

Debt Aversion and the Impact of Natural Borrowing Constraints: Evidence from an Experiment

Jason Ralston

October 5, 2017

ABSTRACT

Consumption theory states that if total lifecycle income remains constant, consumption will be flat regardless of whether the income profile is upward or downward sloping. However, in experimental literature deviations from optimal consumption are larger in lifecycles with upward sloping income profiles than downward. This has come to be referred to as “debt aversion”. Through use of increasing and decreasing income profiles with the same expected total income and same expected total life time utility in a treatment, consumption behavior under borrowing and saving regimes are able to be compared directly. In the experiment, we allow for lengthening of the lifecycle and adding of borrowing constraints. It is found that there are modest reductions in deviations from optimal behavior when in the presence of a borrowing constraint, and that its effects are largest in the longer of the two lifecycle types. This indicates that loss aversion and risk aversion are alleviated by borrowing constraints.

Keywords: Debt aversion, consumption smoothing, loss aversion, prospect theory, uncertainty, natural borrowing constraint

JEL Classifications: C91, D91, E21

1 Introduction

Three hypotheses existence of debt aversion in experimental settings are posited here. First, it is possible that subjects would have learned to optimize properly had the lifecycle been repeated more times. Most experimental papers in this field use long lifecycles, usually twenty or more periods in length, which allows for few repetitions of the consumption task. In contrast, I create much shorter lifecycles which can be repeated many more times in the same amount of time as previous consumption experiments. Secondly, it may have been that these long lifecycles with stochastic income profiles was creating large amounts of uncertainty, which risk and loss averse subjects countered by underconsuming. It may be the case that a shorter lifecycle under a borrowing regime is just as easily optimized as a similar lifecycle under a savings regime since the uncertainty in future income is reduced substantially.

Thirdly, it may be a psychological tendency to avoid losses when possible that prevents agents from borrowing optimally. Such a prediction would be suggested by the prospect theory and loss aversion literature pioneered largely by Kahneman and Tversky in the 1980s and 90s.

2 Theory

2.1 Optimal Consumption

Consumption smoothing results from the combination of three factors. The first is an objective function which the subject must seek to optimize. The second is that the objective function must be over many periods, and these periods must form the lifecycle. The final factor is a budget constraint which sets limits for borrowing and saving.

My own optimization problem will be constructed using a constant relative risk aversion (CRRA) utility function

$$U(c_t) = \theta (1 - e^{-\alpha c_t})$$

In the experiment, c_t is multiplied by a constant α , which maximizes the salience of the point function for subjects. θ is also used in a similar fashion. For two period games, α is set to $\frac{1}{25}$. For five period games, α is set to $\frac{1}{40}$. θ was set such that the two point functions had equivalent monetary payoffs.

Assets held in time zero and assets held in the last period are forced to equal zero. Thus the optimization problem of the agent becomes

$$\max_{c_t} E_{\tau} \sum_{t=\tau}^T U(c_t) \tag{1}$$

$$\text{s.t. } c_t + a_t = w_t \tag{2}$$

$$a_0 = 0, a_T = 0 \tag{3}$$

where τ denotes the current period in the lifecycle, and w_t is defined to be $w_t = y_t + a_{t-1}$. T can take the values 2 or 5 in the experiment. The last thing needed to characterize this problem is an income profile. The two different income profiles will be defined as

$$y_{1,b,t} = 10t + 10 + \varepsilon_t, \text{ for } t = 0, 1 \text{ and } y_{1,s,t} = 40 - 10t + \varepsilon_t, \text{ for } t = 0, 1 \quad (4)$$

and

$$y_{2,b,t} = 10t + 10 + \varepsilon_t, \text{ for } 0 \leq t \leq 5 \text{ and } y_{2,s,t} = 70 - 10t + \varepsilon_t, \text{ for } 0 \leq t \leq 5 \quad (5)$$

The stochastic process ε_t raises income in time t by 10 with probability $\frac{1}{2}$ or lowers income in time t by 10 with probability $\frac{1}{2}$. The essential feature to note is that in the two period lifecycles the expected total lifecycle income is identical regardless of the positive or negative slope of the profile. This is true of the five period lifecycles as well. In standard household consumption theory, the slopes of the income profiles should not matter and both should produce the same consumption decisions, given the specification of the utility function. The standard prediction is that the consumption decisions should produce a flat curve, on average, and this behavior is known as consumption smoothing.

This problem is solved analytically using backward induction. Thus the conditionally and unconditionally optimal values for consumption can be found.

2.2 Loss Aversion

It should be noted that the setup above does not necessarily imply that agents cannot go bankrupt. That is, agents could consume so much during a lifecycle such that they accrue debt that cannot be repaid by income. Should this occur, consumption would become negative in the last period.

Since subjects generally do not have a notion of what negative consumption is, consuming above total life lifecycle income was modeled in a way that was more salient to lab participants. Before a lifecycle began, subjects were given an endowment of points. This endowment was not allowed to be used for consumption decisions. If the subject consumed so much that they could not repay it by the last round, the subject was denied the utility points earned from previous consumption choices and part of the endowment would be removed. This was meant to give the act of overconsumption a little more “bite” than simply using the original CARA function provided. It was felt that this was more in line with how agents would perceive the effect of defaulting on loans. In the full 1400 lifecycles witnessed in the laboratory, an overconsumption like this never occurred, indicating that the disincentive was quite powerful.

To find just how powerful this disincentive was to subjects, borrowing constrained and borrowing *unconstrained* groups were created. These two groups are called “Type A subjects” and “Type B subjects” respectively. The theoretical implications of these two groups are detailed below.

2.2.1 Type A Subjects

If the subject is designated as type A, then a natural borrowing constraint is implemented for the subject. The natural borrowing constraint limits borrowing in the following fashion:

$$a_t \geq \min \sum_{i=t+1}^T -y_i \quad (6)$$

In essence what this constraint promises to a subject is that their consumption choices can never exceed the minimum amount of income the subject will receive in the rest of their lifecycle. Type A subjects will never have to face giving up a portion of their endowment.

2.2.2 Type B Subjects

If the subject is designated to be type B, the subject is allowed to have negative consumption in the last period of the lifecycle. In this event, the subject has their negative consumption converted into a payoff which takes the following form:

$$\text{Lifecycle Payoff} = e^{c_T/5+107025} - 107025 \text{ if } c_T < 0 \quad (7)$$

Thus, as a Type B subject, the possibility of negative consumption in the last round is addressed in the experiment by creating a large disincentive. By the nature of function, negative consumption can produce only negative lifecycle payoffs, and these payoffs asymptotes at $-107,025^1$. This negative value is added to the subject's endowment, which will be discussed in greater detail in the experimental design section.

Being assigned to type A or type B does not affect the optimal consumption path that a rational agent would follow. However, it may have psychological consequences for a subject.

These types are implemented to test for the effects of loss aversion. Loss aversion, a part of prospect theory, states that the reaction to gains and losses relative to some reference point are treated differently. Specifically, this is the tendency that people strongly prefer avoiding losses to acquiring gains (Kahneman & Tversky 1986). The psychological effect from a loss is sometimes cited as being twice that of a gain (Kahneman & Tversky 1992).

If loss aversion does affect the consumption decisions of agents, evidence would be given by a difference in consumption patterns between type A and type B subjects. If that the reference point for subjects is the initial endowment, then subjects given no budget constraint would be averse receiving a final payout of less than this original endowment. Thus, it is expected that type B subjects would under-borrow and under-consume in upward sloping income profiles, because they are more likely to borrow small sums of money (if they borrow at all) to offset against the risk of future income being insufficient to repay their debt.

However, if a subject is type A, it is expected that their consumption and borrowing rates would be higher than their type A counterparts. This is because they do not face the disincentive of losing their endowment as type A subjects do. Since their borrowings can never be so high that they could lose utility points, the type B subjects have, in effect, done away with loss aversion entirely. From this the effect of loss aversion on consumption behavior can be determined.

The borrowing constraint is expected to have a larger effect in the five period rounds than in the two period rounds. This is because the distributions in potential lifetime income in the five period round has a much larger variance compared to that of the two round game.

¹This value was chosen to be equivalent to \$3 after converted from points to dollars.

When faced with greater uncertainty in future income, it incentivizes players who are loss averse to make consumption choices lower than the optimal consumption choice to avoid over consuming and losing their endowment. Therefore, the interaction of round length and borrowing constraint effects may be substantially different between rounds of different lengths.

As previously mentioned, in each lifecycle there are twenty periods. Each period in the lifecycle begins with the subject getting a new draw of income and then the subject deciding how much of this income to consume. The subject is alerted to how the consumption determines savings or borrowing. A list of past values of the main variables of interest (consumption, savings/borrowing, utility points) are provided to the subject so that they may refine their choices. At the end of a completed lifecycle a list of past lifecycle utilities is displayed so the subject can gauge if they are becoming better at optimizing over time, which allows for some learning to occur.

Income is paid out in terms of experimental currency each period. With the income from the current period and what was borrowed/saved in the past, the agents are allowed to buy units of consumption, which are converted into utility points by a function, which is notable because it allows for the possibility of negative consumption. These utility points are used to determine the monetary payoff of the subjects.

3 Experimental Design

3.1 Procedure

Subjects were given all the relevant information regarding the objective function they are attempting to maximize. This is displayed as the graphically and as a table of values generated by the function.

Subjects were made aware of the different income processes that generated experimental currency, called “Tickets” in the experiment. They were exposed to an example which illustrated how income was generated, how consumption in one period affected savings/borrowings, and how a round’s payment was calculated. Subjects were allowed to have scratch paper with which to plan out their consumption decisions or perform any calculation they considered necessary.

Those subjects who were given the type A were provided with the definition and illustrative example of the natural borrowing constraint and how it operated. How subjects were divided into type A and type B, as well as how treatments were randomized, is discussed below.

At the beginning of every period, income was generated and the subjects are endowed with 107,025 points, with which they cannot interact. A full history of all relevant variables (Period, Consumption, Savings/Borrowings, and Money) was presented to the subjects during in each period so they could evaluate how their choices were affecting the round’s payoff. After viewing this information, subjects were allowed to make a consumption choice.

For each combination of treatments, after each round was completed, subjects were shown a history of points from previous rounds, which were not converted to the monetary payoff to increase the salience of the information. This was used to give the subjects an idea of

how well they were performing over time, allowing for learning to take effect.

After the consumption portion of the experiment ended, all subjects were given an incentivized loss elicitation and an incentivized risk elicitation. Following these elicitations, there followed a short questionnaire that included questions about demographics, experience with mathematics, as well as personal experience with borrowing. Following this there was a screen displaying how much the subjects were to be paid. A copy of the choice screens used in the laboratory are provided in the appendix.

3.2 Randomization

The sample of subjects was split into two roughly even groups at random and given the designation of being either type A or type B, representing having the natural borrowing constraint being implemented or not. Both the income profile and round lengths were randomly ordered to protect against order effects. Subjects progressed through the four different combinations of income profile and round length seven times each to allow for learning to take place.

Since the different combinations of treatments are repeated multiple times, a single round from all twenty-eight was selected at random and used as payment. Two period round payments were scaled by a constant to keep amounts earned in two and five period rounds even.

The data was collected at the University of California, Irvine's Experimental Social Science Laboratory (ESSL) using recruits from the pool of ESSL students over three separate sessions. Fifty students participated in total. Most students were undergraduates in social science fields. The minimum payment was set at \$7 and the average payout was \$18.57.

4 Hypotheses

It is useful to quickly sum up the four major hypotheses made in this paper. This will make clear the *a priori* expectations this paper has about the subjects' behavior.

Hypothesis 1 *Longer rounds will lead to higher levels of debt aversion.*

This prediction is informed by predictions of prospect theory. Agents are predicted to be risk averse with respect to gains and risk loving with respect to losses relative to some reference point. This experiment features the possibility of both gain and loss.

On one hand, subjects are exposed to the uncertain gains that they will make from their consumption profile. This is where risk aversion factors in, because the longer a round becomes, the greater the uncertainty there is in future income. The greater that uncertainty, the greater the underconsumption and deviations from optimal behavior. Of course, this will also be affected by the risk preferences of each particular agent. For those who are more risk averse, we expect greater departures from optimal consumption compared to those who are risk neutral.

The other possibility is that subjects will respond to loss aversion. Since the prediction of prospect theory is that people value avoiding losses more than acquiring gains, it would be the case that people who are loss averse would keep consumption low to avoid the possibility

of loss. In this experiment, this would lead to underconsumption and deviations from optimal behavior.

There is slight concern that since this experiment features relatively low amounts of reward ², that the subjects will experience much loss aversion (Ert & Yachman 2008). For this reason, an incentivized loss preference elicitation is performed to categorize subjects as being loss averse or not. When examining the effects of loss aversion, the sample will sometimes be restricted to the set of subjects that were determined to be loss averse using the elicitation.

Hypothesis 2 *The natural borrowing constraint will help those who are loss averse come closer to optimal behavior.*

As mentioned above, loss aversion is a possible driving force behind debt aversion. If we believe subjects are loss averse, it will result in large deviations from optimal consumption choices. Since the borrowing constraint prevents an agent from being able to overconsume to the point where they cannot pay off their debts, loss aversion will contribute to the decisions of subject's less.

Hypothesis 3 *Subjects will learn to play consumption strategies closer to the optimum over many repetitions of lifecycles.*

In previous studies, lifecycles have typically lasted around fifteen to twenty periods. If they are repeated at all, they can only be repeated a few times due to the limited amount of time available in the laboratory. However, since the lifecycles studied in this paper are so short, they are able to be repeated many times more.

Since subjects are able to learn over many more rounds, it is expected that subjects will be able to correlate consumption choices that are closer to optimal with higher lifecycle utility and then use this information to improve their consumption choices.

Hypothesis 4 *Certain qualities of subjects will impact their consumption behavior, especially with their familiarity with borrowing outside the experiment.*

The subjects in Meissner (2015) were all in German schools. For German students attending university there is not much borrowing that must be done to finance their education. This is not the case for American university students, who often must finance their post-secondary education with loans. This may be a difference in experience may lead to differences in comfortability with borrowing between the two studies.

5 Results

5.1 All Agents Analysis

In order to talk about the deviations from optimal behavior and its impacts, two different measures of deviation are introduced. The first, m_1 , is the absolute deviation from conditionally optimal consumption, defined as

²Compared to the hypothetical questions typically proposed to test for loss aversion

$$m_1 = \sum_{t=1}^T |c_t^*(w_t) - c_t|$$

where $c_t^*(w_t)$ is the conditionally optimal consumption depending on the wealth an agent in time t and c_t is the actual value of individual consumption observed in time t in the experiment.

The second measure shows the welfare loss when deviating from conditionally optimal consumption, and is defined as

$$m_2 = \sum_{t=1}^T [u(c_t^*(w_t)) - u(c_t)]$$

This can be thought of as the sum of utility lost in each period from deviating from the conditionally optimal consumption. The unconditional welfare loss results are included for comparison in the appendix.

Figure 1: MEASURES OF DEVIATION FROM OPTIMAL BEHAVIOR

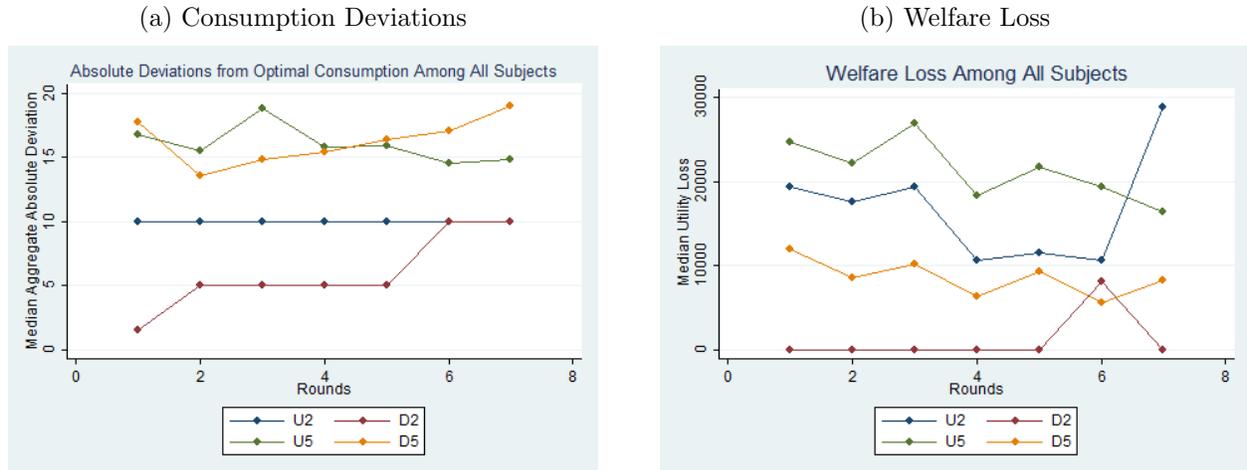


Figure 1 displays the medians³ of the four different round types⁴ against the round number. It must be noted here that five period round deviations are scaled by a factor of $\frac{2}{5}$ so that the deviations between two and five period long rounds can be directly compared. The forms of the utility functions were already scaled so that the utility from a two period game and a five period game would be directly comparable, as mentioned previously.

Three things are illustrated by these graphs. The first is that the length of round seems to have a noticeable impact on deviation from optimal behavior. There are even statistically

³I use the median because my statistical tests are typically done non-parametrically and give information about the median or stochastic dominance instead of the mean. Specific tests used are Mann-Whitney U-test and Kolmogorov-Smirnov. In general, these two tests agreed in significance. Significance in this paper is defined to be anything below $\alpha < 0.05$ and anything above is considered insignificant.

⁴two period with upward sloping income (2U), two period with downward sloping income (2D), five period with upward sloping income (5U), and five period with downward sloping income (5D)

significant differences between our measures not just in all two period rounds and all five period rounds, but also between rounds with upward income and downward income. This makes sense, because, as previously noted, agents face more uncertainty in their future income in a five period round than in a two period round. If an agent is loss averse, this condition will lead to more conservative consumption, and larger deviations from optimal consumption.

Secondly, agents respond differently to playing in upward and downward sloping income profiles. Within each length of round, the upward sloping income profile always has a higher median deviation from optimal behavior. What is more, even between rounds of the same length, the effect of upward and downward income profiles is still statistically significant.

Lastly, it is almost immediately clear that learning is not occurring in any meaningful sense across rounds. Four measures of learning are created to formally test if there is any learning between rounds.

$$\begin{aligned}
 l_1 &= \text{median}(m_1^{r-1} - m_1^r) \\
 l_2 &= \text{median}(m_1^1 - m_1^r) \\
 l_3 &= \text{median}(m_2^{r-1} - m_2^r) \\
 l_4 &= \text{median}(m_2^1 - m_2^r)
 \end{aligned}$$

where m_1^r is the deviation from conditionally optimal consumption in round r and m_1^r is the welfare loss in round r . The first and third measures, l_1 and l_3 , are intuitively the median difference between two consecutive periods. The second and fourth measures, l_2 and l_4 , are the difference between the first round and any subsequent round. Performing non-parametric tests, where the null hypothesis is that the median difference between the pairs in each measure is zero, never results in a p-value that would indicate even marginal significance for any measure. This solidifies the notion that no learning is happening over rounds, which is consistent with previous research.

Next, we attempt to detect if there are differences in consumption behavior between those who have the borrowing constraint and those that do not. First, non-parametric tests were performed to see if there was a difference between the welfare losses between those who did not have the borrowing constraint and those who did. The test resulted in a p-value below the .001 significance level, indicating that the two groups' consumption behaviors were significantly different. Not only this, but the welfare loss looked to be lower with the borrowing constraint than without.

Next, we tried to determine in which rounds the borrowing constraint was more likely to have an affect on welfare loss. Table 1 presents the medians of welfare loss by those who were not constrained by borrowing and by those who were. What is evident is that those in the borrowing constraint treatment seemed to perform better than those who did not have the borrowing constraint except in the case of the two period with downward sloping income. While all results for deviations from optimal consumption are inconclusive, an interesting pattern arises for the welfare loss measure. Only in the five period rounds did the borrowing constraint have a significant effect on welfare loss, which is in line with the hypothesis that loss aversion and risk aversion should affect agents more in longer rounds, and that a borrowing constraint will help alleviate the effects of loss and risk aversion.

The borrowing constraint only lends itself to between group comparisons, thus, how a person’s consumption behavior was changed by switching from unconstrained to constrained borrowing cannot be observed. However, it is noted that the difference between those with the borrowing constraint and those without is *larger* in the context of upward sloping income than with downward sloping income. This suggests that the borrowing constraint has more “bite” when combined with upward sloping income than downward sloping income, indicating that loss aversion plays a roll in the decisions of the subjects.

Table 1: MEDIANS OF ABSOLUTE CONSUMPTION AND WELFARE LOSS

	2U 0	2U 1	2D 0	2D 1	5U 0	5U 1	5D 0	5D 1
Cons Dev	10	10	5	5	18	19.2	15	16.8
<i>(p-value)</i>		<i>(0.8085)</i>		<i>(0.1567)</i>		<i>(0.8051)</i>		<i>(0.8548)</i>
Welfare Loss	16960	13640	.2343	.2343	25211	19848	8994	5676
<i>(p-value)</i>		<i>(0.9187)</i>		<i>(0.1276)</i>		<i>(0.0275)</i>		<i>(0.0071)</i>
<i>N</i>	203	147	203	147	203	147	203	147

0 = subjects without the borrowing constraint. 1 = subjects with the borrowing constraint

5.2 Demographic Analysis

5.2.1 Rational Agents

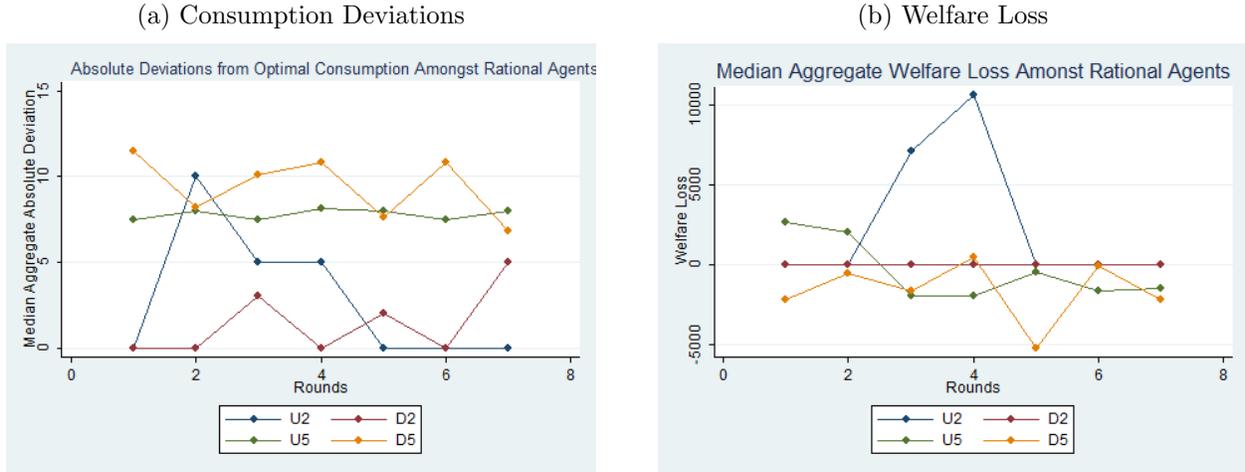
The lack of learning effects in the analysis with all agents seems strange. It is the case that there was a subset of agents in the experiment who never borrowed in any treatment. Indeed, there were some subjects who never even varied their consumption decisions within the different treatment types. That is to say that there were some subjects who acted in an alarmingly irrational way. But what about more rational subjects?

The sample was broken apart to focus on what was defined to be rational agents. The definition for rational agent was a subject who, in a five period round with upward sloping income, borrowed tickets in the first period ⁵. The definition chosen is based on this treatment type, because in a five period round with upward sloping income, it is always conditionally optimal to borrow in the first period. After applying this definition, it was found that of the original 50 subjects, only 13 display “rational” behavior.

Figure 2 is identical to Figure 1, except that only the rational agents are left in the sample. On first glance, the same relationships exist in the rational agents as existed in the full sample: length of round and direction of income profiles still seem to create statistically significant differences in consumption behavior. *However*, when being in an upward or downward sloping income profile is compared in two period rounds, there is no longer a statistically significant effect on absolute deviations from conditionally optimal consumption or welfare loss. This suggests that agents are no longer affected by upward or downward sloping income within two period rounds, which further suggests that under the right conditions, rational subjects can be said to think of savings and borrowings lifecycles in a somewhat equivalent manner.

⁵...and did so in a majority of the rounds. A majority of rounds is preferred to *all* rounds, because it is possible that subjects could learn to be rational after some duration.

Figure 2: MEASURES OF DEVIATION FROM OPTIMAL BEHAVIOR



The sample of rational subjects has only left three that do not have a borrowing constraint. This is too few to have a reasonable amount of power for our between subjects comparisons. Similarly, nonparametric tests could not be used to detect learning as the observations by period in a round type were very small and provided insufficient statistical power.

5.2.2 Other Covariates

During the experiment and following the experiment many demographic questions were asked of the subjects. Some of the more interesting covariates obtained were risk and loss preferences, sex, and borrowing experience. In this section, time is devoted to exploring the relationships between these covariates and consumption choices. Table 2 displays summary statistics of deviations from optimal behavior by different groups⁶.

As noted previously, the variance in future income may be driving the consumption decisions of our subjects. If we posit that subjects are risk averse, we can see that subjects may underconsume to protect against the risk of low income realizations, as theorized in macroeconomic literature (Chetty & Looney 2006).

Before the experiment ended, subjects were given a brief risk elicitation. Subjects' responses then indicated if they were risk averse, risk neutral, or risk loving. Of the 50 subjects, 17 were risk averse, 29 were risk neutral, and the remaining 4 were risk loving.

It would be expected for those who are risk averse to be the same people who underconsume in all periods. Therefore these subjects would also be those whose deviations from optimal consumption might be highest, compared to those who are not risk averse (or at least those who are risk neutral). Differences between medians of those who are risk averse and those who are not are tested for, and it is found that those who are risk averse deviate significantly more than those who are not, using the welfare loss measure. The effect is marginally significant when considering deviations from optimal consumption. The result

⁶From here until the end of the analysis, the full sample is used.

Table 2: ABSOLUTE CONSUMPTION DEVIATION AND WELFARE LOSS BY GROUP

	<u>All Round Types</u>		<u>Medians</u>							
	<u>ACD</u>	<u>WL</u>	<u>2U</u>		<u>2D</u>		<u>5U</u>		<u>5D</u>	
			<u>ACD</u>	<u>WL</u>	<u>ACD</u>	<u>WL</u>	<u>ACD</u>	<u>WL</u>	<u>ACD</u>	<u>WL</u>
Risk Averse	12	12638	10	15853	9	0	16.8	27909	17	9015
Risk Neutral	10.61	10627	10	19314	5	0	15.5	18549	14.9	7452
Risk Loving	10.1	5915	10	5314	4	0	16.2	11091	16.9	11247
Loss Averse	11.35	11316	10	19314	5	0	16.8	24762	16.5	9766
Loss Neutral	10.61	9805	10	10627	8	0	14.9	15829	14.9	7179
w/o experience	15	19314	14	28812	10	0	20.7	35666	22.9	15362
w/ experience	10	7212	10	8204	5	0	13.9	14517	13.2	5234

The total number of females was 32, the total number of males was 18. There were 21 subjects who had borrowing experience and 29 without. ACD = Absolute consumption deviations. WL = Welfare loss. 2U = Two period with upward sloping income. 2D = Two period with downward sloping income. 5U = Five period with upward sloping income. 5D = Five period with downward sloping income.

holds for averse subjects versus neutral subjects at the .001 significance level. What is more, this result seems to be driven by deviations in the five period rounds where borrowing is optimal. In no other types of rounds were the deviations between risk averse and risk neutral subjects statistically significant. It is of note that all but two of the rational agents were risk neutral or risk averse, perhaps explaining why these agents performed better than the group as a whole.

Another interesting pattern that emerges in the data is the interaction between loss preference and borrowing constraint. Before the end of the experiment, and before the risk elicitation, subjects were given a loss preference elicitation. Subjects were placed into two groups, those considered loss averse (risk loving in losses) and loss neutral (risk neutral in losses). The elicitation determined 27 agents to be loss averse and 23 to be loss neutral.

The analysis is started by simply comparing the distributions of deviations, using either measure, of the loss averse group to the loss neutral group. This paper find that there is no significant difference between these groups. However, this may have less to do with how behavior is affected by these preferences and more to do with the fact that many subjects who are loss averse are fairly evenly split between being risk averse and risk neutral. After controlling for risk preference, it is found that those who are loss averse are suffer significantly more welfare loss than those who are loss neutral

Overall, those who were loss averse are likely to make larger deviations from optimal behavior under both types of rounds that encourage borrowing, though not significantly. However, when the sample is restricted to those who are loss averse and look at differences in deviations from optimal behavior between those who do and do not have the borrowing constraint, it is found that an interesting result: risk averse subjects are significantly less likely to deviate from optimal behavior when they are in the presence of a borrowing

constraint than when they are not.

Though many other demographic variables are recorded in my experiment, most do not correlate highly with deviations from optimal behavior. One covariate of interest is borrowing experience. Those who have more borrowing experience have significantly lower deviations from optimal behavior than those who have no experience. This would imply that learning *before* the experiment has a far higher effect than any learning done during the experiment.

5.3 Panel Regression Analysis

After checking for treatment effects that could possibly explain deviations from optimal behavior, and finding evidence that no treatment effect entirely accounted for this behavior, a parametric panel data regressions was used to find what was influencing consumption behavior. Having exhausted most time-invariant controls in the previous analysis, a fixed-effect model is used. The model is specified as follows

$$C_{i,t} = \alpha + \beta_1 PrevAvgCons_{i,t} + \beta_2 COH_{i,t} + \beta_3' Round_t + \beta_4' Period_t + \beta_5' Round_t \times Period_t + \varepsilon_{i,t}$$

where i denotes the subject's index, t denotes the time index, C is the consumption of an agent, $PrevAvgCons$ is a moving average of previous consumption of a subject within round type, COH is the cash on hand of a subject within a round type, $Round$ and $Period$ are both dummy variables for round effects and period effects, and $Round \times Period$ is an interaction between the two variables. Cash on hand is defined as being the wealth at a given period t , or $COH_{i,t} = y_{i,t} + a_{i,(t-1)}$ where $y_{i,t}$ is the income of subject i in period t .

It is worth mentioning here that the consumption decisions of the agents in this model should be determined only by two things: their expectation of future income and their cash on hand. The expectation of future income is determined entirely by the period number, and therefore would be captured by the period fixed effect. However, since on average the income follows a process centered around a deterministic process, *a priori* it would be expected that total future income, on average, to be 50 in the two period rounds and 200 in the five period rounds. Therefore, on average, the expectation of future income should not influence the average consumption decision. Put a different way, on average our income process will follow a deterministic pattern, and therefore the optimal consumption patterns will be perfectly smooth on average. It is therefore expected that period effects would play no role in the consumption decisions of a rational, risk neutral subject on average. The coefficient of each period dummy should be zero for a sample of such subjects. The same prediction should hold for coefficients on round and interaction variables.

Similarly, for cash on hand, on average we expect income to follow a deterministic process. Since cash on hand is simply the sum of income in a period and assets from previous periods, it is expected that the coefficient on these variables to be one. Finally, the moving average of past consumption would be highly collinear if all agents were rational. However, this is not the case in practice, and the variable is included to serve as a measure of behavioral invariance. The results of these panel regressions are shown in Table 3. Each regression is broken down by treatment type.

It is evident from the results that the coefficients do not match the coefficients that would be expected if the subjects acted as expected utility maximizers. What instead is shown is

Table 3: PANEL REGRESSIONS WITH DEPENDENT VARIABLE CONSUMPTION

	2U	2D	5U	5D
consMA	0.618*** (0.0600)	0.725*** (0.0665)	0.674*** (0.0700)	0.760*** (0.0475)
COH	0.831*** (0.0287)	0.758*** (0.0331)	0.442*** (0.0378)	0.363*** (0.0380)
Round 3	-3.210 (1.7098)	- (.)	0.153 (2.1425)	-4.938 (3.0563)
Round 4	-2.605 (1.5893)	1.013 (1.2611)	2.665 (2.4729)	-3.893 (3.0913)
Round 5	-0.996 (2.2825)	3.933* (1.8415)	1.354 (2.3891)	-2.691 (3.4119)
Round 6	0.229 (1.6323)	4.319 (2.2103)	3.657 (2.5110)	-4.182 (2.7641)
Round 7	- (.)	3.856 (2.4742)	4.072 (2.6788)	- (.)
Period 2	-1.100 (1.7861)	7.833*** (1.5298)	8.072* (3.2238)	12.496** (3.9808)
Period 3	- (.)	- (.)	3.404 (2.1951)	9.504** (2.8653)
Period 4	- (.)	- (.)	-1.604 (4.0998)	5.103 (3.0846)
Period 5	- (.)	- (.)	5.231 (5.6277)	23.815*** (5.4531)
Constant	-12.335*** (2.3223)	-25.345*** (2.0927)	-16.484*** (2.4175)	-22.585*** (4.8462)
Period by Round Effects	Yes	Yes	Yes	Yes
R^2	0.788	0.653	0.818	0.703
N	600	600	1500	1500

Standard errors in parentheses. Robust standard errors clustered by subject. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Round 7 is dropped in 2U and 5D rounds due to multicollinearity. Round 3 is dropped in 2D due to multicollinearity.

that subjects, across all treatment types, typically underconsume from their wealth, given that the coefficient is significantly lower than one. Instead, subjects make decisions based on the moving average of consumption and the time period within a round. This suggests that most subjects have a strong degree of habit formation in their consumption decisions and that they base their consumption off their income draw in a particular period instead of their expectations of future income. Interestingly, the point estimates on the period dummies in many of the rounds are positive, indicating that the subjects are, on average, increasing their consumption in every round. Also shown here is that, round by round, subjects do not change their consumption patterns. This matches with previous nonparametric testing for learning effects.

The results in this section lead us to the following findings

Finding 1 *The hypothesis that longer lifecycles lead to higher levels of debt aversion is accepted.*

This is because longer lifecycles are found to significantly increase deviations from optimal consumption in five period rounds with upward sloping income compared to five period rounds with downward sloping income and to two period rounds with upward sloping income. Not only this, but status a subject's status as risk averse is correlated with large deviations from optimal behavior, especially compared to those who are risk neutral.

It is also found that loss aversion plays a role in consumption decisions. Subjects who were determined to be loss averse, like those who are determined to be risk averse, are likely to suffer from higher welfare loss than loss neutral subjects.

Finding 2 *The hypothesis that those who are loss averse and are in the presence of the natural borrowing constraint perform better than their unconstrained peers is accepted.*

Evidence for this claim is given by the constrained group's deviations from optimal behavior being significantly lower than the unconstrained group within loss averse subjects. It is also noted that the difference between median deviations of constrained and unconstrained groups is higher in five period rounds with upward sloping income than downward. This shows even though the borrowing constraint may be having a large impact in both of the five period round types, it makes a larger difference in rounds where borrowing, rather than saving, is the optimal strategy. This shows that loss aversion predominantly affects borrowing decisions.

Finding 3 *The hypothesis that subjects learn over repeated lifecycles is rejected*

Using four different measures of learning with nonparametric testing and a panel data model, it can be never be shown that subjects learn over many periods. That agents do not learn between rounds is a robust results, encountered in many other experimental intertemporal consumption behavior papers. The possible reasons for why learning does not occur is discussed more in the following section.

Finding 4 *The hypothesis that outside learning can improve performance is accepted.*

Though the decision choices of those in this sample did not seem to differ significantly with those in previous literature, an additional question relating to outside borrowing experience allowed for the sample to be split in such a way that allowed for analysis. It was found that those with previous borrowing experience showed markedly lower deviations from optimality than their unexperienced counterparts. This is no big surprise because those who have decided to either use credit cards or who have taken out student loans to finance their own education probably understand the value in smoothing consumption more than those who have never had to personally borrow.

6 Conclusion and Discussion

The problem of debt aversion, that agents deviate further from optimal behavior more in borrowing regimes than savings regimes despite total expected lifecycle income being the same between regimes, has been explored recently in experimental literature. This is the first experimental consumption paper to study two period lifecycles and natural borrowing constraints effects on intertemporal substitution. The major findings are that, among all subjects, debt aversion tends to become worse as we lengthen the number of periods. Specifically it is found that deviations from optimal behavior grow when the number of rounds is increased, due to added complexity and uncertainty in future income. The uncertainty in future income leads to underconsumption, which is driven by risk and loss aversion. That is, the less certain your future income is, the less you would like to borrow to offset against the chance that future income streams will be unable to support large amounts of current consumption relative to current wealth.

That possibility of losing part of an endowment and the great uncertainty in future income is what allows loss aversion to affect consumption decisions and drive them below an optimal point in my experiment. This is explicitly tested for by creating a treatment that prevents subjects from being able to lose their endowment. When examining differences in welfare loss between those who can and those who cannot lose their endowment, it is found that those who can have statistically significant higher welfare loss than those who cannot. Similar patterns are found when examining differences between loss averse subjects under the borrowing constraint to those without the borrowing constraint.

This experiment abstracts from real world considerations, but in many way gives the simplest case for consumption smoothing. In real life, uncertainty in future income is very high, especially for college students or recently graduated college students. This fact can be verified by income distributions of recently graduated college students available from the American Community Survey (ACS) and the National Survey of College Graduates (NSCG). If agents deviate from optimal behavior even in the best of conditions for consumption smoothing, what does it say of their ability to consumption smooth when there are a myriad more confounding factors in the real world?

Of course there are many additions to the dynamic optimization problem that are not present in this experiment. Typical extensions would be inclusion of discounting between periods, an interest rate on borrowing and saving above zero, a variable number of periods in a lifecycle, and different objective functions. It is possible that some objective functions are easier for subjects to maximize compared to others. It is possible that the CARA function

explored in this model may be too difficult for agents to maximize. It might be better to use a simpler objective function, such as log linear.

It may also be the case that as long as agents are within a certain distance of the optimal solution, that they do not care about optimizing any further, for doing so would require additional effort. Knowledge of this “optimization” premium and how it varies by the complexity of the task and also the amount of reward at stake would be valuable. Such an investigation could have wide reaching implications for anything involving study of the optimal behavior of agents.

Something theorized in some preference literature is that behavior will vary widely by the context (Simonson & Tversky 1993). For example, it may be the case that people over borrow when spending on some items and not others (college loans versus loans for luxury goods) and my experiment cannot capture the effects of what context the subject may be imagining, since my experiment is highly abstract. Something like this could be studied by performing an experiment with a different priming treatments meant to guide the subjects to frame this borrowing experiment in a variety of contexts. Such investigations are left to future research projects.

References

Ballinger, T. Parker, Michael G. Palumbo, and Nathaniel T. Wilcox. “Precautionary saving and social learning across generations: an experiment.” *The Economic Journal* 113.490 (2003): 920-947.

Ballinger, T. Parker, et al. “Saving behavior and cognitive abilities.” *Experimental Economics* 14.3 (2011): 349-374.

Brown, Alexander L., Zhikang Eric Chua, and Colin Camerer. “Learning and visceral temptation in dynamic savings experiments.” *Quarterly Journal of Economics* 124.1 (2009): 197-231.

Burdman, Pamela. “The Student Debt Dilemma: Debt Aversion as a Barrier to College Access.” *UC Berkeley: Center for Studies in Higher Education*. (2005).

Callender, Claire, and Jonathan Jackson. “Does the fear of debt deter students from higher education?” *Journal of social policy* 34.04 (2005): 509-540.

Carbone, Enrica, and John D. Hey. “The effect of unemployment on consumption: an experimental analysis.” *The Economic Journal* 114.497 (2004): 660-683.

Carbone, Enrica, and John Duffy. “Lifecycle consumption plans, social learning and external habits: Experimental evidence.” *Journal of Economic Behavior & Organization* 106 (2014): 413-427.

Carroll, Christopher D. “Lecture notes on solution methods for microeconomic dynamic stochastic optimization problems.” manuscript, Department of Economics, Johns Hopkins

University (2009).

Chetty, Raj, and Adam Looney. "Consumption smoothing and the welfare consequences of social insurance in developing economies." *Journal of Public Economics* 90.12 (2006): 2351-2356.

Erev, Ido, Eyal Ert, and Eldad Yechiam. "Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions." *Journal of Behavioral Decision Making* 21.5 (2008): 575-597.

Eckel, Catherine C., and Philip J. Grossman. "Forecasting risk attitudes: An experimental study using actual and forecast gamble choices." *Journal of Economic Behavior & Organization* 68.1 (2008): 1-17.

Field, Erica. "Educational debt burden and career choice: evidence from a financial aid experiment at NYU law school." No. w12282. National Bureau of Economic Research, 2006.

Fischbacher, Urs. "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental economics* 10.2 (2007): 171-178.

Hey, John D., and Valentino Dardanoni. "Optimal consumption under uncertainty: An experimental investigation." *The Economic Journal* (1988): 105-116.

Kahneman, Daniel, and Amos Tversky. "Prospect theory: An analysis of decision under risk." *Econometrica: Journal of the Econometric Society* (1979): 263-291.

Meissner, T. (2015), "Intertemporal consumption and debt aversion: an experimental study", *Experimental Economics*, June 12, pp. 1-18.

Tversky, Amos, and Itamar Simonson. "Context-dependent preferences." *Management science* 39.10 (1993): 1179-1189.

Tversky, Amos, and Daniel Kahneman. "Rational choice and the framing of decisions." *Journal of Business* (1986): S251-S278.

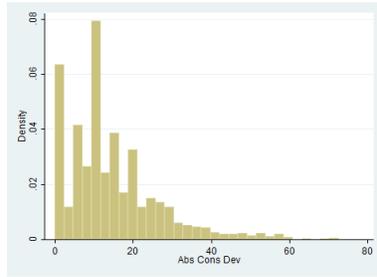
Tversky, Amos, and Daniel Kahneman. "Advances in prospect theory: Cumulative representation of uncertainty." *Journal of Risk and Uncertainty* 5.4 (1992): 297-323.

APPENDIX

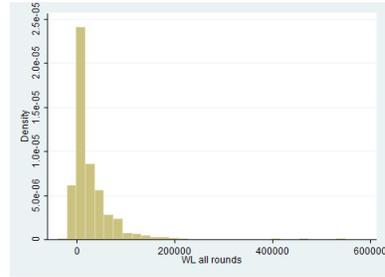
Figure 3: HISTOGRAMS OF DEVIATIONS BY ROUND TYPE

Absolute Consumption Deviation

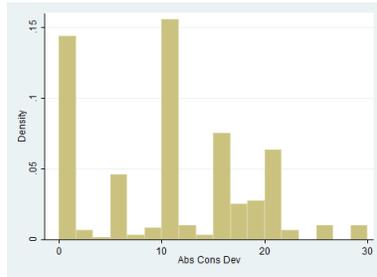
Welfare Loss



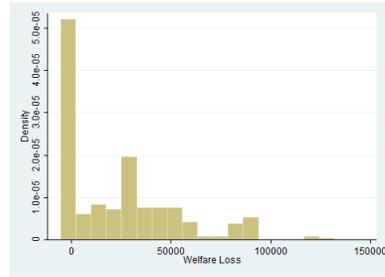
(a) All Rounds



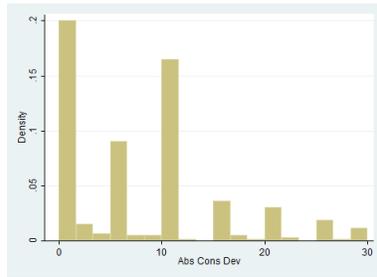
(b) All Rounds



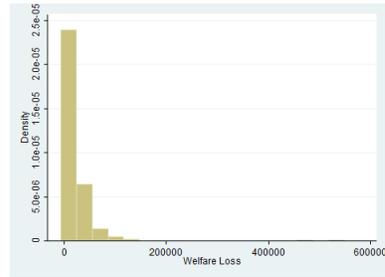
(c) 2U



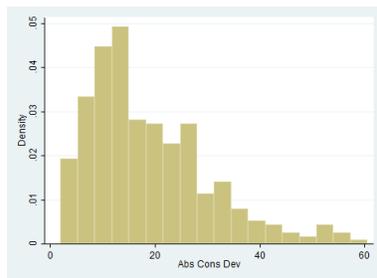
(d) 2U



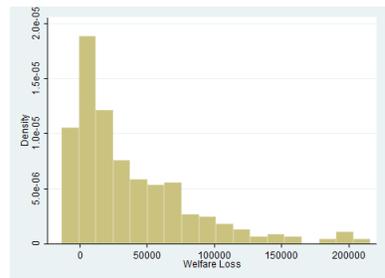
(e) 2D



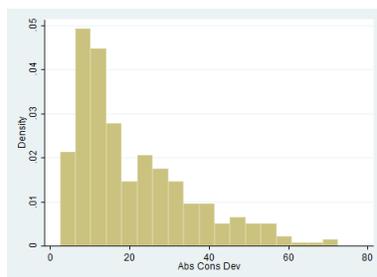
(f) 2D



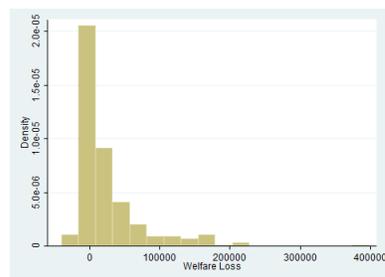
(g) 5U



(h) 5U

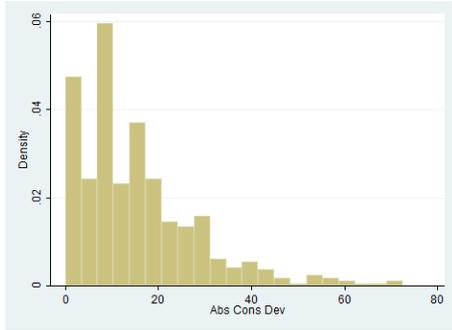


(i) 5D

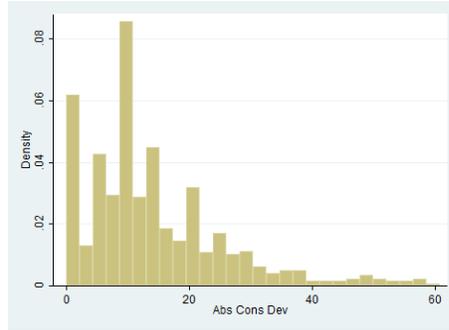


(j) 5D

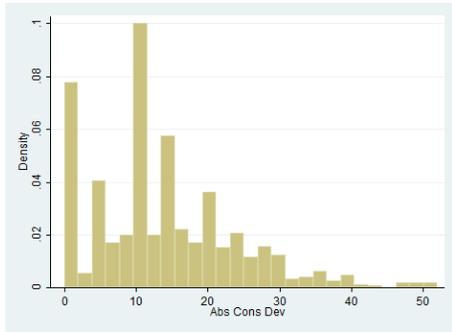
Figure 4: ABSOLUTE CONSUMPTION DEVIATIONS BY GROUP



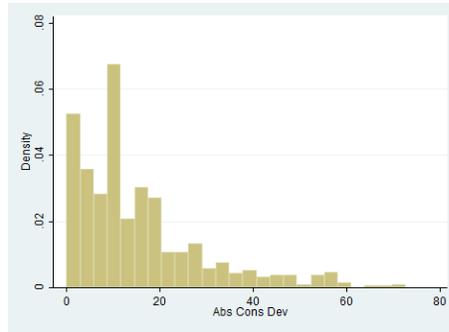
(a) Risk Averse



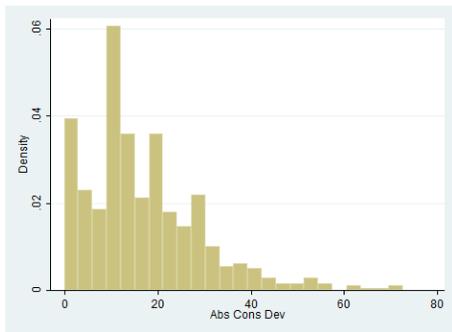
(b) Risk Neutral



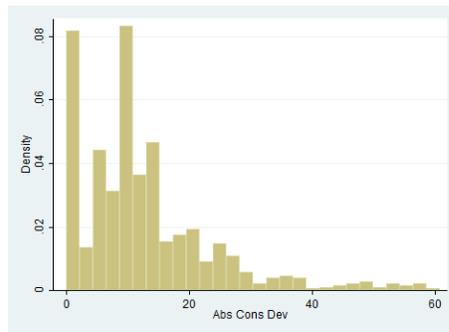
(c) Loss Averse



(d) Loss Neutral

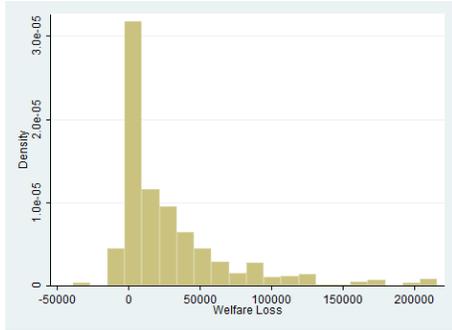


(e) w/o borrow exp

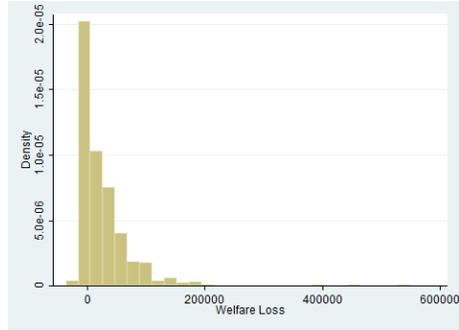


(f) w/ borrow exp

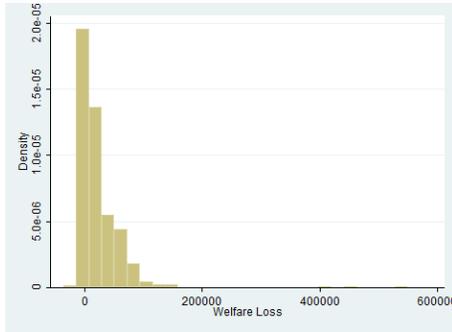
Figure 5: WELFARE LOSS BY GROUP



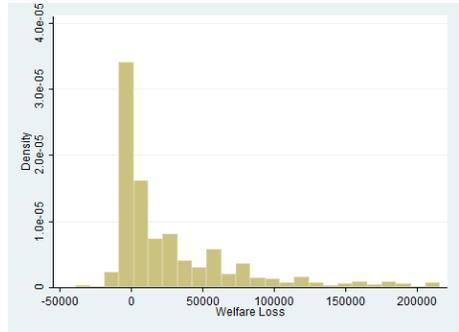
(a) Risk Averse



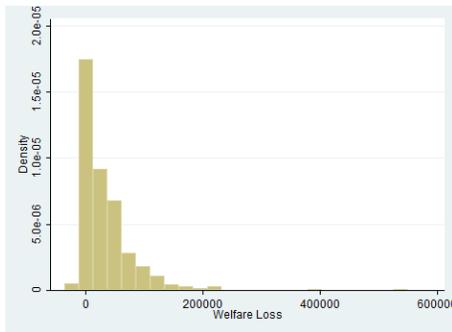
(b) Risk Neutral



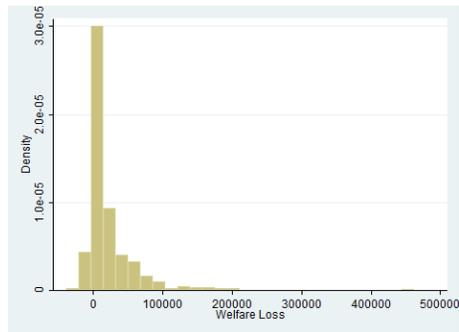
(c) Loss Averse



(d) Loss Neutral



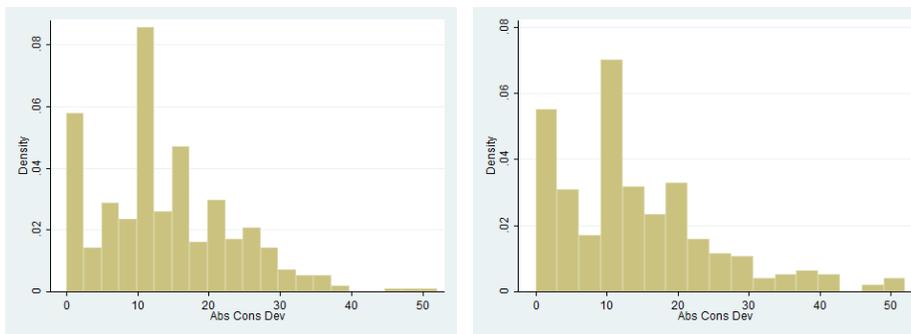
(e) w/o borrow exp



(f) w/ borrow exp

Figure 6: DISTRIBUTION OF DEVIATIONS OF LOSS AVERSE BY BORROWING CONSTRAINT

Absolute Consumption Deviation



Welfare Loss

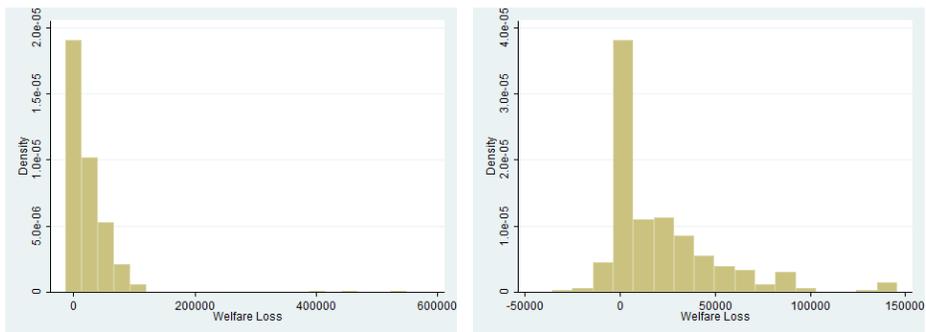
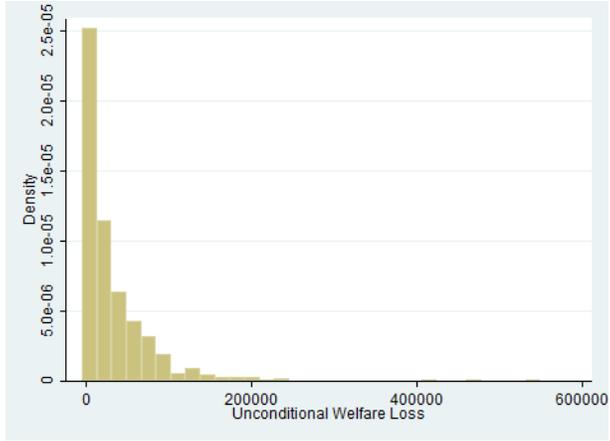
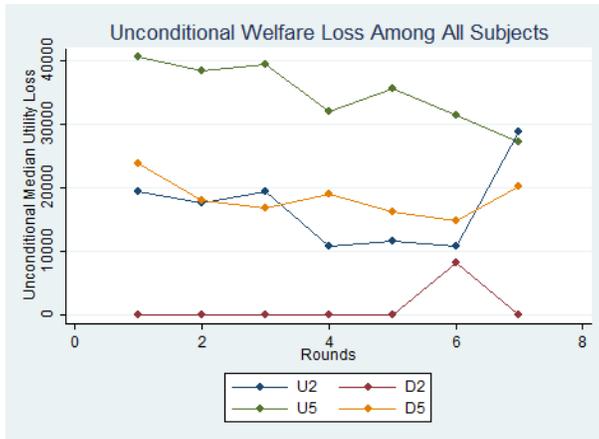


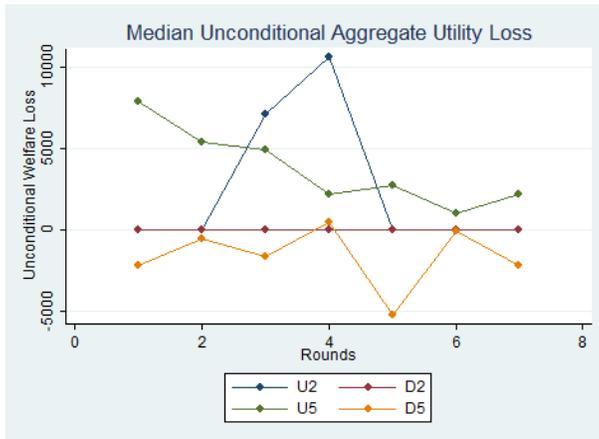
Figure 7: UNCONDITIONAL WELFARE LOSS



(a) Histogram of unconditional welfare loss



(b) Unconditional welfare loss of all agents



(c) Unconditional welfare loss of rational agents

Figure 8: CHOICE SCREENS

Round
(Practice) 1 out of 4

You are in a 2 period round with type 1 income.
This is round number 1.
This is period 1 of 2 in your current lifecycle.

Your Income this period: 10 Tickets

How many Tickets do you want to convert to Points this period?

Confirm

(a) Example of a choice screen

Round
(Practice) 1 out of 4

You are in a 2 period round with type 1 income.
This is round number 1.
This is period 2 of 2 in your current lifecycle.

Period	Income	Consumption	Savings/Borrowings	Points
1	10.00	30.00	-20.00	69881.00
2	20.00	0.00	0.00	0.00

In the final period, your last period income is added to your total savings/borrowing, and that number of Tickets is converted automatically.

Continue

(b) Example of a period review screen

Round

(Practice) 1 out of 4

You are in a 2 period round with type 1 income.
This is round number 1.
This is period 2 of 2 in your current lifecycle.

Round	Total Income	Total Points
1	50.00	154948.00

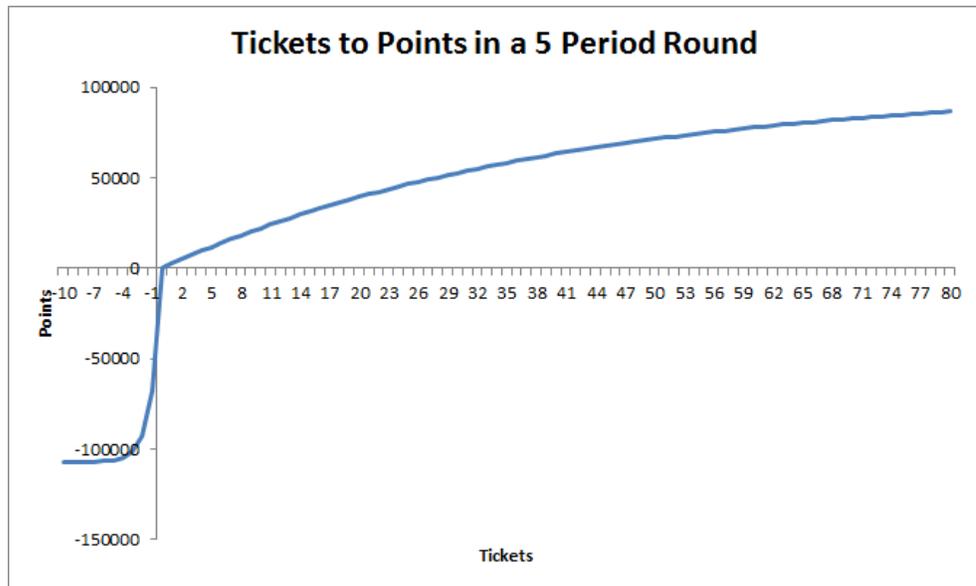
Continue

(c) Example of a round review screen

LABORATORY HANDOUT

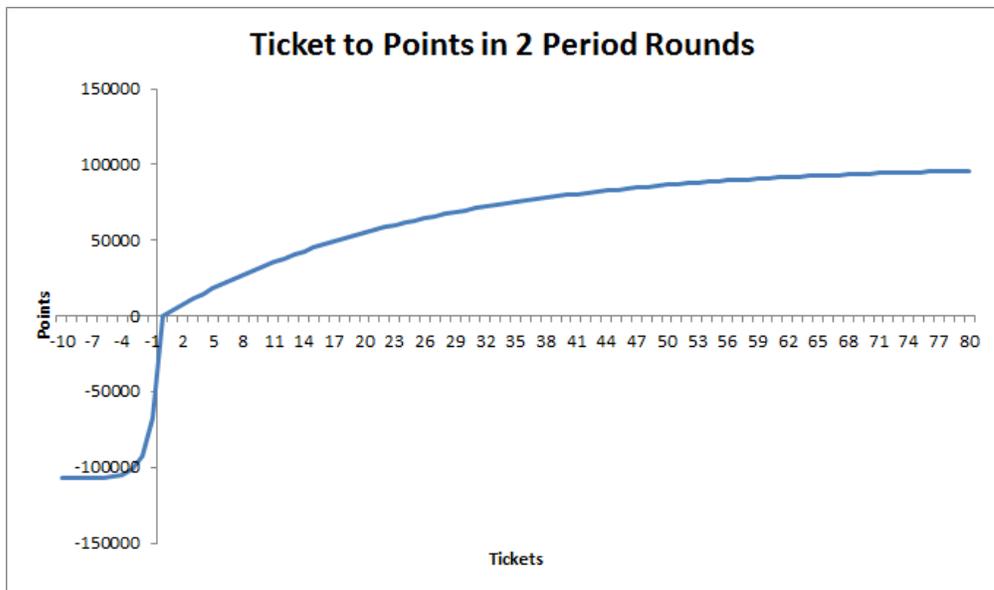
LABORATORY INSTRUCTIONS

Tickets to Points Conversion for Five Period Rounds



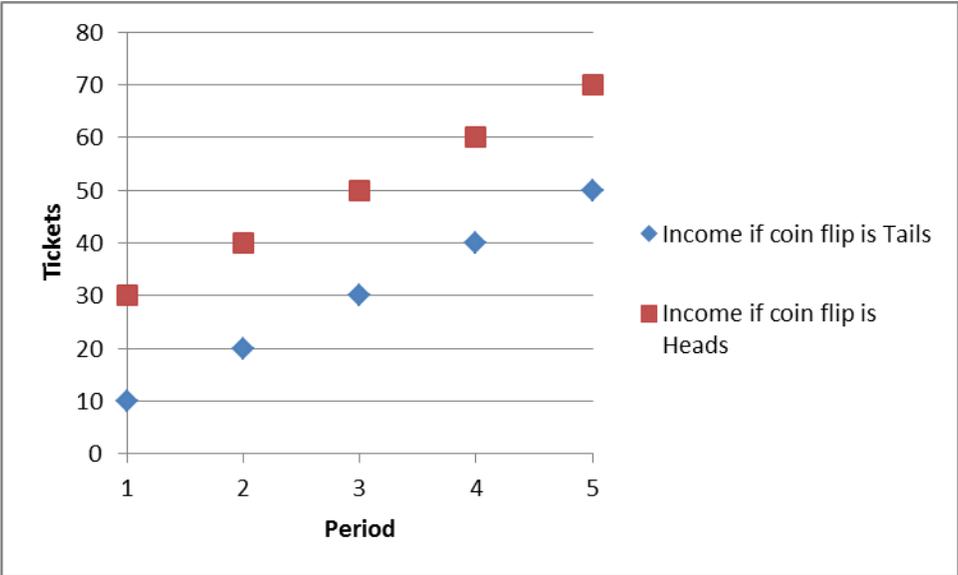
Tickets	Points	Tickets	Points
-10	-107020	6	13929.2
-9	-107012	7	16054.3
-8	-106989	8	18126.92
-7	-106927	9	20148.38
-6	-106760	10	22119.92
-5	-106304	20	39346.93
-4	-105065	30	52763.34
-3	-101697	40	63212.06
-2	-92540.7	50	71349.52
-1	-67652.7	60	77686.98
0	0	70	82622.61
1	2469.009	80	86466.47
2	4877.058		
3	7225.651		
4	9516.258		
5	11750.31		

Tickets to Points Conversion for Two Period Rounds



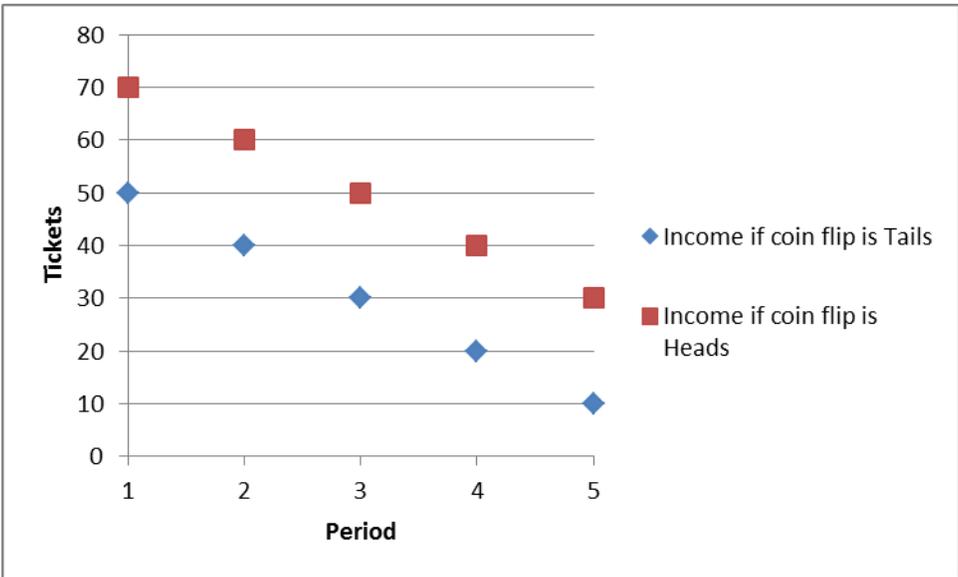
Tickets	Points	Tickets	Points
-10	-107020	6	21337.21
-9	-107012	7	24421.63
-8	-106989	8	27385.1
-7	-106927	9	30232.37
-6	-106760	10	32968
-5	-106304	20	55067.1
-4	-105065	30	69880.58
-3	-101697	40	79810.35
-2	-92540.7	50	86466.47
-1	-67652.7	60	90928.2
0	0	70	93918.99
1	3921.056	80	95923.78
2	7688.365		
3	11307.96		
4	14785.62		
5	18126.92		

5 Period Round with Type 1 Income



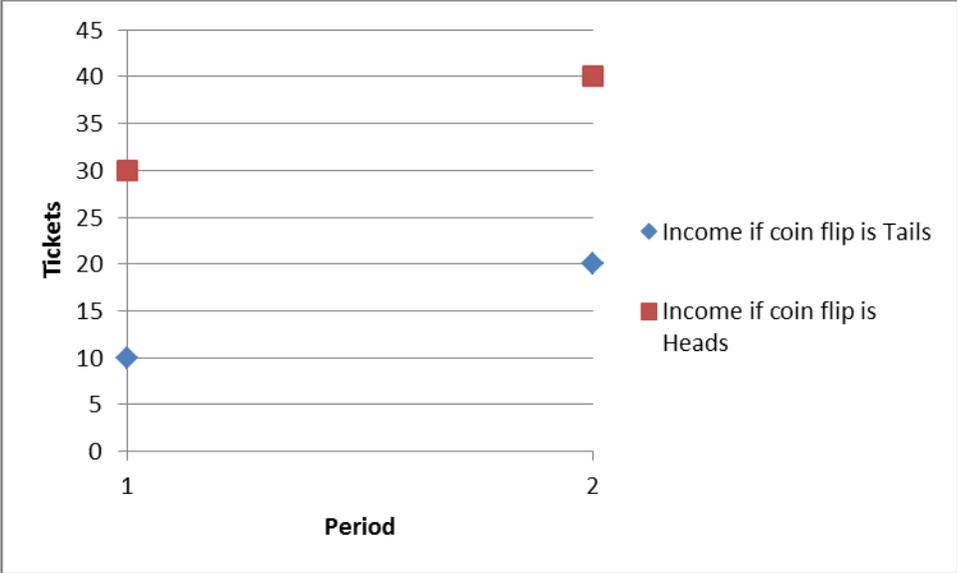
Period	Income if coin flip is Tails	Income if coin flip is Heads
1	10	30
2	20	40
3	30	50
4	40	60
5	50	70

5 Period Round with Type 2 Income



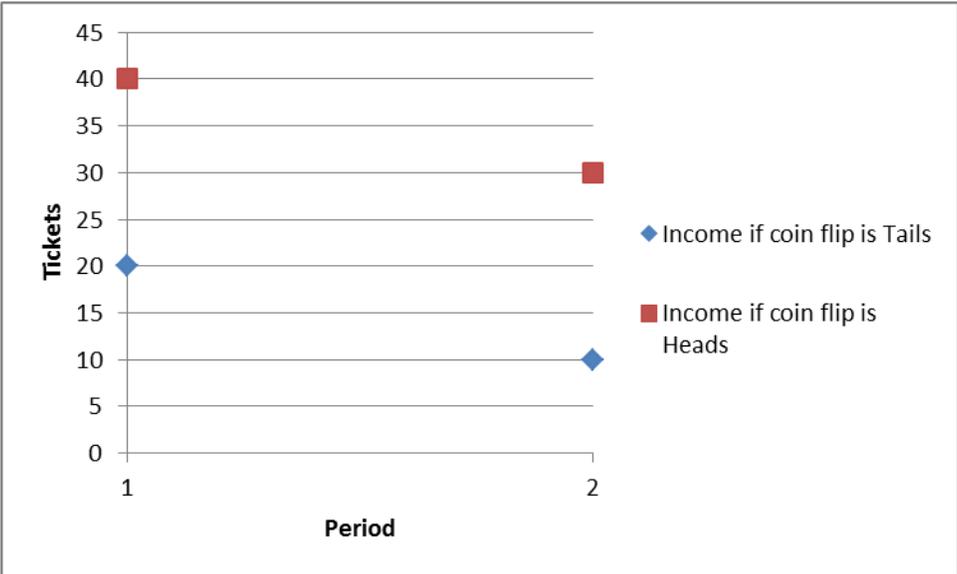
Period	Income if coin flip is Tails	Income if coin flip is Heads
1	50	70
2	40	60
3	30	50
4	20	40
5	10	30

2 Period Round with Type 1 Income



Period	Income if coin flip is Tails	Income if coin flip is Heads
1	10	30
2	20	40

2 Period Round with Type 2 Income



Period	Income if coin flip is Tails	Income if coin flip is Heads
1	20	40
2	10	30