

Meissner, Thomas

Conference Paper

Intertemporal Consumption and Debt Aversion: An Experimental Study

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2014: Evidenzbasierte Wirtschaftspolitik - Session: Capital Accumulation, No. F05-V1

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Meissner, Thomas (2014) : Intertemporal Consumption and Debt Aversion: An Experimental Study, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2014: Evidenzbasierte Wirtschaftspolitik - Session: Capital Accumulation, No. F05-V1, ZBW - Deutsche Zentralbibliothek für Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft, Kiel und Hamburg

This Version is available at:

<http://hdl.handle.net/10419/100522>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Intertemporal Consumption and Debt Aversion: An Experimental Study*

Thomas Meissner[†]

October 4, 2013

Abstract

This paper tests how subjects behave in an intertemporal consumption/saving experiment when borrowing is allowed and whether subjects treat debt differently than savings. Two treatments create environments where either saving or borrowing is required for optimal consumption. Since both treatments share the same optimal consumption levels, actual consumption choices can be directly compared across treatments. The experimental findings imply that deviations from optimal behavior are higher when subjects have to borrow than when they have to save in order to consume optimally, suggesting debt-aversion. Significant under-consumption is observed when subjects have to borrow in order to reach optimal consumption. Only weak evidence is found suggesting that subjects over-consume when saving is necessary for optimal consumption.

JEL-Classification: C91, D81, E21

Keywords: Laboratory Experiment, Intertemporal Consumption, Consumption Smoothing, Debt Aversion

*I am grateful for valuable comments by Frank Heinemann, Dorothea Kübler, Dietmar Fehr and participants of the TU/WZB Colloquium in Berlin. I also thank Satpal Nijjar for help in programming the experiment. Financial support from the Deutsche Forschungsgemeinschaft (DFG) through CRC 649 “Economic Risk” is gratefully acknowledged.

[†]Technische Universität Berlin, Chair of Macroeconomics (H 52), Straße des 17. Juni 135, 10623 Berlin, Germany. E-mail: meissnet@gmail.com

1 Introduction

The question how people cope with stochastic dynamic optimization problems, such as posed by intertemporal consumption/savings problems, has been repeatedly tested in economic experiments. An early reference is Hey and Dardanoni (1988). Ballinger et al. (2003, 2011), Carbone and Hey (2004) as well as Brown et al. (2009) further substantiate the experimental analysis of intertemporal consumption/saving problems. The main finding of this literature is that subjects (somewhat unsurprisingly) deviate from optimal behavior. In particular it is observed that subjects tend to over-consume in early periods, i.e. save too little compared to the optimal solution and that consumption tracks income too closely. When subjects are given the opportunity to learn (either privately through repetition or socially through observing other life-cycle consumption decisions) consumption decisions improve towards optimality. With the exception of Fenig et al. (2013) all existing experimental studies on dynamic intertemporal consumption/savings problems have strict borrowing constraints. This implies that only one side of the problem has been examined, i.e. whether subjects *save* optimally. I believe that the question whether individuals borrow optimally is of similar interest to the understanding of intertemporal consumption decisions.¹ The experiment conducted for this paper addresses two main research questions: first, how do people behave in a life-cycle consumption experiment when borrowing is allowed? Second, do people treat savings systematically different than debt with respect to consumption smoothing?

The experiment is based on a simple discrete-time, finite horizon life-cycle model of consumption with no discounting and no interest paid on savings or debt. In order to compare borrowing and saving behavior this experiment has two treatments, one in which the income stream has an increasing trend and one in which it has a decreasing trend. Since the experimental environment creates an incentive to smooth consumption, the optimal solution with an increasing income stream requires borrowing and the optimal solution with a decreasing income stream requires saving. Moreover, since the specifics of the underlying model were chosen in such a way that optimal consumption is the same in both treatments, this allows to directly compare borrowing and saving behavior. This is a novel feature compared to previous experimental work where mostly

¹The experiment by Fenig et al. (2013) has no borrowing constraints. However they do not specifically analyze borrowing behavior, but introduce and analyze asset markets in an experimental general equilibrium economy. In one treatment they find, however, that ruling out borrowing in order to buy assets does not affect asset prices.

stationary income processes are used, which does not allow for such a comparison.²

With an increasing income stream, the experimental data suggests that subjects significantly underconsume, that is they do not borrow enough compared to the optimal solution. This stands in contrast to evidence from previous experiments on intertemporal consumption/savings problems, where overconsumption is typically observed. In line with previous findings, however, weak evidence is found for overconsumption in the treatment with an increasing income stream, implying that subjects do not save enough relative to the optimal solution. Comparing consumption decisions across treatments reveals that deviations from optimal behavior are higher in the treatment with an increasing income stream than in the treatment with a decreasing income stream. This implies that subjects are less willing to borrow than to save in order to smooth consumption, i.e. they are debt averse.

The experimental literature on intertemporal consumption/savings problems concentrates on two potential explanations for suboptimal behavior. One explanation is bounded rationality or cognitive constraints. Ballinger et al. (2011) for example find that cognitive abilities are the best predictors of saving performance. Moreover there exists substantial evidence that subjects typically learn to improve their consumption decisions over time. This points towards bounded rationality since with perfect rationality there would be no room for improvement. In this experiment cognitive constraints certainly explain some deviations from optimal behavior. However they cannot explain systematic differences in deviations from optimal behavior in the two treatments of this experiment, since the experimental task is equally demanding in both treatments. This is because subjects face essentially the same optimization problem in both treatments, as debt is mathematically equivalent to savings only with a different sign.

Another explanation for suboptimal behavior is a preference for immediacy or present bias. Brown et al. (2009) use immediate rewards in an intertemporal consumption/saving experiment. They find that behavior is consistent with both hyperbolic discounting and dual self models and not with exponential discounting. In this experiment, however, present bias should not play a role since subjects get their reward at only one point in time – at the end of the experiment.

The existing literature on debt aversion primarily analyzes its role in the light of the decision to invest in higher education. Eckel et al. (2007) use survey measures to assess

²Brown et al. (2009) use an income stream that is increasing in expectations. However, their experiment also uses habit formation utility and therefore saving is optimal in early periods even with an increasing income stream.

debt aversion and relate this data to lottery choice experiments. They find that debt aversion has little to no impact on the demand for loans for postsecondary education. In field experiments, Field (2009) as well as Caetano et al. (2011) offer differently labeled loan contracts to students and find that debt aversion does influence career and education investment decisions. However, all these studies analyze debt aversion in the context of student loans. This experiment offers a more general approach, allowing to identify debt aversion in the context of a stochastic dynamic optimization problem, which in one way or another serves as a basis for household behavior in most modern macroeconomic models.

The remainder of the paper is organized as follows. The next section provides the theory and experimental design. Section 3 contains the data analysis and Section 4 concludes.

2 Theory and Experimental Design

The experiment described in the following section aims to implement a particular version of the life-cycle model of consumption. One experimental life-cycle lasts for 20 periods. In each period $t = (1, \dots, 20)$, subjects decide how much to consume (c_t), and implicitly how much to save or borrow. There is no discounting and no interest is paid on savings or debt. Period income y_t follows a trend stationary stochastic process. The intertemporal budget constraint requires period consumption plus period savings (a_t) to equal period wealth, which is defined as $w_t = y_t + a_{t-1}$. Period savings are allowed to be both positive and negative. In the last period, remaining wealth must be consumed, i.e. saving for the (non-existent) next period is not possible. Initial savings a_0 were normalized to zero.³ Induced preferences are given by a time-separable CARA utility function: $u(c_t) = 250(1 - e^{-0.02c_t})$.⁴ The subjects' objective was to choose consumption in every period in order maximize the expected utility of life-time consumption.

³Note that together with the ending condition ($a_{20} = 0$), this also implies that total consumption must equal total income: $\sum_{t=1}^{20} c_t = \sum_{t=1}^{20} y_t$.

⁴CARA utility was chosen because this class of utility functions is defined on the negative domain. Why this is of importance will be explained later in this section. However, other papers in the literature on life-cycle consumption/savings problems also make use of CARA utility. See, for instance Carbone and Hey (2004)

The decision problem subjects face at any period τ can be written as:

$$\max_{c_t} E_\tau \sum_{t=\tau}^{20} u(c_t) \quad (1)$$

$$\text{s.t. } c_t + a_t = w_t, \quad (2)$$

$$a_{20} = 0, \quad a_0 = 0. \quad (3)$$

The solution method to this sort of problem is dynamic programming using backward induction. The specifics of the decision problem were chosen in order to keep the problem as traceable as possible for subjects without making it trivial. However, the solution is still computationally difficult. In particular, due to the finite horizon of the optimization problem, the solution cannot be simplified by exploiting the property of stationarity. The solution to this particular optimization problem was implemented numerically using the software Mathematica and following Carroll's (2009) solution methods for dynamic stochastic optimization problems. Without going into detail, optimal consumption policy functions for every period were calculated via backward induction, as functions of period wealth. With these policy functions optimal consumption levels were calculated starting from the first period and then iterating forward.

The experiment has two base treatments that differ only with respect to the income process. In one treatment the income process is increasing over time, in the other one it is decreasing. Intuitively, the optimal solution with an increasing income process includes borrowing ($a < 0$) whereas with a decreasing income process saving ($a > 0$) is necessary for optimal consumption. Therefore I will refer to the treatment with the increasing income stream as the *Borrowing* treatment and to the treatment with a decreasing income stream as the *Saving* treatment. The income process for the Borrowing treatment is given by: $y_t^B = 10t + \theta_t$, where θ_t takes the values of +10 or -10, each with probability of one half. For the Saving treatment the income process is: $y_t^S = 210 - 10t + \theta_t$. Note that the expected lifetime income over 20 periods is equal across treatments. Since the stochastic component is equivalent at any time for both treatments and income can be shifted costlessly across periods, this implies that optimal consumption plans are equivalent in both treatments. The experimental design therefore allows to directly compare consumption behavior in both treatments (between subject) and thus to examine whether there exist systematic differences between saving and borrowing behavior.

In order to assess learning effects, the experiment is repeated for six experimental life cycles (rounds), with different income realizations in each round. Subjects first play three repetitions of either the Borrowing or the Saving treatment. After three repetitions roles are changed. Subjects who started with the Borrowing treatment now play three repetitions of the Saving treatment and vice versa.⁵ This role change was implemented in order to analyze whether subjects are able to transfer their learning from one treatment to the other. Moreover, it adds a within subject dimension to the analysis of the treatment effect. In the following I will refer to the condition that starts with the Borrowing treatment as *Borrowing First* (BF) and to the condition that starts with the Saving treatment as *Saving First* (SF).

Note that the CARA utility function used in this experiment is defined on the negative domain. In fact, negative consumption was and had to be possible in order to enforce borrowing. However, since income is always positive, negative consumption is never part of the optimal solution. The fact that out of 9120 consumption decisions only 24 or 0.3% were negative, illustrates that this was generally understood by the subjects. How negative consumption was explained to subjects is part of the next subsection.

2.1 Experimental Procedures

The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007). Consumption was explained to the subjects as buying “points” by spending “Taler”, the experimental currency in which income was denoted. Experimental currency was converted to points by the utility function specified above. Subjects were informed about the exact form of the utility function. Furthermore they were given a graph of the function and a table with relevant function values. The advantage of framing consumption as buying points is that negative consumption can be explained as the selling of points in return for experimental currency.

Consumption decisions were input in an interface that displayed period income, savings from the last period along with the history of all previous decisions and relevant values, such as savings, wealth, the number of purchased points, and accumulated points. Before a consumption decision was implemented, subjects were informed about how their chosen consumption level translates into points and the amount of savings that would be available in the next period. After this information was displayed subjects

⁵Subjects did not know about the role change when playing the first three repetitions. They received new instructions after the third round.

had the opportunity to start over, i.e. specify another potential level of consumption, and check the implications of that consumption level. In the final period of each life cycle, the program automatically spent that period's wealth as consumption. At any period during the experiment, subjects knew about the amount of points purchased so far in the current round. At the end of every round subjects were again informed about the amount of points they purchased during the round. At the end of the experiment, two of the six experimental life cycles were randomly chosen to be payoff relevant, one from the first three life cycles and one from the last three life cycles. Subjects' payoffs were determined by a preannounced linear function of the amount of points purchased in the two relevant rounds. The minimal payoff was fixed at 5 €. Subjects earned an average of 19.23 €.

The experiment was conducted at the laboratory of the Technical University of Berlin. A total of 76 subjects participated, 38 in the Borrowing First condition and 38 in the Saving First condition. Most of the subjects were undergraduate students in the field of economics or engineering.

3 Results

The analysis of consumption decisions in the experiment is split into three subsections. In subsection 3.1, I will compare observed consumption and saving decisions to optimal consumption. There are two ways to measure deviations from optimal behavior. Actual consumption levels can be compared to consumption levels that would be optimal, given a subject has always behaved optimally in the past, i.e. has optimal wealth levels in any period. I will refer to this consumption level as *unconditionally* optimal consumption. Another way is to compare actual consumption levels to the consumption level that would be optimal given the actual wealth levels that subjects have at the considered period of the experiment. This wealth level can deviate from the wealth level that would be optimal given the income realizations, if subjects have not behaved optimally in the past. I will refer to this consumption level as *conditionally* optimal consumption. Like Carbone and Hey (2004), I will make use of both concepts, however I prefer the latter one for the following reason: an error made in one period causes the wealth level in the following period to deviate from the unconditionally optimal level. Since optimal consumption is a function of present wealth, optimal consumption in the next period will also deviate from unconditional optimal consumption. An optimal

response to the present wealth level in the next period would thus be counted as a deviation from unconditionally optimal behavior. Furthermore, an error made in any period causes optimal consumption to deviate from unconditionally optimal levels in *all* following periods, even if a subject behaved optimally from that period onwards. Comparing actual consumption levels to unconditionally optimal levels therefore results in a multiple counting of one error and thus leading to a disproportionate weighting of early errors compared to later errors.

In subsection 3.2, I will discuss the theoretical implications of myopia in this experiment and estimate subjects' apparent planning horizons. In subsection 3.3, I will examine learning effects, i.e. whether subjects improve their decisionmaking over the 6 experimental life cycles.

3.1 Deviations from Optimal Behavior

I will start by analyzing directional deviations from optimal consumption, i.e. whether subjects consume too much or too little compared to optimal consumption.⁶ A measure for directional deviations from optimal consumption used by Ballinger et al. (2003) is given by *Measure 1*: $m_1 = \sum_{t=1}^{20} (c_t^*(w_t) - c_t)$, where $c_t^*(w_t)$ is conditionally optimal consumption (depending on current wealth w_t) and c_t is observed consumption in period t . This is the sum of deviations from conditionally optimal consumption for one subject and over one experimental life-cycle. For the sake of legibility I omit indices for subjects and rounds. Figure 1 provides a graphical representation of the median measure m_1 of all subjects from the BF condition (solid line) and the SF condition (dashed line) for rounds 1 to 6.⁷ The grey error bars represent the 95% confidence intervals (from a Wilcoxon signed-rank test).

Whenever subjects faced an increasing income stream, i.e. in round 1 to 3 in the BF condition, as well as in round 4 to 6 in the SF condition, m_1 is significantly positive. This implies that subjects on average consume less than optimal in each of these rounds. Note that this result is contrary to previous experimental studies, where subjects are found to consume too much relative to optimal levels. A potential explanation could be that with an increasing income stream, borrowing is necessary for optimal consumption.

⁶Note that consumption and savings are two sides of one coin: a subject who consumes too little also saves too much and vice versa.

⁷I use the median as a descriptive statistic because I use non-parametric tests, such as the Mann-Whitney U Test to compare treatments. These tests typically provide information about the median of a random sample.




Figure 1: median aggregate deviation from conditional optimal consumption. Solid and dashed lines represent median values from the BF and SF condition respectively.

Since previous experimental studies did not allow borrowing, their findings might not hold for borrowing behavior. The evidence is not as clear concerning deviations from optimality when subjects face a decreasing income stream, i.e. in round 1 to 3 in the SF condition and in round 4 to 6 in the BF condition. All median m_1 values are negative, however only three of them significantly so. This implies that subjects (weakly) overconsume relative to optimal levels, which is in line with previous findings.

The measure introduced above gives account of whether subjects mostly consume more or less than is optimal throughout one experimental life-cycle. However, it is not a very good measure of performance, since it potentially conceals suboptimal behavior. This is the case, for instance, when positive and negative deviations cancel each other out over one experimental life-cycle. Even though the median deviation is close to zero, as observed in rounds 1 to 3 in the SF condition, behavior might still be far from optimal. Therefore I also examine *absolute* deviations from conditional optimal behavior by use of *Measure 2*, $m_2 = \sum_{20}^{t-1} |c_t^*(w_t) - c_t|$. Figure 2(a) illustrates the evolution of Measure 2 over the six experimental life-cycles. Again, the solid line represents the median m_2 over all subjects in the BF condition and the dashed line represents the median m_2 over

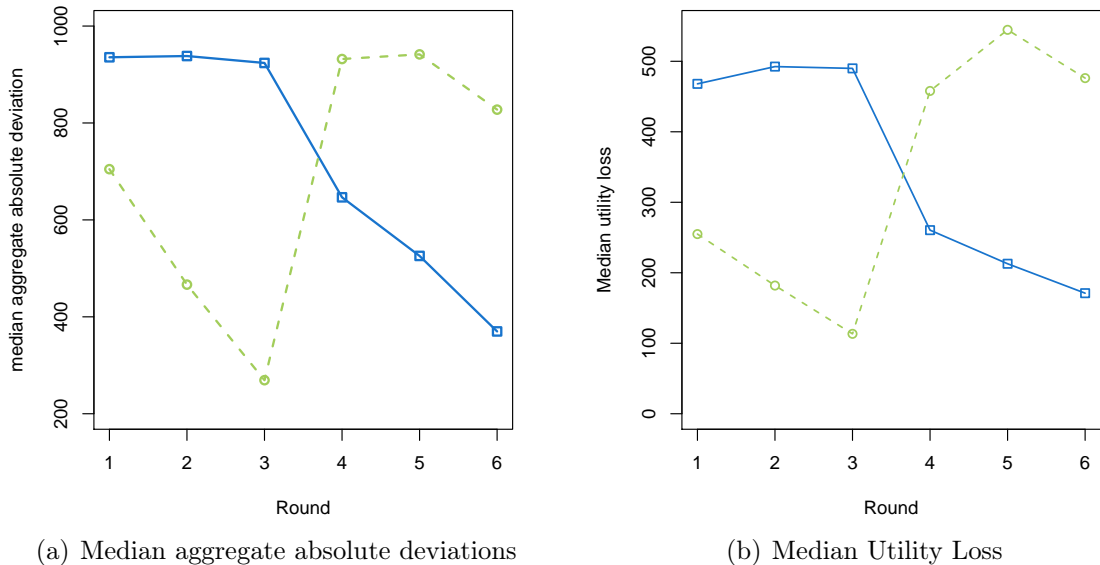


Figure 2: Median aggregate absolute deviations from conditionally optimal consumption and Median utility loss. Solid and dashed lines represent median values from the BF and SF condition respectively.

all subjects in the SF condition.⁸ The median levels of m_2 are higher when subjects face an increasing income stream than when they face a decreasing income stream. This relationship is significant for the first three rounds (p-values from a Mann-Whitney U test are provided in Table 1). This suggests that absolute deviations from conditionally optimal behavior are higher when subjects have to borrow in order to consume optimally than when they have to save in order to consume optimally.

As discussed in Section 2, all subjects observe the same income shock realizations over the course of the experiment. Furthermore subjects from both the BF and SF condition share the same unconditional optimal consumption levels at any period and round during the experiment. Therefore I can directly compare deviations from unconditional optimal behavior across treatments. *Measure 3* is $m_3 = \sum_{t=1}^{20} (u(c_t^*(w_t^*)) - u(c_t))$, where $c_t^*(w_t^*)$ denotes unconditionally optimal consumption at period t as a function of optimal wealth w_t^* . This measure can be interpreted as the utility loss that results from suboptimal consumption in one experimental life-cycle. Using conditionally optimal consumption utility as a benchmark yields similar results, however the interpretation is

⁸Note that for Measure 1 and Measure 2, *conditional* optimal consumption is the only meaningful benchmark. With unconditionally optimal consumption as a benchmark this measure would always be zero, since $\sum_{t=1}^{20} c_t^*(w_t^*) = \sum_{t=1}^{20} c_t$, where $c_t^*(w_t^*)$ denotes unconditionally optimal consumption at period t as a function of optimal wealth w_t^* , always holds.

not as straightforward. Figure 2(b) summarizes the median utility loss for the BF and SF conditions. Similar to the previous results, utility loss is higher for subjects in the Borrowing treatments compared to subjects in the Saving treatments. This difference is significant in rounds one and three and marginally significant ($p=0.058$) in round two.⁹

Summing up, the data suggests that there are substantial deviations from optimal behavior. In particular, subjects consume less than is optimal when they face an increasing income stream. Only weak evidence exists for overconsumption when the income stream is decreasing. There appears to be a difference in deviations between the Borrowing and Saving treatments. Deviations from optimal consumption are higher in the Borrowing treatment than in the Saving treatment. This implies that subjects are less willing to borrow than they are willing to save in order to smooth consumption.

3.2 Myopia

Some authors suggest that sub-optimal consumption choices can be explained with myopia, a particular form of bounded rationality. In order to behave optimally in this experiment, subjects have to solve a backward induction problem that has 20 induction steps. Since this is a very complex problem it comes as no surprise that subjects fail to behave optimally. Ballinger et al. (2003) suggest that subjects may not take the whole experimental time horizon in account, but simplify the problem by shortening their planning horizon, i.e. they act as if only h periods remain. In this experiment, the short-horizon (myopic) optimal solution generally deviates from the full horizon optimal solution, unless the remaining periods are less or equal the planning horizon.

Myopia is an appealing model of behavior since it can not only explain deviations from optimal behavior, but in this experimental environment also makes clear predictions concerning the direction of deviations from optimal behavior. Myopic subjects should over-consume in the Saving treatment and under-consume in the Borrowing treatment. In the Saving treatment, a shorter planning horizon implies that subjects do not take later periods into account, where income is lower. Hence they underestimate the need to save for later periods and consume more than is optimal. Analogously, in the Borrowing treatment a shorter planning horizon implies that later, higher income

⁹Subjects may perceive the role change after round three to reveal new information about the structure and purpose of the experiment, which may lead them to reconsider their strategy. Therefore, the fact that the differences in m_2 and m_3 cease to be significant after the role change is not overly surprising.

measure	condition	Round					
		1	2	3	4	5	6
median m_1	BF	922.323	932.144	890.734	-67.704	-290.327	-260.765
	SF	-133.471	-90.488	-89.290	929.913	940.457	761.357
mean m_1	BF	1047.17	860.227	790.387	-168.05	-236.162	-246.414
	SF	-6.668	31.040	-68.498	848.954	838.205	676.664
		<i><0.01</i>	<i><0.01</i>	<i><0.01</i>	<i><0.01</i>	<i><0.01</i>	<i><0.01</i>
median m_2	BF	935.539	938.128	923.855	646.596	525.661	369.850
	SF	704.687	466.496	269.237	931.966	941.558	827.562
mean m_2	BF	1183.6	936.5	884.476	721.94	628.343	600.024
	SF	875.578	737.92	590.045	906.195	942.792	798.939
		<i>0.017</i>	<i><0.01</i>	<i><0.01</i>	<i>0.19</i>	<i>0.102</i>	<i>0.226</i>
median m_3	BF	468.003	492.493	489.941	260.467	212.786	171.035
	SF	254.895	181.871	113.383	457.953	544.64	476.158
mean m_3	BF	2993.83	547.786	491.923	412.377	323.953	337.165
	SF	674.768	379.117	302.58	494.125	545.743	479.995
		<i>0.041</i>	<i>0.058</i>	<i><0.01</i>	<i>0.445</i>	<i>0.144</i>	<i>0.403</i>
median h^*	BF	4	4	6	20	18	19
	SF	19.5	18.5	20	4.5	6	9
mean h^*	BF	6.711	7.263	8.711	13.684	12.842	14.316
	SF	13.921	13.895	14.711	8.105	9.605	10.184
		<i><0.01</i>	<i><0.01</i>	<i><0.01</i>	<i><0.01</i>	<i>0.057</i>	<i>0.029</i>

Table 1: Measures 1-3 and apparent planning horizon h^* . P-values come from a Mann-Whitney Test and are printed in italic letters. The null hypothesis is that samples are stochastically equivalent.

periods are neglected. Hence, myopic individuals do not borrow enough and therefore consume less than is optimal.

As discussed earlier, significant underconsumption is observed in the Borrowing treatment and weak evidence is found for overconsumption in the Saving treatment. Myopia can furthermore explain certain differences in deviations from *conditional* optimal behavior between the Borrowing and the Saving treatment. In particular, when a positive income shock occurs, the absolute deviation from (full horizon) optimal behavior of a myopic agent will be higher in the Borrowing treatment than in the Saving treatment. Vice versa, when a negative income shock occurs, the deviation is higher in the Saving treatment than in the Borrowing treatment.¹⁰ One of the main results from subsection 3.1 is that deviations from optimal behavior are higher for subjects in the Borrowing treatment than for subjects in the Saving treatment. It therefore seems necessary to examine whether this result can be explained by myopia. To address this issue I estimate the subjects' apparent planning horizons and check whether planning horizons are consistent across treatments. The apparent planning horizon is estimated as $h^* = \arg \min_{h \in \{1, \dots, 20\}} \sum_{t=1}^{20} (c_t^M(w_t, h) - c_t)^2$, where $c_t^M(w_t, h)$ is conditional optimal (myopic) consumption given the planning horizon h .¹¹ The adopted approach of estimating the planning horizon is identical to that of Ballinger et al. (2003) and has since been used repeatedly in the literature on life-cycle consumption experiments.¹² Figure 3 illustrates the median estimated planning horizons. Estimated planning horizons are lower in the Borrowing treatments than in the Saving treatments. This difference is significant at the 5% level for five rounds and marginally significant for the remaining round ($p = 0.057$). This result is particularly interesting since there exists no a-priori reason for subjects to have different planning horizons in the different treatments. If Myopia was the explanation for the observed deviations from optimality, this explanation should be independent of the treatment, and therefore one should observe a planning horizon that is similar across treatments. This is clearly not the case.

Summing up, while myopia can explain certain characteristics of the experimental data, it fails to consistently explain observed behavior in this experiment. There is no a-priori reason why planning horizons should differ across treatments, since subjects are randomly assigned to treatments. This suggests that some other factor, not captured

¹⁰The appendix contains an example that illustrates how myopia affects deviations from full horizon optimal behavior.

¹¹I also estimated h^* using absolute deviations instead of squared deviations. The results are almost identical.

¹²see, for instance, Carbone and Hey (2004), Ballinger et al. (2011)




Figure 3: Estimated planning horizons. Solid and dashed lines represent median values from the BF and SF condition respectively

by the concept of myopia (e.g. debt aversion) drives behavior and causes the difference in estimated planning horizons.

3.3 Learning

In order to assess whether subjects are able to improve their consumption decisions over time, I look at within subject differences in the absolute deviations from conditional optimal consumption (Measure 2) across rounds. Table 2 summarizes the median differences in measure m_2 between consecutive rounds as well as to the first round. The difference between round r and the first round is denoted by $m_2^1 - m_2^r$ and the difference between two consecutive rounds is denoted by $m_2^{r-1} - m_2^r$. There appears to be a constant improvement of consumption decisions in the BF condition. All median differences between two consecutive rounds are positive and most of these differences are significant. Moreover, all rounds show significant improvement in comparison to the first round. One has to be careful, however, when comparing data from two different

measure	condition	Round					
		1	2	3	4	5	6
median $m_2^{r-1} - m_2^r$	BF	NA	58.214 <i><0.001</i>	18.725 <i>0.057</i>	69.421 <i>0.020</i>	98.312 <i>0.003</i>	19.978 <i>0.229</i>
median $m_2^1 - m_2^r$		NA	58.214 <i><0.001</i>	137.063 <i>0.001</i>	370.405 <i><0.001</i>	439.458 <i><0.001</i>	575.825 <i><0.001</i>
median $m_2^{r-1} - m_2^r$	SF	NA	66.612 <i>0.001</i>	80.909 <i>0.001</i>	-202.804 <i><0.001</i>	62.481 <i>0.143</i>	54.186 <i>0.007</i>
median $m_2^1 - m_2^r$		NA	66.612 <i>0.001</i>	155.412 <i><0.001</i>	-40.330 <i>0.408</i>	37.355 <i>0.750</i>	69.895 <i>0.376</i>

Table 2: Learning. P-values are printed in italic letters and come from a Wilcoxon signed-rank test. The null hypothesis is that the median difference between the pairs is zero.

treatments, i.e. data from the first three rounds and data from the second three rounds. Differences in behavior can be caused by both learning and the treatment effect.¹³ The significant improvement of consumption decisions between round 3 and 4 is therefore probably caused by a combination of learning and the treatment effect.

In the SF condition, consumption decisions improve from round 1 to 2 and from round 2 to 3 before they significantly deteriorate between rounds 3 and 4. After round 3 the differences to the first round are not significantly different from zero, implying no improvement in comparison to the first period. Again both learning and treatment effect play a role here. Generally one would expect the learning effect to positively influence consumption decisions. However, the overall effect between rounds 3 and 4 is significantly negative. This implies that the treatment effect is negative and that it dominates the learning effect.

Note that when comparing rounds 3 and 4 the within subject effect is the same as in the between subject analysis in Section 3.1. When changing from the Borrowing to Saving treatment there is an improvement of consumption decisions (even though here it is not possible to discriminate between the learning and the treatment effect) and when changing from the Saving to Borrowing treatment consumption decisions deteriorate.

¹³The treatment effect here is the change of the income process, which determines whether saving or borrowing is necessary for optimal consumption.

Overall these findings imply that subjects generally learn and improve their consumption decisions over time. However, experience does not seem to eliminate the treatment effect. Even after three periods of learning in the Saving First condition, consumption decisions markedly deteriorate when subjects have to borrow in order to consume optimally. In other words: the finding that subjects perform worse when they have to borrow in order to consume optimally compared to when they have to save in order to consume optimally, prevails even after subjects learn to improve their decisionmaking, i.e. learn to smooth consumption.

4 Discussion

With the experiment presented in this paper I am able to assess saving behavior as well as borrowing behavior in a dynamic optimization context. Previous experiments on life-cycle consumption problems done by others find a general tendency for subjects to save too little. This study suggests that this result does not hold when borrowing is allowed and necessary for optimal consumption. In this experiment, the direction of deviations from optimal behavior is context specific. There exists weak evidence that subjects save too little and consume too much (relative to conditionally optimal consumption) when the income stream is decreasing and saving is necessary for optimal consumption. Subjects borrow too little and thus consume too little when the income stream is increasing and borrowing is necessary for optimal consumption. Furthermore, the experimental data suggests that deviations from optimal behavior are higher when subjects have to borrow than when they have to save in order to reach optimal consumption. This implies that subjects are less willing to engage in debt than in savings, i.e. they are debt-averse. Similar to other experiments on intertemporal consumption/savings problems, learning is observed here. However, subjects seem to be unable to transfer what they have learned in the Saving treatment to the Borrowing treatment. In fact, consumption decisions significantly deteriorate after the role change from the Saving to the Borrowing treatment, suggesting that the effect that debt aversion has on consumption decisions prevails even after three rounds of experience.

Often stated reasons in the literature for sub-optimal consumption choices are bounded rationality, in particular limited planning horizons (or myopia), as well as time-inconsistent preferences. Myopia is an appealing model of behavior, because it can explain overconsumption with a decreasing income stream and underconsumption

with an increasing income stream. However, myopia cannot explain the asymmetry between deviations from optimal consumption with an increasing and with a decreasing income stream. This intuition is confirmed by the fact that significantly different apparent planning horizons are estimated in the different treatments. Since one would expect planning horizons to be equally distributed in both treatments, this implies that myopia alone cannot explain the observed differences in behavior. Other forms of computational constraints, e.g. “cost of thinking”, may explain deviations from optimal behavior but not the difference in behavior across treatments. This is because the difficulty of the optimization problem is the same in both treatments, as debt is mathematically equivalent to savings, only with a different sign. Present biased preferences cannot serve as an explanation for the deviations from optimality either. Subjects in this experiment get their (monetary) reward at only one point in time, which is after the experiment. It seems unlikely that some form of present bias affects consumption decisions during the experiment since there exist no immediate rewards. Brown et al. (2009) use sips of water as immediate rewards for thirsty subjects to emulate an intertemporal consumption/savings environment where present biased preferences can play a role. A simple $\beta - \delta$ model of discounting would predict that agents facing an increasing income stream would borrow too much, which is opposite to the findings in this paper. Analyzing which effect dominates, present bias or debt aversion, appears to be a promising future research project.

References

- Ballinger, T Parker, Michael G Palumbo, and Nathaniel T Wilcox (2003). “Precautionary saving and social learning across generations: an experiment”. In: *The Economic Journal* 113.490, pp. 920–947.
- Ballinger, T Parker, Eric Hudson, Leonie Karkoviata, and Nathaniel T Wilcox (2011). “Saving behavior and cognitive abilities”. In: *Experimental Economics* 14.3, pp. 349–374.
- Brown, Alexander L, Zhikang Eric Chua, and Colin F Camerer (2009). “Learning and visceral temptation in dynamic saving experiments”. In: *The Quarterly Journal of Economics* 124.1, pp. 197–231.

- Caetano, Gregorio, Miguel Palacios, and Harry Patrinos (2011). “Measuring Aversion to Debt: An Experiment Among Student Loan Candidates”. In: *World Bank Policy Research Working Paper Series*.
- Carbone, Enrica and John D Hey (2004). “The effect of unemployment on consumption: an experimental analysis”. In: *The Economic Journal* 114.497, pp. 660–683.
- Carroll, Christopher D (2009). “Lecture notes on solution methods for microeconomic dynamic stochastic optimization problems”. In: *manuscript, Department of Economics, Johns Hopkins University*.
- Duffy, John (2008). “Macroeconomics: a survey of laboratory research”. In: *Handbook of Experimental Economics 2*.
- Eckel, Catherine C, Cathleen Johnson, Claude Montmarquette, and Christian Rojas (2007). “Debt aversion and the demand for loans for postsecondary education”. In: *Public Finance Review* 35.2, pp. 233–262.
- Fenig, Guidon, Mariya Mileva, and Luba Petersen (2013). “Asset Trading and Monetary Policy in Production Economies”. In: *Working Paper*.
- Field, Erica (2009). “Educational Debt Burden and Career Choice: Evidence from a Financial Aid Experiment at NYU Law School”. In: *American Economic Journal: Applied Economics* 1.1, pp. 1–21.
- Fischbacher, Urs (2007). “z-Tree: Zurich toolbox for ready-made economic experiments”. In: *Experimental Economics* 10.2, pp. 171–178.
- Hey, John D and Valentino Dardanoni (1988). “Optimal consumption under uncertainty: An experimental investigation”. In: *The Economic Journal* 98.390, pp. 105–116.

Appendix

A Myopia and deviations from optimal behavior

Example 1: Assume we are in the first period of the 20-period life-cycle consumption environment described in Section 2. Further assume that the decisionmaker is myopic and has a planning horizon of 2, that is when making a decision she acts as if the the next period is the final period. Myopic optimal consumption c^M in period 1 follows from the two-period optimization problem:

$$c^M = \arg \max_{c_1} [u(c_1) + u(y_1 - c_1 + E[y_2])] \quad (4)$$

As explained in Section 3, different optimal myopic consumption levels arise in the different treatments. Since optimal consumption is a function of period wealth, optimal consumption levels also depend on the income shock. Long-term (fully rational) optimal consumption only depends on the income shock and does not differ between treatments.

1. **Borrowing treatment (increasing income stream):** Short-term (myopic) optimal consumption is:

$$c_B^M = \begin{cases} 9.503 & \text{if } y_1 = 0 \text{ (negative income shock)} \\ 19.503 & \text{if } y_1 = 20 \text{ (positive income shock)} \end{cases} \quad (5)$$

2. **Saving treatment (decreasing income stream):** Short-term (myopic) optimal consumption is:

$$c_S^M = \begin{cases} 189.503 & \text{if } y_1 = 190 \text{ (negative income shock)} \\ 199.503 & \text{if } y_1 = 210 \text{ (positive income shock)} \end{cases} \quad (6)$$

Long-term optimal consumption does not differ between treatments and is given by:

$$c^L = \begin{cases} 104.322 & \text{with negative income shock} \\ 105.322 & \text{with positive income shock} \end{cases} \quad (7)$$

It is straight-forward to see that deviations from long-term optimal behavior caused by myopia are higher in the Borrowing treatment than in the Saving treatment when

an negative income shock occurs:

$$\delta_B^{neg} = 104.322 - 9.503 = 94.819 > \delta_S^{neg} = 189.503 - 104.322 = 85.181. \quad (8)$$

Analogously, with a positive income shock deviations in the Saving treatment are higher than in the Borrowing treatment:

$$\delta_B^{pos} = 105.322 - 19.503 = 85.819 < \delta_S^{pos} = 199.503 - 105.322 = 94.181. \quad (9)$$

B Instructions

This section contains the instructions of the experiment.¹⁴ The first part of the instructions were given to subjects in a printed handout and were read aloud by the experimenter. The second part was handed out to subjects when they finished the third round. Subjects in the different conditions received different instructions. The instructions printed here were used for subjects from the Borrowing First condition.

Instructions (Part 1)

The experiment you are participating today is part of a research project. It is meant to analyze economic decisionmaking. The rules and instructions are the same for every participant. Your payoff depends on your decisions during the experiment. Please read the instructions carefully. You are not allowed to talk and exchange information with other participants during the experiment. Please raise your hand if you have a question. An Experimenter will then come to your place and answer your question. Please don't ask your questions out loud. Should you break one of these rules we are obliged to exclude you from participation.

Overview. First you will have time to read the instructions. After that we will go through the instructions together and you answer a quiz in order to make sure you understood the instructions. After that you may ask questions, before the experiment starts. After the experiment you will be asked to fill out a short questionnaire.

¹⁴The instructions printed here are a translation of the original German version.

The experiment consists of 6 rounds, each of which consists of 20 periods. The duration of the experiment is around 1.5 hours. Instructions, quiz and questionnaire will take around 30 minutes. The remaining hour is dedicated to the actual experiment. In every period a countdown of 30 seconds will be displayed. You may take more or less time to reach your decision. The countdown is meant to provide some indication how much time you can take in every period to finish the experiment in one hour. You may finish the experiment even if you play for more than one hour.

The following instructions are for the first three rounds of the experiment. After three rounds the experiment stops and you will be asked to type in a password. You will be handed new instructions for the following 3 rounds, containing the password needed to continue with the experiment. After the last round your experiment payoff will be displayed. Please raise your hand when you have finished the last period. You will then be handed a short survey. After filling out the survey, please raise your hand again. You will then be given your experiment payoff.

You are playing an “investment game” and decide in every period how many points you want to purchase. The sum of all points purchased in one round is that period’s result. Your payoff depends on the results from two randomly drawn rounds.

Income, Savings and Wealth. In every period you obtain a certain income, denoted in the experimental currency “Taler”. Your task is to choose how many Taler to spend in order to purchase points. Thereby you (implicitly) also choose how many Taler you want to save or borrow. The difference between income and spending in one period are called savings. At any period in the experiment, your wealth is defined as the sum of savings from all previous periods. This implies that savings from one period added to the wealth in this period yields the wealth in next period. Please note that the sign of the savings can be both positive and negative. In case you decide to spend less Taler than you have as income, your savings have a positive sign. In this case your wealth in the next period is your wealth in this period plus the absolute amount of savings in this period. Should you decide to spend more Taler than you have as income, your savings have a negative sign. In this case your wealth in the next period is your wealth in this period minus the absolute amount of savings.

Example: assume your income in one period is 50 Taler and you spend 30 Taler to purchase points. Your savings are then 20 Taler. In case you instead spend 70 Taler

with the same income, your savings are -20 Taler. In the first case your wealth in the next period is the wealth in this period plus 20 Taler. In the latter case your wealth in the next period is this period's wealth minus 20 Taler.

Your wealth may take positive or negative values as well, depending on whether your savings from previous periods were positive or negative. Your wealth in the first period is 0 Taler. In the last period, your wealth plus income will be spent automatically in order to purchase points. This implies that the sum of Taler spent in all periods of one round equals the sum of income obtained in all periods of this round. In other words: you may spend more or less than your income in one round. However, over one round, the sum of income always equals the sum of Taler spent.

Determination of Income. Income is randomly determined and follows the random process:

$$y_t = 10t + \epsilon_t$$

The index "t" denotes the periods for which income is determined. Since the slope of the process is 10, it has a positive trend. The expected income is therefore increasing over time. ϵ_t is the random part of the process and can be either +10 or -10, both occurring with equal probability of 50%. Since ϵ_6 is either +10 or -10, your income in period 6 is either 70 or 50. Since one round consists of 20 periods, income in the last period will either be 210 or 190. It is very important to understand that ϵ_t is truly randomly determined. Which value ϵ_t takes in one period does not depend on which values it had in previous periods or how you behaved in previous periods.

Taler and Points. Your task is to decide in every period how many Taler you want to spend in order to purchase points. Taler are transformed to points as follows:

$$\text{Points} = 250(1 - e^{-0.02 * (\text{chosen amount of Taler})})$$

A graph of this function, as well as a table with relevant function values is attached to the instructions. Please note that the above function is defined on the positive as well as the negative domain. If you choose to spend a negative amount of Taler, you receive a negative amount of points. In this case you "sell" points and gain Taler. Should your wealth plus income in the last period be negative, you will have to automatically sell

points in order to make sure that your Taler account is balanced.

Payoff. Your payoff depends on the results from two randomly drawn rounds. One round is randomly drawn from the first three rounds and the other is randomly drawn from the second three rounds. Your payoff is calculated as follows:

$$\text{Payoff in Euro} = \frac{(\text{Result1} - 3000) + (\text{Result2} - 3000)}{100}$$

where Result1 is the first randomly drawn result and Result2 is the second randomly drawn result.

Example: suppose the first randomly drawn result is 4300 points and the second randomly drawn result is 3800 points. Your payoff is then:

$$\frac{(4300 - 3000) + (3800 - 3000)}{100} = \frac{1300 + 800}{100} = 21 \text{ €}.$$

Quiz and Questions. You will now be asked to answer a short Quiz regarding the contents of these instructions. In case you have questions after that, please raise your hand. An experimenter will then come to you and answer your question.

Instructions (Part 2)

In the following three rounds the random process that determines your income will change. Therefore, compared to the first part of the instructions, only the paragraph “Determination of Income” changes. The rest of the instructions is still valid.

Determination of Income. Income is randomly determined and follows the random process:

$$y_t = 210t - \epsilon_t$$

The index “t” denotes the periods for which income is determined. Since the slope of the process is -10, it has a negative trend. The expected income is therefore decreasing over time. ϵ_t is the random part of the process and can be either +10 or -10, both occurring with equal probability of 50%. Since ϵ_6 is either +10 or -10, your income in period 6 is either 140 or 160. Since one round consists of 20 periods, income in the last period will either be 20 or 0. It is very important to understand that ϵ_t is truly

randomly determined. Which value ϵ_t takes in one period does not depend on which values it had in previous periods or how you behaved in previous periods.