decision-making. The second explanation is that there is heterogeneity in subjects' cognitive abilities; some subjects are able to look ahead to optimally respond to liquidity constraints while others are not, and we ask whether such differences in cognitive abilities can explain departures from theoretical predictions. Finally, we consider whether repeated experience with the three-period model, or "learning" plays any role. As there may be some novelty to the environment subjects face in the experiment, it may take subjects some time to learn the optimal response to the Con and Unc environments. With experience, their behavior may be more in line with theoretical predictions.

## 6.1 Debt aversion

To examine debt aversion as an explanation for observed behavior, we first test whether subjects' saving/borrowing decisions have the predicted sign, relative to the optimal solution. Second, we derive a behavioral model employing an extreme version of debt aversion, where subjects avoid borrowing even in period 1, though they were always free to borrow in period 1 of our experiment, and we test that model against the other models derived in Section 2.

### 6.1.1 Evidence from a simple binary model

A simple way to detect whether subjects exhibit debt aversion is by a dichotomization: we introduce the variable 'binary optimal' for each savings decision that each subject makes. Binary optimal takes the value 1 if the decision is in line with theory (the subject saves if it optimal to save or borrows if it is optimal to borrow or does neither if neither saving nor borrowing is optimal) and it takes the value 0 otherwise.[21] With this variable, we only consider whether the *sign* of the savings decision is correct and ignore the magnitude of the deviation from the prediction. That is, in the language of macro and labor economists, we test the *extensive margin* prediction (asking *whether* subjects save or borrow) and *not* the intensive margin prediction (*how much* they save or borrow). We address the latter question later on in our assessment of model fits.

Table 7 shows the shares of binary optimal decisions for the range of income (in period 1) or wealth (in period 2) where (i) subjects should borrow and (ii) subjects should save. (We do not show the second period of the CON treatments as all decisions for that period are binary optimal by design.) We observe that, in the borrowing range, between 51.7 and 61.3% of all decisions have the predicted sign. By contrast, in the saving range, between 68.5 and 85.3% of all decisions are classified as binary optimal. The differences between the shares in the two ranges are positive for all treatments and periods and are significantly different from zero in all but period 1 of the CONLOW treatment. This overall finding is consistent with *debt aversion*: subjects take more binary optimal decisions when the income shock realization requires them to save than when it requires them to borrow.

Next, we consider whether and how the binary optimal classification changes with the income/wealth values in all relevant treatments and periods. Figure 5 displays point estimates of the frequency of binary optimal decisions (along with 95% confidence intervals) for all integer values of income and wealth together with separate linear regression lines for the borrowing and saving range.[22] We observe the same pattern in most treatments and periods (except CONLOW

[21]Here, again, binary optimal for period 2 decisions is *conditionally* binary optimal as period 2 decisions also depend on the previous period's savings decision.

[22]In contrast to the previous figures, we do not show the whole wealth range in Figures 5e and 5f. As savings decisions are very dispersed in the range outside the income range, many point estimates for the frequency of binary optimal decisions and confidence intervals in those ranges only display one observation and are thus either 0 or 1. Extending the figures to the complete wealth range makes them less informative. (Note that the fitted

**Table 7:** Binary optimal behavior with tests on aggregate and local differences

| Treatment | Period | Share (conditionally) binary optimal Should borrow | Should save | Aggregate difference | Local difference |
|---|---|---|---|---|---|
| UncHigh | 1 | 52.9% | 83.8% | 30.9pp ($p < 0.001$) | 35.4pp ($p = 0.001$) [-12.6; +12.6] |
| ConHigh | 1 | 57.9% | 81.0% | 23.1pp ($p < 0.001$) | 30.2pp ($p = 0.013$) [-11.1; +11.1] |
| UncLow | 1 | 51.7% | 75.9% | 24.1pp ($p = 0.002$) | 34.0pp ($p = 0.037$) [-4.1; +4.1] |
| ConLow | 1 | 53.9% | 68.5% | 14.6pp ($p = 0.101$) | 6.9pp ($p = 0.637$) [-3.9; +3.9] |
| UncHigh | 2 | 61.3% | 85.3% | 24.0pp ($p < 0.001$) | 39.7pp ($p < 0.001$) [-16.6; +16.6] |
| UncLow | 2 | 55.1% | 84.8% | 29.7pp ($p < 0.001$) | 23.4pp ($p = 0.043$) [-11.9; +11.9] |

Note: Tests on aggregate differences use two-sided t-tests with standard errors clustered at the subject-level. Tests on local differences use a sharp local linear regression discontinuity estimation with triangular kernel and standard errors clustered at the subject-level (using Stata package `rdrobust` by Calonico et al 2017). Numbers in brackets give the symmetric MSE-optimal bandwidth calculated around the cutoff.

in period 1 in Figure 5d and UncLow in period 2 in Figure 5f): the fitted line in the borrowing range decreases with income/wealth until the cutoff (where the predicted sign of savings changes), and after the cutoff in the saving range, the fitted line 'jumps' to a higher level and further increases with income/wealth. In the borrowing range, subjects take on average less accurate decisions when the income shock is closer to the cutoff; when the income shock requires saving, subjects also are on average less correct when the shock is closer to the cutoff, but they seem to take more correct decisions than in the borrowing range.

lines use the observations outside the income range. However, it does not change the fitted lines qualitatively if we base them only on the observations in the income range.)

**Figure 5:** Share of (conditionally) binary optimal decisions

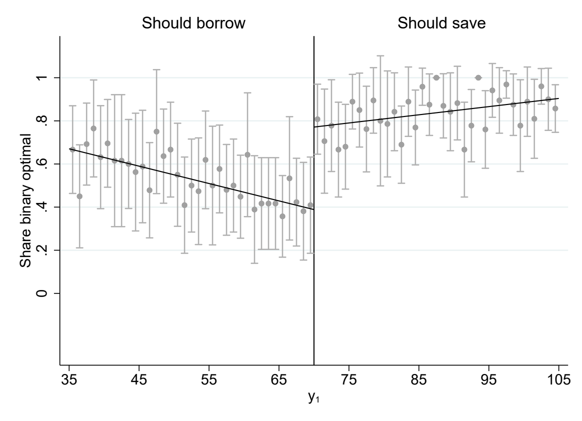


**(a)** Period 1 in UncHigh

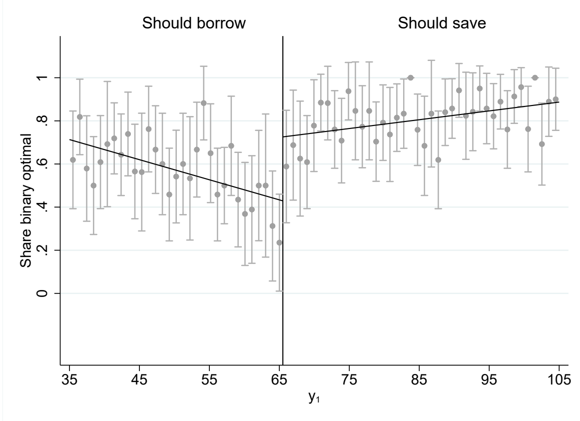


**(b)** Period 1 in ConHigh

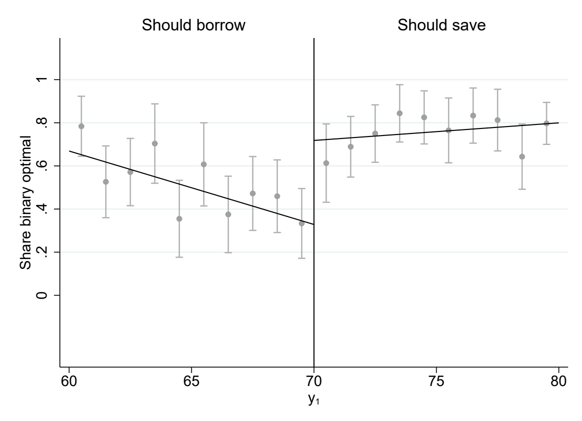


**(c)** Period 1 in UncLow

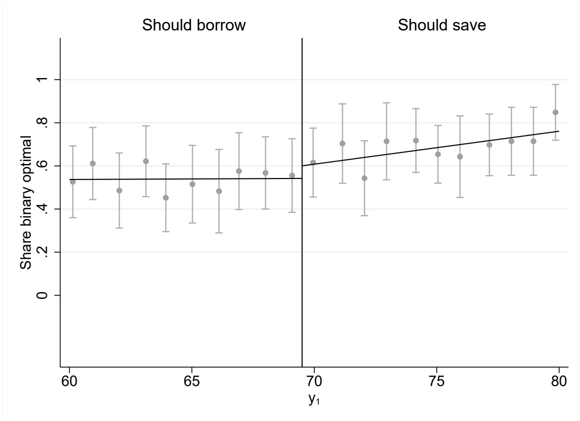


**(d)** Period 1 in ConLow

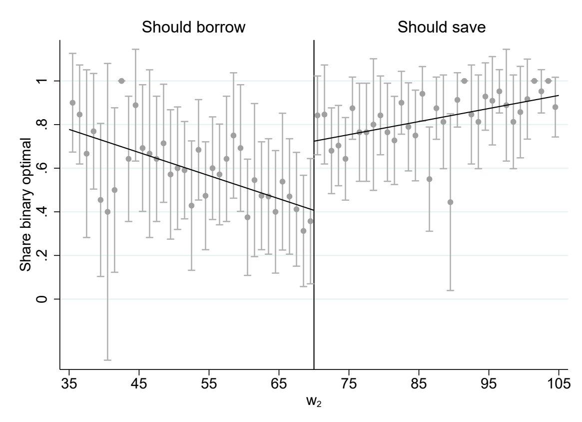


**(e)** Period 2 in UncHigh

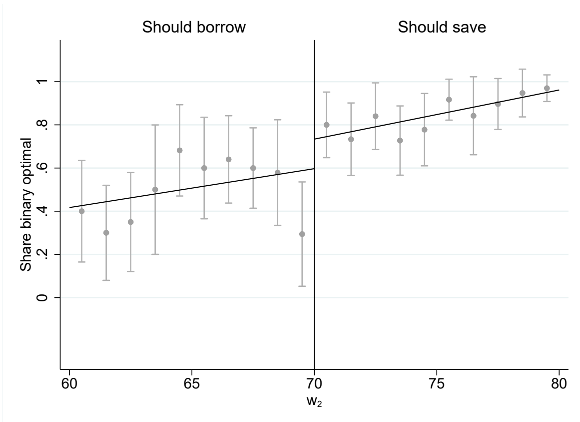


**(f)** Period 2 in UncLow

Note: Point estimates for the wealth ranges in period 2 are calculated for rounded values (as savings decisions in period 1 can have a decimal place). The confidence intervals (at the 95%-level) are based on standard errors clustered at the subject-level. The vertical lines divide the income/wealth range into borrowing/saving regions and are based on theoretical predictions. Graphs created with the Stata package `rdplot` by Calonico et al (2017).

Can we quantify these differences locally around the cutoff? How does the probability to behave in the binary optimal manner change around the cutoff? To answer these questions, we make use of a regression discontinuity analysis[23] of the data where the assignment to the borrowing or saving area is (partly) random and depends deterministically (due to the savings function) on the income realization.[24] Hence, we have a sharp discontinuity in the assignment to either the borrowing or saving range as a function of income/wealth. To improve the power of the test, we do not rely on data just to the left or to the right of the cutoff but instead our estimation makes use of a triangular kernel, which gives decreasing but non-zero weight to observations further from the cutoff.

We estimate the average change of the probability that subjects behave in the binary optimal manner in the saving range as compared with the borrowing range near the cutoff. The results are shown in the final column of Table 7. For all but two treatments/periods (ConLow in period 1 and UncLow in period 2), we observe that the local differences around the cutoff are even higher than the aggregate differences and significantly different from zero.[25] Our interpretation is that the probability to take a binary optimal saving decision is higher than a binary optimal borrowing decision and that this probability difference strongly increases around the cutoff. This bias is in line with debt aversion.[26]

[23]See Lee and Lemieux (2010) for an introduction to this quasi-experimental econometric method.

[24]In the first period, income is completely random across the whole income range; in the second period, subjects, due to their savings from the first period, can partly affect whether they are in the borrowing or in the saving range.

[25]With our approach, we are on the conservative side. Higher-order polynomials estimate greater differences than the linear regressions. (See Gelman and Imbens 2019 for arguments against the use of higher-order polynomials.)

[26]In the Appendix, we provide further evidence supporting debt aversion. Table A.5 shows the ratios of observed savings decisions to the number of predicted savings decisions for all relevant treatments, Table A.6 shows the ratios of observed borrowing decisions to the number of predicted borrowing decisions, both calculated per subject. We observe the following: (i) In almost all treatments, the subjects took on average more savings decisions than predicted. (ii) In all treatments, the subjects took on average fewer borrowing decisions than predicted. (iii) In all treatments, the mean ratio of savings decisions is higher than the borrowing decisions. (iv) In all treatments, the share of subjects who never borrow is higher than the share of subjects who never save.

### 6.1.2 Evidence from model comparisons

In this section, we compare which of the three models best describes observed behavior: the unconstrained model, the liquidity constrained model, or a behavioral, "debt aversion" model. For these comparisons, we derive the savings functions of a hypothetical, perfectly debt-averse individual. The debt-averse individual always seeks to avoid any debt, even if borrowing is possible. She behaves *as if* she faces a no-borrowing liquidity constraint (no matter whether or not a constraint is present). Thus, the savings function in period 2 for a debt-averse individual is the same as for the constrained individual:

$$s_2^{\mathrm{DA}} = s_2^{\textsc{Con}}. \tag{12}$$

Using the same backward induction procedure applied in Section 2 we derive the savings function in period 1 for the debt-averse individual that in addition to $s_2 \geq 0$ also sets $s_1 \geq 0$:

$$s_1^{\mathrm{DA}} = s_1^*(y_1) = \begin{cases} \pm 2\sqrt{\mu}\sqrt{7\mu - 2y_{max} - y_1} + 5\mu - y_{max} & \text{if } \ y_1 \geq \mu \\ 0 & \text{otherwise.} \end{cases} \tag{13}$$

This function is convex: it has a kink at the point where income is equal to mean income. Below this kink, savings are zero (because of the individual's debt aversion in period 1). Above this kink, savings are equal to the constrained case (Equation 9). This part of the savings function is convex because of the individual's anticipation of her debt aversion in period 2 (here, as noted in Section 2, only the negative solution makes economic sense).

We now turn to the question of which of the three models, unconstrained, constrained, or debt-averse, explains our experimental data best. Table 8 shows the root-mean-square errors (RMSE) of the Unc, Con and DA model for the first period (these model predictions are conditioned on realized income $y_1$) and the Unc and Con model for the second period (they are

**Table 8:** Model comparisons using root-mean-square errors

| | Model | Treatment | | | |
|---|---|---|---|---|---|
| | | UncHigh | ConHigh | UncLow | ConLow |
| $s_1$ | Unc | 24.34 | 23.97 | 17.42 | 15.91 |
| | Con | 23.30 | 23.09 | <u>17.37</u> | <u>15.82</u> |
| | DA | <u>22.83</u> | <u>22.98</u> | 17.43 | 15.84 |
| $s_2$ | Unc | <u>27.67</u> | 28.32 | <u>21.17</u> | 18.00 |
| | Con | 28.31 | <u>26.25</u> | 21.45 | <u>15.57</u> |

Note: RMSE $= \sqrt{\frac{1}{N}\sum_{i=1}^{N}(\text{prediction}_i - \text{observation}_i)^2}$. Model with best fit is underlined (perfect score=0).

conditioned on previous savings decision $s_1$ and realized income $y_2$) when compared with the individual-level data from all four treatments. For first-period savings in the High treatment, the DA model provides the best fit to the data (has the lowest RMSE). For first-period savings in the Low treatments, the Con model provides the best fit to the data. For second-period savings, the predicted model, Unc or Con, describes behavior best according to whether the liquidity constraint was not binding or was binding across all four treatments (we can only speculate why behavior in period 2 is closer to predictions—observed oversaving in period 1 might have reduced low wealth observations in period 2 where subjects deviate more from predictions).[27] This confirms our results from the previous section: especially in the High treatments in the first period, the subjects behave in a debt-averse manner. In the Low treatments, subjects are less extreme than the debt averse model would predict, but instead act in a manner closer to the constrained model predictions (as already shown with the good performance of predictions in the ConLow treatment). More generally, notice that the RMSEs across all models are always lower (higher) for the Low (High) variance treatments, suggesting that greater income variance leads to greater mistakes, relative to optimal predictions.

[27] We consider the RMSE an ideal measure to compare observed behavior and model predictions as it ignores the direction of the deviation and punishes deviations disproportionately, which is adequate given the concave utility function. However, we also present results using other verification measures in the Appendix: Table A.7 reports mean absolute errors (in contrast to the RMSE, mean absolute errors punish deviations proportionately), Table A.8 shows mean errors (mean errors take into account the individual deviation's direction but averaging them can cancel out positive and negative deviations; punishment is proportional), and Table A.9 reports a bias measure (which compares the average prediction magnitude with the average observed magnitude). In general, these comparisons also prefer the more constrained models, confirming the results from the RMSE comparisons.

## 6.2 Cognitive abilities and risk preferences

In this section, we examine whether the measures from Part 3 of our experiment, the ToH game, and the lottery task, and the items from our ex-post questionnaire can explain the deviations that we observe from optimal savings predictions.

Table 9 summarizes the subjects' characteristics.[28] The sample is close to being gender-balanced. The subjects' CRT score is more dispersed than is their GPA. Almost all subjects solved the ToH task, almost all of them using only few moves (33 out of 140 subjects solved the task in seven moves, 23 in eight moves, 24 in nine moves). We use a subjects' number of less risky lottery choices in the binary lottery task as a measure of their risk aversion (36 subjects switched more than once or chose the risky option first and then switched to the less risky option; we exclude their observations from the following analysis). The mean and median number of less risky lottery choices of 6 (among those whose started out making such choices) indicates that subjects are, on average, slightly risk-averse.[29]

Table 10 shows results from panel regressions where we try to understand deviations of observed behavior from theoretical predictions in period 1 on the basis of individual characteristics (our primary concern is the deviation from predictions in this first period). Since we expected that GPA, CRT score, and ToH performance would predict a lower deviation, no matter the direction, we use as our dependent variable the absolute deviation. However, we also consider either positive or negative deviations separately, to check whether the cognitive measures' coefficients have the predicted sign and to test if risk aversion increases savings. We also look at simple deviations (to test if risk aversion increases savings) in period 1. These regressions use pooled data on deviations from all treatments; the explanatory variables are the individual

[28]Table A.10 in the Appendix displays a correlation matrix of the variables shown in Table 9.

[29]If subjects are risk neutral, then they will choose the less risky option A for the first four choices and then switch to the more risky option B for the last six choices, corresponding to the measure we use: 4 less risky or safe choices. Subjects who switch later than choice 4 from option A to option B are risk averse making 5 to 10 less risky choices, while those who switch earlier than choice 4 from option A to option B are risk seeking, making 0 to 3 less risky choices.

**Table 9:** Subjects' characteristics

| | N | Mean | Minimum | Median | Maximum |
|---|---|---|---|---|---|
| Share females | 150 | 54.7 % | | | |
| | | (49.9 ) | | | |
| Age | 150 | 20.3 | 18 | 20 | 28 |
| | | (1.5) | | | |
| CRT score | 150 | 2.0 | 0 | 2 | 4 |
| | | (1.3) | | | |
| GPA | 150 | 3.2 | 0 | 3.2 | 4 |
| | | (0.6) | | | |
| Risk aversion[a] | 114 | 6.0 | 0 | 6 | 10 |
| | | (2.7) | | | |
| Share who solved ToH | 150 | 94.0 % | | | |
| | | (23.9 ) | | | |
| Moves if ToH solved[b] | 140 | 10.8 | 7 | 9 | 30 |
| | | (4.7) | | | |

Note: Standard deviations in parentheses. CRT=cognitive reflection test. GPA=grade point average. ToH=Tower of Hanoi.
[a] Risk aversion gives the unique switching point when the subject switches to the risky lottery (the uniqueness criterion reduces observations; subjects switching more than once and switching from risk-seeking to risk-averse are excluded).
[b] Only 140 out of 150 subjects solved the ToH.

characteristics reported on in Table 9.

**Table 10:** Determinants of pooled deviations in period 1

| | \|Deviation\| | Deviation>0 | Deviation<0 | Deviation |
|---|---|---|---|---|
| Female==1 | 1.764 | 1.492 | -2.231 | -3.097 |
| | (2.132) | (2.474) | (1.627) | (2.459) |
| Age | -0.427 | -0.432 | 0.281 | -0.035 |
| | (0.669) | (0.783) | (0.431) | (1.056) |
| CRT score | -3.421*** | -3.106*** | 2.924*** | 0.850 |
| | (0.671) | (0.799) | (0.510) | (0.896) |
| GPA | 0.506 | 0.864 | 0.060 | 3.227 |
| | (1.371) | (1.501) | (1.116) | (1.841) |
| ToH solved==1 | 2.870 | 3.360 | -5.811 | -13.586 |
| | (9.775) | (12.802) | (3.654) | (12.330) |
| ToH solved==1 * Moves | 0.098 | 0.145 | -0.226 | 0.067 |
| | (0.171) | (0.215) | (0.265) | (0.234) |
| Risk aversion | 0.393 | 0.498 | 0.147 | 1.001** |
| | (0.348) | (0.421) | (0.240) | (0.448) |
| constant | 19.961 | 16.777 | -13.678 | -0.219 |
| | (17.298) | (20.993) | (10.587) | (25.414) |
| #obs. | 3,420 | 1,978 | 1,434 | 3,420 |
| #clusters | 114 | 111 | 110 | 114 |
| $R^2$ (within) | — | — | — | — |
| $R^2$ (between) | 0.186 | 0.140 | 0.188 | 0.091 |
| $R^2$ (overall) | 0.103 | 0.102 | 0.154 | 0.039 |

Note: Estimations based on panel regressions. Standard errors in parentheses. ***, **, and * show difference from zero (based on two-sided t-tests) at the 1%-, 5%-, and 10%-level, respectively.

Subjects' performance in the CRT explains these period 1 deviations best. The higher the subjects' ability to override system 1 and employ system 2 thinking (the higher the CRT score), the lower are his or her deviations in all specifications but the last one (where we do not expect that it should have an impact). Variations in GPA and ToH scores (and the number of moves if ToH was solved) do not explain deviations, most likely because almost all subjects solved the ToH task and because there is little variance in GPAs across subjects. The coefficient on the risk aversion measure has the expected sign in all regressions but is only statistically significant if we use all observations together with the direction of the deviation.[30] Our findings here are in line with Ballinger et al 2011, who also found that cognitive measures, as opposed to other demographic factors were the best predictors of behavior in their savings experiment.

### 6.3 Learning over time

Many studies that examine savings in laboratory experiments (e.g., Ballinger et al 2003, Brown et al 2009, and Meissner 2016) observe that subjects improve in their saving behavior when facing savings decisions repeatedly. In this section, we test if this is also the case in our experiment.

Table 11 shows the results from fixed-effect panel data regressions where we regress deviations and absolute deviations of observed behavior from predictions on the number of rounds, separately for the four different treatments and the two periods.[31] Almost all constants in the deviation regressions are positive (pointing to the oversaving already reported in Section 5.1) and significantly different from zero (all but period 1 in ConLow, where observed behavior is very close to predictions, see Table 5 and the results in Section 5.2). The coefficients on the round number are all negative, and thus, reduce the deviation in all treatments/periods (but period 2 in ConHigh). However, the round coefficients in the deviation regression are not significantly different from zero. When we consider absolute deviations as the dependent variable, we observe that all the coefficients on round numbers are negative and significantly different

[30]Table A.11 in the Appendix shows the same regressions as Table 10 for period 2 deviations (conditional on $s_1$ and $y_2$). The results confirm our findings for period 1 (the coefficients of risk aversion are at least significant at the 10%-level in all specifications).

[31]We also test a basic stimulus-response learning model in Section A.6 in the Appendix. To put it shortly: The coefficients of our learning model point into the predicted direction, though many coefficients are not significantly different from zero; the results of this section are confirmed by the stimulus-response model.

from zero. Thus we find evidence that over time, subjects reduce their mistakes; however, the learning over the 15 rounds per treatment of the experiment reduces the oversaving constant by only about 15 to 25%, depending on the treatment/period.[32]

[32] Allen and Carroll (2001) examine individual learning of optimal behavior in the buffer stock savings model. While the optimal solution is a non-linear consumption function, the authors simulate learning by experience with a linear approximation of the consumption function and find that learning takes roughly a million simulation periods. Compared with this result, our subjects learn quickly.

**Table 11:** Learning time as a determinant of (conditional) deviations from optimal behavior

| | | UncHigh | | | ConHigh | | | UncLow | | | ConLow | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Deviation | \|Deviation\| | | Deviation | \|Deviation\| | | Deviation | \|Deviation\| | | Deviation | \|Deviation\| |
| Period 1 | round | -0.152 | -0.211*** | round | -0.134 | -0.183*** | round | -0.117 | -0.189** | round | -0.014 | -0.128* |
| | | (0.097) | (0.068) | | (0.089) | (0.066) | | (0.115) | (0.081) | | (0.102) | (0.074) |
| | constant | 10.639*** | 20.198*** | constant | 6.085*** | 18.722*** | constant | 3.159* | 12.827*** | constant | -0.543 | 11.271*** |
| | | (1.554) | (1.101) | | (1.439) | (1.058) | | (1.852) | (1.306) | | (1.644) | (1.189) |
| | #obs. | 1,500 | 1,500 | #obs. | 1,500 | 1,500 | #obs. | 750 | 750 | #obs. | 750 | 750 |
| | #clusters | 100 | 100 | #clusters | 100 | 100 | #clusters | 50 | 50 | #clusters | 50 | 50 |
| | $R^2$ (within) | 0.002 | 0.007 | $R^2$ (within) | 0.002 | 0.006 | $R^2$ (within) | 0.002 | 0.007 | $R^2$ (within) | 0.000 | 0.004 |
| | $R^2$ (between) | 0.050 | 0.063 | $R^2$ (between) | 0.000 | 0.004 | $R^2$ (between) | 0.013 | 0.005 | $R^2$ (between) | 0.001 | 0.002 |
| | $R^2$ (overall) | 0.024 | 0.038 | $R^2$ (overall) | 0.000 | 0.000 | $R^2$ (overall) | 0.006 | 0.006 | $R^2$ (overall) | 0.000 | 0.003 |
| Period 2 | round | -0.041 | -0.251*** | round | -0.337*** | -0.294*** | round | -0.154 | -0.252*** | round | -0.109 | -0.153** |
| | | (0.110) | (0.081) | | (0.081) | (0.072) | | (0.124) | (0.092) | | (0.076) | (0.067) |
| | constant | 8.477*** | 22.011*** | constant | 16.739*** | 18.799*** | constant | 6.191*** | 15.072*** | constant | 7.999*** | 9.791*** |
| | | (1.778) | (1.307) | | (1.306) | (1.166) | | (2.003) | (1.482) | | (1.225) | (1.081) |
| | #obs. | 1,500 | 1,500 | #obs. | 1,500 | 1,500 | #obs. | 750 | 750 | #obs. | 750 | 750 |
| | #clusters | 100 | 100 | #clusters | 100 | 100 | #clusters | 50 | 50 | #clusters | 50 | 50 |
| | $R^2$ (within) | 0.000 | 0.007 | $R^2$ (within) | 0.012 | 0.012 | $R^2$ (within) | 0.002 | 0.011 | $R^2$ (within) | 0.003 | 0.007 |
| | $R^2$ (between) | 0.058 | 0.051 | $R^2$ (between) | 0.001 | 0.000 | $R^2$ (between) | 0.002 | 0.019 | $R^2$ (between) | 0.028 | 0.030 |
| | $R^2$ (overall) | 0.025 | 0.031 | $R^2$ (overall) | 0.002 | 0.002 | $R^2$ (overall) | 0.002 | 0.016 | $R^2$ (overall) | 0.017 | 0.022 |

Note: Estimations are based on fixed-effect panel regressions. Standard errors in parentheses. ***, **, and * show differences from zero at the 1%-, 5%-, and 10%-level, respectively.

## 7 Summary and conclusion

In this paper, we study the importance of liquidity constraints for savings behavior in the context of buffer stock savings models. Liquidity constraints are thought to be empirically important factors in savings behavior and yet these constraints are often difficult for researchers to directly observe. Hence, we resort to a controlled laboratory test, where we can turn liquidity constraints on or off.

Specifically, we compare two, three-period models, one with a liquidity constraint in the second period, and another version of the same model without this liquidity constraint. We test the comparative statics predictions of these two models in a $2 \times 2$ experimental design where we also manipulate the variance in the known income process to create different degrees of income uncertainty. Our main research question is whether the *anticipation* of a known liquidity constraint increases period 1 savings relative to the environment without that constraint as predicted by theory.

We reject this hypothesis: savings with the liquidity constraint are not higher than without the constraint. Remarkably, however, we cannot reject any of the other hypotheses regarding the comparative statics predictions of our model: (i) savings in the second period of the constrained model are higher than in the treatment without the constraint (because the constraint prohibits borrowing in this period); (ii) savings in the second period are higher than in the first period; (iii) savings are higher when there is increased income uncertainty. Still, in almost all treatments and periods, we observe significant oversaving (except for the two LOW treatments in period 1).

In further analyses, we try to identify why the liquidity constraint result in the predicted anticipation effect. We find that a combination of debt aversion, heterogeneity in cognitive abilities, and/or learning can explain the deviations that we observe from the theoretical predictions. Especially in the experimental treatments with high income uncertainty, debt aversion seems to play an important role in explaining the lack of any anticipation effect; since subjects are already saving much more than they need to in period 1, the effect of the liquidity constraint in period 2 is greatly diminished.

Several empirical studies find that the buffer stock savings model has problems in explaining savings behavior (e.g., Jappelli et al 2008 and Fulford 2015). Despite our evidence regarding the anticipation of the liquidity constraint, we find that all of the other comparative statics of the model work as predicted in our controlled laboratory experiment. In contrast to Fulford (2015), we find that higher income uncertainty increases savings. In our experiment, unexpected expenditure shocks, and other motives for saving are not present. Thus, our experiment provides a very clean and direct test of the underlying theory.

The model underlying our experiment was not previously tested in a laboratory experiment. We think it is promising to adjust that model to address further research questions. For example, one could induce different levels of impatience and test whether subjects react in the predicted way by reallocating consumption between periods. We leave this and other interesting extensions to future research.

## References


Allen T, Carroll C (2001) Individual learning about consumption. Macroeconomic Dynamics 5(2):255–271

Ballinger TP, Palumbo MG, Wilcox NT (2003) Precautionary saving and social learning across generations: An experiment. The Economic Journal 113(490):920–947

Ballinger TP, Hudson E, Karkoviata L, Wilcox NT (2011) Saving behavior and cognitive abilities. Experimental Economics 14(3):349–374

Besley T (1995) Savings, credit and insurance. In: Behrman J, Srinivasan T (eds) Handbook of Development Economics, vol 3, Elsevier, pp 2123–2207

Boar C, Gorea D, Midrigan V (2017) Liquidity constraints in the U.S. housing market. National Bureau of Economic Research Working Paper 23345

Brown AL, Chua ZE, Camerer CF (2009) Learning and visceral temptation in dynamic saving experiments. The Quarterly Journal of Economics 124(1):197–231

Calonico S, Cattaneo MD, Farrell MH, Titiunik R (2017) rdrobust: Software for regression-discontinuity designs. The Stata Journal 17(2):372–404

Campbell JR, Hercowitz Z (2018) Liquidity constraints of the middle class. American Economic Journal: Economic Policy 11(3):130–155

Carroll C (2004) Theoretical foundations of buffer stock saving. National Bureau of Economic Research Working Paper 10867

Carroll CD (1997) Buffer-stock saving and the life cycle/permanent income hypothesis. The Quarterly Journal of Economics 112(1):1–55

Carroll CD, Kimball MS (2001) Liquidity constraints and precautionary saving. National Bureau of Economic Research Working Paper 8496

Carroll CD, Hall RE, Zeldes SP (1992) The buffer-stock theory of saving: Some macroeconomic evidence. Brookings Papers on Economic Activity 1992(2):61–156

Chen DL, Schonger M, Wickens C (2016) oTree—An open-source platform for laboratory, online, and field experiments. Journal of Behavioral and Experimental Finance 9:88–97

Cleveland WS (1979) Robust locally weighted regression and smoothing scatterplots. Journal of the American Statistical Association 74(368):829–836

Cohen J (1988) Statistical power analysis for the behavioral sciences, 2nd edn. Lawrence Erlbaum Associates

Deaton A (1991) Saving and liquidity constraints. Econometrica 59(5):1221–1248

Dhand NK, Khatkar MS (2014) Statulator: An online statistical calculator. URL `http://statulator.com/`, accessed 2019-07-21

Drichoutis AC, Lusk JL (2016) What can multiple price lists really tell us about risk preferences? Journal of Risk and Uncertainty 53(2-3):89–106

Duffy J (2016) Macroeconomics: A survey of laboratory research. In: Kagel JH, Roth AE (eds) Handbook of Experimental Economics, vol 2, Princeton University Press, pp 1–90

Frederick S (2005) Cognitive reflection and decision making. Journal of Economic Perspectives 19(4):25–42

Friedman M (1957) Theory of the consumption function. Princeton University Press

Fuchs-Schündeln N, Schündeln M (2005) Precautionary savings and self-selection: Evidence from the German reunification "experiment". The Quarterly Journal of Economics 120(3):1085–1120

Fulford SL (2015) The surprisingly low importance of income uncertainty for precaution. European Economic Review 79:151–171

Gelman A, Imbens G (2019) Why high-order polynomials should not be used in regression discontinuity designs. Journal of Business & Economic Statistics 37(3):447–456

Gneezy U, Rustichini A, Vostroknutov A (2010) Experience and insight in the Race game. Journal of Economic Behavior & Organization 75(2):144–155

Gross DB, Souleles NS (2002) Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data. The Quarterly Journal of Economics 117(1):149–185

Hansen BE (1999) Threshold effects in non-dynamic panels: Estimation, testing, and inference. Journal of Econometrics 93(2):345–368

Hinz AM, Klavžar S, Milutinović U, Petr C, Stewart I (2013) The Tower of Hanoi – Myths and Maths. Springer

Internal Revenue Service (2020) 2020 filing season statistics. `https://www.irs.gov/newsroom/filing-season-statistics-for-week-ending-may-1-2020`, accessed: 2020-09-01

Jappelli T (1990) Who is credit constrained in the US economy? The Quarterly Journal of Economics 105(1):219–234

Jappelli T, Pistaferri L (2017) The economics of consumption: theory and evidence. Oxford University Press

Jappelli T, Pistaferri L (2020) Permanent income shocks, target wealth, and the wealth gap. Centre for Economic Policy Research Discussion Paper 15145

Jappelli T, Padula M, Pistaferri L (2008) A direct test of the buffer-stock model of saving. Journal of the European Economic Association 6(6):1186–1210

Kimball MS (1990) Precautionary saving in the small and in the large. Econometrica 58(1):53–73

Lee DS, Lemieux T (2010) Regression discontinuity designs in economics. Journal of Economic Literature 48(2):281–355

Love D (2006) Buffer stock saving in retirement accounts. Journal of Monetary Economics 53(7):1473–1492

Ludvigson SC, Michaelides A (2001) Does buffer-stock saving explain the smoothness and excess sensitivity of consumption? The American Economic Review 91(3):631–647

Meghir C, Pistaferri L (2004) Income variance dynamics and heterogeneity. Econometrica 72(1):1–32

Meissner T (2016) Intertemporal consumption and debt aversion: An experimental study. Experimental Economics 19:281–298

Modigliani F, Brumberg R (1954) Utility analysis and the consumption function: an interpretation of cross-section data. In: Kurihara K (ed) Post-Keynesian Economics, Rutgers University Press, pp 388–436

Sahadi J (2015) Nearly 8 out of 10 U.S. taxpayers get refunds. CNN Money URL `https://money.cnn.com/2015/01/13/pf/taxes/taxpayer-refunds/index.html`, accessed: 2020-09-01

Schaner S (2015) Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings. American Economic Journal: Applied Economics 7(2):135–74

Toplak ME, West RF, Stanovich KE (2014) Assessing miserly information processing: An expansion of the cognitive reflection test. Thinking & Reasoning 20(2):147–168

Venti SF, Wise DA (1998) The cause of wealth dispersion at retirement: Choice or chance? The American Economic Review 88(2):185–191

Wang Q (2015) Fixed-effect panel threshold model using Stata. The Stata Journal 15(1):121–134

Zeldes SP (1989) Consumption and liquidity constraints: An empirical investigation. Journal of Political Economy 97(2):305–346

# Appendix

## A.1 Derivation of the savings functions

Here in the appendix, we present the derivation of the two solutions in the unconstrained and constrained model in a slightly more general context (set the initial endowment we introduce here $y = 0$ and the solutions are equal to the ones in Section 2 of the paper). Finally, we also derive the predictions of a model where liquidity is constrained in both periods (it describes behavior that is equivalent to the savings function of a debt-averse individual).
We use the utility function $u(c) = \ln(c)$ (with derivative $u'(c) = \frac{1}{c}$) and an i.i.d. discrete uniform distribution of period incomes $y_1, y_2, y_3 \sim U[y_{min}, y_{max}]$ with mean $\mu = \frac{y_{min}+y_{max}}{2}$ and standard deviation $\sigma = \sqrt{\frac{(y_{max}-y_{min}+1)^2-1}{12}}$.
In all models, the consumption stream with a certain initial endowment $y$ and an uncertain income in all three periods 1, 2 and 3 is given by:

- $c_1 = y + y_1 - s_1$,
- $c_2 = y_2 - s_2 + s_1$, and
- $c_3 = y_3 + s_2$.

The maximization problem of the individual in period 1 is given by:
$\max u(c_1) + \mathbb{E}_{t=1}\, u(c_2) + \mathbb{E}_{t=1}\, u(c_3)$
The only difference between the two models, Unc and Con, is the liquidity constraint: without a constraint, $s_2$ can take negative values (and borrowing is allowed) and with a liquidity constraint in period 2, $s_2 \geq 0$, the individual can only save in period 2 (borrowing in period 1 is possible, though).

### A.1.1 The unconstrained model

The three-period maximization problem in period 1 is given by:
$\max u(c_1) + \mathbb{E}_{t=1}\, u(c_2) + \mathbb{E}_{t=1}\, u(c_3)$
Substitute in the initial endowment, the income stream and the saving variables for consumption:
$\max_{s_1,s_2} u(y + y_1 - s_1) + \mathbb{E}_{t=1}\, u(y_2 - s_2 + s_1) + \mathbb{E}_{t=1}\, u(y_3 + s_2)$
We use backward induction and start with the problem in period 2:
$\max_{s_2} u(y_2 - s_2 + s_1) + \mathbb{E}_{t=2}\, u(y_3 + s_2)$
The FOC is given by $\frac{\partial}{\partial s_2} = 0$:
$-u'(y_2 - s_2 + s_1) + \mathbb{E}_{t=2}\, u'(y_3 + s_2) = 0$
We use the derivative of the utility function and apply the expectation operator to the uncertain income. Then, we solve for $s_2$:
$\frac{1}{y_2 - s_2 + s_1} = \frac{1}{\mu + s_2}$
$s_2^* = s_2^*(y_2, s_1) = \frac{1}{2}(y_2 + s_1 - \mu)$
The amount of saving/borrowing in period 2 depends linearly on the realization of income in period 2 plus previous period's saving (thus, total wealth). We now use $s_2^*$ in the initial periods' problem:
$\max_{s_1} u(y + y_1 - s_1) + \mathbb{E}_{t=1}\, u(y_2 - s_2^* + s_1) + \mathbb{E}_{t=1}\, u(y_3 + s_2^*)$
$\max_{s_1} u(y + y_1 - s_1) + \mathbb{E}_{t=1}\, u\left(y_2 - \frac{1}{2}(y_2 + s_1 - \mu) + s_1\right) + \mathbb{E}_{t=1}\, u\left(y_3 + \frac{1}{2}(y_2 + s_1 - \mu)\right)$
$\max_{s_1} u(y + y_1 - s_1) + \mathbb{E}_{t=1}\, u\left(\frac{1}{2}y_2 + \frac{1}{2}\mu + \frac{1}{2}s_1\right) + \mathbb{E}_{t=1}\, u\left(y_3 + \frac{1}{2}y_2 - \frac{1}{2}\mu + \frac{1}{2}s_1\right)$
The FOC is given by $\frac{\partial}{\partial s_1} = 0$:
$-u'(y + y_1 - s_1) + \frac{1}{2}\,\mathbb{E}_{t=1}\, u'\left(\frac{1}{2}y_2 + \frac{1}{2}\mu + \frac{1}{2}s_1\right) + \frac{1}{2}\,\mathbb{E}_{t=1}\, u'\left(y_3 + \frac{1}{2}y_2 - \frac{1}{2}\mu + \frac{1}{2}s_1\right) = 0$

Again, we use the derivative of the utility function and apply the expectation operator to the uncertain incomes. Then, we solve for $s_1$:

$\frac{-1}{y+y_1-s_1} + \frac{1}{2}\frac{1}{\mu+\frac{1}{2}s_1} + \frac{1}{2}\frac{1}{\mu+\frac{1}{2}s_1} = 0$

$\frac{1}{y+y_1-s_1} = \frac{1}{\mu+\frac{1}{2}s_1}$

$s_1^* = s_1^*(y_1) = \frac{2}{3}(y + y_1 - \mu)$

The amount of saving/borrowing in period 1 depends linearly on the income realization in that period. If current income is above (below) mean income, the individual saves (borrows).

### A.1.2 The model with a constraint in period 2

The problem is the same as before. Additionally, the individual cannot borrow in period 2: $s_2 \geq 0$. The maximization problem is thus not solvable with unconditional maximization:

$\max_{s_1,s_2} u(y + y_1 - s_1) + \mathbb{E}_{t=1} u(y_2 - s_2 + s_1) + \mathbb{E}_{t=1} u(y_3 + s_2) + \lambda\, \mathbb{E}_{t=1} s_2$

Again, the solution procedure starts with backward induction. In period 2, the problem is:

$\max_{s_2} u(y_2 - s_2 + s_1) + \mathbb{E}_{t=2} u(y_3 + s_2) + \lambda s_2$

The FOC in period 2 in case that the constraint does not bind is given by:

$-u'(y_2 - s_2 + s_1) + \mathbb{E}_{t=2} u'(y_3 + s_2) \geq 0$

Using the functional form of the utility function and applying expectations on the income leaves us with the solution for the case that the constraint does not bind:

$\frac{-1}{y_2-s_2+s_1} + \frac{1}{\mu+s_2} \geq 0$

$y_2 - 2s_2 + s_1 - \mu \geq 0$

This implies:

$0 \leq s_2 \leq \frac{1}{2}[y_2 + s_1 - \mu]$

Hence, the solution is convex. It has a kink at $y_2 + s_1 = \mu$, then it increases linearly in $y_2 + s_1$:

$s_2^* = s_2^*(y_2, s_1) = \begin{cases} \frac{1}{2}[y_2 + s_1 - \mu] & \text{if } y_2 + s_1 \geq \mu \\ 0 & \text{otherwise} \end{cases}$

Now, we use $s_2^*$ in the three-period problem:

$\max_{s_1} u(y + y_1 - s_1) + \mathbb{E}_{t=1} u(y_2 - s_2^* + s_1) + \mathbb{E}_{t=1} u(y_3 + s_2^*)$

The FOC is given by $\frac{\partial}{\partial s_1} = 0$:

$-u'(y + y_1 - s_1) + \mathbb{E}_{t=1} u'(y_2 - s_2^* + s_1) = 0$

We use the derivative of the utility function, apply the expectations operator and solve for $s_1$:

$\frac{-1}{y+y_1-s_1} + \frac{1}{\mu-\mathbb{E}_{t=1}s_2^*+s_1} = 0$

$\frac{1}{y+y_1-s_1} = \frac{1}{\mu-\mathbb{E}_{t=1}s_2^*+s_1}$

$y + y_1 - s_1 = \mu - \mathbb{E}_{t=1}s_2^* + s_1$

$s_1 = \frac{1}{2}(y + y_1 - \mu) + \frac{1}{2}\mathbb{E}_{t=1} s_2^*$

What is the expected value of $s_2^*$?

$\mathbb{E}_{t=1} s_2^* = \int_{\mu-s_1}^{y_{max}} \frac{1}{2}[y_2 + s_1 - \mu] \frac{1}{\mu} dy_2$

$\mathbb{E}_{t=1} s_2^* = \left[\frac{1}{4}[y_2 + s_1 - \mu]^2 \frac{1}{\mu} + C\right]_{\mu-s_1}^{y_{max}}$

$\mathbb{E}_{t=1} s_2^* = \left(\frac{1}{4}[y_{max} + s_1 - \mu]^2 \frac{1}{\mu} + C\right) - \left(\frac{1}{4}[\mu + s_1 - s_1 - \mu]^2 \frac{1}{\mu} + C\right)$

$\mathbb{E}_{t=1} s_2^* = \frac{1}{4\mu}[y_{max} + s_1 - \mu]^2$

This expression into $s_1$:

$s_1 = \frac{1}{2}(y + y_1 - \mu) + \frac{1}{2}\mathbb{E}_{t=1} \frac{1}{4\mu}[y_{max} + s_1 - \mu]^2$

$s_1 = \frac{1}{2}(y + y_1 - \mu) + \frac{1}{8\mu}[y_{max} + s_1 - \mu]^2$

We have to solve this for $s_1$:

$8\mu s_1 = 4\mu(y + y_1 - \mu) + [y_{max} + s_1 - \mu]^2$

$8\mu s_1 = 4\mu(y + y_1 - \mu) + \left[y_{max}^2 + \mu^2 + s_1^2 + 2\mu y_{max} + 2\mu s_1 + 2\mu y_{max}\right]$

$$s_1^2 + \underbrace{(2y_{max} - 10\mu)}_{p} s_1 + \underbrace{\left[4\mu\left(y + y_1 - \mu\right) + \left(y_{max} - \mu\right)^2\right]}_{q} = 0$$

This is the quadratic equation in the reduced form. Its solutions are given by:
$s_1^* = s_1^*\left(y_1\right) = \pm 2\sqrt{\mu}\sqrt{7\mu - 2y_{max} - y - y_1} + 5\mu - y_{max}$
Only one of the solutions makes economically sense. This solution is a convex function and only depends on $y_1$.

### A.1.3 The model with constraints in periods 1 and 2

Here, we assume that the individual cannot borrow in periods 1 and 2: $s_1 \geq 0$ and $s_2 \geq 0$. From Section A.1.2 we know the savings function in period 2 when liquidity is constrained:

$$s_2^* = s_2^*\left(y_2, s_1\right) = \begin{cases} \frac{1}{2}\left[y_2 + s_1 - \mu\right] & \text{if } \; y_2 + s_1 \geq \mu \\ 0 & \text{otherwise} \end{cases}$$

Now, we use $s_2^*$ in the three-period problem:
$\max_{s_1} u\left(y + y_1 - s_1\right) + \mathbb{E}_{t=1}\, u\left(y_2 - s_2^* + s_1\right) + \mathbb{E}_{t=1}\, u\left(y_3 + s_2^*\right)$
The FOC in period 1 in case that the constraint does not bind is given by:
$-u'\left(y + y_1 - s_1\right) + \mathbb{E}_{t=1}\, u'\left(y_2 - s_2^* + s_1\right) = 0$
$2s_1 = y + y_1 + \mathbb{E}_{t=1} s_2^* - \mu = 0$
This implies:
$0 \leq s_1 \leq \frac{1}{2}\left[y + y_1 + \mathbb{E}_{t=1} s_2^* - \mu\right]$
We also know the expected value of $s_2^*$ from the previous section and plug it in:
$0 \leq s_1 \leq \frac{1}{2}\left[y + y_1 + \frac{1}{4\mu}\left[y_{max} + s_1 - \mu\right]^2 - \mu\right]$
With this result, we can derive the savings function in period 1:

$$s_1^* = s_1^*\left(y_1\right) = \begin{cases} \pm 2\sqrt{\mu}\sqrt{7\mu - 2y_{max} - y - y_1} + 5\mu - y_{max} & \text{if } \; y_1 \geq \mu \\ 0 & \text{otherwise} \end{cases}$$

The function has a kink and savings in period 1 are thus, for incomes greater than the mean income, identical to the savings in period 1 in the previous section and zero for incomes smaller than mean income.

## A.2 Power calculation

Here, we will calculate aggregate predictions for our treatments, use them to form testable hypotheses, determine the number of required subjects in each treatment, and specify how we test our hypotheses. Table A.1 displays means and standard deviations for the savings in periods 1 and 2 in all treatments. They are based on the savings functions and income and wealth ranges as shown in Figure 2.[33]

The first two hypotheses test the consequences of the liquidity constraint (these two between-treatment hypotheses are the focus of our study). Namely, we expect that the constraint will make individuals save more in both period 1 and 2 (compared to the unconstrained treatment of the same income variance). We use the means and standard deviations in Table 2 to calculate the expected effect sizes between the variables (measured using Cohen's $d$, Cohen 1988) and use the online software Statulator (Dhand and Khatkar, 2014) to calculate the number of required subjects.[34]

**Hypothesis 1a:** In the ConHigh treatment, $s_1$ will be higher than in the UncHigh treatment. (Number of required subjects according to our power calculation: 91 Subjects.)

**Hypothesis 1b:** In the ConLow treatment, $s_1$ will be higher than in the UncLow treatment. (Number of required subjects according to our power calculation: 1,429 subjects.)

**Hypothesis 2a:** In the ConHigh treatment, $s_2$ will be higher than in the UncHigh treatment. (Number of required subjects according to our power calculation: 93 subjects.)

**Hypothesis 2b:** In the ConLow treatment, $s_2$ will be higher than in the UncLow treatment. (Number of required subjects according to our power calculation: 30 subjects.)

The next hypothesis examines savings behavior *within* the treatments:

**Hypothesis 3a:** In the UncHigh treatment, $s_2$ will be higher than $s_1$. (Number of required subjects according to our power calculation: 10 subjects.)

**Hypothesis 3b:** In the ConHigh treatment, $s_2$ will be higher than $s_1$. (Number of required subjects according to our power calculation: 6 subjects.)

**Hypothesis 3c:** In the UncLow treatment, $s_2$ will be higher than $s_1$. (Number of required subjects according to our power calculation: 50 subjects.)

**Hypothesis 3d:** In the ConLow treatment, $s_2$ will be higher than $s_1$. (Number of required subjects according to our power calculation: 5 subjects.)

Finally, the last three hypotheses examine how the variance of the income affects savings behavior[35]:

**Hypothesis 4:** In the ConHigh treatment, $s_1$ will be higher than in the ConLow treatment. (Total number of required subjects according to our power calculation: 244 subjects.)

**Hypothesis 5:** In the UncHigh treatment, $s_2$ will be higher than in the UncLow treatment. (Total number of required subjects according to our power calculation: 86 subjects.)

**Hypothesis 6:** In the ConHigh treatment, $s_2$ will be higher than in the ConLow treatment. (Total number of required subjects according to our power calculation: 66 subjects.)

---

[33] In Section 6.1.2 (and in Section A.1.3 in the Appendix), we derive the savings function $s_1^{\mathrm{DA}}$ where liquidity is constrained in *both* period 1 and period 2. (Notice that $s_2^{\mathrm{DA}} = s_2^{\mathrm{Con}}$.) In High, the mean of $s_1^{\mathrm{DA}}$ is 5.9 ($\sigma = 7.6$). In Low, the mean of $s_1^{\mathrm{DA}}$ is 1.5 ($\sigma = 1.8$). Given these values, we can also test the differences between the Unc, Con, and DA variables with the observations we will collect.

[34] We assume the following for our power calculations: 15 observations for each subject in Con and Unc, observations clustered at subject-level, $\alpha = 0.05$, $1 - \beta = 0.8$, an intraclass correlation coefficient of 0.5. We will test Hypotheses 1-3 with pairwise two-sided t-tests (clustered at the subject-level).

[35] We will test Hypotheses 4-6 with two-sample two-sided t-tests (clustered at the subject-level). We assume 15 observations for each subject in Con and Unc, $\alpha = 0.05$, $1 - \beta = 0.8$, and an intraclass correlation coefficient of 0.5.)

Hypotheses 1b and 4 require too many subjects (independent observations) for our budget. We will test them but keep in mind that our power might not be sufficient. We conduct sessions with 100 subjects in High and with 50 subjects in Low.

We will test our hypotheses using OLS regressions where we restrict the observations to the treatments we want to compare and regress savings on a dummy-variable for the treatments to compare.

## A.3 Additional tables

**Table A.1:** Aggregate predictions, expected effect sizes, and hypotheses

H4, $d = 0.299$

| | | High | | | | | Low | | |
|---|---|---|---|---|---|---|---|---|---|
| | | Unc | Con | | | | Unc | Con | |
| | | H1a, $d = 0.218$ | | | | | H1b, $d = 0.0055$ | | |
| $s_1$ | H3a, $d = 0.74$ | 0.0<br>(13.7) | 2.8<br>(11.9) | H3b, $d = 1.109$ | $s_1$ | H3c, $d = 0.3$ | 0.0<br>(4.0) | 0.2<br>(3.1) | H3d, $d = 1.161$ |
| $s_2$ | | 17.5<br>(30.5) | 23.4<br>(23.4) | | $s_2$ | | 5.0<br>(23.2) | 12.7<br>(14.9) | |
| | | H2a, $d = 0.217$ | | | | | H2b, $d = 0.394$ | | |

H5, $d = 0.461$

H6, $d = 0.545$

Note: The cells in the middle present mean predictions and their standard deviations (in parentheses). The curly brackets show which variables are compared to test our hypotheses (marked with H). The $d$-values present the expected effect size (Cohen's $d$) when the respective variables are compared.

**Table A.2:** Multiple price lists of the risk aversion elicitation

| Lottery A | | | | Lottery B | | | | EV[A] | EV[B] | Difference | Implied CRRA | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| $p$ | $ | $p$ | $ | $p$ | $ | $p$ | $ | ($) | ($) | ($) | interval | |
| 0.5 | 4.67 | 0.5 | 4.44 | 0.5 | 5.58 | 0.5 | 2.78 | 4.556 | 4.181 | 0.375 | $-\infty$ | −1.71 |
| 0.5 | 4.89 | 0.5 | 4.44 | 0.5 | 6.03 | 0.5 | 2.78 | 4.667 | 4.403 | 0.264 | −1.71 | −0.95 |
| 0.5 | 5.11 | 0.5 | 4.44 | 0.5 | 6.44 | 0.5 | 2.78 | 4.778 | 4.611 | 0.167 | −0.95 | −0.49 |
| 0.5 | 5.33 | 0.5 | 4.44 | 0.5 | 6.89 | 0.5 | 2.78 | 4.889 | 4.833 | 0.056 | −0.49 | −0.15 |
| 0.5 | 5.56 | 0.5 | 4.44 | 0.5 | 7.36 | 0.5 | 2.78 | 5.000 | 5.069 | -0.069 | −0.15 | 0.14 |
| 0.5 | 5.78 | 0.5 | 4.44 | 0.5 | 7.94 | 0.5 | 2.78 | 5.111 | 5.361 | -0.250 | 0.14 | 0.41 |
| 0.5 | 6.00 | 0.5 | 4.44 | 0.5 | 8.72 | 0.5 | 2.78 | 5.222 | 5.750 | -0.528 | 0.41 | 0.68 |
| 0.5 | 6.22 | 0.5 | 4.44 | 0.5 | 9.83 | 0.5 | 2.78 | 5.333 | 6.306 | -0.972 | 0.68 | 0.97 |
| 0.5 | 6.44 | 0.5 | 4.44 | 0.5 | 12.50 | 0.5 | 2.78 | 5.444 | 7.639 | -2.194 | 0.97 | 1.37 |
| 0.5 | 6.67 | 0.5 | 4.44 | 0.5 | 13.06 | 0.5 | 2.78 | 5.556 | 7.917 | -2.361 | 1.37 | $+\infty$ |

Note: Last five columns showing expected values and implied CRRA intervals were not shown to subjects.

**Table A.3:** Cognitive reflection test items (Toplak et al, 2014)

| Item | Question, correct answer & intuitive (wrong) answer |
|---|---|
| #1 | If John can drink one barrel of water in 6 days, and Mary can drink one barrel of water in 12 days, how long would it take them to drink one barrel of water together? ___ days [correct answer 4 days; intuitive answer 9]. |
| #2 | Jerry received both the 15th highest and the 15th lowest mark in the class. How many students are in the class? ___ students [correct answer 29 students; intuitive answer 30]. |
| #3 | A man buys a pig for $60, sells it for $70, buys it back for $80, and sells it finally for $90. How much has he made? ___ dollars [correct answer $20; intuitive answer $10]. |
| #4 | Simon decided to invest $8,000 in the stock market one day early in 2008. Six months after he invested, on July 17, the stocks he had purchased were down 50%. Fortunately for Simon, from July 17 to October 17, the stocks he had purchased went up 75%. At this point, Simon has: a. broken even in the stock market, b. is ahead of where he began, c. has lost money [correct answer c, because the value at this point is $7,000; intuitive response b]. |

**Table A.4:** Estimated linear savings functions in period 2 of CON treatments

| | ConHigh | | | | ConLow | | | |
|---|---|---|---|---|---|---|---|---|
| | Before the theoretical threshold at 70 | | | | | | | |
| | | Prediction | | Estimation | | Prediction | | Estimation |
| $s_2$ | $w_2$ | 0.000 | <<< | 0.080*** | $w_2$ | 0.000 | << | 0.071** |
| | | | | (0.018) | | | | (0.032) |
| | constant | 0.000 | | 0.913 | constant | 0.000 | | 1.523 |
| | | | | (0.852) | | | | (1.889) |
| | #obs. | | | 675 | #obs. | | | 399 |
| | #clusters | | | 97 | #clusters | | | 49 |
| | $R^2$ (within) | | | 0.038 | $R^2$ (within) | | | 0.014 |
| | $R^2$ (between) | | | 0.017 | $R^2$ (between) | | | 0.091 |
| | $R^2$ (overall) | | | 0.019 | $R^2$ (overall) | | | 0.021 |
| | After the theoretical threshold at 70 | | | | | | | |
| | | Prediction | | Estimation | | Prediction | | Estimation |
| $s_2$ | $w_2$ | 0.500 | <<< | 0.587*** | $w_2$ | 0.500 | | 0.537*** |
| | | | | (0.033) | | | | (0.064) |
| | constant | -35.000 | << | -26.734*** | constant | -35.000 | | -30.969*** |
| | | | | (3.363) | | | | (5.279) |
| | #obs. | | | 825 | #obs. | | | 351 |
| | #clusters | | | 98 | #clusters | | | 49 |
| | $R^2$ (within) | | | 0.301 | $R^2$ (within) | | | 0.187 |
| | $R^2$ (between) | | | 0.636 | $R^2$ (between) | | | 0.339 |
| | $R^2$ (overall) | | | 0.428 | $R^2$ (overall) | | | 0.376 |

Note: Estimations are based on fixed-effect panel regressions. Standard errors in parentheses. ***, **, and * show differences from zero at the 1%-, 5%-, and 10%-level, respectively. >>>, >>, and > (<<<, <<, <) display the results from two-sided Wald tests where the estimated coefficient is significantly smaller (larger) than the prediction (significant at the 1%-, 5%-, and 10%-level, respectively).

**Table A.5:** Ratios of correct savings decisions (by subject)

| Treatment | Period | Mean | Minimum | Median | Maximum |
|---|---|---|---|---|---|
| UncHigh | 1 | 1.305 | 0.000 | 1.183 | 7.500 |
| | | | (4%) | | |
| | 2 | 0.984 | 0.000 | 1.000 | 2.143 |
| | | | (4%) | | |
| ConHigh | 1 | 1.073 | 0.000 | 1.000 | 5.000 |
| | | | (5%) | | |
| UncLow | 1 | 1.383 | 0.000 | 1.146 | 11.000 |
| | | | (2%) | | |
| | 2 | 1.408 | 0.000 | 1.100 | 14.000 |
| | | | (2%) | | |
| ConLow | 1 | 1.021 | 0.000 | 1.000 | 3.750 |
| | | | (10%) | | |

Note: We calculated for each subject the ratio of correct savings decisions ((number of decisions with savings>0)/(number of decisions where savings>0 is optimal)). A ratio of 0 means that the subject never saved, 1 describes optimal savings behavior, > 1 that the subject saved more frequently than optimal. Numbers in parentheses give the share of subjects who never saved.

**Table A.6:** Ratios of correct borrowing decisions (by subject)

| Treatment | Period | Mean | Minimum | Median | Maximum |
|---|---|---|---|---|---|
| UncHigh | 1 | 0.730 | 0.000 (16%) | 0.714 | 3.000 |
| | 2 | 0.780 | 0.000 (20%) | 0.833 | 3.000 |
| ConHigh | 1 | 0.866 | 0.000 (20%) | 0.775 | 3.750 |
| UncLow | 1 | 0.905 | 0.000 (8%) | 0.800 | 2.500 |
| | 2 | 0.724 | 0.000 (12%) | 0.727 | 2.333 |
| ConLow | 1 | 0.827 | 0.000 (16%) | 1.000 | 3.000 |

Note: We calculated for each subject the ratio of correct borrowing decisions ((number of decisions with savings<0)/(number of decisions where savings<0 is optimal)). A ratio of 0 means that the subject never borrowed, 1 describes optimal borrowing behavior, > 1 that the subject borrowed more frequently than optimal. Numbers in parentheses give the share of subjects who never borrowed.

**Table A.7:** Model comparisons using mean absolute errors

| | Model | Data | | | |
|---|---|---|---|---|---|
| | | UncHigh | ConHigh | UncLow | ConLow |
| $s_1$ | Unc | 16.92 | 16.70 | 9.91 | 9.39 |
| | Con | 16.09 | $\underline{15.89}$ | $\underline{9.81}$ | 9.29 |
| | DA | $\underline{15.98}$ | 16.17 | 9.88 | $\underline{9.27}$ |
| $s_2$ | Unc | $\underline{18.12}$ | 18.98 | $\underline{11.17}$ | 10.45 |
| | Con | 18.71 | $\underline{14.24}$ | 11.21 | $\underline{7.42}$ |

Note: MAE $= \frac{1}{N}\sum_{i=1}^{N} |\text{prediction}_i - \text{observation}_i|$. Model with best fit is underlined (perfect score=0).

**Table A.8:** Model comparisons using mean errors

| | Model | Treatment | | | |
|---|---|---|---|---|---|
| | | UncHigh | ConHigh | UncLow | ConLow |
| $s_1$ | Unc | -8.285 | -6.795 | -1.341 | $\underline{0.557}$ |
| | Con | -5.530 | -4.004 | -1.154 | 0.760 |
| | DA | $\underline{-2.003}$ | $\underline{-0.185}$ | $\underline{0.065}$ | 2.026 |
| $s_2$ | Unc | -7.849 | -16.260 | -3.807 | 9.335 |
| | Con | $\underline{-3.683}$ | $\underline{-11.522}$ | $\underline{-1.172}$ | $\underline{6.306}$ |

Note: ME $= \frac{1}{N}\sum_{i=1}^{N} (\text{prediction}_i - \text{observation}_i)$. Model with best fit is underlined (perfect score=0).

**Table A.9:** Model comparisons using bias

| | Model | Treatment | | | |
|---|---|---|---|---|---|
| | | UncHigh | ConHigh | UncLow | ConLow |
| $s_1$ | Unc | 0.0395 | 0.0254 | 0.0534 | -0.0097 |
| | Con | 0.3588 | 0.4257 | 0.1852 | $\underline{-0.3768}$ |
| | DA | $\underline{0.7678}$ | $\underline{0.9735}$ | $\underline{1.0461}$ | -2.6718 |
| $s_2$ | Unc | 0.3640 | 0.1742 | 0.1593 | -0.0384 |
| | Con | $\underline{0.7015}$ | $\underline{0.4148}$ | $\underline{0.7412}$ | $\underline{0.2985}$ |

Note: Bias $= \frac{\frac{1}{N}\sum_{i=1}^{N} \text{prediction}_i}{\frac{1}{N}\sum_{i=1}^{N} \text{observation}_i}$. Model with best fit is underlined (perfect score=1).

**Table A.10:** Correlation matrix of subjects' characteristics

| | Female | Age | CRT score | GPA | ToH solved | Moves if ToH solved | Risk aversion |
|---|---|---|---|---|---|---|---|
| Female | 1.000 | | | | | | |
| Age | 0.0923 | 1.000 | | | | | |
| | (1.0000) | | | | | | |
| CRT score | -0.2722 | -0.0062 | 1.000 | | | | |
| | (0.0000) | (1.0000) | | | | | |
| GPA | 0.2048 | 0.0124 | -0.0376 | 1.000 | | | |
| | (0.0075) | (1.0000) | (1.0000) | | | | |
| ToH solved | -0.1173 | 0.0999 | 0.1353 | -0.0330 | 1.000 | | |
| | (0.8894) | (1.0000) | (0.3994) | (1.0000) | | | |
| Moves if ToH solved | -0.2224 | 0.0550 | 0.0115 | 0.0115 | . | 1.000 | |
| | (0.0037) | (1.0000) | (1.0000) | (1.0000) | | | |
| Risk aversion | -0.0201 | -0.1723 | -0.0880 | -0.0478 | -0.0121 | -0.1136 | 1.000 |
| | (1.0000) | (0.1922) | (1.0000) | (1.0000) | (1.0000) | (1.0000) | |

Note: Pairwise correlation coefficients of the variables in Table 9. Bonferroni-adjusted significance levels in parentheses. Number of observations for the different coefficients differ slightly. CRT=cognitive reflection test. GPA=grade point average. ToH=Tower of Hanoi.

**Table A.11:** Determinants of pooled conditional deviations in period 2

| | \|Deviation\| | Deviation>0 | Deviation<0 | Deviation |
|---|---|---|---|---|
| Female==1 | 1.102 | 0.692 | -2.736** | -0.472 |
| | (1.927) | (2.662) | (1.355) | (2.049) |
| Age | -0.333 | 0.249 | 0.544 | 0.236 |
| | (0.905) | (1.110) | (0.371) | (0.934) |
| CRT score | -2.369*** | -2.729*** | 2.176*** | -0.190 |
| | (0.650) | (0.831) | (0.487) | (0.717) |
| GPA | 0.376 | 0.782 | 0.392 | 1.173 |
| | (1.721) | (1.966) | (0.788) | (1.890) |
| ToH solved==1 | -3.934 | -9.397 | -4.443 | -17.181** |
| | (8.649) | (9.489) | (5.725) | (6.663) |
| ToH solved==1 * Moves | -0.0167 | 0.089 | 0.029 | 0.137 |
| | (0.176) | (0.272) | (0.120) | (0.200) |
| Risk aversion | 0.836* | 1.078* | 0.484* | 1.457*** |
| | (0.475) | (0.605) | (0.252) | (0.513) |
| constant | 20.279 | 12.417 | -23.142 | 4.621 |
| | (21.220) | (25.306) | (10.089) | (21.139) |
| | | | | |
| #obs. | 3,420 | 2,198 | 956 | 3,420 |
| #clusters | 114 | 114 | 110 | 114 |
| $R^2$ (within) | — | — | — | — |
| $R^2$ (between) | 0.135 | 0.108 | 0.268 | 0.101 |
| $R^2$ (overall) | 0.066 | 0.092 | 0.113 | 0.042 |

Note: Estimations based on panel regressions. Standard errors in parentheses. ***, **, and * show difference from zero (based on two-sided t-tests clustered at the subject-level) at the 1%-, 5%-, and 10%-level, respectively.

### A.4 Additional figures

**Figure A.1:** Estimated linear savings functions, LOWESS filters, and savings predictions in period 1

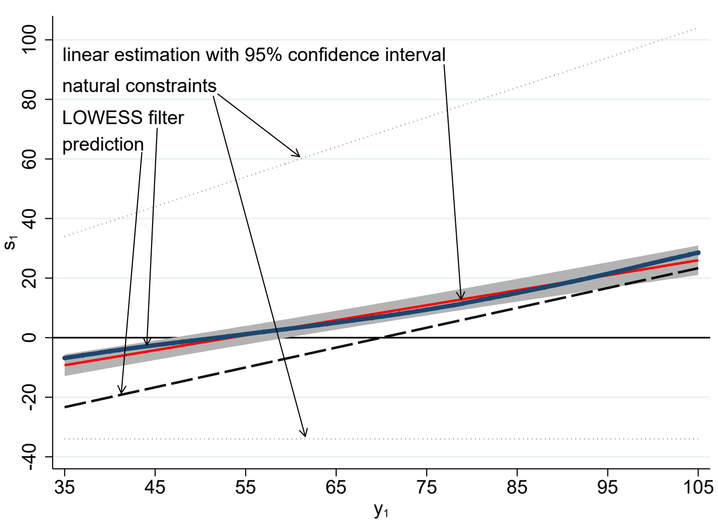


**(a)** Period 1 in UncHigh

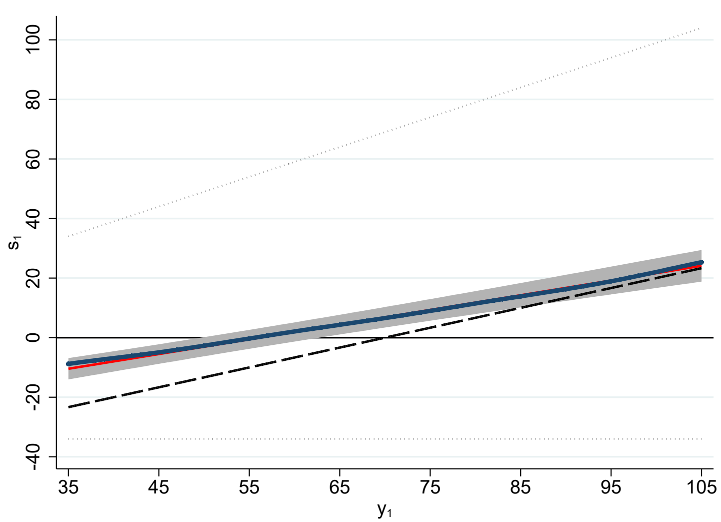


**(b)** Period 1 in ConHigh

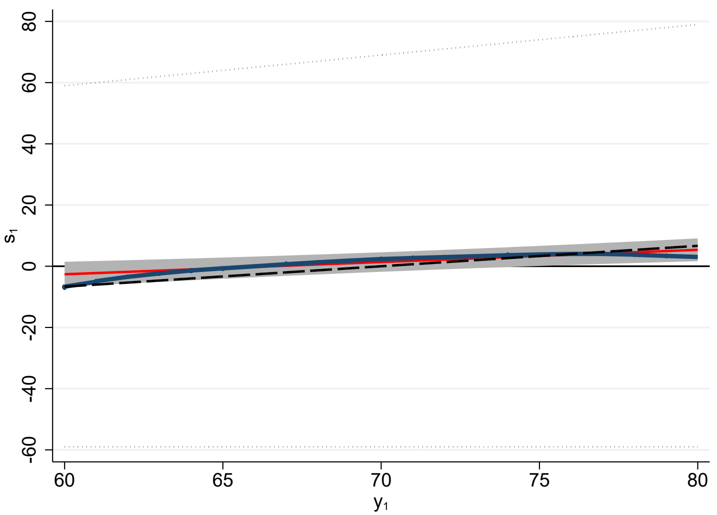


**(c)** Period 1 in UncLow

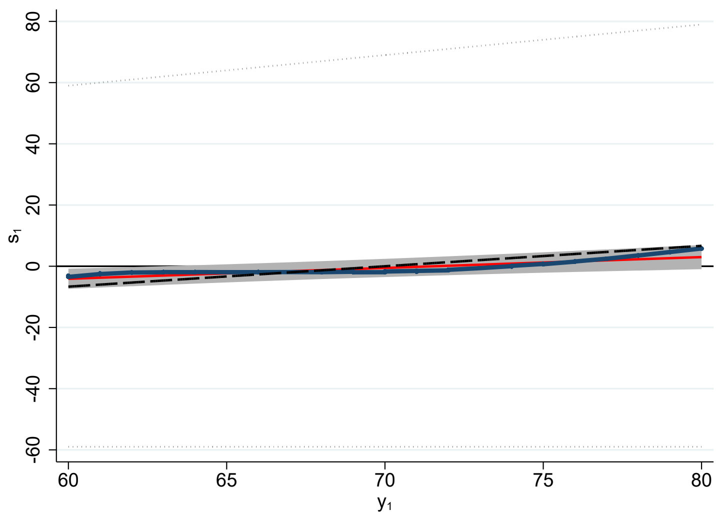


**(d)** Period 1 in ConLow

**Figure A.2:** Linear (threshold) savings functions, LOWESS filters, and savings predictions in period 2

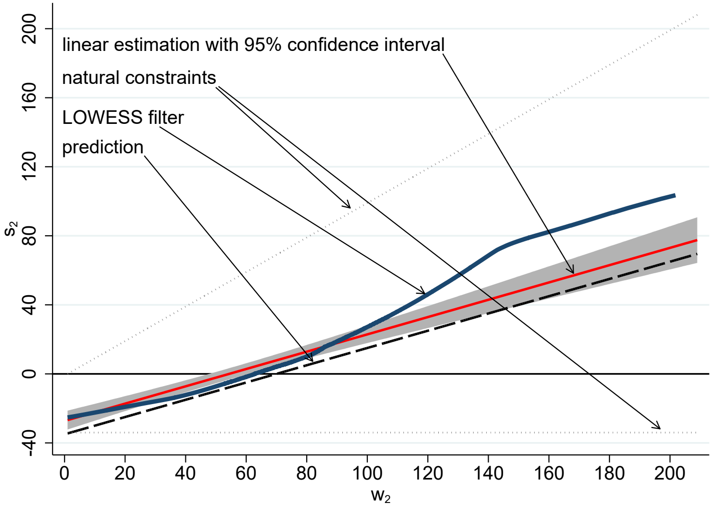


**(a)** Period 2 in UncHigh

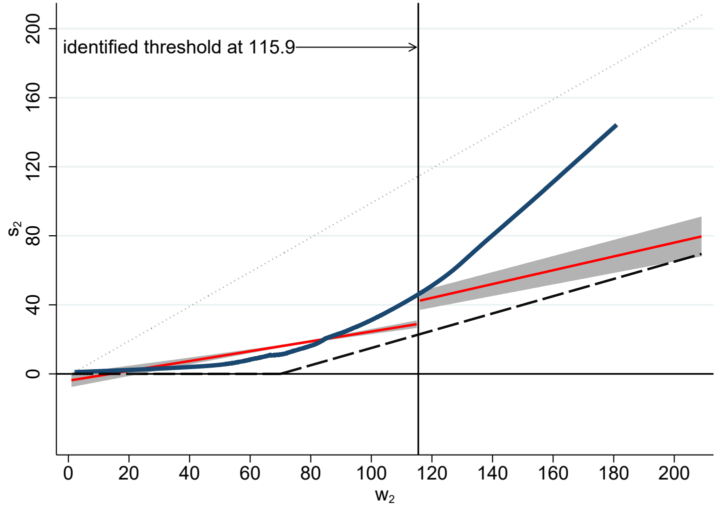


**(b)** Period 2 in ConHigh

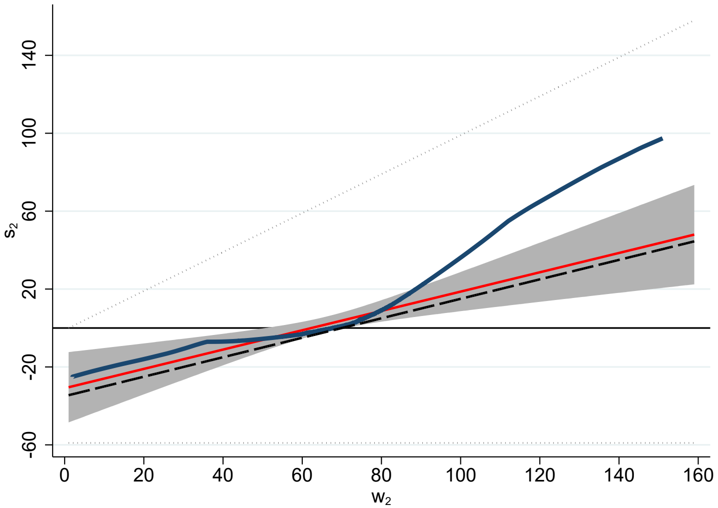


**(c)** Period 2 in UncLow

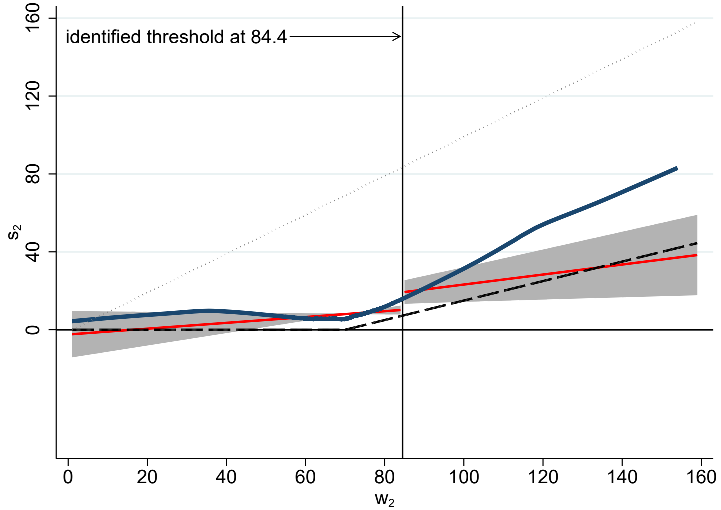


**(d)** Period 2 in ConLow

## A.5 Tower of Hanoi game

We let our subjects solve the Tower of Hanoi (ToH) game which is a one-player mathematical puzzle (first described by the French mathematician Édouard Lucas and explained in detail in Hinz et al 2013). Solving the ToH game requires the ability to think recursively, which is closely related to backward induction.[36]

We let subjects play the ToH game with three different-sized disks and three poles (displayed in Figure A.3). Here, the task is to move the three disks (which initially are stacked up on the first pole) onto the third pole in the same initial order (the second pole is an auxiliary) using the least number of moves possible and obeying the following rules: (i) only one disk can be moved at a time; (ii) each move consists of taking the upper disk from one of the stacks and placing it on top of another stack or on an empty pole; (iii) no larger disk may be placed on top of a smaller disk. In our case, the minimum number of moves is seven.

**Figure A.3:** Illustration of the Tower of Hanoi game


We incentivized the ToH game using the following payoff function that depends on the number of moves needed to solve the game:

$$\text{Payoff in dollars (moves)} = \max\left[10 - 0.5 \cdot \text{moves}, 0\right]. \quad (14)$$

Hence, if a subject solved the task in the minimum number of seven moves, she received \$6.50 and nothing if she exceeded 20 moves. Subjects had three minutes time to read the instructions and solve the task.

[36] The ToH game has some advantages over games like the Race game (Gneezy et al, 2010) as (i) we do not want that two subjects in our experiment play against one another (if only the first mover in a game has the chance to ensure herself a win, we would lose half of our observations); (ii) we do not only want a binary outcome, we want a finer measure as well; and (iii) we do not want to let our subjects play a game as a first-mover against the computer (as it would be hard to communicate the computers' (winning) strategy).

## A.6 Stimulus-response learning

In addition to looking for evidence that subjects were learning to make better decisions, over time, we also evaluate deviations from theoretical predictions using a basic stimulus-response model. The idea behind this model is that subjects not (only) learn over time but also update their decisions based on the feedback they receive, i.e., in response to past income realizations, and sufficiently clear savings/borrowing decision "mistakes". In particular, we introduce two sets of binary variables. The first characterizes each subject's saving decision (in periods 1 and 2) and the second characterizes income realizations in the following period *after* a savings decision has been taken (that is, in periods 2 and 3). The first set, "extremely high saving" ("extremely low saving") takes the value 1 if a subject's saving deviated by more than 50% from the optimal prediction so that they saved either too much or too little. The second set "extremely high income" ("extremely low income") takes the value 1 if the income realization following the savings decision was in the upper (lower) 25% of the treatments-specific income range.

We want to evaluate how subjects adapt their savings behavior in case they deviated very strongly from theoretical predictions (by saving too much or too little) and subsequently received an income that was the very reverse of what this extreme savings amount would require. We expect that subjects who saved very much and then received a high income (saved very little and then received a low income) would adjust their behavior by subsequently lowering (increasing) their deviation from theoretical predictions in subsequent periods (the so-called "hot-stove effect": once burned, you don't touch the hot stove again).

Table A.12 shows fixed-effect panel data regressions where we, in a first model, regress pooled (conditional) deviations from optimal savings of both periods on the interaction term of extremely high saving and extremely high income and the interaction term of extremely low saving and extremely low income. (All models also contain the number of rounds and a constant.) We use lagged variables (of the previous lifecycle), denoted by L., when explaining period 1 deviations and variables from the same lifecycle when explaining period 2 deviations. In a second model, we try to decompose the effects from the saving and the income variables and regress the deviations on the extreme saving variables, the extreme income variables, and the two interaction terms.

The table reveals that all interaction terms in the four regressions have the expected sign, though only two of the coefficients are significantly different from zero. We also observe that extremely high savings in the past predict extremely high savings in the present and that learning over time significantly decreases the positive constant in all specifications.

**Table A.12:** Past experiences as a determinant of (conditional) deviations from optimal behavior

| | | Deviation | Deviation |
|---|---|---|---|
| Period 1 | L.extremely high saving2 * L.extremely high income3 | -0.361 | -2.298* |
| | | (0.699) | (1.194) |
| | L.extremely high saving2 | | 2.837*** |
| | | | (0.705) |
| | L.extremely high income3 | | 1.384 |
| | | | (0.950) |
| | L.extremely low saving2 * L.extremely low income3 | 0.736 | 0.187 |
| | | (1.032) | (1.347) |
| | L.extremely low saving2 | | 1.684** |
| | | | (0.798) |
| | L.extremely low income3 | | 0.522 |
| | | | (0.700) |
| | round | -0.181*** | -0.178*** |
| | | (0.029) | (0.029) |
| | constant | 7.012*** | 4.632*** |
| | | (0.546) | (0.799) |
| | #obs. | 4,350 | 4,350 |
| | #clusters | 150 | 150 |
| | $R^2$ (within) | 0.009 | 0.013 |
| | $R^2$ (between) | 0.118 | 0.032 |
| | $R^2$ (overall) | 0.004 | 0.011 |
| Period 2 | extremely high saving1 * extremely high income2 | -2.015*** | -1.024 |
| | | (0.755) | (1.192) |
| | extremely high saving1 | | 1.299* |
| | | | (0.688) |
| | extremely high income2 | | -2.107** |
| | | | (0.903) |
| | extremely low saving1 * extremely low income2 | 0.455 | 1.859 |
| | | (0.880) | (1.209) |
| | extremely low saving1 | | -3.170*** |
| | | | (0.708) |
| | extremely low income2 | | -0.000 |
| | | | (0.785) |
| | round | -0.258*** | -0.246*** |
| | | (0.029) | (0.029) |
| | constant | 12.378*** | 12.983*** |
| | | (0.531) | (0.781) |
| | #obs. | 4,500 | 4,500 |
| | #clusters | 150 | 150 |
| | $R^2$ (within) | 0.020 | 0.032 |
| | $R^2$ (between) | 0.104 | 0.267 |
| | $R^2$ (overall) | 0.007 | 0.054 |

Note: Estimations are based on fixed-effect panel regressions. L. indicates that an independent variable is from the previous round/lifecycle. Standard errors in parentheses. ***, **, and * show differences from zero at the 1%-, 5%-, and 10%-level, respectively.

# Experimental instructions appendix

## B.1 Instructions for the High treatment (Unc-Con)

# Instructions - Part 1

### Overview

Welcome to this experiment in the economics of decision-making. Funding for this experiment has been provided by the UC Irvine School of Social Sciences. Please do not talk with one another and silence and stow all mobile devices.

For your participation in today's experiment, you will be paid in cash and in private at the end of the experiment. Your money payoff depends partly on your own decisions and partly on chance. There will be a short quiz following the reading of these instructions which you will all need to complete before we can begin the experiment.

The experiment consists of three parts. These instructions are for part 1. At the end of part 1, you will be given instructions for part 2, and thereafter, you will receive instructions for part 3. You will be able to earn payments from all 3 parts of the experiment. All three parts are independent of one another. After the experiment you will be asked to fill out a short questionnaire.

In this first part of the experiment, you face 15 "rounds" of decision-making. Each round consists of three "periods." In periods 1 and 2 of each round, you will receive some information and then enter a choice using your computer workstation. Your decisions in periods 1 and 2 determine the point consumption in each of the three periods and a payoff function will transform your point consumption into dollar earnings, as explained below. At the end of the experiment, we will randomly select one from the 15 rounds played in this first part for payoff. You will receive payments from all three parts together with your $7 show-up payment at the end of the experiment.

### Specific Details

Each round consists of three periods. In each period, you receive some income amount in points, denoted by $y_1$, $y_2$ and $y_3$. Each of the three incomes is a random draw from a uniform distribution over the interval [35, 105], inclusive. This means that any integer number in the set $\{35, 36, 37, ..., 103, 104, 105\}$ is equally likely to be drawn and the mean or expected value for income, $y_i$, $i = 1, 2, 3$, is always 70. Your task is to choose how much to consume in each of the first two periods of the round. In making these two consumption choices, you must consume at least 1 point per period and you can either borrow from future income to increase your current period consumption or you can save current period income for future period consumption, within limits, as explained below. In the final, third period of each round, you automatically consume your period 3 income plus any savings and minus any borrowings.

Your payoff for each round is given by:

$$\text{Round payoff in dollars} = f(c_1) + f(c_2) + f(c_3) \tag{15}$$

where $c_1$, $c_2$ and $c_3$ are the points you choose to consume in periods 1, 2 and 3 and $f(\cdot)$ is a *function* converting point consumption in each period into dollars. Specifically, the function $f(c_i) = .77\ln(c_i)$, for $i$ = 1,2,3. A graph of $f(c_i)$ is shown in Figure 1, Table 1 provides a list of possible integer values for $f(c_i)$.

## Period 1

On the period 1 decision screen for each round, you learn your income for period 1, $y_1$, a random draw from the interval $[35, 105]$. After you learn $y_1$, you choose the amount you wish to consume in period 1, $c_1$.

With your consumption choice, you (implicitly) also choose if and how much you want to save or borrow. If your consumption choice, $c_1$, is less than your period 1 income, $y_1$, then you save the amount $s_1 = y_1 - c_1$ (and $s_1$ is positive). If your consumption choice, $c_1$, is greater than your period 1 income, then you borrow the amount $s_1 = y_1 - c_1$ (and $s_1$ is negative). There is no interest paid on savings nor is there any interest charged on borrowing.

The maximum you can save is $y_1 - 1$ since you must consume at least 1 point (see Figure 1 and Table 1), and the maximum you can borrow is 34. This ensures that your consumption choice in the next period is also in the range of the function $f(c_i)$.

On the period 1 decision screen, below your income, $y_1$, is a slider where you enter a choice for period 1 consumption, $c_1$. Your choice can be any number in this interval up to one decimal place. After you have moved the slider to make your choice for $c_1$, click the submit button. You can change your mind anytime before clicking the submit button.

## Period 2

Following your choice for $c_1$, you will next face the period 2 decision screen. There you will learn your income for period 2, $y_2$, which again is a random draw from the interval $[35, 105]$. You will also be reminded of any savings or borrowings from period 1. The sum of your period 2 income, $y_2$, and your period 1 savings or borrowings, $s_1$, $(y_2 + s_1)$, is available to you for your consumption choice in period 2 and, in addition, as in period 1, you can borrow up to an additional 34 points for period 2 consumption. Again, you must consume at least 1 point.

Any choice of $c_2$ that is less than $y_2 + s_1$ results in period 2 savings, $s_2 = y_2 + s_1 - c_2$ (which is positive), while any choice of $c_2$ that is greater than $y_2 + s_1$ results in period 2 borrowings of $s_2 = y_2 + s_1 - c_2$ (which is negative). Again, there is no interest paid on savings nor is interest charged on borrowings. After you have moved the slider to make your choice for $c_2$, click the submit button. You can change your mind anytime before clicking the submit button.

## Period 3

Following your choice for $c_2$, you will next see the final period 3 screen. There you will learn your income for period 3, $y_3$, which again is a random draw from the interval $[60, 80]$. Since period 3 is the final period of the round, there is no decision for you to make. Your consumption for period 3 is automatically determined for you. Specifically, your consumption for period 3 is given by $c_3 = y_3 + s_2$, where $y_3$ is your period 3 income and $s_2$ is any savings (if positive) or borrowings (if negative) from your period 2 consumption decision. Also, the period 3 screen will show you your total payment in dollars for the round, which depends on your consumption choices in *all three* periods, $c_1$, $c_2$ and $c_3$, and which is determined according to formula (15).

After you have viewed the information on the end-of-round, period 3 screen, click the OK button to continue. If the 15th round has not yet been played, you will move on to the next three-period round, where you will again face the same random process for your income in each of the three periods and where you will complete the same sequence of consumption choices and face the same payment formula for the round. Note that the income draws made in each period and round, $y_1$, $y_2$ and $y_3$, will always be *independent* random draws from the interval $[35, 105]$ and so will most likely differ from round to round and from period to period.

### Feedback

After each round, a history of all your choices and payments in all prior rounds in this part will be shown to you.

### Earnings

Your payment from this first part will equal your payment from 1 of the 15 rounds, chosen randomly, from all 15 rounds. Each round has an equal chance of being chosen. You will learn the round chosen, and your payment from this first part only *after* the completion of the third and final part.

### Questions?

Now is the time for questions. If you have a question about any aspect of these instructions, please raise your hand and an experimenter will answer your question.

# Instructions - Part 2

## Overview

In part 2, as in part 1, you will again participate in 15 rounds of decision-making. Part 2 is very similar to part 1, with only one difference that is explained below. Therefore, compared to the part 1 of the instructions, only the paragraph “Period 2” changes. The rest of the instructions is still valid.

## Period 2

Following your choice for $c_1$, you will next face the period 2 decision screen. There you will learn your income for period 2, $y_2$, which again is a random draw from the interval $[35, 105]$. You will also be reminded of any savings or borrowings from period 1. The sum of your period 2 income, $y_2$, and your period 1 savings or borrowings, $s_1$, $(y_2 + s_1)$, is available to you for your consumption choice in period 2. Again, you must consume at least 1 point. In period 2, in contrast to period 1, you cannot borrow anymore. Thus, your consumption choice, $c_2$, cannot be greater than $y_2 + s_1$.

Any choice of $c_2$ that is less than $y_2 + s_1$ results in period 2 savings, $s_2 = y_2 + s_1 - c_2$ (which is positive). Again, there is no interest paid on savings. After you have moved the slider to make your choice for $c_2$, click the submit button. You can change your mind anytime before clicking the submit button.

## Questions?

Now is the time for questions. If you have a question about any aspect of these instructions, please raise your hand and an experimenter will answer your question.

## B.2 Graph and table handout for the High treatment

**Figure B.1:** A graph of the function $f(c_i) = .77\ln(c_i)$ for all possible values $c_i$

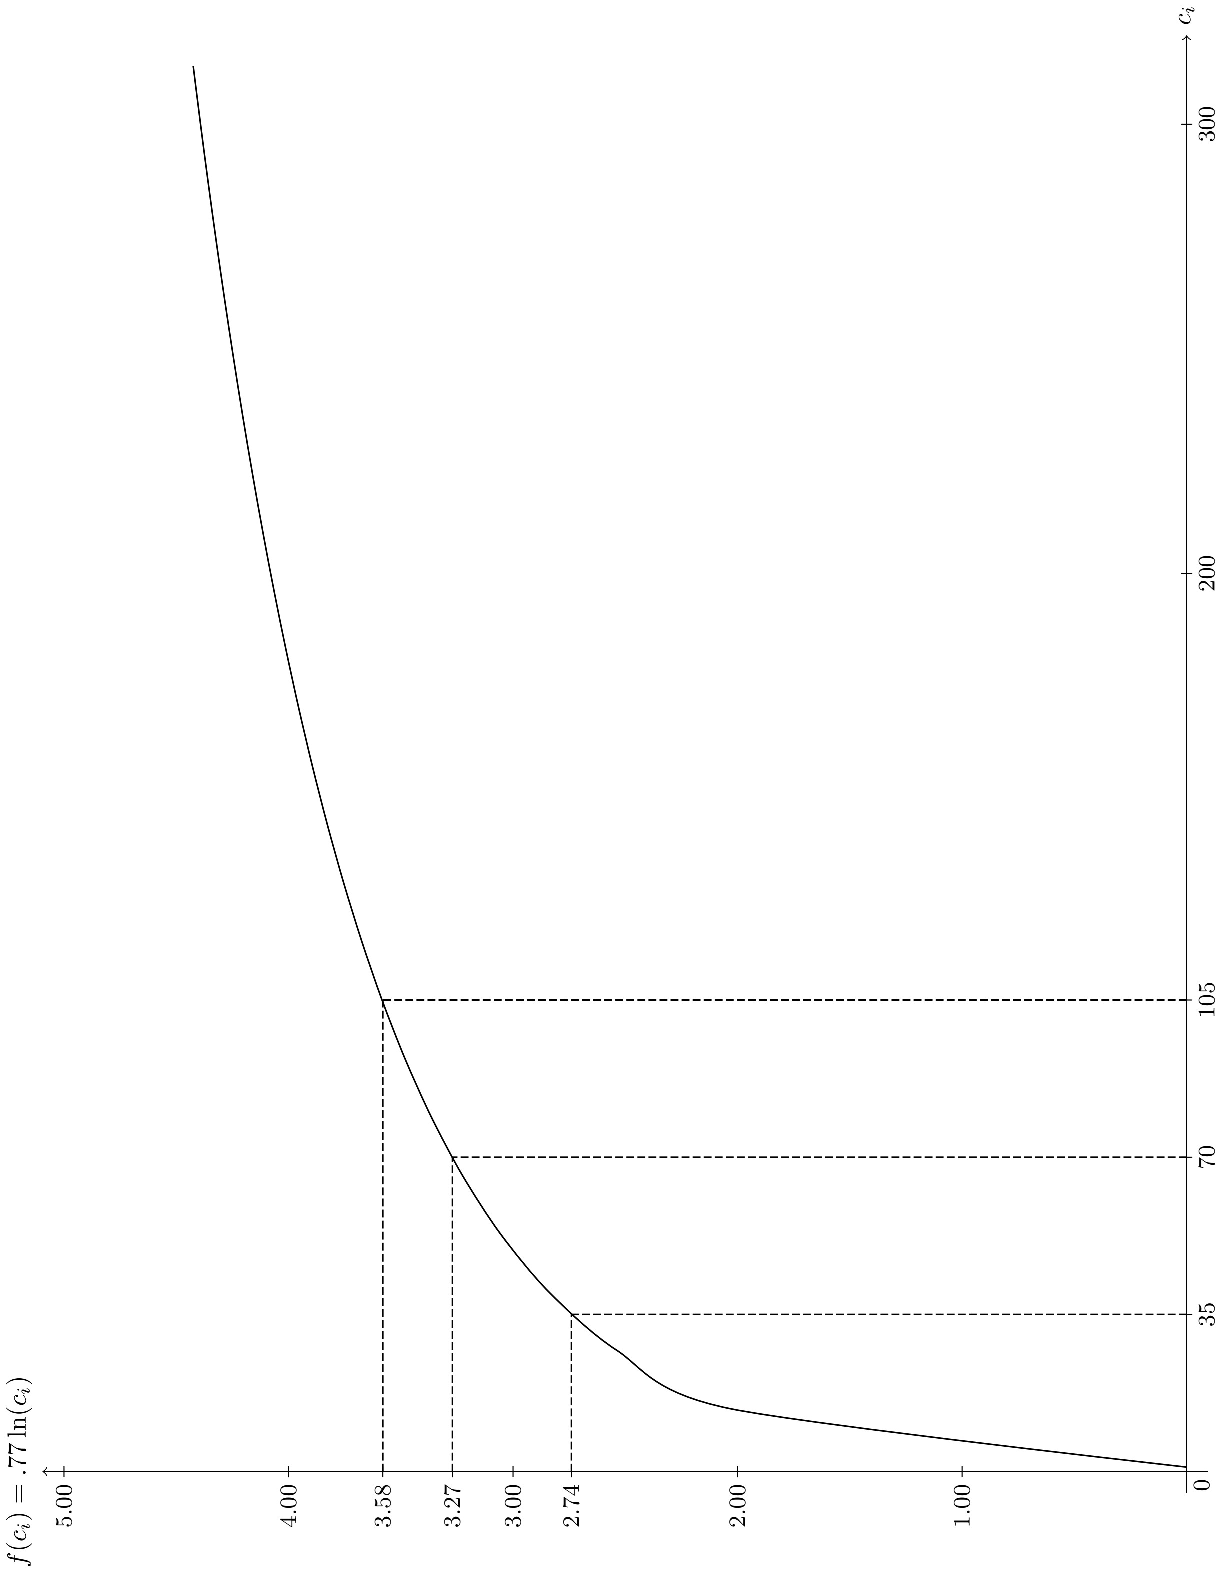

**Table B.1:** A list of all possible integer values for $c_i$ and $f(c_i) = .77\ln(c_i)$

| | | | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| $c_i$ | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 |
| $f(c_i)$ | 0.00 | 0.53 | 0.85 | 1.07 | 1.24 | 1.38 | 1.50 | 1.60 | 1.69 | 1.77 | 1.85 | 1.91 | 1.98 | 2.03 | 2.09 | 2.13 | 2.18 | 2.23 | 2.27 | 2.31 |
| $c_i$ | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 | 32 | 33 | 34 | 35 | 36 | 37 | 38 | 39 | 40 |
| $f(c_i)$ | 2.34 | 2.38 | 2.41 | 2.45 | 2.48 | 2.51 | 2.54 | 2.57 | 2.59 | 2.62 | 2.64 | 2.67 | 2.69 | 2.72 | 2.74 | 2.76 | 2.78 | 2.80 | 2.82 | 2.84 |
| $c_i$ | 41 | 42 | 43 | 44 | 45 | 46 | 47 | 48 | 49 | 50 | 51 | 52 | 53 | 54 | 55 | 56 | 57 | 58 | 59 | 60 |
| $f(c_i)$ | 2.86 | 2.88 | 2.90 | 2.91 | 2.93 | 2.95 | 2.96 | 2.98 | 3.00 | 3.01 | 3.03 | 3.04 | 3.06 | 3.07 | 3.09 | 3.10 | 3.11 | 3.13 | 3.14 | 3.15 |
| $c_i$ | 61 | 62 | 63 | 64 | 65 | 66 | 67 | 68 | 69 | 70 | 71 | 72 | 73 | 74 | 75 | 76 | 77 | 78 | 79 | 80 |
| $f(c_i)$ | 3.17 | 3.18 | 3.19 | 3.20 | 3.21 | 3.23 | 3.24 | 3.25 | 3.26 | 3.27 | 3.28 | 3.29 | 3.30 | 3.31 | 3.32 | 3.33 | 3.34 | 3.35 | 3.36 | 3.37 |
| $c_i$ | 81 | 82 | 83 | 84 | 85 | 86 | 87 | 88 | 89 | 90 | 91 | 92 | 93 | 94 | 95 | 96 | 97 | 98 | 99 | 100 |
| $f(c_i)$ | 3.38 | 3.39 | 3.40 | 3.41 | 3.42 | 3.43 | 3.44 | 3.45 | 3.46 | 3.46 | 3.47 | 3.48 | 3.49 | 3.50 | 3.51 | 3.51 | 3.52 | 3.53 | 3.54 | 3.55 |
| $c_i$ | 101 | 102 | 103 | 104 | 105 | 106 | 107 | 108 | 109 | 110 | 111 | 112 | 113 | 114 | 115 | 116 | 117 | 118 | 119 | 120 |
| $f(c_i)$ | 3.55 | 3.56 | 3.57 | 3.58 | 3.58 | 3.59 | 3.60 | 3.61 | 3.61 | 3.62 | 3.63 | 3.63 | 3.64 | 3.65 | 3.65 | 3.66 | 3.67 | 3.67 | 3.68 | 3.69 |
| $c_i$ | 121 | 122 | 123 | 124 | 125 | 126 | 127 | 128 | 129 | 130 | 131 | 132 | 133 | 134 | 135 | 136 | 137 | 138 | 139 | 140 |
| $f(c_i)$ | 3.69 | 3.70 | 3.71 | 3.71 | 3.72 | 3.72 | 3.73 | 3.74 | 3.74 | 3.75 | 3.75 | 3.76 | 3.77 | 3.77 | 3.78 | 3.78 | 3.79 | 3.79 | 3.80 | 3.81 |
| $c_i$ | 141 | 142 | 143 | 144 | 145 | 146 | 147 | 148 | 149 | 150 | 151 | 152 | 153 | 154 | 155 | 156 | 157 | 158 | 159 | 160 |
| $f(c_i)$ | 3.81 | 3.82 | 3.82 | 3.83 | 3.83 | 3.84 | 3.85 | 3.85 | 3.85 | 3.86 | 3.86 | 3.87 | 3.87 | 3.88 | 3.88 | 3.89 | 3.89 | 3.90 | 3.90 | 3.91 |
| $c_i$ | 161 | 162 | 163 | 164 | 165 | 166 | 167 | 168 | 169 | 170 | 171 | 172 | 173 | 174 | 175 | 176 | 177 | 178 | 179 | 180 |
| $f(c_i)$ | 3.91 | 3.92 | 3.92 | 3.93 | 3.93 | 3.94 | 3.94 | 3.95 | 3.95 | 3.95 | 3.96 | 3.96 | 3.97 | 3.97 | 3.98 | 3.98 | 3.99 | 3.99 | 3.99 | 4.00 |
| $c_i$ | 181 | 182 | 183 | 184 | 185 | 186 | 187 | 188 | 189 | 190 | 191 | 192 | 193 | 194 | 195 | 196 | 197 | 198 | 199 | 200 |
| $f(c_i)$ | 4.00 | 4.01 | 4.01 | 4.02 | 4.02 | 4.02 | 4.03 | 4.03 | 4.04 | 4.04 | 4.04 | 4.05 | 4.05 | 4.06 | 4.06 | 4.06 | 4.07 | 4.07 | 4.08 | 4.08 |
| $c_i$ | 201 | 202 | 203 | 204 | 205 | 206 | 207 | 208 | 209 | 210 | 211 | 212 | 213 | 214 | 215 | 216 | 217 | 218 | 219 | 220 |
| $f(c_i)$ | 4.08 | 4.09 | 4.09 | 4.09 | 4.10 | 4.10 | 4.11 | 4.11 | 4.11 | 4.12 | 4.12 | 4.12 | 4.13 | 4.13 | 4.14 | 4.14 | 4.14 | 4.15 | 4.15 | 4.15 |
| $c_i$ | 221 | 222 | 223 | 224 | 225 | 226 | 227 | 228 | 229 | 230 | 231 | 232 | 233 | 234 | 235 | 236 | 237 | 238 | 239 | 240 |
| $f(c_i)$ | 4.16 | 4.16 | 4.16 | 4.17 | 4.17 | 4.17 | 4.18 | 4.18 | 4.18 | 4.19 | 4.19 | 4.19 | 4.20 | 4.20 | 4.20 | 4.21 | 4.21 | 4.21 | 4.22 | 4.22 |
| $c_i$ | 241 | 242 | 243 | 244 | 245 | 246 | 247 | 248 | 249 | 250 | 251 | 252 | 253 | 254 | 255 | 256 | 257 | 258 | 259 | 260 |
| $f(c_i)$ | 4.22 | 4.23 | 4.23 | 4.23 | 4.24 | 4.24 | 4.24 | 4.25 | 4.25 | 4.25 | 4.25 | 4.26 | 4.26 | 4.26 | 4.27 | 4.27 | 4.27 | 4.28 | 4.28 | 4.28 |
| $c_i$ | 261 | 262 | 263 | 264 | 265 | 266 | 267 | 268 | 269 | 270 | 271 | 272 | 273 | 274 | 275 | 276 | 277 | 278 | 279 | 280 |
| $f(c_i)$ | 4.28 | 4.29 | 4.29 | 4.29 | 4.30 | 4.30 | 4.30 | 4.31 | 4.31 | 4.31 | 4.31 | 4.32 | 4.32 | 4.32 | 4.32 | 4.33 | 4.33 | 4.33 | 4.34 | 4.34 |
| $c_i$ | 281 | 282 | 283 | 284 | 285 | 286 | 287 | 288 | 289 | 290 | 291 | 292 | 293 | 294 | 295 | 296 | 297 | 298 | 299 | 300 |
| $f(c_i)$ | 4.34 | 4.34 | 4.35 | 4.35 | 4.35 | 4.36 | 4.36 | 4.36 | 4.36 | 4.37 | 4.37 | 4.37 | 4.37 | 4.38 | 4.38 | 4.38 | 4.38 | 4.39 | 4.39 | 4.39 |
| $c_i$ | 301 | 302 | 303 | 304 | 305 | 306 | 307 | 308 | 309 | 310 | 311 | 312 | 313 | | | | | | | |
| $f(c_i)$ | 4.39 | 4.40 | 4.40 | 4.40 | 4.40 | 4.41 | 4.41 | 4.41 | 4.41 | 4.42 | 4.42 | 4.42 | 4.42 | | | | | | | |

### B.3 Control questions (with correct answers in bold)

## Quiz

Before we start the first part of today's experiment we ask you to answer the following quiz questions that are intended to check your comprehension of the instructions. The numbers in these quiz questions are illustrative; the actual numbers in the experiment may be quite different. Before starting the experiment we will review each participant's answers. If there are any incorrect answers we will go over the relevant part of the instructions again.

1. True or False: I am limited from borrowing more than I can possibly repay.
   Circle one: **True** False

2. Suppose your income in period 1 is 65 and you choose to consume 75 points.

   a. What is your payoff in dollars for period 1? **$3.32**

   b. Did you: **borrow** save [circle one] in period 1? If so, how much? **10 points**

   c. If your income in period 2 is 71, what is the maximum amount you can consume in period 2?
   **UncHigh: 95 points, UncLow: 120 points, Con: 61 points**

3. Suppose your income in period 2 is 74 and that you have borrowed 25 in period 1.

   a. What is the minimum amount you can consume in period 2?
   **1 point**

   b. What is the maximum amount you can consume in period 2?
   **UncHigh: 83 points, UncLow: 108 points, Con: 49 points**

4. Suppose your income in period 3 is 67 and your savings from period 2 is 6.

   a. What is your consumption in period 3? **73 points**

   b. What is your payoff in dollars for period 3? **$3.30**

5. Suppose in a round you consumed $c_1 = 67$, $c_2 = 96$ and $c_3 = 75$. What is your payoff in dollars for the round? **$10.07**

6. Suppose in a round you consumed $c_1 = 79$, $c_2 = 80$ and $c_3 = 79$. What is your payoff in dollars for the round? **$10.09**

## B.4 Screenshots

**Figure B.2:** Example screen of period 1 in the UNCHIGH treatment

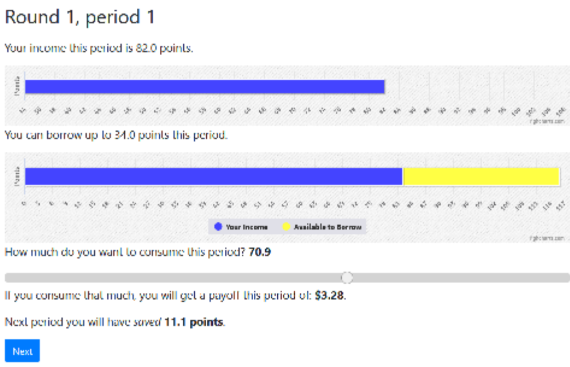


**Figure B.3:** Example screen of period 2 in the UNCHIGH treatment

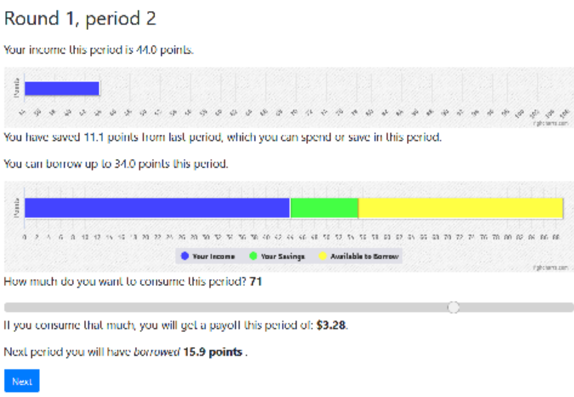

**Figure B.4:** Example screen of period 3 in the UNCHIGH treatment

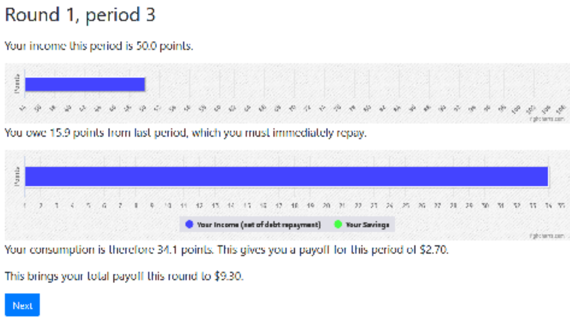


**Figure B.5:** Example screen of period 1 in the CONHIGH treatment

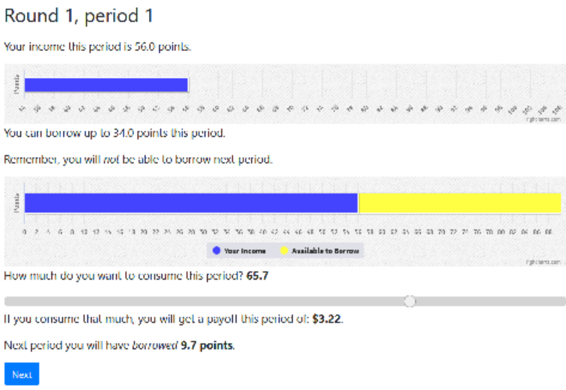


**Figure B.6:** Example screen of period 2 in the CONHIGH treatment

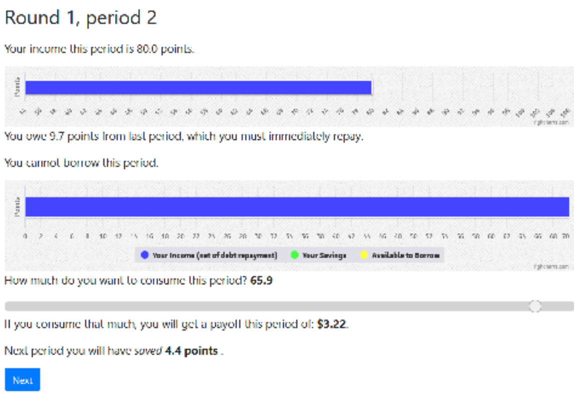

**Figure B.7:** Example screen of period 3 in the ConHigh treatment

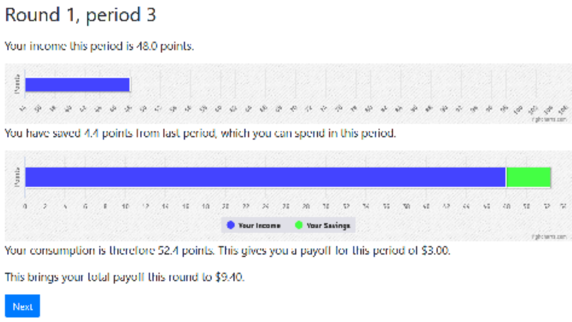


**Figure B.8:** Example screen of the Tower of Hanoi task

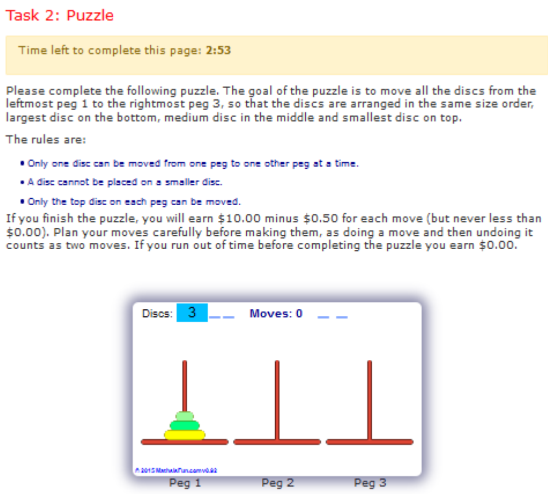

**Figure B.9:** Example screen of the risk aversion elicitation

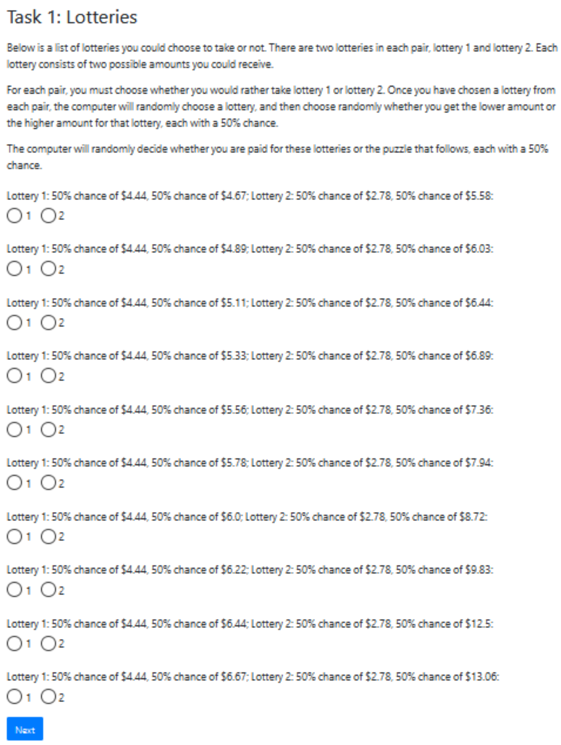

Task 1: Lotteries

Below is a list of lotteries you could choose to take or not. There are two lotteries in each pair, lottery 1 and lottery 2. Each lottery consists of two possible amounts you could receive.

For each pair, you must choose whether you would rather take lottery 1 or lottery 2. Once you have chosen a lottery from each pair, the computer will randomly choose a lottery, and then choose randomly whether you get the lower amount or the higher amount for that lottery, each with a 50% chance.

The computer will randomly decide whether you are paid for these lotteries or the puzzle that follows, each with a 50% chance.

Lottery 1: 50% chance of $4.44, 50% chance of $4.67; Lottery 2: 50% chance of $2.78, 50% chance of $5.58:

1 2

Lottery 1: 50% chance of $4.44, 50% chance of $4.89; Lottery 2: 50% chance of $2.78, 50% chance of $6.03:

1 2

Lottery 1: 50% chance of $4.44, 50% chance of $5.11; Lottery 2: 50% chance of $2.78, 50% chance of $6.44:

1 2

Lottery 1: 50% chance of $4.44, 50% chance of $5.33; Lottery 2: 50% chance of $2.78, 50% chance of $6.89:

1 2

Lottery 1: 50% chance of $4.44, 50% chance of $5.56; Lottery 2: 50% chance of $2.78, 50% chance of $7.36:

1 2

Lottery 1: 50% chance of $4.44, 50% chance of $5.78; Lottery 2: 50% chance of $2.78, 50% chance of $7.94:

1 2

Lottery 1: 50% chance of $4.44, 50% chance of $6.0; Lottery 2: 50% chance of $2.78, 50% chance of $8.72:

1 2

Lottery 1: 50% chance of $4.44, 50% chance of $6.22; Lottery 2: 50% chance of $2.78, 50% chance of $9.83:

1 2

Lottery 1: 50% chance of $4.44, 50% chance of $6.44; Lottery 2: 50% chance of $2.78, 50% chance of $12.5:

1 2

Lottery 1: 50% chance of $4.44, 50% chance of $6.67; Lottery 2: 50% chance of $2.78, 50% chance of $13.06:

1 2

Next