International Economic Review

ORIGINAL ARTICLE OPEN ACCESS

Liquidity Constraints, Income Variance, and Buffer Stock Savings: Experimental Evidence

John Duffy[1,2] | Andreas Orland[3]

[1]Department of Economics, University of California, Irvine, Irvine, California, USA | [2]Institute of Social and Economic Research, Osaka University, Ibaraki, Osaka, Japan | [3]Institute of Economics, Corvinus University of Budapest, Budapest, Hungary

Correspondence: John Duffy (duffy@uci.edu)

Received: 15 May 2023 | **Revised:** 27 September 2024 | **Accepted:** 11 December 2024

Funding: Funding for this project was provided by the UC Irvine School of Social Sciences.

Keywords: buffer stock model | consumption | debt aversion | experimental economics | income variance | intertemporal optimization | liquidity constraints | savings

ABSTRACT

We test the buffer stock model of savings behavior using a three-period intertemporal model. In one treatment, liquidity in the second period is constrained (borrowing not possible), while the unconstrained treatment has no such constraint. The buffer stock model predicts that a second-period liquidity constraint increases first-period savings. We also vary the variance of stochastic income (high or low) in a 2×2 design. While we find no evidence for the predicted liquidity constraint effect, most other predictions hold, for example, income variance effects. Observed departures can be explained by some combination of debt aversion, cognitive heterogeneity, and/or learning.

JEL Classification: C91, D15, D81, D91, E21

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 individuals facing income uncertainty target a certain amount of cash-on-hand (i.e., wealth plus disposable income) to permanent income and use savings as a buffer to smooth out income fluctuations.[1] When individuals are below their target for cash-on-hand, they save (consume less), and they borrow (consume more) when above it. In a world without borrowing constraints, where credit markets are perfect, the only explanation for such buffer stock savings is that individuals have "precautionary" motives in the face of idiosyncratic fluctuations in their income. Such precautionary motives arise naturally if individuals have certain preferences (e.g., constant relative risk aversion [CRRA]) where the marginal utility of consumption is convex so that individuals are "prudent" (Kimball 1990). However, buffer stock savings motives do not *require* prudent preferences. If credit markets are imperfect, so that individuals face liquidity (or borrowing) constraints, buffer stock savings behavior can also arise, even if individuals do not have prudent preferences. As determining the functional form that best characterizes individual preferences is difficult, the liquidity constraint explanation for buffer stock savings seems more generalizable for testing purposes. Our main contribution in this paper is to directly examine this liquidity constraint rationale for buffer stock savings behavior and thus provide a clear and direct empirical evaluation of the buffer stock model of precautionary savings.

Liquidity constraints are thought to be important for various macroeconomic policy questions such as the effectiveness of monetary policy and the magnitude of the fiscal multiplier. Yet the role played by liquidity constraints on savings behavior is difficult to evaluate in the field. We generally do not observe the precise constraints that individuals face, or the shocks to their income. For these reasons, we test the role of liquidity constraints for buffer stock savings in a controlled laboratory setting, using experiments with paid student subjects.[2] While several experiments have been conducted testing the lifecycle/permanent income theory of savings (see Hey and Dardanoni 1988, 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 the first to test whether the presence or absence of future anticipated liquidity constraints, specifically a no-borrowing constraint, affects savings behavior when individuals also face an uncertain income.

The empirical relevance of liquidity constraints is well established. Jappelli (1990, 220) estimates that 19% of US households are liquidity-constrained and discusses the characteristics of these households. Gross and Souleles (2002, 153) estimate the share of potentially liquidity-constrained households in the United States 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 are not restricted to those with little or no wealth; Boar et al. (2022) suggest that 82% of all US *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 whether savings behavior responds to a perfectly anticipated liquidity constraint, as might arise, for example, if a person left a job without having another one lined up.

Several empirical and simulation studies evaluate predictions of the buffer stock savings model. Most consider the buffer stock model *without* liquidity constraints, as such constraints are difficult to observe. On the one hand, some research supports the predictions of the buffer stock saving model. 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 supportive results from a model that includes both unemployment benefits and tax-sheltered retirement accounts; Jappelli and Pistaferri (2025) find that consumers adjust their *target* wealth approximately one-to-one with their permanent income. On the other hand, some research 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, and none do so in a controlled laboratory setting as we do. Zeldes (1989) compares models with and without liquidity constraints and concludes that liquidity constraints affect the consumption of a significant portion of the population; Haliassos and Michaelides (2003) examine the role of liquidity constraints in infinitely lived households' portfolio decisions between stocks and bonds; Campbell and Hercowitz (2018) set up a model where households consume both a standard good and a special good that is not consumed too often (e.g., college education or a new home) and report a good performance of their model.

Our experiment implements a simple, three-period model of intertemporal consumption and savings choices that is close to Besley (1995, 2141–2144) and further detailed in Jappelli and Pistaferri (2017, 115–118). Besley considers the case where individuals 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 consume all remaining resources.[3] In our setting, individuals face a liquidity constraint only in period 2 as in Jappelli and Pistaferri as this setup allows individuals 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 period 2 liquidity constraint case with the case where individuals are unconstrained in their period 2 borrowing. In both cases, the model predicts that individuals target a critical amount of cash-on-hand, defined by the model parameters, to determine their savings/borrowing decisions. If realized income is above this target level, individuals save, while if realized income is below the target level, they borrow. When a liquidity constraint is known to be binding in period 2, individuals should anticipate the impact of that constraint and adjust their savings and borrowing decisions in period 1. Thus, as our main contribution, we examine whether liquidity constraints affect saving and borrowing behavior in the predicted manner. As an additional contribution, we consider two different values for the variance of the uncertain income realizations, as theory predicts that savings are positively associated with the volatility of income. 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 2×2 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 experiment that compares buffer stock savings behavior in situations with and without liquidity constraints.[4]

To preview our results, we find mixed support for the predictions of the buffer stock savings model. In our analysis, subjects behave in line with most of the comparative statics predictions of the model. They are more deliberative in the choices they are making relative to random or heuristic decision rules. However, we do not find support for the main prediction 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 preregistered three behavioral explanations for this observed departure. The first and most important explanation is *debt aversion*. While the theoretical literature emphasizes the importance of *borrowing* constraints for saving decisions, Thaler

(1990, 202) long ago speculated that "another important source of liquidity constraints are self-imposed rules used by households who simply do not like to be in debt." This explanation for savings decisions has not received much attention in the buffer stock savings literature. A contribution of this paper is to show that debt aversion *is* an important factor and that it may greatly attenuate the role played by borrowing constraints in understanding savings decisions.[5]

Specifically, we find strong evidence for debt aversion both at the extensive margin (of whether to borrow or to save given a particular income realization) and at the intensive margin of how much to borrow or to save. This debt aversion is found in the first period of all treatments where there is never any borrowing constraint and is also found in the second period of the unconstrained treatments where there is also no borrowing constraint. We derive the savings predictions for a perfectly debt-averse individual and test the performance of this model relative to the constrained and unconstrained buffer stock savings models. We find that the debt-aversion model yields the best fit to our experimental data when the income variance is high. We also examine behavioral reference-dependent savings functions, and find that in all treatments and periods, the median decision exhibits partial debt aversion.

Second, we find that cognitive abilities also play a role; subjects who score better on a cognitive reflection test (CRT) generally behave closer to theoretical predictions. Finally, there is some evidence for learning over time; subjects reduce their deviations from theoretical predictions by about 13%–25% throughout 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 model that we test in our experiment. Specifically, we calculate the savings predictions for the unconstrained and liquidity-constrained models.[6] 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.[7] 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, μ. At the start of the first and the second periods, the individual first learns the income realization for that period. Then, they make a consumption decision for the period, which implies 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, the risk-free rate of interest is set to zero.[8] In the third period, the individual learns her income for that period and consumes all remaining resources, thus making no decisions. We assume that the individual evaluates per-period consumption using the period utility function $u(c) = \ln(c)$.[9] Figure 1 illustrates the timing of the model.

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 in period 1 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 apply backward induction and calculate the savings predictions in period 2. We start by rewriting the individual's two-period problem in period 2. As s_2 in Equation (5) is a function of previous savings s_1 and current income y_2, we substitute these two terms by wealth, $w_2 = y_2 + s_1$. For the *unconstrained* model, where the individual can save *and* borrow, the problem is given by

$$\max_{s_2} \ln(w_2 - s_2) + \mathbb{E}_{t=2} \ln(y_3 + s_2), \tag{6}$$

while, for the *constrained* model, it is

$$\max_{s_2 \geq 0} \ln(w_2 - s_2) + \mathbb{E}_{t=2} \ln(y_3 + s_2). \tag{7}$$

An analytical solution of these two problems requires us to form the first derivative of the respective problem with respect to savings s_2, and to form rational expectations of the (nonlinear) derivation of the log term.[10] We solve this problem numerically for s_2 for both models.[11] Figure 2b shows a schematic representation of optimal savings in period 2, depending on wealth w_2 (for the HIGH treatments; see the next section for the parameterization and other treatments). Without the constraint, savings increase (almost linearly and slightly convex) over the wealth range. With the constraint, savings are defined piecewise: We observe no borrowing for low wealth values: up to the point where unconstrained savings become positive, constrained savings are zero; after this point, constrained savings are identical to unconstrained savings.

When taking the savings decision in period 1, optimal savings in period 2 (with or without the constraint) are anticipated. Again, we solve this problem numerically. Figure 2a shows optimal savings in period 1, now depending on income y_1. Unconstrained savings increase (almost linearly and slightly convex) over the income range. Compared to the situation without the constraint, savings with the constraint are higher but increase less steeply

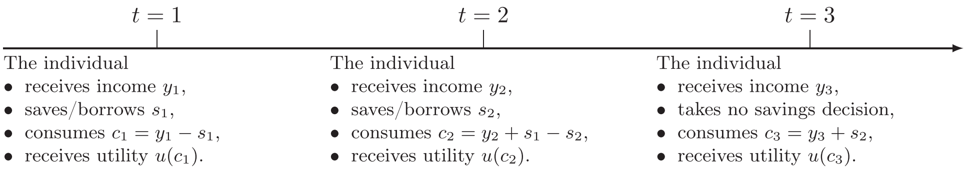

FIGURE 1 | Timing in the three-period lifecycle model.

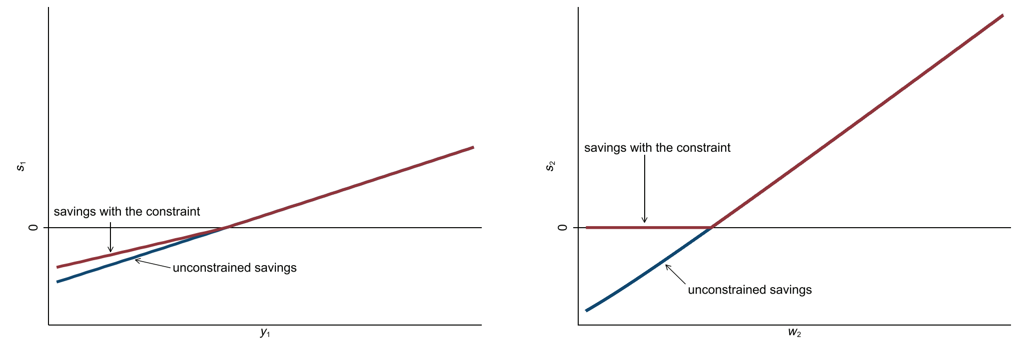

(a) Optimal savings in period 1; negative values indicate borrowing

(b) Optimal savings in period 2; negative values indicate borrowing

FIGURE 2 | Schematic representation of optimal savings.

up to the point where both savings predictions become positive. After that point, both constrained and unconstrained savings coincide.[12] In the next section, we introduce the parameterization of our experimental treatments and display the unconstrained and constrained savings predictions for those parameterizations.

3 | Experimental Design and Procedures

3.1 | Experimental Design

Our experiment employs a 2×2 design, where 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). The rationale for the latter treatment variable is that differences in the income variance affect savings predictions in periods 1 and 2 differently, as detailed in Section 4.

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 savings decision in a low variance environment). In the High treatments, we sample incomes uniformly from a discrete uniform distribution over integer values with a lower bound $y_{min} = 35$ and an upper bound $y_{max} = 105$ (with $\mu = 70$ and $\sigma = 20.5$). In the Low treatments, we sample integer income levels in the same manner, but where $y_{min} = 60$ 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).[13]

Using our parameterization of the model, Figure 3 presents the two models' savings predictions for all possible income levels in period 1 and for all possible wealth levels in period 2 in both the High and Low treatments. In Figure 3a,b, one can see that optimal savings in the constrained case, s_1^{CON}, is always greater than in the unconstrained case, s_1^{UNC}, for all period 1 income levels up to $y_1 \cong 63.5$ in High and for all period 1 income levels up to $y_1 \cong 69.7$ in Low. Figure 3c,d displays the optimal savings predictions for period 2. Here, the predictions are defined for the range of possible wealth levels (defined from zero up to twice the maximum income level, y_{max}, minus 1; 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^{CON}, and the unconstrained optimal period 2 savings, s_2^{UNC}, for wealth levels below $w_2 \cong 63.7$ in High, and for wealth levels below $w_2 \cong 69.5$ in Low, and no difference above it. By comparing the High and Low savings predictions in Figure 3, the role of income variance becomes clear: with a higher level of income uncertainty, saving (instead of borrowing) becomes optimal for smaller income and wealth levels in both periods. Later, we use the predictions shown here for the different treatments to analyze the experimental data both at the aggregate and individual levels.

In all four treatments, we induce the same per-period utility function $u(c) = 0.77\ln(c)$, transforming experimental consumption amounts in each period into money payoffs.[14] This utility function requires strictly positive consumption. As all loans are one-period, we impose natural saving and borrowing constraints (in addition to the liquidity constraint explicitly spelled out in the model). The natural *saving* constraints ensure that current consumption is strictly nonnegative in every period: $s_1 < y_1$ and $s_2 < w_2$. The natural *borrowing* constraints ensure nonnegative future consumption in every period: $s_1 > -y_{min}$ 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

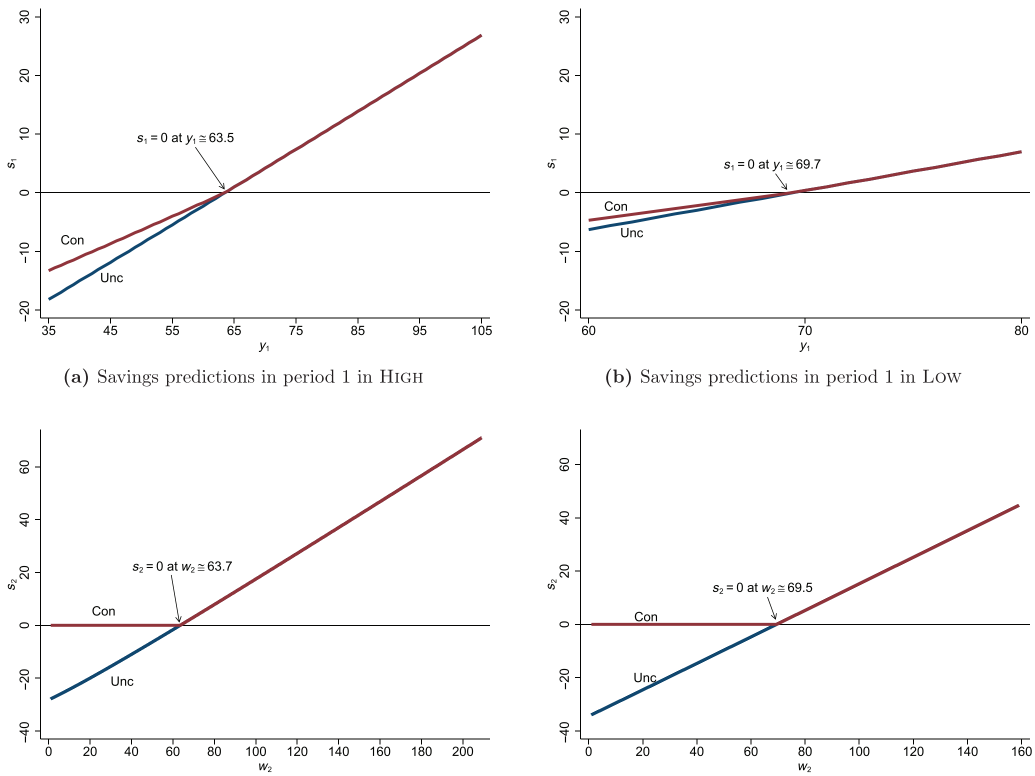

(a) Savings predictions in period 1 in HIGH

(b) Savings predictions in period 1 in LOW

(c) Savings predictions in period 2 in HIGH

(d) Savings predictions in period 2 in LOW

FIGURE 3 | The savings predictions in all experimental treatments.

constraints prevent subjects from losing money in any period of the experiment and do not affect the theoretical predictions in either the HIGH or LOW treatment parameterization.

3.2 | Experimental Procedures

We combine a within-subject design (where each subject completes 15, three-period lifecycles of *both* CON and UNC) with a between-subject approach (each subject was randomly assigned to either the HIGH or 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 we called lifecycles in the instructions) in the CON treatment and another one in the UNC treatment for payoff and this fact was known to subjects.

Given our experimental design and model parameterization, we considered the statistical power of hypothesis tests in choosing the number of subjects; details are provided in Section A.2 of the online Appendix). Based on these considerations, we recruited 100 subjects for the HIGH treatment (50 in the CON–UNC order, 50 in the UNC–CON order) and 50 subjects for 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. Our power analysis in Section A.2 of the online Appendix reveals that these numbers yield sufficient power for nearly all of our hypotheses, except for Hypotheses 1a–1b, where we would need a far greater number of subjects than our budget allows achieving the desired power level of 80%.

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 in 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 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) 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.3 in the online Appendix shows the 10 pairs of lotteries each subject had to choose between. (If the risk elicitation task was chosen for payoff, then one of the 10 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 to solve the game, the lower the subject's payoff.) Details on the ToH game can be found in Section A.6 of the online Appendix.

After the subjects had made all payoff-relevant decisions (but before feedback about their total dollar earnings was shown), they filled out a postexperimental questionnaire asking for their age, gender, field of study, and grade point average (GPA). They then completed an unincentivized CRT by Toplak et al. (2014). The CRT score (the total number 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. The CRT we use asks four questions (shown in Table A.4 in the online 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, so 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 received printed instructions, including 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 to proceed.[15] The experiment was computerized and programmed using oTree (Chen et al. 2016). Each session lasted about 90 min. After the experiment, all subjects were paid in cash and in private. We conducted 12 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. The 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 (the average for the savings-consumption task is \$9.62 for each of the two randomly chosen lifecycles; minimum \$5.10, maximum \$10.50).

4 | Aggregate Predictions and Hypotheses

In this section, we present aggregate predictions for all treatments and use these to form testable hypotheses. Table 1 displays means and standard deviations for optimal savings in periods 1 and 2 of all four treatments. These predictions are based on the savings predictions, and income and wealth ranges as shown in Figure 3. As we formed our hypotheses prior to data collection without knowing how subjects would behave in the experiment, and also, more specifically, how they would behave in period 1 (which then influences wealth in period 2), we chose the complete possible wealth range in period 2 (with wealth ranging from 1 to $2y_{max} - 1$) for the predictions.[16]

The first two sets of between-treatment, within-subject hypotheses (H1a–H1b and H2a–H2b) test the consequences of the liquidity constraint. 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 the anticipation of the liquidity constraint should increase their savings. Thus, in period 1, we are testing the additional prudence arising from the liquidity constraint in period 2, in addition to the prudence induced by the log utility function (which is the same in all treatments and periods).

Hypothesis 1a. *In the* CONHIGH *treatment, s_1 is higher than in the* UNCHIGH *treatment.*

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 examines 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 four between-treatment between-subject hypotheses examine how income variance affects savings behavior: higher-income uncertainty should cause higher savings.

Hypothesis 4. *In the* UNCHIGH *treatment, s_1 is higher than in the* UNCLOW *treatment.*

Hypothesis 5. *In the* CONHIGH *treatment, s_1 is higher than in the* CONLOW *treatment.*

Hypothesis 6. *In the* UNCHIGH *treatment, s_2 is higher than in the* UNCLOW *treatment.*

Hypothesis 7. *In the* CONHIGH *treatment, s_2 is higher than in the* CONLOW *treatment.*

5 | Results

5.1 | Comparative Statics

To give a complete picture of our data for all treatments and periods, we show our main aggregate-level findings in Table 2, which maps directly to Table 1. Table 2 reports mean savings (and standard errors) for the two periods (s_1, s_2) of all four treatments.

In a first step, we test the hypotheses formulated in Section 4 using one-sided t-tests to determine whether the difference between savings in the compared treatments/periods has the predicted

TABLE 1 | Aggregate predictions and hypotheses.

		HIGH UNC	HIGH CON				LOW UNC	LOW CON	
		H4							
			H5						
		H1a	H1a				H1b	H1b	
s_1^*	H3a	4.2 (13.2)	5.2 (12.0)	H3b	s_1^*	H3c	0.4 (4.0)	0.8 (3.6)	H3d
s_2^*	H3a	20.4 (28.9)	24.7 (23.6)	H3b	s_2^*	H3c	5.3 (22.7)	12.7 (14.7)	H3d
		H2a	H2a				H2b	H2b	
		H6							
			H7						

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).

TABLE 2 | Aggregate results.

		HIGH UNC	HIGH CON				LOW UNC	LOW CON	
		H4: ✓✓✓, **							
			H5: ✓✓✓, ***						
		H1a: ✗, x	H1a: ✗, x				H1b: ✗, x	H1b: ✗, x	
s_1	H3a: ✓✓, *	8.62 (1.66)	6.97 (1.76)	H3b: ✓✓✓, ***	s_1	H3c: ✗, x	1.42 (1.64)	−0.55 (1.51)	H3d: ✓✓✓, ***
s_2	H3a: ✓✓, *	12.34 (2.61)	19.69 (2.52)	H3b: ✓✓✓, ***	s_2	H3c: ✗, x	4.53 (2.75)	8.99 (1.91)	H3d: ✓✓✓, ***
		H2a: ✓✓✓, ***	H2a: ✓✓✓, ***				H2b: ✓✓, ***	H2b: ✓✓, ***	
		H6: ✓, x							
			H7: ✓✓✓, *						

Note: Means with standard errors clustered at the subject level in parentheses. ✓✓✓, ✓✓, and ✓ mark hypotheses where we reject that the savings difference between treatments/periods is less than or equal to zero at the 1%, 5%, and 10% level, respectively (based on one-sided *t*-tests clustered at the subject level). ***, **, and * show the significance of equivalent sign rank and rank sum tests at the 1%, 5%, and 10% level, respectively (based on the R package `clusrank` by Jiang et al. 2020). ✗ (x) marks hypotheses where we do not find evidence that the difference is less than or equal to zero based on the *t*-tests (based on the sign rank and rank sum tests). *p*-Values are corrected for multiple hypotheses tests (details in Table 3).

sign (where we adjust the *p*-values for multiple hypothesis testing using a Bonferroni correction for three comparisons). Most importantly, we do not find evidence for the main hypothesis set H1a–H1b. In both HIGH and LOW, observed period 1 savings in CON is *smaller* than in UNC. Thus, the anticipation of a liquidity constraint in the second period does not increase first-period savings compared to the setting without a liquidity constraint. As noted earlier, these are our only underpowered hypotheses, so a null result here might not be surprising. We *cannot* reject hypotheses set H2a–H2b, that in both treatments, HIGH and LOW, period 2 savings in CON are greater than period 2 savings in UNC.

As a robustness check, in Table A.5 of the online Appendix, we also report results for hypothesis sets H1a–H1b and H2a–H2b, separately for income/wealth levels where either borrowing or savings is predicted. Period 1 savings in CON are not significantly different from UNC savings in both ranges. Period 2 savings in CON are not significantly different from UNC savings only in the borrowing range, but they are significantly different in the savings range. The last two findings regarding period 2 savings are in line with theoretical predictions.[17]

Beyond hypotheses H1a–H1b and H2a–H2b, we cannot reject all other hypotheses; that is, we reject the null of no difference in savings between the treatments/periods in favor of theoretical directional predictions with the sole exception of hypothesis H3c.[18]

Table 3 provides the more detailed findings. There, we also report more conservative nonparametric test results as robustness checks: We use clustered signed-rank tests for hypotheses 1a–3d

TABLE 3 | Hypotheses tests.

Hypothesis				Difference of mean savings between treatments/periods		Significance of difference of means	Significance of paired differences
Result 1a	s_1^{ConHigh}	$>$	s_1^{UncHigh}	$s_1^{\text{ConHigh}} - s_1^{\text{UncHigh}} =$	−1.65	$p = 1.000$	$p = 1.000$
Result 1b	s_1^{ConLow}	$>$	s_1^{UncLow}	$s_1^{\text{ConLow}} - s_1^{\text{UncLow}} =$	−1.97	$p = 1.000$	$p = 1.000$
Result 2a	s_2^{ConHigh}	$>$	s_2^{UncHigh}	$s_2^{\text{ConHigh}} - s_2^{\text{UncHigh}} =$	7.35	$p = 0.002$	$p < 0.001$
Result 2b	s_2^{ConLow}	$>$	s_2^{UncLow}	$s_2^{\text{ConLow}} - s_2^{\text{UncLow}} =$	4.46	$p = 0.019$	$p < 0.001$
Result 3a	s_2^{UncHigh}	$>$	s_1^{UncHigh}	$s_2^{\text{UncHigh}} - s_1^{\text{UncHigh}} =$	3.72	$p = 0.017$	$p = 0.091$
Result 3b	s_2^{ConHigh}	$>$	s_1^{ConHigh}	$s_2^{\text{ConHigh}} - s_1^{\text{ConHigh}} =$	12.72	$p < 0.001$	$p < 0.001$
Result 3c	s_2^{UncLow}	$>$	s_1^{UncLow}	$s_2^{\text{UncLow}} - s_1^{\text{UncLow}} =$	3.11	$p = 0.125$	$p = 0.329$
Result 3d	s_2^{ConLow}	$>$	s_1^{ConLow}	$s_2^{\text{ConLow}} - s_1^{\text{ConLow}} =$	9.54	$p < 0.001$	$p < 0.001$
Result 4	s_1^{UncHigh}	$>$	s_1^{UncLow}	$s_1^{\text{UncHigh}} - s_1^{\text{UncLow}} =$	7.21	$p = 0.004$	$p = 0.011$
Result 5	s_1^{ConHigh}	$>$	s_1^{ConLow}	$s_1^{\text{ConHigh}} - s_1^{\text{ConLow}} =$	7.52	$p = 0.002$	$p = 0.007$
Result 6	s_2^{UncHigh}	$>$	s_2^{UncLow}	$s_2^{\text{UncHigh}} - s_2^{\text{UncLow}} =$	7.81	$p = 0.060$	$p = 0.127$
Result 7	s_2^{ConHigh}	$>$	s_2^{ConLow}	$s_2^{\text{ConHigh}} - s_2^{\text{ConLow}} =$	10.70	$p = 0.001$	$p = 0.096$

Note: The p-values of the difference of means are based on (one-sided) t-tests that the observed difference of means is less than or equal to zero. The p-values of the paired differences are based on (one-sided) signed-rank and rank-sum tests (based on the R package `clusrank` by Jiang et al. 2020). All tests are clustered at the subject level, and adjusted for three comparisons using a Bonferroni correction.

and clustered rank-sum tests for hypotheses 4–7. The results confirm our findings using t-tests; only Result 6 is insignificant. As noted earlier, our test of hypothesis set H1a–1b is underpowered given our sample size, though all other hypotheses are sufficiently powered. The reader should, therefore, condition on these facts in evaluating our empirical findings.

In a second step, we compare observed savings with predictions (this cannot be done by just comparing predicted savings in Table 1 and observed mean savings in Table 2 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 income realizations in that period). Table 4 shows the results of our analysis of *conditional* savings behavior by comparing it to optimal savings, s^*. In the first period, we observe significant oversaving in the UncHigh treatment and significantly higher oversaving in UncLow compared to ConLow. In the second period, we observe significant *oversaving* in all treatments (with savings up to five times the predicted level).[19] Furthermore, deviations from predicted period 2 savings are larger than for period 1 in all treatments. As the wealth ranges in period 2 are larger than the income ranges in period 1 (and as period 1 oversaving in three treatments shifts actual wealth in period 2 to the right), this makes larger mistakes in period 2 possible (also see Figures 4 and 5 to compare ranges and natural constraints). Finally, in the Con treatments, deviations in period 2 are larger than in the Unc treatments. This is a direct consequence of the constraint (due to which, for low wealth realizations, borrowing is not possible): here, subjects in Con can only take decisions that result in positive deviations (whereas in Unc, for low wealth realizations in period 2, subjects can also take decisions that result in negative deviations).[20]

In a third step, we test whether the observed behavior results from subjects' *deliberate* decision-making using three different heuristics. We first compare observed savings with the expected savings that would result if subjects randomly chose values from the allowed savings intervals, s^R.[21] Second, we test whether subjects made use of a boundedly rational consumption-smoothing heuristic, s^{H1}, to determine their savings. In particular, we consider the case where the individual simplifies the uncertainty about future income using $\mathbb{E}y_2$ and $\mathbb{E}y_3$ in place of y_2 and y_3 and then solves for s_1 and s_2 using only the budget constraints in Equations (2)–(4), and ignoring the induced utility function and the future liquidity constraint.[22] This nonoptimizing approach results in perfect consumption-smoothing across periods, $c_1 = c_2 = c_3 = c$. The third approach, s^{H2}, is similar to the second approach, except that we assume that agents ignore the presence of period 3 (and the future liquidity constraint in period 2), smoothing consumption only in the first two periods.[23] All three heuristics discussed in this section are shown for period 1 in online Appendix Figures A.4 (for UncHigh) and A.5 (for ConHigh). There one can see that expected random savings are higher than optimal savings for all income levels while the other two heuristics generally predict less than optimal savings for all income levels.

Table 4 reports on deviations and root-mean-square errors (RMSEs) of observed behavior from these three heuristics in all treatments and periods. We see that with one exception (period 2 of UncLow), all deviations from random choices, s^R, are significantly negative. Furthermore, the magnitude of all deviations from random behavior is larger than the corresponding deviations from optimal behavior s^*, with the sole exception of ConLow. In terms of RMSE, the data are generally closer to s^* than to s^R. The fit of the consumption smoothing heuristics s^{H1} and s^{H2}, are generally better than the fit of s^R in terms of the absolute deviations and RMSEs except for period 2 of UncLow. Except for period 1 of ConLow, the absolute deviations and RMSEs for s^{H1} and s^{H2} are somewhat worse than for the predicted savings, s^*. Overall, we conclude that by comparison with random choices and the consumption-smoothing heuristics, the predicted savings

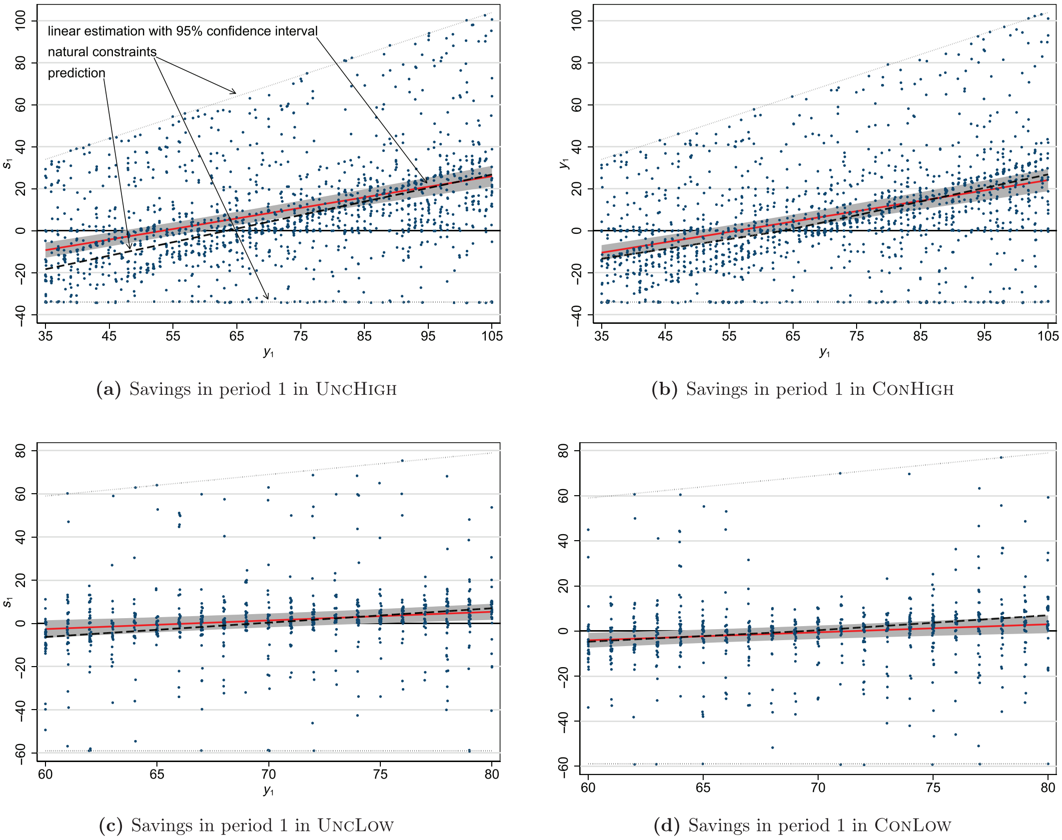

(a) Savings in period 1 in UncHigh

(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 by 0.5 points), estimated linear savings functions, and optimal predictions in period 1.

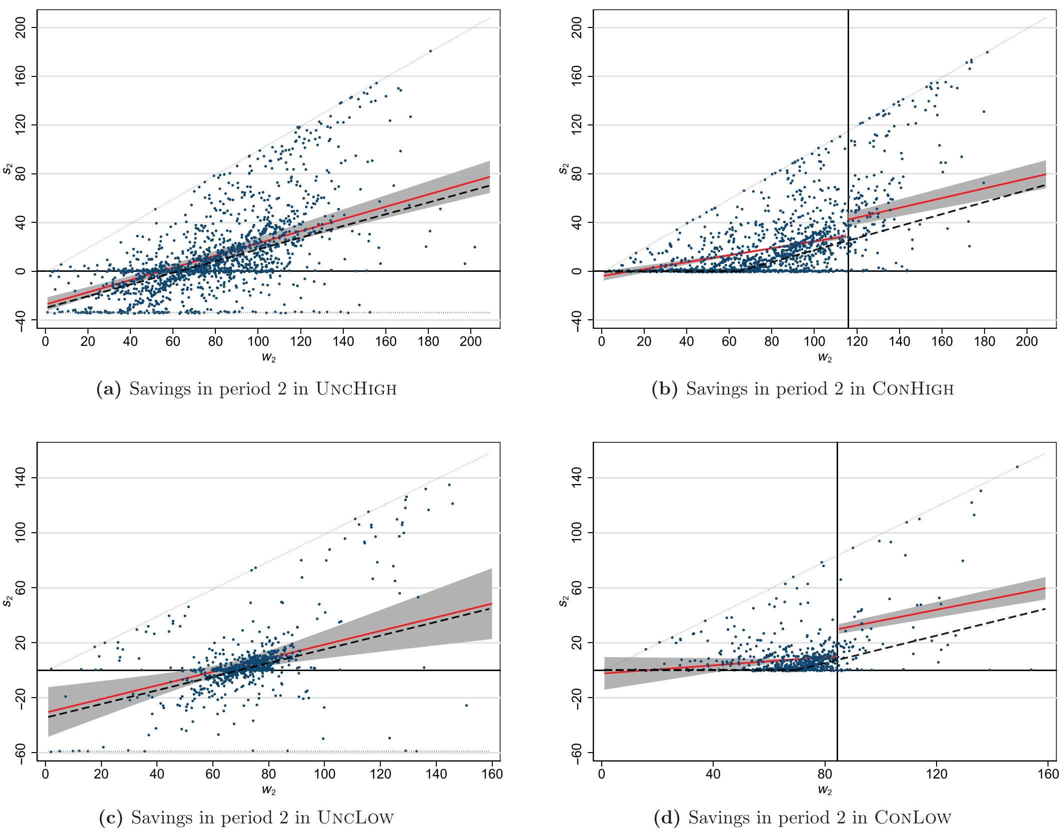

(a) Savings in period 2 in UncHigh

(b) Savings in period 2 in ConHigh

(c) Savings in period 2 in UncLow

(d) Savings in period 2 in ConLow

FIGURE 5 | Savings decisions (jittered by 0.5 points), linear (threshold) savings functions, and optimal predictions in period 2.

TABLE 4 | Deviations of observed (s) from optimal (s^*), and heuristic behavior (s^R, s^{H1}, and s^{H2}).

	HIGH		LOW	
	UNC	CON	UNC	CON
Period 1				
Deviation $s_1 - s_1^*(y_1)$	+ 4.07**	+ 1.53	+ 0.99	− 1.32
	(1.68)	(1.76)	(1.62)	(1.52)
	[23.17]	[22.81]	[17.39]	[15.87]
	difference=−2.54 ($p = 0.124$)		difference=−2.31** ($p = 0.047$)	
Deviation $s_1 - s_1^R(y_1)$	− 9.13***	−10.66***	− 3.64**	− 5.56***
	(1.68)	(1.76)	(1.62)	(1.52)
	[24.35]	[25.03]	[17.70]	[16.74]
	difference=−1.53 ($p = 0.353$)		difference=−1.92* ($p = 0.096$)	
Deviation $s_1 - s_1^{H1}(y_1)$	− 8.29***	− 6.80**	+ 1.34	+ 0.56
	(1.68)	(1.77)	(1.62)	(1.53)
	[24.34]	[23.97]	[17.41]	[15.91]
	difference=1.49 ($p = 0.367$)		difference=1.89 ($p = 0.101$)	
Deviation $s_1 - s_1^{H2}(y_1)$	+ 8.37***	+ 6.84***	+ 1.36	− 0.56
	(1.67)	(1.76)	(1.62)	(1.52)
	[24.07]	[23.66]	[17.37]	[15.80]
	difference=1.53 ($p = 0.353$)		difference=−1.92* ($p = 0.096$)	
Period 2				
Deviation $s_2 - s_2^*(y_2, s_1)$	+ 6.57***	+10.90***	+ 3.70*	+ 6.25***
	(2.00)	(2.00)	(2.17)	(1.61)
	[27.50]	[26.00]	[21.13]	[15.55]
	difference=4.34** ($p = 0.014$)		difference=2.52 ($p = 0.120$)	
Deviation $s_2 - s_2^R(y_2, s_1)$	− 9.65***	−18.24***	− 1.19	−25.16***
	(1.97)	(1.85)	(2.17)	(1.72)
	[28.23]	[29.50]	[20.85]	[29.50]
	difference=−8.59*** ($p = 0.000$)		difference=−23.97*** ($p = 0.000$)	
Deviation $s_2 - s_2^{H1}(y_2, s_1)$	+ 7.85***	+11.52***	+ 3.81*	+ 6.31***
	(1.97)	(1.96)	(2.17)	(1.61)
	[27.67]	[26.25]	[21.16]	[15.57]
	difference=3.67** ($p = 0.034$)		difference=2.50 ($p = 0.128$)	

Note: Deviations from optimal savings (conditioned on income realization) and conditionally optimal savings (conditioned on income realization and previous period's savings decision), s_1^* and s_2^*, are based on the predictions in Section 2. Deviations from random behavior, s_1^R and s_2^R, use the savings functions from Footnote 21. Deviations from heuristic behavior, s_1^{H1} and s_2^{H1}, use the savings functions from Footnote 22. Deviation from heuristic behavior ignoring period 3, s_1^{H2}, use the savings function from Footnote 23. Standard errors are clustered at the subject level in parentheses. Root-mean-square errors of s and either s^* or s^R in square brackets. ***, **, and * show the difference from zero (based on two-sided t-tests clustered at the subject level) at the 1%, 5%, and 10% level, respectively.

functions generally provide a closer fit to the observed data, providing evidence for deliberative behavior by subjects in our experiment.

In a fourth step, we test whether there is persistence in savings choices between lifecycles due to anchoring. Specifically, we regress period 1 savings: (i) on period 1 savings of the previous lifecycle and (ii) on period 1 savings of the previous lifecycle *and* current period 1 income. In Table 5, we report the results of fixed-effect regressions in two panels. Examining the upper and lower panel, we see that the previous lifecycle's period 1 savings decision does not significantly determine the current lifecycle's behavior as it might under anchoring. However, the lower panel shows the significant positive effect that current period 1 income has on current period 1 savings. The within-subject coefficients of determination R^2 confirm this interpretation: They are small for the first model (just a constant and the previous lifecycle's s_1 decision do not explain current savings behavior well). When we add current income, within-subject R^2 increases in all four treatments as the subjects react to current income (and save more with higher income; this increase of within-subject R^2 is stronger in HIGH). Finally, in Table 6, we report absolute (dollar) losses and the share of optimal dollar earnings achieved by the subjects. We report these per-period losses only for period 1 savings decisions (the cleanest test, since in period 1 there are no wealth effects from previous savings decisions; also, in period 1, the natural constraints in the UNC and the CON treatments are identical while they differ—by design—in period 2) for all treatments. We see that absolute losses (shares of optimal) are slightly smaller (larger) in CON compared to UNC, and smaller (larger) in LOW compared to HIGH.[24] Our interpretation from these six tests is that the observed behavior of subjects involves deliberate decision-making based on (possibly biased) optimizing reasoning.

TABLE 5 | The impact of the previous-lifecycle's savings decision on current savings behavior in period 1 (with and without current income).

	UNCHIGH	CONHIGH	UNCLOW	CONLOW
Lagged s_1	−0.040	0.007	−0.019	−0.061
	(0.044)	(0.040)	(0.106)	(0.070)
Constant	8.830**	6.933**	1.292**	−0.667**
	(0.381)	(0.282)	(0.148)	(0.038)
#obs.	1450	1450	725	725
#clusters	100	100	50	50
R^2 (within)	0.002	0.000	0.000	0.004
R^2 (between)	0.999	0.989	0.992	0.985
R^2 (overall)	0.178	0.239	0.183	0.171
Lagged s_1	−0.016	0.010	−0.021	−0.058
	(0.043)	(0.036)	(0.101)	(0.070)
y_1	0.513**	0.488**	0.411**	0.332**
	(0.041)	(0.044)	(0.113)	(0.113)
Constant	−27.529**	−27.374**	−27.531**	−23.911**
	(2.894)	(3.036)	(7.940)	(7.913)
#obs.	1450	1450	725	725
#clusters	100	100	50	725
R^2 (within)	0.303	0.320	0.035	0.032
R^2 (between)	0.000	0.010	0.025	0.597
R^2 (overall)	0.151	0.165	0.011	0.005

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.

TABLE 6 | Absolute losses due to nonoptimal behavior and share of optimal in period 1.

	Dollar loss	Share of optimal
UNCHIGH	0.170	0.947
	(0.041)	(0.013)
CONHIGH	0.125	0.960
	(0.041)	(0.013)
UNCLOW	0.060	0.982
	(0.032)	(0.010)
CONLOW	0.021	0.994
	(0.020)	(0.008)

Note: Standard errors, clustered at the subject level, in parentheses.

5.2 | Estimated Savings Functions

In this section, we estimate savings *functions* based on subjects' decisions and compare those estimated savings functions with predictions. We first consider the savings functions for period 1. We estimate a period 1 savings function by regressing savings decisions s_1 by each subject i in lifecycle "round" on current income y_1 using a panel regression estimator (with u being the individual effect and e the disturbance term):

$$s_{1,i,\text{round}} = \text{constant} + \beta y_{1,i,\text{round}} + u_i + e_{i,\text{round}}. \quad (8)$$

Figure 4 shows the fitted values 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 income, y_1, as well as predicted intertemporally optimal period 1 savings, s^*. We observe the following: (i) considerable heterogeneity in observed savings decisions in all treatments, as the scatterplots make clear;[25] (ii) in the UNCHIGH treatment, we see that for low period 1 incomes the estimated period 1 savings function is significantly higher than predicted but becomes indistinguishable from the prediction for higher period 1 incomes; and (iii) in the CONHIGH treatment and the LOW treatments, the estimated and predicted period 1 savings overlap.

Next, we consider the savings behavior for period 2, s_2, in a manner similar to period 1, but where s_2 is a 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 5. In those same figures, we report linear estimated savings functions for the UNC treatments that are similar in specification to Equation (8) except that w_2 replaces y_1, and linear *threshold* estimated savings functions for the CON treatments.[26] Since period 2 savings predictions in the CON treatments are piecewise defined, 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, γ, at which period 2 savings behavior changes:

$$s_{2,i,\text{round}} = \text{constant} + w_{2,i,\text{round}}\left(w_{2,i,\text{round}} < \gamma\right)\beta_1 + w_{2,i,\text{round}}\left(w_{2,i,\text{round}} \geq \gamma\right)\beta_2 + u_i + e_{i,\text{round}}. \quad (9)$$

Thus, to estimate the period 2 savings function for the CON treatments, we use Equation (9) and estimate the slope β_1 (which is, in theory, equal to zero), before the estimated threshold γ (theoretically 63.7 in HIGH and 69.5 in LOW), and the slope β_2 (theoretically 0.492 in HIGH and 0.499 in LOW), after that estimated threshold. In the threshold regression approach, the threshold γ is *endogenously* determined by minimizing the residual sum of squares; u is the individual effect, and e is the disturbance term.

In Figure 5, 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 predictions for all period 2 wealth levels; and (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 63.7 (CONHIGH) and 69.5 (CONLOW); the estimated slope before the endogenously determined threshold is greater than the prediction of zero, while after the threshold the estimated slope is even greater than before the threshold, but less steep than the prediction for these period 2 wealth levels.

Table 7 shows the estimated coefficients from the savings functions displayed in Figures 4 and 5. As the savings predictions in the UNC treatments are nearly linear, the predictions shown here

TABLE 7 | Estimated linear (threshold) savings functions.

	UNCHIGH				CONHIGH				UNCLOW				CONLOW		
	Prediction		Estimation		Prediction		Estimation		Prediction		Estimation		Prediction		Estimation
y_1	0.644	>>>	0.503***	y_1	0.582	>>>	0.494***	y_1	0.665	>>>	0.399***	y_1	0.588	>>	0.355***
			(0.021)				(0.019)				(0.084)				(0.074)
Constant	−40.890	<<<	−26.860***	constant	−35.499	<<<	−27.731***	constant	−46.198	<<<	−26.597***	constant	−40.400	<<<	−25.390***
			(2.905)				(2.207)				(6.124)				(5.415)
#obs.			1500	#obs.			1500	#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
w_2	0.481		0.501***	$w_2 < 115.9$			0.285***	w_2	0.498		0.496***	$w_2 < 84.4$			0.150*
			(0.018)				(0.026)				(0.138)				(0.087)
				$w_2 \geq 115.9$	0.492	>>>	0.400***					$w_2 \geq 84.4$	0.499	>>>	0.257***
							(0.035)								(0.103)
Constant	−30.086		−27.261***	constant			−3.948***	constant	−34.568	<<<	−30.915***	constant			−2.434**
			(2.040)				(2.030)				(9.343)				(6.152)
#obs.			1500	#obs.			1500	#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: The upper panel shows regressions with s_1 as the dependent variable, the lower panel with s_2 as the dependent variable. 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 are derived from linear regressions of optimal savings on income/wealth (all $R^2 > 0.99$ of these regressions).

also result from linear regressions of optimal savings decisions on income.[27] 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 and coefficients are close to predictions (though the UNCLOW intercept is significantly greater than predicted). For the period 2 savings coefficients in the CON treatments, we only have predictions for the coefficients *above* the threshold (they lie entirely 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). In the estimated postthreshold region, these savings slope coefficients are significantly smaller than predicted.[28], [29], [30] Summarizing, we generally find that, relative to theoretical predictions in period 1, 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 UNCHIGH. 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 the model predicts, indicating a greater precautionary motive than predicted.

6 | Behavioral Explanations

In Table 3, we have seen that first-period savings in the constrained case are *smaller* than in the unconstrained case, contrary to theory. In this section, we consider three behavioral explanations that help to rationalize these departures from theoretical predictions. 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 the optimal choice. Consequently, the presence or absence of a constraint on borrowing 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 heterogeneity in cognitive abilities can explain the departures we observe from theoretical predictions. Finally, we consider whether repeated experience with the 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 behavior. With experience, their behavior might be more in line with theoretical predictions.

6.1 | Debt Aversion

To examine debt aversion as an explanation for the observed behavior, we first test whether subjects' savings 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 in period 1, even though they were always free to borrow in period 1 of our experiment, and we test that model against the optimal solution. Finally, we posit a reference-dependent savings model to assess the extent of subjects' debt aversion.

6.1.1 | Evidence From a Simple Binary Model

For this analysis, 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 is optimal to save, borrows if it is optimal to borrow, or does neither otherwise). It takes the value 0 otherwise.[31] With this variable, we only consider whether the *sign* of the savings decision is correct. That is, 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 later address the latter question in Section 6.1.2.

Table 8 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 63.1% of all decisions have the predicted sign. By contrast, in the saving range, between 68.5% and 84.9% of all decisions are classified as binary optimal. The differences between the shares in the two ranges are positive for all treatments and periods. They differ significantly 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 predicts they should save than when it predicts they should borrow.

Next, we consider whether and how the binary optimal classification changes with income/wealth values in all relevant treatments and periods. Figure 6 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.[32] We observe the same pattern in most treatments and periods except period 1 of CONLOW) in Figure 6d and period 2 of UNCLOW in Figure 6f): 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 make less accurate decisions, on average, when the income shock is closer to the cutoff. In the savings range, subjects are also less correct on average when the shock is closer to the cutoff, but overall, they take more correct decisions than in the borrowing range.

To answer the questions of whether we can quantify these differences locally around the cutoff and how the probability of behaving in a binary optimal fashion changes around the cutoff, we make use of a regression discontinuity analysis[33] of the data where the assignment to the borrowing or saving area is (partly) random and depends deterministically (due to the savings predictions) on the income realization.[34] Hence, we have a sharp discontinuity in the assignment to either the borrowing or saving range as a function of the running variables income and wealth (where our estimation uses a triangular kernel, which gives decreasing but nonzero weight to observations further from the cutoff).

TABLE 8 | Binary optimal behavior with tests on aggregate and local differences.

Treatment	Period	Share (conditionally) binary optimal		Aggregate difference	Local difference
		Should borrow	Should save		
UNCHIGH	1	55.6%	80.1%	24.5pp ($p < 0.001$)	16.7pp ($p = 0.117$) [−12.5; +12.5]
CONHIGH	1	59.5%	80.7%	21.2pp ($p = 0.001$)	24.8pp ($p = 0.046$) [−11.0; +11.0]
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 ($p = 0.101$)	8.4pp ($p = 0.612$) [−3.7; +3.7]
UNCHIGH	2	63.1%	84.5%	21.3pp ($p < 0.001$)	10.3pp ($p = 0.478$) [−8.6; +8.6]
UNCLOW	2	55.5%	84.9%	29.4pp ($p < 0.001$)	10.2pp ($p = 0.495$) [−4.0; +4.0]

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 a 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.

We estimate the average change of the probability that subjects behave in the binary optimal manner in the saving range compared to the borrowing range near the cutoff. The results are shown in the final column of Table 8. The local differences around the cutoff are positive for all treatments and periods. For CONHIGH and UNCLOW in period 1, we observe that the local differences are greater than the aggregate differences and significantly different from zero.[35]

6.1.2 | Evidence from Model Comparisons

In this section, we compare which of these three models best describes observed behavior: the unconstrained, the liquidity-constrained, or a behavioral, "debt aversion" model. For these comparisons, we derive the savings predictions 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 a constraint is present). Thus, the savings function in period 2 for a debt-averse individual is the same as for the constrained individual:

$$s_2^{\mathrm{DA}} = s_2^{\mathrm{CON}}. \tag{10}$$

Using the same backward induction procedure applied in Section 2, we derive the savings prediction in period 1 for a debt-averse individual that, in addition to $s_2 \geq 0$, also sets $s_1 \geq 0$ (thus, the individual has higher savings in period 1, compared to UNC and CON because due to her debt aversion, she does not borrow when the two models predict borrowing):

$$s_1^{\mathrm{DA}} = \begin{cases} s_1^{\mathrm{CON}} & \text{if } s_1^{\mathrm{CON}} \geq 0 \\ 0 & \text{otherwise.} \end{cases} \tag{11}$$

TABLE 9 | Model comparisons using root-mean-square errors.

	Model	Treatment			
		UNCHIGH	CONHIGH	UNCLOW	CONLOW
s_1	UNC	23.17	23.06	17.39	15.92
	CON	22.84	<u>22.81</u>	<u>17.36</u>	<u>15.87</u>
	DA	<u>22.64</u>	22.90	17.43	15.90
s_2	UNC	28.14	27.63	<u>21.13</u>	17.94
	CON	<u>27.50</u>	<u>26.00</u>	21.42	<u>15.55</u>

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

We now turn to which of the three models explains our experimental data best. Table 9 shows the RMSEs of the UNC, CON, and DA models relative to the data 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 conditioned on the previous savings decision s_1 and realized income y_2) when compared with the individual-level data from all four treatments.[36] For first-period savings, UNC decisions are best explained by the more debt-averse models DA and CON while CON decisions are best explained by the correct model. For second-period savings in LOW, the predicted model, CON or UNC, describes behavior best according to whether the liquidity constraint was binding or not; in HIGH, the CON model explains behavior best. 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. These findings confirm our results from the previous section: especially in the UNCHIGH treatment in the first period, subjects behave in a very debt-averse manner.

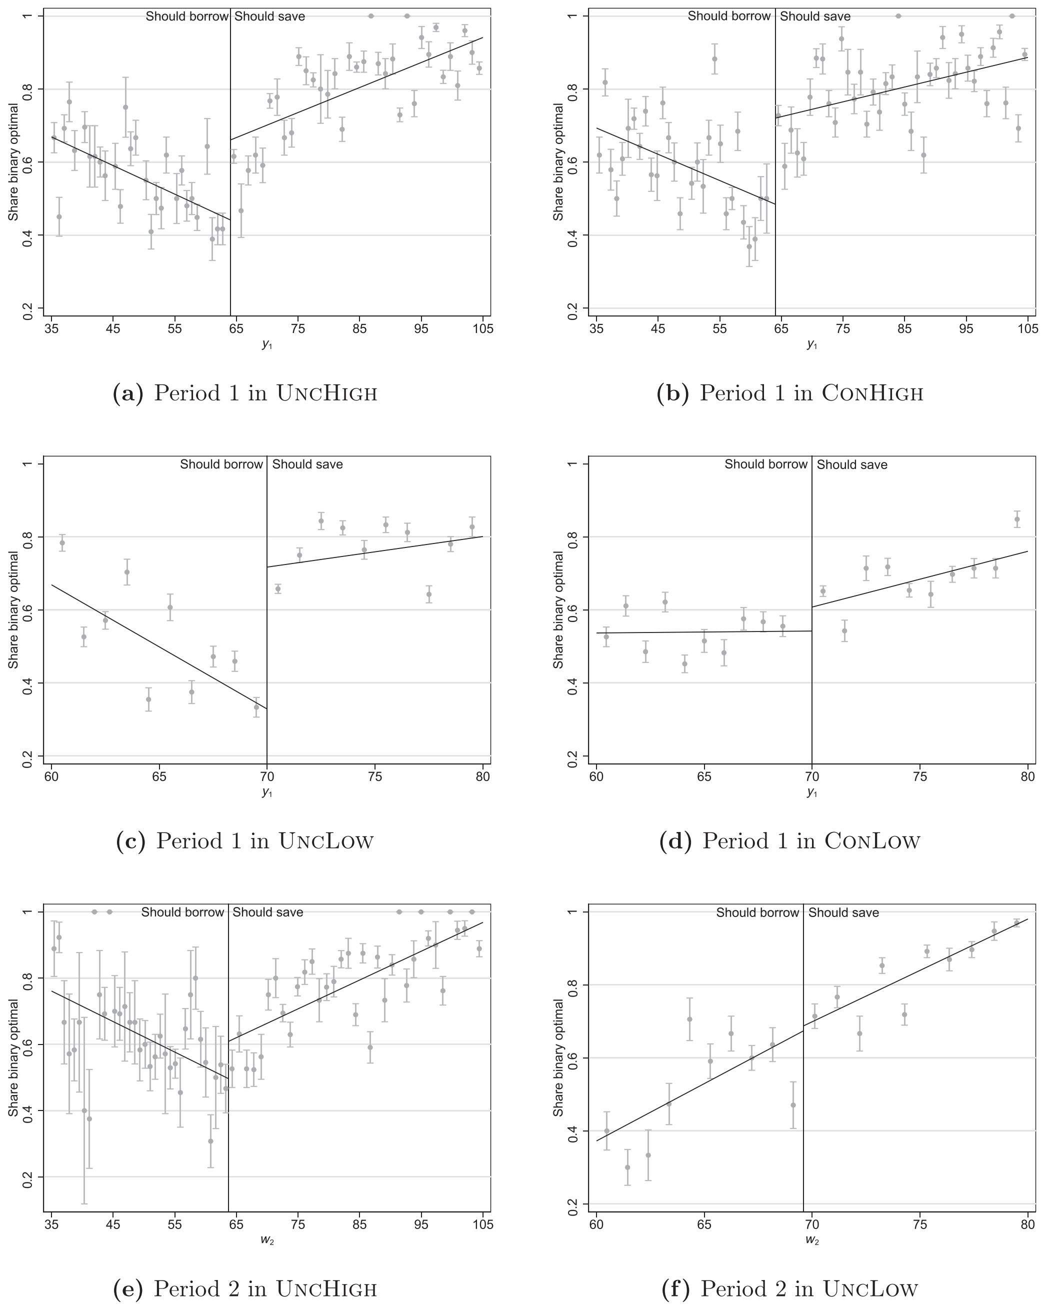

(a) Period 1 in UncHigh

(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

FIGURE 6 | Share of (conditionally) binary optimal decisions.

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 based on theoretical predictions. Graphs created with the Stata package `rdplot` by Calonico et al. (2017).

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 based on theoretical predictions. Graphs created with the Stata package `rdplot` by Calonico et al. (2017).

6.1.3 | Classification of Saving Behavior in the Borrowing Range

The behavioral model proposed in the previous section involves an extreme form of debt aversion—it assumes that individuals avoid debt altogether. While this model performs very well (especially in the High treatments), this assumption might seem too strong. In this section, we set up a more general model to evaluate our data.

We propose reference-dependent savings functions: we assume that in the borrowing range, individuals multiply the optimal savings amount for each period by a debt-aversion parameter, λ, when setting their *behaviorally* optimal savings amount (while

TABLE 10 | Summary of the calculated λ parameter, by observations.

	HIGH		LOW	
	UNC	CON	UNC	CON
Period 1				
Mean	−0.70	0.41	−0.91	2.75
	(0.84)	(0.66)	(0.91)	(1.43)
Median	0.35	0.93	0.14	0.41
# obs.	597	615	344	356
Period 2				
Mean	−0.51	—	0.05	—
	(1.37)		(0.52)	
Median	0.73	—	0.62	—
# obs.	548	—	321	—

Note: Standard errors clustered at the subject level in parentheses.

they set optimal savings according to Section 2 in the savings range).[37] Thus, the behavioral savings function, allowing for λ, $s_t^{\lambda,j}$, is a piecewise function:

$$s_t^{\lambda,j} = \begin{cases} \lambda^j s_t^{*,j} & \text{if } s_t^{*,j} < 0 \\ s_t^{*,j} & \text{if } s_t^{*,j} \geq 0. \end{cases} \quad (12)$$

Here, j designates the treatment, t the period, and $s_t^{*,j}$ is the optimal solution according to Section 2. Given an observed savings decision, s_t^j, in the borrowing range, we can calculate λ (for decisions in the savings range, we cannot calculate λ; by construction, this is not possible in period 2 in the CON treatments):

$$\lambda_t^j = \begin{cases} \frac{s_t^j}{s_t^{*,j}} & \text{if } s_t^{*,j} < 0 \\ \text{not defined} & \text{if } s_t^{*,j} \geq 0. \end{cases} \quad (13)$$

Thus, λ can be considered a (potential) wedge between the borrowing amount predicted by the standard savings functions and the observed decision (where λ can push the decision toward either saving more or less than predicted, depending on the sign). We can interpret the debt-aversion parameter λ in the following way:[38]

- $\lambda \leq 0$ indicates *perfect debt aversion* (the subject does not borrow, even though it is predicted),
- $0 < \lambda < 1$ indicates *partial debt aversion* (the subject borrows, though less than predicted),
- $\lambda = 1$ indicates *no debt aversion* (the subject borrows *exactly* as much as predicted),
- $\lambda > 1$ indicates *overborrowing* (the subject borrows more than predicted).

We calculate the debt-aversion parameter for each decision in each treatment of our experiment. Table 10 shows summary statistics. While the medians of all treatments and periods lie in the range of partial debt aversion, the means differ by the treatments: In period 1 of all UNC treatments and period 2 of UNCHIGH, the means indicate perfect debt aversion; with a lower income variance and with the constraint, the mean savings decisions are less debt-averse.

6.2 | Cognitive Abilities and Risk Preferences

In this section, we examine whether the measures from Part 3 of our experiment, the ToH game, the lottery task, and the items from our ex post questionnaire can explain the deviations we observe from optimal savings predictions.[39]

Table 11 shows results from panel regressions, where we try to understand deviations of observed behavior from theoretical predictions in period 1 based on 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 the absolute deviation as our first dependent variable. 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 strict risk aversion increases savings. We also look at simple deviations (to test if strict risk aversion increases savings) in period 1.

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 their deviations are in all specifications but the last one (where we do not expect it to 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 strict risk-aversion measure has the expected sign in all regressions. However, it is only statistically significant if we use all observations together with the direction of the deviation.[40] Our findings 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

Other laboratory studies of savings behavior (e.g., Ballinger et al. 2003, Brown et al. 2009, Meissner 2016) find that subjects improve their saving decisions with repetition. While in our experiment, one lifecycle consists of only two active periods, we give subjects ample time to learn by collecting data on 15 lifecycles for each of the CON/UNC treatments, resulting in 60 decisions per subject. In this section, we test whether there is evidence for learning in our experiment.

Table 12 reports results from fixed-effect panel data regressions, regressing deviations and absolute deviations of behavior from predictions on the round number in each of the four treatments and for each of the two periods. The coefficients on the round number are all negative and, thus, reduce the deviation in all treatments/periods. However, the round coefficients in the deviation regression are not significantly different from zero. Using

TABLE 11 | Determinants of pooled deviations in period 1.

	|Deviation|	Deviation >0	Deviation <0	Deviation
Female==1	2.335	1.291	−2.563	−3.493
	(2.152)	(2.555)	(1.922)	(2.639)
Age	−0.389	−0.478	0.170	0.0729
	(0.744)	(0.876)	(0.659)	(0.912)
CRT score	−3.493***	−3.120***	2.869***	0.798
	(0.802)	(0.969)	(0.706)	(0.984)
GPA	0.257	0.842	0.340	3.089
	(1.695)	(2.032)	(1.520)	(2.079)
ToH solved==1	4.527	5.665	−5.815	−15.055
	(9.570)	(11.086)	(8.620)	(11.735)
ToH solved==1 * Moves	0.090	0.129	−0.246	0.086
	(0.245)	(0.286)	(0.220)	(0.301)
Strict risk aversion	0.274	0.490	0.158	0.946**
	(0.378)	(0.450)	(0.332)	(0.464)
Constant	18.750	14.706	−12.783	−2.546
	(18.876)	(22.110)	(16.833)	(23.147)
#obs.	3420	1704	1701	3420
#clusters	114	110	111	114
R^2 (within)	—	—	—	—
R^2 (between)	0.201	0.143	0.182	0.092
R^2 (overall)	0.112	0.106	0.151	0.040

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

absolute deviations as the dependent variable, we observe that all the coefficients on round numbers are negative and significantly different from zero. Thus, we find evidence that over time, subjects reduce their mistakes; the learning over the 15 rounds per treatment of the experiment reduces the oversaving constant by about 13.4%–25.7%, depending on the treatment/period.[41]

7 | Summary and Conclusion

In this study, we have explored the impact of liquidity constraints on savings behavior in the buffer stock savings model. Liquidity constraints are considered empirically important factors in savings behavior, yet these constraints are often difficult for researchers to observe directly. Hence, we resort to a controlled laboratory test to turn liquidity constraints on or off in a simple three-period model. In our experiment, unexpected expenditure shocks and other savings motives are absent. Thus, our experiment provides a clean and direct test of the underlying theory.

Contrary to theoretical expectations, we find that the anticipation of liquidity constraints in the second period of our model does not lead to increased savings in the first period. However, our findings do corroborate the model's predictions in several other dimensions: (i) savings in the second period of the constrained model are higher than in the treatment without the constraint; (ii) savings in the second period are higher than in the first period except in one treatment; and (iii) savings are higher when there is increased income uncertainty. Still, in almost all treatments and periods, we observe significant oversaving relative to the theoretical, rational actor model predictions.[42]

In further analyses, we try to identify why the liquidity constraint does not result in higher period 1 savings. We find that a combination of debt aversion, heterogeneity in cognitive abilities, and/or learning can explain the deviations 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 more than they need to in period 1, the effect of the liquidity constraint in period 2 is greatly diminished.

Our results have several implications for macroeconomic modeling and policy design. First, our experimental evidence suggests that traditional buffer stock models could be adjusted to account for a different kind of liquidity constraint, namely, the behavioral bias against debt, along with heterogeneity in the extent of such biases. As we show, a model with partial debt aversion fits the experimental data better than the standard buffer stock model with or without real constraints on borrowing. Our results suggest a broader conceptualization of savings dynamics that incorporates behavioral factors and heterogeneous types.

TABLE 12 | 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	−0.153	−0.211***	round	−0.133	−0.164**	round	−0.117	−0.189**	round	−0.008	−0.120
		(0.096)	(0.066)		(0.090)	(0.065)		(0.115)	(0.081)		(0.102)	(0.074)
	constant	6.436***	19.329***	constant	3.583**	18.273***	constant	2.808	12.786***	constant	−1.196	11.194***
		(1.546)	(1.062)		(1.440)	(1.052)		(1.852)	(1.307)		(1.648)	(1.190)
	#obs.	1500	1500	#obs.	1500	1500	#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	0.005	R^2 (within)	0.002	0.008	R^2 (within)	0.000	0.004
	R^2 (between)	0.050	0.052	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	0.032	R^2 (overall)	0.000	0.000	R^2 (overall)	0.006	0.006	R^2 (overall)	0.000	0.002
Period 2	round	−0.041	−0.251***	round	−0.337***	−0.296***	round	−0.150	−0.261***	round	−0.109	−0.153**
		(0.110)	(0.081)		(0.081)	(0.072)		(0.125)	(0.092)		(0.076)	(0.067)
	constant	8.477***	22.011***	constant	16.512***	18.749***	constant	6.112***	15.217***	constant	7.995***	9.791***
		(1.778)	(1.307)		(1.307)	1.160)		(2.009)	(1.482)		(1.225)	(1.080)
	#obs.	1500	1500	#obs.	1500	1500	#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.

Regarding the policy implications of our findings, fiscal or monetary policy design might consider our finding that liquidity constraints play a more nuanced role than theory suggests. For instance, policies aimed at alleviating liquidity constraints could inadvertently affect savings behavior, potentially leading to oversaving and underconsumption. Moreover, our results indicate that income uncertainty significantly influences savings and consumption decisions, underscoring the importance of income stabilization policies to mitigate the adverse effects of income volatility on savings decisions.

Adjusting our experimental design to address further research questions would be promising. For example, in the settings we study, optimal first period savings is always positive; it would be of interest to study cases where optimal first savings is negative, to see whether this reduced the incidence of debt aversion. It would also be interesting to examine the role of income uncertainty further. By comparing savings behavior in treatments with an uncertain income (drawn from a given income range) with treatments with perfect foresight (with income from the same range), one could further disentangle the role of uncertainty in buffer stocks savings (an idea building on experiments by Carbone and Duffy 2014 and Duffy and Li 2019). It would also be interesting to consider the behavior of other subject pools that varied more in age, wealth, or other demographic factors, or to correlate behavior with subjects' experience with borrowing and borrowing constraints. We view our experimental design and findings as an important first step, so we leave these and other interesting extensions to future research.

Acknowledgments

We thank the editor, Maskai Aoyagi, and three anonymous referees for helpful comments and suggestions. Andreas Orland wishes to thank the Corvinus Institute for Advanced Studies, Corvinus University of Budapest, for their support of this research. Funding for this project was provided by the UC Irvine (UCI) School of Social Sciences. The experimental protocol used in this study was approved by UCI's Institutional Review Board, HS#2011-8378. A preregistration of this study can be found at: https://osf.io/hxfrj.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availbility Statement

The data that support the findings of this study are openly available on the Open Science Framework (OSF) website at: https://osf.io/827xq/files/osfstorage Reference: Liquidity constraints, income variance, and buffer stock savings: Experimental evidence.

Endnotes

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

[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). The savings/consumption context is especially suitable for experimental testing since almost every person engages as a consumer in intertemporal consumption and savings decisions (in contrast to selling goods in oligopoly markets or bidding in auctions)—here a student sample has many advantages over a sample from the general population. Students are intelligent and learn fast (making for a more conservative test of the theory), the usual monetary stakes in experiments are more salient for student subjects than for others with more lucrative outside options. Finally, student recruitment is relatively cheap, as students are already on campus. We consider laboratory experiments an important complement to empirical work on savings behavior using field data.

[3] Carroll (1997)'s buffer stock model has the possibility of a zero-income state, but delivers results similar to the case of there being a liquidity constraint.

[4] Some experiments have studied the role of liquidity constraints in intertemporal decision-making but without a buffer stock savings motive. Stahl (2013) examines preferences over the timing of two hypothetical payments, early and later, and finds that the share of subjects who prefer the earlier payment is higher when the payment dates are before and after Christmas (in comparison to two neutral dates in the future). Stahl assumes that this behavior difference is driven by the subjects' anticipation of their liquidity constraints due to Christmas shopping. Di Laurea and Ricciuti (2003) conduct overlapping-generations experiments to examine how subjects react to tax cuts. They find stronger support for Ricardian equivalence in a treatment without liquidity constraints than in a treatment with a liquidity constraint.

[5] There is some empirical literature documenting debt aversion outside the context of buffer stock savings models. Meissner (2016) found evidence for debt aversion in an experiment without liquidity constraints. Meissner and Albrecht (2022) develop a model of debt aversion that disentangles it from risk, time, and loss preferences. Using this framework, they report on an experiment involving real debt and find that 89% of their subjects can be identified as debt-averse. Caetano et al. (2019) study how the framing of debt matters for debt aversion among student loan candidates, for example, as a loan or a human capital contract. Field evidence for debt aversion comes mainly from research on higher education investments, where debt aversion may hinder borrowing or the choice of career paths—see, for example, Callender and Jackson (2005), Field (2009), Oosterbeek and van den Broek (2009), and Gopalan et al. (2024). By contrast, Eckel et al. (2007) report that debt aversion does not affect the take-up of loans for higher education in Canada. Elsewhere, evidence of debt aversion has been found among homeowners, Pavan and Barreda-Tarrazona (2020), Schleich et al. (2021) and in the firm financing decisions of small entrepreneurs, Nguyen et al. (2021) and Paaso et al. (2022).

[6] The derivation of the analytical solutions of similar models can be found in Besley (1995, 2141–2144), Carroll et al. (2021, 36–38), and Jappelli and Pistaferri (2017, 115–118).

[7] Furthermore, 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, they would unlikely fare better in a many-period lifecycle model.

[8] 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).

[9] Note that with log preferences, $u'''(c) = c^{-4} > 0$. This means the individual should prudently respond to risk, that is, an operative precautionary savings motive exists. With a quadratic utility function, $u(c) = ac - \frac{b}{2}c^2$, instead of the log function, we would avoid inducing any precautionary motives in the unconstrained treatments. However, this would come at a cost: The quadratic function increases almost linearly in consumption around the solution for many experimental parameterizations that we calculated (when designing the experiment, we faced a trade-off of having a solution where the curvature of the utility function is sufficiently strong and the optimal solution not violating the natural borrowing constraint, discussed in the next section). Thus, the log function has the important advantage that we induce more salient incentives than with the quadratic function.

[10] For example, the rearranged first-order condition of the problem in Equation (6) is given by $\frac{1}{w_2 - s_2} = \mathbb{E}_{t=2}\left[\frac{1}{y_3 + s_2}\right]$.

[11] The MATLAB code is provided in Section A.1 of the online Appendix.

[12] What happens up to the point where savings predictions in period 1 (with and without the constraint) coincide? Technically, in period 1, we need to calculate expected savings s_2 (but these depend on realized y_1, optimal s_1, and expected y_2). Depending on the presence of the constraint, these expected savings can or cannot be negative. At the point where expected period 2 savings become identical with and without the constraint, optimal period 1 savings coincide. Thus, for the expected period 2 savings where the constraint is binding, the individual needs to compensate by saving more in period 1 compared to the unconstrained model.

[13] In our high variance income treatment, if we increased the mean-preserving spread further, it would result in a violation of the natural borrowing constraints explained at the end of this section (the solution would predict more borrowing than allowed by the constraint).

[14] The constant of 0.77 in the utility function only scales the payoffs and does not affect the solutions.

[15] The online 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).

[16] We could have also shifted the income range in period 1 by the mean period 1 savings of the same treatment (which would have resulted in the same comparative statics predictions shown in Table 1, only with different levels). See Section A.3 in the online Appendix for details.

[17] In Table A.6 of the online Appendix, we also report tests of Hypotheses 4–7, separately for income/wealth levels where either borrowing or saving is predicted. All differences between HIGH and LOW have the predicted sign (in period 2 in CON, no savings are predicted); they are only significantly different from zero in the savings range but not in the borrowing range. This finding is consistent with debt aversion. (Note that in this comparison, the thresholds that divide the income/wealth range into savings/borrowing range are different between treatments.)

[18] The difference in means test of H3c has the predicted sign but is not significantly different from zero; thus, our result is inconclusive. We note that this result may be sensitive to whether we extend the prediction for period 2 savings in UNCLOW to the wealth range or stick with the income range. As observed oversaving in period 1 of UNCLOW is low (see Table 4), this might have affected the relevant wealth range in period 2.

[19] Oversaving, especially for potentially liquidity-constrained individuals, might not seem plausible at first sight. However, according to Sahadi (2015), many US 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).

[20] We report tests for order effects in Table A.7 of the online Appendix. We do not observe strong evidence of order effects: The differences between period 1 savings in UNCHIGH are only significant at the 10% level, and the difference in period 2 savings in UNCHIGH, even though significant at the 5% level, might be driven by previous oversaving in period 1.

[21] Under the assumption that subjects choose randomly from the interval of allowed savings values (those between the natural savings and the natural borrowing constraints), expected savings functions are described by the midpoint within the allowed interval and are as follows: In the UNC treatments $s_1^R(y_1) = \frac{1}{2}(y_1 - y_{min})$ and $s_2^R(w_2) = \frac{1}{2}(w_2 - y_{min})$, and in the CON treatments: $s_1^R(y_1) = \frac{1}{2}(y_1 - y_{min})$ and $s_2^R(w_2) = \frac{1}{2}(w_2 - 1)$.

[22] In that case, the savings function in period 1 is: $s_1 = (2y_1 - \mathbb{E}_{t=1}y_2 - \mathbb{E}_{t=1}y_3)/3$ and the savings function in period 2 is: $s_2 = (w_2 - \mathbb{E}_{t=2}y_3)/2$. Since expected income is 70 and the same in all treatments, we have $s_1^{H1}(y_1) = \frac{2}{3}(y_1 - 70)$ for both treatments and $s_2^{H1}(w_2) = \frac{1}{2}(w_2 - 70)$ for UNC and $s_2^{H1}(w_2) = \begin{cases} \frac{1}{2}(w_2 - 70) & \text{if } w_2 \geq 70 \\ 0 & \text{otherwise} \end{cases}$ for CON.

[23] In that case, the savings function for period 1 is: $s_1^{H2}(y_1) = \frac{1}{2}(y_1 - \mathbb{E}_{t=1}y_2) = \frac{1}{2}(y_1 - 70)$ for both treatments. Under this heuristic, there is no need to borrow or save in period 2.

[24] We test for treatment differences between UNC and CON using two-sided t-tests (clustered at the subject level) of absolute and shares of optimal. Absolute losses ($p = 0.060$) and shares of optimal ($p = 0.061$) of LOW are significantly different from zero at the 10% level; in HIGH, the differences are not significant (both $p = 0.221$).

[25] Such heterogeneity can also be found outside the lab. An enormous dispersion of wealth and savings, even among US households with similar socioeconomic characteristics, has been reported by Venti and Wise (1998) with data from the Health and Retirement Survey.

[26] The idea of threshold estimations is that they specify that individual observations can be divided into classes based on the value of an observed variable (Hansen 1999). This allows for the fact that, in our case, the savings predictions in the CON treatments have different slopes, depending on the wealth realization. We use the Stata package `xthreg` for panel regressions (due to Wang 2015; the article also gives an introduction to threshold panel regressions).

[27] The nonlinearity of the predictions is negligible, and the linear approximations provide a very good fit with $R^2 > 0.99$ in both the HIGH and LOW treatment.

[28] 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).

[29] Table A.8 in the online 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 7 (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 7. The coefficient in CONHIGH is significantly larger than the prediction, and the intercept is larger than predicted.

[30] To complement the results in Table 7, we also show Tobit regressions for all periods and treatments (but period 2 in the CON treatments) in Table A.9 in the online Appendix to account for the observations on

the natural savings constraints. The results shown are not qualitatively different from those reported in Table 7.

[31] 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.

[32] In contrast to the previous figures, we do not show the whole wealth range in Figure 6e,f. As savings decisions are very dispersed 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 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.)

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

[34] 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 the saving range. Note that it is not necessary for subjects to know the cutoff where they should switch from borrowing to saving (and we do not inform them about this cutoff) to apply a regression discontinuity analysis. For example, Hoekstra (2009) examines the effect of attending a flagship state university on future earnings. He uses applicants' SAT scores as the running variable, and the required SAT score to be accepted (i.e., the cutoff) was set according to an admission rule. This rule was not publicly known, and the cutoff score changed (moderately) over time.

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

[36] 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 online Appendix: Table A.10 reports mean absolute errors (in contrast to the RMSE, mean absolute errors punish deviations proportionately), Table A.11 shows mean errors (mean errors consider the individual deviation's direction but averaging them can cancel out positive and negative deviations; punishment is proportional), and Table A.12 reports a bias measure (which compares the average prediction magnitude with the average observed magnitude). Generally, these comparisons also prefer more constrained models, confirming the results from the RMSE comparisons.

[37] This seems a reasonable assumption. The estimated linear savings functions in Figures 4 and 5 show that, in the aggregate, deviations from optimal savings in the borrowing range increase with the distance from the savings range.

[38] Notice that the denominator on the right-hand side of Equation (13), $s_t^{*,j}$, is negative by definition.

[39] Table A.13 in the online Appendix summarizes subjects' characteristics and Table A.14 in the online Appendix displays a correlation matrix of the variables. The sample is close to being gender-balanced. Subjects' CRT scores are more dispersed than their GPAs. Almost all subjects solved the ToH task using only a few moves (33 out of 140 subjects solved the task in seven moves, 23 in eight, and 24 in nine moves). We use the binary lottery task to derive two risk-aversion measures: strict and proxy. Strict risk aversion is gauged by the number of less risky choices made by a subject and excludes 36 subjects who switched choices more than once or initially chose a risky option. Proxy risk aversion is measured similarly but includes all subjects, regardless of the number of choice switches or whether they started with a risky choice. Figures A.6–A.7 in the online Appendix show histograms of subjects' strict and proxy risk aversion. Both distributions closely align, and the average measures, with a mean and median of around 6, suggest mild risk aversion among subjects.

[40] Table A.15 in the online Appendix shows the same regressions as Table 11 for period 2 deviations (conditional on s_1 and y_2). The results confirm our findings for period 1 (the risk-aversion coefficients are at least significant at the 10% level in all specifications). Tables A.16 and A.17 in the online Appendix show regressions that use proxy risk aversion instead of strict risk aversion. All risk-aversion coefficients have a positive sign, though many observations are needed for significance in these regressions.

[41] Allen and Carroll (2001) simulate learning in the buffer stock savings model using a linear approximation of a nonlinear consumption function and find that learning takes roughly a million simulation periods. In contrast, our subjects learn quickly.

[42] Several empirical studies report that the buffer stock savings model has problems in explaining savings behavior (e.g., Jappelli et al. 2008 and Fulford 2015). In contrast to Fulford (2015), higher income uncertainty increases savings.

References

Allen, T. W., and C. D. Carroll. 2001. "Individual Learning About Consumption." *Macroeconomic Dynamics* 5: 255–271.

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

Ballinger, T. P., M. G. Palumbo, and N. T. Wilcox. 2003. "Precautionary Saving and Social Learning Across Generations: An Experiment." *Economic Journal* 113: 920–947.

Besley, T. 1995. "Savings, Credit and Insurance." In *Handbook of Development Economics*, edited by J. Behrman and T. Srinivasan, Vol. 3, 2123–2207. Elsevier.

Boar, C., D. Gorea, and V. Midrigan. 2022. "Liquidity Constraints in the U.S. Housing Market." *Review of Economic Studies* 89: 1120–1154.

Brown, A. L., Z. E. Chua, and C. F. Camerer. 2009. "Learning and Visceral Temptation in Dynamic Saving Experiments." *Quarterly Journal of Economics* 124: 197–231.

Caetano, G., M. Palacios, and H. A. Patrinos. 2019. "Measuring Aversion to Debt: An Experiment Among Student Loan Candidates." *Journal of Family and Economic Issues* 40: 117–131.

Callender, C., and J. Jackson. 2005. "Does the Fear of Debt Deter Students From Higher Education?" *Journal of Social Policy* 34: 509–540.

Calonico, S., M. D. Cattaneo, M. H. Farrell, and R. Titiunik. 2017. "Rdrobust: Software for Regression-Discontinuity Designs." *Stata Journal* 17: 372–404.

Campbell, J. R., and Z. Hercowitz. 2018. "Liquidity Constraints of the Middle Class." *American Economic Journal: Economic Policy* 11: 130–155.

Carbone, E., and J. Duffy. 2014. "Lifecycle Consumption Plans, Social Learning and External Habits: Experimental Evidence." *Journal of Economic Behavior & Organization* 106: 413–427.

Carroll, C. D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112: 1–55.

Carroll, C. D. 2004. "Theoretical Foundations of Buffer Stock Saving." National Bureau of Economic Research Working Paper 10867.

Carroll, C. D., R. E. Hall, and S. P. Zeldes. 1992. "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity* 1992: 61–156.

Carroll, C. D., M. B. Holm, and M. S. Kimball. 2021. "Liquidity Constraints and Precautionary Saving." *Journal of Economic Theory* 195: 105276.

Chen, D. L., M. Schonger, and C. Wickens. 2016. "oTree—An Open-Source Platform for Laboratory, Online, and Field Experiments." *Journal of Behavioral and Experimental Finance* 9: 88–97.

Deaton, A. 1991. "Saving and Liquidity Constraints." *Econometrica* 59: 1221–1248.

Di Laurea, D., and R. Ricciuti. 2003. "An Experimental Analysis of Two Departures From Ricardian Equivalence." *Economics Bulletin* 8: 1–11.

Drichoutis, A. C., and J. L. Lusk. 2016. "What Can Multiple Price Lists Really Tell Us About Risk Preferences?" *Journal of Risk and Uncertainty* 53: 89–106.

Duffy, J. 2016. "Macroeconomics: A Survey of Laboratory Research." In *Handbook of Experimental Economics*, edited by J. H. Kagel and A. E. Roth, Vol. 2, 1–90. Princeton University Press.

Duffy, J., and Y. Li. 2019. "Lifecycle Consumption Under Different Income Profiles: Evidence and Theory." *Journal of Economic Dynamics and Control* 104: 74–94.

Eckel, C. C., C. Johnson, C. Montmarquette, and C. Rojas. 2007. "Debt Aversion and the Demand for Loans for Postsecondary Education." *Public Finance Review* 35: 233–262.

Field, E. 2009. "Educational Debt Burden and Career Choice: Evidence From a Financial Aid Experiment at NYU Law School." *American Economic Journal: Applied Economics* 1: 1–21.

Frederick, S. 2005. "Cognitive Reflection and Decision Making." *Journal of Economic Perspectives* 19: 25–42.

Friedman, M. 1957. *A Theory of the Consumption Function*. Princeton University Press.

Fuchs-Schündeln, N., and M. Schündeln. 2005. "Precautionary Savings and Self-Selection: Evidence From the German Reunification 'Experiment'." *Quarterly Journal of Economics* 120: 1085–1120.

Fulford, S. L. 2015. "The Surprisingly Low Importance of Income Uncertainty for Precaution." *European Economic Review* 79: 151–171.

Gelman, A., and G. Imbens. 2019. "Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs." *Journal of Business & Economic Statistics* 37: 447–456.

Gopalan, R., B. H. Hamilton, J. Sabat, and D. Sovich. 2024. "Aversion to Student Debt? Evidence From Low-Wage Workers." *Journal of Finance* 79: 1249–1295.

Gross, D. B., and N. S. Souleles. 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence From Credit Card Data." *Quarterly Journal of Economics* 117: 149–185.

Haliassos, M., and A. Michaelides. 2003. "Portfolio Choice and Liquidity Constraints." *International Economic Review* 44: 143–177.

Hansen, B. E. 1999. "Threshold Effects in Non-Dynamic Panels: Estimation, Testing, and Inference." *Journal of Econometrics* 93: 345–368.

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

Hoekstra, M. 2009. "The Effect of Attending the Flagship State University on Earnings: A Discontinuity-Based Approach." *Review of Economics and Statistics* 91: 717–724.

Internal Revenue Service. 2020. "2020 Filing Season Statistics." Accessed September 1, 2020. https://www.irs.gov/newsroom/filing-season-statistics-for-week-ending-may-1-2020.

Jappelli, T. 1990. "Who is Credit Constrained in the US Economy?" *Quarterly Journal of Economics* 105: 219–234.

Jappelli, T., M. Padula, and L. Pistaferri. 2008. "A Direct Test of the Buffer-Stock Model of Saving." *Journal of the European Economic Association* 6: 1186–1210.

Jappelli, T., and L. Pistaferri. 2017. *The Economics of Consumption: Theory and Evidence*. Oxford University Press.

Jappelli, T., and L. Pistaferri. 2025. "Permanent Income Shocks, Target Wealth, and the Wealth Gap." *American Economic Journal: Macroeconomics* 17: 102–125.

Jiang, Y., X. He, M.-L. T. Lee, B. Rosner, and J. Yan. 2020. "Wilcoxon Rank-Based Tests for Clustered Data With R Package clusrank." *Journal of Statistical Software* 96: 1–26.

Kimball, M. S. 1990. "Precautionary Saving in the Small and in the Large." *Econometrica* 58: 53–73.

Lee, D. S., and T. Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature* 48: 281–355.

Love, D. 2006. "Buffer Stock Saving in Retirement Accounts." *Journal of Monetary Economics* 53: 1473–1492.

Ludvigson, S. C., and A. Michaelides. 2001. "Does Buffer-Stock Saving Explain the Smoothness and Excess Sensitivity of Consumption?" *American Economic Review* 91: 631–647.

Meissner, T. 2016. "Intertemporal Consumption and Debt Aversion: An Experimental Study." *Experimental Economics* 19: 281–298.

Meissner, T., and D. Albrecht. 2022. "Debt Aversion: Theory and Measurement." https://doi.org/10.48550/arxiv.2207.07538.

Modigliani, F., and R. Brumberg. 1954. "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data." In *Post-Keynesian Economics*, edited by K. Kurihara, 388–436. Rudgers University Press.

Nguyen, H. T., H. M. Nguyen, M. Troege, and A. T. Nguyen. 2021. "Debt Aversion, Education, and Credit Self-Rationing in SMEs." *Small Business Economics* 57: 1125–1143.

Oosterbeek, H., and A. van den Broek. 2009. "An Empirical Analysis of Borrowing Behaviour of Higher Education Students in the Netherlands." *Economics of Education Review* 28: 170–177.

Paaso, M., V. Pursiainen, and S. Torstila. 2022. "Entrepreneur Debt Aversion and Financing Decisions: Evidence From COVID-19 Support Programs." In *Proceedings of Paris December 2021 Finance Meeting EUROFIDAI-ESSEC*.

Pavan, M., and I. Barreda-Tarrazona. 2020. "Should I Default on My Mortgage Even If I Can Pay? Experimental Evidence." *Journal of Economic Dynamics & Control* 110: 103733.

Sahadi, J. 2015. "Nearly 8 Out of 10 U.S. Taxpayers Get Refunds." Accessed September 1, 2020. https://money.cnn.com/2015/01/13/pf/taxes/taxpayer-refunds/index.html.

Schleich, J., C. Faure, and T. Meissner. 2021. "Adoption of Retrofit Measures Among Homeowners in EU Countries: The Effects of Access to Capital and Debt Aversion." *Energy Policy* 149: 112025.

Stahl, D. O. 2013. "Intertemporal Choice With Liquidity Constraints: Theory and Experiment." *Economics Letters* 118: 101–103.

Thaler, R. H. 1990. "Anomalies: Saving, Fungibility, and Mental Accounts." *Journal of Economic Perspectives* 4: 193–205.

Toplak, M. E., R. F. West, and K. E. Stanovich. 2014. "Assessing Miserly Information Processing: An Expansion of the Cognitive Reflection Test." *Thinking & Reasoning* 20: 147–168.

Venti, S. F., and D. A. Wise. 1998. "The Cause of Wealth Dispersion at Retirement: Choice or Chance?" *American Economic Review* 88: 185–191.

Wang, Q. 2015. "Fixed-Effect Panel Threshold Model Using Stata." *Stata Journal* 15: 121–134.

Zeldes, S. P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97: 305–346.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.

Data S1 Data S1