The Burden of Household Debt

By ALEJANDRO MARTÍNEZ-MARQUINA and MIKE SHI *

January 14, 2021
[Link to the latest version]

Abstract

We propose that holding debt causes worse financial decisions using a novel experimental design where we randomly assign debt. Our design isolates the consequences of holding debt while controlling for potential confounding factors such as initial wealth levels, selection, risk, and time preferences. Our findings show that debt causes behavioral biases detrimental to subjects' financial payoffs. However, subjects' strategies are not random but instead debt-biased, consistent with a model of debt aversion. We refer to the financial losses caused by debt as the Burden of Debt *and provide evidence that, under certain circumstances, these behavioral biases can compound and lead to substantial losses. Furthermore, we show in additional treatments how these debt-biased behaviors can also deter subjects from borrowing and miss profitable opportunities.*

Debt is pervasive in the United States. Whether you buy a car or a house, use a credit card, face a surprise medical bill, or attend college, debt is taken as a given for Americans. According to a recent Pew Report, 8 in 10 Americans hold debt, and nearly 70% view debt as necessary even if they prefer not to have it.[1] Recent empirical evidence shows that holding debt is correlated with suboptimal decision-making and worse financial outcomes.[2] Despite this evidence, the mechanisms at play are not clear and under-studied relative to the ubiquity of debt in American life. Clearly, debt should be factored into decision-making as it speaks to commonly understood aspects

*Martínez-Marquina: Department of Economics, Stanford University, Stanford, CA 94305, (amm15@stanford.edu). Shi: Department of Economics, Stanford University, Stanford, CA 94305, (mikeshi@stanford.edu). We are thankful to Ran Abramitzky, Doug Bernheim, Muriel Niederle, Kirby Nielsen, Luigi Pistaferri, Collin Raymond, Al Roth, Frank Schilbach, Colin Sulivan, Dmitry Taubinsky, and numerous seminar participants for their invaluable comments and feedback. This research is supported by the Stanford Institute of Research in Social Sciences, the Center for Effective Global Action (CEGA), and the Bradley Graduate Fellowship through a grant to the Stanford Institute for Economic Policy Research.

[1]The Complex Story of American Debt, Pew (2015).

[2]For example, Gathergood et al. (2019) shows how people with several credit cards do not always minimize interests payments. Azmat and Macdonald (2020) finds evidence of sub-optimal house mortgages repayments in Pakistan.

like liquidity constraints, wealth effects, or default risk. Yet, could it be that the impact of holding debt goes above and beyond these direct effects? Given the totality of debt among American households, it is essential to understand all of its consequences.

In this paper, we show that debt affects decision making through behavioral biases consistent with subjects deriving disutility from holding negative balances. Using a novel experimental design, we are able to show how debt-biased behaviors translate into worse financial outcomes, which we call the *Burden of Debt*. Even stripped of its direct consequences on credit scores or bankruptcy, debt is more than just negative savings. In a setting with this simplified and stylized version of debt, we find that it still causes significant deviations from typical maximizing behavior.

To illustrate how debt can cause biased behavior that negatively impacts financial outcomes consider the following simplified decision where an agent has to allocate some funds. In our experiment a participant has several investment opportunities available, some with substantial upside potential, that she can devote her funds to. Alternatively, she can also use this money to repay her outstanding debt, even if it has lower interest rates. From the perspective of a non-behavioral agent who only focuses on maximizing payoffs, there is no reason to prioritize debt repayments. On the other hand, if the agent is averse to holding debt balances, she might still decide to devote her funds to repaying debt. Our experiments finds causal evidence of these debt-biased behaviors and the financial losses they can cause.

In our main treatments, subjects own virtual accounts with different interest rates and balances that generate returns over time. During a week-long experiment, they allocate the returns these accounts generate and make a total of four allocation decisions. To maximize returns, and hence the final payoff from the experiment, subjects should always allocate everything to the highest interest rate account. While accounts differ in their interest rate and initial balances, only interest rates matter when maximizing returns. In our first treatment, subjects can only allocate points to accounts that start with positive balances, providing a baseline for what fraction of participants follows the return-maximizing strategy in the absence of debt. Our main treatment variation changes the accounts' initial balances, such that two of them now have a negative starting balance. These negative balances are such that subjects can fully repay them during the experiment. We keep the available interest rates constant across these two treatments, and thus the return-maximizing action remains the same regardless of the presence of debt.

We quantify deviations from the return-maximizing strategy between treatments to assess the *Burden of Debt*. We find that subjects with debt are three times less likely to maximize returns across all decisions compared to subjects in our baseline group. This difference is mainly driven by 38 percent of subjects who fully repay at least one outstanding negative balance and miss the opportunity of higher returns from other

accounts. Thus subjects with negative balances focus on repaying their outstanding debt at the expense of lower monetary payoffs. Next, we test if the prevalence of this strategy depends on the size, and not just the sign, of the debt balances by running an additional treatment where we increase the initial negative balances. We find a heterogeneous response: more subjects maximize total returns, but those who focus on repaying debt end up with larger financial losses. Altogether, our results indicate that many subjects are attempting to repay their debt as quickly as possible, which in our experimental context is financially sub-optimal.

In additional treatments, we increase subjects' agency by allowing them to redistribute balances across accounts. If redistribution is also debt-biased, we can see if the *Burden of Debt* compounds and causes larger financial losses. This additional option could benefit subjects as it allows them to further increase their returns by consolidating towards the highest interest account. When subjects can redistribute towards debt accounts, we find that 34 percent of participants partially exhaust their high interest account towards a low interest one, dampening their returns. In contrast, when redistribution towards debt accounts is not possible, only 14 percent of the subjects follow such strategy, and twice as many subjects now increase their payoffs by redistributing towards the highest interest account. By running these additional treatments, we show that redistribution decisions, and not just allocation decisions, are biased towards early debt repayments, which we take as further evidence of the *Burden of Debt* and its negative financial consequences.

After finding that debt distorts allocation and redistribution decisions, we then explore if the *Burden of Debt* also manifests in borrowing decisions. In two additional treatments, we introduce borrowing opportunities and vary whether they incur debt or not. These borrowing opportunities allow subjects to further benefit from the highest interest rate account. Nevertheless, only 34 percent of subjects borrow the maximum amount when borrowing involves debt compared to 63 percent when borrowing does not incur debt. This stark difference persists even among subjects who otherwise maximize returns. Therefore, we show that debt also hinders borrowing behavior and prevents subjects from undertaking profitable investment opportunities.

While debt is pervasive, our understanding of how holding debt affects decisions is still limited. Previous models in the behavioral literature focus on understanding borrowing decisions as opposed to the behaviors of those already in debt. For example present bias (Laibson, 1997; O'Donoghue and Rabin, 1999; Laibson et al., 2007; Meier and Sprenger, 2010; Ikeda and Kang, 2011) or self-control motives (Cadena and Keys, 2013; Allcott et al., 2020) could explain excessive borrowing but they do not imply that debt will cause financial mistakes after it is acquired.[3] Using a model of prospective

[3]Other potential drivers of borrowing decisions present in the literature are lack of information (Bertrand and Morse, 2011; Burke et al., 2016) and debt aversion (Callender and Jackson, 2005; Caetano et al., 2011; Meissner, 2016).

accounting, Prelec and Loewenstein (1998) rationalizes aversion to financing purchases through credit–i.e., debt aversion–because of the "pain of payment". Our findings, in line with Prelec and Loewenstein (1998), suggest that subjects perceive negative balances differently and act upon these perceptions, consistent with debt aversion.

Previous work in the experimental and psychology literature provides evidence of subjects struggling to make debt repayment decisions that minimize accrued interest. Experimental subjects exhibit a preference for closing out small debt accounts (Amar et al., 2011; Besharat et al., 2014) and concentrating repayments as opposed to spreading them out (Kettle et al., 2016). Furthermore, Besharat et al. (2015) shows how the timing and the type of debt, hedonic vs. utilitarian, can amplify these effects. In recent work, Ozyilmaz and Zhang (2019) documents how subjects struggle to minimize interest payments in an all-debt environment compared to a non-debt setting. In our design, as opposed to previous studies, subjects also have high return investment opportunities in addition to the option to repay debt. This allows us to show the severe opportunity costs of debt repayment, even if it is done optimally. Furthermore, the experimental literature has also found evidence of sub-optimal behavior in borrowing decisions. Caetano et al. (2011) finds a reluctance to accept contracts presented under a debt frame in a field experiment. Along these lines, Meissner (2016) shows that subjects are reluctant to borrow to increase present consumption in an inter-temporal consumption experiment. Our work combines both findings present in the literature and shows how debt aversion can lead to large opportunity costs through both the borrowing and debt repayment decisions.

One key feature of our design is the introduction of debt independent of financial hardship. While people in precarious financial situations can potentially be more exposed to indebtedness, it is important to separate debt from poverty. Prior work on scarcity has found lacking money or time leads to poorer decisions (Shah et al., 2012; Mani et al., 2013), that people under financial strain are less productive (Kaur et al., 2019) and have higher cognitive load (Haushofer and Fehr, 2014; Schilbach et al., 2016; Ridley et al., 2019), and that wealth impacts the relative perception of assets and debt balances (Sussman and Shafir, 2012). Debt is independent of wealth in our setting, yet it still causes financial mistakes. In contrast to scarcity, we find evidence of debt-biased financial mistakes which are not necessarily random or erratic as limited cognitive capacity would imply.

This paper relates to an empirical strand of the literature on understanding the consequences of indebtedness. Azmat and Macdonald (2020) provides evidence of sub-optimal home mortgage repayments with borrowers in Pakistan, where many choose to pay additional fees to repay their loans faster, even when this does not reduce interest payments. Further evidence on credit card repayments shows that conditional on repaying debt, many borrowers do it sub-optimally (Stango and Zinman, 2009; Keys et al., 2016; Ponce et al., 2017; Gathergood et al., 2019). Research on students loans

finds that debt can lead to higher paid jobs (Field, 2009; Rothstein and Rouse, 2011; Luo and Mongey, 2016), deter graduate school enrollment (Fos et al., 2017), or reduce labor search effort (Ji et al., 2016). At the same time, there has also been documented evidence of a reluctance of borrowing for college, even among qualified students (Callender and Jackson, 2005).[4] Our controlled experimental setting allows us to provide causal evidence on the impacts of debt and to isolate the mechanism of debt aversion. This mechanism can provide a potential behavioral explanation as to why borrowers would want to repay their loans faster and why student loans could have such a significant impact on students' careers.

Future research will assess the economic relevance of our findings on debt aversion, particularly in sub-optimal financial decisions outside the lab. As with any experimental evidence, further work should also assess the external validity and replicability of our results. Behavioral measures have been shown to correlate with choices in non-experimental settings. For example, (Meier and Sprenger, 2010) show that present-biased individuals are more likely to have credit card debt. We believe debt aversion could also become a behavioral trait. In that case, experimental measures of debt aversion have the potential to shed light on financial decisions.

In the next section we describe the experimental design as well as our hypotheses and predictions. Section II contains our main results on the existence of the *Burden of Debt*. Section III presents evidence on how these financial mistakes can compound and lead to substantial losses. In section IV we show how the consequences of the *Burden of Debt* can also affect borrowing decisions. We then outline a theoretical model of debt aversion and how it relates to previous literature. Finally, we summarize and discuss our findings.

I. Experimental Design

A. Isolating the Burden of Debt: Benefits and Limitations of the Lab

In our experimental design, we aim to show that debt causes financial mistakes by randomly assigning debt. Outside the lab, there are many different types of debt with varying financial consequences, but in our context, debt denotes a negative balance. We show that even this simplified and stylized version of debt, with no consequences in terms of risk or wealth, can still cause behavioral biases. As with any lab experiment, control and simplification come at the expense of generalization. In contrast to previous empirical work, our environment allows us to find causal effects and rule out confounding factors like income uncertainty or the selection of borrowers. While we are aware of the experimental tradeoff, showing the existence of the *Burden of Debt* in the lab is a necessary step towards a better understanding of how holding debt impacts

[4] Avery and Turner (2012) argues that given the high returns to education, the claim that student borrowing is "too high" can clearly be rejected in most cases.

decisions.

We focus on understanding financial mistakes caused by debt in two scenarios: first, how subjects choose to repay their outstanding debt, and second, the decision to borrow in the first place. In the former, we study whether subjects focus on debt repayments at the expense of higher returns from other accounts. This debt-biased behavior will show how holding debt can be detrimental to financial wealth. In the latter scenario, we aim to show that the *Burden of Debt* can hinder financial decisions, even for subjects without any outstanding debt. Subjects can obtain more substantial returns by borrowing from other accounts, but they might be reluctant to do so when it involves incurring debt.

B. Basic Setting

Subjects own virtual accounts with different interest rates and balances that generate returns over time. Balances change based on subjects' decisions, while interest rates are constant throughout the experiment. Accounts with positive balances generate positive returns and accounts with negative balances generate negative returns. For example, an account with a starting balance of 1100 points and an interest rate of 20% generates a positive return of 220 points. Similarly, an account with an initial balance of -900 points and an interest rate of 10% generates a negative return of -90 points. We label accounts with positive starting balances as Savings accounts and those with negative starting balances as Debt accounts. In the main setting, subjects can only impact balances by allocating points and not by redistributing balances, but we relax this in additional treatments. Points allocated to an account increase its balance for all subsequent decisions and affect the returns it generates.

Subjects make four decisions over a week-long period. The timeline works as follows. For their first decision, participants must allocate an initial endowment of points. Two days after the first allocation decision, returns materialize, and these returns constitute the new endowment that subjects must allocate across their accounts. These subsequent allocation decisions likewise affect the balances and hence the returns that accounts generate. This process repeats for a week until subjects have made four allocation decisions in total.

When deciding how to allocate points, the only relevant factor is the interest rate each account provides. In all treatments of our experiment, interest rates are kept constant, and hence, all treatments have the same set of opportunities for allocating points. For example, all treatments have accounts with 20% and 5% interest rates. While allocating 500 points to the 20% interest rate accounts generates a return of 100 points, allocating those points to the 5% account would only generate 25 points. It does not matter what the starting or current balances are or whether these accounts are labeled as Savings or Debt accounts; only interest rates matter for maximizing returns.

In this setting, since returns are known, and certain and final payments are made only at the conclusion of the experiment,[5] the action that maximizes returns is unique regardless of risk or time preferences. Unlike other empirical settings where risk or return uncertainty can impact debt decisions, our experiment rules out these confounding factors; subjects cannot default on their debt, and investment returns do not fluctuate. Accounts accrue interest over time, and final payments depend on the outstanding account balances. We incentivize the final decision by also paying subjects the returns that their final balances generate. Paying subjects at the end of the experiment allows us to rule out other potential confounding effects like self-control (Cadena and Keys, 2013; Allcott et al., 2020) or present-bias (Laibson, 1997; O'Donoghue and Rabin, 1999; Laibson et al., 2007; Meier and Sprenger, 2010). In addition to the returns generated by the accounts, subjects have several opportunities to make additional gains throughout the experiment, discussed in detail in subsection G.

C. Decisions with Debt: Main Treatments

We introduce debt to subjects by giving them accounts with initial negative balances while keeping the highest interest rate still associated with a savings account. Using additional accounts, we keep initial wealth and returns equal across treatments. Subjects own six accounts in total, four in which they can allocate points to and two that are locked throughout the experiment. Locked accounts still generate returns, but participants cannot allocate any points to them. Varying the starting balances of these locked accounts across treatments allows us to introduce debt while still keeping total wealth and returns equal.

No Debt Treatment

In this treatment, subjects can allocate points only to savings accounts, providing a baseline for return-maximizing behavior. Since the action that maximizes returns is on the boundary of the action set–i.e., allocating all points to the highest interest rate account–we do not expect subjects to all choose that action.[6] The locked accounts correspond to Debt 1 and Debt 2, both with an initial balance of zero points that generate no returns. This feature allows us to control for the possibility that mentioning the word "debt" could affect choices. Hence, accounts labeled as Debt are present in all the main treatments. Subjects accounts are as follows:

[5]This excludes additional payments based on answers for the elicited risk and time preferences which we discuss further in subsection G.

[6]Previous research has shown that when the payoff maximizing action is on the boundary, it is selected less often, e.g., in charity donation games. For a detailed discussion, see Vesterlund (2016).

Table 1: Accounts in Main Treatments

No Debt:

	Savings 1	**Savings 2**	**Savings 3**	**Savings 4**	**Debt 1**	**Debt 2**
Interest Rate	**20%**	**10%**	**15%**	**5%**	15%	5%
Balance	**1100**	**700**	**900**	**1500**	0	0

Low Debt:

	Savings 1	**Savings 2**	**Debt 1**	**Debt 2**	**Savings 3**	**Savings 4**
Interest Rate	**20%**	**10%**	**15%**	**5%**	15%	5%
Balance	**1100**	**700**	**-900**	**-1500**	1800	3000

Notes: In both cases, the sum of the balances is 4200 points and the returns of these six accounts sum up to 500 points, which is the initial endowment that subjects must allocate.

As discussed earlier, allocating all points to Savings 1 maximizes returns and hence, final payoffs.[7] Savings 1 has the highest interest rate but not the highest initial balance, but again the initial balances should not be considered when maximizing returns.

Low Debt Treatment

In this treatment, subjects now have debt accounts with initial negative balances that generate negative returns in contrast to *No Debt*. These two accounts, Debt 1 and Debt 2, are not locked and subjects can allocate points to them; to keep the same number of opportunities, we now lock two savings accounts, Savings 3 and Savings 4, with parallel interest rates. We redistribute initial balances such that the sum of the balances of Savings 3 and Debt 1, both with the same interest rate, is equal across treatments, and we do the same for Savings 4 and Debt 2. This redistribution allows us to introduce negative balances while keeping the net sum of balances, the number of actions and the available interest rates equal across our main treatments. Despite having negative balances, the return-maximizing strategy does not change. The accounts in this treatment are as follows:

Parameters are such that subjects can obtain enough points during the experiment to repay both outstanding debt balances entirely. All subjects can zero-out (fully repay) Debt 1 by the end of the 2nd day, while Debt 2 requires three days. Due to additional opportunities to obtain more points, subjects may be able to repay them earlier. Al-

[7]Since subjects make four allocation decisions in total, the interest difference can compound up to four times. The return of the initial 500 points when entirely allocated to a 20% interest account, including the returns it generates, will provide a $1.20^4 \approx 2.07$ return compared to the $1.05^4 \approx 1.22$ return from a 5% interest rate account.

though subjects have negative starting balances, their initial wealth is positive, and their earnings can only grow.

We interpret deviations from the payoff-maximizing allocation decision as evidence of the *Burden of Debt*. While some deviations could be due to the inherent complexity of the setting, *No Debt* and *Low Debt* together allow us to isolate those caused by holding debt. Because treatments have the same number and types of accounts and the same available interest rates, subjects go through the same set of instructions and understanding checks. Therefore we rule out that treatment differences are due to differential understanding or treatment complexity.

D. Increasing the Negative Balances: High Debt Treatment

In this treatment, we increase the initial negative balances so subjects cannot fully repay them. We hypothesize that debt repayment strategies are contingent on the feasibility of repaying debt entirely. On the one hand, more debt could exacerbate deviations from payoff maximizing behavior by devoting even more points to debt repayments. However, on the other hand, it could lead to better financial decisions in our experiment. If subjects only repay their debt if they can do it fully, we might observe more subjects focusing on maximizing their returns instead–which again means allocating all points to the highest interest rate account, Savings 1. Table 2 shows the available accounts and balances in the *High Debt* treatment.

Table 2: Accounts in *High Debt* Treatment

	Savings 1	**Savings 2**	**Debt 1**	**Debt 2**	**Savings 3**	**Savings 4**
Interest Rate	**20%**	**10%**	**15%**	**5%**	15%	5%
Balance	**1100**	**700**	**-2900**	**-3500**	3800	5000

In order to increase debt balances while maintaining a constant initial wealth and returns, balances of the debt accounts (and both locked savings accounts) are increased by 2000 points each. These higher balances imply that most subjects will not be able to fully repay one debt account, even if they allocate all their earnings throughout the experiment.[8]

E. Redistribution Treatments

In addition to our main treatments, we run two redistribution treatments where subjects can reallocate the balances from Savings 1 and Savings 2. This additional option gives

[8]Only subjects that manage to obtain substantial extra points from the additional questions can manage to repay one debt account entirely.

subjects more control over their accounts and can benefit subjects who consolidate points into the highest interest account. On the other hand, subjects could reallocate points to accounts with lower interest rates. Our redistribution treatments vary the accounts associated with these lower interest rates. Thus, we can test if redistribution decisions are also debt-biased. In that case, the *Burden of Debt* could compound and cause even larger financial losses than in our main treatments.

Redistribution Debt

We now introduce the option to redistribute Savings 1 (20%) and Savings 2 (10%) balances into other accounts. The available accounts are the same as in *Low Debt* but with different starting balances. While potentially beneficial, redistributing points could be detrimental to the subjects' returns. By depleting their savings balances, subjects increase the amount of points they allocate in a given decision. If they consolidate points in Savings 1, their returns will be higher as they will shift balances from a 10% account into a 20% account. However, they could instead use those balances to repay their outstanding debt. Moving points from Savings 1 to Debt 2 lowers the debt balance at the expense of reducing the return on those points by a factor of four since they generate 5% interest rather than the original 20%. Parameters are such that if subjects reallocate all points towards Debt 2, subsequent returns become so low that debt balances cannot be fully repaid during the experiment. If instead subjects maximize returns, by day 4 they will accumulate enough points to fully repay all debt if they want to. Hence, subjects have an additional option that is potentially beneficial but will severely decrease their final payoffs if misused.

Table 3: Accounts in Redistribution Treatments

Redistribution Debt:

	Savings 1	**Savings 2**	**Debt 1**	**Debt 2**	Savings 3	Savings 4
Interest Rate	**20%**	**10%**	**15%**	**5%**	15%	5%
Balance	**2000 (0)**	**400(0)**	**-600**	**-4300**	300	6400

Redistribution No Debt:

	Savings 1	**Savings 2**	**Savings 3**	**Savings 4**	Debt 1	Debt 2
Interest Rate	**20%**	**10%**	**15%**	**5%**	15%	5%
Balance	**2000(0)**	**400(0)**	**600**	**4300**	-900	-2200

Notes: Minimum balance requirement in parenthesis.

Redistribution No Debt

We control for the additional option to redistribute balances by running a *Redistribution* treatment without debt, analogous to our main treatments. Like in *Redistribution Debt*, subjects in *Redistribution No Debt* have the option to redistribute points from the 20% and 10% accounts, potentially increasing or reducing their returns. In contrast to the prior treatment, balances cannot be redistributed to offset an outstanding negative balance: all accounts start with positive balances. As table 3 shows, available interest rates and initial net balances are comparable across treatments. To compensate for the positive balances, locked accounts are now debt accounts with initial negative balances. This also ensures that, of the six accounts, four are Savings and two are Debt in both treatments. Just as in *Redistribution Debt*, subjects have an additional option that is potentially beneficial, but if it is misused it will severly decrease their final payoffs.

F. Borrowing Treatments

In our next two treatments, we modify our basic setting to give subjects the opportunity to increase their allocation endowments by borrowing from their locked accounts. This is in contrast to the redistribution treatments where subjects moved points from non-locked accounts. While our previous treatments are concerned with how the presence of debt affects behavior after exogenously assigning debt, these borrowing treatments allow us to see if the *Burden of Debt* also manifests in the decision of going into debt.

In both of our borrowing treatments, the amount of points subjects may borrow is equal, and the only difference across the two is the initial balance of the accounts and thus their labelling, Savings or Debt. Subjects make their borrowing decision each day, before the allocation decision. They can continue to borrow as long as the cumulative amount borrowed does not exceed the account caps, which are constant and equal across treatments. While subjects can borrow from these accounts freely up to the cap, these accounts are otherwise locked and no points can be allocated to them.

Because subjects have the option to move points from accounts with lower interest rates to accounts with higher interest rates, they can profit by borrowing and properly investing these points, regardless of which borrowing treatment they belong to. However, this profitable opportunity is not mechanical, it is possible for borrowing to reduce payoffs if points are invested in lower-interest accounts. As in our main treatments, we maintain constant initial wealth, returns, and available interest rates. While one concern may be that we are adding an additional action that subjects may or may not take for other, non-debt reasons, we only compare subjects across the two borrowing treatments and not against the main treatments discussed earlier.

Borrowing from Debt (Borrow Debt)

This treatment allows us to study if subjects are reluctant to borrow to invest when doing so requires going into debt. As discussed earlier, the only difference between our *Borrow Debt* and *Borrow Savings* treatments are the initial balances and labelling of the two borrowing accounts. In this case, the two accounts are labeled as Debt accounts, Debt 1 and Debt 2, and start with initial balances of zero. Thus if a subject borrows any amount of points, their balances become negative. Table 4 shows the initial balances for all accounts subjects face as well as the borrowing caps for the two relevant accounts, which are 900 and 1500 respectively. While there may be concern about one treatment starting at zero and borrowing involving going into negative values, that is precisely the effect of debt that we are after.

Table 4: Accounts in Borrowing Treatments

Borrow Debt:

	Savings 1	**Savings 2**	**Savings 3**	**Savings 4**	Debt 1	Debt 2
Interest Rate	**20%**	**10%**	**15%**	**5%**	15%	5%
Balance	**1100**	**700**	**900**	**1500**	0 (-900)	0 (-1500)

Borrow Savings:

	Savings 1	**Savings 2**	**Savings 3**	**Savings 4**	Savings 5	Savings 6
Interest Rate	**20%**	**10%**	**15%**	**5%**	15%	5%
Balance	**1100**	**700**	**900**	**1500**	900 (0)	1500 (0)

Notes: Minimum balance requirement in parenthesis.

Borrowing from Savings (Borrow Savings)

To provide a baseline for how many subjects are willing to take this additional borrowing action, subjects now can borrow from accounts with positive starting balances. In parallel with *Borrow Debt*, subjects have two locked accounts, now labeled as Savings accounts, with positive starting balances that they can borrow from. These starting balances are the same as the borrowing caps in *Borrow Debt*, and subjects can only borrow until these accounts are zeroed out, thus resulting in the same borrowing cap as the prior treatment. Again, this additional action may or may not be a profitable opportunity depending on the flow of points from different accounts with different interest rates. The exact starting parameters are shown in Table 4. As we can see, the only differences across treatments is the labeling and starting balances of some of the accounts, which do not affect the payoff-maximizing strategy. Our design thus isolates

the effect of debt on borrowing decisions and subsequent point-allocating behavior.

Altogether, these two treatments allow us to see if subjects are willing to hold debt for a few additional periods to make additional gains. Our lab setting provides a situation where there is no downside to borrowing; there is no risk or uncertainty involved, and borrowed points can only be invested. It is clear then that reluctancy to borrow only leads to missing a profitable opportunity. Thus we can study if the *Burden of Debt* causes subjects to fail to take a strictly payoff-maximizing action.

G. Additional Questions

During the experiment, subjects have the opportunity to make additional earnings by answering additional questions. While these questions may affect the earnings subjects make, they do not alter the main implications of our experiment as they do not change the payoff maximizing action.

Initial Survey

To prevent attrition and control for baseline risk and time preferences, subjects must complete an initial survey before their first allocation decision. For our week-long experiment, and as with all longitudinal studies, attrition is an issue. The initial survey ensures that subjects who fail to check their emails, necessary to get the links to the subsequent decisions, are dropped before we randomize them into treatment groups. Furthermore, this reduces concerns about selective attrition by treatment.

Beyond our concerns with attrition, our initial survey also allows us to elicit starting risk and time preferences through a BDM mechanism following the guidelines in Healy (2016). For each question, subjects are shown a price list where they choose between two options. In the risk preferences case, subjects are asked whether they prefer dollars for sure versus a 50 percent chance at earning \$1. In the time preferences case, subjects face a tradeoff between dollars today versus \$1 next week.[9] In addition, subjects are asked a series of understanding questions beforehand that they have to get correct before they can respond to the lists to ensure understanding of the mechanism. This price list BDM mechanism is used again in our main decisions, so this is also a way to introduce subjects to these questions beforehand.

[9]Both price lists have 100 versions of this question, with the dollars for sure ranging from one cent to \$1. Rather than have subjects answer all 100 questions, we ask them for the spot on the list where they would switch from preferring one option to the other. We then fill in assumed answers for all other questions based on their switching point. One question from one list is randomly selected and implemented.

Eliciting Time and Risk Preferences

After the initial allocation decisions each day, subjects have the opportunity to make additional gains by answering a series of risk and time preferences questions. Using the same BDM mechanism as in our initial survey, we ask subjects four risk and time tradeoff questions, two for risk preferences and two for time preferences (details in Table 5). These questions are shown in random order within the risk or time block. Of these four lists, one question from one list will be randomly selected and implemented. To ensure that all subjects have additional points to allocate, we give everyone 100 additional points regardless of the implemented question. Thus after completing the additional questions subjects will again be able to allocate their earned points to their four available accounts.

Table 5: Additional Questions Time and Risk Preferences

	Option A		**Option B**
Risk Question #1:	50% chance of 500 points	vs.	X points for sure
Risk Question #2:	50% chance of 500 points	vs.	X dollars paid today
Time Question #1:	500 points for the next allocation decision	vs.	X points for the current one
Time Question #2:	500 points for the next allocation decision	vs.	X dollars paid today

One-Shot Allocations

In the last allocation decision, right before finishing the experiment, subjects are presented with three simplified one-shot scenarios mimicking the main treatments to assess the robustness of our week-long findings. Each one-shot scenario corresponds to one of the three main treatments, *No Debt*, *Low Debt* and *High Debt*, but with only one allocation decision instead of four. Subjects face these three decisions in random order, knowing that only one will count for payment. In all of these scenarios, subjects must allocate 1000 points among the four available accounts.[10] Like in the week-long experiment, the payoff-maximizing action is to allocate all points to the account with the highest interest rate (20%).

H. Procedures

Subjects are recruited on Amazon Mechanical Turk (MTurk) and asked to complete an online survey for a week-long study. All participants are required to reside in the US, to have completed at least 50 HITs with a 90% approval rate, and to not have

[10] See Appendix Table A.10 for details on each one-shot scenario.

Table 6: Summary of Experimental Design

	Day 1	**Day 2**	**Day 3**	**Day 4**
Part 0	Initial Survey	–	–	–
Part 1	Allocation Decision	Allocation Decision	Allocation Decision	Allocation Decision
Part 2	Risk and Time Elicitation	Risk and Time Elicitation	Risk and Time Elicitation	Risk and Time Elicitation*
Part 3	–	–	–	One-shot
Part 4	–	–	–	End Survey

* Only Risk Question #1

taken a past or similar version of this experiment. To prevent attrition, reminders were sent every 6 hours on the days of the allocation decisions. We recruited a total of 578 participants with unique IP addresses, and a completion rate of 85%, with nonsignificant difference in attrition by treatment.[11] Treatment assignment only happens after responding to the follow-up email from the initial survey, ensuring that participants receive the notifications and minimizing differential attrition by treatment.

Since all treatments involve allocating points, instructions are identical in all of them. During the instruction period, participants see several examples and understanding questions that they must correctly answer before proceeding to the main decisions.[12] We track the number of errors subjects make and use this as a control in our analysis. Each decision day, subjects must go through the instructions again, including the understanding questions. So by the last day, subjects would have gone through the instructions four times altogether. We provide the instructions for all of our treatments in Appendix A.

Subjects who finish the entire experiment are paid a $10 participation fee and a bonus based on their performance. Payments are determined as follows: after the last allocation decision, all point balances are added up and converted into dollars with a 500 point to $1 conversion rate. To incentivize the last allocation, we calculate the corresponding returns for the final balances and add them to the total point count. This is also paid out for the selected one-shot treatment. The median subject made $33 in the main treatments and $35 in the borrowing treatments and took 2 hours to complete all parts of the survey.

At the conclusion of the experiment, we elicit a series of demographic and feedback questions. We obtain information on subjects' general characteristics such as age or

[11] On average, each treatment has 82 participants, ranging from 77 to 86.

[12] At the end of the first allocation decision, we ask subjects if any part of the instructions or the survey were confusing. 95% of the subjects mention no problems in understanding all parts.

gender as well as information on their finances like student loan exposure or outstanding debt. Beyond eliciting demographics, we also ask subjects to describe the reasoning behind their decisions in their own words. Subjects describe their own strategy and how they would behave in hindsight given their knowledge of how the experiment has played out. In addition subjects evaluate, in a ranking from one to five, the relevance of different aspects–balances, interest rates and debt–on their decision-making process. While these questions were not incentivized, they still provide suggestive evidence of their behavior.

I. Empirical Roadmap for the Main Results

We use our main treatments to establish the *Burden of Debt* in a controlled experimental setting. Since *No Debt* and *Low Debt* have identical initial conditions, the same number of actions, and the same return-maximizing strategy, we interpret differences in behavior as evidence of financial mistakes caused by debt biases. At a first pass, we restrict attention to quantifying the portion of participants who maximize returns in all their allocation decisions. We proceed next to analyze if, when in debt, subjects behave more randomly or erratically (i.e., choosing strategies that do not take into consideration interest rates) or if they follow a debt-specific behavior that does not maximize returns. Our one-shot games serve as robustness checks of our main results in a shorter time horizon.

We use our redistribution treatments to show that the *Burden of Debt* does also impact redistribution decisions and can exacerbate financial losses. *Redistribution No Debt* and *Redistribution Debt* both allow subjects to redistribute balances from two savings accounts. However, in one case this option can allow subjects to repay their outstanding debt balances even faster. We therefore compare how many subjects take advantage of these redistribution options and how many use them to maximize their returns or to frontload debt repayments. In the latter case, the consequences of debt-biased allocation decisions can be exacerbated by the misuse of redistribution opportunities.

After establishing the detrimental consequences of holding debt, we want to assess if the *Burden of Debt* is also present for borrowing decisions. In that case, it would imply that the *Burden of Debt* can also impact subjects without any outstanding debt. The two borrowing treatments, *Borrow Debt* and *Borrow Savings*, differ in the type of account subjects can borrow from and thus allow us to observe if participants are more likely to forgo profitable investment opportunities when they require incurring debt. At the same time, we can directly test if those who decide to borrow are positively selected or not by comparing their subsequent allocation decisions.

II. The Consequences of Holding Debt

We assess the *Burden of Debt* by comparing deviations from the return-maximizing strategy between *No Debt* and *Low Debt*. When considering participants who maximize returns in all four allocation decisions, we find that almost three times as many subjects in *No Debt* use this strategy compared to subjects in *Low Debt*. We then show that subjects who hold debt focus on repaying their outstanding balances, with a large fraction zeroing them out entirely. Using our *High Debt* treatment, we show that when this strategy is no longer feasible effects are heterogenous: more subjects maximize returns while those who pay off debt end up repaying larger amounts. Finally, we explore if debt also impacts risk and time tradeoffs, and we also replicate our main results in the one-shot scenarios.

A. Evidence of The Burden of Debt

In this subsection, we show evidence that subjects with debt are more likely to deviate from the return-maximizing strategy. As a first step, consider the allocation decision on day 1 where everyone has the same starting wealth and number of points to allocate. Accounts start with different initial balances between *No Debt* and *Low Debt*, but the available interest rates are the same, and hence subjects have the same set of opportunities and financial incentives. Figure 1 shows that subjects in both treatments allocate the largest share of points to the account with the highest interest rate (Savings 1, 20%). However, while in *No Debt* subjects allocate 73 percent of their initial 500 points to Savings 1, subjects in *Low Debt* allocate less than half (47 percent) to that same account, a 26 percentage point difference (p.value<0.000). This difference is mainly driven by subjects in *Low Debt* allocating 18 percentage points (p.value<0.000) more to the debt account with the highest interest rate (Debt 1, 15%) compared to the account in *No Debt* with the same interest rate (Savings 3, 15%). There is also an increase in the other savings account (Savings 2, 10%), but of a smaller magnitude (6 percentage points, p.value=0.002). Thus we find that subjects with debt are less likely to allocate points to the accounts that would benefit them the most. These differences in behavior translate into financial losses; in *Low Debt* returns from the initial allocation are 9.3 percent lower. On average subjects in *No Debt* obtain a 17.5 percent average return on their first allocation decision, as compared to 15.8 percent in *Low Debt* (p.value<0.000).

Figure 1: Allocation shares of the initial endowment in day 1

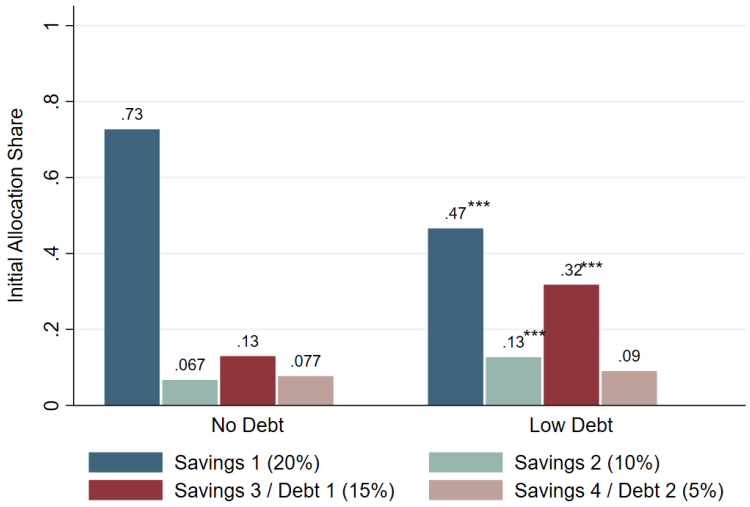

Taking into consideration all four days of our experiment, we still find more deviations from return-maximizing strategies for subjects who hold debt. Subjects now maximize returns if they allocate all their points to the highest interest rate account across all allocations. We find that when holding debt, subjects are almost three times less likely to maximize returns (Figure 2). In *No Debt*, 38 percent of subjects maximize returns in all decisions while in *Low Debt* the rate drops down to 13 percent, a 25 percentage point difference (p.value<0.000).

Figure 2: Percent of subjects that maximize returns in all decisions

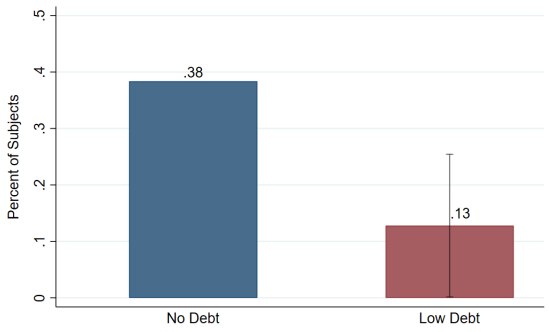

Notes: Error bars correspond to the 95% confidence interval on the difference across treatments.

We interpret these results as evidence that subjects are not behaving randomly and instead making debt-biased financial mistakes. Despite a lower share of subjects

maximizing returns in *Low Debt*, most subjects instead allocate points to the account with the 2nd highest interest rate. In fact, the ranking of the allocation shares directly follows the ranking of the interest rates in *Low Debt*. Furthermore, we do not find evidence that subjects use other strategies present in the literature such as equal split or balance-matching heuristics.[13] Since Debt 1 has a lower balance in absolute value than Debt 2, such balance strategies cannot explain why subjects predominantly allocate points to Debt 1. In the next subsection, we further explore what strategies subjects follow when holding debt.

B. Repayment Strategies when Holding Debt

Given the evidence on how subjects with debt deviate more often from return-maximizing behavior, we now explore what strategies are more prevalent when holding debt. One hypothesis that we aim to rule out is that subjects with debt behave more erraticaly i.e. allocating points randomly without taking interest rates in consideration. Another possibility is that subjects focus first on repaying their debt and then switch to maximizing returns.

We start by showing evidence that subjects' strategies are contingent on outstanding negative balances. In Figure 3 we show the total amount of points allocated to debt accounts in *Low Debt* and their equivalent savings accounts in *No Debt*. While in *No Debt* 38 percent of subjects allocate no points whatsoever to the second highest interest account, only 13 percent of the subjects follow the same strategy in *Low Debt*, the exact same numbers as those who maximize returns in both treatments. We observe a similar pattern for the lowest interest rate account, with 51 percent of subjects allocating no points in *No Debt* versus 28 percent in *Low Debt*. Furthermore, we find the biggest difference when isolating subjects who allocate a non-zero amount to the debt accounts and their equivalents in *No Debt*. A large fraction of subjects, 34 percent and 17 percent, zero-out the outstanding negative balances of Debt 1 and Debt 2 respectively. These spikes are not mechanical as subjects can still allocate points after fully repaying the debt balances; in fact, Figure 3a shows that a small number of subjects end up with debt accounts with positive balances.[14] We do not observe the same pattern for the equivalent accounts in *No Debt*.

[13] In Gathergood et al. (2019) debt repayments are consistent with a balance-matching heuristic under which the share of repayments is matched to the share of the balance.

[14] In that case the accounts move from generating negative to positive returns.

Figure 3: Total points allocated to 15% and 5% accounts

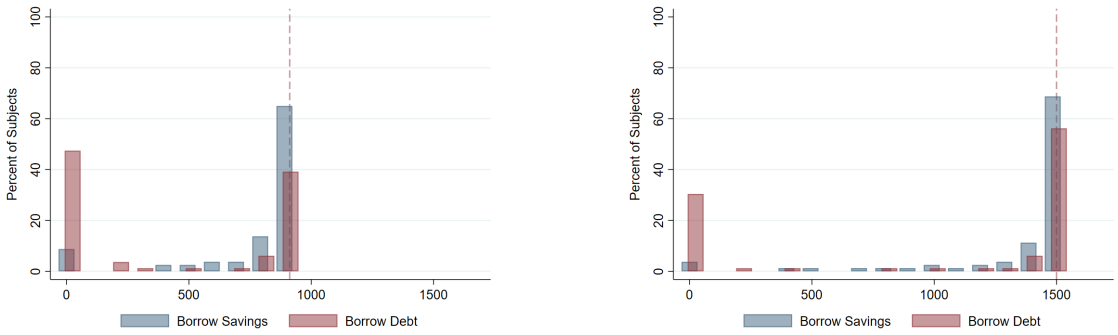

(a) Savings 3/Debt 1 (15% interest)

(b) Savings 4/Debt 2 (5% interest)

Notes: Vertical dashed lines indicate the outstanding negative balances for each debt account.

This evidence indicates that negative balances are perceived differently for a large share of subjects. Debt causes debt-specific strategies rather than erratic behavior. A large fraction of subjects in *Low Debt* fully repaid at least one outstanding negative balance. We provide additional evidence using the strategy descriptions that subjects provide on the last day. In their descriptions, 24 percent of the subjects in *Low Debt* explicity mention using a strategy of first repaying their outstanding debt and then focusing on maximizing returns. This hindsight was not incentivized and thus should be interpreted carefully; however, most descriptions are consistent with the observed behavior in the experiment. See Appendix E for example answers that subjects gave.

In the next subsection, we use our *High Debt* treatment to explore what strategies subjects follow when negative balances cannot be fully repaid. Do subjects make even larger financial mistakes? Or on the other hand, do they ignore debt accounts altogether and focus on maximizing returns?

C. The Effects of Larger Debt Balances: High Debt

We observe heterogenous effects in behavior when negative balances cannot be fully repaid: more subjects maximize returns in all decisions, but those who focus on repaying debt end up allocating even more points to debt than in *Low Debt*. Now 26 percent of subjects in *High Debt* maximize returns in all decisions which is larger than the 13 percent in *Low Debt* (p.value=0.042), but it is still lower than the 38 percent in *No Debt* (p.value=0.063). Compared to the *Low Debt* treatment, we find fewer subjects allocating points to either debt account. Figure 4 shows that when initial negative balances are larger, 27 percent of subjects do not allocate any points to Debt 1 and more than 57 percent do not allocate any points to Debt 2.

Figure 4: Total points allocated to debt accounts

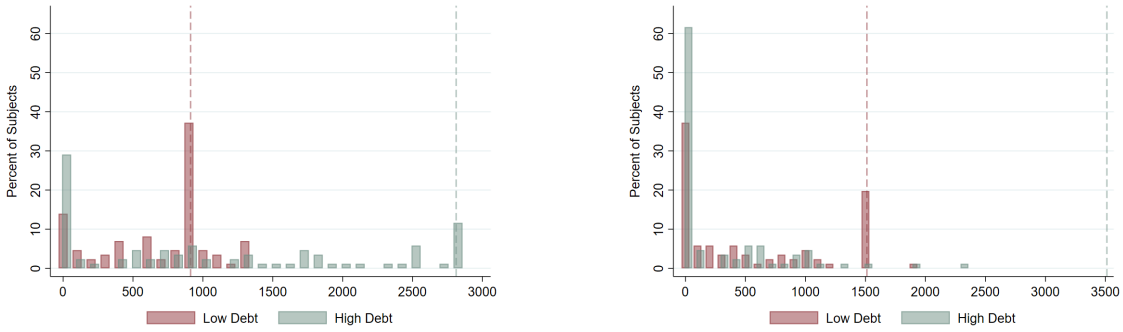

(a) Debt 1/Debt 1 (15% interest)

(b) Debt 2/Debt 2 (5% interest)

Notes: Vertical dashed lines indicate the outstanding negative balances for each debt account.

The fact that we observe lower deviations from return-maximizing behavior in *High Debt* could lead us to conclude that larger debt balances reduce the *Burden of Debt*. However, we instead find that the average measure masks the heterogeneity of our results. Despite the increase in return-maximizing behavior, there are subjects who allocate a large amount of points to debt repayment. Subjects who do not maximize returns in *High Debt* allocate a larger share (8.9 percentage points, p.value=0.019) of their total points to debt repayment compared to those who do not maximize returns in *Low Debt*. Furthermore, 10 percent fully repay the debt account with the highest interest rate (Debt 1) which has an initial outstanding balance of 2,900 points. This implies that these subjects have allocated most, if not all, of their returns to this account.[15]

Our evidence shows that subjects' strategies depend on the size of the initial debt balances. When debt balances cannot be fully repaid, more subjects maximize returns but those that do not perform even worse. If subjects only considered interest rates, behavior in *Low Debt* and *High Debt* should not differ as the only difference is the starting balances that do not affect the payoff maximizing strategies. This evidence suggests that balances should be taken into consideration when modeling behavioral responses to debt. Furthermore, this could indicate that some subjects might be partially sophisticated with respect to debt repayments, and only allocate points if they can repay these outstanding balances completely.[16]

[15]The fact that we do not observe a similar spike around 900 or 1,500 points as in *Low Debt* also indicates that subjects' decisions depend on the size of the negative balances.

[16]This reasoning is consistent with previous experimental evidence showing that subjects exhibit a premium for closing negative balances, see Amar et al. (2011); Besharat et al. (2014); Kettle et al. (2016). In addition, Zhang et al. (2020) shows how people disengage from debt when full repayment seems difficult.

D. Risk and Time Preferences Elicitation

Our measures of risk and time preferences do not show systematic differences between treatments with the exception of risk preferences in one domain. As expected from an experimental population, our subjects are on average risk averse, see Table 7.[17] Using all the elicitations after the allocation decisions, we find that subjects with debt exhibit more risk taking behavior but only when both options involve points–even after controlling for initial responses and allocation decisions. In *Low Debt*, subjects require 5 percent more points to forego the risky prospect of the lottery. Larger debt balances aggravate this effect, with subjects in *High Debt* requiring 7 percent more. However, these effects are no longer present when the trade-off involves dollars for sure versus a lottery of points. Similarly, we find no significant differences for time trade-offs when one option involves money. For the time tradeoffs between points vs points, we find a similar pattern as in the risk question for that same domain: subjects with debt discount future payments more heavily and more so when debt balances are higher, although these differences are only significant for *High Debt*.[18]

While we may be concerned that people holding debt behave more erratically, our evidence suggests this is not the case. For most tradeoffs, subjects with and without debt answer similarly. Despite the large differences in behavior from the previous subsections, these effects do not seem to extrapolate to risk or time choices, except for risk in the point domain. The latter suggests that when measuring risk preferences, the domain of the trade-offs matters for people who hold debt.

E. Robustness: One-shot

We replicate our main result using the three one-shot scenarios that subjects face after the last allocation decision. Here subjects face three one-shot versions of the main allocation decisions, all on the same day, in random order. Regardless of their previous experience and treatment assignment, we find that subjects are more likely to maximize returns in *One-shot No Debt* than in any of the other one-shot scenarios with debt.

[17] Experimental subjects are in general risk averse, and often excessively so given the low stakes, see Rabin (2000).

[18] These results are in line with the findings of Meier and Sprenger (2010) where people with credit card debt tend to be more present biased. However, our measure captures time discounting but not present bias.

Table 7: Main Treatments: estimation output using risk and time elicitation questions

	(1) Risk 1 Points vs Points	(2) Risk 2 Money vs Points	(3) Time 1 Points vs Points	(4) Time 2 Money vs Points
Mean of dep. var	0.240***	0.202***	0.249***	0.194***
	(0.058)	(0.066)	(0.063)	(0.063)
Low Debt	0.049**	-0.018	0.035	-0.005
	(0.025)	(0.030)	(0.025)	(0.028)
High Debt	0.070***	0.015	0.055**	-0.006
	(0.024)	(0.028)	(0.027)	(0.026)
Initial Risk	0.182***	0.200***	0.085*	0.059
	(0.056)	(0.060)	(0.050)	(0.058)
Initial Time	0.080	0.219***	0.262***	0.308***
	(0.053)	(0.065)	(0.045)	(0.065)
Observations	1032	774	774	774

Notes: Results from a linear regression with clustered standard errors at the individual level in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is an index of risk and time preferences that ranges from 0 to 1. Higher numbers imply higher risk-seeking and time-discounting behavior. Low Debt is a treatment dummy that equals 1 if the subject participated in *Low Debt*. Similarly for High Debt dummy. The regression also includes responses to the initial survey, controls for order effects, dummy if participant maximize returns in all decisions, the number of errors in the instructions, demographic controls (Gender, Ethnicity, Age, and Schooling), controls for whether they hold debt or student loans, and the personal impacts from Covid-19. The full output is presented in Table A.16 in the Appendix.

Among subjects that completed the *No Debt* treatment, 59 percent maximize returns in *One-shot No Debt*. In contrast, when facing the *One-shot Low Debt* scenario, the percentage of subjects who maximize returns is 15 percentage points lower. A similar difference arises in the *One-shot High Debt* scenario.[19]

[19]Note that unlike our main treatments, we do not observe a significant difference between *One-shot Low Debt* and *One-shot High Debt*.

Figure 5: Percent of subjects that maximize returns in one-shot scenarios

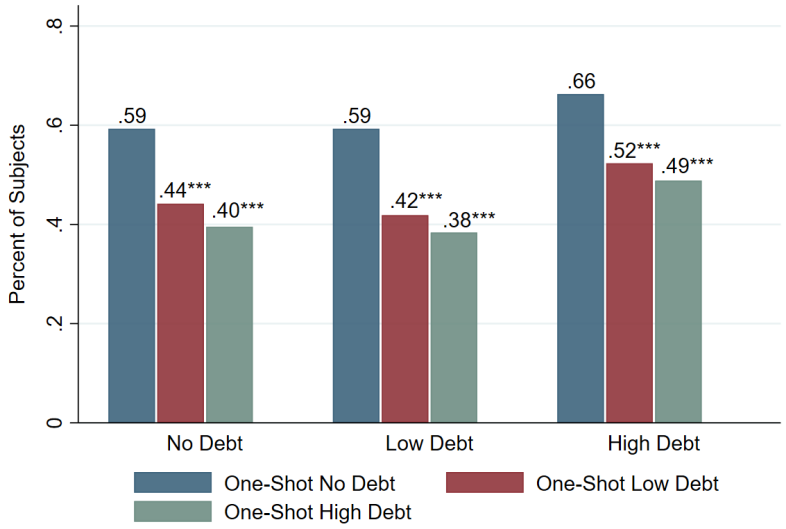

Notes: Significance test compare *One-shot No Debt* with the corresponding one-shot scenario.

When looking at the subjects who completed any of the other treatments, *Low Debt* and *High Debt*, a similar pattern arises. More subjects maximize returns in the *One-shot No Debt* scenario compared to *One-shot Low Debt* and *One-shot High Debt*. Despite previous experience with debt accounts, these subjects behave similarly, with no statistically significant differences based on initial treatment assignment.

III. Increased Agency and the *Burden of Debt*

We find that when given the option of reallocating their savings balances, many subjects do not use this option to increase their returns and instead use it to repay debt. Under our *Redistribution* treatments, subjects have more agency over their accounts and can reallocate points from their Savings 1 and Savings 2 balances. When we introduce this option, we see that some subjects actually get lower returns compared to the worst possible outcome without any redistribution. Debt-biased behavior in redistribution and allocation decisions compounds and exacerbates financial losses. Furthermore, we find that subjects who attempt to repay their debt do so sub-optimally. Altough feasible, no subject manages to fully repay their debt by the end of the experiment.

A. Redistribution Decisions

Similar to our main results, we find that redistribution decisions are also debt-biased. In both *Redistribution* treatments, most subjects redistribute balances in the first day.

However, despite similar uptakes, there are large differences in how subjects employ these redistribution options. Since subjects have an initial endowment of 500 points, any allocation above 500 indicates that subjects are taking advantage of the redistribution of balances. Hence, we define a subject as consolidating towards an account if they allocate more than 500 points to it on day 1. Note that this is only feasible through redistributing balances. Table 8 shows that in *Redistribution No Debt* subjects are twice as likely to redistribute balances towards the high interest rate account (37 percent vs 17 percent, p.value=0.004), hence maximizing their returns on day 1. On the other hand, 34 percent of subjects in *Redistribution Debt* consolidate towards Debt 1 (15% interest, -600 initial balance). Therefore, while in both conditions subjects use the additional option to redistribute balances, subjects with debt are less likely to use it for maximizing their returns.

Table 8: Redistribution Decisions: Percentage of subjects that consolidate on day 1

	Redistribution No Debt	**Redistribution Debt**	P.value
Savings 1	37.04	16.88	0.004
Savings 2	3.7	2.60	0.692
Savings 3 / Debt 1	13.58	33.77	0.003
Savings 4 / Debt 2	8.64	12.99	0.384

Notes: Subjects are assigned to one category if they allocate more than the initial endowment in Day 1 to that account. Allocating more than the initial endowment of 500 points is only feasible if subjects redistribute balances from Savings 1 or Savings 2. With the exception of Savings 1, subjects can be classified in more than one category.

In addition to differences in redistribution decisions, we also replicate our allocation results from the main treatments: subjects holding debt accounts are less likely to allocate points towards the account with the highest interest rate. When we classify subjects based on their allocation decisions, similarly as in the main treatments, we find that 43 percent of subjects exclusively allocate points to the highest interest account in *Redistribution No Debt* compared to only 14 percent in *Redistribution Debt* in all decisions (p.value<0.000).

Despite this difference in allocation decisions, subjects who only allocate to Savings 1 exhibit a similar redistribution behavior. In both treatments, approximately three-quarters of the subjects who maximize returns in their allocation decisions also redistribute as much as possible towards Savings 1 (69 percent vs 81 percent, p.value=0.597), hence maximizing their returns. This is not necessarily surprising, as in both treatments this involves redistributing points between savings accounts without affecting any negative balance, thus there are no treatment differences in where the redistributed points

originate from.

B. Compounding Debt-biased Behavior

Debt-biased redistribution and allocation decisions compound and lead to subjects obtaining lower returns in *Redistribution Debt*. When looking at the total returns that subjects make throughout the experiment (Figure 6) two main differences arise: First, subjects in *Redistribution Debt* are more than twice as likely to obtain the maximum attainable return (30 percent vs 12 percent, p.value=0.005).[20] Second, subjects with debt are more likely to obtain lower returns than what they would have obtained under the worst possible allocation, although this difference is marginally significant (21 percent vs 10 percent, p.value=0.059).[21] For this latter group, removing the option to redistribute balances would have actually increased their payoffs.

Figure 6: Total returns

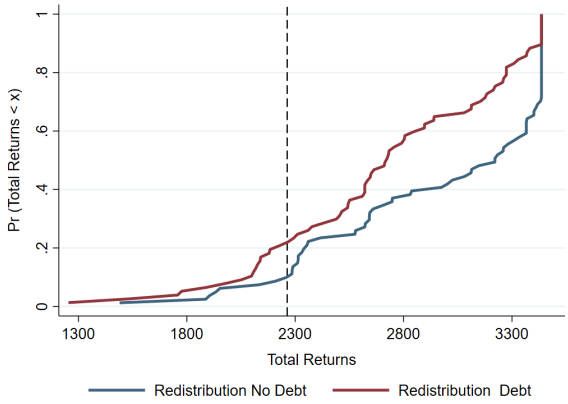

Notes: Vertical dashed line indicates the returns obtained from the worst allocation decision without any redistribution.

Although feasible, we find no subject fully repaying their outstanding debt balances. For our given parameters, a subjects that consolidates into Savings 1 and allocates all their points to it would start day 4 with an endowment of 922 points and a balance of 4308 points in Savings 1. These amounts are enough to completely repay all outstanding debt (4900 points) before the end of the experiment. However, we find no subjects following such strategy.

[20]This return corresponds to following the strategy of consolidating towards Savings 1 in the first day and allocating all the endowment and subsequent returns also to Savings 1, which results in 3,436 points.

[21]The total returns for the worst possible allocation decision corresponds to 2,263 points. This is based on following the strategy of allocating the initial endowment and subsequent returns to the account with the lowest interest rate (5%), in addition to not redistributing any points.

IV. The *Burden of Debt* in Borrowing Decisions

When given the opportunity to borrow, subjects could make two financial mistakes: first, they could be reluctant to borrow and miss a profitable opportunity, and second, they could borrow but misallocate the funds. We find evidence that subjects are reluctant to borrow from debt accounts, but those who take advantage of the borrowing opportunities are not negatively selected. As discussed earlier, the available borrowing opportunities, regardless of treatment, can always provide a profitable gain: by moving points from lower interest accounts (15% and 5%) to a 20% interest rate account. Thus subjects who choose not to borrow end up missing out on these potential gains. We therefore interpret differences in borrowing behavior as evidence of the *Burden of Debt*.

A. The Decision to Borrow

Debt severely impacts borrowing decisions: 34 percent of subjects in *Borrow Debt* borrow the maximum amount compared to 63 percent in *Borrow Savings* (29 percentage point difference, p.value<0.000). These differences are also present for subjects that, apart from borrowing decisions, maximize returns in their allocation decisions. Figure 8 shows that 64 percent of subjects in *Borrow Savings* borrow the maximum amount from the 15% account, and 66 percent borrow the maximum amount from the 5% account. Clearly, subjects take advantage of the borrowing opportunities. In contrast, subjects in *Borrow Debt* are almost half as likely (38 percent) to borrow the maximum from their 15% account, and 52 percent fully borrow from the 5% account. Furthermore, only 8 percent of subjects in *Borrow Savings* borrow nothing from the 15% account, but 46 percent of subjects borrow nothing from the equivalent account in *Borrow Debt*. That same qualitative difference persists for the 5% account, though not as dramatically. Rather than a steady increase in points borrowed up to the cap, we instead find a bimodal distribution. Subjects in *Borrow Debt* either borrow the maximum amount or they borrow nothing at all. It is worth emphasizing that the only difference across treatments is the initial balances, and thus the labeling, of the borrowing accounts. Hence, we interpret these results as evidence that subjects are reluctant to borrow from debt accounts.

Figure 7: Percent of subjects that borrow the maximum amount from both accounts

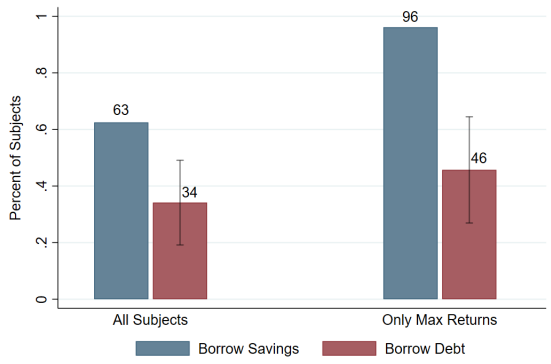

Notes: Only Max Returns restricts to subjects who allocate all points to the highest interest account in all allocation decisions. Error bars correspond to the 95% confidence interval on the difference across treatments.

Figure 8: Total points borrowed from each account

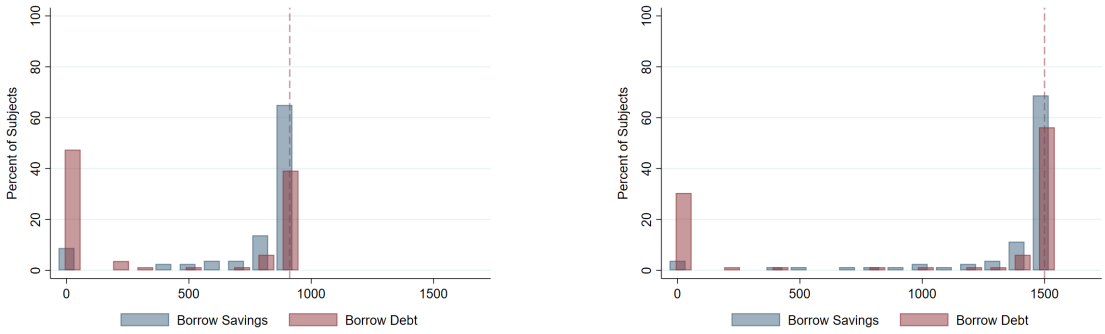

(a) Savings 5/Debt 1 (15% interest)

(b) Savings 6/Debt 2 (5% interest)

Notes: Vertical dashed lines indicate the borrowing limits for each account.

While these results are for our entire sample, when we restrict to subjects who understand how to maximize their returns the difference in borrowing behavior persists. Figure 7 shows that within this restricted sample, 96 percent borrow the maximum amount from both accounts in *Borrow Savings*. In contrast, in *Borrow Debt* not even half of the subjects (46 percent), who otherwise maximize returns, borrow the maximum amount from both accounts. Thus we see that debt still causes financial mistakes even for those subjects who are otherwise financially sophisticated in the context of our experiment.

Given the difference in borrowing behavior, we also find differences in the average returns subjects make each day. By the end of the experiment, subjects in *Borrow Savings* have 5 percent higher returns than those in *Borrow Debt* (p.value=0.061). These differences are entirely explained by different borrowing behavior and not by different allocation decisions. Once we control for subjects borrowing the maximum amount, there is no difference in returns across treatments (p.value=0.586).
blablalblablala check the p.values and update the table

Table 9: Borrowing Treatments: returns and payments estimation output

	(1) Log Total Returns	(2) Log Total Returns	(3) Log Total Returns
Sample	All Subjects	Max Returns	All Subjects
Mean of dep.var.	8.618***	8.676***	8.490***
	(0.078)	(0.090)	(0.052)
Borrow Debt	-0.050*	-0.078**	0.013
	(0.027)	(0.031)	(0.023)
Borrow Max			0.186***
			(0.022)
Observations	162	61	162

Notes: Results from a linear regression with robust standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is the log of the cumulative returns in the four allocation decisions and the final payments that subjects obtained without the participation fee. Borrow Max is a dummy that equals 1 if the subject borrowed the maximum amount by the end of the experiment. The regression also includes the number of errors in the instructions, demographic controls (Gender, Ethnicity, Age, and Schooling), controls for whether they hold debt or student loans, and the personal impacts from Covid-19. The full output is presented in Table A.17 in the Appendix.

One possible explanation for our results is that subjects anticipate that after borrowing, and hence having debt, they will make worse allocation decisions similar to our previous main treatments. Nonetheless, we rule out this possibility by design since subjects cannot allocate any points to the accounts they borrow from. Even in *Borrow Debt*, subjects are only allowed to allocate points to savings accounts so concerns about anticipated debt-biased mistakes after borrowing cannot explain the reluctance to borrow that we observe.

Taken altogether, we find that the *Burden of Debt* prevents subjects from undertaking a profitable opportunity. These financial mistakes apply to all subjects, even those who otherwise maximize returns in their allocation decisions. These differences cannot be explained by participants anticipating debt-biased behavior, as by design we

rule out the possibility to repay debt. Whereas in our main treatments the *Burden of Debt* manifests as financial mistakes in debt repayments, we now show that financial mistakes are also present in borrowing decisions, in line with previous findings in the debt aversion literature.[22]

B. The Selection of Borrowers

Though in our setting borrowing opportunities should be profitable, it is still possible for borrowers to make additional financial mistakes by misallocating the additional funds. We find that this is not the case: borrowers are positively selected. Subjects who borrow the maximum amount from both savings accounts allocate 87 percent of their total points to the highest interest rate account, and subjects who borrow from debt allocate 79 percent. This result also holds when we look at subjects who maximize returns in all allocation decisions; 50 and 57 percent of subjects who borrow the maximum also maximize returns in all decisions in *Borrow Savings* and *Borrow Debt* respectively, with no significant difference after adding our standard controls (p.value=0.108).

While our main treatment groups relied on initially endowing subjects with debt to find treatment effects, we now have again found evidence of the *Burden of Debt* when indebtedness is now a choice instead of inherent. In this second scenario, subjects do not start with debt. Yet, we still find reluctancy to borrow from debt accounts rather than savings accounts which leads to differences in final payoffs.

V. Discussion

We have provided evidence of the *Burden of Debt* in a novel experimental design that has shown financial losses caused by debt-biased behavior. Allocation, redistribution and borrowing decisions are all affected by debt. In this section, we discuss the intuition for a stylized theoretical model of debt aversion that can rationalize our results and how it relates to other models in the behavioral literature. We then discuss the potential implications of debt aversion outside the lab.

One potential way to rationalize our results is through a model where subjects derive disutility from holding debt balances. For example, one could consider a utility function that incorporates a penalty $C(d) < 0$ for any outstanding debt balance d, where $C(\cdot)$ is decreasing in d, so larger balances impose a larger penalty. Under this simple model of debt aversion, subjects might focus on repaying negative balances, even if their interest rates are low, and turn down profitable borrowing opportunities to avoid incurring the penalty from holding debt. Our findings also suggest that if subjects dislike holding debt, this occurs even when outstanding balances are small

[22]For example Callender and Jackson (2005) and Caetano et al. (2011) find reluctance to borrow among prospective college students and bank customers respectively.

and can be easily repaid. In the *Low Debt* treatment and one-shot scenarios with debt, many subjects focus on repaying their outstanding debt balances and completely zero out their balances. Hence, a model of debt aversion should incorporate a discontinuous jump for any non-zero debt amount (i.e., $C(0) = 0$ and for any $d > 0$, $C(d) < \alpha$ where $\alpha < 0$). This last assumption is in line with prior experimental findings where subjects exhibit a premium for closing out small debt accounts (Amar et al., 2011; Kettle et al., 2016).

Our proposed model of debt aversion closely relates to the concept of loss aversion in Tversky and Kahneman (1991) where losses are overweighted, but in our case it should apply to negative balances. However, there are three additional assumptions required to rationalize our results using a model of loss aversion. First, we must assume that subjects do not integrate over all their accounts; if they did, they would always be in the gain domain and hence, would not incur any loss penalty. Therefore, we need to assume narrow framing in evaluating account balances.[23] Second, we assume a reference point at zero than does not change overtime. Models of loss aversion consider losses with respect to a reference point, with different proposals in the literature for determining these reference points.[24]. Since we define a debt balance as any balance lower than zero, our reference point is always at zero. Finally, as previously mentioned, many subjects fully repay their outstanding debt balances, even when these amounts are small. Thus, our third assumption is a discontinuous jump for any non-zero loss. With all these assumptions, a model of loss aversion can also rationalize our findings.

Debt aversion provides a different approach to think about the unintended consequences and benefits of debt related policies. Relief and debt forgiveness programs could provide additional benefits above and beyond their direct wealth impact as they would also reduce the *Burden of Debt*. In these cases, reducing debt could create a virtuous cycle, where lower debt improves subsequent financial decisions and prevents indebtedness in the future. On the other hand, if indebtedness causes sub-optimal financial decisions, even short-term borrowing in times of financial hardship could have long-term consequences. Gelman et al. (2015) shows that a small shift in income timing can drive extended indebtedness for highly-constrained consumers. Similarly, Leary and Wang (2016) shows that a large fraction of the demand for payday loans is caused by sub-optimal savings and consumption decisions, and (Carvalho et al., 2019) finds that "misfortunes" also play an important role in demand for payday loans.

How we present and frame debt could help alleviate debt aversion. For example, debt aversion suggests that financial aid programs in the form of grants might be more effective than loans in encouraging uptake. Prospective college students might be more

[23] A common finding in the mental accounting and narrow framing literature is that financial behaviors not always align with considerations of net worth (Thaler, 1999). See Barberis et al. (2006) for a discussion on narrow framing and its implications.

[24] For example, Kőszegi and Rabin (2006, 2009) propose using expectations as a reference point.

likely to enroll after receiving a grant than a loan, even if the return on higher education compensates for the amount borrowed. Field (2009) provides evidence of such framing effects by showing that students who receive a grant instead of a financially equivalent student loan are more likely to pursue a career in the public sector.[25] Varying how debt is framed and presented could help reduce the consequences of debt aversion and alleviate the *Burden of Debt*.

VI. Conclusion

With over 80 percent of adults in the US holding some form of debt (Pew, 2015), debt has become a pervasive aspect of Americans' lives. Despite the empirical evidence showing that holding debt correlates with sub-optimal decision-making, how debt impacts decisions has been understudied in behavioral economics. In this paper we show that debt causes behavioral biases that can lead to lower financial wealth. We refer to the financial losses caused by this debt-biased behavior as the *Burden of Debt* and provide experimental evidence on its existence and its implications.

To show that debt causes financial mistakes, we develop a new experimental design where subjects are assigned debt randomly, and actions and payoff opportunities are identical across treatments. In this controlled environment, where factors like selection or uncertainty play no role, we are able to quantify the opportunity cost of subjects' financial mistakes. First, we show in an allocation decision problem that subjects with debt are less likely to maximize returns as they focus on repaying negative balances. Once we increase the outstanding negative balances, and they become harder to be fully repaid, we find a heterogeneous effect: more subjects do not attempt to repay debt and maximize returns instead, but those who try to repay debt end up with larger financial losses. This highlights the complexity of debt behavior, and how its impacts can vary widely across subjects. We believe understanding the determinants of this heterogeneity and what factors predicts the *Burden of Debt* is a promising avenue for future research.

Of particular interest is how these financial mistakes can prevent profitable investments when they require access to credit. In two additional treatments, we show that the *Burden of Debt* is also relevant for borrowing decisions. In a setting where profitable opportunities are available, we find a widespread reluctance to borrow from debt accounts. This effect persists even when we restrict to subjects who otherwise maximize returns, and hence show that our effects are also present in more sophisticated subjects. Our evidence on the reluctance to borrow can help us understand why in some cases people do not borrow enough. While in general one could argue that people might be borrowing too much, this may not be the case for some types of debt such as

[25] In that setting, both were equivalent because taking a job in the public sector resulted in forgiveness of the loan and a job in the private sector required returning the grant.

student loans. Avery and Turner (2012) argues that given the high returns to education *"the claim that student borrowing is too high across the board can–with the possible exception of for profit colleges–clearly be rejected"*. Along these lines, Callender and Jackson (2005) shows that, in England, the fear of debt can prevent qualified students from attending college especially for those from low socioeconomic backgrounds.

We leave to future research potential avenues to alleviate the *Burden of Debt*. Evidence from our *High Debt* treatment suggests that higher debt balances are not necessarily detrimental for everyone. Though subjects who are debt averse do worse, more subjects also end up maximizing returns. Much of the evidence we find is consistent with subjects trying to get rid of their debt as quickly as possible. However, this does not have to come at the expense of maximizing earnings. We believe that alternative debt repayment methods like income-based repayment also have the potential to mitigate the adverse consequences of debt by better aligning the incentives for maximizing earnings and repaying debt faster. Similarly, policies that allow for debt repayment deferral, such as seen with federal student loans, could also mitigate the negative consequences of the *Burden of Debt*. Along these lines, Hershfield et al. (2015) advocates for incorporating behavioral and psychological biases into the policy discussion around indebtedness.

While our experiment studies the consequences of a stylized version of debt, we believe the *Burden of Debt* still has interesting implications outside a lab setting. Many Americans live in precarious situations that debt could exacerbate. Lusardi et al. (2011) finds that one quarter to one half of households report being unable to come up with $2,000 within the next month to cope with an unexpected expense. Financially constrained households could end up borrowing at high interest rates after adverse wealth shocks. If debt causes worse financial decisions and prevents wealth accumulation, the *Burden of Debt* has the potential to explain why debt is so prevalent. To capture the full repercussions of debt, we must incorporate all of its consequences.

VII. Bibliography

Hunt Allcott, Joshua Kim, Dmitry Taubinsky, and Jonathan Zinman. Are high-interest loans predatory? theory and evidence from payday lending. Technical report, Working paper, 2020.

Moty Amar, Dan Ariely, Shahar Ayal, Cynthia E Cryder, and Scott I Rick. Winning the battle but losing the war: The psychology of debt management. *Journal of Marketing Research*, 48(SPL):S38–S50, 2011.

Christopher Avery and Sarah Turner. Student loans: Do college students borrow too much–or not enough? *Journal of Economic Perspectives*, 26(1):165–92, 2012.

Saad Azmat and Isabel H. Macdonald. The psychological cost of debt: Evidence from housing mortgages in pakistan. 2020.

Nicholas Barberis, Ming Huang, and Richard H Thaler. Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American economic review*, 96(4):1069–1090, 2006.

Marianne Bertrand and Adair Morse. Information disclosure, cognitive biases, and payday borrowing. *The Journal of Finance*, 66(6):1865–1893, 2011.

Ali Besharat, François A Carrillat, and Daniel M Ladik. When motivation is against debtors' best interest: The illusion of goal progress in credit card debt repayment. *Journal of Public Policy & Marketing*, 33(2):143–158, 2014.

Ali Besharat, Sajeev Varki, and Adam W Craig. Keeping consumers in the red: Hedonic debt prioritization within multiple debt accounts. *Journal of Consumer Psychology*, 25(2):311–316, 2015.

Kathleen Burke, Jesse B Leary, and Jialan Wang. Information disclosure and payday lending in texas. *Unpublished mimeo*, 2016.

Brian C Cadena and Benjamin J Keys. Can self-control explain avoiding free money? evidence from interest-free student loans. *Review of Economics and Statistics*, 95(4): 1117–1129, 2013.

Gregorio Caetano, Miguel Palacios, and Harry Anthony Patrinos. *Measuring aversion to debt: An experiment among student loan candidates*. The World Bank, 2011.

Claire Callender and Jonathan Jackson. Does the fear of debt deter students from higher education? *Journal of social policy*, 34(4):509–540, 2005.

Leandro Carvalho, Arna Olafsson, and Dan Silverman. Misfortune and mistake: The financial conditions and decision-making ability of high-cost loan borrowers. Technical report, National Bureau of Economic Research, 2019.

Erica Field. Educational debt burden and career choice: Evidence from a financial aid experiment at nyu law school. *American Economic Journal: Applied Economics*, 1 (1):1–21, 2009.

Vyacheslav Fos, Andres Liberman, and Constantine Yannelis. Debt and human capital: Evidence from student loans. *Available at SSRN 2901631*, 2017.

John Gathergood, Neale Mahoney, Neil Stewart, and Jörg Weber. How do individuals repay their debt? the balance-matching heuristic. *American Economic Review*, 109 (3):844–75, 2019.

Michael Gelman, Shachar Kariv, Matthew D Shapiro, Dan Silverman, and Steven Tadelis. *How individuals smooth spending: Evidence from the 2013 government shutdown using account data*. National Bureau of Economic Research Cambridge, MA, 2015.

Johannes Haushofer and Ernst Fehr. On the psychology of poverty. *Science*, 344(6186): 862–867, 2014.

Paul J Healy. Explaining the bdm—or any random binary choice elicitation mechanism—to subjects. Technical report, mimeo, 2016.

Hal E Hershfield, Abigail B Sussman, Rourke L O'Brien, and Christopher J Bryan. Leveraging psychological insights to encourage the responsible use of consumer debt. *Perspectives on Psychological Science*, 10(6):749–752, 2015.

Shinsuke Ikeda and Myong-Il Kang. Generalized hyperbolic discounting, borrowing aversion, and debt holding. 2011.

Yan Ji et al. Job search under debt: Aggregate implications of student loans. *Massachusetts Institute of Technology Unpublished Working Paper*, 2016.

Supreet Kaur, Sendhil Mullainathan, Suanna Oh, and Frank Schilbach. Does financial strain lower productivity? Technical report, Working Paper, 2019.

Keri L Kettle, Remi Trudel, Simon J Blanchard, and Gerald Häubl. Repayment concentration and consumer motivation to get out of debt. *Journal of Consumer Research*, 43(3):460–477, 2016.

Benjamin J Keys, Devin G Pope, and Jaren C Pope. Failure to refinance. *Journal of Financial Economics*, 122(3):482–499, 2016.

Botond Kőszegi and Matthew Rabin. A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165, 2006.

Botond Kőszegi and Matthew Rabin. Reference-dependent consumption plans. *American Economic Review*, 99(3):909–36, 2009.

David Laibson. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478, 1997.

David Laibson, Andrea Repetto, and Jeremy Tobacman. Estimating discount functions with consumption choices over the lifecycle. Technical report, National Bureau of Economic Research, 2007.

Jesse Leary and Jialan Wang. Liquidity constraints and budgeting mistakes: Evidence from social security recipients. *Unpublished*, 2016.

Mi Luo and Simon Mongey. Student debt and job choice: Wages vs. job satisfaction. *New York University*, 2016.

Annamaria Lusardi, Daniel J Schneider, and Peter Tufano. Financially fragile households: Evidence and implications. Technical report, National Bureau of Economic Research, 2011.

Anandi Mani, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao. Poverty impedes cognitive function. *science*, 341(6149):976–980, 2013.

Stephan Meier and Charles Sprenger. Present-biased preferences and credit card borrowing. *American Economic Journal: Applied Economics*, 2(1):193–210, 2010.

Thomas Meissner. Intertemporal consumption and debt aversion: an experimental study. *Experimental Economics*, 19(2):281–298, 2016.

Ted O'Donoghue and Matthew Rabin. Doing it now or later. *American economic review*, 89(1):103–124, 1999.

Hakan Ozyilmaz and Guangli Zhang. Suboptimal credit card repayments: A laboratory experiment. *Available at SSRN 3430746*, 2019.

Trust Charitable Pew. The complex story of american debt: Liabilities in family balance sheets. *Washington, DC: Pew Charitable Trusts*, 2015.

Alejandro Ponce, Enrique Seira, and Guillermo Zamarripa. Borrowing on the wrong credit card? evidence from mexico. *American Economic Review*, 107(4):1335–61, 2017.

Drazen Prelec and George Loewenstein. The red and the black: Mental accounting of savings and debt. *Marketing science*, 17(1):4–28, 1998.

Matthew Rabin. Risk-aversion for small stakes: A calibration theorem. *Econometrica*, 68:1281–1292, 2000.

Matthew Ridley, Gautam Rao, P Vikram, and F Schilbach. Poverty and mental illness: Causal evidence. Technical report, Mimeo, 2019.

Jesse Rothstein and Cecilia Elena Rouse. Constrained after college: Student loans and early-career occupational choices. *Journal of Public Economics*, 95(1-2):149–163, 2011.

Frank Schilbach, Heather Schofield, and Sendhil Mullainathan. The psychological lives of the poor. *American Economic Review*, 106(5):435–40, 2016.

Anuj K Shah, Sendhil Mullainathan, and Eldar Shafir. Some consequences of having too little. *Science*, 338(6107):682–685, 2012.

Victor Stango and Jonathan Zinman. What do consumers really pay on their checking and credit card accounts? explicit, implicit, and avoidable costs. *American Economic Review*, 99(2):424–29, 2009.

Abigail B Sussman and Eldar Shafir. On assets and debt in the psychology of perceived wealth. *Psychological science*, 23(1):101–108, 2012.

Richard H Thaler. Mental accounting matters. *Journal of Behavioral decision making*, 12(3):183–206, 1999.

Amos Tversky and Daniel Kahneman. Loss aversion in riskless choice: A reference-dependent model. *The quarterly journal of economics*, 106(4):1039–1061, 1991.

Lise Vesterlund. Using experimental methods to understand why and how we give to charity. *Handbook of experimental economics*, 2:91–151, 2016.

Yi Zhang, Ronald T Wilcox, and Amar Cheema. The effect of student loan debt on spending: The role of repayment format. *Journal of Public Policy & Marketing*, 39 (3):305–318, 2020.

A. Appendix

A. Instructions

- In this experiment, you will have accounts that generate positive and negative returns over the duration of the experiment
- At each decision, you will allocate points between different accounts
- There is a total of 4 decisions, one today, and another on Friday, Monday and next Wednesday
- You will only receive your payments if you successfully complete all four decisions

Accounts

- Accounts generate returns, based on their balance and interest rates
- Balances change over time based on your decisions
- Interest rates are constant throughout the experiment
- Accounts with positive balances generate positive returns
- Accounts with negative balances generate negative returns
- We label accounts with positive starting balances as Savings accounts
- We label accounts with negative starting balances as Debt accounts

For example, a Savings account with a balance of 150 points and an interest rate of 30% would generate a positive return of 150 x 0.30 = 45 points

Similarly, a Debt account with a balance of -120 points and an interest rate of 40% would generate a negative return of -120 x 0.40 = -48 points

Endowment

- At each decision, you will receive an endowment of points that you will have to allocate between the accounts
- The endowment is calculated as the sum of the returns of all your accounts
- You can allocate the points as you wish between your available accounts as long as the sum is equal to the endowment
- Points allocated to an account increase its balance for all subsequent decisions

Timeline of the experiment

Decision 1: Wednesday (Today)

- You receive an endowment of points based on the initial balances and interest rates
- You must allocate this endowment between the available accounts
- The accounts generate returns, based on their balances and interest rates

Decision 2: Friday (Two days from today)

- Balances on Friday will reflect your previous decisions
- The sum of the returns from Wednesday is your endowment for this decision
- You will allocate the new endowment between the available accounts

Decision 3: Monday (Five days from today)

- Balances on Monday will reflect your previous decisions
- The sum of the returns from Friday is your endowment for this decision
- You will allocate the new endowment between the available accounts

Decision 4: Next Wednesday (A week from now)

- Balances on next Wednesday will reflect your previous decisions
- The sum of the returns from Monday is your endowment for this decision
- You will allocate the new endowment between the available accounts
- The accounts will generate returns and the experiment will end

Payments:

- You will be paid a $10 participation fee as well as a bonus based on your decisions
- We will sum all your points, including your final balances and the returns on next Wednesday. You will earn $1 for every 500 points
- Payments will be disbursed within 2 days of the end of the study
- Throughout the study there will be additional opportunities to make earnings, so please pay close attention to all the instructions

You will only receive your final payment if you finish the experiment

B. Initial Survey/Additional Questions

- You have the opportunity to earn monetary payoffs by answering lists of questions
- In each list, you will indicate what you prefer between two options. For example, a 50% chance of earning $1 versus $0.40 for sure
- One of these lists will be randomly selected and one question from it implemented
- Before you answer these questions lists, you will go through an explanation of the setting and how to answer these questions

- This is an example to help you understand how the questions work
- Imagine that you are given the choice between a 50% chance to get $1 or dollars for sure and you have a list of questions like this one:

Q#		Option A		Option B
1	Would you rather have:	50% chance of $1	or	$0.01 for sure
2	Would you rather have:	50% chance of $1	or	$0.02 for sure
3	Would you rather have:	50% chance of $1	or	$0.03 for sure
...	...	...	...	...
98	Would you rather have:	50% chance of $1	or	$0.98 for sure
99	Would you rather have:	50% chance of $1	or	$0.99 for sure
100	Would you rather have:	50% chance of $1	or	$1 for sure

In each question, you pick either Option A (50% chance of $1) or Option B (dollars for sure)

- After you answer all 100 questions, I will randomly pick one question and pay you the option you chose on that one question
- Each question is equally likely to be chosen for payment. Obviously, you have no incentive to lie on any question, because if that question gets chosen for payment then you'd end up with the option you like less.

I assume you're going to choose Option A in at least the first few questions, and I assume you're going to choose Option B in at least the last few questions. So at some point, you will switch from preferring Option A to Option B. To save time, just tell when you would switch from preferring Option A to Option B.

- I can then 'fill out' your answers to all 100 questions based on your switch point (choosing Option A for all questions before your switch point, and Option B for all questions at and after your switch point).

– I'll still draw one question randomly for payment. Again, if you lie about your true switch point you might end up getting paid an option that you like less.

At which dollar value would you switch from Option A to Option B?

C. Redistribution/Borrowing Decisions

Redistribution Treatments:

– Before you decide how to allocate your initial endowment of points, you have the opportunity to withdraw points from the Savings 1 and Savings 2 accounts

– You can withdraw up to 2000 points from Savings 1 and up to 400 points from Savings 2

– Any amount that you withdraw will be added to your endowment of points and will reduce the balances of Savings 1 and Savings 2 accounts

Borrowing Treatments:

– Before you decide how to allocate your initial endowment of points, you have the opportunity to withdraw [borrow] points from any of the locked Savings [Debt] accounts

– You can withdraw [borrow] up to 900 points from Savings 5 [Debt 1] and up to 1500 points from Savings 6 [Debt 2]

– Any amount that you withdraw [borrow] will be added to your endowment of points and will affect the returns that locked Savings [Debt] accounts generate

– You will still not be able to allocate any points to the locked Savings [Debt] accounts. [Any amount that you borrow will be repaid at the end of the experiment]

D. One-shot Parameters

After the last allocation decision, subjects face these one-shot scenarios in random order. In each decision, they must allocate 1000 points between the four available accounts. In this simplified version of our main experiment, there are no locked accounts and all decisions are done back-to-back. Parameters are chosen to mimic the main treatments of our experiment. Initial wealth is kept constant across the three one-shot scenarios, but net returns differ. For *One-shot Low Debt*, subjects can choose to fully repay at least one debt account, but this is not feasible in the *One-shot High Debt*.

Table A.10: Accounts in One-shot Scenarios

One-shot No Debt:

	Savings 1	**Savings 2**	**Savings 3**	**Savings 4**
Interest Rate	**20%**	**10%**	**15%**	**5%**
Balance	**200**	**100**	**300**	**200**

One-shot Low Debt:

	Savings 1	**Savings 2**	**Debt 1**	**Debt 2**
Interest Rate	**20%**	**10%**	**15%**	**5%**
Balance	**1000**	**1200**	**-600**	**-800**

One-shot High Debt:

	Savings 1	**Savings 2**	**Debt 1**	**Debt 2**
Interest Rate	**20%**	**10%**	**15%**	**5%**
Balance	**2000**	**2400**	**-1700**	**-1900**

E. Hindsight Examples

At the end of the final day, subjects are asked to describe how they would behave if they could do the experiment again with the benefit of hindsight. Specifically, we ask "After completing the experiment and with the benefit of hindsight, what strategy would you follow in order to make as many points as possible and obtain a high payoff?". This elicitation is not incentivized. We present some examples of the responses that subjects wrote.

We categorize the hindsight into three categories. Based on subjects provided hindsight and performance during their allocation decision, we categorize subjects as either "maximizing returns", "debt then maximize", and "other". Subjects who mention maximizing their returns or focusing on the high interest account are put into the first category. These subjects also have allocation strategies that generally match their given hindsight. Similar to this group, there is also another group of subjects who mention first being concerned with debt and wanting to pay if off and then wanting to follow the return maximizing strategy. Interestingly, there are even some subjects in the *No Debt* treatment who bring up similar concerns with debt even though in their cases debt has 0 balances and is only in the locked accounts. Finally, there are other subjects who do not describe either such type of strategy. These subjects either focus solely on debt, did not have cogent strategies, or describe behavior that is not necessarily reflected by their actions. When comparing *No Debt* with the two *Debt* treatments, we find a similar fraction of subjects who eventually maximize returns, i.e. in the *Debt* treatments, the number of subjects who act optimally from the start and those who say they focused on debt and then acted optimally roughly corresponds to the same fraction of subjects in the *No Debt* who simply described optimizing.

Maximizing returns:

> "I would always put all the endowment in the account with the highest interest rate."
>
> "Allocate all to the savings account with highest interest rate. Of course if it has bigger percentage than my debt accounts."
>
> "I would go with the same strategy again, which was putting as much as I could in my highest interest account."

Paying debt first then maximize returns:

> "I think I would still pay off all debt first and then put all the rest of my points into the 20 percent allocation."
>
> "Get rid of the high interest debt and stick any gains in high interest savings."

"I probably could've completely ignored the low percentage debt, eliminated the high percentage debt, then just put everything into the 20% return account"

Other strategies:

"I would invest in the account that had the least amount of debt, but gained the most interest."

"Make sure no accounts are falling into debt, and make sure there is some even spreading of the money even if it means less interest"

"I ensure 40% of earns on paying debts "

F. Tables

Subjects were only randomized into treatments after completing the initial survey and responding to the follow-up email. At the end of the last day, subjects were asked a series of demographic questions on sex, race, and education. We also surveyed them on their past experiences with debt as well as their experiences during the Covid-19 pandemic. Below we present balance tables comparing means for these elicited characteristics across the *No Debt*, *Low Debt*, and *High Debt* treatments. To create comparable categories, we collapsed the race question into White and non-White, and we also collapsed the education question into college-plus and non-college-plus. We also compare subjects initial risk and time preferences from the initial survey before they were sorted into treatment groups. The final total duration across all days is shown as well, for both the mean and median.

Table A.11: Main Treatments: balance table

	No Debt	Low Debt	High Debt	
	Mean	Mean	Mean	P.value
Age	37.56	38.58	37.87	0.82
Male	0.60	0.56	0.57	0.82
White	0.73	0.85	0.73	0.11
College	0.73	0.67	0.72	0.68
Hold Student Loan	0.54	0.51	0.41	0.27
Hold Debt	0.66	0.61	0.60	0.67
Impact Covid	2.86	2.90	2.86	0.97
Initial Risk	0.62	0.64	0.63	0.87
Initial Time	0.85	0.88	0.89	0.52
Duration (hours)	2.31	2.17	2.13	0.79
Median Duration (hours)	1.68	1.70	1.60	0.38
Observations	86	86	86	

Notes: This table shows results from a balance test between our main treatments. We report the p.values of an F-test of equivalence of the three treatment means.

We also present the balance table for the two *Redistribution* and two *Borrowing* treatments. We compare the same elicited characteristics as from the three *Main* treatments.

Table A.12: Redistribution Treatments: balance table

	Redistribution No Debt	Redistribution Debt	
	Mean	Mean	P.value
Age	39.48	36.03	0.05
Male	0.54	0.60	0.49
White	0.77	0.75	0.86
College	0.80	0.77	0.58
Hold Student Loan	0.45	0.43	0.80
Hold Debt	0.61	0.74	0.09
Covid	3.05	3.17	0.51
Initial Risk	0.64	0.68	0.36
Initial Time	0.85	0.84	0.87
Duration (hours)	1.38	1.41	0.87
Median Duration (hours)	1.18	1.08	0.63
Observations	81	77	

Notes: This table shows results from a balance test between our redistribution treatments. We report the t-test p.values of equivalence of the two means.

Table A.13: Borrowing Treatments: balance table

	Borrowing No Debt	Borrowing Debt	
	Mean	Mean	P.value
Age	37.10	35.83	0.44
Male	0.70	0.56	0.07
White	0.77	0.74	0.65
College	0.68	0.79	0.09
Hold Student Loan	0.49	0.52	0.70
Hold Debt	0.66	0.67	0.82
Covid	3.04	3.10	0.70
Initial Risk	0.61	0.66	0.13
Initial Time	0.85	0.84	0.78
Duration (hours)	2.91	2.55	0.49
Median Duration (hours)	1.56	1.65	0.53
Observations	80	82	

Notes: This table shows results from a balance test between our borrowing treatments. We report the t-test p.values of equivalence of the two means.

Using the first allocation decision in our *Main* Treatments, we present the results from a regression on the share that subjects allocate to each account. In both *Low Debt* and *High Debt* subjects allocate less points to Savings 1 since they allocate a larger share to Debt 1.

Table A.14: Main Treatments: estimation output using initial allocation

	(1) Allocation Share Savings 1	(2) Allocation Share Savings 2	(3) Allocation Share Savings 3/Debt 1	(4) Allocation Share Savings 4/Debt 2
Mean of dep.var	0.483***	0.160*	0.268**	0.268**
	(0.102)	(0.0769)	(0.100)	(0.100)
Low Debt	-0.259***	0.0553**	0.192***	0.0115
	(0.0457)	(0.0174)	(0.0399)	(0.0187)
High Debt	-0.233***	0.0130	0.224***	-0.00324
	(0.0511)	(0.0180)	(0.0465)	(0.0218)
Errors Instructions	-0.0101***	0.00806***	-0.00264	0.00472***
	(0.00183)	(0.000876)	(0.00160)	(0.000758)
Above Median Age	0.0313	-0.0167	-0.0290	0.0143
	(0.0447)	(0.0166)	(0.0424)	(0.0159)
Male	0.0451	-0.00328	-0.0232	-0.0186
	(0.0427)	(0.0143)	(0.0406)	(0.0169)
White	0.104*	-0.0309	-0.0119	-0.0616*
	(0.0509)	(0.0198)	(0.0467)	(0.0259)
College Education	0.0706	-0.00560	-0.0648	-0.000168
	(0.0505)	(0.0169)	(0.0478)	(0.0154)
Hold Student Loan	0.0140	-0.0268	0.0186	-0.00584
	(0.0485)	(0.0172)	(0.0467)	(0.0166)
Hold Debt	-0.0413	0.00975	0.0228	0.00874
	(0.0488)	(0.0159)	(0.0467)	(0.0143)
Covid - Little Impact	0.156	-0.112	-0.0221	-0.0212
	(0.0825)	(0.0664)	(0.0873)	(0.0225)
Covid - Moderate	0.100	-0.114	0.00526	0.00838
	(0.0836)	(0.0657)	(0.0868)	(0.0226)
Covid - A lot	0.0748	-0.113	0.0202	0.0178
	(0.0928)	(0.0672)	(0.0928)	(0.0291)
Covid - Great	0.236*	-0.120	-0.0753	-0.0401
	(0.0950)	(0.0671)	(0.0955)	(0.0277)
Batch FE	Y	Y	Y	Y
Observations	258	258	258	258

Notes: Results from a linear regression with robust standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is the share of the initial endowment of 500 points that subjects allocate to each account.

Next, we present the results from a regression on the total amount of points that subjects allocate to each account by the end of the experiment. In both *Low Debt* and *High Debt* subjects allocate less points to Savings 1, although the difference is marginally significant. For *High Debt*, subjects allocate 616 more points to the Debt 1 account. In *Low Debt*, we also find a significant increase in the amount of points that subjects allocate to the Debt 2 account. This effect is mainly driven by subjects who repay Debt 2 after fully repaying Debt 1, which is an infeasible strategy in *High Debt*.

Table A.15: Main Treatments: estimation output using total allocation

	(1) Total Allocation Savings 1	(2) Total Allocation Savings 2	(3) Allocation Share Savings 3/Debt 1	(4) Allocation Share Savings 4/Debt 2
Mean of dep.var	3573.6***	642.4	479.9	213.1
	(509.3)	(344.7)	(254.3)	(135.9)
Low Debt	-417.6*	76.02	175.6*	241.8**
	(185.0)	(63.39)	(74.17)	(81.65)
High Debt	-451.4*	-34.97	616.0***	-26.85
	(214.4)	(71.79)	(128.5)	(68.48)
Errors Instructions	-82.11***	32.57***	11.51**	24.15***
	(9.089)	(3.509)	(4.007)	(3.752)
Above Median Age	80.01	-67.67	133.2	-13.55
	(174.8)	(59.85)	(95.12)	(68.29)
Male	137.7	-32.94	-38.43	-30.60
	(166.2)	(50.25)	(91.15)	(64.72)
White	482.6*	-172.3*	-127.8	-213.7*
	(202.0)	(69.64)	(111.3)	(91.48)
College Education	-0.0267	-41.46	-82.63	29.99
	(187.3)	(63.04)	(108.8)	(76.27)
Hold Student Loan	234.5	-7.336	-84.34	-136.7
	(182.6)	(62.63)	(108.2)	(78.73)
Hold Debt	-172.1	-90.02	176.2	107.9
	(186.0)	(57.94)	(103.9)	(74.73)
Covid - Little Impact	120.8	-366.4	7.356	92.57
	(429.4)	(295.2)	(212.3)	(73.19)
Covid - Moderate	-91.34	-247.9	23.99	135.3
	(433.6)	(292.5)	(212.5)	(77.53)
Covid - A lot	27.65	-281.5	-39.83	106.5
	(457.2)	(298.1)	(223.8)	(93.06)
Covid - Great	124.9	-255.8	-117.9	-55.60
	(452.8)	(296.7)	(236.7)	(102.1)
Batch FE	Y	Y	Y	Y
Observations	258	258	258	258

Notes: Results from a linear regression with robust standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is the total amount of points that subjects allocate to each account by the end of the experiment.

Table A.16: Main Treatments: estimation output using risk and time elicitation questions

	(1) Risk 1 Points vs Points	(2) Risk 2 Money vs Points	(3) Time 1 Points vs Points	(4) Time 2 Money vs Points
Mean of dep.var	0.240***	0.202***	0.249***	0.194***
	(0.058)	(0.066)	(0.063)	(0.063)
Low Debt	0.049**	-0.018	0.035	-0.005
	(0.025)	(0.030)	(0.025)	(0.028)
High Debt	0.070***	0.015	0.055**	-0.006
	(0.024)	(0.028)	(0.027)	(0.026)
Initial Risk	0.182***	0.200***	0.085*	0.059
	(0.056)	(0.060)	(0.050)	(0.058)
Initial Time	0.080	0.219***	0.262***	0.308***
	(0.053)	(0.065)	(0.045)	(0.065)
Max Returns	0.111***	0.027	-0.019	-0.079***
	(0.024)	(0.028)	(0.027)	(0.023)
Risk First	0.009	-0.010	0.001	0.029*
	(0.011)	(0.018)	(0.015)	(0.016)
Errors BDM Instructions	-0.021	0.001	0.033**	0.063***
	(0.014)	(0.018)	(0.016)	(0.017)
Above Median Age	0.005	0.006	-0.008	0.016
	(0.021)	(0.024)	(0.022)	(0.025)
Male	-0.011	-0.031	-0.018	-0.017
	(0.021)	(0.023)	(0.022)	(0.023)
White	-0.007	-0.007	-0.012	-0.043
	(0.026)	(0.028)	(0.028)	(0.027)
College Education	0.036	0.036	-0.007	0.004
	(0.023)	(0.027)	(0.024)	(0.026)
Hold Student Loan	0.011	-0.003	0.015	0.013
	(0.023)	(0.026)	(0.026)	(0.027)
Hold Debt	-0.012	-0.034	-0.021	-0.047*
	(0.022)	(0.026)	(0.026)	(0.027)
Covid - Little Impact	0.008	-0.000	-0.104*	-0.064
	(0.048)	(0.059)	(0.055)	(0.055)
Covid - Moderate	-0.020	-0.022	-0.067	-0.030
	(0.047)	(0.058)	(0.056)	(0.055)
Covid - A lot	-0.020	-0.024	-0.107*	-0.082
	(0.053)	(0.063)	(0.060)	(0.058)
Covid - Great	0.005	-0.004	-0.094	-0.058
	(0.061)	(0.065)	(0.066)	(0.067)
Day FE	Y	Y	Y	Y
Observations	1032	774	774	774

Notes: Results from a linear regression with clustered standard errors at the individual level in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is an index of risk and time preferences that ranges from 0 to 1. Higher numbers imply higher risk-seeking and time-discounting behavior.

Table A.17: Borrowing Treatments: returns and payments estimation output

	(1) Log Total Returns	(2) Log Total Returns	(3) Log Total Returns	(4) Final Payment	(5) Final Payment	(6) Final Payment
Sample	All Subjects	Max Returns	All Subjects	All Subjects	Max Returns	All Subjects
Mean of dep.var.	8.578***	8.847***	8.439***	23.131***	25.217***	21.793***
	(0.090)	(0.092)	(0.059)	(0.896)	(1.667)	(0.639)
Borrow Debt	-0.090***	-0.094**	-0.008	-0.860**	-1.279*	-0.069
	(0.031)	(0.043)	(0.027)	(0.355)	(0.667)	(0.335)
Borrow Max			0.184***			1.771***
			(0.027)			(0.358)
Errorts Instructions	-0.011***	0.001	-0.007***	-0.082***	0.111	-0.044**
	(0.002)	(0.008)	(0.002)	(0.021)	(0.113)	(0.020)
Above Median Age	-0.006	-0.013	0.002	-0.335	-0.782	-0.266
	(0.028)	(0.039)	(0.024)	(0.351)	(0.588)	(0.314)
Male	0.025	-0.114***	0.039	0.253	-1.115*	0.392
	(0.032)	(0.041)	(0.025)	(0.367)	(0.652)	(0.316)
White	0.013	-0.050	0.018	0.409	0.383	0.454
	(0.042)	(0.048)	(0.032)	(0.422)	(0.819)	(0.353)
College Education	-0.040	-0.045	-0.050*	-0.306	-0.513	-0.407
	(0.031)	(0.042)	(0.028)	(0.408)	(0.662)	(0.387)
Hold Student Loan	0.023	0.009	0.051*	0.555	0.284	0.816**
	(0.035)	(0.047)	(0.027)	(0.419)	(0.714)	(0.361)
Hold Debt	-0.001	-0.008	-0.015	0.020	-0.235	-0.112
	(0.032)	(0.040)	(0.027)	(0.440)	(0.675)	(0.401)
Covid - Little Impact	-0.046	-0.053	-0.084**	-0.434	0.251	-0.796*
	(0.075)	(0.066)	(0.041)	(0.733)	(1.133)	(0.444)
Covid - Moderate	-0.032	-0.090	-0.067*	-0.455	-0.477	-0.787**
	(0.075)	(0.059)	(0.038)	(0.714)	(0.884)	(0.395)
Covid - A lot	-0.013	-0.062	-0.051	-0.268	-0.312	-0.641
	(0.080)	(0.059)	(0.045)	(0.821)	(0.980)	(0.569)
Covid - Great	-0.083	-0.140	-0.093	-0.949	-1.855	-1.037*
	(0.096)	(0.089)	(0.065)	(0.810)	(1.237)	(0.527)
Observations	117	47	117	117	47	117

Notes: Results from a linear regression with robust standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is the log of the cumulative returns in the four allocation decisions and the final payments that subjects obtained without the participation fee. Borrow Max is a dummy that equals 1 if the subject borrowed the maximum amount by the end of the experiment. Columns 2 and 6 focus on subjects who maximize returns in all allocation decisions, regardless of their borrowing behavior.