Title
Essays on Bounded Rationality in Repayments and Learning

Author
Ozyilmaz, Hakan

Publication Date
2022

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Essays on Bounded Rationality in Repayments and Learning

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Economics

by

Hakan Özyılmaz

Committee in charge:

Professor Erik Eyster, Chair
Professor Gary Charness
Professor Ryan Oprea
Professor Emanuel Vespa
Professor Sevgi Yuksel

December 2022

The Dissertation of Hakan Özyılmaz is approved.

Professor Gary Charness

Professor Ryan Oprea

Professor Emanuel Vespa

Professor Sevgi Yuksel

Professor Erik Eyster, Committee Chair

September 2022

Essays on Bounded Rationality in Repayments and Learning

For my boon;
without him I would be a fellow with no collective significance

Acknowledgements

While the completion of my dissertation required hard work and dedication on my part, it would not have been possible without the help of many others. I am heavily indebted to my dissertation committee - Erik Eyster, Emanuel Vespa, Gary Charness, Ryan Oprea, and Sevgi Yuksel- for their guidance and encouragement over the years. I want to thank Erik for expanding the limits of what I believe I can do and providing me with constant support. I want to thank Manu for always grounding my thinking, showing me the way forward, and encouraging me to take action. I want to thank Gary for teaching me the art of communication and negotiation. I want to thank Ryan and Sevgi for introducing me to new possibilities whenever I felt stuck. I also would like to thank Hani Mansour for guiding me into the next chapter of my career.

I spent most of my time at UCSB alongside my co-author and friend Guangli Zhang. It was an absolute joy to learn and grow together as researchers. I am very thankful to my office mates Jeffrey Cross and Yongwook Kim, whom I learned so much from over the years. I am also extremely grateful to my friends who gave me the strength to keep moving: Maria Kogelnik, Baris Ince, Elif Kubilay, Alec McQuilkin, Charlie Nusbaum, Alina Harper, Matthew Wayne Gonzales, and Lucas Reddinger.

My family was the most significant source of support on this journey. I would feel completely lost without the support of my parents Ümmühan Özyılmaz, Patrice Ryan and Stephen Randolph, my siblings Sema Özyılmaz, Furkan Özyılmaz, Megan Cameron, and my aunt Nagihan Atasoy.

Finally, I would like to thank my father, Erkan Özyılmaz, whom I did not get to share much time with. I always felt proud to be your son, and I hope I made you proud.

Curriculum Vitæ

Hakan Özyılmaz

Education

2022	Ph.D. in Economics (Expected), University of California, Santa Barbara.
2016	M.A. in Economics, University of California, Santa Barbara.
2015	M.A. in Economics, Koc University, Istanbul.
2013	B.Sc. in Mathematics, Koc University, Istanbul.
2012	Exchange Student, Sciences Po, Paris.

Research Fields

Behavioral and Experimental Economics, Household Finance, Microeconomic Theory

Working Papers

The Debt Payment Puzzle: An Experimental Investigation (with Guangli Zhang)

Restoring Rational Choice in Repayments: Disclosures or Advice? (with Guangli Zhang)

Mental Models and Endogenous Learning

Work in Progress

Two-Sided Financial Technology Underadoption (with Sabrine Bair)

Teaching Experience

Instructor 2018 - 2019

University of California, Santa Barbara, CA

- Math Camp (PhD) 2019
- Math Camp (PhD) 2018

Teaching Assistant 2016 - 2021

University of California, Santa Barbara, CA

- History of Economic Thought, Intermediate Microeconomics 2021
- Behavioral Finance, Game Theory, 20th Century US Economic History 2020
- Game Theory, Principles of Microeconomics 2019
- Intermediate Microeconomics I and II, Intermediate Macroeconomics Principles of Microeconomics 2018

- Intermediate Microeconomics II, Principles of Microeconomics Principles of Macroeconomics — 2017
- Principle of Microeconomics, Game Theory — 2016

Teaching Assistant — 2013 - 2015
Koç University

- Forecasting — 2015
- Econometrics I, Game Theory — 2014
- Game Theory — 2013

Awards and Scholarship

Outstanding Undergraduate TA Award, UCSB	2020
Humanities and Social Sciences Research Grant ($2,640), UCSB	2020
Departmental Research Grant ($2,770), UCSB	2018
Departmental Fellowship, UCSB	2015 - 2016
Academic Scholarship for M.A. Economics, Koç University	2013 - 2015
Vehbi Koç Scholar, Koç University	2009 - 2013
Science Fellowship for B.Sc. Mathematics, TUBITAK	2008 - 2013
Academic Scholarship for B.Sc. Mathematics, Koç University	2008 - 2013
Global Exchange Scholarship, Koç University	2012

Permissions and Attributions

The content of Chapters 1 and 2 their appendices is the result of a collaboration with Guangli Zhang.

Abstract

Essays on Bounded Rationality in Repayments and Learning

by

Hakan Özyılmaz

This dissertation uses controlled experiments to understand why people fall short of making rational decisions in simple financial decision-making situations, how we can restore rationality through simple interventions, and how providing people with substantial opportunities to learn about their environment might lead them *away* from making rational decisions.

In the first chapter, joint with Guangli Zhang, we study the sources of suboptimal allocations observed in credit card repayments using a diagnostic laboratory experiment. We find that optimization ability and limited attention are jointly insufficient to explain the puzzle. Moving beyond existing results, we find that the inherent negative frame of the debt payment problem interferes with subjects' ability to optimize and hinders learning. We show that subjects predominantly rely on the irrelevant balance information while forming their decisions, regardless of how vividly the balance information is displayed. Using additional treatments, we find that the debt frame increases subjects' focus on the irrelevant balance information.

In the second chapter, joint with Guangli Zhang, we study what type of interventions would be effective in eliminating simple arbitrage failures in repayments. We construct a simple repayment environment in the laboratory and test the role of a set of behavioral mechanisms that would directly inform the design of consumer protection policies. We find that providing salient interest rate disclosure has no effect while disclosing the interest rate in a fee format has modest effects. On the other hand, providing an opportunity

to purchase automated financial advice reveals that subjects are predominantly aware of their choice inefficiencies and are relatively good at gauging the extent of their mistakes and using financial advice. Our results suggest that promoting and subsidizing consumer financial technology applications that provide automated financial advice would be a substantially more effective way of protecting consumers from simple arbitrage failures than conventional disclosure policies.

In the final chapter, I study how people learn about their environment when their subjective understanding of the environment, their mental model, is misspecified. I use people's tendency to hold optimistic beliefs about their abilities to generate a significant amount of model misspecification and investigate the implications of overconfidence as a misspecified mental model on learning about own ability and a fundamental. Consistent with the theoretical predictions, overconfident subjects develop pessimistic beliefs about the fundamental and take growingly suboptimal actions. Inconsistent with the theoretical predictions, I find endogenous feedback does not exacerbate the extent of suboptimal behavior. Investigating how subjects learn about their own ability reveals that abundant feedback "weakens" misspecified mental models. The "weakening" of mental models is more pronounced with endogenous feedback and explains why endogenous feedback may not exacerbate the extent of suboptimal behavior.

Contents

Chapter 1

The Debt Payment Puzzle: An Experimental Investigation

1.1 Introduction

Borrowing households frequently make decisions that appear inconsistent with models of rational choice. Recent examples include insufficient search effort while choosing a mortgage contract, failure to refinance a mortgage contract when market conditions improve, and borrowing on a higher interest rate credit card while there is available credit limit on a lower interest rate credit card ((1), (2), (3)). Understanding the sources of suboptimal borrowing behavior is fundamental to developing informed consumer financial protection policies and improving the descriptive success of boundedly rational models of decision making.

In this paper, we use a diagnostic laboratory experiment to study how people make financial decisions when the decision involves a debt frame. Specifically, we investigate *the debt payment puzzle* where people pay down debt on a lower interest rate credit card while forgoing the opportunity to pay down debt on a higher interest rate credit card.[1] A distinct advantage of *the debt payment problem* over other "problematic" debt settings is that the optimal payment rule is unambiguously determined without any assumption on time and risk preferences.

Two recent studies, (3) and (4), show that the average credit card holder misallocates 50% of her payment to the card with lower interest rate and leaves a significant amount of money on the table annually.[2] Moreover, both studies show that suboptimal repayments cannot be rationalized with various plausible explanations that can be tested with observational data.[3] Despite the strength and persistence of the evidence on suboptimal repayments, it is still an open question why consumers behave inconsistently with the presumption of welfare maximization.

This paper studies the potential sources of suboptimal credit card repayments. Specifically, we design a diagnostic laboratory experiment that aims to answer what features of the debt payment problem make it hard for consumers to solve correctly. There are a number of potential explanations for this puzzling behavior. Two immediate explanations are financial literacy and limited attention. Researchers in household finance have

[1]Consider a cardholder with revolving debt on two credit cards who cannot afford to pay off both cards at the end of the month. The uniquely optimal rule would prescribe one pays the card with the higher interest rate while making the minimum required payment on each card.

[2]This type of allocation decision is common and costly. 1) The revolving credit card debt reached $1.3-trillion in the US in the last quarter of 2019, constituting almost 6% of the US GDP (NY Fed, Consumer Credit Panel). 2) 61% of the Americans have at least one credit card and the average card holder has four credit cards (according to the credit reporting agency Experian's nationally representative data, 2019). 3) (4) calculate that 71.5% of credit card holders in the U.S. market have two or more cards, and this group accounts for 91.8% of balances. Moreover, (4) find the average annual cost of misallocation to be $85 for individuals who hold two cards and $325 for individuals who hold five cards. The authors further document that the degree of misallocation does not decline in stakes: the cost of misallocation at the 90th percentile rises from $218 in the two-card sample to $1,213 in the five-card sample.

[3](3) document that the following explanations are at best able to account for small variations: 1) Differences in due dates 2) Differences in the ease of payment 3) Differences in unobserved characteristics 4) Strategic manipulation of interest rates and credit limits. (4) show that the following explanations do not account for the observed behavior: 1) Consumers face a fixed cost of optimization due to time, psychological or cognitive costs. 2) Consumers learn over time to make correct payments but the cross-sectional data masks this learning behavior.

long emphasized the role of financial literacy ((5), (6)). It is plausible that consumers who self-select into having revolving credit card debt are not sufficiently financially literate to optimally manage their repayments given the plethora of evidence linking financial literacy and suboptimal household behavior ((7), (8)). The behavioral economics literature has emphasized the role of limited attention in consumer choice ((9), (10), (11), (12)). In the context of credit card repayments, consumers might not know their interest rates or even if they do, they might not remember what the rates are at the time of decision making. A common feature of these explanations is that their identification often requires more detailed information of consumers and their choice processes than what is available in a typical administrative data set. However, developing informed consumer financial protection policies and improving the descriptive success of boundedly rational models of decision making crucially depend on identifying mechanisms that underlie such puzzling repayment behavior.[4][5] A controlled laboratory environment allows us to circumvent the identification challenges faced by observational studies, and to study how consumers make their allocations and how the quality of their decisions are affected by their choice environment.

We begin our investigation by establishing suboptimal allocation behavior in an extremely simple decision environment where potential confounds that exist in the field are minimized. Moreover, we show that suboptimization is not specific to people who lack the skills to solve an optimization problem or the knowledge of their interest rates at the time of decision making. We show that the share of optimal allocations in our baseline treatment - where the decision environment captures the essential features of a

[4](13) is an excellent reference on why people might not use readily available information to make better decisions and the importance of mechanisms for developing descriptive theories of decision making.

[5]In particular, if consumers struggle with their repayments due to their inability to solve simple optimization problems, this would necessitate promoting financial literacy education. On the other hand, if consumers' struggles are related to a lack of attention to their interest rates, this would make the case for information disclosure policies. Indeed, the current policy debates regarding consumer protection revolve around financial literacy education and information disclosures.

typical online payment screen - is only 18.8% despite the fact that 82% of our subjects can solve simple optimization problems and 93% of our subjects actively seek interest rate information before making their decisions.[6] Our findings clearly indicate that even the combination of optimization ability and the knowledge of interest rates is insufficient to explain this puzzle. We further show that subjects do not learn to make better decisions nor do they respond to higher incentives, corroborating the findings of (3) and (4). Finally, we show that allocation behavior causally moves with balance information. Specifically, subjects allocate higher amounts to an account with higher balances without regard to interest rate information - a finding that is consistent with the balance matching heuristic documented in (4).

The fact that we are able to replicate the field findings in a tightly controlled environment with an algebraically sophisticated subject pool deepens this puzzle and urges us to investigate mechanisms that underlie this suboptimal behavior. Although our baseline findings suggest that people pay attention to interest rate information, psychology experiments suggest that this might not be sufficient to make optimal allocations as choices are influenced by *salience* of information; that is if one part of the environment attracts more attention, then the information contained in that part is reflected more in the choices.

We move beyond existing findings by examining the role of information salience. Specifically, we examine two potential channels that could affect the salience of interest rate information: *information vividness* and *framing*. The reason that we focus on channels that revolve around salience is that it is an established cognitive mechanism that guides choice behavior in various contexts ((14), (15)). Its applications in behavioral economics have been particularly fruitful in capturing deviations from rational choice in simple environments ((16), (17)).

[6] (3) find the share of optimal allocations to be approximately 15% among people who hold two comparable credit cards using observational data. (4) find this rate to be 11.8% .

A critical aspect of the credit card repayment environment is the predominant display of balance information. A typical credit card statement or an online account displays balance information more vividly than any other information. The vivid display of balance information might increase the salience of balance information, leading consumers to form their allocation decisions by relying on irrelevant balance information. This would indeed justify the suboptimality of allocations as irrelevant balance information is incorporated into the decision process.[7] Interestingly, our result suggests that subjects' allocation decisions are not affected by the vividness of balance information. Compared to our baseline treatment with vividly displayed balance information, maximizing the vividness of interest rate information surprisingly has a null effect on the share of optimal allocations.

Another way the salience mechanism might operate in the credit card repayment environment is through the framing of the allocation problem. The credit card payment environment is inherently a negative situation. Specifically, the balance information indicates how much a person owes on an account – an amount that affects the welfare of the decision maker negatively. Psychologists document that such inherent negativity of a piece of information changes the amount of attention that information attracts ((18), (19), (20)). If balance information attracts more differential attention due to its inherent negativity, this creates another channel for the salience mechanism to interfere with the decision process and lead to suboptimal allocations. We confirm this hypothesis and find that the inherent debt frame of the problem interferes with subjects' decisions. Compared to a subject who faces this allocation problem under an otherwise identical debt frame, a subject who faces the investment frame has a 24.2 percentage point higher probability of making an optimal allocation -this is equivalent to a 128% increase in the share of optimal allocations.

[7] *Irrelevant* in the sense that objectively optimal allocation does not depend on balances.

To further investigate why we observe such an asymmetry in the share of optimal allocations across frames, we conduct two additional treatments. Our results hint at two explanations that are not necessarily mutually exclusive: asymmetric attention and asymmetric heuristic use. First, we document an asymmetry in measured attention across two frames. We show that an average subject spends significantly more time on balance information compared to interest rate information under the debt frame; under the investment frame, there is no difference in time spent on the interest rate and balance information. Second, we document an asymmetry in heuristic use across frames. Under the debt frame, we find subjects' allocations are mostly consistent with a balance matching heuristic i.e. they seem to make their allocations roughly proportional to their balances. Under the investment frame, a majority of the subjects' allocations are consistent with an *interest matching heuristic* i.e. they seem to make their allocations roughly proportional to interest rates.

We contribute to the growing body of evidence showing that people seem to struggle with correctly resolving simple trade-offs with financial frames ((3), (4)). It is hard to establish that deviations from the rational benchmark are *mistakes* using observational data since we do not know the exact trade-off people face in the field. They must solve a dynamic allocation problem with varying income streams, due dates, card limits, cash rewards, and alike where their attention to this allocation problem is limited. A critical point here is that consumers with multiple accounts might not even be aware of the fact that they face a simple trade-off regarding their repayments. Using the power of a controlled environment where such concerns are brought to a minimum, we show that people indeed struggle with simple trade-offs with financial frames as severely and persistently in the field. This finding has a broader implication on the case for consumer protection as people seem to suffer pecuniary losses by deviating from normative prescriptions given their preferences.

We also contribute to the policy discussion regarding how to improve consumer financial decisions using empirically informed interventions ((21)). Our results have implications on the performance of two popular policy alternatives: mandating disclosure policies and promoting financial education.[8] A common finding in previous studies that investigate financial behavior in the debt domain is that conventional disclosure policies are ineffective in improving financial outcomes ((22), (23)). We find evidence aligning with previous findings. We show that vividly disclosing interest rate information has no significant effect on the share of optimal allocations compared to our baseline treatment where interest rate information is disclosed *non-vividly*. This does not mean to say that every potential disclosure policy will fall short of restoring rational choice. We think that non-conventional disclosures of interest rate information might prove useful in improving the quality of decisions in this repayment context.

A popular policy alternative to information disclosure policies is financial education. Financial literacy surveys indicate that many households struggle with algebraic calculations related to interest rates ((24), (5)). While confirming that optimization ability is associated with improved decision making, we find a significant majority of subjects capable of solving simple optimization problems fail to make their allocations optimally during the experiment. Our finding suggests that an effective financial education program should acknowledge the mental gaps between real-life financial decision problems and algebraic counterparts, and focus on training people how to translate these problems into simple optimization problems.

Our final contribution is to the vast framing literature in behavioral economics. We show that many subjects have a harder time making optimal allocations under a debt frame despite exhibiting similar optimization abilities on the algebraic version of the prob-

[8]Figuring whether to implement information disclosure policies or to bolster financial education programs is particularly important as neither of them comes without a trade-off. See (7) for a discussion of these trade-offs.

lem. Our further investigation into the asymmetry in the share of optimal allocations across frames hints at systematic differences in how attention is allocated under different frames. The asymmetric attention allocation pattern that we observe is inconsistent with optimal allocation of attention ((25)), models of salience ((16)), focusing ((17)) and selective attention ((26)). This suggests that exploring how frames affect attention allocation might be worthwhile. We also document how different frames may trigger different heuristics. Although the use of heuristics in financial decision making has long been documented ((27), (4)), we present systematic evidence on how an algebraically identical allocation problem under different frames induces different distributions of heuristic use over subjects.

1.2 Evidence for Suboptimal Repayments

The purpose of the baseline experiment is two-fold. First, it helps us documenting the severity and persistence of suboptimal repayments even in extremely simple environments, corroborating the field findings. Second, it documents that the combination of limited attention and optimization ability is not sufficient to explain this puzzling behavior.

1.2.1 Baseline Design

Decision Environment

Our experiment interface captures the essential features of the decision environment faced by credit card consumers who make their repayments in the field (See Figure 2.1). Each subject is endowed with two hypothetical credit card accounts and a hypothetical checking account. The experiment consists of multiple periods. At the beginning of each period, we deposit a fixed amount of 500 *Experimental Currency Unit*s (ECU) into their checking account. Subjects' task in each period is to make repayments toward their credit cards using their deposit. During a period, subjects face a screen that is split into two halves. Each half represents a credit card account. At the top part of each half of the screen, subjects see the current balance information. At the center of the screen, subjects see a list of other account attributes that are typically displayed on a credit card statement. These attributes are interest rate, interest charged, previous balance and previous repayment. The information on each of these attributes is presented simultaneously and singularly to a subject once she clicks on the *information button* that carries the name of that attribute.[9] Clicking on information buttons is costless and subjects are allowed to click freely. Each period ends once a subject submits an allocation decision.

It is important to emphasize that subjects **always** see how much they owe on an account at the top part of the screen and they *need not* click any button to acquire balance information while they *need* to click the information buttons to see other attributes. We describe the information that is always displayed at the top part of the screen and that does not require the click of subjects as *vividly displayed* - an important point that we will revisit in Section 1.3. Hence in the baseline design current balance information is vividly displayed.

Our interface allows us to sidestep many confounding features of the actual decision environment and focus on the allocation problem that lies at the core of this repayment situation. An essential feature of our design is that interest rate information is readily available at the time of decision making at a cost as low as clicking a button. Indeed in

[9] For instance, a subject who wants to find out the interest rate information on both accounts needs to click the *Interest Rate* button. Once she clicks the interest rate button, she sees the interest rate information on both cards at the same time and does not see any other information until she clicks on some other information button.

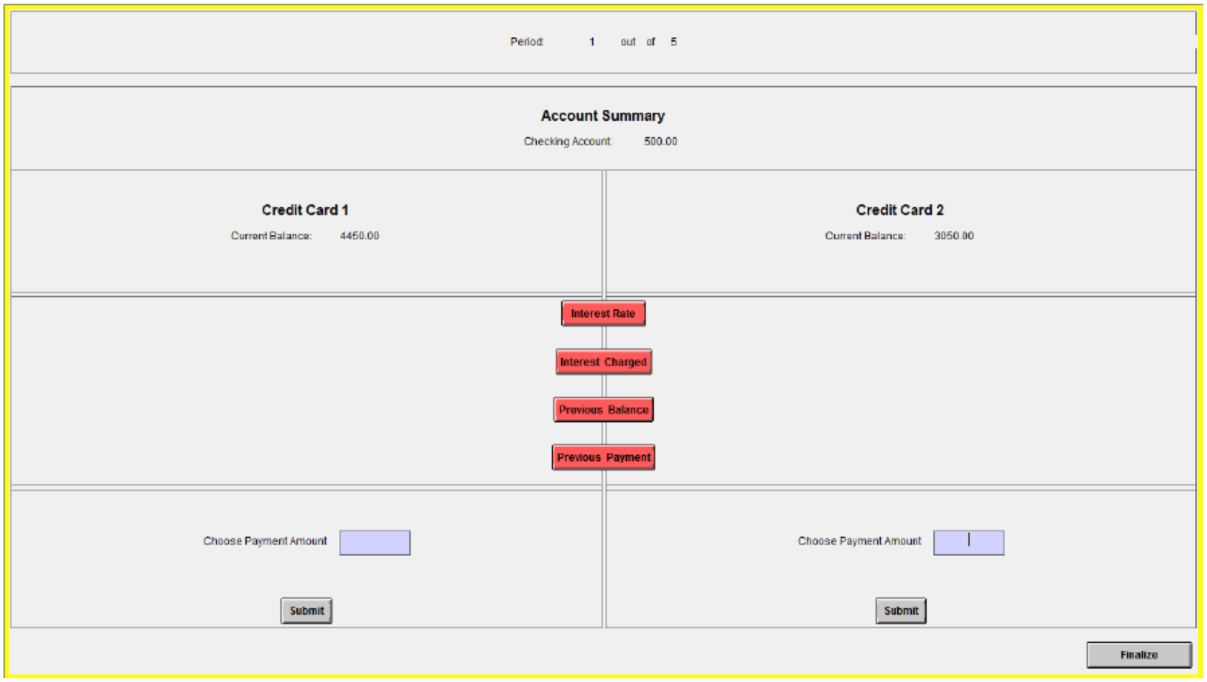

Figure 1.1: Experiment Interface

all of our sessions, an overwhelming majority of subjects clicks the interest rate button and acquires their interest rate information.[10] Other important simplifications we make include no minimum required repayment, simultaneity of repayments and no previous purchase decision.[11][12]

A crucial aspect of this repayment problem in the field is that consumers do not get feedback on the quality of their decisions. The only feedback consumers get is the amount of interest charged on each account which is then incorporated in the total debt they owe to each card in the subsequent period. We recreate this implicit feedback mechanism in the laboratory by employing a block design where we combine decision periods into

[10]Knowledge of interest rate information at the time of repayment is a significant source of variation in the actual decision environment as the interest rate information is complexly disclosed.

[11]See the online appendix of (3) for a larger set of potential confounds that exist in the actual credit card repayment environment.

[12]Empirical studies ((28), (29)) have documented robust findings on how minimum required payments could create anchoring on the required amount. Our experiment eliminates the use of minimum payment in order to remove any potential anchoring that is induced from making the minimum payment.

stages. Each stage consists of five decision periods.[13] In the first period of each stage, we determine the amount of debt on each card. In the subsequent periods, each subject's debt on each card is endogenously determined by their previous allocation decisions in that stage. Since subjects are assigned some debt at the beginning of each stage, we endow subjects with a fixed positive amount in order for each subject to make some money in the experiment. We determine a subject's payoff for a stage by their end of stage balance on each card subtracted by the fixed endowment. We then convert their stage payoffs into US dollars and randomly choose one of their stage payoffs for their actual payment.

We employ six stages with different balance and interest rate configurations. The first four stages of the experiment have the same structure, and together they constitute the first part of the experiment. The parameter choices for the first period of these stages are presented in Table 2.1. We choose the interest rate difference to be 1.5% as a plausible upper bound of the the observed monthly interest rate differences in the field.[14] We keep the interest rate difference across stages fixed to keep the incentives the same across these stages. We choose the initial balances to be consistent with the average credit card debt observed in the field and keep the balance difference around 1,500 ECU in order to separate potential balance-matching behavior from naively allocating equal amounts to each account ($1/N$ heuristic).[15] To provide causal evidence for the impact of higher interest rate and higher balances on allocation decisions, we design our stages so

[13] We choose five periods per stage to have a sense of subjects' within stage learning and to keep the duration of the experiment reasonable.

[14] (4) document that the observed annual interest rate difference is 15% at the 90th percentile corresponding to a monthly interest rate difference of 1.25%. (3) find the average monthly interest rate gap to be 1.1% in their data. Update: The levels of interest rates are chosen to be consistent with the APR in Gathergood paper: Average APR is around 20%. For a 5-period stage, an average "APR" is then equivalent to 4% per period interest rate. The average "APR" in our experiment for a stage is 19.5% for the low interest rate card.

[15] According to Experian's 2019 data, the average American owes $6,200 on their credit cards and 80% of credit card holders owe less than $10,000.

that each credit card account carries observations under each potential balance/interest rate configuration. The shaded stages in Table 2.1 represent *aligned stages*: a higher interest rate account is also assigned a higher initial balance. In contrast, non-shaded stages represent *misaligned stages*: a higher interest rate account is assigned a lower initial balance.

Table 1.1: Parameter Choices and Balance Reallocation

Stage	Account	Interest Rate (per period)	Initial Balance	Balance Reallocation
1	1	4.9%	4,450	No
	2	3.4%	3,050	
2	3	5.7%	2,950	No
	4	4.2%	4,350	
3	5	3.7%	4,550	No
	6	5.2%	2,950	
4	7	3.9%	2,850	No
	8	5.4%	4,450	
5	9	5.3%	4,650	Yes
	10	3.8%	3,150	
6	11	5.9%	3,050	Yes
	12	4.4%	4,550	

In the second part of the experiment, subjects face the remaining stages, namely 5 and 6. These stages differ from the first four stages in one important way - there is an additional period at the end of each stage.[16] In the last period of stage 5 and 6, subjects are asked to reallocate their balances between the two accounts. This intervention

[16]See Figure A7.1 for a screenshot of these periods.

tightens the screws on the potential suboptimal repayment behavior as it simplifies the allocation problem even further and increases the incentives to optimize.[17]

Timeline

Upon arrival, each subject is provided with instructions where the rules of the experiment and how their payment is determined are clearly explained.[18] After the experimenter goes through the instructions, the experiment starts with an explanation phase where subjects are familiarized with the interface. When the explanation phase ends, subjects move on the first part of the experiment. The first part of the experiment contains four stages. Subjects are provided ten minutes for the first two stages and seven minutes for the subsequent stages. Subjects are advanced to the next stage if they complete a stage or if they exceed the maximum allotted time.[19]

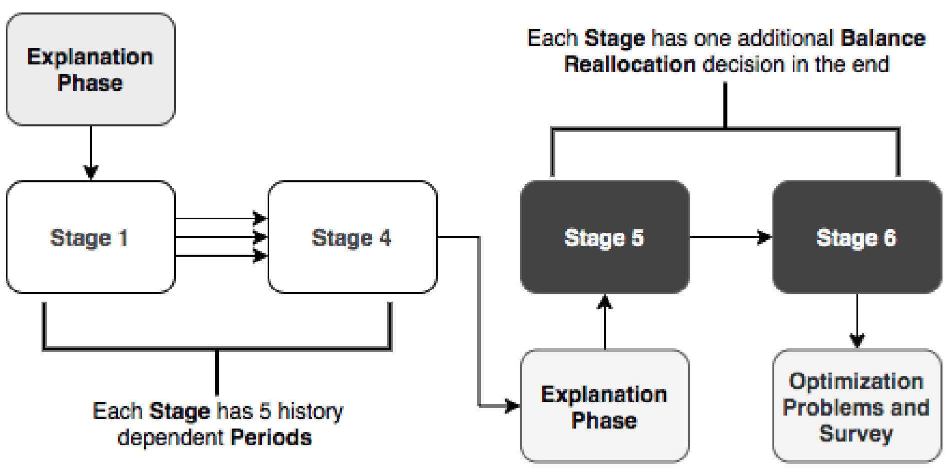

Figure 1.2: Experiment Timeline

[17]Given these parameter choices, the payoff difference for a subject who allocates all her deposit into the high interest rate account throughout a stage makes $5 more than a subject who allocates all her deposit into the lower interest rate account throughout a stage. In the last two stages, we increase this payoff difference to $12 by introducing the balance reallocation period.

[18]Experiment Instructions are located in Appendix A.7.

[19]Only 2 out of 44 subjects used up the maximum time in a given stage. We discard these auto-advanced periods in our analysis.

Upon completing the first part of the experiment, subjects are provided with instructions on balance reallocation. After the experimenter goes through the balance reallocation instructions, subjects face an explanation phase where they learn how to reallocate their balances using the interface. Once the explanation phase is over, subjects go through Stages 5 and 6. Subjects are provided ten minutes for each stage in this part of the experiment.

Once the main parts of the experiment ends, subjects are asked four incentivized optimization problems represented in algebraic expressions. These problems correspond to algebraic versions of the allocation problems subjects go through in the main part of the experiment.[20] We use subjects' scores on these problems as a proxy for their optimization ability. An important design choice here is that we do not ask optimization problems at the beginning of the experiment as it might affect subjects' ability to optimize in the experiment. The experiment ends with subjects answering exiting survey questions that record basic demographic information and subjects' justification for their allocation behavior.

Procedural Information

We conducted our experiment at the UCSB Experimental and Behavioral Economics Laboratory. The experiment was coded using z-Tree software ((30)). A total of 44 subjects, recruited through ORSEE (Online Recruitment System For Economic Experiments), participated in the baseline experiment . The average payment per subject was $13.2 including a $5 show-up fee. The average duration of a session was 75 minutes.

[20]The four optimization problems that we ask the participants are: i) $\min_{x,y} 3(1000 - x) + 2(2000 - y)$ ii) $\max_{x,y} 3(1000 + x) + 2(2000 + y)$ iii) $\min_{x,y} -3x - 2y$ iv) $\max_{x,y} 3x + 2y$ all subject to $x + y = 300, \ x, y \geq 0$

1.2.2 Baseline Results

Do subjects know their interest rates?

An important question that arises from previous studies is "Do people actually know their interest rates? And if they do, do they recall the interest rate information at the time of decision making?" Since we track the information buttons that a subject clicks, we can answer this question with our baseline treatment. Figure 1.3 shows the proportion of subjects acquiring the interest rate information by the first period of each stage.[21] In the first period of the first stage, 100% of the subjects click the interest rate button to acquire the interest rate information. Although this proportion decreases in later stages, on average 93.2% of the first period decisions are made after acquiring the interest rate information. Moreover, we find that the average response time for the first period decisions is 38.7 seconds and 11.3 of these seconds are spent on the interest rate information. In light of these findings, we conclude that an overwhelming majority of our subjects know their interest rates at the time of decision making.

Can subjects solve optimization problems?

Another potential explanation for suboptimal repayments is that people are not good at solving optimization problems. In order to see if inability to solve optimization problems drives this mistake, we ask subjects four incentivized optimization problems after the main experiment. We find that 82% of our subjects are able to solve at least one of the four simple optimization problems. Hence we conclude that a significant majority of our subjects can solve simple optimization problems.

[21]Recall that the interest rate on each card is fixed within a stage.

Figure 1.3: Proportion of Subjects Acquiring Interest Rate Information

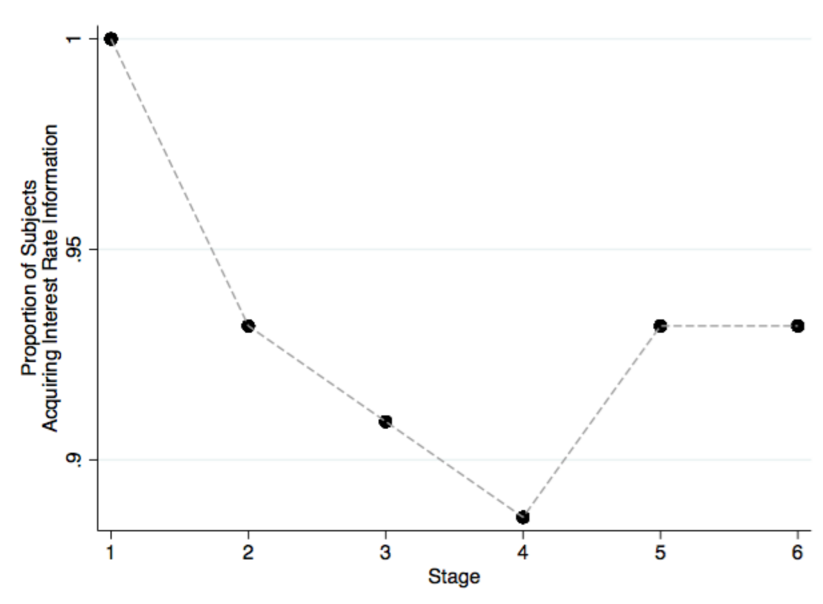

Note: Figure shows the proportion of subjects acquiring interest rate information by the *first period* of each stage.

How do subjects make their payments?

Now that we know most of our subjects do look at the interest rate at the time of decision making, and they can deal with simple optimization problems, we turn to the main analysis of our baseline treatment. For the remainder of this chapter, we restrict the sample to the first period decisions while excluding observations from subjects who do not acquire interest rate information or fail to answer any optimization question correctly. Most of our results are qualitatively similar when we extend our analyses to include all observations. We indicate and discuss when our results depend on the sample restrictions.

Result 0 *Suboptimal allocations persist when the potential confounds that exist in the field are removed, knowledge of interest rates and optimization ability are ensured.*

Theoretically, subjects should allocate 100% of their assigned deposit to the card with the higher interest rate. However, as illustrated by Figure 1.5, only 22.4% of the

Figure 1.4: Distribution of Subjects' Optimization Abilities

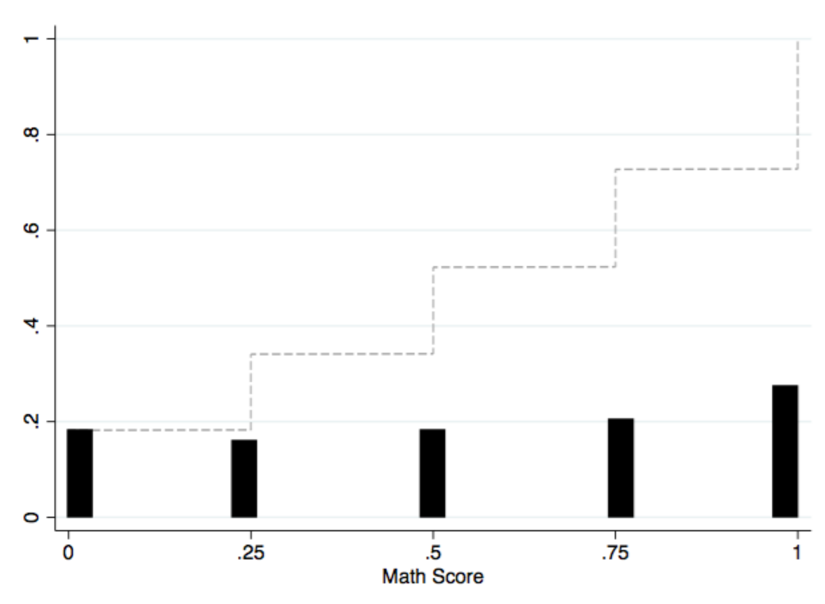

Note: Figure shows the distribution of subjects' optimization abilities. *Math Score* represents the fraction of correctly answered optimization problems. Each bar represents the fraction of subjects achieving a certain score. The dotted line represents the empirical cumulative distribution function of math scores.

repayments are allocated toward the card with the higher interest rate. The distribution of optimal repayments is significantly different than the observed repayments (clustered Wilcoxon signed-rank test, $p < 0.001$). The optimality rate decreases to 18.8% when we do not impose any sample restriction. Our results corroborate the field findings: despite the simplifications we make in the decision making environment, subjects seem to make similar levels of optimal allocations compare to the field findings. (3) find the share of optimal allocations to be approximately 15% among people who hold two comparable credit cards. (4) find this rate to be 11.8%.

Since optimality seems to be a stringent test on how well subjects make their payments, we also report the fraction of misallocated repayments - the fraction of repayment that is incorrectly allocated to the lower interest card. We find that 33.5% of the repay-

Figure 1.5: Distribution of Allocations - Period 1 Decisions

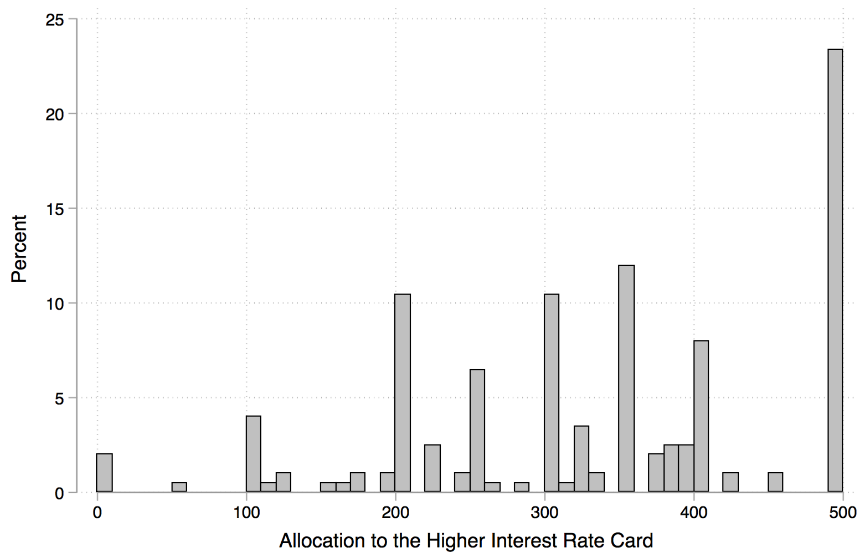

Note: Figure shows the distribution of payments subjects make toward the high interest rate card in the first period of each stage. The sample excludes (1) those who fail to correctly answer at least one out of four math questions and (2) those who do not acquire interest rate information. This eliminates 63 out of 264 observations at the *subject* $\times$ *stage* level. The histogram contains 50 equally sized bins. The rational choice theory predicts a distribution with full mass located at 500.

ments is misallocated.[22] (3) report that consumers misallocate 50% of their repayments to the low interest rate card and (4) report a misallocation level of 48.5%.[23] The difference in the misallocation rate between our experiment and the field studies, combined with the similarity in the share of optimal allocations, suggest that our participants deviate less from the rational benchmark given that there is a deviation. Nevertheless, our participants' allocation behavior is still far from the rational benchmark despite the fact that they actively seek interest rate information and they can solve simple optimization problems.

To get a sense of how subjects make their repayments, we first show the distribution

[22]The misallocation rate is 36.3% when we do not impose any sample restriction.

[23]These numbers are the amount of misallocation in excess of the minimum required payments for consumers who hold two credit cards.

Figure 1.6: Allocation Patterns Across Stages - Period 1 Decisions

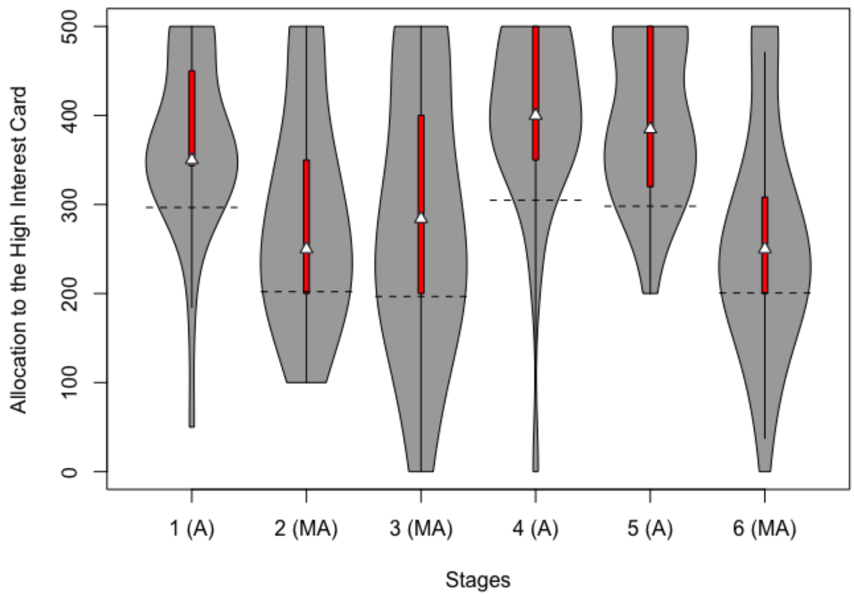

Note: The violin plot shows the distribution of repayments subjects make toward high interest rate card in the first period of each stage. The center white dot represents the median allocation towards the higher interest rate card in a given stage. The thick bars around the median represents allocations within the interquartile range. The end of the whisker represents the maximum and the minimum allocation. The violin shape visualizes the kernel density distribution of the allocation patterns - the wider sections of the violin represents a higher likelihood of allocating in the corresponding value. The letters A and MA next to stage numbers represent if that stage is aligned or misaligned. The dotted horizontal reference lines represent the hypothetical allocation under an exact balance matching heuristic towards the higher interest card in the first period of each stage. The rational choice theory predicts a distribution with full mass located at 500 for all stages.

of allocations made to the high interest rate card by stage. Figure 1.6 provides some suggestive evidence on subjects' tendency to allocate more towards the card with higher balances. In *aligned stages* where the high interest rate card comes with higher initial balances (Stages 1, 4 and 5), the median allocation is well above 250 ECU (more than half of their assigned deposit). We find that 94% of the subjects allocate more than 250 ECU to the high interest rate card indicating that an overwhelming majority of the subjects are at least partially responsive to interest rates.[24] However, this interpretation overstates the extent that subjects' decisions are influenced by the high interest rate as the effect of high interest rate on the allocations made is confounded with the effect of high balances. In order to discuss the impact of high interest rate separate from the impact of high balances, we present our findings from the *misaligned stages* where the high interest rate card comes with lower initial balances (Stages 2, 3 and 6). We find in each of the *misaligned stages*, the median allocation is 250 ECU which is virtually indistinguishable from a baseline where subjects are completely unresponsive to interest rates.[25] Taken together, we interpret our findings from aligned and misaligned stages as subjects being responsive to the irrelevant balance information as well as the relevant interest rate information. In particular, subjects' allocations seem to move away from the high interest rate card when it comes with lower initial balances.[26]

We solidify this interpretation by quantifying the effect of having a higher interest rate on a card (and a higher balance) on the allocation made towards that card. We are able to provide causal evidence on these effects using a simple linear regression on our subjects' first period decisions in each stage since we exogenously and independently assign the interest rates and debt levels to be high or low on a single card. We choose, without loss of generality, the left card on our subjects' screens for our analysis. We call the left card "treated" with a higher interest rate if the assigned interest rate on the left card is greater than the assigned interest rate on the right card, and we denote this "treatment" with the dummy variable *Higher Interest Rate*. Similarly, we call the left card treated with a higher balance if the assigned current balance on the left card is

[24]The proportion of subjects who allocate at least 250 ECU to the high interest rate account in each aligned stage is exactly 94%.

[25]The proportion of subjects who allocate at least 250 ECU to the high interest rate account in Stages 2,3 and 6 is respectively 50%, 52% and 50%.

[26]The results are nearly identical when we do not impose any sample restriction. The proportion of subjects who allocate at least 250 ECU to the high interest rate account in each stage is respectively 93%, 50%, 50%, 88%, 88% and 50%.

Table 1.2: OLS Estimation of Repayments

	(1) Left Card Allocation	(2) Left Card Allocation
Higher Interest Rate	164.0	184.5
	(25.80)	(31.79)
Higher Balance	109.7	80.83
	(16.69)	(16.89)
Constant	117.2	111.4
	(14.03)	(16.45)
Observations	201	645
R^2	0.423	0.406
Period	First	All

Note: Column 1 represents a model of repayments made in the first period of each stage. The dependent variable is the amount of allocation made on the left card which takes a value in between 0 and 500. The regressor *Higher Interest Rate* is a dummy variable that takes the value 1 when interest rate on the left card is higher compared to the right card. The regressor *Higher Balance* is another dummy variable that takes the value 1 when balance on the left card is higher compared to the right card. The rational choice theory requires that *Higher Interest Rate* to perfectly predict all allocation behavior and give no predictive power to *Higher Balance*. Column 2 extends the analysis by including repayments for all periods. Standard errors in parentheses. Errors are clustered at the subject level.

greater than the assigned current balance on the right card and we denote this treatment with the dummy variable *Higher Balance*.[27]

A rational decision maker's allocation behavior should solely be guided by the interest rate information, giving no predictive power to the normatively irrelevant balance information. Table 1.2 provides the regression results. In Column 1, we see that subjects take both the relevant interest rate information and the irrelevant balance information into account while determining their allocations. On average, subjects allocate 164 ECU

[27]One caveat here is that whenever the left card has a higher balance, it also has a higher interest charge and a higher previous balance by design. In other words, higher current balance perfectly correlates with higher interest charges and higher previous balances. Hence the "treatment" *Higher Balance* captures an aggregate effect of all normatively irrelevant information presented to the subjects.

more to the card with a higher interest rate and 109.7 ECU more to the card with a higher balance. These effects are significant ($p = 0.0000$ for both) and statistically equal in magnitude ($p = 0.13$). These results suggest that subjects are indeed responsive to a higher interest rate although the effect's magnitude is less than the prescription of rational choice. However, we see that subjects are similarly responsive to the irrelevant balance information, which indicates that the deviations from the rational choice are not random errors but systematic mistakes that are governed by the irrelevant balance information. In Column 2, we extend the analysis to all periods. Although this analysis loses the causal interpretation, we see that both higher interest rates and higher balance information predict allocation behavior in all periods significantly ($p = 0.0000$ for both) yet the effect of higher interest rate is greater in magnitude ($p = 0.03$).

These results corroborate the field findings that people take irrelevant balance information into account while making their payments. (4) find, using various machine learning algorithms, that balance information has the highest variable importance, which is 3 to 40 times larger than the variable importance of interest rates in predicting allocation behavior. (3), using regression analysis, find that fraction of outstanding balances on a card explains almost 10 times greater variation in the allocation behavior than the variation explained by the interest rate difference. Although our findings are consistent with the field results, we find no difference in the predictive power of higher balances and higher interest rates on allocation behavior. We see this improvement in the predictive power of interest rates relative to the field findings as a manifestation of subjects' increased appreciation toward the importance of interest rates due to the simplifications we make in the decision environment and our subject pool's relatively higher algebraic sophistication.

Do subjects learn to make better decisions?

Subjects are not provided any feedback between periods or stages. In addition, there is no explicit intervention in the first part of the experiment that would potentially induce them to change their allocation decisions. The only source of learning in the first part of the experiment is repetition which is similar to how such decisions are made in the field. However, once subjects complete the first part of the experiment, we inform them that the remaining stages have a balance reallocation period, which might induce subjects to re-evaluate their decision making strategies.

Figure 1.7: Measures of Optimality Within and Between Stages

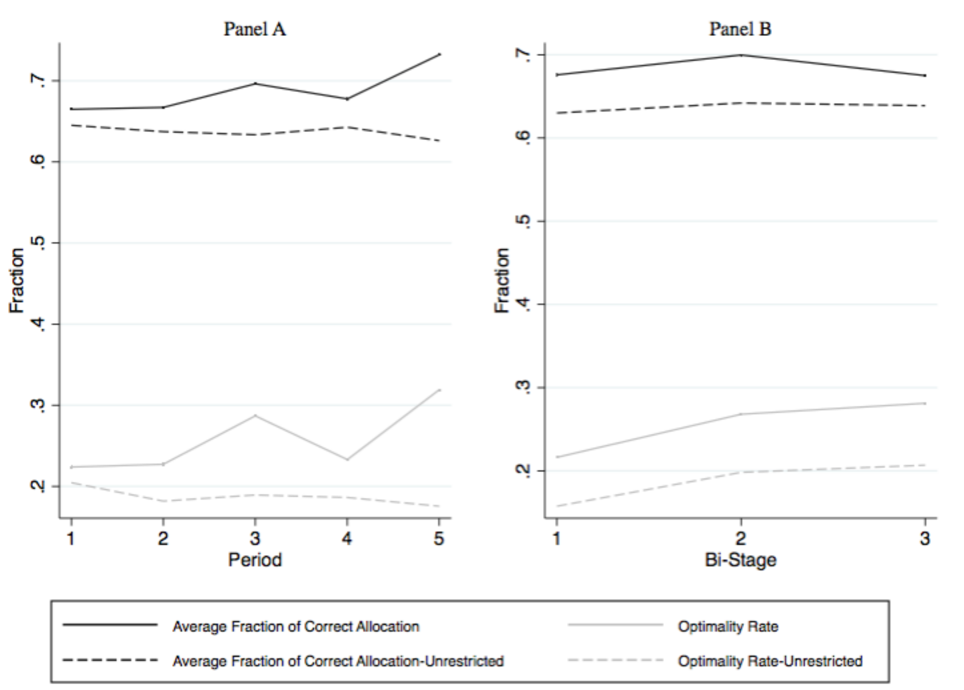

Note: Panel A shows both the average fraction of correctly made allocations and the share of optimal allocations by periods within a stage. Panel B shows the same optimality measures by bi-stages. A bi-stage consists of two consecutive stages with one aligned and one misaligned stage. The solid lines indicate the optimality measures 1) for allocations made after acquiring interest rate information, 2) for the subjects who solve at least one optimization question correctly. The dashed lines indicate the optimality measures without imposing any sample restriction.

We find that subjects do not learn to make better decisions within a stage or between stages. Figure 1.7 shows the average fraction of correctly made allocations and the share of optimal allocations within and between stages.[28] Although subjects' average fraction of correctly made allocation increases from 66% to 73% within a stage corresponding to a 1.4% per period increase, this effect is insignificant ($p = 0.068$). Similarly, the share of optimal allocations increase from 22.4% to 31.9% within a stage corresponding to a 1.9% per period increase yet the effect is insignificant ($p = 0.18$). Moreover, we do not find any significant evidence that subjects' allocations improve between bi-stages ($p = 0.12$ for the share of optimal allocations, $p = 0.96$ for the average fraction of correctly made allocations).[29]

The results are consistent with previous findings and serve as direct evidence regarding the difficulty of learning to avoid interest charges in the context of debt payment even for people who pay attention to interest rates and who are equipped with sufficient optimization ability.[30]

Do subjects respond to higher incentives?

An important class of economic models explain the deviations from rational choice by arguing cost-benefit considerations of making an optimal decision ((31), (25)). In particular, if our subjects face a fixed cost of optimization due to time, psychological or cognitive costs of making an optimal payment, the reduction in interest charges due to optimization may not be high enough to justify to incur this fixed cost. Therefore, one might expect an increase in the incentive to optimize would improve subjects' allo-

[28]The fraction of correctly made allocation refers to the fraction of the deposit that is assigned to the high interest rate card. For instance, the fraction of correct allocation for an allocation that assigns 400 ECU to the high interest rate card is 0.8.

[29]The results are qualitatively similar when we do not impose any sample restriction, the regressions can be found in Appendix A.3.

[30]Both (4) and (3) find that the fraction of correctly made allocations do not increase with the length of account tenure.

cation decisions. The balance reallocation periods in our design allows us to test this explanation as we effectively increase the incentives to optimize from \$1 per period to \$7 while simplifying the problem even further by directly asking subjects how much debt they would like to have on each card. As illustrated in Figure 1.8, the drastic increase in incentives to optimize do not lead to any improvement in the share of optimal allocations. In fact, the share of optimal balance reallocations is 16.7% - which is lower than the share of optimal allocations observed in the main part of the experiment. Our findings from balance reallocation is consistent with previous findings ((4), (3)), which have documented the degree of misallocation is virtually invariant to the economic stakes.

Figure 1.8: Distribution of Balance Reallocation Decisions

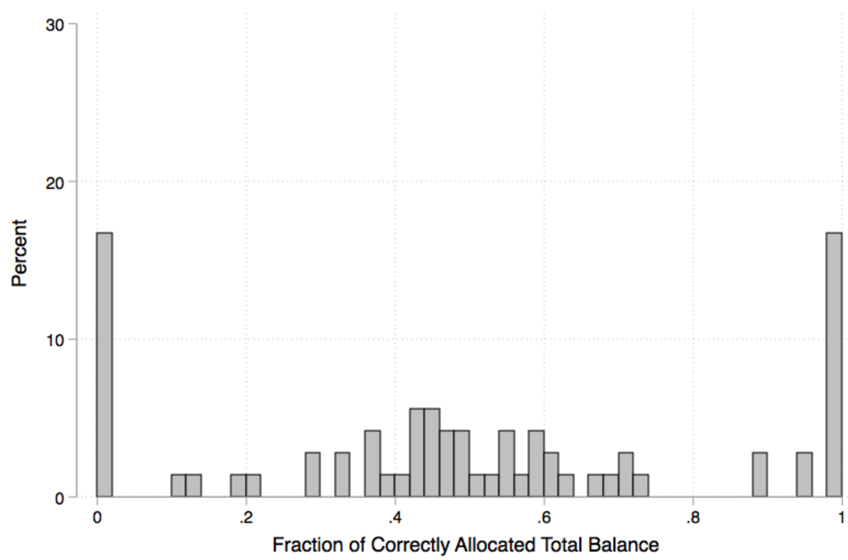

Note: Figure shows the distribution of fraction of total balances subjects reallocate toward the high interest rate card in balance reallocation periods. The distribution is represented with 50 equally sized bins. The sample is restricted to subjects who can solve at least one of the optimization problems. This restriction removes 8 out of 44 subjects, and leaves us 72 subject $\times$ period observations. The rational choice theory predicts a distribution with full mass located at 1.

1.3 Mechanisms

After establishing the suboptimality of allocation behavior and characterizing the suboptimal repayments as balance-dependent, we extend our baseline design to include further treatments with the goal of understanding what features of the decision environment leads to suboptimal repayments. Although the suboptimality of choices has no justification from the perspective of rational choice and hence standard economic theory, substantial research in psychology documents departures from normative models of decision making and investigate various mechanisms that could explain such departures.[31] Moreover, there has been significant advances in behavioral economics literature that incorporates these insights from psychology to develop descriptive theories of financial decision making ((16), (17), (25), (39), (13)).

In the context of credit card repayments, one way such suboptimization can arise is through the *vivid* display of balance information. A typical credit card statement or an online account displays balance information more vividly than any other information. Psychologists argue that vividly displayed information has more impact on judgments compared to other information ((14)) and they think such vividness effects to be generated through differential attention to one portion of the environment ((15)).[32] Comparing to interest rate information, the vividly displayed balance information might therefore attract greater attention and influence the subsequent decisions more heavily.

Another way such suboptimality can arise is through the debt frame of the decision problem. The credit card repayment problem has an intrinsic negative frame: it is an optimization problem over *balances that affect utility negatively*. A parsimonious

[31] These mechanisms include selective attention ((14), (32)), mental models ((33), (34)), dual process theories ((35), (36)) and heuristics ((37), (38)).

[32] We use the word attention to indicate *observable attention* which is simply the amount of time spent. Although how observable attention relates to attention is an open question, measuring observable attention is an established way of measuring attention. See (40) for a detailed discussion.

explanation for why the debt frame might yield balance-dependence is the *valence* of information. Psychologists define valence as the intrinsic attractiveness and aversiveness possessed by events, objects and situations ((41)).[33] Although the negative valence of balance information should play no role in the decisions made by consumers from the perspective of rational choice, there is substantial research in psychology that documents that negative information attracts greater attention and contributes more strongly to the observed choices ((18), (19), (20)).

In order to motivate our experimental design and show how our manipulations in the decision environment might lead to different payment behavior, we outline a simple framework in Appendix A.6 where we conceptualize a behavioral decision maker whose decisions are influenced by the salience of information that is presented to her. It is important to emphasize that we think of salience mechanism as a psychologically founded way of generating context-dependent choice behavior within optimizing agent paradigm that could unify our hypotheses, while acknowledging that there might be other mechanisms that could lead to differences in payment behavior across the decision environments we create in the laboratory.

In the next subsection, we describe our treatments that aim to change the salience of interest rate information.

1.3.1 Mechanism Treatments

We extend our baseline design to test if certain features of the decision environment plays a role in driving suboptimal allocations. In the extended design, we vary two main factors: the information that is vividly displayed and the frame of the decision problem.

[33](42) discusses how differences in valence of information can trigger different cognitive processes that lead to different decisions. The idea of valence-dependent encoding is far from being strange to the field of economics. (20) was a critique of expected utility theory that is based on framing of outcomes as gains and losses which lead to subsequent development of an immense literature on reference-dependent preferences and its applications.

Table 1.3 presents an overview of our treatments.[34] It is important to note that the *Debt Balance* treatment is exactly our baseline treatment. In treatment *Debt Interest Rate*, we decrease the vividness of balance information while increasing the vividness of interest rate information. We implement this manipulation by displaying the information that we call *vivid* at the top part of the experiment interface while keeping every other feature of the design unchanged. In treatment *Investment Balance*, we manipulate the frame of the allocation problem by reframing the credit card repayment problem as a mutual fund investment problem. The allocation problems that subjects face under each frame are algebraically identical and offer the same incentives to optimize. Similarly, the interface under both frames is identical in all respects except for the language that we use: treatments under the debt frame feature a *checking account* and two *credit cards*; treatments under the investment frame feature an *investment account* and two *mutual funds*.[35] In treatment *Investment Interest Rate*, we manipulate both the vividness of interest rate information and the frame of the allocation problem to capture any interaction between these two factors.

Table 1.3: Overview of Mechanism Treatments

Treatments	**Design Features**	**Sample Size**
Debt Balance [**DB**]	Debt Frame, Vivid Balance	44
Debt Interest Rate [**DR**]	Debt Frame, Vivid Interest Rate	43
Investment Balance [**IB**]	Investment Frame, Vivid Balance	38
Investment Interest Rate [**IR**]	Investment Frame, Vivid Interest Rate	40

Role of Information Vividness. If the vividness of information plays a role in

[34]See Appendix A.7 for the screenshots of the interface of these new treatments.

[35]Another semantic difference across frames is the substitution of the words *charged* and *earned*; and *payment* and *investment*.

driving the suboptimal repayments, a decrease in the vividness of balance information and an increase in vividness of interest rate information should increase the salience of interest rate information. The increase in salience of interest rate information increases the probability that a behavioral decision maker accounts for interest rate information and makes the objectively optimal allocation.

Prediction 1 *An increase in vividness of interest rate information increases the share of optimal allocations and the average allocation to the high interest rate account.*

Role of Framing. If the framing of the decision problem plays a role in driving the suboptimal repayments, a positive frame of the decision problem (and hence an increase in valence of balance information) should lead to less attention being allocated to balance information and increase the salience of interest rate information. The increase in salience of interest rate information increases the probability that a behavioral decision maker accounts for interest rate information and makes the objectively optimal allocation.

Prediction 2 *A positive framing of the decision problem increases the share of optimal allocations and the average allocation to the high interest rate account.*

1.3.2 Results from Mechanism Treatments

Role of Information Vividness

In Figure 1.9, Panel A shows the share of optimal repayments made across treatments and Panel B shows the average allocation made to the high interest rate card for subjects who can solve optimization problems and who acquire interest rate information before making their decision in the first period of each stage. We see that there is no significant increase, on average, in any of the optimality measures. The share of optimal allocations increases by 3.4 percentage points -from 22.4% in **DB** to 25.8% in **DR** ($p = 0.68$). The average allocation to the high interest rate account goes in the opposite direction of our prediction, and decreases by 13 ECU - from 332 ECU to 319 ECU (p = 0.46). The results are qualitatively similar when we relax our sample restrictions and control for demographic information (See Tables A1.1 and A1.2 in Appendix).

Figure 1.9: Optimality Measures Across Debt Treatments - Period 1 Decisions

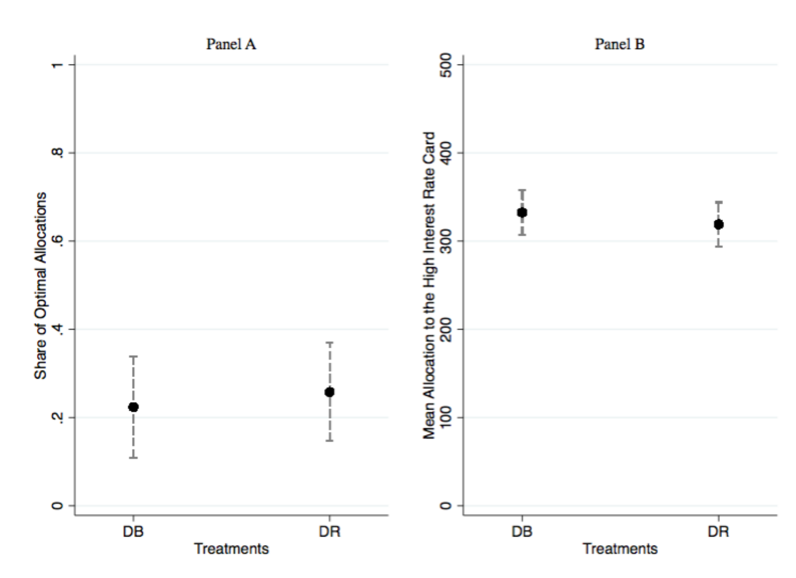

Note: Panel A shows the share of optimal allocations made under **DB** and **DR**. The whiskers indicate 95% confidence interval calculated using subject-level clusters. Panel B shows the average allocation made to the high interest rate card under **DB** and **DR**.

Figure 1.10: Allocation Patterns Across Debt Treatments - Period 1 Decisions

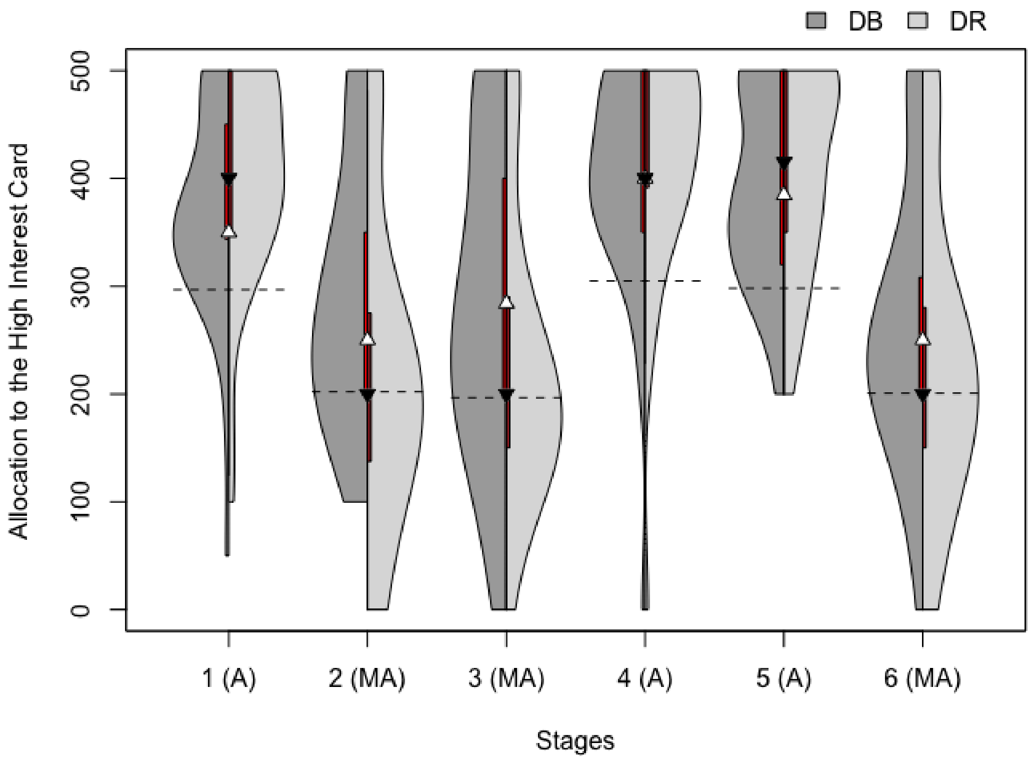

Note: The violin plots show the distribution of repayments subjects make toward the high interest rate card in the first period of each stage. The upward white triangle and the downward black triangle represent the median allocation towards the higher interest rate card in a given stage for **DB** and **DR**, respectively. The thick red and blue bars around the median represents allocations within the interquartile range for **DB** and **DR**, respectively. The violin shape visualizes the kernel density distribution of the allocation patterns - the wider sections of the violin represents a higher likelihood of allocating in the corresponding value. The letters A and MA next to stage numbers represent if that stage is aligned or misaligned. The dotted horizontal reference lines represent the hypothetical allocation under an exact balance matching heuristic towards the higher interest card in the first period of each stage. The rational choice theory predicts a distribution with full mass located at 500 for all stages.

Figure 1.10 documents further evidence that allows us to compare the allocation patterns across treatments. The patterns seem mostly similar. We find that in all aligned stages 94% of the subjects allocate more than half of their deposit into the high interest rate card which is identical to the same measure calculated in our baseline treatment. However, the percentage of subjects' that allocate more than half of their deposit into the high interest rate card in misaligned stages is respectively 26%, 29% and 36% which is lower than the same measure calculated in the baseline treatment. This finding is particularly striking given that subjects can achieve a higher payoff by simply uniformly randomizing their payments in misaligned stages. Taken together, these patterns suggest that subjects in **DR** are responsive to both interest rate and balance information, yet their decisions seem to be more responsive to balance information compared to the decisions of the subjects in our baseline treatment. Indeed, we surprisingly find that subjects are significantly more responsive to balance information in **DR** compared to **DB** ($p = 0.02$) whereas there is no difference in responsiveness to interest rate information across treatments ($p = 0.47$). Although subjects in **DR** are more responsive to balance information compared to the subjects in **DB**, they are not significantly more responsive to balance information compared to interest rate information ($p = 0.13$). These findings are robust to relaxing our sample restrictions and including demographic controls (See Tables A1.3 and A1.4 in Appendix).

Result 1 *Neither the share of optimal allocations nor the average allocation to the high interest rate account improves with an increase in the vividness of interest rate information.*

As a final note, we show that subjects in **DR** do not seem to learn to make better decisions within or between stages, similar to the subjects in **DB**. These results suggest that subjects in **DR** also struggle with learning how to make their allocations correctly.

Role of Framing

Figure 1.11: Comparison of Balance Treatments

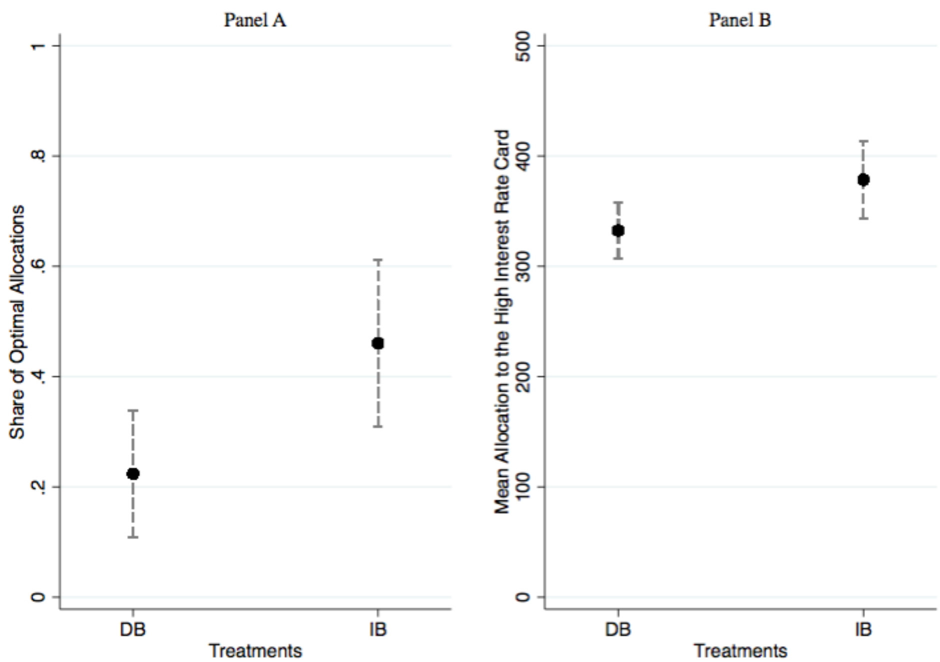

Note: Panel A shows the share of optimal allocations made in **DB** and **IB**. The whiskers indicate 95% confidence interval calculated using subject-level clusters. Panel B shows the average allocation made to the high interest rate card.

In Figure 1.11, Panel A shows the share of optimal repayments made across treatments and Panel B shows the average allocation made to the high interest rate card for subjects who can solve optimization problems and who acquire interest rate information before making their decision in the first period of each stage. We see that there is a significant increase, on average, in each optimality measure. The share of optimal allocations more than doubles -increases from 22.4% in **DB** to 46.1% in **IB** (p = 0.0166). The average allocation to the high interest rate account increases by 46.14 ECU - from 332.4 ECU to 378.54 ECU (p = 0.038). The results are qualitatively similar when we relax our sample restrictions and control for demographic information (See Tables A1.5 and A1.6

in Appendix).

Figure 1.12: Allocation Patterns Across Vivid Balance Treatments - Period 1 Decisions

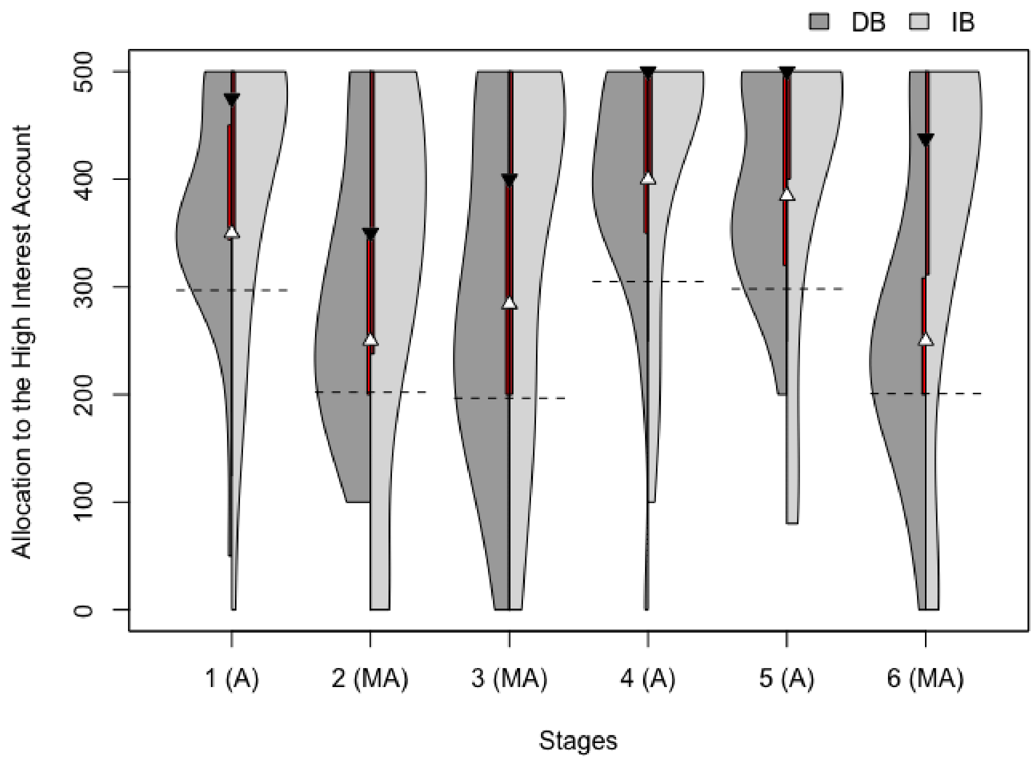

Note: The violin plots show the distribution of repayments subjects make toward the high interest rate card in the first period of each stage. The upward white triangle and the downward black triangle represent the median allocation towards the higher interest rate card in a given stage for **DB** and **IB**, respectively. The thick red and blue bars around the median represents allocations within the interquartile range for **DB** and **IB**, respectively. The violin shape visualizes the kernel density distribution of the allocation patterns - the wider sections of the violin represents a higher likelihood of allocating in the corresponding value. The letters A and MA next to stage numbers represent if that stage is aligned or misaligned. The dotted horizontal reference lines represent the hypothetical allocation under an exact balance matching heuristic towards the higher interest card in the first period of each stage. The rational choice theory predicts a distribution with full mass located at 500 for all stages.

Figure 1.12 documents further evidence that allows us to compare the allocation patterns across treatments. There are stark differences in the distribution of allocations made across treatments. We find that in all aligned stages 85% of the subjects allocate more than half of their deposit into the high interest rate account which is lower than the same measure calculated in our baseline treatment. However, the percentage of subjects that allocate more than half of their deposit into the high interest rate card in misaligned stages is respectively 71%, 68% and 81% which is significantly higher than the same measure calculated in the baseline treatment. The fact that the mass of allocations that are made in the correct direction is high and do not move much across aligned and misaligned stages suggest that subjects in **IB** are more responsive to interest information than balance information. We confirm this intuition statistically: we find that subjects in **IB** are more responsive to interest rate information compared to the balance information ($p = 0.01$). Moreover, we find that subjects in **IB** take interest rate information more into account while making their decisions compared to the subjects in **DB** ($p = 0.04$) and there is no difference in the extent that balance information is taken into account across **IB** and **DB** ($p = 0.21$). These findings are robust to relaxing our sample restrictions and including demographic controls (See Tables A1.7 and A1.8 in Appendix).

Result 2 *Subjects make significantly better allocations under the investment frame. There is a 23.7 percentage point increase - more than doubling - in the share of optimal allocations from **DB** to **IB**.*

Furthermore, we find that subjects in **IB** exhibit small yet significant learning which stands in contrast to the subjects' behavior in **DB**. This suggests that the debt frame of the problem do not only interfere with subjects' ability to optimize but also hinders learning.

1.3.3 Role of Vividness under the Investment Frame

We find that, similar to our finding under the debt frame, neither the share of optimal allocations nor the average allocation to the high interest rate account improves with an increase in the vividness of interest rate information across investment frames. The comparison between the treatments Investment Debt and Investment Interest Rate can be found in Appendix A.2.

1.3.4 Information Acquisition Patterns and Use of Allocation Heuristics

The results presented in this subsection have implication for models of bounded rationality. In particular, we present evidence towards two channels that pertain to models of attention and salience, and the literature on the use of heuristics. First, we find a sharp asymmetry in the way subjects acquire information across frames and we show how this asymmetric pattern correlates with allocation behavior. Second, we document an asymmetry in the response times and link this with the use of allocation heuristics across frames.

Information Acquisition Patterns

To understand the cognitive channels that lead to an asymmetric optimality rate across decision frames, we introduce two new treatments (*Debt No-Vivid* and *Investment No-Vivid*) where we do not display any information vividly, and thus require subjects to actively click on information buttons to reveal the corresponding piece of information before making their decisions. This representation-neutral information environment allows us to capture how subjects allocate their attention in a clear way. Specifically, we keep track of how many times a subject clicks on an information button, how much time they spend on each information button and in which order they decide to acquire information.[36]

Table 1.4: Overview of Information Acquisition Treatments

Treatments	**Design Features**	**Sample Size**
Debt No-Vivid [**DN**]	Debt Frame, No Vivid Attribute	15
Investment No-Vivid [**IN**]	Debt Frame, No Vivid Attribute	22

In Figure 1.13, Panel A shows the average click rates on current balance and interest rate buttons in each period by stage for the subjects in **DN**. We see that subjects consistently click more on the current balance button than interest rate button ($p = 0.0000$). Panel B documents the same measures for **IN**. In sharp contrast with the click patterns in **DN**, we find that subjects in **IN** click on the current balance and interest rate buttons at similar rates ($p = 0.44$). Using additional analysis, we find that a subject who is assigned to **IN** clicks, on average, 0.6 times less on the current balance compared to a subject who is assigned to **DN** ($p = 0.005$) while the click rates on interest rate information is similar across treatments ($p = 0.87$). See Table A4.1. When we analyze the time spent on each information button and the order in which subjects click on the information buttons, we find a similar *balance-focusedness* under the debt frame that does not exist under the investment frame.[37]

[36]See Figures A7.5 and A7.6 for the screenshots of the interface.

[37]When we compare the time spent on information buttons across treatments, we find that subjects in **DN** spend significantly more time on the current balance information compared to the interest rate information ($p = 0.0000$) while there is no such difference in the behavior of subjects in **IN** ($p = 0.07$). Moreover, we find that subjects in **IN** treatment spend significantly less time on the current balance information compared to the subjects in **DN** ($p = 0.001$) although there is no difference in the time spent on interest rate information across these two treatments ($p = 0.62$). See Table A4.2.

When we look at the click order, we see that the mode of first information button a subject clicks within a period is the current balance button if the subject is assigned to **DN** and interest rate button if the subject is assigned to **IN**. Figure **??** presents the click order data.

Figure 1.13: Average Click Rates Across No-Vivid Treatments

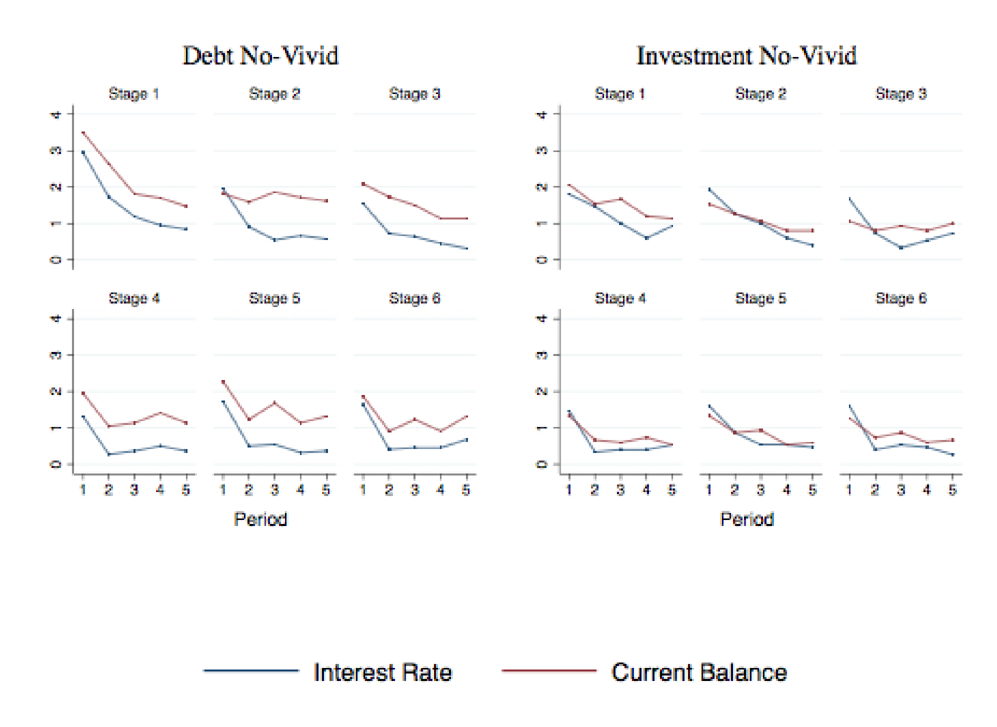

Note: Panel A documents the difference in average click rates on interest rate and current balance button for each period in each stage under DN treatment. Panel B presents the same measures for IN treatment.

Result 3 *Subjects pay significantly less attention to the irrelevant balance information under the investment frame. Compared to the debt frame, subjects click significantly less to the current balance button and spend significantly less time on the current balance button under the investment frame.*

We further show that clicking and spending more time on current balance information are tightly correlated with making lower quality decisions. See Appendix A.4.

Use of Allocation Heuristics

An alternative way balance-dependent allocations could occur is through the use of heuristics. In order to uncover potential regularities in allocation decisions, we investigate the following set of heuristics that we see as the most relevant:

1. **Optimal (OPT)**: Allocate optimally.[38]
2. **Balance Matching (BM)**: Allocate more into the account with higher balances.[39]
3. **Interest Matching (IM)**: Allocate more into the account with higher interest rates.[40]

Panel A of Table 1.5 shows the heuristic distribution across frames under a fairly strict classification requirement. According to this classification, a subject is classified as a certain heuristic type *i)* if her allocation is consistent with the same heuristic for at least 8 out of 10 periods in a given bi-stage, and *ii)* the assigned heuristic is a strictly better fit than any other heuristic. Using this approach we are able to classify around 60% of the subjects in each frame. The distribution of heuristic types is drastically different across the two frames. Under the debt frame, the number of subjects classified as the balance matching type is strictly greater than the number of subjects classified as the other two heuristic types. However, this is reversed under the investment frame: there is always a greater number of subjects who are classified as the interest matching or the optimal type compared to the number of subjects who are classified as the balance matching type. In Panel B of Table 1.5 we show the heuristic distribution under each frame when we weaken

[38] We allow for a 5% margin for error. Hence a subject is considered to be an *Optimal* type in a given period if she allocates at least 475 ECU to the high interest rate account in that period.

[39] Our definition of the balance matching heuristic is less strict than (4) although it still captures the same intuition that greater balances on an account lead to greater allocations on that account.

[40] Specifically, a subject who allocates between 250 ECU and 475 ECU into the higher interest account in a given period is considered to be an *Interest Matching* type for that period. Recall that we classify those who allocate at least 475 ECU to the high interest rate account as an *Optimal* type.

Table 1.5: Distribution of Heuristic Types Across Frames at Bi-Stage Level

Panel A: Strict Classification - 80% of Periods

	OPT	BM	IM	Other	Total
Debt (Bi-Stage 1)	6	**46**	13	44	109
Debt (Bi-Stage 2)	5	**49**	19	36	109
Debt (Bi-Stage 3)	9	**39**	18	43	109
Investment (Bi-Stage 1)	**19**	15	**25**	34	93
Investment (Bi-Stage 2)	**27**	11	**22**	33	93
Investment (Bi-Stage 3)	**24**	16	**28**	25	93

Panel B: Weak Classification - 60% of Periods

	OPT	BM	IM	Other	Total
Debt (Bi-Stage 1)	7	**61**	32	9	109
Debt (Bi-Stage 2)	5	**60**	33	11	109
Debt (Bi-Stage 3)	10	**55**	31	13	109
Investment (Bi-Stage 1)	**21**	20	**34**	18	93
Investment (Bi-Stage 2)	**27**	14	**33**	19	93
Investment (Bi-Stage 3)	**24**	17	**37**	15	93

Note: The table documents the number of subjects that are classified as a certain heuristic type under each frame at the bi-stage level. Panel A documents the distribution of heuristic types when the classification requires a subject to be consistent with a heuristic type for at least 8 out of 10 periods in a bi-stage. Panel B executes the same analysis by requiring a subject to be consistent with a heuristic type for at least 6 out of 10 periods in a bi-stage. Since there is no significant difference in the way that subjects make their allocations within the debt treatments and within the investment treatments, we conduct the heuristic analysis at the frame level by grouping subjects across the debt treatments **DB, DR, DN** and across the investment treatments **IB, IR, IN**.

the classification requirement.[41] This approach allows us to classify a significantly higher portion of the subjects and the results remain qualitatively similar.

Result 4 *A significant majority of the subjects are classified as the balance matching type under the debt frame. In contrast, the majority of the subjects are classified as either optimal or the interest matching type under the investment frame.*

[41]Now a subject is classified as a heuristic type *i)*when her allocation is consistent with that rule for at least 6 out of 10 periods in a given bi-stage *ii)* and the assigned rule is a strictly better fit than any other rule.

In addition to the asymmetry in the distribution of heuristic types across two frames, we find that subjects' assigned heuristic types to be persistent over time. In both debt and investment treatments, subjects whose allocations are consistent with the dominating heuristic in a given bi-stage (**BM** under the debt frame, and **IM** or **OPT** under the investment frame) are highly likely to be classified as the same heuristic type in the following bi-stage. We report the heuristic transition matrices in Appendix A.5.

Summary

To sum up this subsection, the asymmetry we document in information acquisition patterns is directly associated with the asymmetry in the share of optimal allocations and consistent with the distribution of heuristic types across frames. In particular, the tight connection between higher click rates/longer time spent on balance information and the share of optimal allocations is consistent with the salience mechanism. This suggests that frames can systematically affect decision makers' attention allocation and information processing while improving or worsening outcomes depending on the normative relevance of the information that the decision maker is drawn to.

1.4 Discussion

1.4.1 Policy Implications

Many researchers studying household finance have gathered an abundance of evidence toward departures from rational choice in the last three decades. These departures are not specific to one branch of financial decision making but cover every aspect of household finance. Credit card markets, being one of these domains, have offered various suboptimal consumer behavior and inefficient market outcomes ((7), (8)). The welfare consequences of such departures for the households have alerted policy makers to consider the tools available to them in order to restore the choices that consumers would make if they were rational and well informed.[42] Two widely discussed policies that aim to improve consumer financial decision making are mandating disclosure policies and promoting financial education.

A common finding in previous studies that investigate financial behavior in the debt domain is that conventional disclosure policies are ineffective in improving financial outcomes ((22), (23)). We find evidence aligning with previous findings. We show that *vividly* disclosing interest rate information has no significant effect on the misallocation rate compared to our baseline treatment where we *non-vividly* disclosure the interest rate information. We consider the quality of decisions in the vivid interest rate treatment (**DR**) to be an upper bound of the quality of decisions that can be obtained through conventional disclosure policies in the field. This is due to our removal of potential confounds that exist in the field and relatively high optimization ability of our subjects. This does not mean to say that any potential disclosure policy will fall short of restoring rational choice. We think that non-conventional disclosures of interest rate information might prove useful in improving the quality of decisions in this repayment context.[43]

A widely discussed alternative to information disclosure policies is financial education. According to recent financial literacy surveys, an important aspect of financial decision making that many households seem to struggle is the capacity to undertake algebraic calculations related to interest rates ((24), (5)). While confirming that optimization ability is associated with improved decision making, we find that a significant majority

[42] In the United States, the Truth in Lending Act of 1968 standardized the format of interest rate and other financial charge disclosures. The CARD Act of 2009 increased the amount of notice consumers receive in their credit terms. The Dodd-Frank Act of 2010 established the Consumer Financial Protection Bureau (CFPB) with the goal of protecting consumers from unfair, deceptive, or abusive practices of lenders.

[43] Both (22), (23) explore psychology-guided disclosures in similar borrowing situations and find them to have modest effects.

of subjects who are capable of solving simple optimization problems fail to make their allocations optimally during the experiment. We think the reason for this discrepancy is subjects' inability to translate the credit card repayment problem into a simple algebraic problem that they are clearly better at thinking through.[44] Our finding suggests that an effective financial education program should acknowledge the mental gaps between real-life financial decision problems and algebraic counterparts, and focus on training people how to translate these problems into simple optimization problems as well as solving algebraic problems.

A critical insight that arises from our findings is that people with similar levels of optimization ability struggle managing their allocations more as borrowers than investors. The welfare consequences of such mismanagement are particularly strong if we think of the allocation problems that we investigate as a simplified version of a larger allocation problem across various types of debt and investment accounts with differing interest rates. This insight has a direct implication on the evolution of wealth inequality. Households that have similar levels of optimization ability yet extensively borrow rather than invest will end up with lower overall wealth over their lifetime simply due to the greater mismanagement of their allocations that follows from the psychology of being in debt.[45] This is especially concerning for young adults as their mismanagements are amplified through compounding over their lifetime and they tend to be more on the borrowing than investment side. We believe that the incorporation of this mechanism into life-cycle models where people endogenously determine their level of financial education (an excellent example is (46)) should enhance the descriptive power of these models and the accuracy of policy evaluations obtained under these models.

[44] There is a substantial educational psychology literature that discusses mechanisms that underlie errors in algebraic thinking and methods to overcome these errors ((43), (44)).

[45] A related psychology and economics literature investigates how scarcity might affect various cognitive functions and lead to suboptimal behavior in many domains (e.g. (45)).

1.4.2 Implications for Models of Attention

In the last decade, one of the exciting developments in the behavioral economics literature is the increasing number of theoretical accounts of attention. We present evidence on how attention to various attributes systematically changes across frames and we further relate those findings to allocation behavior.

According to the salience theory proposed by (16), a salient thinker allocates strictly greater attention to balance information compared to interest rate information since the balance information shows greater variability.[46] Similar to salience theory, both (17)'s model of focusing and (25)'s model of sparsity predict greater attention to balance information as the range of outcome utilities differ more in that attribute compared to interest rate information. Our results on time spent on each attribute justify this prediction under the debt frame. However, we observe our subjects allocating similar levels of attention toward balance and interest rate information under the investment frame which stands in contrast to the predictions of these models. This suggests that accounting for the valence of information might improve the descriptive success of these theories.

These models' consequent predictions on the choices that agents make do not help us explain subjects' choices in our experiment. (16) is constructed to accommodate additively separable utility functions in attributes, and do not capture the richer interaction in attributes in the allocation problems that we investigate. Although (17) and (25)'s models allow for a more general class of utility functions, their predictions align with rational choice, which is clearly inconsistent with our results.

Our results on asymmetric attention allocation are also inconsistent with models of

[46] In order to obtain predictions from these models, we think of our subjects' choice as a discrete choice problem with 501 choice objects. Each choice object c is a four-tuple that lays out the balance on the left account after allocating $x \in \{0, 1, ..., 500\}$ to the left account, balance on the right account after allocating x to the left account, interest rate on the left account, and interest rate on the right account.

selective attention where people derive direct utility from attending to information (e.g. (26)). In this class of models people optimally choose to avoid information that negatively affects their welfare. Although such models predict an asymmetry in attention allocation to balance information across debt and investment frames, the direction of the asymmetry is in contrast to our findings.

1.5 Conclusion

This paper provides clear evidence regarding people's struggle with correctly solving simple trade-offs with financial frames. We move beyond existing findings in the literature by examining the sources of such suboptimal behavior using a diagnostic laboratory experiment. We show that standard explanations for consumer mistakes such as optimization ability and limited attention fall short of explaining the observed misallocations. We document the role of information salience by examining two channels that could affect allocation behavior. We find that vividness of balance information plays no role in driving the suboptimal allocations. Instead, we show that people's ability to solve such simple trade-offs is substantially hindered by the instrinsic negative frame of the debt payment situation.

Our findings have both applied and theoretical implications. On the policy side, we show limited effectiveness of traditional disclosure policies. We think that further research in psychology-guided disclosure policies is needed to establish their overall effectiveness as a way to restore rational choice. We also show that optimization ability does not pin down our subjects' ability to correctly resolve such simple trade-offs. We think that the mixed results that are obtained on the effectiveness of financial education programs might be partially due to the differences in the content of such programs. Specifically, we think that financial education programs that acknowledge the mental gaps between algebraic problems and real-world counterparts might be more effective in improving financial outcomes of the decision makers.

On the theory side, we show that existing models of attention are not able to fully capture the way that attention affects choice behavior across frames. We think that a valence-based approach to attention might be fruitful in generating insights regarding the richness of consumer behavior.

Chapter 2

Restoring Rational Choice in Repayments: Disclosures or Advice?

2.1 Introduction

With the renewed interest in household balance sheets in the aftermath of the Great Recession, household finance researchers uncovered novel inefficiencies in how borrowing individuals manage their finances. In the most dramatic versions of these inefficiencies, consumers fail to exploit differences in prices for identical products and incur significant welfare losses. Alerted by the burden of these inefficiencies on the households, policy-makers seek ways to empower individuals to make more informed financial decisions. Starting from the Truth in Lending Act of 1968, a significant amount of policy focus has been placed on information disclosure policies where financial institutions are required to provide consumers with salient and standardized information on their cost of borrowing. On the other hand, the advent of financial technologies reduced the cost of providing advice and provided an alternative way of achieving informed decision making through educating consumers on the spot. Despite our growing knowledge on the type of mistakes consumers seem to make, we know little about what type of interventions would serve as effective ways of informing and protecting consumers in the marketplace.

The goal of this study is to compare the effectiveness of two alternative approaches to restoring rational choice in financial markets: information disclosure policies and providing automated advice. In order to achieve this goal, we focus on a novel choice inefficiency where recent studies document a significant welfare loss for consumers. (3) and (4), document how credit card holders with revolving debt fail to exploit price differences in their accounts while making their repayments and incur additional interest charges that would otherwise be avoidable. In particular, individuals with multiple credit cards do not choose to fully repay their card with the highest interest rate. Instead, they seem to follow heuristic rules where they make their repayments proportional to their balances on each card.

A striking feature of this repayment problem is that one can identify an optimal allocation rule without making assumptions on consumers' preferences. This is because rationality dictates one to minimize their cost of borrowing irrespective of their time and risk preferences. The fact that one does not need to estimate preferences to identify an optimal decision drastically simplifies our welfare analysis as the extent of choice inefficiency serves as a direct measure of welfare loss.

Using the debt repayment problem as a proof-of-concept, we aim to answer what type of interventions would help consumers better exploit simple arbitrage opportunities on the liability side of their balance sheets. We construct a simple choice environment that mimics an online payment screen in the laboratory where we can clearly identify an optimal repayment rule. Identifying an optimal repayment rule is harder to achieve in the field as the list of potential confounds is large: 1) One credit card might provide additional benefits that distort incentives 2) One card provider might make it easier to make payments through digital banking and density of the branch network 3) Individuals might care about their credit scores which take balance information on each card into account and so on.[1] Another advantage of a laboratory experiment is that it allows us to sidestep the endogeneity issues that arise while studying the effectiveness of advice and it allows us to control for various measures of financial literacy while comparing the effectiveness of alternative consumer protection approaches.

Our experiments are inspired by the underlying principles of two alternative approaches to restoring rational choice in financial markets: information disclosures and advice. Information disclosures are highly popular around the world as a remedy to consumers' financial mistakes due to their simplicity and often low-cost of implementation ((47)). Alternatively, providing individuals with financial advice is another way of educating consumers to make informed financial decisions. With the advent of financial technologies, the cost of providing unbiased financial advice is lower than ever.

Although both disclosure policies and advice aim to empower individuals to make better financial decisions without directly intervening in the markets, we would expect their effectiveness to differ based on the underlying reason for the inefficient choice behavior. If financial mistakes stem from attending and processing relevant information, we would expect providing individuals with salient and easy-to-process information might improve outcomes. Following this line of reasoning, we conduct two treatments where we manipulate the display of interest rate information. In our first treatment, we increase the salience of interest rate information. In a typical credit card statement or online payment screens, interest rates are presented as non-salient attributes and instead providers make balance information predominantly salient. Increasing the salience of interest rates ensure that individuals pay attention to the price information of their credit cards. Hence this treatment aims to zero out the cost of attending the interest rate information. We find that increasing the salience of interest rate information does not improve the extent

[1] Online Appendix of (3) provides a larger set of confounds.

of choice inefficiencies. This suggests that attentional constraints play a limited role in the observed choice inefficiencies and a conventional disclosure policy where interest rate is displayed saliently would have a limited impact in reducing the choice inefficiencies.

In our second treatment, we change the percentage format of the interest rate to a fee format. Research in probabilistic reasoning identifies a mental hurdle for processing percentage information and documents that this hurdle is much less present when the probability information is presented in frequencies (48). A more directly relevant study (49) documents consumers struggle with converting the percentage interest rates into fees. We find that displaying the interest rate information in a fee format helps alleviate the choice inefficiencies to a certain extent although this does not lead to a change in behavior in conventional significance levels. This suggests that processing interest rate information can be a hurdle for some consumers to exploit simple arbitrage opportunities and hence a behaviorally informed disclosure policy where the percentage format of the interest rates is substituted with a fee format would help consumers make better decisions.

An equally actionable policy alternative to information disclosures is the provision of automated financial advice. If the observed choice inefficiencies result from cognitive limitations or psychological biases, direct provision of financial advice should be more effective than disclosure policies. A key challenge with providing financial advice to individuals as a solution to fix their financial mistakes is that they might not understand that they are committing a financial mistake and might not demand financial advice. Therefore, we conduct a treatment where we elicit subjects' willingness-to-pay for automated financial advice and test if purchasing financial advice improves outcomes. We find that over 90% of our subjects are willing to pay some amount for financial advice. If we take the amount of money our subjects leave on the table as a benchmark for their willingness-to-pay, we find that subjects are slightly yet significantly under-demanding financial advice. This suggests that subjects are somewhat aware of the extent of their choice inefficiencies and respond to an opportunity to purchase advice that would help them reduce their inefficiencies.

Next, we find subjects who purchase advice dramatically reduces the extent of their choice inefficiencies. Using the random BDM prices as an instrument for subjects' willingness to pay, we find that purchasing advice generates a 25 p.p. increase in the share of optimal choices, corresponding to 100% increase compared to the baseline optimality rate. Although subjects benefit from financial advice, we find that they do not completely eliminate their choice inefficiencies after purchasing advice. The inefficient use of advice generates a situation where our subjects are willing to pay more than what they should pay given how well they use advice and ends up making our subjects who purchase financial advice as the worst performing group in our welfare calculations. Our results on the provision of financial advice shows both a promise and a pitfall. Although consumers benefit significantly more from the provision of financial advice, their tendency to "overpay" for financial advice might make them worse off in a free market environment. This suggests that not only promoting but also subsidizing consumer financial technology applications where consumers could be provided with automated financial advice would be a significantly more effective way of protecting consumers from simple arbitrage failures.

Our study has three main contributions. First, we elicit individuals' willingness-to-pay for financial advice that would help them move beyond their cognitive limitations to tackle a choice inefficiency using an incentive-compatible mechanism. There is a nascent literature studying the effect of unbiased financial advice on the optimality of financial decisions. (50) find that retail investors who would benefit from obtaining financial advice the most are less likely to demand it and the advice is hardly followed among those who accept the advice. (51) find that adopters of a robo-advisor that constructs portfolios tailored to investors' holdings and preferences are less likely to exhibit behavioral biases and they exhibit the same set of behavioral biases to a smaller extent after the adoption of the robo-advising tool. (52) find that a robo-advisor induces investors to pay more attention to their portfolios, to increase their investment and exposure to equity, and it results in higher risk-adjusted returns. To our knowledge, this is the first study that investigates the demand for financial advice in borrowing decisions and directly measuring willingness to pay for financial advice.

Second, we compare the effectiveness of conventional disclosure policy against a behaviorally informed disclosure policy. There is a growing experimental literature studying the effects of information disclosure policies within the context of credit card markets.[2] (57) studies if credit card borrowers in Brazil respond to interest rate salience when offered a menu of payment plans that would allow them to pay down their balances in fixed installments. He finds that salient disclosure of interest rates has no effect on take-up rates and enrollment interest rate elasticities. (23) investigates what type of information disclosures would be effective in improving various financial outcomes for indebted individuals in the Mexican credit card market using a randomized controlled trial. They test the effect of salient information disclosure policies alongside direct warnings about the consumer's probability of default and peer comparisons in terms of the amount of debt individuals carry. They find that salient information disclosure has no effect while direct warnings and peer comparison treatments have modest effects.[3] Our study complements these studies by reinforcing their findings in a controlled environment where attentional constraints of decision makers are minimized.

Third, we compare the effectiveness of two distinct approaches to restoring rational choice in financial decisions. We are not aware of any study that uniformly investigates

[2]There is a broader household finance literature that looks at the effect of information provision on consumer welfare in the context of micro-loan take-up (53), mutual fund choice (54), payday loans (22), retirement contributions (55) and savings accounts (56).

[3]An earlier set of studies directly investigates the role of disclosure reforms in the US on various market outcomes in consumer loan and credit card markets. (58) test the impact of Truth in Lending Act of 1968 on lenders' ability to price discriminate in consumer installment loans and (59) test the impact of CARD Act of 2009 on borrowing outcomes.

the effectiveness of financial advice and disclosure policies. The closest to our study in this sense is (60). They compare the effect of information disclosures and contract standardization on delinquency and default rates in Chilean consumer loan market. They find that information disclosures are substantially more effective than contract standardization in decreasing delinquency and default rates.

Finally, this study complements our earlier investigation (61) where we document people's struggle to exploit simple arbitrage opportunities in a repayment context replicating the findings of (3) and (4) in a controlled environment. While (61) focuses on robustly replicating the field results and testing various explanations that inform theories of boundedly rational behavior by contrasting the repayment situation to an algebraically identical investment situation, this paper focuses on the performance of two directly actionable consumer protection policy approaches to fix such simple arbitrage failures.

2.2 Experimental Design

The goal of our design is three fold. First, we want to test if manipulating the display of crucial information would induce a change in behavior and if so, what type of manipulation would work more effectively. Second, we want to understand how people respond to an opportunity to purchase financial advice and test if the purchase of financial advice would induce a change in behavior. Third, we want to evaluate the welfare of our participants under these two alternative consumer protection measures.

In order to clearly answer these questions, we create a simple choice environment that is easy for our subjects to navigate while sharing similarities with actual credit card statements in the way that essential information is communicated. We build on our earlier design in our companion paper (61) to answer our research questions.

2.2.1 Baseline Design

The experiment captures the essential features of the decision environment faced by credit card consumers who make their repayments in the field (See Figure 2.1). Each subject is endowed with two hypothetical credit card accounts and a hypothetical checking account. The experiment consists of multiple periods. At the beginning of each period, we deposit a fixed amount of 500 *Experimental Currency Units* (ECU) into their checking account. Subjects' task in each period is to make repayments toward their credit cards using their deposit. During a period, subjects face a screen that is split into two halves. Each half represents a credit card account. At the top part of each half of the screen, we saliently display the current balance information. At the center of the screen, subjects see a list of other account attributes that are typically displayed on a credit card statement. These attributes are interest rate, interest charged, previous balance and previous repayment. The information on each of these attributes is presented simultaneously and singularly to a subject once she clicks on the *information button* that carries the name of that attribute. Clicking on information buttons is costless and subjects are allowed to click freely. Each period ends once a subject submits an allocation decision. It is important to emphasize that subjects always see how much they owe on an account at the top part of the screen and they do not need to click any button to acquire balance information while they need to click the information buttons to see other attributes. The salient display of balance information in our design mirrors the salient display of balance information on actual credit card statements and online accounts, and is an aspect of the decision environment we manipulate in other treatments.

Figure 2.1: Experiment Interface

A crucial aspect of this repayment problem in the field is that consumers do not get explicit feedback on the quality of their decisions. The only feedback consumers get is the amount of interest charged on each account which is then incorporated in the total debt they owe to each card in the subsequent period. We recreate this implicit feedback mechanism in the laboratory by employing a block design where we combine decision periods into stages. Each stage consists of five decision periods.[4] In the first period of each stage, we determine the amount of debt on each card. In the subsequent periods, each subject's debt on each card is endogenously determined by their previous allocation decisions in that stage. Since subjects are assigned some debt at the beginning of each stage, we endow subjects with a fixed positive amount in order for each subject to make some money in the experiment. We determine a subject's payoff for a stage by their end of stage balance on each card subtracted by the fixed endowment. We then convert their

[4]We choose five periods per stage to have a sense of subjects' within stage learning and to keep the duration of the experiment reasonable.

stage payoffs into US dollars and randomly choose one of their stage payoffs for their actual payment.

We employ six stages with different balance and interest rate configurations. The parameter choices for the first period of these stages are presented in Table 2.1. We choose the interest rate difference to be 1.5% as a plausible upper bound of the the observed monthly interest rate differences in the field.[5] We keep the interest rate difference across stages fixed to keep the incentives the same across these stages. We choose the initial balances to be consistent with the average credit card debt observed in the field.

[5](4) document that the observed annual interest rate difference is 15% at the 90th percentile corresponding to a monthly interest rate difference of 1.25%. (3) find the average monthly interest rate gap to be 1.1% in their data.

Table 2.1: Parameter Choices

Stage	Account	Interest Rate	Initial Balance
1	1	4.90%	4,450
	2	3.40%	3,050
2	3	5.70%	2,950
	4	4.20%	4,350
3	5	3.70%	4,450
	6	5.20%	2,950
4	7	3.90%	2,850
	8	5.40%	4,450
5	9	5.30%	4,650
	10	3.80%	3,150
6	11	5.90%	3,050
	12	4.40%	4,550

Opportunity to Purchase Advice

Once subjects go through the first four stages of the experiment, we provide them with new instructions and an opportunity to purchase advice. We frame the decision to purchase advice as hiring a robo-advisor that will help subjects how to minimize their total debt in the remaining two stages of the experiment and allow them to obtain the maximum possible earnings from the experiment which is \$15. Subjects are then asked how much they would be willing to pay to hire a robo-advisor on a range from \$0 to \$15.

B provides the instructions.

We incentivize subjects to truthfully state their willingness to pay using the Becker-DeGroot-Marschak mechanism (62). The mechanism is frequently used in laboratory and field studies to elicit willingness-to-pay due to its ability to provide an exact WTP measure under the expected utility hypothesis. After a subject states their willingness to pay, we randomly draw a price from the specified range. The subject gets to hire the robo-advisor if the drawn price is lower than or equal to subject's stated WTP. The mechanism operates as a second-price auction and the dominant strategy for a subject is to state their true WTP irrespective of their risk preferences. In addition to explaining the subjects the details of the mechanism, we explicitly tell subjects that stating their true WTP is a dominant strategy for them (63). Screenshots of the procedure to hire a robo-advisor can be found in B.3.

Once the random BDM prices are realized, subjects whose valuations are weakly greater than the realized prices acquire robo-advisor. These subjects complete the remaining stages of the experiment with an advice at the top left hand side of the screen that tells them exactly what they need to do to minimize their interest charges. An example of the screen is provided in Figure B3.7. Subjects who do not end up with a robo-advisor complete the remaining stages of the experiment which are identical to the first four stages except for the parameter choices.

Despite the strong incentive compatibility properties of the mechanism, recent studies document a lack of training on the mechanism hinders these properties (64). We address this issue by providing an understanding quiz on the mechanism that provides subjects with detailed feedback to their answers and a practice period before the actual elicitation. The understanding quiz can be found in B. Another critique of the BDM mechanism is that the price range provided by the experimenter affects the valuation of the item to be purchased although such effect is not predicted by the theory (65). In our experiment, rationality puts natural bounds on what this range can be since the value of advice should rationally be at most how much subjects can improve their earnings. We set the upper bound of the range of our BDM prices to this natural bound, that is, the maximum amount of earnings a subject can make.

Measuring Optimization Ability and Financial Literacy

At the end of the six stages that are described above, subjects are asked four incentivized optimization problems represented in algebraic expressions. The optimization problems correspond to algebraic versions of the allocation problems subjects go through in the main part of the experiment. We use the subjects' scores on these problems as a proxy for their optimization ability. We consider "optimization ability" as a context-specific measure of cognitive ability. An important design choice here is that we do not ask optimization problems at the beginning of the experiment as it might affect subjects' ability to optimize in the experiment.

In addition to the optimization problems, we ask subjects the big three financial literacy questions that has now become the global standard for measuring financial literacy (66). The experiment ends with a simple survey where we collect demographic information (gender, years in college, credit card account ownership) and ask subject their reasoning for the choices that have made in the main part of the experiment.

2.2.2 Understanding the Design

The controlled laboratory environment allows us to remove many confounding features of the actual decision environment, clearly define a simple arbitrage situation between the two accounts and incentivize our subjects to exploit this price difference.

First, the sequential nature of due dates in the field might lead consumers to narrowly bracket their payment decisions to each card and induce them to ignore the interdependency between their payments ((67), (68)). Such narrow bracketing naturally incentivizes consumers to make a decision between how much cash to hold and how much payment to make at each due date rather than exploiting the price differences between the two cards. We eliminate the possibility of narrow bracketing by requiring subjects to make simultaneous payment decisions to each card.

Second, credit card heterogeneity in the field confounds the incentives to pay off the more expensive credit card. The cheaper credit card might provide greater additional benefits to consumers in the form of cash rewards and miles. The "credit cards" we endow our subjects with in our experiment do not provide such additional benefits.

Third, credit availability on each card is an important component of credit score calculations.[6] A consumer might then have an incentive to reduce the amount she owes on her cheaper card if she owes a significant amount on that card. Our experiment eliminates the possibility of this confound by not featuring credit scores.

Fourth, minimum payments required on each card leads consumers to anchor on this amount ((28), (29)). This suggests a consumer who has a higher minimum payment required on her cheaper card might allocate a greater proportion of her payments to the cheaper card. We eliminate this possibility by not requiring a minimum payment amount. In addition to removing these essential confounding factors, we simplify the repayment problem further by providing our subjects easy access to their interest rate information and plenty of time to think about their decisions. The cost of accessing interest rate information is as low as clicking a button and subjects, on average, has 100 seconds to make a payment decision.[7]

[6]The amount of debt determines 30% of the commonly used FICO score: https://www.myfico.com/credit-education/credit-scores/amount-of-debt

[7]Knowledge of interest rate information at the time of repayment is a significant source of variation in the actual decision environment as the interest rate information is complexly disclosed.

2.2.3 Additional Treatments

We conduct two treatments in addition to the **Baseline** treatment to test if manipulating the display of interest rates changes behavior. We eliminate the opportunity to purchase advice in both of these additional treatments.

In treatment **Salience**, all features of the design is identical to **Baseline** except that 1) we display interest rate information very saliently at the top of the screen instead of current balance information and 2) we eliminate the opportunity to purchase advice. Although our **Baseline** treatment substantially reduces the cost of attending to interest rate information by requiring subjects to simply click on the interest rate information to reveal it, **Salience** treatment goes a step further to completely eliminate the cost of attending to interest rate information by not requiring subjects to reveal their interest rate information and making it the only piece of information on their screen that they can see without any effort. Through this treatment, we take the Truth in Lending Act[8] type of disclosure requirement to an extreme where the cost of borrowing is very salient and easy to compare between the two credit cards. A screenshot of the interface for this treatment is provided in Figure 2.2.

[8]See a summary of the Act here: https://www.debt.org/credit/your-consumer-rights/truth-lending-act/

Figure 2.2: Interface for Treatment **Salience**

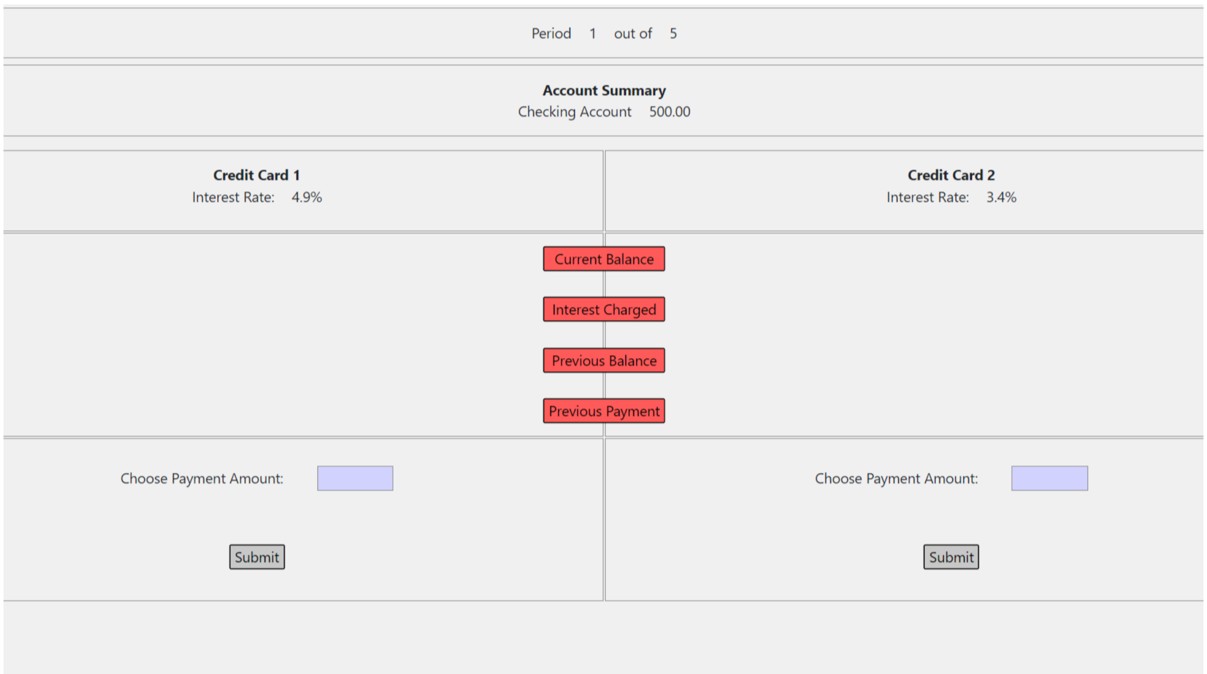

In treatment **Fee Format**, all features of the design is identical to **Baseline** except that 1) we change the percentage format of the interest rate to a fee format and 2) we eliminate the opportunity to purchase advice. **Fee Format** aims to test if people struggle with processing the interest rate information as a price and hence fail to make an accurate cost comparison across the two cards. This treatment aims to test the effect of a potential behaviorally informed Truth in Lending Act type of disclosure policy. Although the Act standardized the display of cost of borrowing through Annual Percentage Rate, people might fail to make accurate price comparisons if they struggle with percentages. A screenshot of the interface for this treatment is provided in Figure 2.3.

Figure 2.3: Interface for Treatment **Fee Format**

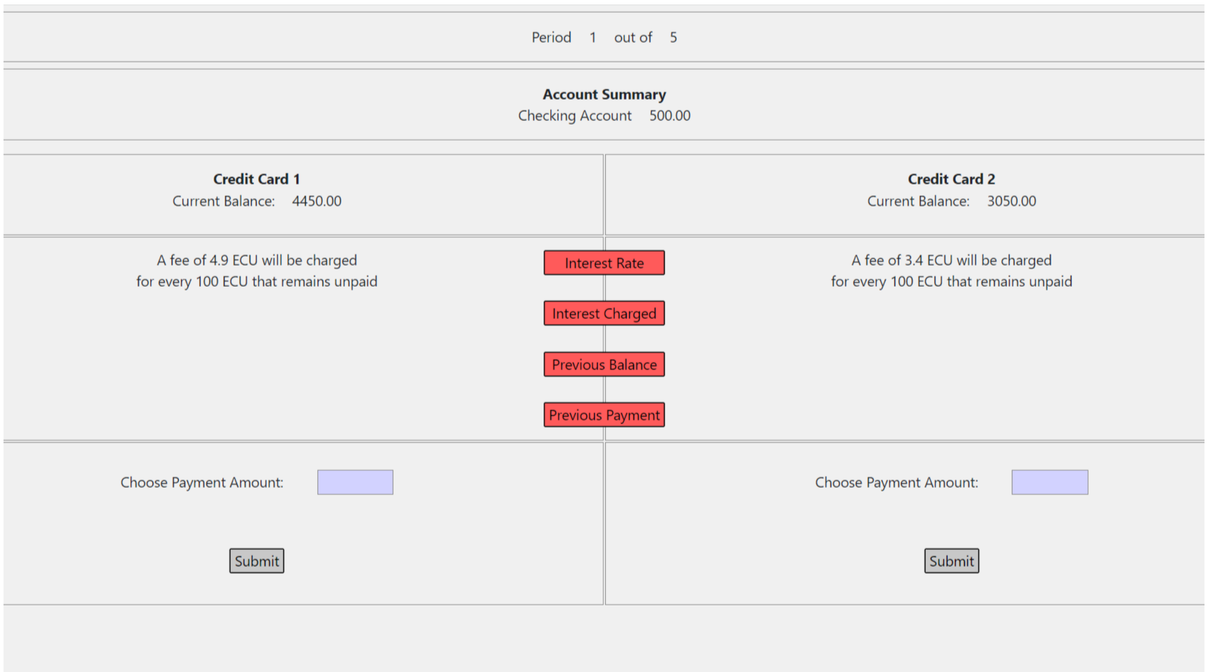

2.2.4 Procedural Information

We conducted our experiment online using the subject pool of UCSB Experimental and Behavioral Economics Laboratory. The experiment was coded using o-Tree software ((69)). A total of 199 subjects, recruited through ORSEE (Online Recruitment System For Economic Experiments) (70). The average payment per subject was $15.3 including a $5 show-up fee. Sessions for the Baseline treatment lasted for 90 minutes and the sessions for other treatments lasted for 60 minutes.

2.3 Results

2.3.1 Do subjects make optimal payments?

We first replicate our previous findings on the suboptimality of payments in our simple environment. We consider two measures of optimality. The first one is a binary variable that takes the value 1 if a payment is made fully made toward the higher interest rate card. We call this measure the optimality rate. Rational choice theory prescribes an optimality rate of 100%. However, as illustrated by Figure 2.4, only 25.4% of the repayments are fully made toward the card with the higher interest rate. The average optimality rate we calculate is statistically different than the average optimality rate of 100% prescribed by rational choice theory (clustered Wilcoxon signed-rank test, $p < 0.001$).

Our alternative measure of optimality, which we call as the misallocation rate, is a continuous variable that reflects the percentage of a payment made toward the lower interest rate card. Similar to the optimality rate, rational choice would necessitate zero percentage of a payment to be allocated towards the low interest rate card. However, we find an average misallocation rate of 33.4%. The average misallocation rate we measure is statistically different than the average misallacation rate of 0% prescribed by rational choice theory (clustered Wilcoxon signed-rank test, $p < 0.001$). Both measures of optimality are similar to our findings in our companion paper where we find an optimality rate of 22.4% and a misallocation rate of 35.6%.

Optimality Measures by Optimization Ability and Financial Literacy

We find that the ability to solve an algebraic version of the credit card repayment problem is a significant predictor of making optimal payments. Indeed, we find an average optimality rate of 9.4% among our subjects who are unable to solve the algebraic version of the problem. However, the average optimality rate goes up by 34.6 percentage points

Figure 2.4: Distribution of Payments in the First Four Stages of the Baseline

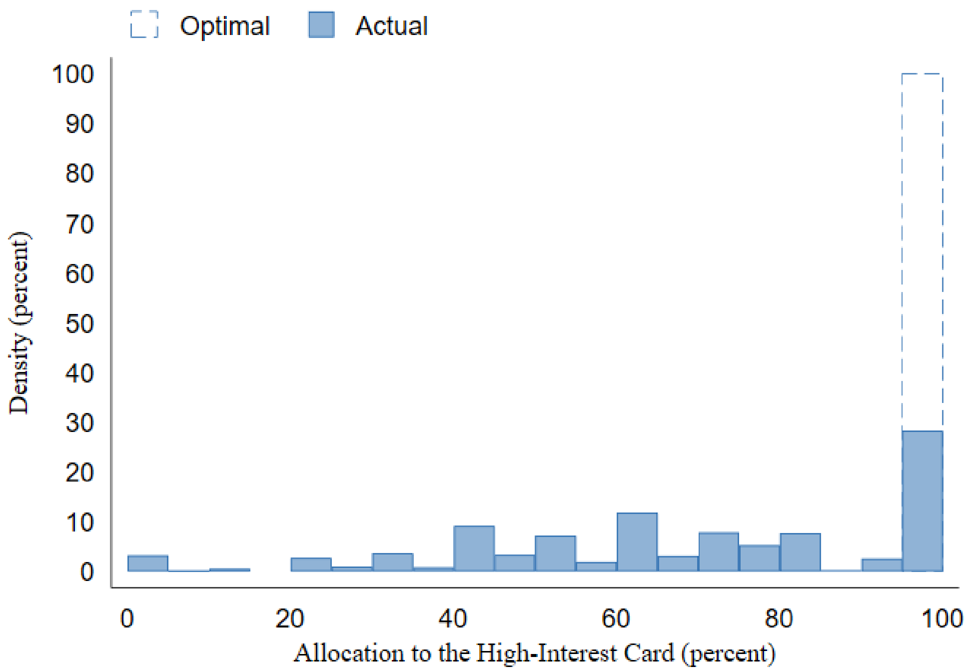

Notes: Figure shows the distribution of payments made towards the high interest rate card as a percent of the available deposit in the first four stages of the experiment. The unit of observation is subject-by-period (91x20). The bin with dotted lines indicates the degenerate distribution prescribed by the rational choice theory. The histogram contains 50 equally sized bins.

to 44% among our subjects who are able to solve an algebraic version of the problem. Moreover, we find that only 42 out of 91 subjects can solve an algebraic version of the problem which is statistically indistinguishable from half of our subject population unable to solve the algebraic version of the credit card repayment problem (binomial test, $p = 0.53$).

Unlike optimization ability, a traditional measure of financial literacy, the big three questions, does not predict making optimal payments. We find an average optimality rate of 18.6% among our subjects who are unable to solve at least one of the big three questions we ask. Although the subjects who can solve all the big three questions has a higher average optimality rate of 29.8%, the increase in the optimality rate is statistically indistinguishable from 0 ($p = 0.13$). Similar to the optimization ability, we find that only 55 out of 91 subjects can solve all big three questions suggesting that half of our subject population unable to solve all big three financial literacy questions.

Figure 2.5: Optimality Rate by Optimization Ability and Financial Literacy

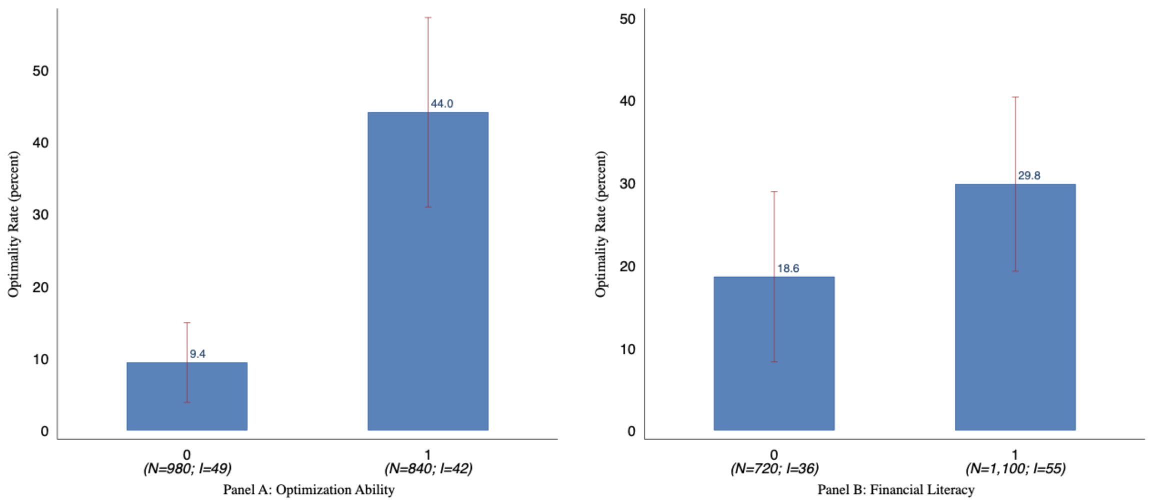

Notes: Panel A shows the average optimality rate by our subjects' ability to solve an algebraic version of the credit card repayment problem. Subjects who are unable to solve the algebraic version of the credit card repayment problem are indicated by the group O. Panel B shows the average optimality rate by our subjects' ability to show the Big Three financial literacy questions. Subjects who are able to solve *all* big three questions are indicated by group 1. The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

The results are qualitatively similar when we consider the misallocation rate instead of the optimality rate as our measure of optimality. Subjects who do not have "optimization ability" has an average misallocation rate of 39.4% whereas subjects who have "optimization ability" has an average misallocation rate of 26.2%, a significant lower misallocation rate ($p < 0.001$). On the other hand, subjects who fail to solve one of the big three financial literacy questions have an average misallocation rate of 35.9% whereas subjects who successfully solve all big three financial literacy questions have an average misallocation rate of 31.7%, displaying an insignificant decrease ($p = 0.17$). Figure B1.1 visualizes these results in Appendix B.1.

Result 1 *A significant majority of subjects do not make their payments optimally. Subjects who can solve an algebraic version of the credit card repayment problem are significantly better at making optimal payments. A traditional measure of financial literacy does not predict making optimal repayments.*

2.3.2 Measuring Demand for Financial Advice

A natural benchmark for subjects' willingness-to-pay for financial advice is the amount of money they leave on the table by not making their payments optimally. A rational subject who is aware of the fact that he has left some amount of money on the table on the first four stages of the experiment should be willing to pay an equivalent amount for financial advice unless he has a compelling reason to believe that his allocations will improve in the remaining periods. We find that in the first four stages of the experiment, misallocation rate reduces by 0.4% over time ($p = 0.49$), suggesting there is no learning and hence no compelling reason for rational subjects to anticipate an improvement over their misallocation rate.

We find that 90.1% of subjects should be willing to pay some amount for financial advice ($p < 0.001$). The average willingness-to-pay for financial advice should be 33.4% of \$15 which is the maximum earnings from the experiment that would be obtained if subjects were to make no misallocation in their payments in the remaining periods. However, our subjects are willing-to-pay 29.1% of the maximum potential earnings from the experiment. The difference in rational and actual willingness-to-pay amounts is statistically significant ($p = 0.017$). This suggests that at the aggregate our subjects demand financial advice less than they should. Interestingly, we find that 97.8% of subjects are willing-to-pay some amount for financial advice. The difference in the percentage of subjects who should pay and who pay for financial advice is statistically significant ($p = 0.007$). The results show that an overwhelming majority of our subjects value financial advice to some extent. However, aggregate demand for financial advice is less than what they should "rationally" demand.

Figure 2.6: Rational and Actual Willingness-to-Pay for Financial Advice

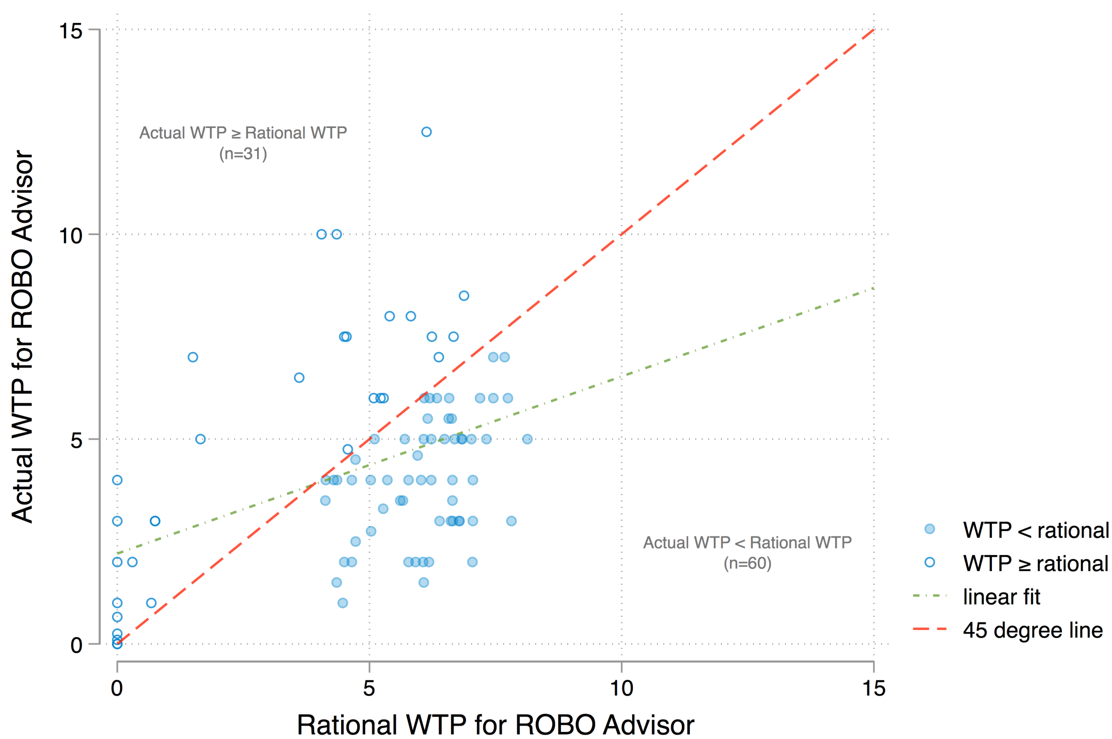

Figure 2.6 provides an individual level analysis of how much our subjects should rationally pay for financial advice and how much they do pay. We find that 60 out of 91 subjects are willing-to-pay less than what they should pay for financial advice. We reject the hypothesis that only half of our subjects "under-demand" financial advice (binomial test, $p = 0.003$). Joint with our aggregate analysis, the evidence suggests that subjects

do not value financial advice as much as they should.

Demand for Financial Advice by Optimization Ability and Financial Literacy

Table 2.2 disaggregates the actual and rational willingness-to-pay by levels of optimization ability and financial literacy. We find that subjects who are unable to solve an algebraic version of the credit card repayment problem has a higher willingness-to-pay compared to the subjects who can. However, subjects who fail to solve an algebraic version of the problem still exhibit a significantly lower willingness-to-pay compared to what they should rationally be willing to pay ($p = 0.005$). On the other hand, there is no significant difference in actual and rational WTP amounts for subjects who can solve an algebraic version of the problem. This suggests that a lack of optimization ability is not only associated with making suboptimal repayments but also it interferes with subjects' ability to accurately demand financial advice.

The results are somewhat similar for financial literacy. Although subjects who fail to solve one of the big three questions do not demand financial advice more than those who successfully answer all big three questions, they exhibit a significantly less willingness-to-pay compared to what they should rationally be willing to pay ($p = 0.02$).

Table 2.2: WTP for Financial Advice by Optimization Ability and Financial Literacy

	Low		High	
	Actual WTP	Rational WTP	Actual WTP	Rational WTP
Optimization Ability	32.2%	39.4%	25.5%	26.2%
Financial Literacy	29.0%	35.9%	29.2%	31.7%

Notes: The table presents the averages of actual and rational willingness-to-pay for financial advice by optimization ability and financial literacy. Actual WTP is the amount elicited through the BDM mechanism indicated as the percentage of USD 15. Rational WTP is the percentage of USD 15 subjects leave on the table by not making their payments optimally before the opportunity to purchase financial advice is presented to them. Optimization ability is a dummy variable that is *Low* for subjects who fail to solve an algebraic version of the credit card repayment problem. Financial literacy is a dummy variable that is *Low* for subjects who fail to solve at least one of the big three financial literacy questions.

We summarize our findings on the demand for financial advice in the following result:

Result 2 *An overwhelming majority of subjects appreciate financial advice yet a significant majority of subjects are willing to pay less than what they should rationally be willing to pay. The extent of "under-demand" for financial advice is significantly greater among subjects who fail to solve an algebraic version of the problem. Similarly, the extent of "under-demand" for financial advice is significantly greater among subjects who fail to achieve a full score on a traditional measure of financial literacy.*

2.3.3 Effectiveness of Financial Advice

Before describing our results on the effectiveness of providing computerized financial advice, we discuss which of our subjects end up with a robo-advisor. We find that 23.1% end up with a robo-advisor after the realization of the BDM prices. This is statistically indistinguishable from our anticipation on the percentage of subjects ending up with robo-advisor before the realization of the BDM prices, which is 29.1% (binomial test, $p = 0.25$). Among the subjects who end up with a robo-advisor, 42.9% can solve an algebraic version of the credit card payment problem while 57.1% has perfect score on financial literacy questions.

Figure 2.7 shows the average optimality rate before and after the opportunity to purchase financial advice is presented. Average optimality rate sharply increases for subjects who purchase financial advice in the remaining part of the experiment. Subjects who do not purchase financial advice has an average optimality rate of 27.4% in the first four stages before the opportunity to purchase advice is presented and this rate moves up to 27.7% in the last two stages. On the contrary, subjects who purchase financial advice has an average optimality rate of 18.6% in the first four stages and this rate moves up to 40.5%.

Figure 2.7: Effect of Financial Advice on Making Optimal Repayments

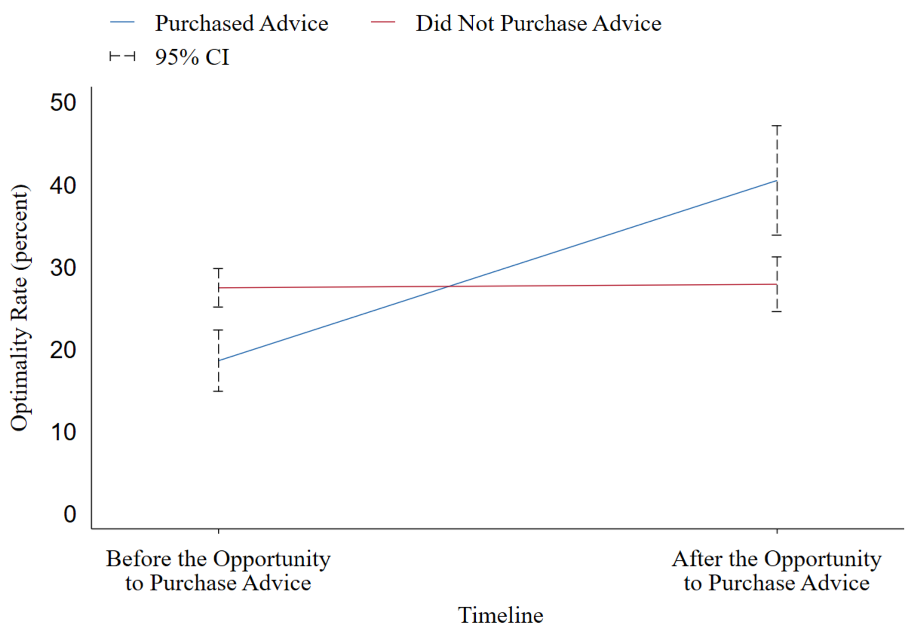

Notes: The figure presents the average optimality rate for purchasers and non-purchasers of financial advice before and after the opportunity to buy a financial advice is presented. A repayment is optimal if a subject has allocated all of her deposit into the card with higher interest rate. The opportunity to purchase financial advice is presented at the end of the fourth stage of the experiment. Subjects go through two more stages after the opportunity to purchase financial advice. The whiskers represent the 95% confidence interval for the average optimality rate. Errors are clustered at the individual level.

We estimate the causal effect of purchasing financial advice on optimality rate using the randomly assigned BDM prices. Table 2.3 presents the results. In the first stage estimation in Column (2), the BDM price strongly predicts purchasing financial advice. Column (3) shows that purchasing advice significantly increases the probability of making an optimal repayment by 23.9 percentage points while controlling for how much subjects are willing-to-pay (actual WTP) and how much they should rationally be willing to pay (rational WTP). Column (4) shows that the BDM prices continue to strongly predict purchasing advice when we include optimization ability, financial literacy and gender as additional controls in the first stage. Column (5) shows that the effect remains unchanged when reiterate the IV estimation using aforementioned additional controls.

We similarly estimate the causal effect of financial advice on misallocation rate and find that purchasing financial advice strongly reduces the misallocation rate by 16.7 percentage points, that is, subjects allocate 16.7 percentage points less of their deposits towards the card with lower interest rate. Table B1.1 in Appendix B.1 presents the estimation results.

Table 2.3: Causal Effect of Financial Advice on Making Optimal Repayments

	OLS	First Stage	IV Estimate	First Stage	IV Estimate
	(1)	(2)	(3)	(4)	(5)
	Optimal Repayment	Purchased Advice	Optimal Repayment	Purchased Advice	Optimal Repayment
Purchased Advice	24.26**		23.93*		24.95*
	(8.934)		(10.44)		(10.06)
Actual WTP	-0.628	0.0724***	-0.607	0.0732***	-0.505
	(1.403)	(0.0167)	(1.499)	(0.0157)	(1.429)
Rational WTP	-13.83***	-0.00867	-13.83***	-0.00661	-12.79***
	(1.024)	(0.0138)	(1.019)	(0.0165)	(1.368)
BDM Price		-0.0654***		-0.0664***	
		(0.00710)		(0.00708)	
Observations	910	2730	910	2730	910
Additional Controls	No	No	No	Yes	Yes

Notes: Results from an instrumental variables regression that uses the (randomly assigned) BDM price as an instrument for purchasing financial advice to estimate the causal impact of financial advice on probability of making optimal repayments. Columns (1) and (2) present the OLS and first-stage estimates, respectively. Column (3) and (5) use optimality of repayments as the outcome variable, that is, if a repayment is fully allocated towards the card with the high interest rate. Column (4) presents the first-stage estimates using optimization ability, financial literacy and gender as additional controls. Column (5) presents the IV estimates with the aforementioned additional controls. Standard errors in parentheses. Errors are clustered at individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

We summarize our findings on the effectiveness of financial advice in the following result:

Result 3 *Purchasing financial advice significantly improves the optimality of allocations. Subjects who purchase advice are 25 p.p. more likely to make optimal repayments, corresponding to an approximately 100% increase compared to our baseline optimality rate.*

Moreover, subjects who purchase advice misallocate 17 p.p. less of their deposits towards the lower interest rate card, corresponding to an approximately 100% decrease in the misallocation rate.

2.3.4 Manipulating Information Representations

Do subjects respond to salient interest rates?

Figure 2.8 displays that the optimality rate decreases when we increase the salience of interest rate. The 4.6 p.p. decrease in optimality rate is insignificant ($p = 0.44$), but it is surprising as it goes against our predictions. When we use misallocation rate as our outcome measure, the results are qualitatively similar. We find an insignificant 3.6 p.p. increase in the amount of deposit allocated toward the low interest rate card ($p = 0.21$). Figure B1.2 provides more information.

Figure 2.8: Effectiveness of Interest Rate Salience

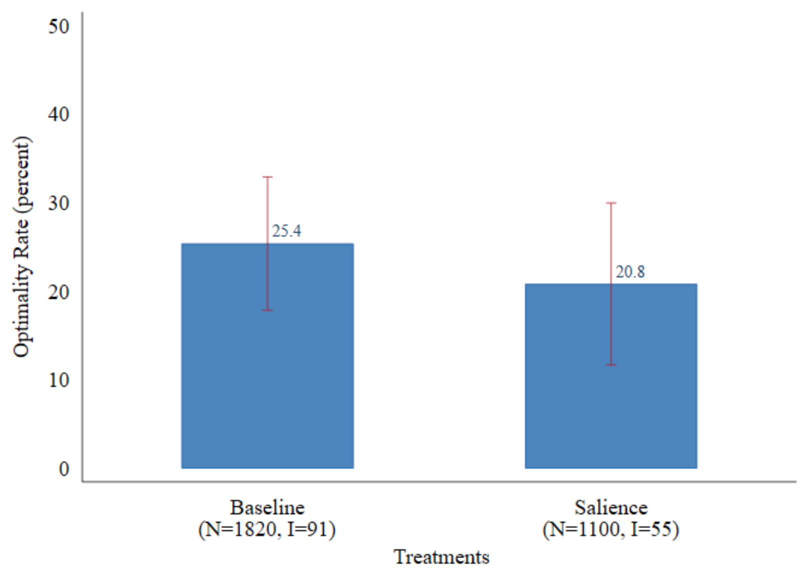

Notes: Figure shows the optimality rate in treatments **Baseline** and **Salience**. Optimality rate is the share of payments that are made *fully* towards the more expensive card among all payments. The number of observations and the number of individuals in each treatment is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Disaggregating the data by optimization ability gives us further insights on the surprising negative effect of interest rate salience on the optimality rate. The negative effect seems to be driven by subjects who have "low optimization ability", that is, subjects who fail to solve an algebraic version of the repayment problem. When we estimate the average treatment effect for each subgroup, we find that the optimality rate decreases by 4.3 percentage points among subjects with low optimization ability ($p = 0.20$). On the other hand, it increases among subjects with high optimization ability by a similar amount ($p = 0.71$). It is important to note that these effects are not significant at conventional levels. Results are qualitatively similar when we consider the misallocation rate. Figure B1.3 provides further information.

Figure 2.9: Effectiveness of Interest Rate Salience by Optimization Ability

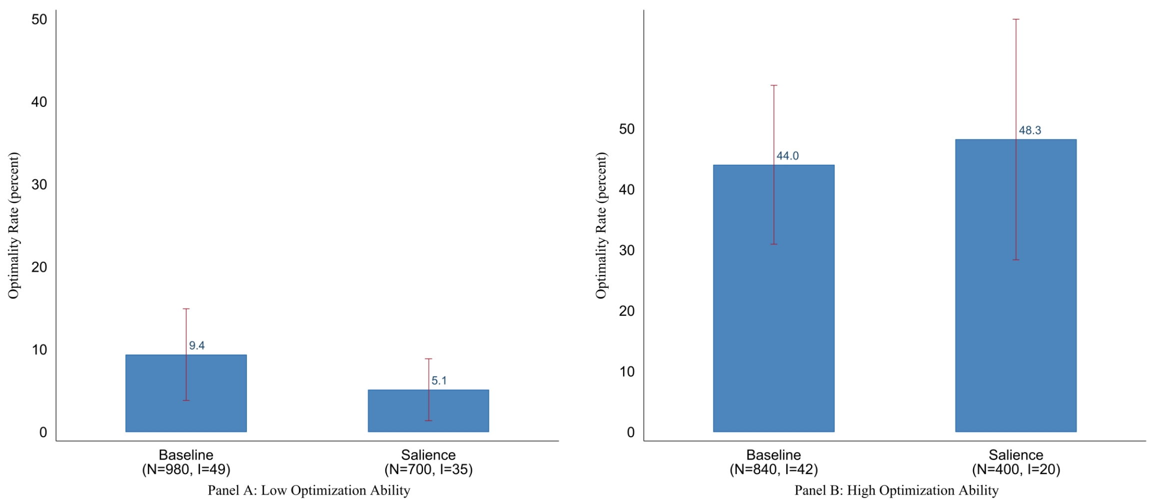

Notes: Panel A shows the optimality rate for **Baseline** and **Salience** treatments among subjects who fail to solve an algebraic version of the credit card repayment problem. Panel B shows the same rate across the same treatments for subjects who solve an algebraic version of the credit card repayment problem. The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

When we look at the effects of interest rate salience by financial literacy, we find that the optimality rate decreases by 14.5 p.p. among subjects with low financial literacy ($p = 0.009$) whereas there is no change in optimality rate among subjects with high financial literacy ($p = 0.85$). Results are qualitatively similar when we consider the misallocation rate. Figure B1.4 provides further information.

Figure 2.10: Effectiveness of Interest Rate Salience by Financial Literacy

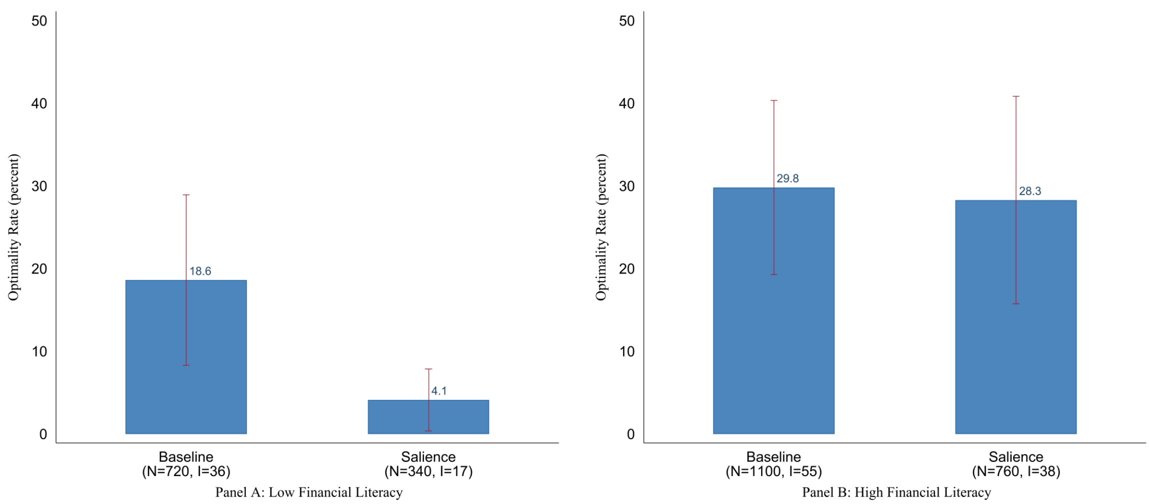

Notes: Panel A shows the optimality rate for **Baseline** and **Salience** treatments among subjects who fail to solve one of the big three financial literacy questions. Panel B shows the same rate across the same treatments for subjects who solve all big three financial literacy questions. The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Result 4 *Increasing the salience of interest rate information does not generate an increase in the share of optimal repayments. Contrary to our hypothesis, the intervention seems to decrease the optimality rate among subjects with low financial literacy.*

Do subjects respond to a change in interest rate format?

Figure 2.11 shows that the fee format helps subjects make better decisions. The optimality rate under the treatment **Fee Format** increases by 6.9 p.p., yet the effect is insignificant ($p = 0.23$). We find a qualitatively similar result when we use the misallocation rate as our outcome measure (Figure B1.5).

Figure 2.11: Effectiveness of Fee Format

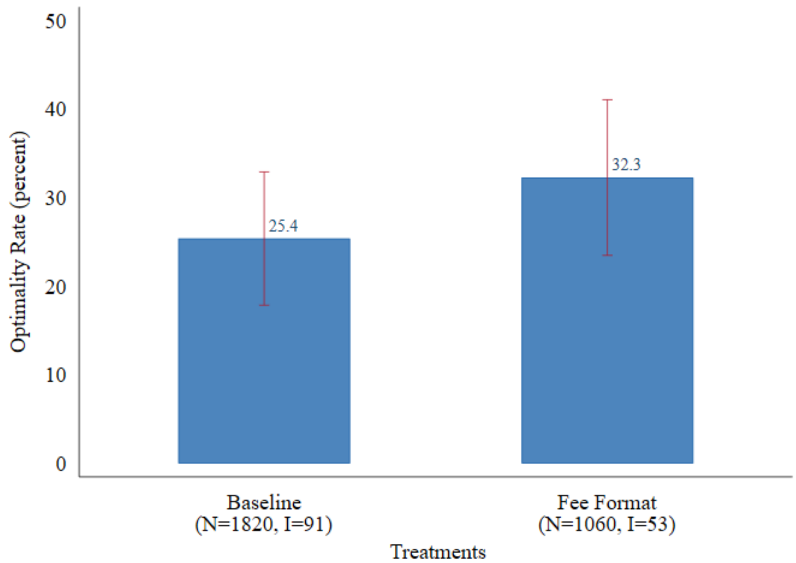

Notes: Figure shows the optimality rate in treatments **Baseline** and **Fee Format**. Optimality rate is the share of payments that are made *fully* towards the more expensive card among all payments. The number of observations and the number of individuals in each treatment is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Figure 2.12 presents heterogenous treatment effects when we disaggregate the data by subjects' optimization ability. The optimality rate among subjects who fail to solve an algebraic version of the problem increases by 12.9 p.p. ($p = 0.017$). On the other hand, the optimality rate among subjects who solve an algebraic version of the problem is virtually unchanged ($p = 0.82$). Hence subjects with low optimization ability benefit more from the intervention compared to the subjects with high optimization ability.

The results are qualitatively similar for misallocation rates. The misallocation rate among subjects who fail to solve an algebraic version of the problem decreases by 3.8 p.p. ($p = 0.076$). On the other hand, the misallocation rate among subjects who solve an algebraic version of the problem is virtually unchanged ($p = 0.95$). See Figure B1.6.

Figure 2.12: Effectiveness of Fee Format by Optimization Ability

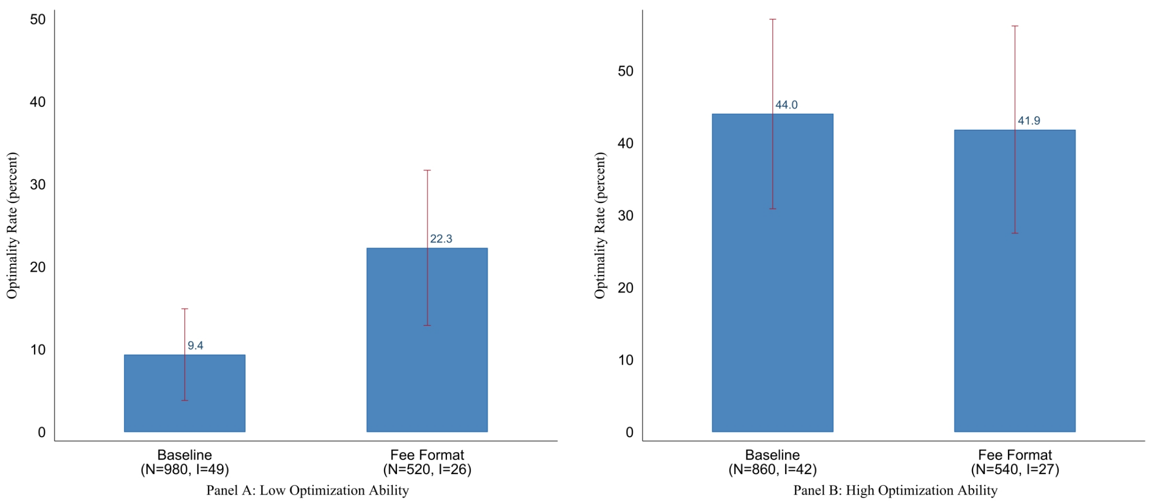

Notes: Panel A shows the optimality rate for **Baseline** and **Fee Format** treatments among subjects who fail to solve an algebraic version of the credit card repayment problem. Panel B shows the same rate across the same treatments for subjects who solve an algebraic version of the credit card repayment problem. The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Next, we disaggregate the data by our subjects' ability to pass a standard test of financial literacy in Figure 2.13, we find that the optimality rate among subjects who fail to solve at least one of the big three questions increases by 2.8 p.p. ($p = 0.74$). Whereas the optimality rate among subjects with "high financial literacy‘" increases by 6.3 p.p. ($p = 0.40$). Results are similar for misallocation rates. See Figure B1.7. Unlike optimization ability, the effect of Fee Format treatment does not seem to depend on subjects' level of financial literacy.

Figure 2.13: Effectiveness of Fee Format by Financial Literacy

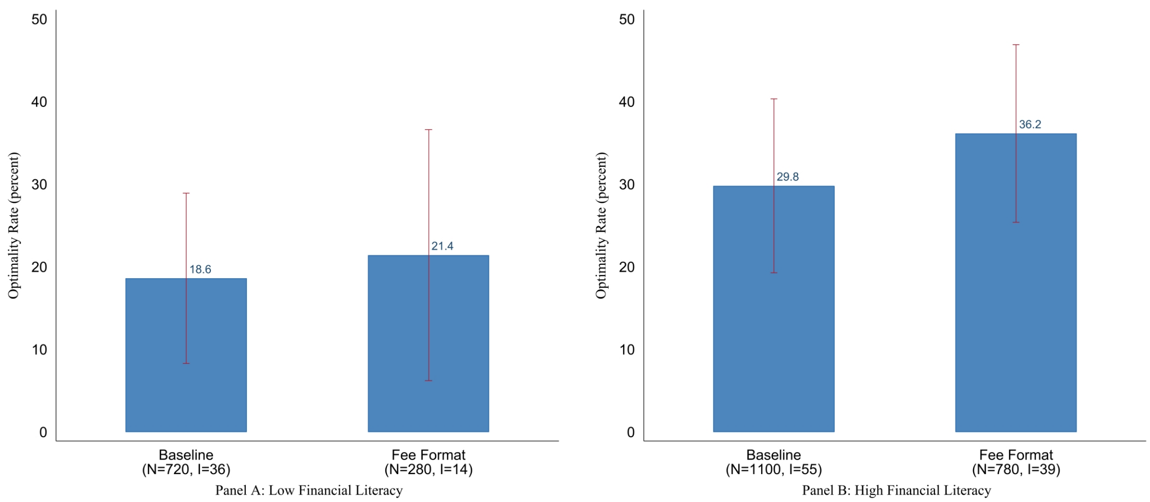

Notes: Panel A shows the optimality rate for **Baseline** and **Fee Format** treatments among subjects who fail to solve one of the big three financial literacy questions. Panel B shows the same rate across the same treatments for subjects who solve all big three financial literacy questions. The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Result 5 *Presenting the interest rate information in a fee format generates a modest increase in the share of optimal repayments. The intervention is most effective among subjects with low optimization ability. The effectiveness of intervention does not depend on subjects' level of financial literacy.*

2.4 Welfare Analysis Across Treatments

In this subsection, we compare the welfare gains from providing our subjects the opportunity to purchase advice and manipulating the information environment. The fact that the choice problem subjects face is a simple arbitrage situation that has an optimal allocation rule irrespective of subjects' preferences allows us to use the amount of deposit allocated towards the high interest rate card as a direct welfare measure.

Figure 2.14 documents the percentage of all payments that is allocated towards the high interest rate card and corresponds to exactly the percentage of the maximum earnings subjects can achieve in the experiment. Subjects who purchase advice significantly increases their gross earnings from the experiment by 11 p.p. compared to **Baseline** ($p = 0.0017$). Those who do not purchase advice do not see an increase in their earnings compared to **Baseline** ($p = 0.43$), neither subjects who participate in **Salience** ($p = 0.60$) nor **Fee Format** treatments ($p = 0.53$).

Figure 2.14: Gross Earnings Across Treatments

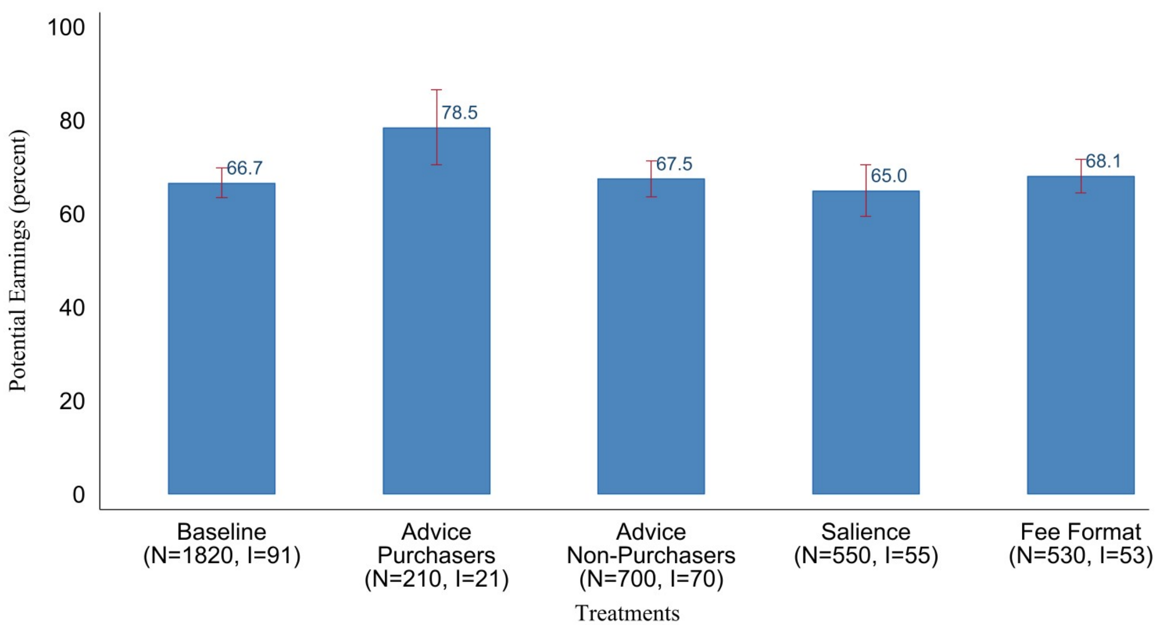

Notes: The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Figure 2.15 displays net earnings across treatments. Although purchasers of advice see an increase in their earnings, the above analysis does not take the percentage of earnings that are forgone to purchase advice. We see a dramatic change in earnings comparisons across treatments when we account for the cost of advice: subjects who purchase advice see a decrease in their earnings ($p = 0.047$) and become the worst performing group in the experiment in terms of earnings. This point emphasizes that an inability to use financial advice effectively generates a situation where subjects overpay for financial advice.

Figure 2.15: Net Earnings Across Treatments

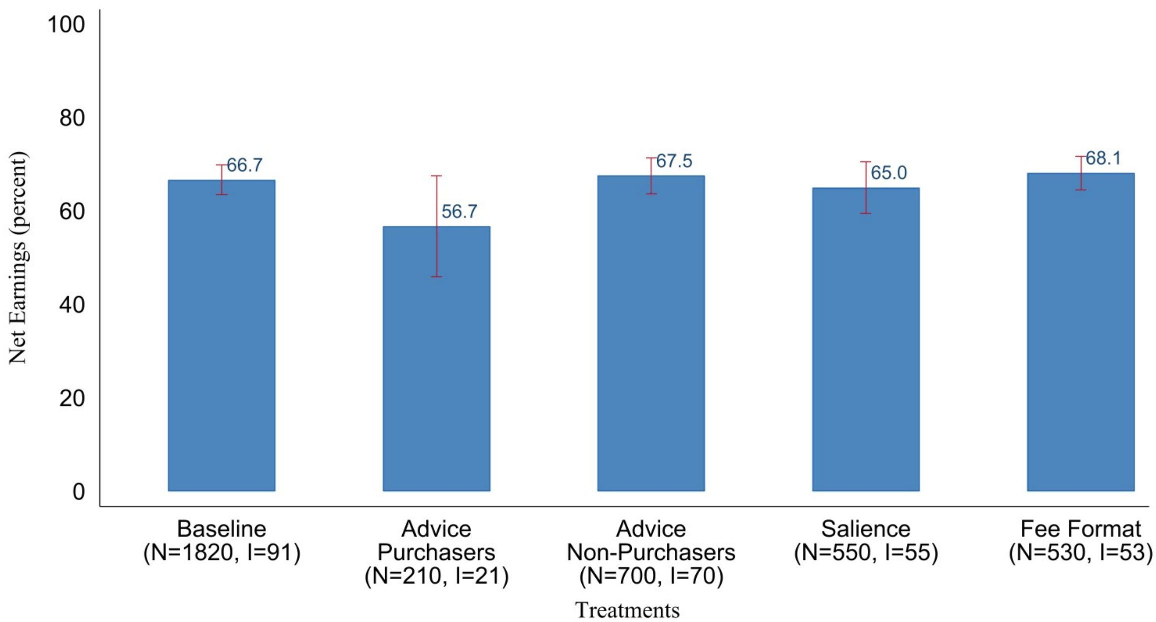

Notes: The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

2.5 Conclusion

This study compares the effectiveness of two different approaches to consumer protection policies on individuals' ability to exploit simple arbitrage opportunities in their repayments. We construct a simple repayment situation that mimics an online payment screen and measure people's willingness to pay for financial advice. We find that people are willing to pay for financial advice less than they should when compared to the amount of money they leave on the table by not purchasing advice. Once purchased, advice dramatically decreases the choice inefficiencies. One caveat is that people do not effectively use the advice that they have purchased. This creates a situation where subjects overpay for advice given how effectively they use the advice.

On the other hand, we manipulated the representation of interest rate information. We find that neither increasing the salience of interest rate nor changing the percentage format of the interest rate to a fee format achieves to reduce choice inefficiencies. Our results suggest that advice is a more effective tool in helping consumers tackle simple arbitrage opportunities on their balance sheets. Moreover, in addition to promoting consumer technology applications that would provide consumers with automated advice, subsidizing such applications should increase the welfare gains for consumers.

Chapter 3

Mental Models and Endogenous Learning

3.1 Introduction

People regularly construct mental models of their environment to guide their reasoning, to make inferences and to understand how their actions map into outcomes. Accumulating evidence from psychology and economics documents that people frequently create mental models that fall short of being an accurate representation of their decision environment, as they struggle gathering, attending, and processing crucial information in their environment ((71), (72)). An important question is then if the constructed models do not admit the possibility of truth, what are the implications of this "misspecification" for how people learn about their environment and subsequently make decisions that are informed by their learning? In particular, do more data necessarily mean that people will correctly learn a fundamental variable and take the first-best action?

We set out to answer this question using a carefully designed laboratory experiment. In the experiment, subjects repeatedly make investment decisions on a fixed project over

1000 periods, and in each period observe a noisy signal on the profit they generate. Subjects do not know the quality of their assigned project, but correctly know that the assignment is random. Moreover, there are complementarities between the investment amount and the expected project quality – projects with higher expected quality require higher investment at the optimum. Another key determinant of the profit is an ability parameter for the subject that positively affects the profit. In particular, we assign each subject an ability score based on their ranking on an "IQ test" that they take at the beginning of the experiment. Tying an ego-relevant ability parameter to the profit and not resolving the uncertainty over the ability parameter provides scope for model misspecification in how subjects perceive the profit equation. In our controlled environment, subjects frequently create misspecified models of the profit equation because of their overestimation of their ranking on the IQ test. We find that of 34% our subjects assign 100% likelihood to ability scores that are *strictly* above their true abilities.

The theory makes sharp predictions about how subjects with misspecified models of the profit equation should take actions throughout the experiment. (73) show that if agents with misspecified models do not update their prior on their ability and learn from feedback in a Bayesian fashion, they should on average make growingly suboptimal investments in this environment. Intuitively, this is because agents regularly experience less-than-expected profit because of their overconfidence and explain the less-than-expected output by developing pessimistic beliefs about the project quality. Consistent with this prediction, we find that overconfident subjects with misspecified models of the profit equation make growingly suboptimal investments throughout the experiment's time horizon. By the last period of the experiment, overconfident subjects significantly under-invest relative to the first-best in their projects compared to the subjects who have a correctly specified model of the profit equation.

A further theoretical insight in (73) that directly applies to our setting is that endogenous learning exacerbates the extent of suboptimal behavior when agents have misspecified models. This is because suboptimal behavior generated through model misspecification further depresses profit. Agents rationalize the additionally depressed profits by even more pessimistic beliefs about the project quality. These more pessimistic beliefs about the project quality then lead overconfident subjects to curb investment even further because of the complementarities in project qualities and investment amounts. In order to test this prediction we create two treatments **Exogenous** and **Endogenous** where we manipulate the endogeneity of feedback. While subjects in **Endogenous** see their investment decisions immediately implemented in each period and receive feedback that comes from the "profit distribution" they select in that period, subjects in **Exogenous** do not see their investment decisions immediately implemented and instead receive feedback from a fixed pre-announced "profit distribution" throughout the experiment. Although overconfident subjects in both treatments make growingly suboptimal investments, we do not find that endogenous feedback exacerbates the suboptimal investment behavior. Although subjects in **Exogenous** take actions that are statistically indistinguishable from the action a myopically optimizing Bayesian agent would take on average by the last period of the experiment, subjects in **Endogenous** sharply deviate from this Bayesian benchmark.

Investigating how subjects in **Exogenous** and **Endogenous** learn about their abilities reveals insights into the deviation from the theoretical prediction and the Bayesian benchmark. Although Bayesian learners would learn virtually nothing about their abilities in our experiment, a comparison of elicited prior and posterior means clearly indicate that overconfident subjects have become less overconfident by the end of the experiment. This "weakening" of mental models, including the truth within the set of possibilities, in the face of abundant objective feedback is consistent with previous work ((74)). Interestingly, we find that this "weakening" is more pronounced for subjects who face endogenous feedback although this effect is not significant at conventional levels.[1]

Substantial evidence in psychology suggests that on average people have unrealistically positive views of their traits (e.g. (75), (76), (77)). A large literature in economics investigates how such overconfident beliefs about own traits and prospects lead people to make suboptimal decisions (excess entry decisions in the laboratory, (78); over-trading by retail investors, (79); over-investment by CEOs, (80), (81)). As the evidence on the material costs of overconfidence accumulates, theoretical and laboratory studies started exploring how people produce and maintain overconfident beliefs about their abilities. While the theoretical literature documented how overconfidence may arise as a result of biased memory, ego utility , motivational and signalling values ((82), (83)), laboratory studies confirmed these mechanisms ((84), (85), (86)) and further documented how biased processing of objective noisy feedback prevent people from learning their true abilities ((87), (88), (89)).

Our investigation pushes the literature on overconfident agents' learning processes forward in three fundamental ways. First, while the previous literature focuses on settings where there is a single source of uncertainty (e.g. the agent's ability) that generates noisy feedback, we focus on a setting where the noisy feedback features two sources of uncertainty (e.g. the agent's ability and an external stable fundamental). Many economically important settings feature such multidimensional sources of uncertainty and a lower dimensional feedback. Examples include an employee not knowing her marginal return to effort, deciding how hard she wants to work and observing the output of her efforts; a team member not knowing the ability of his teammate, deciding how much of the work to delegate and observing the joint output produced by his team. Second, while the focus of the previous literature is on how overconfident people learn about their abilities through

[1]However, the greater reduction in overconfidence is consistent with our main finding that endogenous learning does not exacerbate suboptimal investment

noisy feedback, we focus on how overconfident people learn not about their ability, but about an external decision-relevant variable. There is ample evidence documenting that people are good at supplying overconfident beliefs about their own abilities, however, in the environments that we are investigating there is little evidence if the supply of such beliefs biases the way people learn about their environment. Third, while the previous literature is interested in exogenous learning situations where individuals learn about their abilities without choosing actions, our focus is on endogenous learning situations where individuals are provided opportunities to take actions that allow them to sample feedback from distinct distributions. It is clear that most learning environments have this "experimentation" feature (e.g. an employee might try working very hard or very little to get a more precise signal of her own ability) and hence, arguably, is of greater economic relevance than exogenous learning situations.

3.2 Theoretical Framework

In this section, we present our theoretical framework and discuss its main predictions. The model has two main goals. First, it illustrates an environment where overconfident agents who have a misspecified model of their environment underestimate an external fundamental while such underestimation does not necessarily arise for agents who have a correctly specified model. Second, it highlights how endogenous learning might exacerbate the extent of underestimation for overconfident agents.

3.2.1 Overview

We focus on a simple decision environment where the agent is uncertain about her ability and an external fundamental that affects the output she generates. The agent periodically takes an action and receives a noisy feedback on the output she generates.

More specifically, let $a \in A = \{20, 40, 60, 80, 100\}$ represent the agent's unchanging ability and $\phi \in= [0, 100]$ represent the unobservable unchanging fundamental. We assume the fundamental is randomly drawn from the uniform distribution $\pi_0 :\to \mathbb{R}_+$ before the agent starts making her decisions and is independent of the agent's ability a. The agent chooses an action $e_t \in E = [0, 100]$ in each period and produces an output $y(e_t, a, \phi)$ with partial derivatives $y_a \geq 0, y_\phi > 0$. In particular, we assume that the output has a simple functional form

$$y(e_t, a,) = (a + e_t) - \frac{e_t^2}{2}$$

After each action, the agent observes a noisy feedback f_t on the output she generates. The feedback is distributed Bernoulli with mean $\mu(e_t, a, \phi)$ that corresponds to the normalization of the output function:

$$\mu(e_t, a,) = \frac{y(e_t, a, \phi) - \underline{y}}{\bar{y} - \underline{y}}$$

where $\bar{y} = \max_{e,a,\phi} y(e, a, \phi)$ and $\underline{y} = \min_{e,a,\phi} y(e, a, \phi)$.

3.2.2 Objective Model

Fix an ability level for the agent, a_o, and her teammate, ϕ_o. For each effort level e_t, there is an objective feedback distribution $Q_o(\cdot|e_t)$ that is a Bernoulli density with mean $\mu(e_t, a_{o,o})$.

3.2.3 Mental Model

The mental model represents the set of feedback distributions the agent considers possible a priori. For a fixed objective decision problem Q_o, a mental model is a tuple

$$\mathcal{Q} = \langle \Theta, (Q_\theta)_{\theta \in \Theta} \rangle$$

where $\Theta \subset A \times$ is the agent's parameter set and $Q_\theta(\cdot|e_t)$ is the action-dependent feedback distributions parametrized by $\theta = (a, \phi) \in \Theta$. While the action-dependent objective feedback distribution $Q_o(\cdot|e_t)$ represents the true environment, the mental model represents the agent's perception of their environment.

We assume that the agent correctly believes that the map from actions to probability distributions over feedback is fixed and depends only on their current action, but they are uncertain about the distribution each action induces. The agent's uncertainty about what the true environment $Q_o(\cdot|e_t)$ is captured by their mental model $\langle \Theta, (q_\theta)_{\theta \in \Theta} \rangle$ and a joint density function $\Pi_0 : A \times \to \mathbb{R}_+$ that describes the agent's prior belief. Following the previous literature, we call the agent's mental model correctly specified if the true parameter vector lies in the support of the agent's prior beliefs ($\theta_o = (a_o, \phi_o) \in \Theta$), and otherwise call it misspecified.

We assume that the agent correctly believes that the fundamental is independently drawn from their own ability and from the uniform distribution π_0 with the support . Due to independence, we can decompose the agent's prior density $\Pi_0(a,) = p_0(a)\pi_0()$ where p_0 is a probability mass function that describes agent's prior belief about own ability. In this environment, we assume that the agent chooses myopically optimal actions at each period, aiming to maximize the probability of getting a "positive" feedback and learns from feedback using Bayes' rule.

3.2.4 Overconfidence as a Misspecified Mental Model

Since the agents in this framework have potentially misspecified models of their environment, we apply the solution concept proposed in (90) to derive the set of possible limit points of the agent's learning process.

The equilibrium requires agents' beliefs to put probability 1 on the set of subjective feedback distributions that are "closest" to the objective distribution. Building on (91), (90) shows that the correct notion of "distance" is the Kullback-Leibler divergence in statistics. It represents a "distance" between the objective output distribution $Q_o(\cdot|e)$ and the family of parametrized subjective distributions $(Q_\theta(\cdot|e))_{\theta\in\Theta}$ for a fixed action e:

$$K(e,\theta) = E_{Q_o(\cdot|e)} log\Big[\frac{Q_o(f|e)}{Q_\theta(f|e)}\Big]$$

Both objective and subjective mental models belonging to the family of Bernoulli distributions, the KL divergence in our context is simply

$$\begin{aligned} K(e,\theta) &= E_{Q_o(\cdot|e)}\Big[f log\frac{\mu(e,a_o,\phi_o)}{\mu(e,a,\phi)} + (1-f) log\frac{1-\mu(e,a_o,\phi_o)}{1-\mu(e,a,\phi)}\Big] \\ &= \mu(e,a_o,\phi_o) log\frac{\mu(e,a_o,\phi_o)}{\mu(e,a,\phi)} + (1-\mu(e,a_o,\phi_o)) log\frac{1-\mu(e,a_o,\phi_o)}{1-\mu(e,a,\phi)} \end{aligned}$$

The set of closest parameter values for the agent given an effort decision e can then be described as

$$\hat{\Theta}(e) =_{\theta\in\Theta} K(e,\theta)$$

The interpretation is that $\hat{\Theta}(e) \subset \Theta$ is the set of parameter values that the agent can believe to be possible after observing feedback consistent with the effort decision e.

A pure strategy Berk-Nash equilibrium of a single agent problem is then a pair of action and a belief (e^*, Π^*) that satisfies

i) $e^* \in_{e \in E} E_{\bar{Q}_{\Pi}(\cdot|e)} f$ where $\bar{Q}_{\Pi} = \int Q_{\theta} \Pi(d\theta)$

ii) $\Pi^* \in \Delta(\hat{\Theta}(e^*))$

Defining Overconfidence. In this environment, we define an overconfident agent to be one whose prior on their own ability assigns zero mass on their true ability and is supported by abilities that are greater than their own ability i.e. $a > a_o$ for any $a \in \text{supp} p_0(a)$. Hence an overconfident agent's mental models are misspecified as $\theta_o = (a_0, \phi_o) \notin \Theta$.

Exogenous Learning

We first look at how an overconfident agent learns the fundamental when their action is fixed and when they are provided with infinite feedback. Fix an effort decision $\bar{e}$. Assume that for each $a \in \text{supp} p_0(a)$, there is $\phi_a \in \text{supp} \pi_0(\phi)$ such that $\mu(\bar{e}, a_{o,o}) = \mu(\bar{e}, a, \phi_a)$. This assumption ensures that for any fixed action, the agent can always find a fundamental that explains the observed distribution of feedback irrespective of what they believe their own ability level to be. This implies KL divergence is minimized at 0 for all (a, ϕ_a) where $a \in \text{supp} p_0(a)$ and generates the following set of KL minimizers

$$\hat{\Theta}(\bar{e}) = \{(a, \phi_a = \frac{a_o + \bar{e}}{a + \bar{e}} \phi_o) | a \in \text{supp} p_0(a)\}$$

Suppose that the agent takes a fixed action $\bar{e}$ in all periods and for each $a \in \text{supp} p_0(a)$, there exists $\phi_a \in \text{supp} \pi_0(\phi)$ such that

$$\mu(\bar{e}, a_o, \phi_o) = \mu(\bar{e}, a, \phi_a)$$

Then the agent's beliefs on (a, ϕ) almost surely converges and concentrates on $\hat{\Theta}(\bar{e})$.

Proof: See (91). ■

Let $\Pi^{\bar{e}}_{\infty}$ be the limiting posterior distribution on (a, ϕ) when the agent repeatedly chooses effort $\bar{e}$. Then $\Pi_{\infty}(a, \phi) = p_0(a)\mathbb{1}_{\hat{\Theta}(\bar{e})}(a, \phi)$.

Proof: By Lemma 3.2.4, $\Pi^{\bar{e}}_{\infty}(a, \phi) = 0$ whenever $(a, \phi) \notin \hat{\Theta}(\bar{e})$. Moreover, $\hat{\Theta}(\bar{e})$ is finite since $p_o(a)$ has finite support. Take (a_1, ϕ_{a_1}) and $(a_2, \phi_{a_2}) \in \hat{\Theta}(\bar{e})$,

$$\begin{aligned}
\lim_{t\to\infty} \frac{\Pi_t(a_1, \phi_{a_1} \mid f_1, .., f_t)}{\Pi_t(a_2, \phi_{a_2} \mid f_1, .., f_t)} &= \lim_{t\to\infty} \frac{\Pi_0(a_1, \phi_{a_1})}{\Pi_0(a_2, \phi_{a_2})} \frac{Q_{\theta_1}(f_1, .., f_t \mid \bar{e})}{Q_{\theta_2}(f_1, .., f_t \mid \bar{e})} \\
&= \lim_{t\to\infty} \frac{\Pi_0(a_1, \phi_{a_1})}{\Pi_0(a_2, \phi_{a_2})} \frac{\mu(\bar{e}, a_1, \phi_{a_1})^{tf}(1 - \mu(\bar{e}, a_1, \phi_{a_1}))^{t(1-f)}}{\mu(\bar{e}, a_2, \phi_{a_2})^{tf}(1 - \mu(\bar{e}, a_2, \phi_{a_2}))^{t(1-f)}} \\
&= \lim_{t\to\infty} \frac{\Pi_0(a_1, \phi_{a_1})}{\Pi_0(a_2, \phi_{a_2})} \frac{\mu(\bar{e}, a_o, \phi_o)^{tf}(1 - \mu(\bar{e}, a_o, \phi_o))^{t(1-f)}}{\mu(\bar{e}, a_o, \phi_o)^{tf}(1 - \mu(\bar{e}, a_o, \phi_o))^{t(1-f)}} \\
&= \lim_{t\to\infty} \frac{\Pi_0(a_1, \phi_{a_1})}{\Pi_0(a_2, \phi_{a_2})} \\
&= \frac{\Pi_0(a_1, \phi_{a_1})}{\Pi_0(a_2, \phi_{a_2})} \\
&= \frac{p_0(a_1)\pi_0(\phi_{a_1})}{p_0(a_2)\pi_0(\phi_{a_2})} \\
&= \frac{p_0(a_1)}{p_0(a_2)}
\end{aligned}$$

Hence $\Pi^{\bar{e}}_{\infty}(a, \phi_a) = p_0(a)$ for each $(a, \phi_a) \in \hat{\Theta}(\bar{e})$ ■

A overconfident agent's learning process leads him to underestimate the fundamental i.e. whenever min supp$p_0(a) > a_o$, $E_{\Pi^{\bar{e}}_{\infty}}[\phi] < \phi_o$. *Proof:* For any $a \in$ supp$p_0(a)$, $\phi_a = \frac{a_o+\bar{e}}{a+\bar{e}}\phi_o < \phi_o$. Hence $E_{\Pi^{\bar{e}}_{\infty}}[\phi] < \phi_o$. ■

Endogenous Learning

We now look at what the overconfident agent comes to believe about the fundamental when he is allowed to change his ction in each period in response to his beliefs. In particular, we are interested in if his inferences about the fundamental improve when he chooses myopically optimal actions in each period. A Berk-Nash equilibrium for an overconfident agent with prior $p_0(a)$ is a pair $(\Pi_\infty(a,\phi), e^*)$ such that

$$i)\ \ e^* = E_{\Pi_\infty}[\phi] \tag{3.1}$$

$$ii)\ \ \Pi_\infty(a,\phi) = p_0(a)\mathbb{1}_{\hat{\Theta}^J(e^*)}(a,\phi) \tag{3.2}$$

where

$$\hat{\Theta}(e^*) = \{(a, \phi_a = \frac{a_o + e^*}{a + e^*}\phi_o)|a \in \text{supp}p_0(a)\} \tag{3.3}$$

Assume the equations (3.1) and (3.2) have a solution. Then $e^* = E_{\Pi_\infty}[\phi] < \phi_o$.

Proof: Note that $\frac{a_o+e^*}{a+e^*} < 1$ for any $a \in \text{supp}p_0(a)$. Then

$$e^* = E_{\Pi_\infty}[\phi] = \sum_{a \in \text{supp}p_0(a)} p_0(a)\frac{a_o + e^*}{a + e^*}\phi_o < \phi_o \tag{3.4}$$

■

The above lemma shows that the equilibrium action is lower than the fundamental.

Let $\Pi^{\bar{e}}_\infty$ be the limiting posterior distribution when the agent exogenously learns from data with fixed action $\bar{e}$ and let Π_∞ be the limiting posterior distribution when the agent learns endogenously with the optimal action. Then $E_{\Pi_\infty}[\phi] < E_{\Pi^{\bar{e}}_\infty}[\phi]$ whenever $\bar{e} \geq \phi_o$.

Proof: Note that for $\phi_a(\bar{e}) = \frac{a_o+\bar{e}}{a+\bar{e}}\phi_o$

$$sgn(\frac{\partial \phi_a}{\partial \bar{e}}) = sgn(a - a_o) > 0 \tag{3.5}$$

Hence for an overconfident agent (min supp$p_0(a) > a_o$), learning from a higher fixed effort decision leads to a higher belief on ϕ for each KL minimizing pair (a, ϕ_a). Since $e^* < \phi_o$, whenever $\bar{e} \geq \phi_o$

$$\frac{a_o + e^*}{a + e^*}\phi_o < \frac{a_o + \phi_o}{a + \phi_o}\phi_o \leq \frac{a_o + \bar{e}}{a + \bar{e}}\phi_o \quad \forall a \in \text{supp} p_0(a) \tag{3.6}$$

Therefore,

$$\sum_{a \in \text{supp} p_0(a)} p_0(a) \frac{a_o + e^*}{a + e^*}\phi_o < \sum_{a \in \text{supp} p_0(a)} p_0(a) \frac{a_o + \bar{e}}{a + \bar{e}}\phi_o \tag{3.7}$$

$$E_{\Pi_\infty}[\phi] < E_{\Pi_\infty^{\bar{e}}}[\phi] \tag{3.8}$$

■

Proposition 3.2.4 implies that if the agent starts off with an effort level at or above the optimal effort level, then the opportunity to change his effort decision in response to his inferences leads to more *incorrect* long-run expectations than if he could not change his effort decision. We call this type of learning *self-defeating.*

When $\bar{e} \geq \phi_o$, an overconfident agent's long-run optimal effort decision e^* is more inaccurate when he is allowed to change his effort decision in response to his inferences than a hypothetical optimal effort decision $\bar{e}^*$ he would like to take when he could not change his effort decision. *Proof:* Since the optimal effort level is equal to the expectation of ϕ under the joint feedback, we have $e^* = E_{\Pi_\infty}[\phi] < E_{\Pi_\infty^{\bar{e}}}[\phi] = \bar{e}^* < \phi_o$ ■

Intuitively, the overconfident agent is "surprised" by the negative feedback he observes when he collects sufficient data to identify the feedback distribution that he faces when he repeatedly takes an action greater than the first-best action, that is, the action he would take if he were to know the fundamental. The reason is that the feedback he receives increases in his ability and hence he expects higher feedback than actually realized. Once he identifies a feedback distribution that is lower than his expectation, he attributes the low output he generated to the fundamental being lower than his expectation. The beliefs that the agent develop lead him to exert lower than the first best action as his incentives to take higher actions increase in the fundamental. Since the action he takes increases the probability of receiving positive feedback when he chooses an action that is lower the first-best level, he decreases his probability of receiving positive feedback by choosing a lower action. This provides further negative feedback to the overconfident agent that "surprises" him, he explains these further negative feedback by lowering his expectations about the fundamental even further. This process continues until the overconfident agent is no longer "surprised" about the feedback he receives.

3.2.5 Hypotheses

Hypothesis 1. When the actions are exogenous and fixed at a level $\bar{e}$ greater than the first-best level and if beliefs converge:

- *an overconfident agent's expectation on the fundamental converges to a point that is less than the fundamental i.e.* $E_{\Pi_\infty^{\bar{e}}}[\phi] < \phi_o$

Our second hypothesis is that learning is *self-defeating* for overconfident agents i.e. when provided an opportunity to revise their actions in response to their inferences, overconfident agents' expectations are further away from the truth.

Hypothesis 2. When the actions are endogenous and if beliefs and actions converge:

- *an overconfident agent's expectation converges to* $E_{\Pi_\infty}[\phi]$ *that satisfies* $E_{\Pi_\infty}[\phi] < E_{\Pi_\infty^{\bar{e}}}[\phi] < \phi_o$

- *an overconfident agent's action converges to a point that is less than the first-best level*

Our third hypothesis involves learning about own ability.

Hypothesis 3. Irrespective of the endogeneity of actions, if beliefs converge:

- *an overconfident agent's expectation on his own ability converges to a point that is identical to his prior expectation on his own ability i.e.* $E_{\Pi_\infty}[a] = E_{\Pi_0}[a]$

3.2.6 Correctly Specified Mental Models

When the agent's prior about his own ability assigns some mass to his true ability a_o, his mental model is correctly specified. In this instance, the agent's beliefs do not need to converge to his true ability and the true fundamental as he can only exactly identify the feedback distribution he faces while unable to pin down the underlying parameters (a and ϕ) of that distribution. The predictions for such agents are ambiguous and prior-specific. The two examples below show that an agent who has a correctly specified mental model yet who expects his ability to be greater than his actual ability might 1) grow pessimistic or optimistic about the fundamental and 2) endogenous learning might either exacerbate or alleviate the extent of mislearning.

Example 1 (*Almost Misspecified*). Suppose the agent's prior on his own ability $p_0(a)$ is such that $a \geq a_o$ for any $a \in \mathrm{supp} p_0(a)$ with $p_0(a_o) \in (0,1)$. Clearly, $E_{p_0}[a] > a_o$. So the agent has a correctly specified model where he expects his ability to be strictly better than his actual ability. It is easy to see that Lemma 3.2.4 and Proposition 3.2.4 are still valid for this agent. Hence the agent is pessimistic about the fundamental when he learns under a fixed action and he exhibits self-defeating learning when he is allowed to change his action in response to his inferences.

Example 2 (*Almost Symmetric Prior Around* a_0). Suppose the agent's prior on his own ability $p_0(a)$ is such that supp$p_0(a) = \{a_+, a_o, a_-\}$ where $a_+ - a_o = a_o - a_-$ with $p_0(a_+) = \bar{p} + \varepsilon$, $p_0(a_o) = \bar{p}$, $p_0(a_-) = \bar{p} - \varepsilon$. Again, $E_{p_0}[a] > a_o$. Since the KL minimizing belief $\phi_a = \frac{a_o+\bar{e}}{a+\bar{e}}\phi_o$ is strictly decreasing and convex in a, $E_{\Pi^{\bar{e}}_{\infty}}[\phi] > \phi_o$. Hence the agent becomes optimistic about his teammate's ability when he learns under a fixed effort decision. We can use a little bit more algebra to show that the agent's expectations about his teammate move closer to the truth when he learns endogenously compared to the situation where he learns exogenously at the first-best effort level. Hence his endogenous learning is self-correcting.

3.3 Experimental Design

The objective of the design is to construct a decision environment in which i) agents are likely to form misspeficied mental models and ii) see if and how these mental models misguide learning about payoff-relevant decision variables.

Overview

The experiment consists of five parts. At the beginning of each part of the experiment, we provide subjects with the instructions, familiarize them with the interface and test their understanding of the rules of the experiment through a series of understanding quizzes. In the first part of the experiment, we measure the "ability" of our subjects using Raven matrices framed as an IQ test. In the second part of the experiment, we elicit subjects' beliefs about their relative performance on the IQ test compared against 19 randomly selected participants who participated in a pilot session.

The third part of the experiment is the main part where each subject is randomly assigned to one of our treatments. At the beginning of this part of the experiment, we assign subjects different fundamentals. More specifically, we frame the decision environment as subjects acting as project managers for a company where projects correspond to fundamentals.[2] Subjects' relative ranking on the IQ test and their assigned fundamentals jointly determine the probability of receiving a positive feedback. The assigned fundamental for each subject remains constant until the end of this part. In each period, subjects are required to submit actions framed as investment recommendations on their assigned projects. In order to help them with their decisions, subjects are provided with special calculators that take their beliefs about their own ability parameter "a" as an input and calculate the myopically optimal actions. In treatment **Exogenous**, subjects actions are *not* implemented to generate feedback but they are implemented for their payment. In treatment **Endogenous**, subjects' actions are implemented to generate feedback and also for the calculation of their payments.

In the fourth part of the experiment, we re-elicit subjects' beliefs about their ranking on the IQ test they have taken at the beginning of the experiment. In the fifth and the final part of the experiment, subjects complete a survey where they are asked to provide basic demographic information. In both treatments, subjects' payoffs are determined by the sum of the amount they made in a randomly selected part (either \$25 or \$0) and a show-up fee of \$10.

3.3.1 Part 1: Establishing Ability Parameters

The goal of this part of the experiment is to establish an ability parameter for each subject. We measure subjects' ability parameters using Raven's matrices. Subjects are

[2]Subjects' task is to recommend investment decisions (actions) to the company that is to be invested into their assigned projects (fundamentals). Their goal is to maximize their profit (output) from the project. In each period after they make an investment recommendation, they get an evaluation from the company if their profit for that period beats the company's profit expectations or not (Bernoulli feedback on the output). We choose to frame our decision environment to increase subjects' understanding of our relatively complicated decision environment (92).

introduced to Raven's matrices as a test of intelligence to accentuate the ego-relevance of the task and to provide scope for overconfidence. We ask each subject to solve the same 10 Raven's matrices, present them in the same order and provide the subjects with 10 minutes to finish the test. Once subjects finish the test, we compare the number of correctly answered questions to the performance of 19 randomly selected subjects who took the exact same IQ test in a pilot session of the experiment. Each subject is then assigned an "IQ rank score" depending on their ranking within their assigned group of 19 other participants with random tie-breaking. Specifically, a subject that ranks within the i^{th} quintile is assigned an IQ rank score of $20i$. We then establish subjects' true ability parameters in the main part of the experiment as their IQ rank scores, i.e. $a_0 = 20i$. We incentivize subjects by paying them $25 if a randomly selected answer in the IQ test is correct.

3.3.2 Part 2: Establishing Mental Models

The main goal of this part of the experiment is to elicit prior beliefs on own ability that we use to establish mental models in the third part of the experiment. In order to achieve this goal, we ask subjects how they think they rank in their randomly constructed group of 20 people based on their IQ test scores. The reason we choose to measure relative overconfidence (or "overplacement") rather than absolute overconfidence (or "overestimation") is that previous experiments find a greater scope for overconfidence when it is measured in relative terms.[3] Another important design choice is that we ask our subjects to state their beliefs over quintiles rather than the more conventional way of measuring overconfidence using 2-quantiles. The reason we ask our subjects to state their full belief distribution over quintiles is that we want to provide scope for model misspecification

[3] In particular, research in psychology documents that people "overplace" themselves in easy tasks ((93), (94)). We specifically choose the Raven matrices to benefit from this "easy" effect. Indeed, the average number of correct answers in our experiment is 6.78 out of 10.

while limiting the complexity of the belief elicitation procedure. The idea is that eliciting beliefs using smaller quantiles would create situations in which subjects predominantly assign positive probability to each quantile more frequently which would then render a majority of our subjects as correctly specified agents. On the other hand, if we elicit belief using larger quantiles, we complicate the belief elicitation procedure as subjects are required to state their full belief distribution over each quantile.

In order to simplify the belief elicitation procedure over quintiles, we use five sliders. Each quintile is associated with a slider. Subjects assign a total likelihood of 100% over five different quintiles through associated sliders at a precision of two decimal points. We use a standard incentive compatible mechanism to pay for the belief elicitation ((95)). A critical design choice here is that we elicit beliefs over quintiles as full belief distribution. Eliciting full belief distribution with high precision allows us to sharply draw a line between subjects with misspecified mental models and correctly specified mental models. In this regard, our belief elicitation procedure is a key element of the design as it allows us to strictly follow the theoretical conceptualization of overconfidence as a model misspecification.

3.3.3 Part 3: Learning Environment

After establishing an ability parameter for each subject and eliciting subjects' beliefs over their ability parameters, the only ingredient that is missing to construct a decision environment that is identical to our theoretical framework is the assignment of a fundamental to each subject. At the beginning of this part of the experiment, we randomly draw a fundamental for each subject using a discrete uniform distribution. The assignment of fundamental being independent from subjects' ability parameters is clearly communicated to subjects. Once subjects are assigned fundamentals, each subject faces an objective decision environment and they have mental models of their decision environments. As the researchers, we can observe both the objective and mental models of each subject.

Feedback Design and Minimizing Problems Related to Bayesianism

A crucial part of the experiment is the feedback that we provide to our subjects. Since our predictions are valid under Bayesian learning and the decision environment of our experiment is fairly complex, we help our subjects substantially make accurate inferences using the feedback regarding the fundamental. Consistent with the idea that people's learning about non-ego relevant variables is more in line with Bayesianism compared to learning about ego-relevant variables, we completely rule out that possibility that subjects' learning about the fundamental is inconsistent with Bayesian learning. We implement this critical feature of the design by providing subjects with a simple report that we frame as "the Statistician's Report" where we show subjects the Bayesian posterior mean of the fundamental conditional on each ability level. The report is updated in every period based on the feedback generated by the subject up until that period. Figure 3.1 presents an example of these reports.

Figure 3.1: The Statistician's Report

The Statistician's Report	
Your IQ Rank Score	**Project Quality**
20	75
40	71
60	69
80	66
100	65

Notes: Figure shows an example report. Each row shows how a particular IQ rank score corresponds to an expected project quality where the expectation is taken over the Bayesian posterior conditional on the IQ rank score.

Endogeneity of Feedback across Treatments

The only difference between our treatments **Exogenous** and **Endogenous** is if the feedback we provide to subjects is exogenous or endogenous to their actions. In treatment **Exogenous**, we provide subjects feedback based on the highest possible action,which is 100, not their actual actions. In contrast, we provide subjects feedback based on their actual actions in treatment **Endogenous**. We frame the lack of endogeneity of feedback to actions in treatment **Exogenous** as the company not being able to implement the subjects' recommended investment decisions immediately and instead investing an originally planned investment amount of 100 throughout the experiment. On the other hand, we frame the endogeneity of feedback in treatment **Endogenous** as the company implementing subjects' investment recommendations immediately instead of implementing their originally planned investment amount of 100. Note that we still mention the fact that there is an originally planned investment amount of 100 in treatment **Endogenous** to control for the potential anchoring effects.

The reason we choose the fixed action in treatment **Exogenous** as the highest possible action for each subject is two-fold. First, the predictions on self-defeating learning requires an exogenous action that is at or above the fundamental in our environment. Choosing the maximum possible action ensures that this requirement is satisfied irrespective of the realization of the fundamental. Second, the difference between predictions in **Exogenous** and **Endogenous** treatments in terms of expected fundamentals and chosen actions increases with the fixed action chosen in **Exogenous** treatment. Thus, choosing the highest possible action as the fixed action generates the largest possible treatment effect in theory.

Myopically Optimal Actions and the Calculator

A crucial element of the literature on learning with misspecified models is that agents take optimal actions (myopic or dynamic) in each period using their subjective expectations on the decision variables. In order to create an environment that allows our subjects to easily take myopically optimal actions, we choose a strictly concave output function that has a unique and simple optimal decision rule that only depends on the fundamental: "match your action to your expectation of the fundamental." We communicate this simple optimal action rule to our subjects as well as going through the details of how subjects can arrive at this conclusion on their own. We further test subjects' understanding of the optimal action rule through understanding quizzes.

We go a step further to make it even easier for our subjects to take myopically optimal actions by providing them with a calculator that takes their beliefs on their ability parameter as input and produces the myopically optimal action for that period as output. Hence for any subjective belief the subjects may have on their own ability parameter, they can accordingly calculate a myopically optimal action.

We attach the calculator to the Statistician's Report and ask subjects to enter their beliefs about their IQ rank score in corresponding rows. Once subjects enter their beliefs, the calculator produces the corresponding myopically optimal action using the Statistician's Report. More specifically, the calculator calculates a myopically optimal action by taking a weighted average of the expected project qualities with weights coming from the subjects' assigned likelihoods on each IQ rank score. Figure 3.2 provides an example.

Figure 3.2: Calculator

Calculator	The Statistician's Report	
Enter Likelihood (out of 100)	Your IQ Rank Score	Project Quality
	20	50
	40	50
	60	50
	80	50
	100	50
CALCULATE		

Notes: Figure shows an example report with the calculator attached to it. Subject are asked to enter their beliefs about their IQ rank score in corresponding rows as input. The calculator then calculates a myopically optimal action by taking a weighted average of the expected project qualities with weights coming from the subjects' assigned likelihoods on each IQ rank score.

We choose to help subjects calculate myopically optimal actions rather than dynamically optimal ones for several reasons. The first and main reason is that myopically optimal actions are significantly easier to explain to our subjects. Second, investigating how people learn under myopically optimal actions is of empirical relevance as previous research documents many instances where people narrowly bracket their decisions. Third, assuming that subjects' actions are consistent with myopic optimization allows us to interpret their choices reflecting their mean beliefs on the fundamental in our **Endogenous** treatment.[4]

[4]Note that subjects' actions in treatment **Exogenous** directly correspond to their expectation of the fundamental as the difference between myopic and dynamically optimal actions vanishes due to the fact that there is no scope for experimentation in treatment **Exogenous**, that is, subjects sample from the same distribution throughout the experiment irrespective of their actions.

Once a subject provides an input to the calculator and calculates an optimal action, a decision box appears where subjects are allowed to submit their actions. Figure 3.3 provides an example.

Figure 3.3: Decision Screen - Period 1

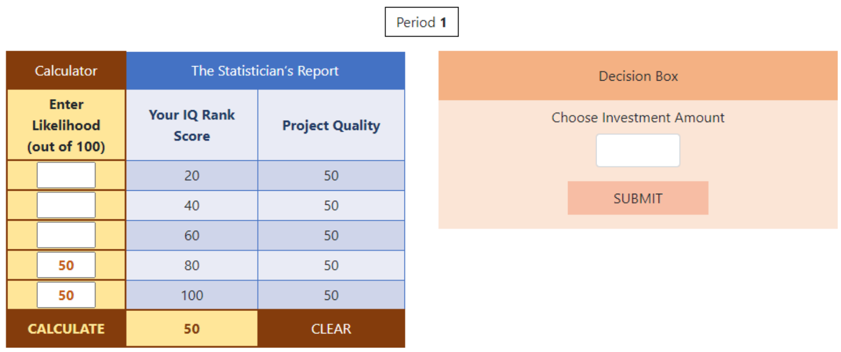

Notes: Figure shows an example decision screen. Once a subject provides an input to the calculator and calculates an optimal action, a decision box appears where subjects are allowed to submit their actions.

Although we require subjects to enter their beliefs about their IQ rank scores and calculate a myopically optimal action in each period, we still allow our subjects to submit actions that are distinct from what they obtain from the calculator as they might appreciate actions that are higher than the myopically optimal ones due to their informational value. The reason we require subjects to use the calculator in each period is to get a sense of the evolution of their beliefs on their own ability throughout the experiment.

The Amount of Opportunities to Learn

Since our goal is to investigate how subjects with misspecified models learn about an external fundamental, we design the experiment so that subjects have plenty of opportunities to take actions, generate feedback and learn from the feedback that they generate.

We create three connected subparts for this part of the experiment. Although the number of actions subjects take are similar in each subpart, the amount of feedback that is generated through the implemented action gradually increases.

Subpart 1: Periods 1 to 10

Subjects start this part of the experiment by taking 10 actions. After each action, subjects get a binary feedback on the implemented action.

Subpart 2: Periods 11 to 100

Starting from the 11th period up to 100th period, subjects take actions every 10 periods, that is, in periods 11,21,31,...,91. We count the actions subjects take in each of these periods towards the following 9 periods and provide aggregate feedback for every 10 periods. For instance, when a subject takes an action in period 11, the same action also counts as the action the subject takes for periods 12 to 20. The subject is then provided aggregate feedback on implemented actions from periods 11 to 20. Subjects take a total of 9 actions in this subpart and get feedback from 90 periods.

Subpart 3: Periods 101 to 1000

When subjects reach the 101st period, they start taking actions every 100 periods until period 1000, that is, in periods 101,201,301,...,901. Similar to the previous subpart, we count the actions subjects take in each of these periods towards the following 99 periods and provide aggregate feedback for every 100 periods. For instance, when a subject takes an action in period 101, the same action also counts as the action the subject takes for periods 102 to 200. The subject is then provided aggregate feedback on implemented actions from periods 101 to 200. Subjects take a total of 9 actions in this subpart and get feedback from 900 periods.

Subjects continuously move from the first subpart to the third subpart and are informed of the beginning of a new subpart along the way. The subpart structure we implement follows from earlier designs carefully studying learning ((74)) and allows us to generate a significant amount of feedback without increasing the duration of the experiment.

Subjects' Payments

We incentivize our subjects by paying them a fixed reward of $25 if the feedback in a randomly chosen period is positive in treatment **Endogenous**. In treatment **Exogenous**, we re-draw a feedback for each period that is generated through subjects' actions in the experiment. This creates an incentive compatible mechanism for subjects to take optimal actions in a manner that is equivalent to binarized scoring rule (95).

3.3.4 Part 4: Re-Examining Mental Models

In the fourth part of the experiment, we elicit subjects' beliefs about their ability parameter for a second time. Eliciting the full posterior belief distribution after subjects receive 1000 periods worth of feedback on their own ability allows us to answer if subjects retain their initial mental models or switch to alternative models.

The belief elicitation procedure is identical to the second part of the experiment. Subjects use five sliders to indicate their beliefs about which quintile their rank in their randomly constructed group of 20 people based on their IQ test scores. We use binarized scoring rule to incentivize subjects to truthfully report their beliefs about their IQ rank score.

3.3.5 Part 5: Exit Survey

We finalize the experiment by asking subjects control questions about their gender, their year of study, if they are enrolled in a STEM major and if they have taken a college-level statistics class.

3.3.6 Procedural Details

We conducted our experiment online using the subject pool of UCSB Experimental and Behavioral Economics Laboratory. The experiment was coded using o-Tree software ((69)). A total of 124 subjects, recruited through ORSEE (Online Recruitment System For Economic Experiments) (?). The average payment per subject was $27.6 including a $10 show-up fee. Each session lasted for 105 minutes.

3.4 Results

3.4.1 Identifying Misspecified Mental Models

We define an overconfident agent as one whose prior belief assigns zero mass on their true and all lower level IQ rank scores. Our clear identification strategy directly follows from the theoretical conceptualization of overconfidence as a misspecified mental model. Note that our conceptualization of overconfidence is more stringent than the typical conceptualization of overconfidence as having mean or median beliefs laying above the actual "ability" parameter. Figures 3.4 and 3.5 respectively display examples of overconfident subjects and subjects with correctly specified mental models.

Figure 3.4: Overconfidence as a Misspecified Mental Model

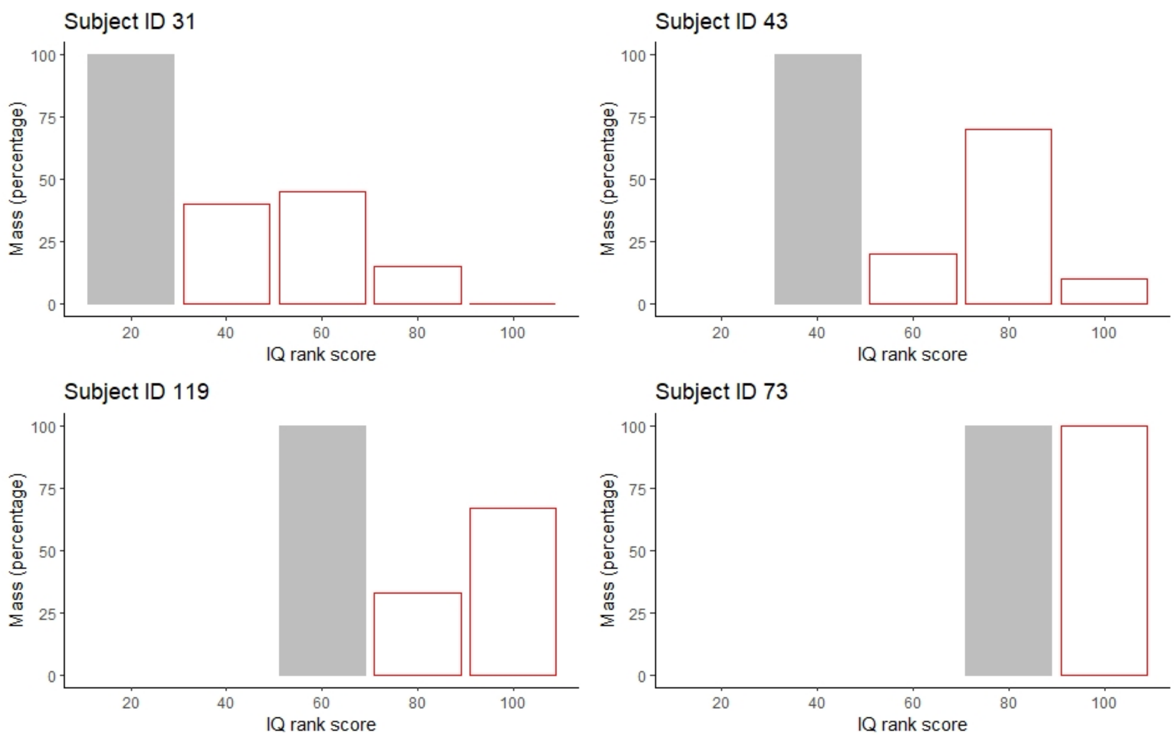

Notes: Figure shows the prior beliefs of selected subjects on their IQ rank scores. The grey bars display subjects' true IQ rank scores. The red bars display subjects' priors as probability mass functions.

We identify a total of 42 overconfident subjects (out of 124) distributed almost evenly across our two treatments **Exogenous** and **Endogenous**. The share of overconfident subjects in treatment **Exogenous** is 31.25% whereas the share of overconfident subjects in treatment **Endogenous** is 34.4%. The fact that we generate a significant amount of model misspecification through a simple task lends support to the main premise of the theoretical literature that people might form priors that sharply exclude the possibility of truth. We do not identify underconfidence as a model misspecification in our data, all remaining subjects in our experiment assigns some positive mass on their true IQ rank score.

Figure 3.5: Correctly Specified Mental Models

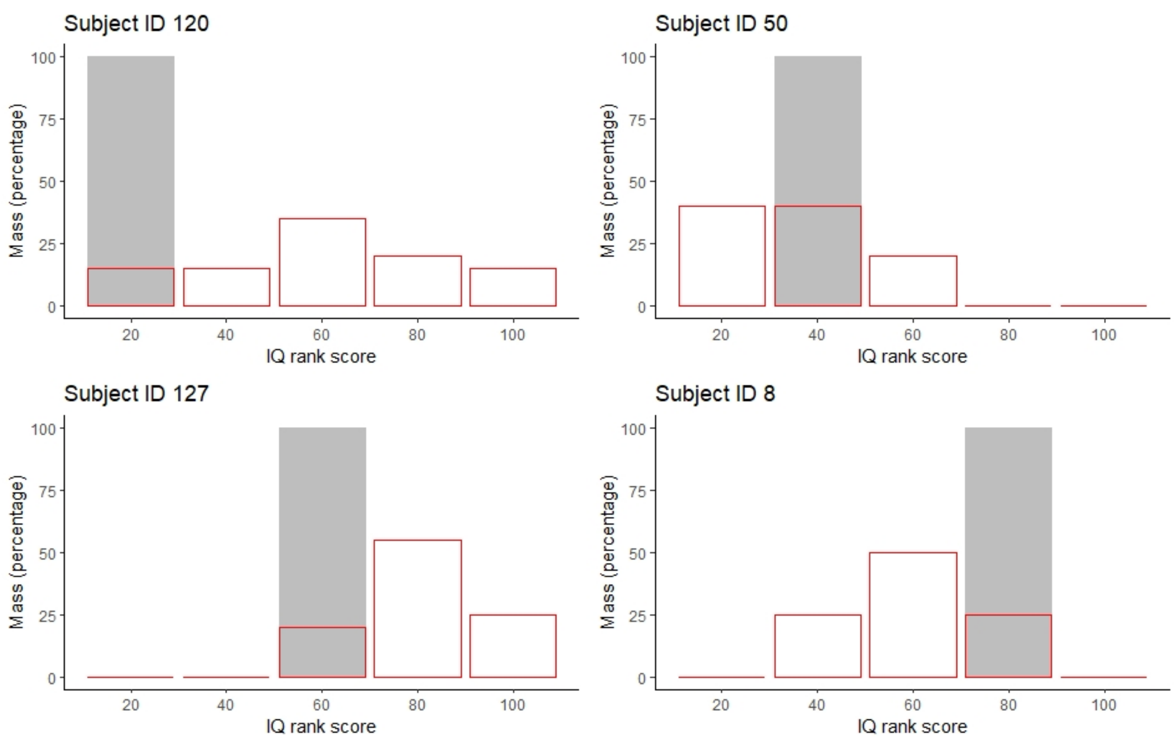

Notes: Figure shows the prior beliefs of selected subjects on their IQ rank scores. The grey bars display subjects' true IQ rank scores. The red bars display subjects' priors as probability mass functions.

3.4.2 Do Misspecified Mental Models Generate Suboptimal Behavior?

In this subsection, we aim to answer two questions. First, we ask if overconfident subjects' learning processes lead them to take suboptimal actions. Second, we ask how overconfident and correctly specified subjects' actions compare against the myopically optimizing Bayesian benchmark.

A First Look at How Overconfidence Generates Growingly Suboptimal Behavior

Figure 3.6 shows the evolution of actions for overconfident and correctly specified subjects when we aggregate the data from treatments **Exogenous** and **Endogenous**. Since subjects' fundamentals are drawn uniform randomly over integers from 0 to 100, we expect the fundamentals and hence the first-best optimal actions to average around 50. Indeed, Panel A of Figure 3.6 shows that average first-best optimal action, that is the action a subject would take if they were to know the true fundamental, for overconfident subjects is 53 and the average first-best optimal action for correctly specified subjects is 50.73. We find that correctly specified subjects' learning process do not lead them away from the first-best optimal action, the average action for correctly specified agents remain around the first-best optimal action throughout the experiment and moves closer to the first-best optimal action starting from period 11.

The behavior of overconfident subjects is dramatically different from correctly specified subjects. We find that overconfident subjects start out by taking higher actions than the first-best optimal actions reflecting their optimistic beliefs about the fundamental. After 10 periods, we find overconfident subjects start taking actions that are significantly lower than the first-best optimal action, reflecting their pessimistic beliefs about the fundamental and persistently keep doing so for the remainder of the experiment. By Period 901, we find a clear difference in behavior among overconfident and correctly specified subjects: overconfident subjects are on average taking actions that are significantly lower than optimal whereas correctly specified subjects on average are taking optimal actions.

Panel B of Figure 3.6 shows that the stark difference in behavior we observe is consistent with the theoretical predictions. When we simulate Bayesian learning with myopically optimal actions for each subject taking their mental models as given, we find that correctly specified subjects should take an average action of 49.05 in Period 901, whereas overconfident subjects should take an average action of 36.37. The difference in simulated average actions between correctly specified and overconfident subjects is significant ($p < 0.01$). We find that correctly specified agents take actions that are consistent with Bayesian learning. On the other hand, although overconfident subjects' actions are moving towards the simulated Bayesian action, the average action is still far away from the simulated Bayesian action.

Figure 3.6: Evolution of Actions

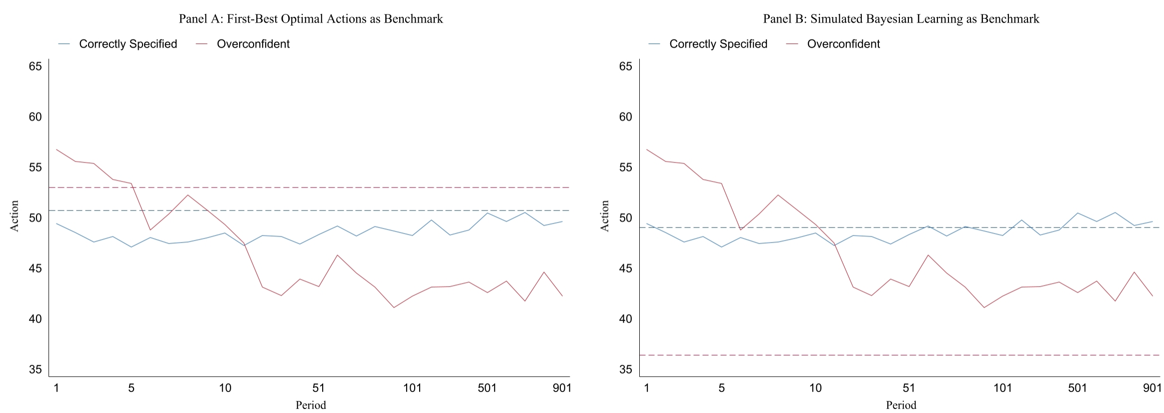

Notes: Both Panel A and Panel B show the average action for correctly specified and overconfident subjects across periods. Panel A presents the first-best optimal action as a benchmark. The blue dashed line in Panel A represents the average first-best optimal action for correctly specified subjects. The red dashed line in Panel A represents the average first-best optimal action for overconfident subjects. Panel B presents the average action a myopically optimizing Bayesian agent would take in the last period of the experiment as a benchmark. The simulations are conducted using each subject's prior beliefs about their abilities. The blue dashed line in Panel B represents the average simulated Bayesian action for correctly specified subjects whereas the red dashed line in Panel B represents the average simulated Bayesian action for overconfident subjects.

Statistical Differences in Learning with and without a Misspecified Model

We use a displacement measure $\Delta_{OPT} = e - e^*(\phi_0)$ to capture how far each action e is relative to the first-best optimal action. Note that this measure is positive for actions that are greater than the first-best optimal action and negative for actions that are smaller than the first-best optimal action. We then estimate the following displacement-from-benchmark regression in different periods: $\Delta_{OPT} = \alpha + \beta M + \varepsilon$ where α captures displacement from first-best optimal action for correctly specified subjects, M is a dummy variables that takes the value 1 for overconfident subjects, β captures the difference in displacement-from-benchmark for overconfident subjects and ε is an error term.

Table 3.1 presents the estimation results. First, note that correctly specified subjects do not significantly displace themselves from the first-best optimal action neither at the beginning nor towards the end of the experiment. On the other hand, we find that overconfident subjects start out positively yet not significantly displacing themselves from the first-best optimal action ($p = 0.48$). However, we find that initial positive displacement of overconfident subjects turn significantly and persistently negative in the later periods of the experiment. In particular, throughout the last 400 periods of the experiment, overconfident subjects displace themselves around 10 points away from the first-best optimal action. Figure C1.1 in Appendix C.1 presents these patterns in detail.

Table 3.1: Estimation of the Effect of Overconfidence on Displacement Relative to the First-Best Optimal Action

	Dependent Variable: Δ_{OPT}			
	(1)	(2)	(3)	(4)
β	5.087	-10.13**	-11.03***	-9.662**
	(6.348)	(3.231)	(3.032)	(2.887)
α	-1.326	-0.293	-0.217	-1.113
	(3.556)	(2.184)	(2.195)	(1.746)
Observations	128	128	128	128
Period	1	501	701	901

Notes: The table presents the average displacement relative to the first-best optimal action for correctly specified and overconfident agents. Each column conducts the estimation $\Delta_{OPT} = \alpha + \beta M + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action. The observations within periods aggregated across treatments **Exogenous** and **Endogenous**.

We use a second displacement measure $\Delta_{BAYES} = e - e^*(E_{\Pi_{sim}}[\phi])$, where Π_{sim} is the simulated posterior distribution on ϕ in the last period of the experiment, to capture how far each action e is relative to the simulated Bayesian action for the last period in the experiment.[5][6] We separately estimate a displacement-from-benchmark regression for correctly specified and overconfident subjects in different periods: $\Delta_{BAYES} = \alpha + \varepsilon$ where α captures displacement from the simulated Bayesian action for either correctly specified

[5]To be more precise, $e^*(E_{\Pi_{sim}}[\phi])$ is the action that a myopically optimizing Bayesian agent would take in the last period of the experiment.

[6]Similarly, this measure is positive for actions that are greater than the simulated Bayesian action and negative for actions that are smaller than the simulated Bayesian action.

or overconfident subjects, and ε is an error term.

Table 3.2: Estimation of Displacement Relative to the Simulated Bayesian Action

	Panel A: Correctly Specified				Panel B: Overconfident			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
α	0.357	1.389	1.465	0.570	20.40***	6.208***	5.389**	5.861*
	(3.427)	(1.955)	(1.965)	(1.475)	(4.296)	(1.687)	(1.567)	(2.214)
Observations	86	86	86	86	42	42	42	42
Period	1	501	701	901	1	501	701	901

Notes: The table presents the average displacement from the simulated Bayesian action relative to the last period for correctly specified and overconfident agents. Each column conducts the estimation $\Delta_{BAYES} = \alpha + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action. The observations within periods are aggregated across treatments **Exogenous** and **Endogenous**.

Table 3.2 presents the estimation results. Note that correctly specified subjects on average do not displace themselves from the simulated Bayesian action. On the other hand, overconfident subjects are systematically over the simulated Bayesian action. However, we see that the extent of displacement for the overconfident subjects is getting smaller as subjects move along the experiment's time horizon. Figure C1.2 in Appendix C.1 presents these patterns in detail.

We summarize our findings in this subsection in the following result:

Result 1 *Overconfident subjects' learning processes lead them to take growingly suboptimal actions throughout the experiment's time horizon. On the other hand, correctly specified subjects' learning processes do not generate a systematic deviation from the first-best optimal action throughout the experiment. Moreover, overconfident subjects' learning processes yield outcomes that are* ***less*** *suboptimal than the Bayesian prediction while correctly specified subjects' learning process is fully consistent with the Bayesian* *prediction.*

3.4.3 Does Endogenous Learning Exacerbate Suboptimal Behavior?

In this subsection, we aim to answer if endogenous learning exacerbate suboptimal behavior when the feedback subjects receive is endogenous to their actions. Theoretical predictions are such that endogenous learning should exacerbate overconfident agents' suboptimal behavior. For correctly specified agents, the theory's predictions are ambiguous. However, using simulations, we find that endogenous learning should not lead to a change in behavior for correctly specified subjects by the end of the experiment. We start this subsection by comparing the behavior of overconfident subjects in **Exogenous** and **Endogenous**. We then turn to correctly specified subjects and compare their behavior across treatments.

Behavior of Overconfident Subjects

Figure 3.7 shows the evolution of displacement relative to the first-best optimal action for overconfident subjects in **Exogenous** and **Endogenous**. We find that overconfident subjects in both treatments start out with actions that are close to the first-best optimal action. In both treatments, subjects exhibit negative displacement over time and we find that overconfident subjects in **Endogenous** exhibit greater negative displacement starting with Period 10. However, the difference in negative displacement vanishes towards the end of the experiment. Panel B provides further insights as to why we see the difference between the treatments vanish towards the end of the experiment. Overconfident subjects in **Exogenous** persistently move closer to the Bayesian prediction and meets the Bayesian prediction in the final period of the experiment whereas overconfident subjects in **Endogenous** decelerate their move towards the Bayesian prediction towards the end of the experiment.

Figure 3.7: Evolution of Displacement for Overconfident Subjects

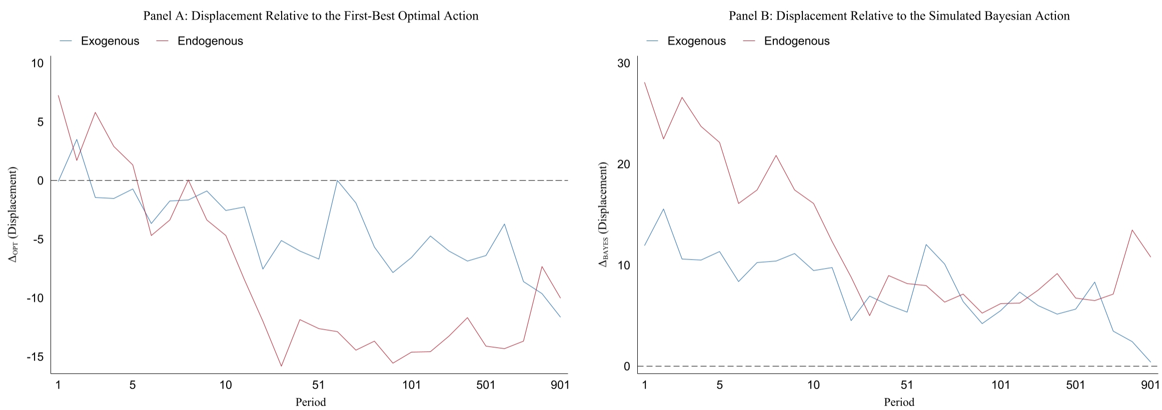

Notes: Panel A shows the average displacement relative to the first-best optimal action for overconfident subjects in treatment **Exogenous** and **Endogenous** across periods. Panel B shows the average displacement relative to the simulated Bayesian action for overconfident subjects in treatment **Exogenous** and **Endogenous** across periods. Each observation in a period corresponds to an individual action.

Table 3.3 provides the estimates of the treatment effect using displacement relative to the first-best optimal action as a benchmark. We estimate the regression $\Delta_{OPT} = \alpha + \theta T + \varepsilon$ where α captures the average displacement-from-benchmark for subjects in **Exogenous**, T is a dummy variable that takes the value 1 for subjects in **Endogenous**, θ captures the treatment effect and ε is an error term. The estimates for α across periods clearly show that overconfident subjects' actions in **Exogenous** growingly and significantly moves away from to the first-best optimal action exhibiting negative displacement from Period 1 to 901. On the other hand, the estimates for θ across periods show that there is no significant exacerbation of displacement for overconfident subjects in **Endogenous**. Table C1.1 provides statistical evidence that overconfident subjects in Exogenous act consistent with the Bayesian prediction by Period 901, while the behavior of overconfident subjects remain markedly different from the Bayesian prediction by Period 901.

Table 3.3: Estimation of the Treatment Effect for Overconfident Subjects

	Dependent Variable: Δ_{OPT}			
	(1)	(2)	(3)	(4)
θ	7.373	-7.683	-5.076	1.627
	(10.51)	(4.633)	(4.127)	(4.677)
α	-0.100	-6.404*	-8.588**	-11.63**
	(6.604)	(2.889)	(2.564)	(3.410)
Observations	42	42	42	42
Period	1	501	701	901

Notes: The table presents the average displacement relative to the first-best optimal action for overconfident agents. Each column conducts the estimation $\Delta_{OPT} = \alpha + \beta T + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action.

Behavior of Correctly Specified Subjects

Panel A of Figure 3.8 displays the evolution of displacement relative to the first-best optimal action for overconfident subjects in **Exogenous** and **Endogenous**. There is no discernible difference between the actions of correctly specified subjects in our treatments. We see that in both treatments behavior remains close to the first-best optimal benchmark throughout the experiment. Panel B of Figure 3.8 displays the evolution of displacement relative to the simulated Bayesian actions. Again, we do not see any systematic deviation from the Bayesian benchmark for the duration of the experiment.

Figure 3.8: Evolution of Displacement for Correctly Specified Subjects

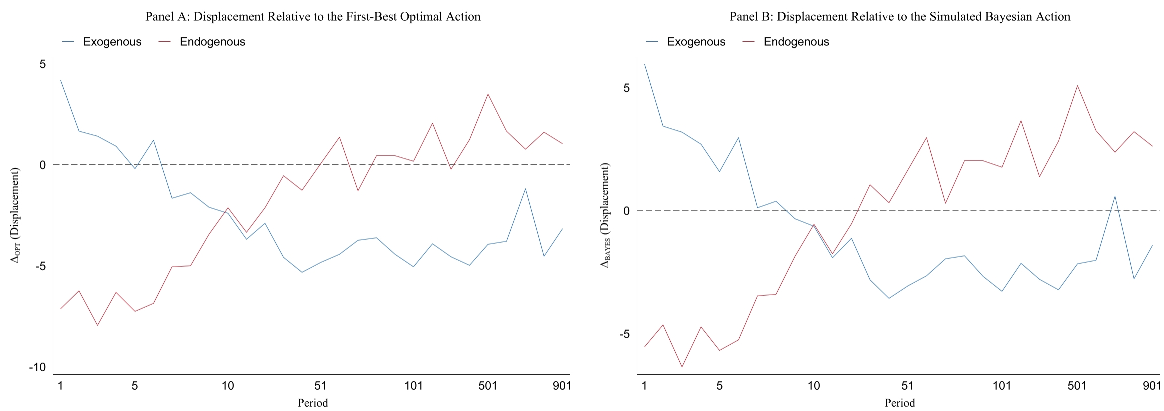

Notes: Panel A shows the average displacement relative to the first-best optimal action for correctly specified subjects in treatment **Exogenous** and **Endogenous** across periods. Panel B shows the average displacement relative to the simulated Bayesian action for overconfident subjects in treatment **Exogenous** and **Endogenous** across periods. Each observation in a period corresponds to an individual action.

Table 3.4 provides the estimates of the treatment effect using displacement relative to the first-best optimal action as a benchmark. We estimate the identical displacement-from-benchmark regression $\Delta_{OPT} = \alpha + \theta T + \varepsilon$ for correctly specified subjects. The estimates for α show that subjects in Exogenous exhibit negative displacement towards the end of the experiment although the magnitude of this move in each period is insignificant. The estimates for θ in each period indicate that subjects in **Endogenous** do not take significantly different actions compared to the subjects in **Exogenous**. Table C1.2 further documents that the subject behavior in both **Exogenous** and **Endogenous** are consistent with the simulated Bayesian action by the end of the experiment.

Table 3.4: Estimation of the Treatment Effect for Correctly Specified Subjects

	Dependent Variable: Δ_{OPT}			
	(1)	(2)	(3)	(4)
θ	-11.32	7.426	1.965	4.209
	(7.025)	(4.298)	(4.370)	(3.462)
α	4.205	-3.920	-1.177	-3.168
	(5.038)	(3.212)	(3.535)	(2.626)
Observations	86	86	86	86
Period	1	501	701	901

Notes: The table presents the average displacement relative to the first-best optimal action for correctly specified subjects. Each column conducts the estimation $\Delta_{OPT} = \alpha + \theta T + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action.

We summarize our findings from this subsection in the following result:

Result 2 *Contrary to the theoretical prediction, endogenous learning does not exacerbate the extent of suboptimal behavior for overconfident subjects. Similarly, although now consistent with the theory, we do not detect a change in behavior for correctly specified subjects when feedback is endogenous to their actions. Moreover, we find that overconfident subjects' behavior deviates from the Bayesian benchmark when feedback is endogenous but not when feedback is exogenous. On the other hand, correctly specified subjects behavior do not deviate from the Bayesian benchmark irrespective of the endogeneity of feedback.*

3.4.4 Learning About One's Self

We have so far investigated how subjects learn about the external decision variable in their environment. In this subsection, we turn to how subjects learn about their "ability" parameters. A Berk-Nash equilibrium of the single agent problem we investigate is one in which beliefs about the ability parameter concentrate on the prior where there is no self-learning as we have discussed earlier. Simulating a Bayesian learning model for the finite duration of our experiment yields posteriors consistent with this equilibrium by the end of the experiment, that is Period 1000. Simulations also confirm that there should be no self-learning for the duration of the experiment in subjects in **Exogenous**. We start this subsection by first looking at how overconfident subjects learn about their ability compared to correctly specified subjects at the aggregate. We then look at how endogenous learning affects learning about self for overconfident and correctly specified subjects.

Throughout this subsection, we use a displacement measure $\Delta_a = E_p[a] - a_0$ to capture how far each expected ability level $E_p[a]$ is relative to the true ability a_0 where the expectation is taken using the probability mass function p on a. Note that this measure is positive for beliefs that generate an expected ability level that is greater than the true ability level.

Differences in Self-Learning between Overconfident and Correctly Specified Subjects

Figure 3.9 presents a comparison of the displacement of expected abilities relative to the true ability using three different probability mass functions: subject's elicited *prior* on their ability, subject's elicited *posterior* on their ability and the simulated *Bayesian posterior* for the subject. Panel A shows that overconfident agents' expectations of their abilities move towards their true abilities after receiving 1000 periods worth of feedback. This move is significant at conventional levels ($p = 0.02$). The significant reduction in displacement relative to the true ability is also inconsistent with Bayesian learning. We find that subjects' posterior means average 5.96 lower than the simulated Bayesian posteriors ($p = 0.04$). Panel B documents that correctly specified agents' posterior expectations about their abilities do not significantly differ from their prior expectations ($p = 0.16$) or the Bayesian benchmark ($p = 0.21$).

Figure 3.9: Displacement of Expected Ability Relative to the True Ability

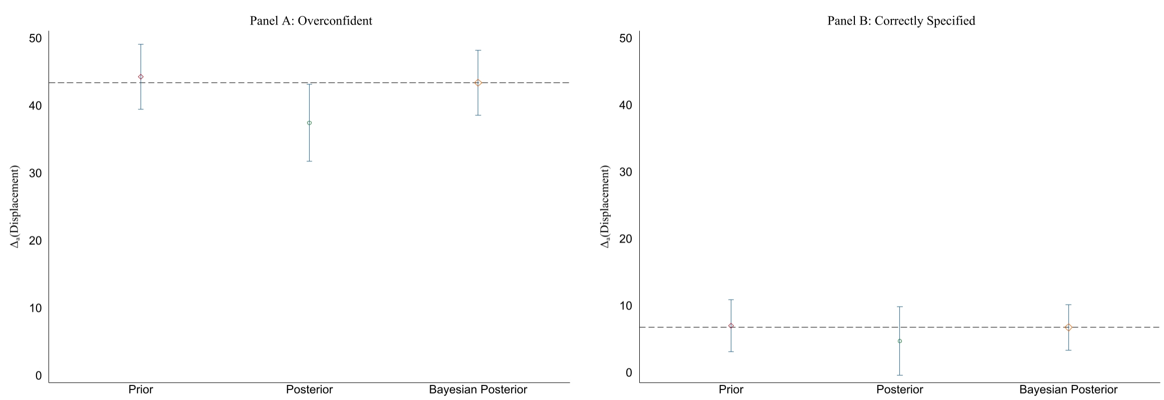

Notes: Figure shows the average displacement of expected ability levels relative to the true abilities using subjects' priors, posteriors, and simulated Bayesian posteriors. Panel A focuses on overconfident subjects where as Panel B focuses on correctly specified subjects. The black dashed line indicates the Bayesian benchmark. Whiskers indicate 95% confidence intervals.

One might then be curious if overconfident subjects move their expectations on their abilities towards their true abilities within the confines of their initial mental models. Table 3.5 provides evidence on how subjects' learning processes may lead them to completely switch their mental models. Although the majority of overconfident subjects stick with their initial mental models, we find that 22% of overconfident subjects end up assigning some probability to their true ability level after receiving feedback for 1000 periods. On the other hand, 13% of correctly specified subjects end up assigning no probability their true ability level at the end of their learning process.

Table 3.5: Switching Mental Models

		Posterior Models		
		Overconfident	Correctly Specified	Underconfident
Prior	Overconfident	78%	22%	0%
Models	Correctly Specified	13%	85%	2%

How Does Endogeneity of Feedback Affect Self-Learning for Overconfident Subjects?

According to the Bayesian benchmark, there is no difference in self-learning depending on the endogeneity of feedback. We expect overconfident subjects to exhibit virtually no self-learning in both **Exogenous** and **Endogenous** after 1000 periods. Panel A of Figure 3.10 documents that subjects in **Exogenous** somewhat learn their true ability after 1000 periods as the expected posterior beliefs show a smaller displacement from the true ability. The difference in prior and posterior means is insignificant ($p = 0.10$) and subjects' posterior mean is consistent with Bayesian posterior mean ($p = 0.17$). On the other hand, Panel B documents that posterior means of subjects in **Endogenous** considerably move towards their true abilities, yet the move is not significant at conventional levels ($p = 0.07$).

Figure 3.10: Displacement of Expected Ability Relative to the True Ability - Overconfident Subjects

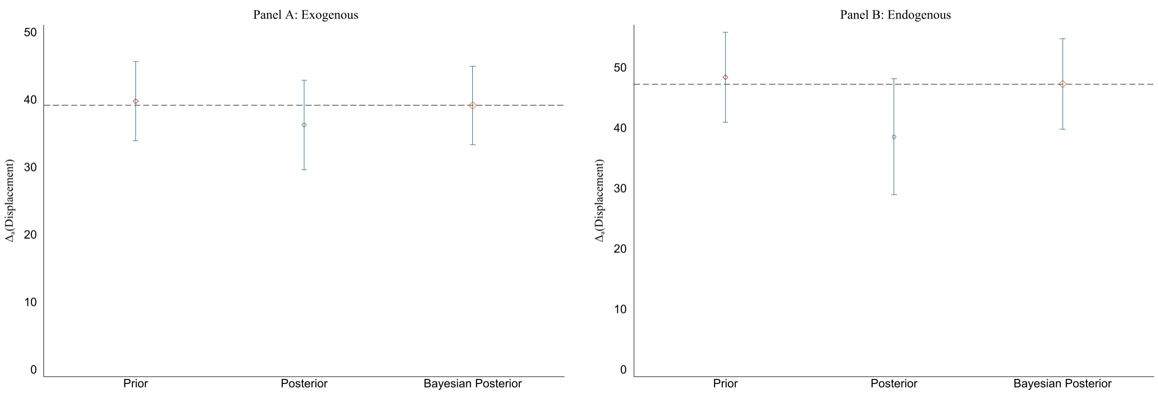

Notes: Figure shows the average displacement of expected ability levels relative to the true abilities using subjects' priors, posteriors, and simulated Bayesian posteriors. Panel A focuses on overconfident subjects in **Exogenous** where as Panel B focuses on overconfident subjects in **Endogenous**. The black dashed line indicates the Bayesian benchmark. Whiskers indicate 95% confidence intervals.

An important point that is worth emphasizing here is the increased self-learning with endogenous feedback is consistent with our earlier finding that endogenous learning does not exacerbate suboptimal behavior. If endogeneity of feedback leads subjects to better learn their own abilities, then subjects should take actions that are closer to the first-best optimal action in the main part of the experiment.

How Does Endogeneity of Feedback Affect Self-Learning for Correctly Specified Subjects?

As in the case of overconfidence, Bayesian learning does not predict subjects' self-learning to depend on the endogeneity of feedback. We expect no difference in mean prior and posterior beliefs for subjects in **Exogenous** and **Endogenous**. Panel A of Figure 3.11 confirms the Bayesian prediction for subjects in **Exogenous**: there is virtually no difference in prior and posterior means ($p = 0.92$). On the other hand, we find posterior means to move significantly closer to the subjects' true abilities in **Endogenous** ($p = 0.04$). The difference in posterior and the Bayesian posterior means in **Endogenous** is also significant ($p = 0.03$).

Figure 3.11: Displacement of Mean Beliefs Relative to the True Ability - Correctly Specified Subjects

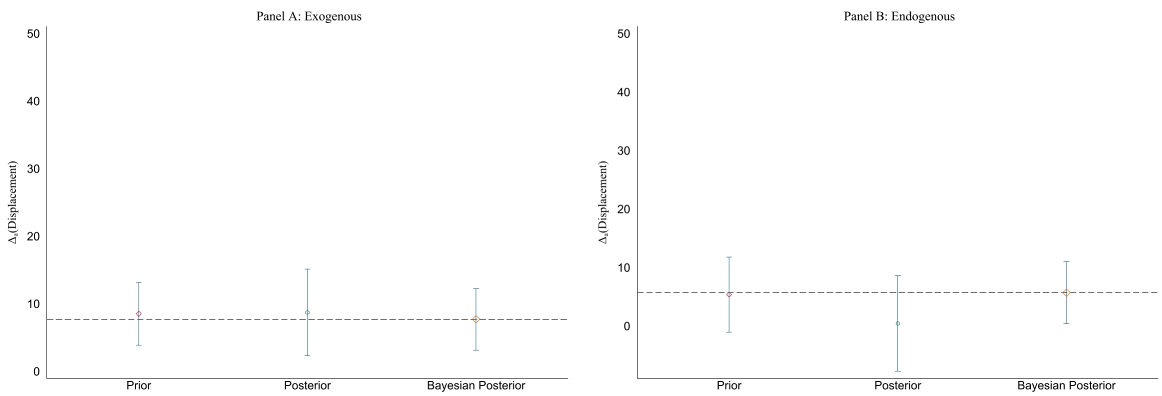

Notes: Figure shows the average displacement of expected ability levels relative to the true abilities using subjects' priors, posteriors, and simulated Bayesian posteriors. Panel A focuses on correctly specified subjects in **Exogenous** where as Panel B focuses on correctly specified subjects in **Endogenous**. The black dashed line indicates the Bayesian benchmark. Whiskers indicate 95% confidence intervals.

We summarize our findings from this subsection in the following result:

Result 3 *Inconsistent with Bayesian learning, overconfident subjects' posterior expectations about their ability significantly move closer to their true abilities. On the other hand, consistent with Bayesian learning, correctly specified subjects' prior and posterior expectations about their own ability do not differ. Moreover, both overconfident and correctly specified subjects exhibit greater self-learning when feedback is endogenous to their actions.*

3.5 Conclusion

In this paper we use people's tendency to hold optimistic beliefs about their abilities to generate model misspecification and investigate the implications of overconfidence as a misspecified mental model on learning about own ability and a fundamental. We find that overconfident subjects develop pessimistic beliefs about the fundamental and take growingly suboptimal actions. On the other hand, we find that endogenous feedback does not exacerbate the extent of suboptimal behavior: a result that is inconsistent with the theoretical prediction. When we look at how subjects learn about their own ability, we find that 1000 periods' worth of objective feedback lead some overconfident subjects to open their models to the possibility of truth. The "weakening" of mental models we observe is consistent with previous evidence. Complementing the nascent experimental literature on learning with misspecified mental models, we find that the "weakening" of mental models is more pronounced with endogenous feedback, explaining why endogenous feedback may not exacerbate the extent of suboptimal behavior.

Appendix A

Appendix for *The Debt Payment Puzzle*

A.1 Additional Results

Table A1.1: Differences in Optimality Measures Across Debt Treatments

	Optimality Rate			Correct Allocation Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
Debt Interest Rate	0.0342	-0.0303	-0.0121	-13.41	-20.90	-4.441
	(0.0810)	(0.0848)	(0.0579)	(17.98)	(18.23)	(12.89)
Constant	0.224	0.251	0.188	332.4	341.4	318.5
	(0.0583)	(0.0688)	(0.0435)	(12.76)	(15.20)	(10.29)
Observations	387	1573	2605	387	1573	2605
R^2	0.002	0.001	0.000	0.002	0.006	0.000
Period	First	All	All	First	All	All
Restrict to Optimizers	Yes	Yes	No	Yes	Yes	No
Restrict to Interest Rate Acquirers	Yes	Yes	No	Yes	Yes	No

Note: Each column reports the effect of being assigned to *Debt Interest Rate* treatment on some optimality measure using an OLS regression. In Columns 1,2 and 3, the dependent variable is a dummy that takes the value 1 if the allocation made is optimal. In Columns 3,4 and 5, the dependent variable is the amount of allocation made to the high interest rate card which takes a value between 0 and 500. Columns 1 and 4 restrict the sample to observations from the first period in each stage where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 2 and 5 restrict the sample to observations where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 3 and 6 execute the same analysis without imposing any sample restrictions. Standard errors in parentheses. Errors are clustered at the subject level.

Table A1.2: Differences in Optimality Measures Across Debt Treatments with Demographic Controls

	Optimality Rate			Correct Allocation Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
DR	-0.0121	-0.00259	-0.0742	-4.441	-2.351	-13.76
	(0.0579)	(0.0531)	(0.0512)	(12.89)	(11.87)	(12.50)
Math Score		0.265	0.126		58.36	36.65
		(0.0760)	(0.0597)		(16.84)	(13.98)
Gender			-0.180			-30.81
			(0.0649)			(15.42)
STEM/Economics			0.163			22.55
			(0.0532)			(12.86)
Constant	0.188	0.0563	0.213	318.5	289.6	317.3
	(0.0435)	(0.0392)	(0.0709)	(10.29)	(9.457)	(17.29)
Observations	2605	2605	2605	2605	2605	2605

Note: Column 1 to 3 represent the differences in the share of optimal allocations between *Debt Balance* and *Debt Interest Rate* treatments. The dependent variable *Optimal* is a dummy variable that takes the value 1 if the allocation is made optimally. Column 4 to 6 represent the differences in the amount of correctly made allocations between **DB** and **DR**. The dependent variable is the amount of allocation made on the high interest rate card which takes a value in between 0 and 500. The unit of observation is *subject x period*. The term **DR** is a dummy variable that takes the value 1 for observations made under Debt Interest Rate treatment. *Math Score* is a discrete variable that takes values [0,0.25,0.5,0.75,1] representing the percentage of correct answers to four optimization problems. *Gender* is a dummy variable that takes the value 1 for female subjects. *STEM/Economics* is a dummy variable that takes the value 1 for subjects whose majors are either STEM or Economics. Standard errors in parentheses. Errors are clustered at the subject level.

Table A1.3: Estimation of Repayments Across Debt Treatments

	(1)	(2)	(3)	(4)	(5)	(6)
Higher Interest Rate	137.8	151.0	135.2	164.0	184.5	140.7
	(25.40)	(21.79)	(16.80)	(25.67)	(31.58)	(21.27)
Higher Balance	182.8	147.6	136.7	109.7	80.83	91.95
	(25.65)	(15.32)	(12.24)	(16.61)	(16.78)	(14.97)
DR x Higher Interest Rate				-26.21	-33.47	-5.442
				(35.98)	(38.26)	(27.05)
DR x Higher Balance				73.09	66.80	44.75
				(30.39)	(22.64)	(19.29)
DR				-24.23	-4.473	-13.70
				(21.19)	(20.45)	(16.85)
Constant	93.01	106.9	118.5	117.2	111.4	132.2
	(16.06)	(12.40)	(10.10)	(13.96)	(16.34)	(13.53)
Observations	186	928	1288	387	1573	2605
R^2	0.477	0.445	0.433	0.452	0.430	0.370
Period	First	All	All	First	All	All
Restrict to Optimizers	Yes	Yes	No	Yes	Yes	No
Restrict to Interest Rate Acquirers	Yes	Yes	No	Yes	Yes	No
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.27$	$p = 0.90$	$p = 0.94$			
$\beta_{\mathbf{DR}xHigherInterestRate} = 0$				$p = 0.47$	$p = 0.39$	$p = 0.84$
$\beta_{\mathbf{DR}xHigherBalance} = 0$				$p = 0.02$	$p = 0.0044$	$p = 0.02$

Note: Columns 1 to 3 estimate, using OLS, how having a higher interest rate and a higher balance on a card affects the allocations made towards that card in *Debt Interest Rate* treatment. The dependent variable is the amount of allocation made on the left card (without loss of generality) which takes a value in between 0 and 500. The regressors *Higher Interest Rate* and *Higher Balance* are two dummy variables that takes the value 1 whenever the interest rate and the balance on the left card, respectively, is higher compared to the right card. Columns 4 to 6 estimate, using OLS, how having a higher interest rate and a higher balance on a card affect the allocations made towards that card using observations from both *Debt Interest Rate* and *Debt Balance* treatments. The term **DR** is a dummy variable that takes the value 1 if the allocation is made under *Debt Interest Rate* treatment. The terms **DR** x Higher Interest Rate and **DR** x Higher Balance are interaction variables. *Period* indicates if the analysis is limited to the first period decisions or not. *Restrict to Optimizers* indicate if the analysis is limited to subjects who can solve optimization problems. *Restrict to Interest Rate Acquirers* indicate if the analysis is limited to observations where the subjects acquired interest rate information before making their decisions. The last part of the table reports the parametric test results on estimated coefficients through associated p-values. Standard errors in parentheses. Errors are clustered at the subject level.

Table A1.4: Estimation of Repayments Across Debt Treatments with Demographic Controls

	(1)	(2)	(3)	(4)	(5)	(6)
Higher Interest Rate	137.4	150.5	135.1	164.5	184.7	140.9
	(25.36)	(21.41)	(16.57)	(25.70)	(31.64)	(21.27)
Higher Balance	182.8	147.7	136.8	109.8	80.92	91.95
	(25.79)	(15.29)	(12.23)	(16.65)	(16.78)	(14.97)
Gender	-12.68	-19.39	-8.734	13.26	0.723	-1.794
	(13.20)	(7.063)	(7.937)	(12.27)	(9.171)	(7.340)
STEM/Economics	-8.656	-3.355	8.926	11.66	4.803	7.012
	(13.32)	(7.674)	(7.837)	(10.64)	(8.624)	(5.495)
DR x Higher Interest Rate				-26.38	-33.72	-5.477
				(36.10)	(38.29)	(26.95)
DR x Higher Balance				73.00	66.70	44.78
				(30.46)	(22.63)	(19.29)
DR				-22.06	-4.875	-15.15
				(21.54)	(21.04)	(17.45)
Constant	104.8	119.1	119.1	101.3	108.6	131.0
	(21.16)	(13.01)	(13.10)	(21.32)	(19.00)	(15.61)
Observations	186	928	1288	387	1573	2605
R^2	0.479	0.449	0.435	0.453	0.430	0.370
Period	First	All	All	First	All	All
Restrict to Optimizers	Yes	Yes	No	Yes	Yes	No
Restrict to Interest Rate Acquirers	Yes	Yes	No	Yes	Yes	No
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.27$	$p = 0.92$	$p = 0.93$			
$\beta_{\mathbf{DR}xHigherInterestRate} = 0$				$p = 0.47$	$p = 0.38$	$p = 0.84$
$\beta_{\mathbf{DR}xHigherBalance} = 0$				$p = 0.02$	$p = 0.0044$	$p = 0.02$

Note: The table executes the analysis in Table A1.3 with demographic controls. *Gender* is a dummy variable that takes the value 1 for female subjects. *STEM/Economics* is a dummy variable that takes the value 1 for subjects whose majors are either STEM or Economics. Standard errors in parentheses. Errors are clustered at the subject level.

Table A1.5: Differences in Optimality Measures Across Balance Treatments

	Optimality Rate			Correct Allocation Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
Investment Balance	0.237	0.247	0.242	46.14	50.13	48.19
	(0.0960)	(0.110)	(0.0759)	(21.78)	(23.91)	(18.53)
Constant	0.224	0.251	0.188	332.4	341.4	318.5
	(0.0583)	(0.0689)	(0.0435)	(12.77)	(15.21)	(10.29)
Observations	353	1095	2452	353	1095	2452
R^2	0.063	0.065	0.069	0.026	0.031	0.028
Period	First	All	All	First	All	All
Restrict to Optimizers	Yes	Yes	No	Yes	Yes	No
Restrict to Interest Rate Acquirers	Yes	Yes	No	Yes	Yes	No

Note: Each column reports the effect of being assigned to *Investment Balance* treatment on some optimality measure using an OLS regression. In Columns 1,2 and 3, the dependent variable is a dummy that takes the value 1 if the allocation made is optimal. In Columns 3,4 and 5, the dependent variable is the amount of allocation made to the high interest rate account which takes a value between 0 and 500. Columns 1 and 4 restrict the sample to observations from the first period in each stage where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 2 and 5 restrict the sample to observations where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 3 and 6 execute the same analysis without imposing any sample restrictions. Standard errors in parentheses. Errors are clustered at the subject level.

Table A1.6: Differences in Optimality Measures Across Balance Treatments with Demographic Controls

	Optimality Rate			Correct Allocation Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
IB	0.242	0.268	0.162	48.19	55.54	40.24
	(0.0759)	(0.0725)	(0.0868)	(18.53)	(17.47)	(20.13)
Math Score		0.221	0.191		61.97	59.84
		(0.0964)	(0.0854)		(24.05)	(23.05)
Gender			-0.287			-49.35
			(0.0874)			(21.66)
STEM/Economics			0.0432			-3.269
			(0.0771)			(18.99)
Constant	0.188	0.0643	0.300	318.5	283.9	326.7
	(0.0435)	(0.0520)	(0.0969)	(10.29)	(13.89)	(23.08)
Observations	2452	2452	2452	2452	2452	2452
R^2	0.069	0.102	0.182	0.028	0.053	0.076

Note: Column 1 to 3 represent the differences in the share of optimal allocations between *Debt Balance* and *Investment Balance* treatments. The dependent variable *Optimal* is a dummy variable that takes the value 1 if the allocation is made optimally. Column 4 to 6 represent the differences in the amount of correctly made allocations between **DB** and **IB**. The dependent variable is the amount of allocation made on the high interest rate card which takes a value in between 0 and 500. The unit of observation is *subject x period.* The term **IB** is a dummy variable that takes the value 1 for observations made under Debt Interest treatment. *Math Score* is a discrete variable that takes values [0,0.25,0.5,0.75,1] representing the percentage of correct answers to four optimization problems. *Gender* is a dummy variable that takes the value 1 for female subjects. *STEM/Economics* is a dummy variable that takes the value 1 for subjects whose majors are either STEM or Economics. Standard errors in parentheses. Errors are clustered at the subject level.

Table A1.7: Estimation of Repayments Across Balance Treatments

	(1)	(2)	(3)	(4)	(5)	(6)
Higher Interest Rate	255.1	271.1	221.3	164.0	184.5	140.7
	(35.20)	(34.66)	(29.02)	(25.70)	(31.62)	(21.28)
Higher Balance	62.05	71.58	90.21	109.7	80.83	91.95
	(33.79)	(27.56)	(23.99)	(16.63)	(16.81)	(14.97)
IB x Higher Interest Rate				91.08	86.64	80.63
				(43.27)	(46.62)	(35.82)
IB x Higher Balance				-47.62	-9.246	-1.741
				(37.33)	(32.02)	(28.14)
IB				-10.44	-16.79	-25.69
				(29.68)	(33.27)	(23.84)
Constant	106.8	94.58	106.5	117.2	111.4	132.2
	(26.47)	(29.30)	(19.76)	(13.98)	(16.36)	(13.54)
Observations	152	450	1135	353	1095	2452
R^2	0.430	0.502	0.414	0.428	0.461	0.374
Period	First	All	All	First	All	All
Restrict to Optimizers	Yes	Yes	No	Yes	Yes	No
Restrict to Interest Rate Acquirers	Yes	Yes	No	Yes	Yes	No
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.01$	$p = 0.0001$	$p = 0.0027$			
$\beta_{\mathbf{IB}xHigherInterestRate} = 0$				$p = 0.04$	$p = 0.07$	$p = 0.03$
$\beta_{\mathbf{IB}xHigherBalance} = 0$				$p = 0.21$	$p = 0.77$	$p = 0.95$

Note: Columns 1 to 3 estimate, using OLS, how having a higher interest rate and a higher balance on a fund affects the allocations made towards that fund in *Investment Balance* treatment. The dependent variable is the amount of allocation made on the left fund (without loss of generality) which takes a value in between 0 and 500. The regressors *Higher Interest Rate* and *Higher Balance* are two dummy variables that takes the value 1 whenever the interest rate and the balance on the left fund, respectively, is higher compared to the right account. Columns 4 to 6 estimate, using OLS, how having a higher interest rate and a higher balance on an account affect the allocations made towards that account using observations from both *Investment Balance* and *Debt Balance* treatments. The term **IB** is a dummy variable that takes the value 1 if the allocation is made under *Investment Balance* treatment. The terms **IB** x Higher Interest Rate and **IB** x Higher Balance are interaction variables. *Period* indicates if the analysis is limited to the first period decisions or not. *Restrict to Optimizers* indicate if the analysis is limited to subjects who can solve optimization problems. *Restrict to Interest Rate Acquirers* indicate if the analysis is limited to observations where the subjects acquired interest rate information before making their decisions. The last part of the table reports the parametric test results on estimated coefficients through associated p-values. Standard errors in parentheses. Errors are clustered at the subject level.

Table A1.8: Estimation of Repayments Across Balance Treatments with Demographic Controls

	(1)	(2)	(3)	(4)	(5)	(6)
Higher Interest Rate	253.6	271.4	220.5	164.9	184.7	140.9
	(35.18)	(33.80)	(28.93)	(25.69)	(31.44)	(21.23)
Higher Balance	61.83	70.80	90.01	109.8	81.04	91.95
	(33.91)	(27.14)	(23.96)	(16.69)	(16.75)	(14.97)
Gender	-5.178	-16.12	0.0103	16.47	10.04	4.796
	(18.33)	(17.28)	(11.35)	(13.38)	(12.74)	(8.899)
STEM/Economics	18.50	-2.535	16.45	21.19	4.828	9.604
	(17.72)	(16.25)	(11.58)	(10.55)	(10.96)	(6.967)
IB x Higher Interest Rate				88.85	85.81	79.83
				(43.24)	(46.29)	(35.74)
IB x Higher Balance				-48.03	-9.068	-1.849
				(37.37)	(31.88)	(28.12)
IB				-7.739	-15.08	-26.00
				(30.67)	(33.48)	(24.64)
Constant	99.55	105.3	97.09	94.27	101.5	124.7
	(32.96)	(34.75)	(21.96)	(20.80)	(20.09)	(15.86)
Observations	152	450	1135	353	1095	2452
R^2	0.433	0.503	0.416	0.432	0.462	0.375
Period	First	All	All	First	All	All
Restrict to Optimizers	Yes	Yes	No	Yes	Yes	No
Restrict to Interest Rate Acquirers	Yes	Yes	No	Yes	Yes	No
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.001$	$p = 0.0001$	$p = 0.003$			
$\beta_{\mathbf{IB}xHigherInterestRate} = 0$				$p = 0.04$	$p = 0.07$	$p = 0.03$
$\beta_{\mathbf{IB}xHigherBalance} = 0$				$p = 0.20$	$p = 0.78$	$p = 0.95$

Note: The table executes the analysis in Table A1.3 with demographic controls. *Gender* is a dummy variable that takes the value 1 for female subjects. *STEM/Economics* is a dummy variable that takes the value 1 for subjects whose majors are either STEM or Economics. Standard errors in parentheses. Errors are clustered at the subject level.

Table A1.9: Differences in Optimality Measures Across Investment Treatments

	Optimality Rate			Correct Allocation Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
Investment Interest Rate	0.0673	0.0994	-0.0224	8.864	10.09	-10.85
	(0.108)	(0.114)	(0.0889)	(27.19)	(26.26)	(22.71)
Constant	0.461	0.498	0.429	378.5	391.5	366.7
	(0.0765)	(0.0863)	(0.0623)	(17.69)	(18.48)	(15.42)
Observations	296	1170	2335	296	1170	2335
R^2	0.005	0.009	0.001	0.001	0.001	0.001
Period	First	All	All	First	All	All
Restrict to Optimizers	Yes	Yes	No	Yes	Yes	No
Restrict to Interest Rate Acquirers	Yes	Yes	No	Yes	Yes	No

Note: Each column reports the effect of being assigned to *Investment Interest Rate* treatment on some optimality measure using an OLS regression. In Columns 1,2 and 3, the dependent variable is a dummy that takes the value 1 if the allocation made is optimal. In Columns 3,4 and 5, the dependent variable is the amount of allocation made to the high interest rate fund which takes a value between 0 and 500. Columns 1 and 4 restrict the sample to observations from the first period in each stage where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 2 and 5 restrict the sample to observations where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 3 and 6 execute the same analysis without imposing any sample restrictions. Standard errors in parentheses. Errors are clustered at the subject level.

Table A1.10: Differences in Optimality Measures Across Investment Treatments with Demographic Controls

	Optimality Rate			Correct Allocation Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
IR	-0.0224	-0.0264	0.00209	-10.85	-11.86	-7.052
	(0.0889)	(0.0823)	(0.0765)	(22.71)	(21.05)	(20.34)
Math Score		0.373	0.377		94.74	97.58
		(0.0952)	(0.0935)		(26.10)	(26.95)
Gender			-0.300			-52.58
			(0.0773)			(20.74)
STEM/Economics			-0.0776			-18.91
			(0.0846)			(22.36)
Constant	0.429	0.265	0.460	366.7	325.0	361.4
	(0.0623)	(0.0729)	(0.0949)	(15.42)	(19.44)	(25.82)
Observations	2335	2335	2335	2335	2335	2335
R^2	0.001	0.094	0.188	0.001	0.060	0.089

Note: Column 1 to 3 represent the differences in the share of optimal allocations between *Investment Balance* and *Investment Interest Rate* treatments. The dependent variable *Optimal* is a dummy variable that takes the value 1 if the allocation is made optimally. Column 4 to 6 represent the differences in the amount of correctly made allocations between **IB** and **IR**. The dependent variable is the amount of allocation made on the high interest rate fund which takes a value in between 0 and 500. The unit of observation is *subject x period*. The term **IR** is a dummy variable that takes the value 1 for observations made under Investment Interest Rate treatment. *Math Score* is a discrete variable that takes values [0,0.25,0.5,0.75,1] representing the percentage of correct answers to four optimization problems. *Gender* is a dummy variable that takes the value 1 for female subjects. *STEM/Economics* is a dummy variable that takes the value 1 for subjects whose majors are either STEM or Economics. Standard errors in parentheses. Errors are clustered at the subject level.

Table A1.11: Estimation of Repayments Across Investment Treatments

	(1)	(2)	(3)	(4)	(5)	(6)
Higher Interest Rate	275.0	286.5	204.4	255.1	271.1	221.3
	(42.15)	(36.58)	(31.82)	(34.93)	(34.32)	(28.83)
Higher Balance	89.21	93.43	56.45	62.05	71.58	90.21
	(27.65)	(22.13)	(18.16)	(33.53)	(27.30)	(23.83)
IR x Higher Interest Rate				19.93	15.35	-16.89
				(54.43)	(49.88)	(42.80)
IR x Higher Balance				27.16	21.85	-33.76
				(43.29)	(35.00)	(29.90)
IR				-40.46	-42.92	7.296
				(35.05)	(34.45)	(27.94)
Constant	66.35	51.66	113.8	106.8	94.58	106.5
	(23.43)	(18.76)	(20.00)	(26.27)	(29.02)	(19.63)
Observations	144	720	1200	296	1170	2335
R^2	0.483	0.533	0.327	0.458	0.524	0.371
Period	First	All	All	First	All	All
Restrict to Optimizers	Yes	Yes	No	Yes	Yes	No
Restrict to Interest Rate Acquirers	Yes	Yes	No	Yes	Yes	No
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.005$	$p = 0.0006$	$p = 0.0001$			
$\beta_{\mathbf{IR}xHigherInterestRate} = 0$				$p = 0.72$	$p = 0.76$	$p = 0.7$
$\beta_{\mathbf{IR}xHigherBalance} = 0$				$p = 0.53$	$p = 0.54$	$p = 0.26$

Note: Columns 1 to 3 estimate, using OLS, how having a higher interest rate and a higher balance on a fund affects the allocations made towards that card in *Investment Interest Rate* treatment. The dependent variable is the amount of allocation made on the left card (without loss of generality) which takes a value in between 0 and 500. The regressors *Higher Interest Rate* and *Higher Balance* are two dummy variables that takes the value 1 whenever the interest rate and the balance on the left fund, respectively, is higher compared to the right fund. Columns 4 to 6 estimate, using OLS, how having a higher interest rate and a higher balance on a fund affect the allocations made towards that card using observations from both *Investment Interest Rate* and *Investment Balance* treatments. The term **DR** is a dummy variable that takes the value 1 if the allocation is made under *Investment Interest Rate* treatment. The terms **IR** x Higher Interest Rate and **IR** x Higher Balance are interaction variables. *Period* indicates if the analysis is limited to the first period decisions or not. *Restrict to Optimizers* indicate if the analysis is limited to subjects who can solve optimization problems. *Restrict to Interest Rate Acquirers* indicate if the analysis is limited to observations where the subjects acquired interest rate information before making their decisions. The last part of the table reports the parametric test results on estimated coefficients through associated *p*-values. Standard errors in parentheses. Errors are clustered at the subject level.

Table A1.12: Estimation of Repayments Across Investment Treatments with Demographic Controls

	(1)	(2)	(3)	(4)	(5)	(6)
Higher Interest Rate	274.5	286.6	205.1	253.9	272.7	221.3
	(42.34)	(36.46)	(31.77)	(34.87)	(33.87)	(28.80)
Higher Balance	89.21	93.08	56.14	61.87	71.11	90.20
	(27.85)	(22.35)	(18.25)	(33.58)	(27.06)	(23.83)
Gender	-5.804	-14.20	-7.793	-5.467	-14.96	-1.986
	(22.91)	(14.65)	(12.28)	(14.35)	(10.98)	(8.325)
STEM/Economics	9.392	-18.30	-16.19	14.64	-11.25	0.105
	(30.58)	(22.15)	(11.81)	(16.39)	(14.14)	(8.333)
IR x Higher Interest Rate				20.48	13.72	-16.93
				(54.34)	(49.85)	(42.79)
IR x Higher Balance				27.33	21.92	-33.80
				(43.43)	(34.81)	(29.92)
IR				-43.15	-40.17	7.539
				(35.37)	(34.79)	(28.04)
Constant	62.69	73.19	127.6	101.8	108.6	107.4
	(40.73)	(31.60)	(21.61)	(31.74)	(32.07)	(21.20)
Observations	144	720	1200	296	1170	2335
R^2	0.484	0.535	0.329	0.459	0.525	0.371
Period	First	All	All	First	All	All
Restrict to Optimizers	Yes	Yes	No	Yes	Yes	No
Restrict to Interest Rate Acquirers	Yes	Yes	No	Yes	Yes	No
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.005$	$p = 0.0006$	$p = 0.0001$			
$\beta_{\mathbf{IR}xHigherInterestRate} = 0$				$p = 0.71$	$p = 0.78$	$p = 0.69$
$\beta_{\mathbf{IR}xHigherBalance} = 0$				$p = 0.53$	$p = 0.53$	$p = 0.26$

Note: The table executes the analysis in Table A1.11 with demographic controls. *Gender* is a dummy variable that takes the value 1 for female subjects. *STEM/Economics* is a dummy variable that takes the value 1 for subjects whose majors are either STEM or Economics. Standard errors in parentheses. Errors are clustered at the subject level.

A.2 Role of Vividness under the Investment Frame

In Figure A2.1, Panel A shows the share of optimal repayments made across treatments and Panel B shows the average allocation made to the high interest rate fund for subjects who can solve optimization problems and who acquire interest rate information before making their decision in the first period of each stage. We see that there is no significant increase, on average, in any of the optimality measures. The share of optimal allocations increases by 6.7 percentage points -from 46.1% in **IB** to 52.8% in **IR** ($p = 0.54$). The average allocation to the high interest rate account increases by 8.9 ECU - from 378.5 ECU to 387.4 ECU ($p = 0.75$). The results are qualitatively similar when we relax our sample restrictions and control for demographic information (See Tables A1.9 and A1.10 in Appendix).

Figure A2.1: Comparison of Investment Treatments

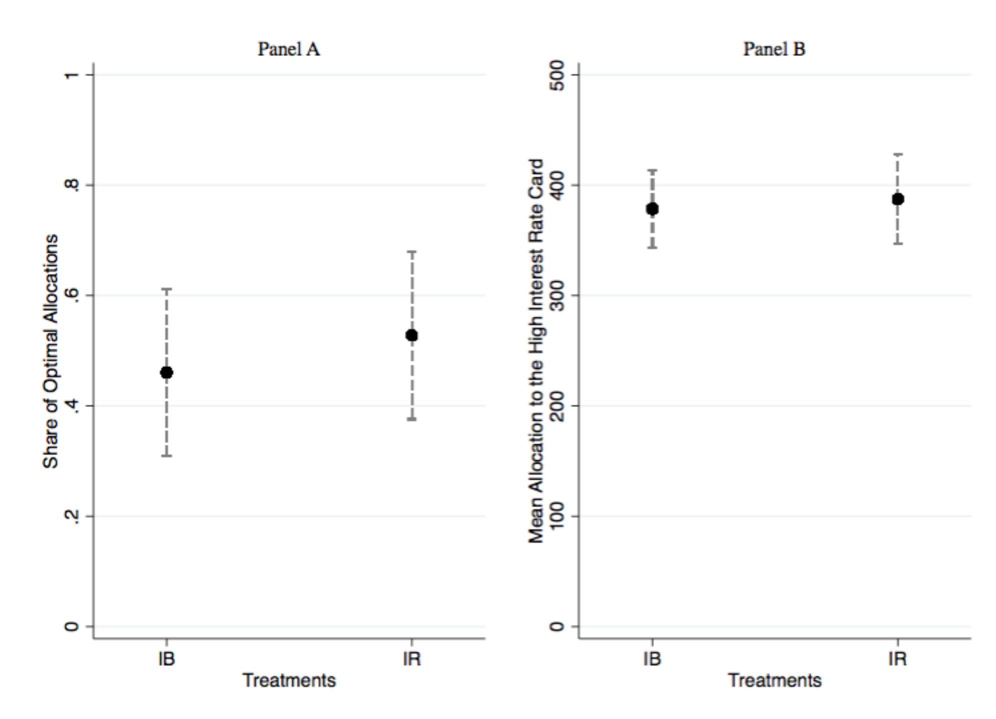

Note: Panel A shows the share of optimal allocations made in **IB** and **IR**. The whiskers indicate 95% confidence interval calculated using subject-level clusters. Panel B shows the average allocation made to the high interest rate card.

Figure A2.2 documents further evidence that allows us to compare the allocation patterns across treatments. The patterns seem mostly similar. We find that in aligned stages 92%, 84% and 96% (respectively) of the subjects allocate more than half of their deposit into the high interest rate fund which are similar to the rates calculated in Interest Balance treatment. Moreover, the percentage of subjects that allocate more than half of their deposit into the high interest rate fund in misaligned stages is respectively 63%, 75% and 55% which are, again, similar to the rates calculated in **IB**. Overall, we find no statistical difference in responsiveness to interest rate and balance information across subjects in **IR** and **IB** ($p = 0.71$ and $p = 0.53$, respectively). These findings are robust to relaxing our sample restrictions and including demographic controls (See Tables A1.11

Figure A2.2: Allocation Patterns Across Investment Treatments - Period 1 Decisions

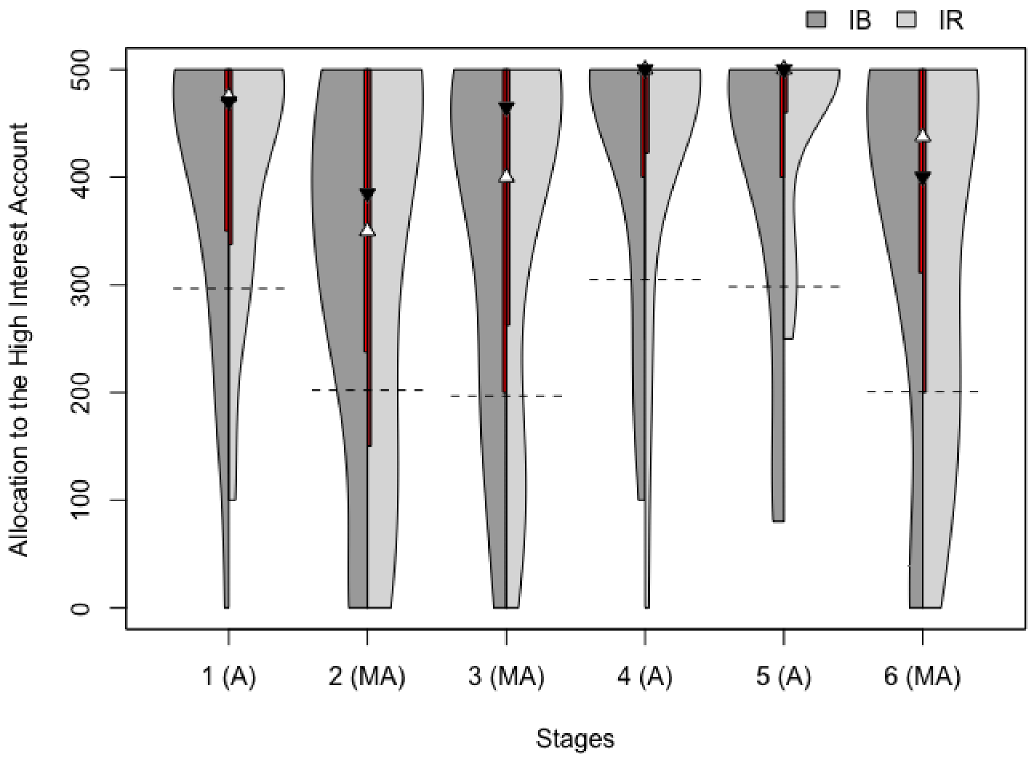

Note: The violin plots show the distribution of repayments subjects make toward the high interest rate fund in the first period of each stage. The upward white triangle and the downward black triangle represent the median allocation towards the higher interest rate card in a given stage for **IB** and **IR**, respectively. The thick red and blue bars around the median represents allocations within the interquartile range for **IB** and **IR**, respectively. The violin shape visualizes the kernel density distribution of the allocation patterns - the wider sections of the violin represents a higher likelihood of allocating in the corresponding value. The dotted horizontal reference lines represent the hypothetical allocation under an exact balance matching heuristic towards the higher interest card in the first period of each stage. The rational choice theory predicts a distribution with full mass located at 500 for all stages.

and A1.12 in Appendix).

Result A2.1 *Similar to the debt frame, neither the share of optimal allocations nor the average allocation to the high interest rate account improves with an increase in the vividness of interest rate information across investment frames.*

A.3 Learning

Table A3.1: Within Stage Learning in **DB**

	Optimal		Correct Fraction	
	(1)	(2)	(3)	(4)
Period	0.0189	-0.00534	0.0136	-0.00324
	(0.0138)	(0.00553)	(0.00723)	(0.00346)
Constant	0.202	0.204	0.647	0.647
	(0.0599)	(0.0488)	(0.0266)	(0.0221)
Observations	645	1317	645	1317
R^2	0.004	0.000	0.006	0.000

Note: Each column reports the coefficients from an OLS regression of some optimality measure on the decision period within a stage denoted with the variable *Period*. In Columns 1 and 2, the dependent variable is a dummy that takes the value 1 if the allocation made is optimal. In Columns 3 and 4, the dependent variable is the fraction of correctly made allocation that takes a value between 0 and 1. Columns 1 and 3 restrict the sample to observations where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 2 and 4 execute the same analysis without imposing any sample restrictions. Standard errors in parentheses. Errors are clustered at the subject level.

Table A3.2: Between Stage Learning in **DB**

	Optimality Rate		Mean Correct Fraction	
	(1)	(2)	(3)	(4)
Bi-Stage	0.0334	0.0246	0.000657	0.00441
	(0.0212)	(0.0136)	(0.0125)	(0.00848)
Constant	0.188	0.138	0.682	0.628
	(0.0680)	(0.0426)	(0.0339)	(0.0235)
Observations	645	1317	645	1317
R^2	0.004	0.003	0.000	0.000

Note: Each column reports the coefficients from an OLS regression of some optimality measure on the bi-stages. A bi-stage consists of two consecutive stages with one aligned and one misaligned stage, and takes an integer value in between 1 and 3. In Columns 1 and 2, the dependent variable is a dummy that takes the value 1 if the allocation made is optimal. In Columns 3 and 4, the dependent variable is the fraction of correctly made allocation that takes a value between 0 and 1. Columns 1 and 3 restrict the sample to observations where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 2 and 4 execute the same analysis without imposing any sample restrictions. Standard errors in parentheses. Errors are clustered at the subject level.

Table A3.3: Within Stage Learning in **DR**

	Optimal		Correct Fraction	
	(1)	(2)	(3)	(4)
Period	-0.0153	-0.00717	0.00345	0.00634
	(0.00570)	(0.00519)	(0.00466)	(0.00397)
Constant	0.267	0.197	0.631	0.609
	(0.0589)	(0.0461)	(0.0266)	(0.0214)
Observations	928	1288	928	1288
R^2	0.003	0.001	0.000	0.001

Note: Each column reports the coefficients from an OLS regression of some optimality measure on the decision period within a stage denoted with the variable *Period*. In Columns 1 and 2, the dependent variable is a dummy that takes the value 1 if the allocation made is optimal. In Columns 3 and 4, the dependent variable is the fraction of correctly made allocation that takes a value between 0 and 1. Columns 1 and 3 restrict the sample to observations where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 2 and 4 execute the same analysis without imposing any sample restrictions. Standard errors in parentheses. Errors are clustered at the subject level.

Table A3.4: Between Stage Learning in **DR**

	Optimality Rate		Mean Correct Fraction	
	(1)	(2)	(3)	(4)
Bi-Stage	0.0155	0.0147	0.0139	0.00929
	(0.0130)	(0.00966)	(0.00689)	(0.00610)
Constant	0.190	0.146	0.613	0.609
	(0.0494)	(0.0376)	(0.0209)	(0.0167)
Observations	928	1288	928	1288
R^2	0.001	0.001	0.002	0.001

Note: Each column reports the coefficients from an OLS regression of some optimality measure on the bi-stages. A bi-stage consists of two consecutive stages with one aligned and one misaligned stage, and takes an integer value in between 1 and 3. In Columns 1 and 2, the dependent variable is a dummy that takes the value 1 if the allocation made is optimal. In Columns 3 and 4, the dependent variable is the fraction of correctly made allocation that takes a value between 0 and 1. Columns 1 and 3 restrict the sample to observations where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 2 and 4 execute the same analysis without imposing any sample restrictions. Standard errors in parentheses. Errors are clustered at the subject level.

Table A3.5: Within Stage Learning in **IB**

	Optimal		Correct Fraction	
	(1)	(2)	(3)	(4)
Period	0.0192	0.0207	0.0183	0.0107
	(0.0176)	(0.00832)	(0.00794)	(0.00381)
Constant	0.449	0.367	0.737	0.701
	(0.0741)	(0.0620)	(0.0390)	(0.0317)
Observations	450	1135	450	1135
R^2	0.003	0.003	0.008	0.002

Note: Each column reports the coefficients from an OLS regression of some optimality measure on the decision period within a stage denoted with the variable *Period*. In Columns 1 and 2, the dependent variable is a dummy that takes the value 1 if the allocation made is optimal. In Columns 3 and 4, the dependent variable is the fraction of correctly made allocation that takes a value between 0 and 1. Columns 1 and 3 restrict the sample to observations where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 2 and 4 execute the same analysis without imposing any sample restrictions. Standard errors in parentheses. Errors are clustered at the subject level.

Table A3.6: Between Stage Learning in **IB**

	Optimality Rate		Mean Correct Fraction	
	(1)	(2)	(3)	(4)
Bi-Stage	0.00438	0.0369	0.0311	0.0270
	(0.0246)	(0.0148)	(0.0181)	(0.0125)
Constant	0.489	0.355	0.724	0.679
	(0.0844)	(0.0601)	(0.0527)	(0.0366)
Observations	450	1135	450	1135
R^2	0.000	0.004	0.008	0.005

Note: Each column reports the coefficients from an OLS regression of some optimality measure on the bi-stages. A bi-stage consists of two consecutive stages with one aligned and one misaligned stage, and takes an integer value in between 1 and 3. In Columns 1 and 2, the dependent variable is a dummy that takes the value 1 if the allocation made is optimal. In Columns 3 and 4, the dependent variable is the fraction of correctly made allocation that takes a value between 0 and 1. Columns 1 and 3 restrict the sample to observations where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 2 and 4 execute the same analysis without imposing any sample restrictions. Standard errors in parentheses. Errors are clustered at the subject level.

Table A3.7: Within Stage Learning in **IR**

	Optimal		Correct Fraction	
	(1)	(2)	(3)	(4)
Period	0.0215	0.0163	0.00247	0.00333
	(0.00934)	(0.00666)	(0.00597)	(0.00459)
Constant	0.533	0.358	0.796	0.702
	(0.0785)	(0.0644)	(0.0401)	(0.0357)
Observations	720	1200	720	1200
R^2	0.004	0.002	0.000	0.000

Note: Each column reports the coefficients from an OLS regression of some optimality measure on the decision period within a stage denoted with the variable *Period*. In Columns 1 and 2, the dependent variable is a dummy that takes the value 1 if the allocation made is optimal. In Columns 3 and 4, the dependent variable is the fraction of correctly made allocation that takes a value between 0 and 1. Columns 1 and 3 restrict the sample to observations where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 2 and 4 execute the same analysis without imposing any sample restrictions. Standard errors in parentheses. Errors are clustered at the subject level.

Table A3.8: Between Stage Learning in **IR**

	Optimality Rate		Mean Correct Fraction	
	(1)	(2)	(3)	(4)
Bi-Stage	0.0167	0.0150	0.0102	0.00929
	(0.0277)	(0.0171)	(0.0243)	(0.0155)
Constant	0.564	0.377	0.783	0.693
	(0.0989)	(0.0741)	(0.0696)	(0.0495)
Observations	720	1200	720	1200
R^2	0.001	0.001	0.001	0.001

Note: Each column reports the coefficients from an OLS regression of some optimality measure on the bi-stages. A bi-stage consists of two consecutive stages with one aligned and one misaligned stage, and takes an integer value in between 1 and 3. In Columns 1 and 2, the dependent variable is a dummy that takes the value 1 if the allocation made is optimal. In Columns 3 and 4, the dependent variable is the fraction of correctly made allocation that takes a value between 0 and 1. Columns 1 and 3 restrict the sample to observations where the subject acquires interest rate information and solves at least one optimization question correctly. Columns 2 and 4 execute the same analysis without imposing any sample restrictions. Standard errors in parentheses. Errors are clustered at the subject level.

A.4 Information Acquisition and the Measures of Optimality

Table A4.1: Click Rates on Information Buttons across No-Vivid Treatments

	(1) Interest Rate	(2) Current Balance	(3) Other	(4) Total
IN	-0.0168	-0.597	-0.333	-0.947
	(0.102)	(0.200)	(0.277)	(0.458)
Constant	0.863	1.595	1.526	3.985
	(0.0776)	(0.133)	(0.216)	(0.358)
Observations	1102	1102	1102	1102

Note: The table documents the differences in average click rates on various information buttons between *Debt No-Vivid* and *Investment No-Vivid* treatments. The unit of observation is *subject* × *period* × *click rate*. The regressor **IN** is a dummy variable that takes the value 1 for observations under *Investment No-Vivid* treatment. The dependent variables *Interest Rate*, *Current Balance*, and *Total* take non-negative integer values that respectively indicate the number a subject click on interest rate button, current balance button, and any information button. Similarly, the dependent variable *Other* in Column 3 takes non-negative integer values that indicates the total number a subject clicks on either interest charged/earned button, previous payment/investment button and previous balance button. Standard errors in parentheses. Errors are clustered at the subject level.

In Section 3, we present evidence that there is a wedge in the share of optimal allocations across frames as well as in the click rates and time spent on current balance information button. Here, we tie these pieces of evidence together by presenting how clicking and spending time on certain information buttons are correlated with consequent choices of the subjects.

Table A4.3 shows how our measures of optimality are correlated with click rates and time spent on information buttons. Column 1 indicates that each click to interest rate button is correlated with 5.6% increase in optimal allocations ($p = 0.058$) whereas each click to current balance button is correlated with a 6.8% decrease ($p = 0.04$). The difference in magnitude of these changes is significant ($p = 0.03$). Column 2 indicates that each click to interest rate button is correlated with an increase of 25.2 ECU in correctly made allocations ($p = 0.02$) whereas each click to current balance button is

Table A4.2: Time Spent on Information Buttons across No-Vivid Treatments

	(1) Interest Rate	(2) Current Balance	(3) Other	(4) Total
IN	-0.262	-4.498	-1.961	-6.721
	(0.529)	(1.225)	(1.406)	(2.227)
Constant	3.404	9.562	7.832	20.80
	(0.373)	(0.947)	(1.006)	(1.716)
Observations	1110	1110	1110	1110

Note: The table documents the differences in time spent on various information buttons between Debt No-Vivid and Investment No-Vivid treatments. The unit of observation is *subject x period*. The regressor **IN** is a dummy variable that takes the value 1 for observations under *Investment No-Vivid* treatment. The dependent variable *Interest Rate* in Column 1 takes a positive real value that indicates the time (in seconds) a subject spends on interest rate button within a period. The dependent variable *Current Balance* in Column 2 takes a positive real value that indicates the time (in seconds) a subject spends on current balance button within a period. The dependent variable *Other* in Column 3 takes a positive real value that indicates the total time (in seconds) a subject spends on interest charged/earned button, previous payment/investment button and previous balance button within a period. The dependent variable *Total* in Column 4 takes a positive real value that indicates the total time (in seconds) a subject spends on all information buttons. Standard errors in parentheses. Errors are clustered at the subject level.

correlated with a decrease of 20.4 ECU ($p = 0.05$). The difference in magnitude of these changes is significant ($p = 0.02$).

Columns 3 and 4 show how time spent correlates with our measures of optimality. Here we find that each additional second spent on interest rate button has no impact on either the share of optimal allocations or on the amount of allocation correctly made ($p = 0.7$). However, we find that each additional second that is spent on current balance button correlates with a 0.59 percentage point decrease in the level of optimality ($p = 0.005$). Similarly, each additional second spent on other information buttons correlates with a 0.48 percentage point decrease ($p = 0.01$) in the share of optimal allocations. The amount of correctly made allocation decreases by 2 ECU for each second spent on current balance button ($p = 0.005$) and decreases by 1.48 ECU for each second spent on other information ($p = 0.04$).

Result A4.1 *Each click to interest rate button is correlated with an increase in the correctly allocated amount whereas each click to current balance button correlates with*

Table A4.3: Click Rates, Time Spent and Measures of Optimality

	Click Rate		Time Spent	
	(1)	(2)	(3)	(4)
	Optimal	Allocation	Optimal	Allocation
Interest Rate	0.0563	25.23	0.000939	1.086
	(0.0288)	(9.985)	(0.00236)	(0.919)
Current Balance	-0.0683	-20.39	-0.00593	-2.069
	(0.0323)	(9.864)	(0.00197)	(0.696)
Other	-0.0219	-8.497	-0.00484	-1.483
	(0.0148)	(4.336)	(0.00185)	(0.695)
IN	0.101	37.18	0.111	39.54
	(0.0910)	(25.37)	(0.0986)	(26.14)
Math Score	0.347	59.99	0.343	59.24
	(0.117)	(33.34)	(0.125)	(35.78)
Constant	0.159	295.8	0.159	300.3
	(0.0591)	(17.43)	(0.0604)	(17.66)
Observations	1102	1102	1102	1102
R^2	0.171	0.092	0.161	0.083

Note: The table documents how click rates and time spent on information buttons are correlated with making an optimal allocation. The regressors *Interest Rate, Current Balance* and *Other* represent click rates (in Columns 1 and 2) and time spent (in Columns 3 and 4) on the respective buttons. The regressor **IN** is a dummy variable that takes the value 1 for observations under *Investment No-Vivid* treatment. *Math Score* is a discrete variable that takes values [0,0.25,0.5,0.75,1] representing the percentage of correct answers to four optimization problems. The dependent variable *Optimal* is a dummy that takes the value 1 for optimal payments. The variable *Allocation* indicates the amount of correctly made allocation by a subject in a period. Standard errors in parentheses. Errors are clustered at the subject level.

a decrease. Moreover, time spent on the interest rate button does not correlate with the correctly allocated amount whereas each second spent on current balance information correlates with a decrease.

A.5 Use of Heuristics - Heuristic Transition Matrices

Table A5.1: Debt Frame: Bi-Stage 1 to Bi-Stage 2

	$Other_2$	IM_2	Opt_2	BM_2
$Other_1$	4	2	0	3
IM_1	3	18	1	10
Opt_1	0	3	4	0
BM_1	4	10	0	47

Table A5.2: Debt Frame: Bi-Stage 2 to Bi-Stage 3

	$Other_3$	IM_3	Opt_3	BM_3
$Other_2$	4	2	0	5
IM_2	4	13	5	11
Opt_2	0	0	5	0
BM_2	5	16	0	39

Table A5.3: Investment Frame: Bi-Stage 1 to Bi-Stage 2

	$Other_2$	IM_2	Opt_2	BM_2
$Other_1$	10	5	1	2
IM_1	3	23	5	3
Opt_1	0	3	17	1
BM_1	6	2	4	8

Table A5.4: Investment Frame: Bi-Stage 2 to Bi-Stage 3

	$Other_3$	IM_3	Opt_3	BM_3
$Other_2$	10	8	0	1
IM_2	4	21	2	6
Opt_2	0	6	19	2
BM_2	1	2	3	8

Note: The tables describe the share of subjects who are assigned to a heuristic type in a certain bi-stage by the heuristic type they are assigned in the consecutive bi-stage. In order to construct these matrices, we employ the weak classification requirement. Under the weak classification, a subject is considered as a balance matching (BM) type if she allocates at least 50% of her deposit to the account with the higher balances for at least 6 out of 10 periods within a bi-stage. Similarly, a subject is considered as an interest matching (IM) type if she allocates between 50% to 95% of her deposit to the account with the higher interest rate for at least 6 out of 10 periods. A subject is considered as an optimal type if she allocates at least 95% of her deposit to the account with the higher interest rate for at least 6 out of 10 periods. When the criteria for both BM and IM are satisfied, we give the tie breaker to BM.

A.6 Conceptual Framework

There is a unit mass of identical decision makers who allocate a fixed amount of income M to two accounts with differing interest rates $r = (r_1, r_2) \in [0, 1]^2$ and balances $b = (b_1, b_2) \in \mathbb{R}^2$. We assume for simplicity $r_1 > r_2$. The decision maker i chooses $c^i \in [0, M]^2$ where each dimension represents an allocation made to an account and each choice satisfies $c_1^i + c_2^i = M$. A decision maker's outcome-based utility if she chooses the allocation (c_1^i, c_2^i) is given by $U(c^i; r, b) = \sum_{j=1}^{2}(1 + r_j)(c_j^i + b_j)$ which simply states that the utility from a choice is the sum of total balances after both accounts accrue interest. Hence the outcome-based utility strictly increases in c_1^i and decreases in c_2^i. However, instead of maximizing outcome-based utility, the decision maker maximizes the salience-adjusted utility function

$$\tilde{U}(c^i; r, b) = \sum_{j=1}^{2}(1 + w_r r_j)(c_j^i + b_j)$$

where $w_r \in \{0, 1\}$ is the salience adjustment on interest rate information.

Our model's central assumption concerns how salience adjustment w_r is determined. We model the decision maker's salience to interest rate information as a function of attention to interest rate and balance information. The decision maker i's attention to interest rate and balance information are respectively given by the parameters $a_r^i \in \mathbb{R}_+$ and $a_b^i \in \mathbb{R}_+$. Following (15), we define the salience of interest rate information $\sigma_r^i \in \mathbb{R}$ as the attention differential between interest rate information and balance information

$$\sigma_r^i = a_r^i - a_b^i$$

We assume that σ_r^i follows a normal distribution with mean μ and variance σ_ε^2, and is independent and identical across decision makers. The decision maker obtains a realization of σ_r^i and uses the salience adjustment rule $w_r = \mathbb{1}(\sigma_r^i \geq 0)$. This stylized salience adjustment rule that we assume is consistent with the view of many psychologists and economists that information that attracts greater attention contributes more strongly to the observed choices ((16), (17), (25)). The model captures how salience of interest rate information affects the decision maker's choices in a simple fashion: If the decision maker obtains a non-negative realization of salience of interest rate information, then her optimal decision overlaps with the optimal decision of a rational decision maker. Otherwise she does not take the interest rate information into account and her optimal decision involves uniformly randomizing over choices that are available to her.

Given this salience adjustment rule, we expect the allocation to the high interest rate account to be

$$\mathbb{E}[\bar{c}_1] = \left(1 + \Phi\left(\frac{\mu}{\sigma_\varepsilon}\right)\right) M/2$$

where $\Phi(\cdot)$ represents the standard normal cumulative distribution function. A critical observation here is that the expected allocation to high interest rate account is strictly increasing in the mean attention differential to interest rate μ. Hence any change in the decision environment that increases the salience of interest rate information should lead to an increase in the average allocation made to the high interest rate account.

A.7 Experiment Interface and Instructions

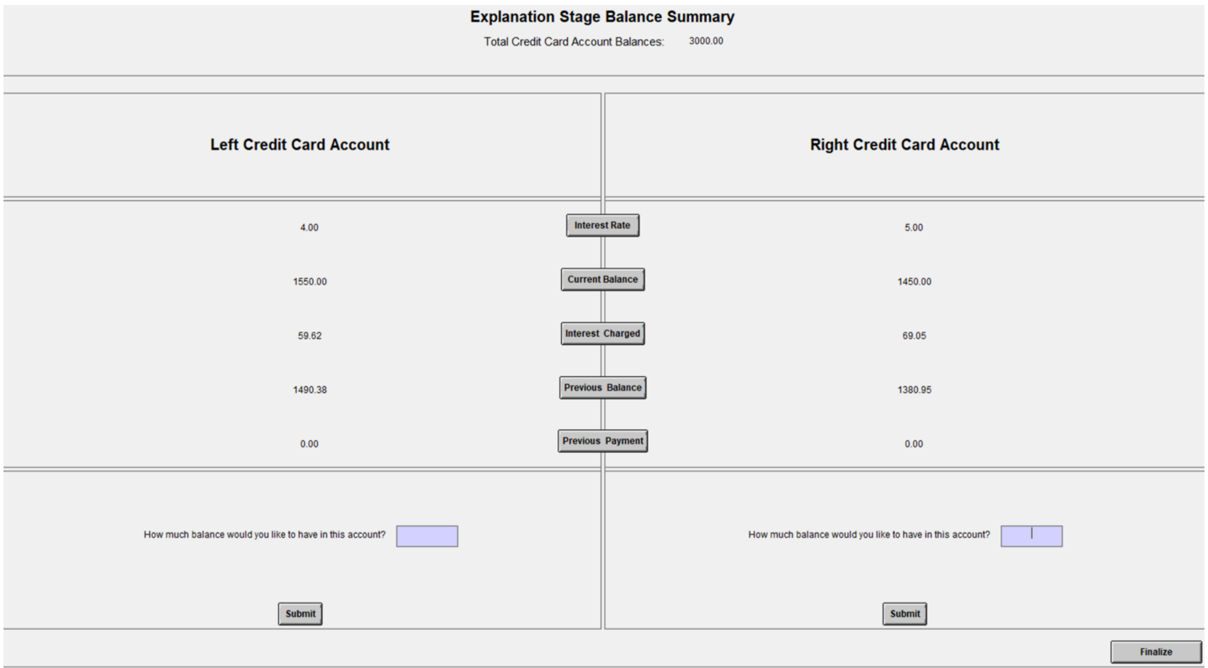

Figure A7.1: Experiment Interface for the treatment **DB** in Balance Reallocation Periods

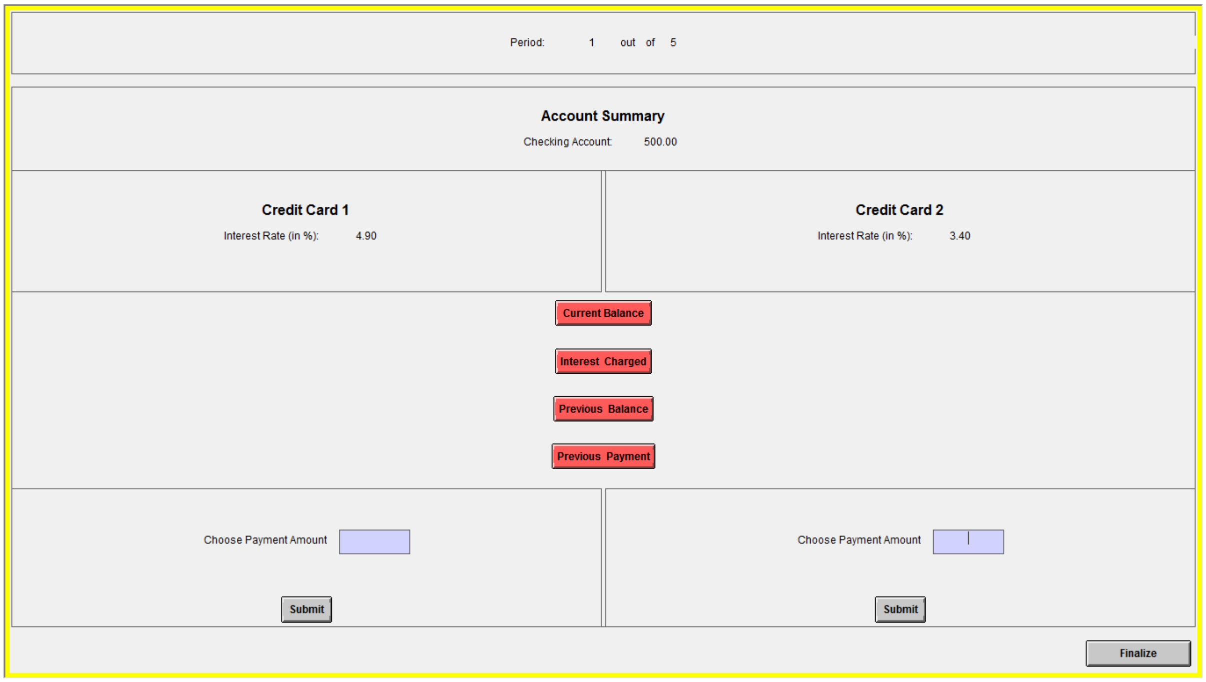

Figure A7.2: Experiment Interface for the treatment **DR**

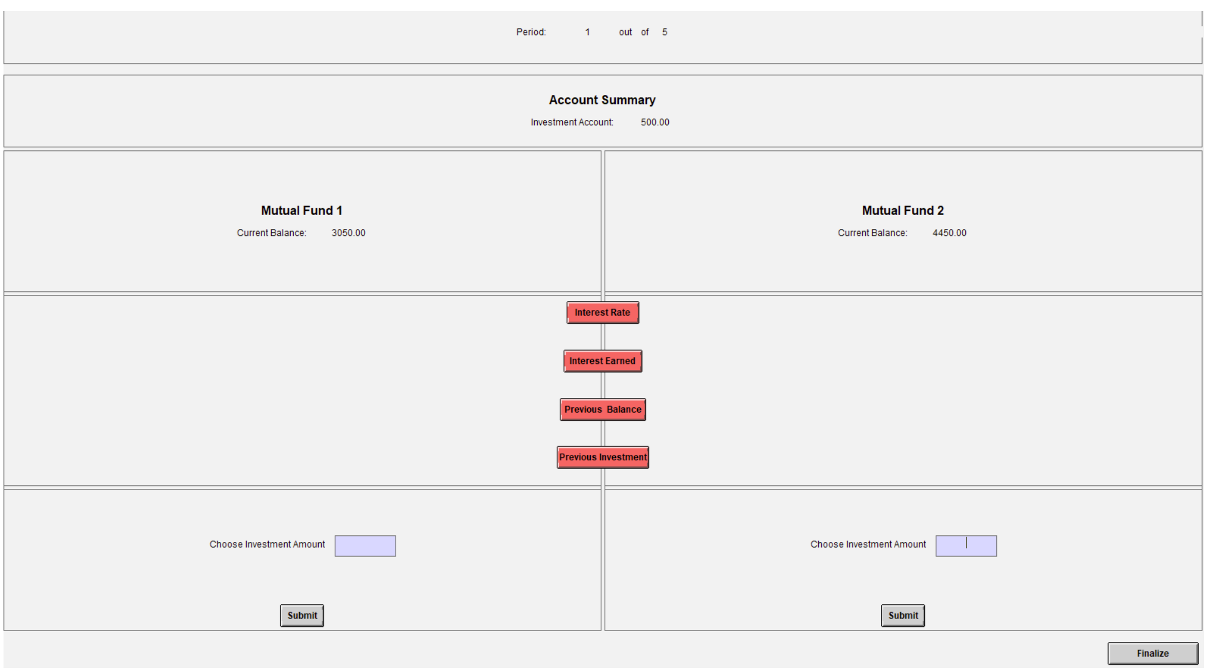

Figure A7.3: Experiment Interface for the treatment **IB**

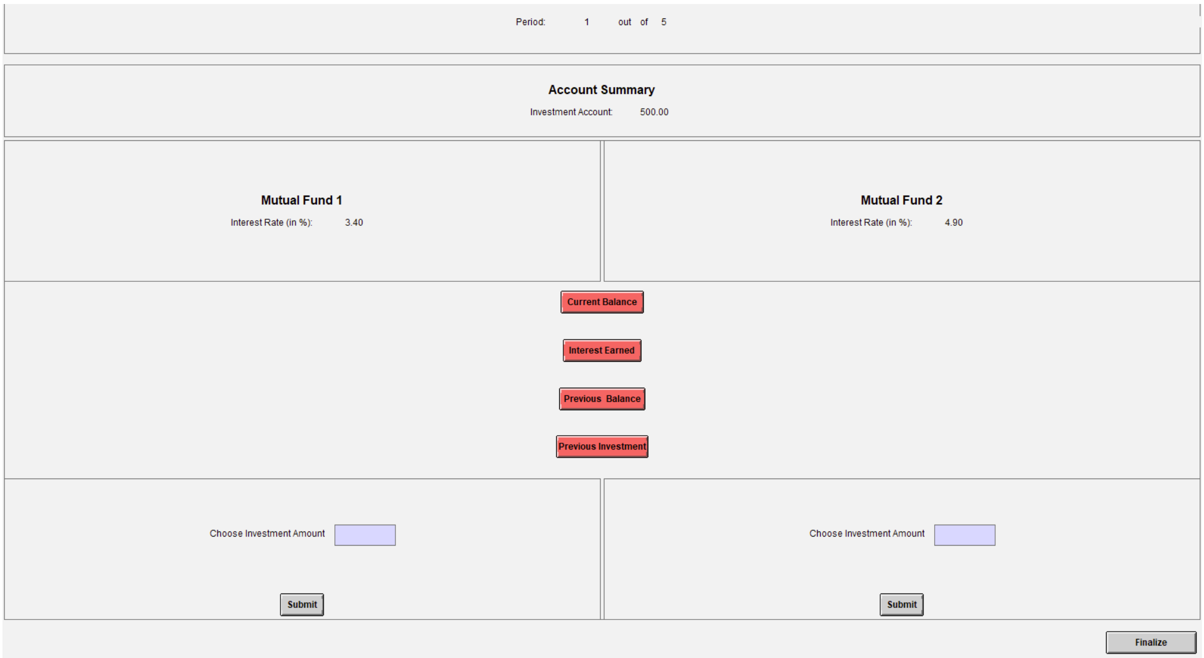

Figure A7.4: Experiment Interface for the treatment **IR**

Figure A7.5: Experiment Interface for the treatment DN

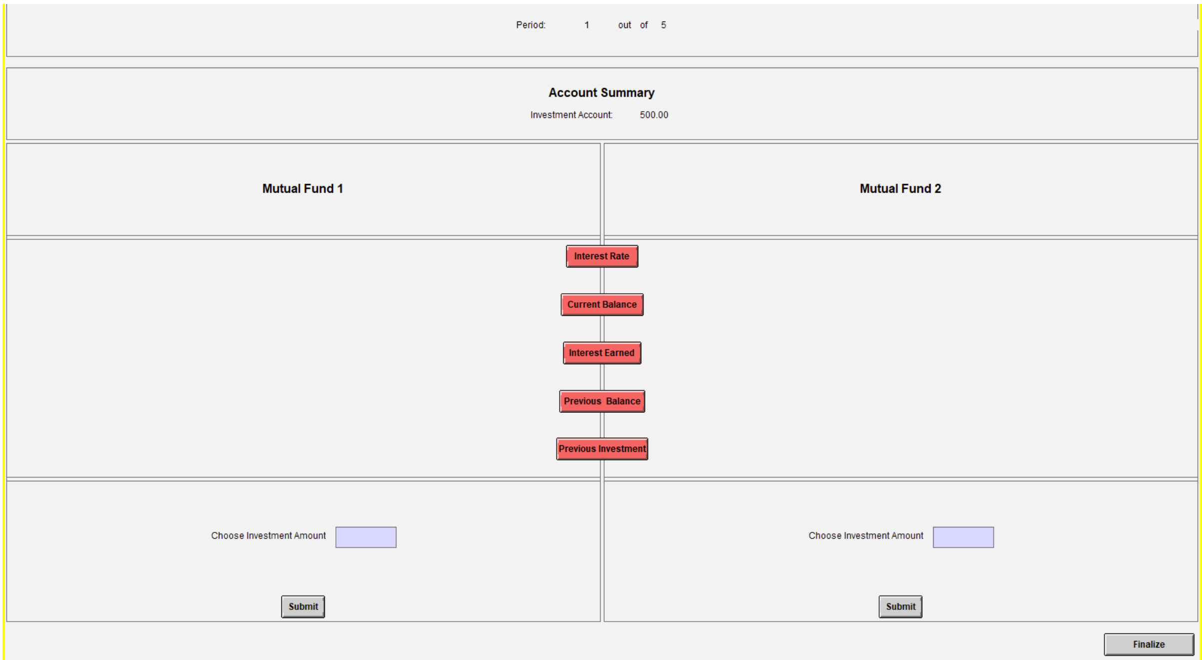

Figure A7.6: Experiment Interface for the treatment IN

Experiment Instructions for Debt Treatments

INSTRUCTIONS

Welcome

You are about to participate in a decision making experiment. In this experiment, you have the ability to earn a considerable amount of money, which will be paid to you in cash at the end of the experiment. The amount of money you earn will depend partly on your decisions. Therefore, it is in your best interest that you read these instructions carefully in order to have a clear understanding of the rules of the experiment. If you need assistance, please raise your hand quietly. Someone will come and answer your question in private.

This experiment is going to be conducted through computer terminals. The information provided to you on your terminal is private and it belongs only to you. It is very important that you do not communicate with other participants for the duration of the experiment. All necessary decision making information will be provided to you through your terminal. Please turn off your cell phone now, and refrain from opening any other programs or browsers on your computer during the experiment.

Economics experiments have a strict policy against deception. The rules you are going to read next will be implemented just as they are written.

The experiment should take no more than 60 minutes.

Background

This is a financial decision making experiment. In this experiment, you will be assigned two credit card accounts and a checking account. The experiment will be divided into stages and periods where you will be asked to make payments toward these credit card accounts.

Experiment Roadmap

The main experiment contains 6 Independent Stages. Each stage consists of 5 payment periods. You will be presented with different credit cards in each stage.

Your Task

At the beginning of each period, you will receive a fixed amount of money, called a deposit, in your checking account. Your task in each period is to make credit card payment decisions, using the amount of money you have available in your checking account.

A Period

There will be multiple periods in the experiment. An experimental period starts when you receive your deposit, and ends when you finalize your payments to each card for that period.

Level of Debt

At the beginning of the first period, each credit card will be assigned a level of debt. From the second period onward, the level of debt will be determined by two factors: interest rates and your previous period's payment decisions for each card. To illustrate this point, consider the following example:

Suppose that you have two credit cards, Left and Right. Your Left Card has a 4% per period interest rate and you owe 2,000 on that card. Your Right Card has a 5% per period interest rate and you owe 1,000 on that card. After you determine your payments on each card, your *Total Credit Card Debt in the following period* will be calculated as

(1+ 4%) (2,000 - Payment to Left Card) + (1+5%) (1,000 - Payment to Right Card)

Your *End of Stage Total Credit Card Debt* will be calculated as above once you make your last payment decision in that stage.

Your Payment

You will have an initial endowment of 6,500 experimental currency units (ECUs) at the beginning of each stage. To determine a *Stage Payoff*, we will subtract your End of Stage Total Credit Card Debt from your initial endowment. Your stage payoff will then be converted into US Dollars at the rate of 25 ECUs=$1. Only one stage payoff will be randomly selected as your cash payment in the end. All stage payoffs have the same chance of being selected.

Thank you for your participation in this experiment.

Key Features Recap

Setting:	Two credit card accounts
Task:	Make payment decisions on both cards
Duration:	5 periods per stage, 6 stages
Time:	No strict time restriction (as long as total time < 60 mins)
Payoff:	The less the total debt you have at the end of each stage, the more money you will make from the experiment

We will explain how the to use the interface next, please wait for further instructions.

Experiment Instructions for Balance Reallocation Periods

Instructions for Balance Reallocation

In this part of the experiment, you will go through the remaining two stages. The first 5 periods of these stages will be exactly the same as before. However, there is going to be an additional, sixth, period at the end of each stage. We will call these additional periods *Balance Reallocation Periods*. During these periods you will not be assigned a deposit, nor be asked to make a payment decision. Instead, your task will be reallocating your total debt between two cards.

Your stage payoff will be calculated similar to previous stages. We will subtract your End of Stage Total Credit Card Debt from your initial endowment. In this part of the experiment, we change your initial endowment to be 7,390 ECUs. Consider the following example:

Suppose that at the beginning of a Balance Reallocation period, your Left Card has 4% interest rate and you owe 2,000 on that card. Your Right Card has 5% interest rate and you owe 1,000 on that card. After you determine your new debt level on each card, your *End of Stage Total Credit Card Deb*t will be calculated as

(1 + 4%)(New Debt Level on Left Card) + (1 + 5%) (New Debt Level on Right Card)

To determine a Stage Payoff, we will subtract your End of Stage Total Credit Card Debt from your initial endowment of 7,390 ECUs. Your stage payoff will then be converted into US Dollars at the rate of 25 ECUs=$1 as before. Remember that each stage is equally likely to be selected for your payment.

You will go through an explanation period before you start making your decisions.

This explanation period will not count for money.

What Has Changed?

- Each stage has an additional Balance Reallocation period as a 6th period
- Your task in those periods is to adjust your balance levels on each card
- Your initial endowment is 7,390 ECUs

Experiment Instructions for Investment Treatments

INSTRUCTIONS

Welcome

You are about to participate in a decision making experiment. In this experiment, you have the ability to earn a considerable amount of money, which will be paid to you in cash at the end of the experiment. The amount of money you earn will partly depend on your decisions. Therefore, it is in your best interest that you read these instructions carefully in order to have a clear understanding of the rules of the experiment. If you need assistance, please raise your hand quietly. Someone will come and answer your question in private.

This experiment is going to be conducted through computer terminals. The information provided to you on your terminal is private and it belongs only to you. It is very important that you do not communicate with other participants for the duration of the experiment. All necessary decision making information will be provided to you through your terminal. Please turn off your cell phone now, and refrain from opening any other programs or browsers on your computer during the experiment.

Economics experiments have a strict policy against deception. The rules you are going to read next will be implemented just as they are written.

The experiment should take no more than 60 minutes.

Background

This is a financial decision making experiment. In this experiment, you will be assigned two mutual funds and an investment account. The experiment will be divided into stages and periods where you will be asked to make investment decisions toward these mutual funds.

Experiment Roadmap

The main experiment contains 6 Independent Stages. Each stage consists of 5 investment periods. You will be presented with different mutual funds in each stage.

Your Task

At the beginning of each stage, you will be given a loan to be repaid so that you have some amount of money to invest. At the beginning of each period, you will receive a fixed amount of money, called a deposit, in your investment account. Your task in each period is to make investment decisions, using the amount of money you have available in your investment account.

A Period

There will be multiple periods in the experiment. An experimental period starts when you receive your deposit, and ends when you finalize your investment decisions on each fund for that period.

Level of Investment

At the beginning of the first period, each mutual fund will be assigned a level of investment. From the second period onward, the level of investment will be determined by two factors: interest rates and your previous period's investment decisions on each fund. To illustrate this point, consider the following example:

Suppose that you have two mutual funds, Left and Right. Your Left Fund has a 4% per period interest rate and you own 2,000 in that fund. Your Right Fund has a 5% per period interest rate and you own 1,000 in that fund. After you determine your investment decisions on each fund, your *Total Investment in the following period* will be calculated as

(1+4%) (2,000 + Investment to Left Fund) + (1+5%) (1,000 + Investment to Right Fund)

Your *End of Stage Total Investment* will be calculated as above once you make your last investment decision in that stage.

Your Payment

To determine a *Stage Payoff*, we will subtract a loan repayment of 12,000 experimental currency units (ECUs) from your End of Stage Total Investment. Your stage payoff will then be converted into US Dollars at the rate of 25 ECUs=$1. Only one stage payoff will be randomly selected as your cash payment in the end. All stage payoffs have the same chance of being selected.

Thank you for your participation in this experiment.

Key Features Recap

Setting:	Two mutual funds
Task:	Make investment decisions on both funds
Duration:	5 periods per stage, 6 stages
Time:	No strict time restriction (as long as total time < 60 mins)
Payoff:	The higher the total investment you have at the end of each stage, the more money you will make from the experiment

We will explain how to use the interface next, please wait for further instructions.

Appendix B

Appendix for *Restoring Rational Choice in Repayments*

B.1 Additional Results

Figure B1.1: Misallocation Rate by Optimization Ability and Financial Literacy

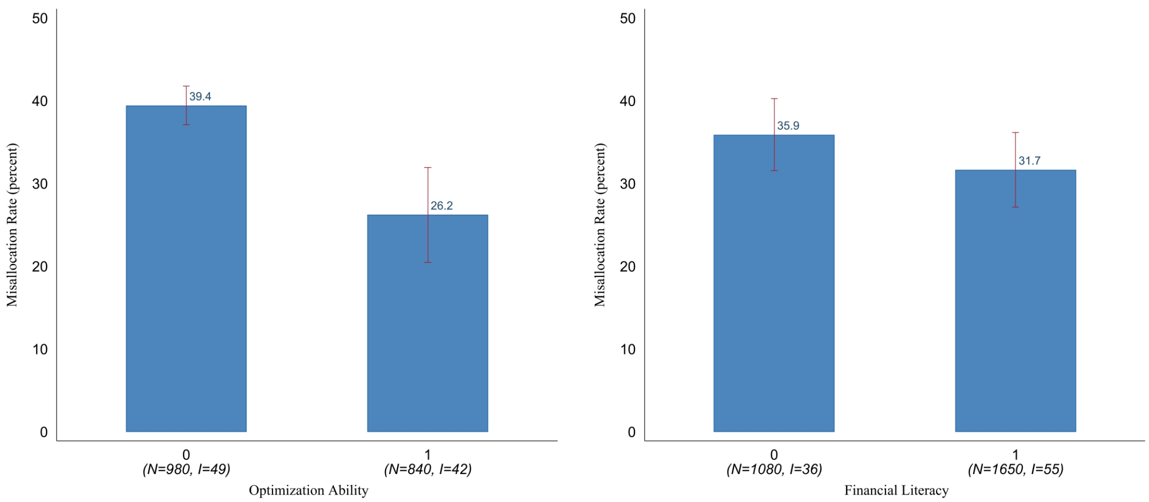

Notes: Panel A shows the average misallocation rate by our subjects' ability to solve an algebraic version of the credit card repayment problem. Subjects who are unable to solve the algebraic version of the credit card repayment problem are indicated by the group O. Panel B shows the average misallocation rate by our subjects' ability to show the Big Three financial literacy questions. Subjects who are able to solve *all* big three questions are indicated by group 1. The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Table B1.1: Causal Effect of Financial Advice on Misallocation Rate

	OLS	First Stage	IV Estimate	First Stage	IV Estimate
	(1)	(2)	(3)	(4)	(5)
	Misallocation Rate	Purchased Advice	Misallocation Rate	Purchased Advice	Misallocation Rate
Purchased Advice	-15.52***		-16.30***		-16.66***
	(3.527)		(3.956)		(3.839)
Actual WTP	-0.00702	0.0724***	0.0435	0.0732***	0.0420
	(0.543)	(0.0167)	(0.571)	(0.0157)	(0.545)
Rational WTP	6.057***	-0.00867	6.056***	-0.00661	5.765***
	(0.487)	(0.0138)	(0.486)	(0.0165)	(0.563)
BDM Price		-0.0654***		-0.0664***	
		(0.00710)		(0.00708)	
Observations	910	2730	910	2730	910
Additional Controls	No	No	No	Yes	Yes

Notes: Results from an instrumental variables regression that uses the (randomly assigned) BDM price as an instrument for purchasing financial advice to estimate the causal impact of financial advice on misallocation. Columns (1) and (2) present the OLS and first-stage estimates, respectively. Column (3) and (5) use optimality of repayments as the outcome variable, that is, if a repayment is fully allocated towards the card with the high interest rate. Column (4) presents the first-stage estimates using optimization ability, financial literacy and gender as additional controls. Column (5) presents the IV estimates with the aforementioned additional controls. Standard errors in parentheses. Errors are clustered at individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Figure B1.2: Effectiveness of Interest Rate Salience using Misallocation Rate

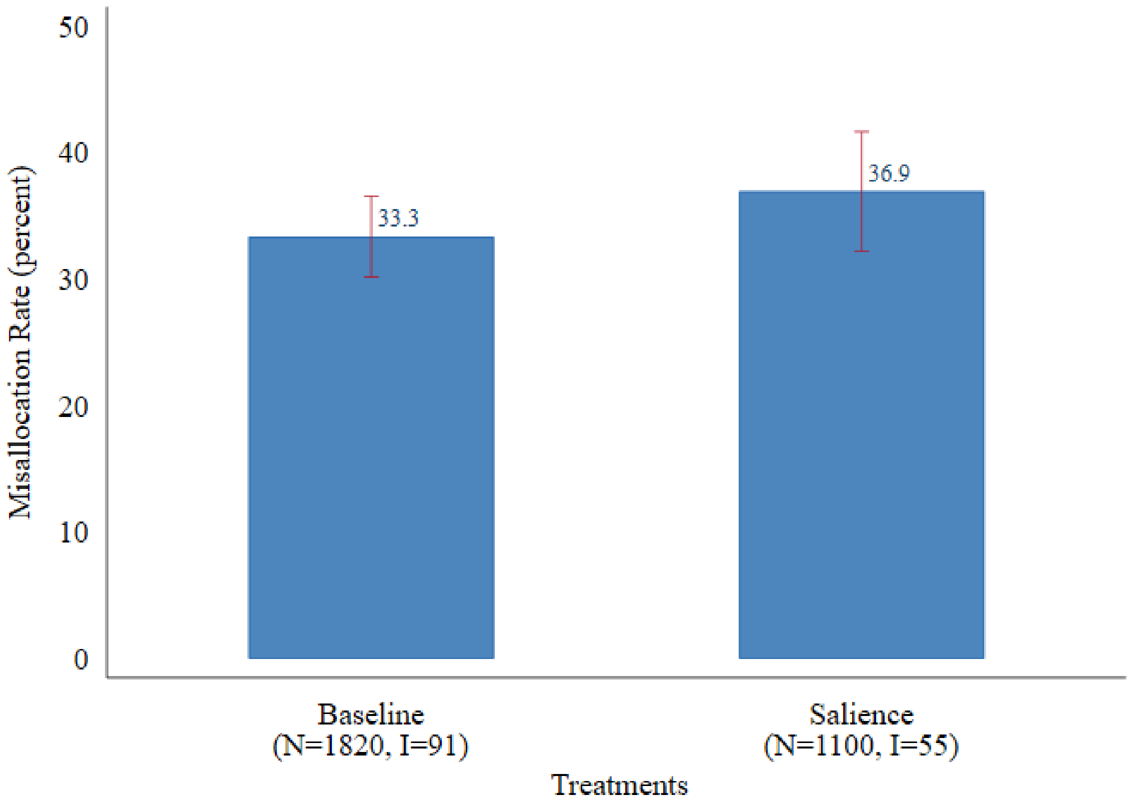

Notes: Panel A shows the average optimality rate by our subjects' ability to solve an algebraic version of the credit card repayment problem. Subjects who are unable to solve the algebraic version of the credit card repayment problem are indicated by the group O. Panel B shows the average optimality rate by our subjects' ability to show the Big Three financial literacy questions. Difference is insignificant ($p = 0.21$). Subjects who are able to solve *all* big three questions are indicated by group 1. The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Figure B1.3: Effectiveness of Interest Rate Salience by Optimization Ability using Misallocation Rate

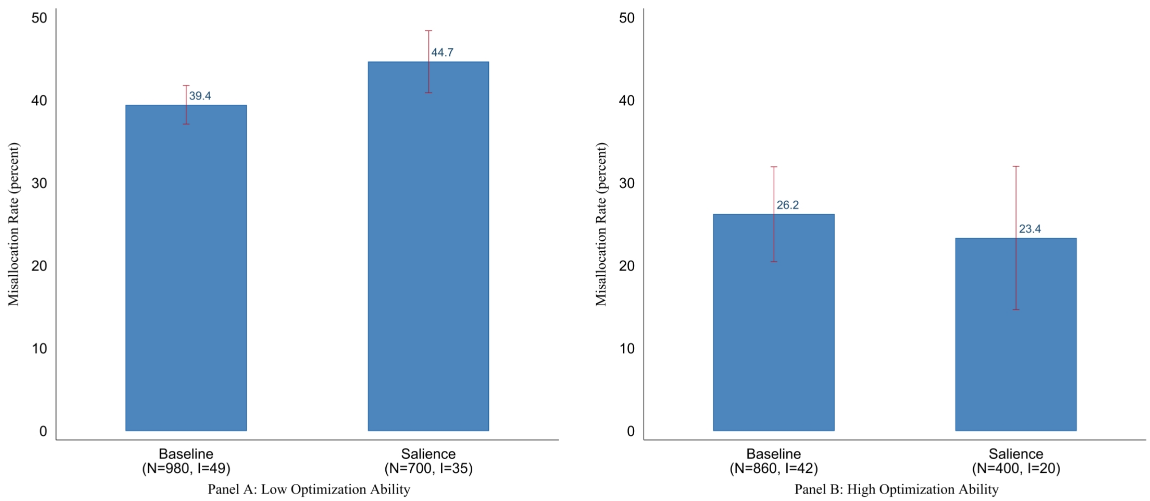

Notes: Panel A shows the misallocation rate for **Baseline** and **Salience** treatments among subjects who fail to solve an algebraic version of the credit card repayment problem. Panel B shows the same rate across the same treatments for subjects who solve an algebraic version of the credit card repayment problem. Difference is significant in Panel A but not in B ($p = 0.017$, $p = 0.63$) The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Figure B1.4: Effectiveness of Interest Rate Salience by Financial Literacy using Misallocation Rate

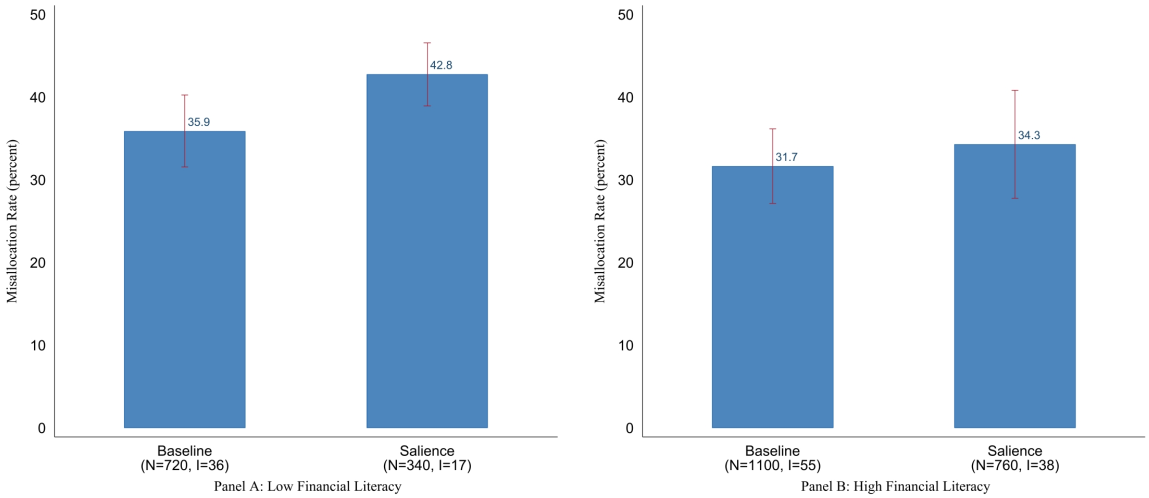

Notes: Panel A shows the misallocation rate for **Baseline** and **Fee Format** treatments among subjects who fail to solve one of the big three financial literacy questions. Panel B shows the same rate across the same treatments for subjects who solve all big three financial literacy questions. Difference is significant in Panel A but not in B ($p = 0.016$, $p = 0.51$). The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Figure B1.5: Effectiveness of Fee Format using Misallocation Rate

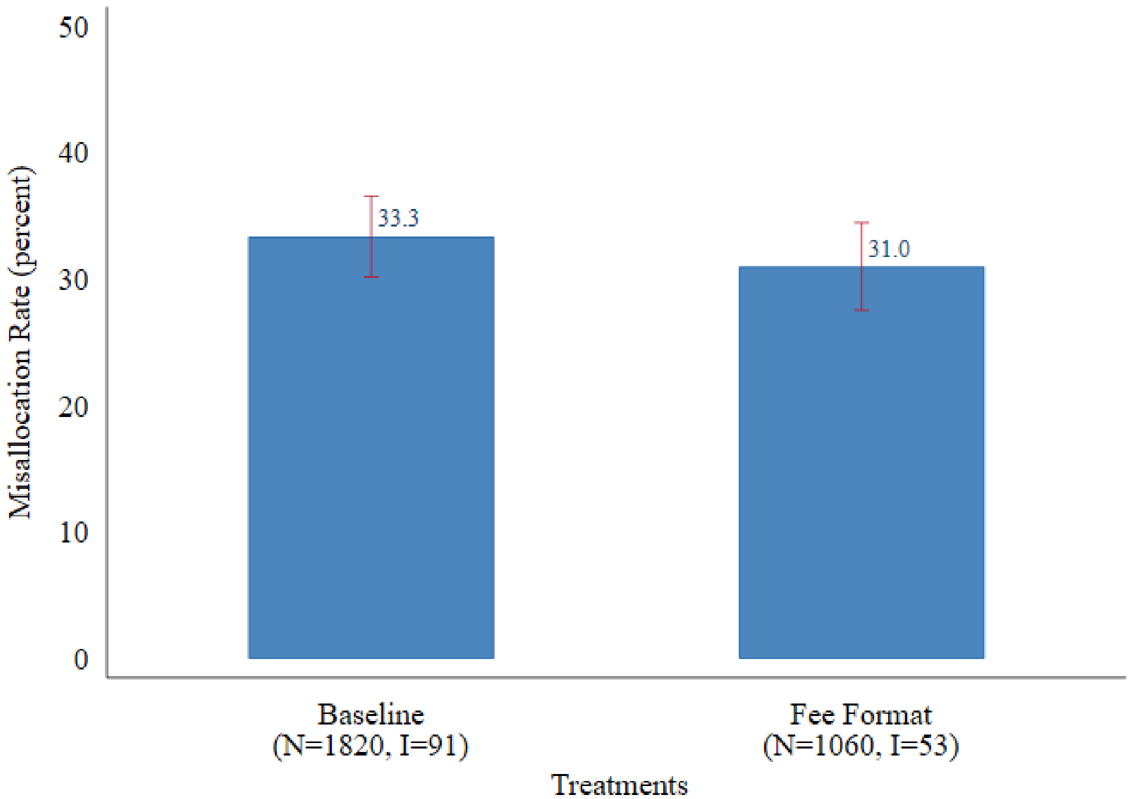

Notes: Panel A shows the misallocation rate by our subjects' ability to solve an algebraic version of the credit card repayment problem. Subjects who are unable to solve the algebraic version of the credit card repayment problem are indicated by the group O. Panel B shows the average optimality rate by our subjects' ability to show the Big Three financial literacy questions. Difference is insignificant ($p = 0.32$). Subjects who are able to solve *all* big three questions are indicated by group 1. The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Figure B1.6: Effectiveness of Fee Format by Optimization Ability using Misallocation Rate

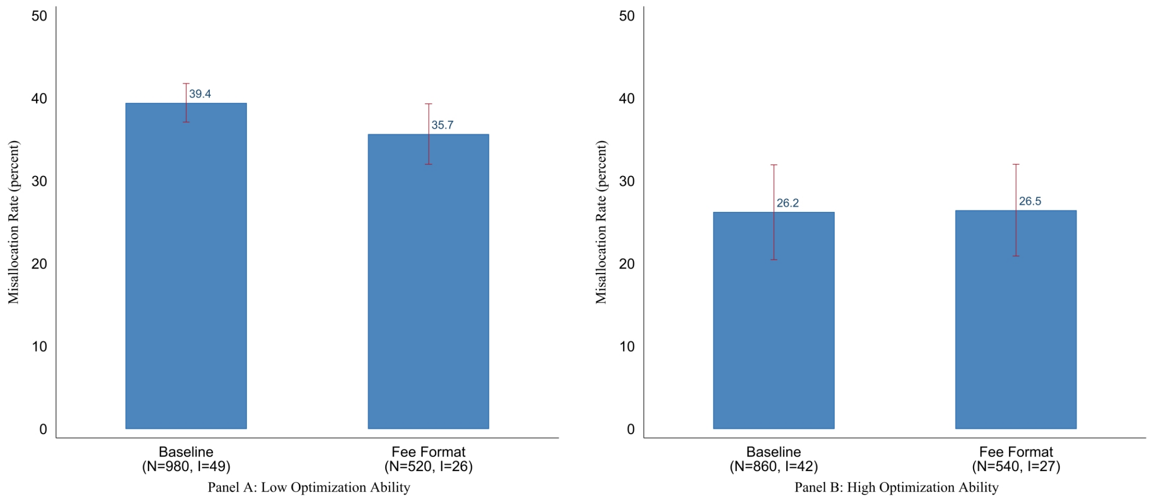

Notes: Panel A shows the average optimality rate by our subjects' ability to solve an algebraic version of the credit card repayment problem. Subjects who are unable to solve the algebraic version of the credit card repayment problem are indicated by the group O. Panel B shows the average optimality rate by our subjects' ability to show the Big Three financial literacy questions. Differences in misallocation rates are insignificant ($p = 0.08$, $p = 0.95$, respectively). Subjects who are able to solve *all* big three questions are indicated by group 1. The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

Figure B1.7: Effectiveness of Fee Format by Financial Literacy using Misallocation Rate

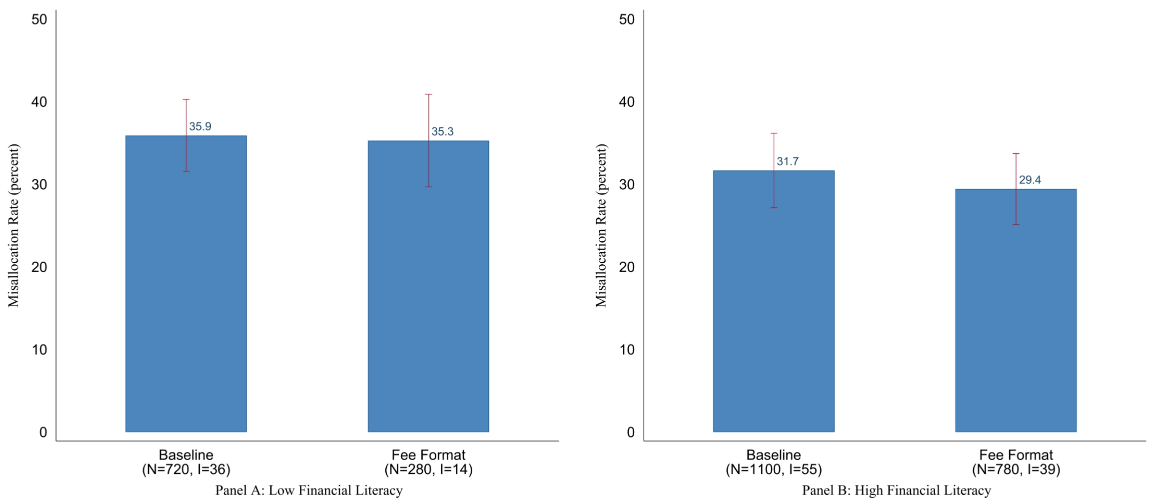

Notes: Panel A shows the average optimality rate by our subjects' ability to solve an algebraic version of the credit card repayment problem. Subjects who are unable to solve the algebraic version of the credit card repayment problem are indicated by the group O. Panel B shows the average optimality rate by our subjects' ability to show the Big Three financial literacy questions. Differences in misallocation rates are insignificant ($p = 0.85$, $p = 0.47$, respectively). Subjects who are able to solve *all* big three questions are indicated by group 1. The number of observations and the number of individuals in each group is indicated by N and I respectively. The unit of observation is subject-by-period. The whiskers indicate 95% confidence intervals. Errors are clustered at the subject level.

B.2 Experiment Instructions

INSTRUCTIONS

Welcome

You are about to participate in a decision-making experiment. In this experiment, you have the ability to earn a considerable amount of money, which will be paid to you through Venmo at the end of the experiment. The amount of money you earn will depend on your decisions. Therefore, it is in your best interest that you read these instructions carefully in order to have a clear understanding of the rules of the experiment. If you need assistance, please raise your hand through the Zoom app. The experimenter will answer your question in a private chat.

All necessary information will be provided to you through your computer. Please turn off your cell phone now, and refrain from opening any other programs or browsers on your computer during the experiment.

Economics experiments have a strict policy against deception. The rules you are going to read next will be implemented just as they are written.

The experiment should take no more than 60 minutes.

--

Background

This is a financial decision-making experiment. In this experiment, you will be assigned two credit card accounts and a checking account. The experiment will be divided into stages and periods where you will be asked to make payments toward these credit card accounts.

Experiment Roadmap

The main experiment contains 6 Independent Stages. Each stage consists of 5 payment periods. You will be presented with different credit cards in each stage.

Your Task

At the beginning of each period, you will receive a fixed amount of money, called a deposit, in your checking account. Your task in each period is to make credit card payment decisions, using the amount of money you have available in your checking account.

A Period

There will be multiple periods in the experiment. An experimental period starts when you receive your deposit, and ends when you finalize your payments to each card for that period.

Level of Debt

At the beginning of the first period, each credit card will be assigned a level of debt. From the second period onward, the level of debt will be determined by two factors: interest rates and your previous period's payment decisions for each card. To illustrate this point, consider the following example:

Suppose that you have two credit cards, Left and Right. Your Left Card has a 4% per period interest rate and you owe 2,000 on that card. Your Right Card has a 5% per period interest rate and you owe 1,000 on that card. After you determine your payments on each card, your *Total Credit Card Debt in the following period* will be calculated as

(1+ 4%) (2,000 - Payment to Left Card) + (1+5%) (1,000 - Payment to Right Card)

Your *End of Stage Total Credit Card Debt* will be calculated as above once you make your last payment decision in that stage.

Your Payment

You will have an initial endowment assigned to you at the beginning of each stage. To determine a *Stage Payoff*, we will subtract your End of Stage Total Credit Card Debt from your initial endowment. Your stage payoff will then be converted into US Dollars at the rate of 12.5 ECUs=$1. Only one stage payoff will be randomly selected as your cash payment in the end. All stage payoffs have the same chance of being selected.

Thank you for your participation in this experiment.

Key Features Recap

Setting:	Two credit card accounts
Task:	Make payment decisions on both cards
Duration:	5 periods per stage, 6 stages
Time:	No strict time restriction (as long as total time < 60 mins)
Payoff:	The less the total debt you have at the end of each stage, the more money you will make from the experiment

We will explain how to use the interface next, please wait for further instructions.

Instructions for Hiring a Robo-advisor

In this part of the experiment, you have the opportunity to hire a robo-advisor that will help you with your decisions in the remaining stages for a fee.

If you decide to **HIRE** the robo-advisor, it will give you advice on how to make your payments in a way that minimizes your total debt in the *remaining two stages* and you will earn the ***maximum*** possible stage payoff for those stages if you follow the advice. However, the fee of the robo-advisor will be deducted from your stage payoff for those stages.

If you decide **NOT TO HIRE** the robo-advisor, you will not be provided with advice while making your payments in the *remaining two stages,* and you will ***not*** be charged for the robo-advisor.

--

An Example

Suppose that the maximum possible stage payoff for each of the remaining stages is $15.

Assume that the fee of the robo-advisor is $*F*.

- If you **HIRE** the robo-advisor at this fee, your stage payoff will be $15 - *F*.
- If you **DO NOT HIRE** the robo-advisor, your stage payoff will be an amount between $0 and $15 depending on how you make your payments.

You will now answer some questions that test your understanding of these instructions. If you have any questions, please ask one of the experimenters through the chat.

Instructions for Hiring a Robo-advisor - Continued

Now we would like to give you the opportunity to hire the robo-advisor, but the fee is **NOT FIXED** yet. It will be determined by chance in a game we are about to play.

You will not spend on the robo-advisor any more than you really want to.

You may even be able to hire the robo-advisor for less than you'd be willing to pay.

Here is how the game works:

- The computer will ask you to tell the **HIGHEST** fee you are willing to pay for the robo-advisor.
- Once you enter the fee that you are willing to pay, you will see a pricemeter.
- The range of the pricemeter represents the range of the fees for the robo-advisor.
- Then you will click a button to start the pricemeter and it will **RANDOMLY** stop at a fee.
- The fee where the pricemeter stops on is the fee for the robo-advisor!

--

Here is some more important information on this game:

- If the pricemeter stops at a fee that is **less than or equal to** what you are willing to pay, you will **HIRE** the advisor and you will pay the fee where the pricemeter stopped.
- If the pricemeter stops at a fee that is **more than** what you are willing to pay, then you will **NOT HIRE** the robo-advisor.
- You will only have one chance to play this game and hire a robo-advisor.
- You cannot change how much you are willing to pay after seeing where the pricemeter stopped.

HINT:

- The payoff-maximizing strategy for you in this game is to truthfully tell the computer how much you would like to pay for the robo-advisor.

We will now play a practice round to see how this game works.

Understanding Quiz for Hiring a Robo-Advisor

[Below are the four questions that we ask before the Practice Round begins.]

1. Suppose that the maximum possible stage payoff is $15.
   If the fee of the robo-advisor is $2 and if you decide to hire the robo-advisor at this fee, how much will your stage payoff be in each of the remaining two stages if you follow the advice?

   [Feedback when correct:] Correct! If you hire the robo-advisor and follow the advice, you will earn the highest possible stage payoff which is $15 and you will need to pay the fee which is $2. Hence your stage payoff for each of the remaining stages will be $13.

   ]Feedback when wrong:] Wrong.
   Hint: If you hire the robo-advisor and follow its advice, you will earn the highest possible stage payoff which is $15 and you will need to pay the fee which is $2.

2. *Check all that are true.* Suppose that the fee of the robo-advisor is $2. If you think that you will make $8 in each of the remaining stages **WITHOUT** hiring the robo-advisor and if the maximum possible stage payoff is $15.
   - you will make $13 if you hire the robo-advisor and follow its advice
   - you will make $8 if you hire the robo-advisor and follow its advice
   - it is more profitable to hire the robo-advisor and follow its advice
   - it is less profitable to hire the robo-advisor and follow its advice

   [Feedback when correct:] Correct! If you hire the robo-advisor and follow its advice, you will earn the highest possible stage payoff which is $15 and you will need to pay the fee which is $2. Hence you will make $13 if you hire the robo-advisor and follow its advice.
   Therefore, it is **MORE** profitable to hire the robo-advisor and follow its advice as you will make $13 if you hire and $8 if you do not hire.

   [Feedback when wrong:] Wrong.
   Hint: Remember that you will make $13 if you hire the robo-advisor at this fee and follow its advice.

3. Suppose that the maximum possible stage payoff is $15.
   If the fee of the robo-advisor is $9 and if you decide to hire the robo-advisor at this fee, how much will your stage payoff be in each of the remaining two stages?

[Feedback when correct:] Correct! If you hire the robo-advisor, you will earn the highest possible stage payoff which is $15 and you will need to pay the fee which is $9. Hence your stage payoff for each of the remaining stages will be $6.

[Feedback when wrong:] Wrong.
Hint: If you hire the robo-advisor and follow its advice, you will earn the highest possible stage payoff which is $15 and you will need to pay the fee which is $2.

4. *Check all that are true.* Suppose that the fee of the robo-advisor is $9. If you think that you will make $8 in each of the remaining stages **WITHOUT** hiring the robo-advisor and if the maximum possible stage payoff is $15.
   - you will make $8 if you hire the robo-advisor and follow its advice
   - you will make $6 if you hire the robo-advisor and follow its advice
   - it is more profitable to hire the robo-advisor and follow its advice
   - it is less profitable to hire the robo-advisor and follow its advice

[Feedback when correct:] Correct! If you hire the robo-advisor and follow its advice, you will earn the highest possible stage payoff which is $15 and you will need to pay the fee which is $9. Hence you will make $6 if you hire the robo-advisor.
Therefore, it is **LESS** profitable to hire the robo-advisor as you will make $6 if you hire and $8 if you do not hire.

[Feedback when wrong:] Wrong.
Hint: Remember that you will make $6 if you hire the robo-advisor at this fee and follow its advice.

[Below are the questions that we ask after the practice round.

All questions have the below text displayed at the top of the screen.]

In the practice round for hiring a robo-advisor,
You stated you are willing to pay at most $Z for the robo-advisor.
The pricemeter stopped at $Y and hence determined the fee of the robo-advisor as $Y.

1. Did you get to hire a robo-advisor?
    - Yes
    - No

[For those that did not hire a robo-advisor:]
[Feedback if yes:] Wrong. The fee of the robo-advisor was higher than what you were willing to pay.
[Feedback if no:] Correct!

[For those that hired a robo-advisor:]
[Feedback if yes:] Correct!

[Feedback if no:] Wrong. The fee of the robo-advisor was lower than what you were willing to pay.

2. [For those that did not hire a robo-advisor:]

A) Do you wish that you had hired the robo-advisor at this fee?
- Yes, I wish I had hired the robo-advisor at this fee.
- No, I am happy that I did not hire the robo-advisor at this fee.

[Feedback if yes:] Next time, you should choose the maximum amount you really want to pay!

[Feedback if no:] Great!

[For those that hired a robo-advisor:]

B) Do you wish that you had **NOT** hired the robo-advisor at this fee?
- Yes, I wish I didn't hire the robo-advisor at this fee.
- No, I am happy that I hired the robo-advisor at this fee.

[Feedback if No:] Great!
[Feedback if Yes:] Next time, you should not choose an amount that is greater than what you really want to pay!

3. ***[For those that did not hire a robo-advisor:]***

   A) If the pricemeter stopped on $X instead of $Y (X should be chosen by the computer uniform-randomly to something **less than or equal to** Z) would you have had hired the robo-advisor?
      - Yes
      - No

      [Feedback if yes:] Correct!

      [Feedback if no:] Wrong. If the pricemeter stopped on $Y, this would mean that the fee of the robo-advisor is now $Y which is cheap enough for you to hire it.

[***For those that did hire a robo-advisor:]***

   B) If the pricemeter stopped on $X instead of $Y (X should be chosen by the computer randomly to something **greater than** Z) would you have had hired the robo-advisor?
      - Yes
      - No

      Feedback if yes: Wrong. If the pricemeter stopped on $Y, this would mean that the fee of the robo-advisor is now $Y which is more expensive than what you are willing to pay!

      Feedback if no: Correct!

B.3 Experiment Interface

Figure B3.1: Introduction to Hire a Robo-Advisor

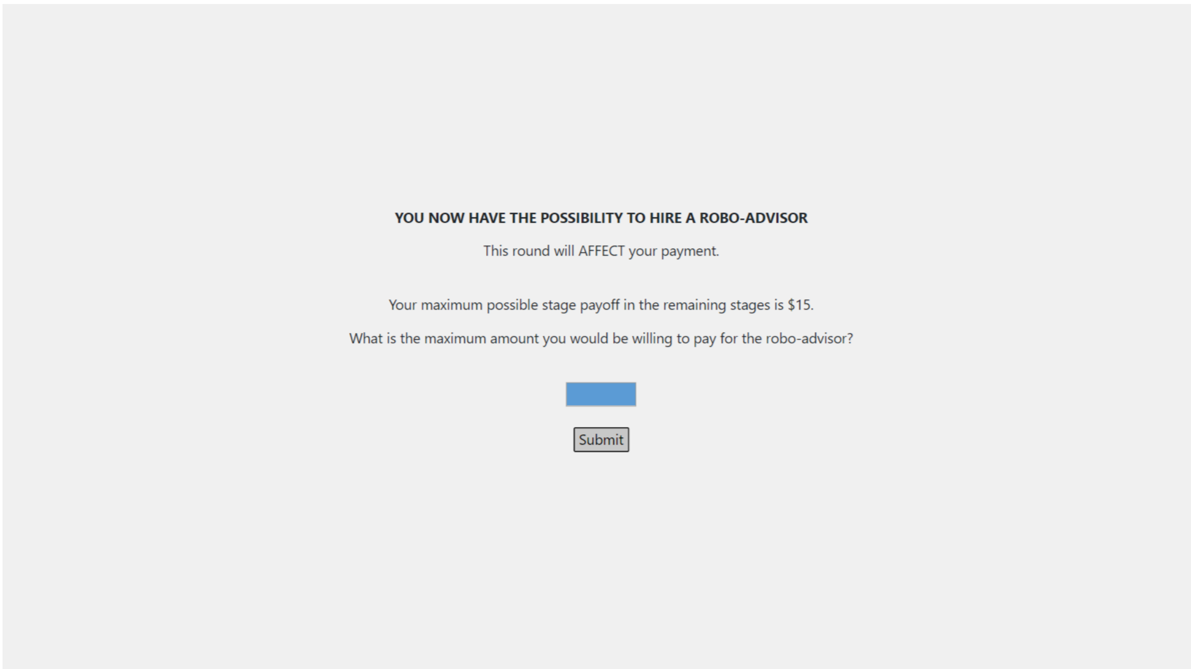

Figure B3.2: Before the Random Price Realization

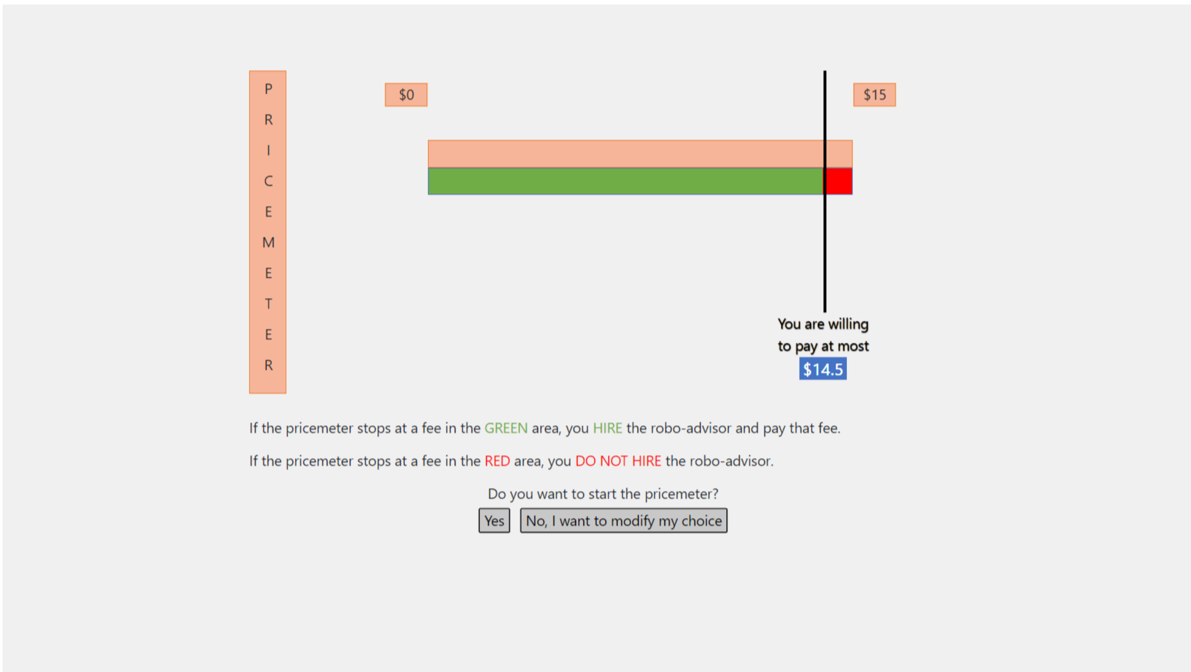

Figure B3.3: Random Price Realization in Real Time

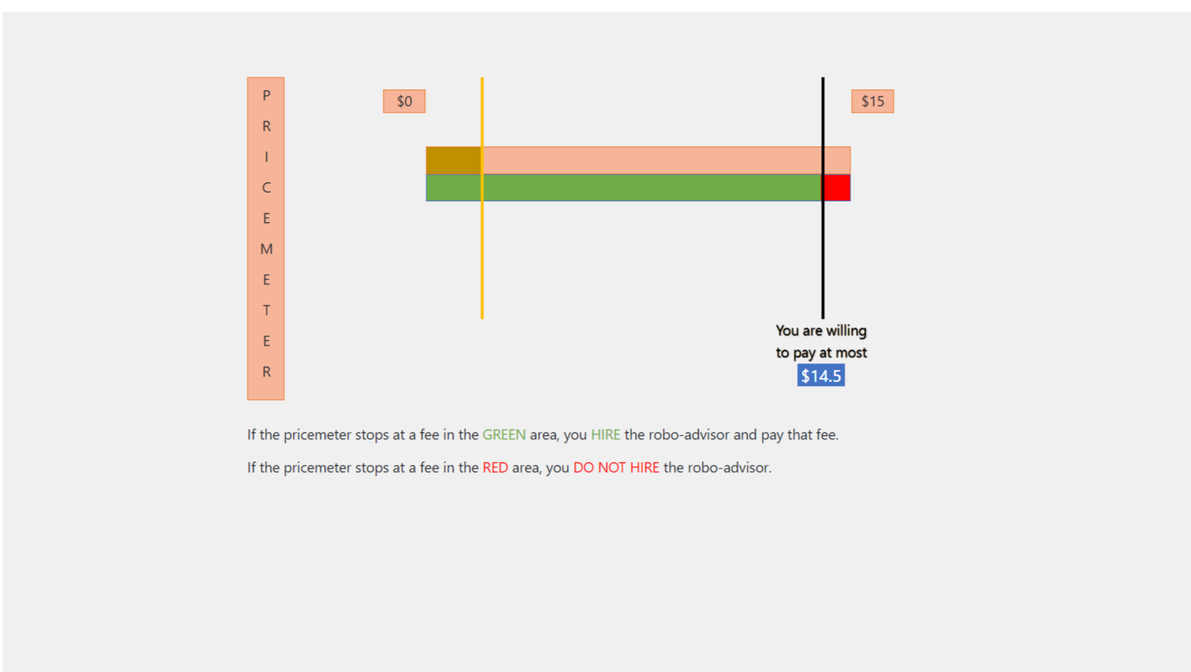

Figure B3.4: After the Random Price Realization - Hired

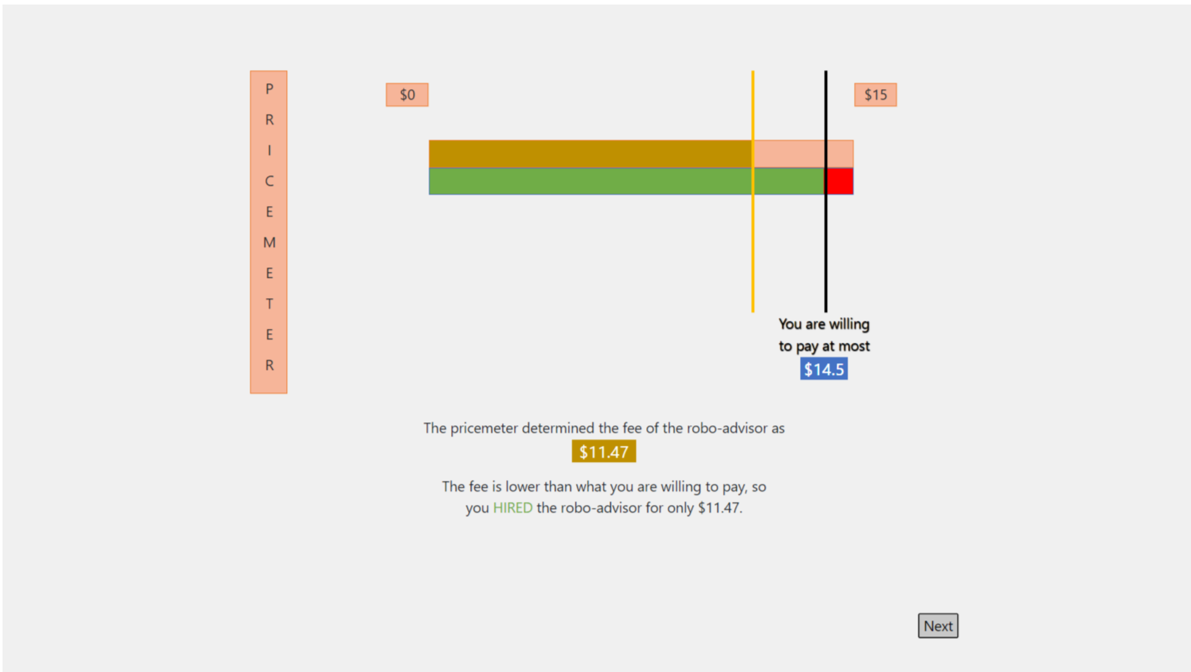

Figure B3.5: After the Random Price Realization - Not Hired

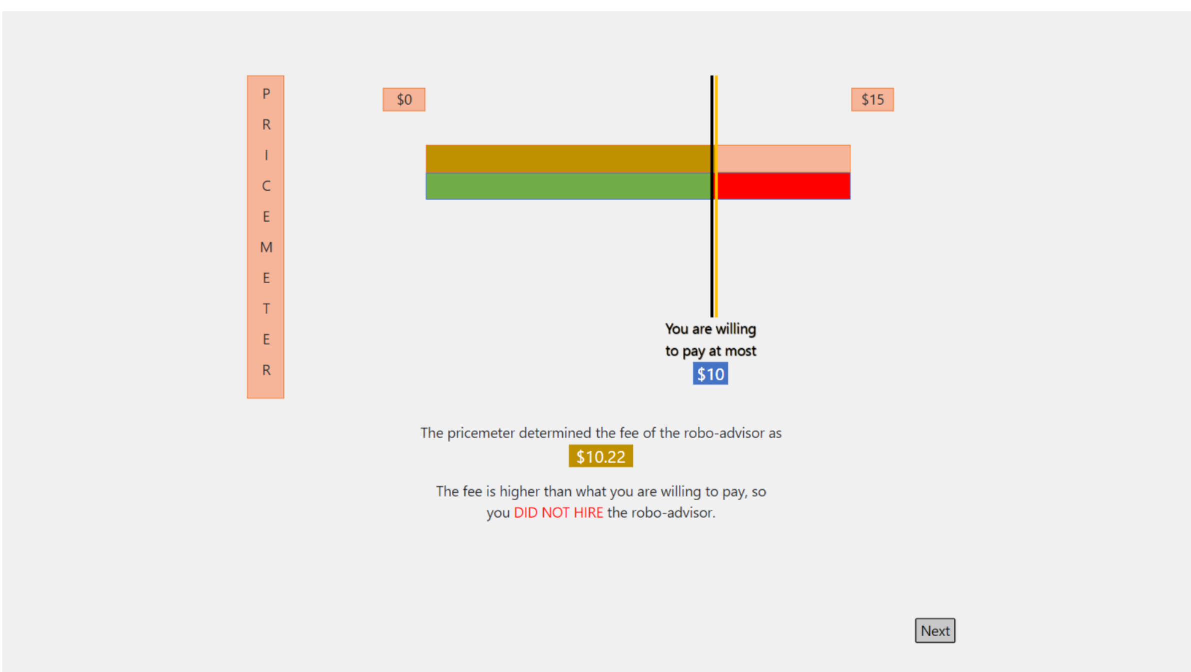

Figure B3.6: After the Random Price Realization - Hired

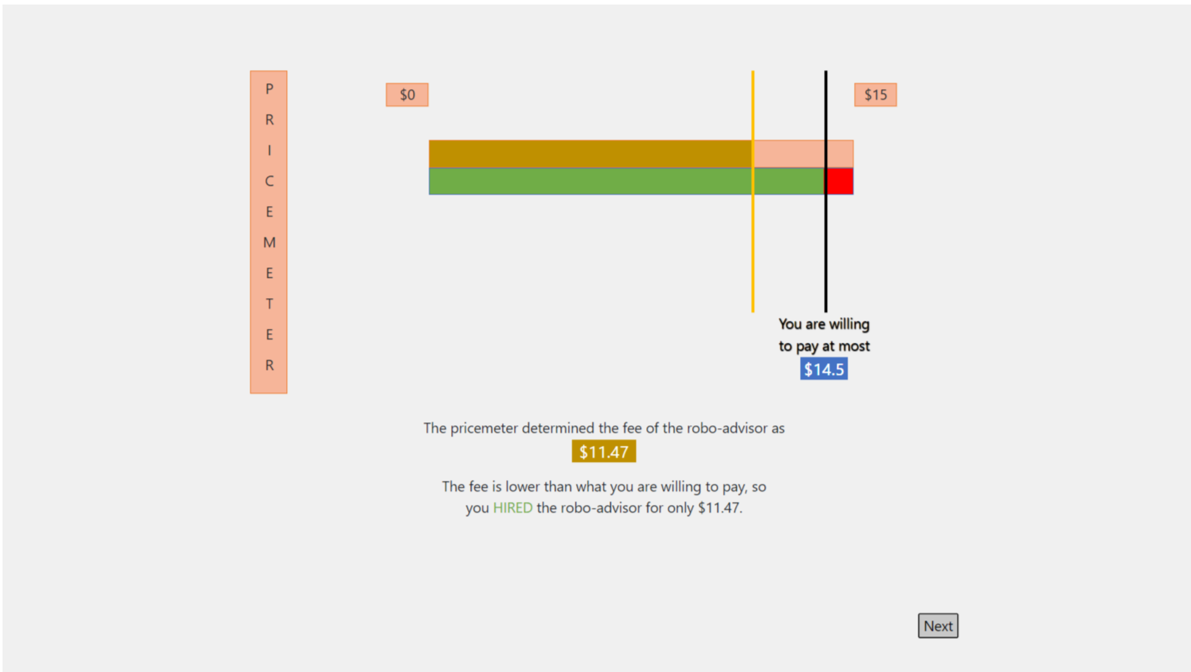

Figure B3.7: Making Decisions with a Robo-advisor

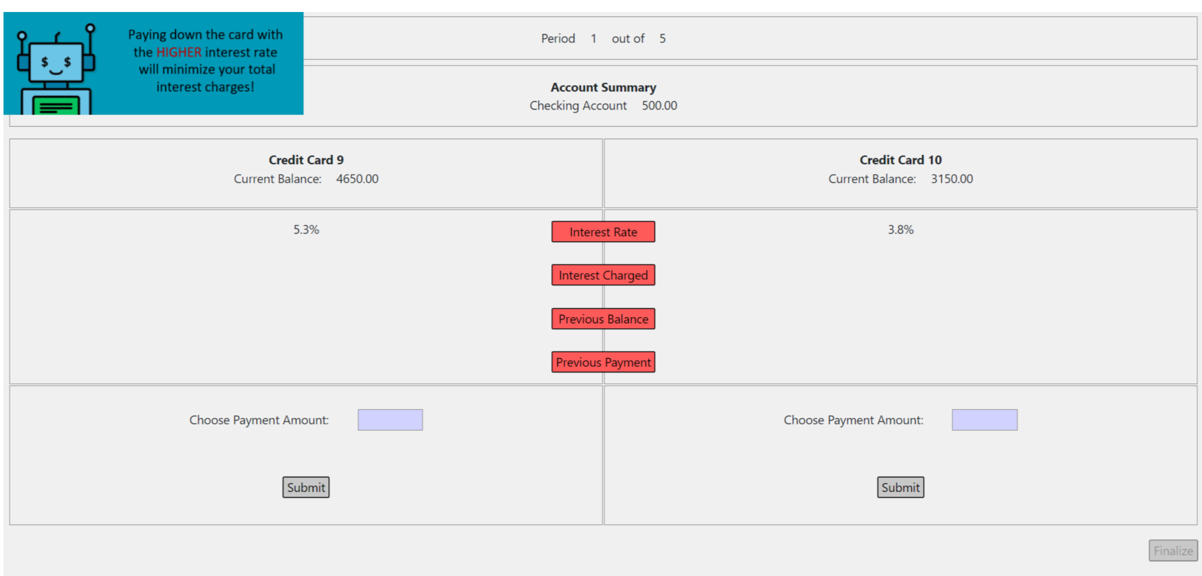

Appendix C

Appendix for *Mental Models and Endogenous Learning*

C.1 Additional Results

Figure C1.1: Evolution of Displacement Relative to the First-Best Optimal Action

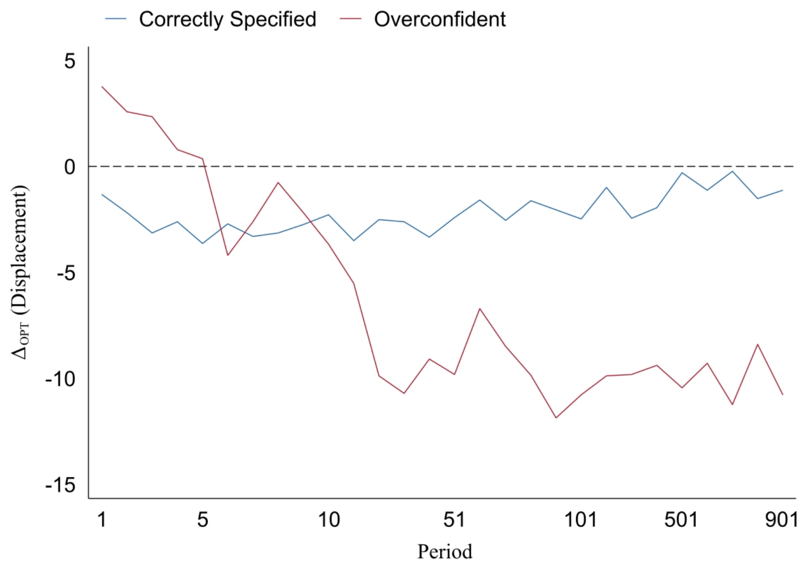

Notes: Figure shows the average displacement relative to the first-best optimal action for correctly specified and overconfident agents across periods. Each observation in a period corresponds to an individual action. The observations within a period are aggregated across treatments **Exogenous** and **Endogenous.**

Figure C1.2: Evolution of Displacement Relative to the Simulated Bayesian Action

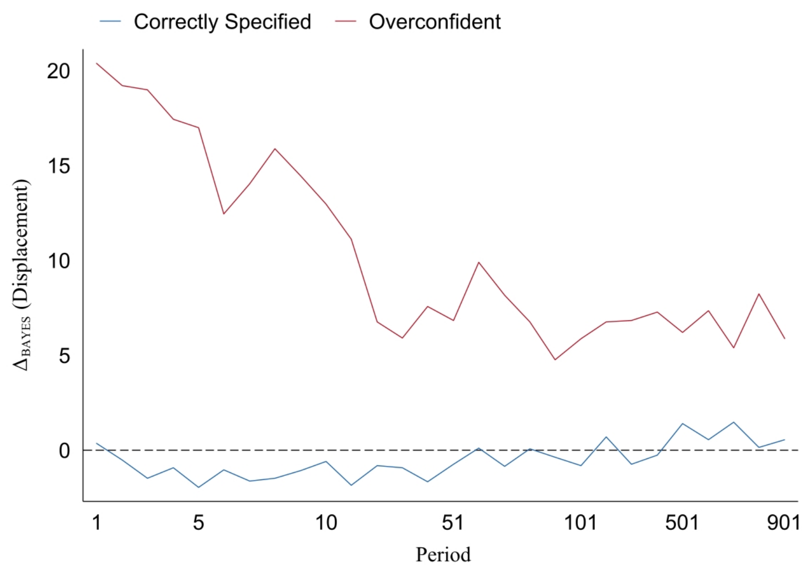

Notes: Figure shows the average displacement relative to the simulated Bayesian action for correctly specified and overconfident agents across periods. Each observation in a period corresponds to an individual action. The observations within a period are aggregated across treatments **Exogenous** and **Endogenous.**

Figure C1.3: Evolution of Actions for Overconfident Subjects

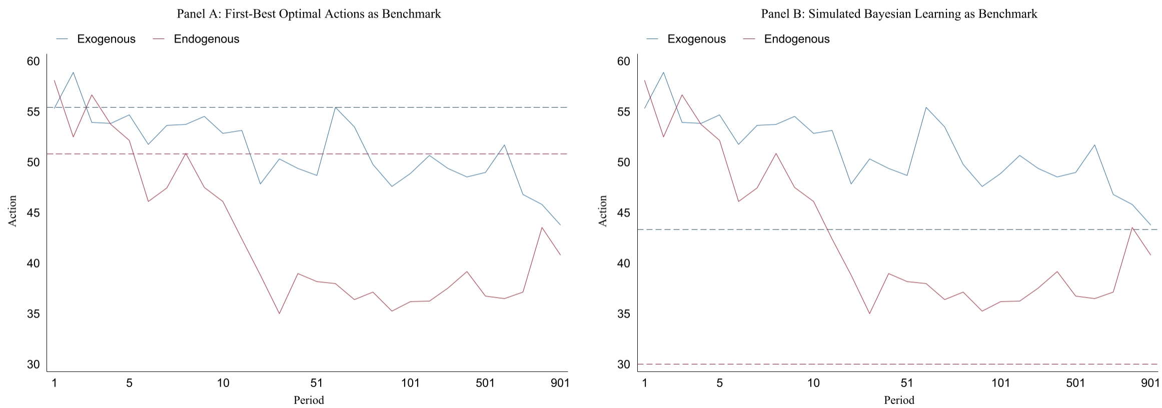

Notes: Both Panel A and Panel B show the average action separately for overconfident subjects in **Exogenous** and **Endogenous** across periods. Panel A presents the first-best optimal action as a benchmark. The blue dashed line in Panel A represents the average first-best optimal action for overconfident subjects in **Exogenous**. The red dashed line in Panel A represents the average first-best optimal action for overconfident subjects in **Endogenous**. Panel B presents the average action a myopically optimizing Bayesian agent would take in the last period of the experiment as a benchmark. The simulations are conducted using each subject's prior beliefs about their abilities. The blue dashed line in Panel B represents the average simulated Bayesian action for overconfident subjects in **Exogenous** whereas the red dashed line in Panel B represents the average simulated Bayesian action for overconfident subjects in **Endogenous**.

Figure C1.4: Evolution of Actions for Correctly Specified Subjects

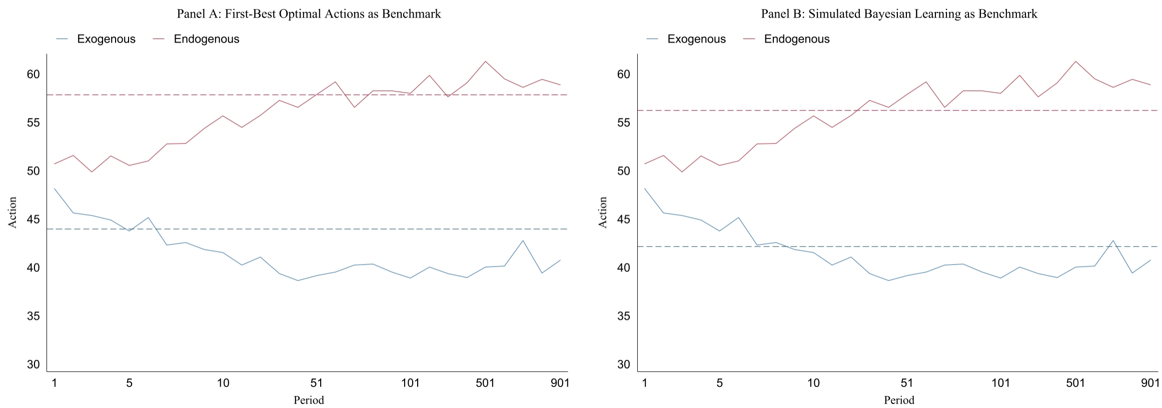

Notes: Both Panel A and Panel B show the average action separately for correctly specified subjects in **Exogenous** and **Endogenous** across periods. Panel A presents the first-best optimal action as a benchmark. The blue dashed line in Panel A represents the average first-best optimal action for correctly specified subjects in **Exogenous**. The red dashed line in Panel A represents the average first-best optimal action for correctly specified subjects in **Endogenous**. Panel B presents the average action a myopically optimizing Bayesian agent would take in the last period of the experiment as a benchmark. The simulations are conducted using each subject's prior beliefs about their abilities. The blue dashed line in Panel B represents the average simulated Bayesian action for correctly specified subjects in **Exogenous** whereas the red dashed line in Panel B represents the average simulated Bayesian action for correctly specified subjects in **Endogenous**.

Table C1.1: Estimation of Displacement Relative to the Simulated Bayesian Actions - Overconfident Subjects

	Panel A: Exogenous				Panel B: Endogenous			
	Dependent Variable: Δ_{BAYES}				Dependent Variable: Δ_{BAYES}			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
α	11.95*	5.642*	3.458	0.418	28.08***	6.722*	7.145*	10.81**
	(5.249)	(2.190)	(1.676)	(2.886)	(6.341)	(2.576)	(2.556)	(2.995)
Observations	20	20	20	20	22	22	22	22
Period	1	501	701	901	1	501	701	901

Notes: The table presents the average displacement relative to the first-best optimal action for overconfident agents in **Exogenous** and **Endogenous**. Each column conducts the estimation $\Delta_{BAYES} = \alpha + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action.

Table C1.2: Estimation of Displacement Relative to the Simulated Bayesian Actions - Correctly Specified Subjects

	Panel A: Exogenous				Panel B: Endogenous			
	Dependent Variable: Δ_{BAYES}				Dependent Variable: Δ_{BAYES}			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
α	5.972	-2.152	0.590	-1.401	-5.526	5.099*	2.381	2.634
	(4.842)	(2.978)	(3.318)	(2.364)	(4.736)	(2.415)	(2.060)	(1.700)
Observations	44	44	44	44	42	42	42	42
Period	1	501	701	901	1	501	701	901

Notes: The table presents the average displacement relative to the first-best optimal action for correctly specified agents in **Exogenous** and **Endogenous**. Each column conducts the estimation $\Delta_{BAYES} = \alpha + \varepsilon$ for the indicated period. Each observation in a period corresponds to an individual action.

C.2 Experiment Instructions

INSTRUCTIONS

Welcome

You are about to participate in a decision-making experiment. In this experiment, you can earn a considerable amount of money, which will be paid to you through Venmo at the end of the experiment. The amount of money you earn will depend on your decisions. Therefore, it is in your best interest that you read these instructions carefully. If you need assistance, please raise your hand through the Zoom app. The experimenter will answer your question in a private chat.

The experiment consists of four parts. One of these parts will be randomly selected for payment at the end of the experiment. In the part that is randomly selected for payment, you can make either $25 or $0. In addition to your earnings from the experiment, you will receive a show-up fee of $10 for participating in the experiment. This means that at the end of the experiment you will receive either a payment of $35 (if in the randomly selected part you made $25) or $10 (if in the randomly selected part you made $0).

For each part of the experiment, you will be precisely instructed about your task.

Please put away your cell phone and do not interact with other participants throughout the experiment.

Instructions for Part 1

You will go through an IQ test in this part of the experiment. Tests similar to this are frequently used to measure intelligence.

The test consists of 10 questions, and you have 10 minutes to solve them. You should solve as many of the 10 questions as possible. Your earnings from this part of the experiment will be either $25 or $0. At the end of the experiment, we will randomly select one of your answers to the IQ test. If the selected answer is correct, you will earn $25 from this part of the experiment. This means the higher the number of correct answers, the more likely you will make $25 in this part of the experiment.

Instructions for Part 2

We conducted the exact same IQ test with other participants who previously, exactly like you, participated in an experiment at UCSB Experimental and Behavioral Economics Laboratory. We randomly selected 19 of these participants. Together with these 19 participants, you now form a group of 20 participants.

We constructed a ranking of this group based on the IQ test scores. The group member that scored the highest on the IQ test obtained rank 1. The group member with the second-highest score obtained rank 2, and so on. The group member with the worst performance on the IQ test got rank 20. In case of a draw between group members, the computer randomly decided who received the higher rank.

The computer then assigned you a color based on your ranking in your group. The top scoring members with ranking 1 to 4 are assigned Dark Green, the members with ranking 5 to 8 are assigned Light Green, the members with ranking 9 to 12 are assigned yellow, the members with ranking 13 to 16 are assigned light red, and the bottom scoring members with ranking 17 to 20 are assigned dark red.

How do you think you ranked on the IQ test?

In this part of the experiment, we are interested in how you think you ranked based on your IQ test score within your group of 20.

Your task is to submit your belief about how likely it is that you are assigned the color dark green, light green, yellow, light red or dark red based on the IQ test score rankings.

To indicate your beliefs, you will use a slider. Where you move the slider will represent your best assessment of the likelihood (expressed as a chance out of 100) that you are assigned one of these colors.

We will now go through an Explanation Stage to understand how the sliders work.

In this Explanation Stage, I would like you to enter a hypothetical subject’s beliefs into the system.

Let’s call this subject Ash.

Ash believes that their ranking is somewhere from 1st to 8th. However, they think it is more likely that they ranked from 1st to 4th rather than 5th to 8th. So, Ash believes they are more likely to be assigned the color dark green rather than light green.

Suppose, specifically, that Ash believes they are assigned dark green with a likelihood of 60 out of 100, and they are assigned light green with a likelihood of 40 out of 100.

Let us now enter Ash's beliefs into the system using sliders.

Now I would like you to do another example on your own.

In this exercise, you enter another hypothetical subject's beliefs into the system. Let's call this subject Blake.

Blake believes their likelihood of ranking from 13^{th} to 16th is 40 out of 100, and ranking from 17^{th} to 20th is 60 out of 100.

Please move the sliders to indicate Blake's beliefs and finalize.

Your Payment

You will be paid based on the accuracy of your belief. Your earnings from this part of the experiment will be either 25 or 0 USD, depending on how accurate your belief is about your color assignment based on the IQ test scores. This means the higher the likelihood your belief assigns to your actual color, the more likely you will receive $25.

If you understand this, you can click directly "Next". If you want to know the details of how we calculate your payments, please click "Details".

[Details on Your Payment]
After you state your belief, the computer will randomly draw a number *k*. This number is between 0 and 20,000. (More precisely, this number is drawn from a discrete uniform distribution on the interval from 0 to 20,000.) You will then receive $25 if the sum *S* is smaller or equal to *k* where *S* is the sum of the following elements:

- The squared deviation between the likelihood (out of 100) that you allocated to your actual color and 100 points.
- For *each* possible color that is not your actual color: The squared deviation between 0 points and the number of points that you allocated to that color.

The exact formula that we use to determine S is

$$S = \sum_{c \in C} (I(YourColor = c) \times 100 - L_c)^2$$

Where $I(YourRanking = c)$ is an indicator function that takes the value 1 if your color is c, $C = \{DarkGreen, LightGreen, Yellow, LightRed, DarkRed\}$ is the set of colors and L_i is the likelihood (out of 100) that you assign to color c.

While this formula might look complicated, the basic idea is very simple: you can secure the largest chance of winning $25 by reporting your most accurate belief about your assigned color.

Instructions for Part 3

Welcome to the main part of the experiment!

Introduction

At the beginning of this part of the experiment, we will assign a project to you. Your job is to act as a project manager for a company. Specifically, we will ask you to repeatedly recommend investment decisions to the company on your assigned project to maximize total profits from this project over multiple periods. The higher the profit you generate from this project, the more likely you will earn $25 from this part of the experiment.

Information on Projects

Projects have different qualities. Higher quality projects generate more profits.

Some projects are of higher quality than others. The project qualities can be any whole number between 0 and 100. The lowest possible project quality is 0, and the highest is 100. The higher the quality of your assigned project, the higher your profits are from the project.

You cannot choose or change your assigned project.

Although a high-quality project increases your profits, you cannot choose your project or its quality. We will randomly assign a project to you at the beginning of the experiment. You will be working on the same project that we assigned you until the end of the experiment.

You will not know your assigned project's quality.

The qualities of the projects vary between 0 and 100. All you will know about your assigned project's quality is that it can be any whole number between 0 and 100. You will not know your assigned project's quality until the end of the experiment.

Recap:

- You act as a project manager for a company in this part of the experiment
- We randomly assign you a project at the beginning of this part of the experiment
- You do not get to choose the project or its quality
- You work on the same project until the end of the experiment
- You repeatedly recommend to the company how much to invest into your assigned project

- Your assigned project's quality can be any whole number between 0 to 100, each number equally likely
- You do not know your project's quality until the end of the experiment
- Your goal is to maximize profits from the project you are assigned to

How do you maximize the profit from your assigned project in each period?

To make things easy for you, we designed the experiment so that it is straightforward to maximize the profit from your assigned project. You maximize your profit in each period by **recommending an investment amount that exactly matches what you think your project's quality is**.

Example.

Suppose you believe that your project's quality is 50. Then, you should recommend an investment amount of 50 to maximize the profit from the project in that period. Similarly, suppose you believe your project's quality is 74. In that case, you should recommend an investment amount of 74 to maximize the profit from the project in that period.

You can make sense of this simple profit-maximizing rule in the following way. If you have a high-quality project, you are better off investing a lot into that project as the return on that project is high. On the other hand, if you have a low-quality project, you are better off not investing too much into the project as the return on that project is low. Therefore, the profit-maximizing strategy is to ***match your recommended investment amount with the quality of the project***.

\---

Now we will go through the details of calculating your profit when you recommend an investment decision. The instructions we will go through in the following pages might seem complex. However, please remember that we will <u>NOT</u> ask you to solve complex equations to maximize your profit during the experiment. The reason we provide these details is to ensure that you have a complete understanding of the experiment's rules.

The idea behind profit maximization is straightforward: <u>**recommend the investment amount that matches what you believe your project's quality is.**</u> You do not need to worry about maximizing your profit as long as you match your investment amount to what you think your project's quality is.

Please feel free to ask any questions along the way.

Details of Profit Calculations

The way we calculate profit in each period is straightforward. First, we will calculate the income you generate from the project and subtract the investment cost to calculate the profit. Then, we add a bonus of 5000 to ensure that no one ends up with a negative profit in the experiment.

$$Profit = ProjectIncome - InvestmentCost + 5000$$

As you see, profits have two main components: project income and investment cost. We will now go through each of these components individually.

Step 1: Project Income

1. ***Project Quality*** refers to the intrinsic quality of the project:
   - It will be a whole number between 0 and 100 in the experiment.
   - Higher quality projects generate higher incomes
2. ***Investment Amount*** refers to the amount you recommend the company to invest into the project:
   - You can choose any number between 0 and 100 as your Investment Amount
   - The higher the amount that is invested into your project, the higher the project income you generate
3. ***Your IQ Rank Score*** refers to your ranking in the IQ test you have completed in the previous part of the experiment. The Blue Table below describes how each ranking corresponds to a score:
   - The higher your ranking in the IQ test you have completed at the beginning of the experiment, the higher the project income you generate

Your Ranking	1-4	5-8	9-12	13-16	17-20
Your IQRANKSCORE	100	80	60	40	20

Specifically, we calculate the project income using the following equation:

$$ProjectIncome = ProjectQuality \times (InvestmentAmount + YourIQRankScore)$$

Step 2: Investment Cost

Investments you recommend to be made into the project have costs. The higher the amount you recommend to be invested, the higher the investment cost.

Specifically, we calculate the investment cost using the following equation:

$$InvestmentCost = \frac{(InvestmentAmount)^2}{2}$$

We can rewrite the full profit equation as

$$Profit = ProjectQuality \times (InvestmentAmount + YourIQRankScore) - \frac{(InvestmentAmount)^2}{2} + 5000$$

The green part of the profit equation is the income from the project and the red part of the profit equation is the investment cost.

If you have taken calculus, you can verify that the profit is maximized when you match the investment amount to the project's quality

$$InvestmentAmount = ProjectQuality$$

Recap:

- You maximize your profit in each period by matching your recommended investment amount to what you believe the project's quality is
- Higher quality projects generate higher profits
- We calculate an IQ rank score for you based on your ranking on the IQ test
- Higher IQ rank score generates higher profits

Will the company follow your investment recommendations immediately?

The company originally planned to invest an amount of 100 in each period on your project before your assignment. However, the company will immediately implement your recommended investment decisions and choose the amount you recommend in each period rather than the originally planned investment amount of 100.

Will you know how much profit you make after each investment decision?

A crucial point in the experiment is that you will not know how much profit you make after each investment decision. Instead, you will get an evaluation from your company if your profit is above or below your company's profit expectation. Since the company immediately implements your recommendations in each period during the experiment, the evaluations you get from the company will be based on your investment recommendations, not based on the company's originally planned investment amount of 100.

Recap:

- The company immediately implements your recommended investment decisions
- You will not know how much profit you make after each investment decision
- ...but you will know if you beat your company's profit expectation or not
- The evaluation you get from the company during the experiment is based on your recommended amount

How does the company determine its profit expectation?

The lowest possible profit you can generate in the experiment is 0, and the highest is 20,000. In each period, the company will randomly choose a profit amount, call it X, from the lowest possible profit amount (0) to the highest one (20,000) to expect from your project. If your profit is at or above X in a period, you beat your company's profit expectation. If your profit is below X, you do not meet your company's profit expectation.

Note that the higher your profit, the more likely you beat your company's profit expectation.

Will you know your company's profit expectations while making your decisions?

You will not know your company's profit expectation X before or after making your investment decision. The only information we will provide is if the profit you generate is above or below this profit expectation X.

What happens when you beat your company's profit expectation?

Once you make your last decision in the experiment, we will randomly select a period. If the profit based on your recommended investment decision beats the company's profit expectation in the randomly selected period, you earn $25 from this part of the experiment!

Recap:

- Your company chooses a number X between 0 and 20,000 as its profit expectation in each period, you will not know what X is
- The higher your profit, the more likely you beat your company's profit expectation
- If you beat your company's profit expectation in a randomly selected period, you earn $25 from this part of the experiment

We have established that you maximize profit in a period by matching your recommended investment amount with your project's quality. However, you do not know what your project's quality is! We will now go through how you can make some sophisticated guesses about your project's quality.

How can you make sophisticated guesses about your project's quality?

You can use your company's profit feedback to help you better understand your IQ rank score and your project's quality. Remember that your profits increase with your IQ rank score and your project's quality. <u>Hence any feedback that tells you that you beat the company's profit expectation is good news for your IQ rank score and your project's quality</u>.

<u>On the other hand, any feedback that tells you that you did not beat the company's expectations is bad news for your IQ rank score and your project's quality.</u>

To help you interpret the feedback that you get from the company, we will provide you with an expert statistician. In each period, the statistician will prepare a report for you, which you can use to make sophisticated guesses about your project's quality.

The Statistician's Report

These reports are going to look like the one on this page:

The Statistician's Report	
Your IQ Rank Score	**Project Quality**
20	75
40	71
60	69
80	66
100	65

The report is very straightforward to read. The statistician tells you:

- If your IQ rank score is 20, you should expect your project's quality to be 75.
- If your IQ rank score is 40, you should expect your project's quality to be 71, and so on.

Depending on what you believe your IQ rank score is, you may then make a sophisticated guess about your project's quality.

The statistician will update the report in each period incorporating the evaluations you receive from your company up until that period.

Instructions for Part 4

Please remember that at the beginning of the experiment, we assigned each participant in this session to a group with 19 other randomly selected people who had previously taken the same IQ test at UCSB Experimental and Behavioral Economics Laboratory.

We then constructed a ranking of each group based on the IQ test scores, and the computer assigned you a color based on your ranking in your group of 20.

The top scoring members with ranking 1 to 4 are assigned Dark Green, the members with ranking 5 to 8 are assigned Light Green, the members with ranking 9 to 12 are assigned yellow, the members with ranking 13 to 16 are assigned light red, and the bottom scoring members with ranking 17 to 20 are assigned dark red.

In this part of the experiment, we are again interested in how you think you ranked based on your IQ test score within your group of 20.

Your task is to submit your belief about how likely it is that you are assigned the color dark green, light green, yellow, light red or dark red based on the IQ test score rankings.

To indicate your beliefs, you'll use a slider exactly as before.

Your Payment

You will be paid based on the accuracy of your belief. Your earnings from this part of the experiment will be either 25 or 0 USD, depending on how accurate your belief is about your color assignment based on the IQ test scores. This means the higher the likelihood your belief assigns to your actual color, the more likely you will receive $25.

If you understand this, you can click directly "Next". If you want to know the details of how we calculate your payments, please click "Details".

[Details are identical to Part 2's Payment Details]

Understanding Quiz

Question 1.

Do you get to choose your project or its quality in the experiment?

- Yes, I choose both the project and its quality
- No, I do not get to choose either the project or its quality
- I only choose the project, but I don't get to choose its quality
- I do not get to choose the project, but I choose its quality

Question 2.

What type of decisions do you make on the project you are assigned to?

- I repeatedly give recommendations on how much the company should invest in the project
- I repeatedly give recommendations on how many projects the company should undertake
- I repeatedly give recommendations on whom to delegate the project
- I repeatedly give recommendations on which company should be responsible for the project

Question 3.

When do you learn your assigned project's quality?

- At the beginning of the experiment
- After my first investment decision
- Before my last investment decision
- At the end of the experiment

Understanding Quiz II

Question 1.

How do you maximize your profit in a period in the experiment?

- I match my recommended investment amount to the project's quality
- I match my recommended investment amount to my IQ Rank score
- I match my recommended investment amount to investment cost
- None of the above

Question 2.

Suppose you believe your project's quality is 62 in a period.

What investment amount maximizes your profits in that period?

- 31
- 62
- 93
- Not enough information to answer this question

Question 3.

Which of the below factors increase profits? [Multiple choice available.]

- Project's Quality
- My IQ Rank Score
- Investment Cost

Question 4.

What is your IQ rank score if you rank 1st in your group on the IQ test you have previously taken?

- 0
- 20
- 60
- 100

Question 5.

What is your IQ rank score if you rank 5th in your group on the IQ test you have previously taken?

- 0
- 20
- 60
- 100

Understanding Quiz III

Question 1.

When will the company implement your recommended investment decisions for each period?

- Immediately after I make my decisions
- Once I make my last decision
- At the beginning of the experiment, before I make any decision
- After I make my first decision, but before my last decision

Question 2.

What will we tell you after each investment decision you make?

- How much profit I make
- An evaluation from the company if my profit is higher than the company's profit expectation or not
- My assigned project's quality
- My IQ rank score

Question 3.

<u>Before you make your last decision in the experiment</u>, the evaluations you get from the company are based on which investment decisions?

- My recommended investment decisions
- The company's originally planned investment amount of 100
- Neither my recommended investment decisions nor the company's originally planned investment amount of 100
- Both my recommended investment decisions and the company's originally planned investment amount of 100

Understanding Quiz IV

Question 1.

How does the company choose its profit expectation in each period?

- It randomly chooses a number between the lowest and highest possible profit amounts (0 and 20,000)
- It uses historical data
- It uses investment costs
- It uses project's quality

Question 2.

How do we decide to pay you $25 in this part of the experiment?

- If my recommended investment decision generates a profit that beats my company's profit expectation in the first period
- If my recommended investment decision generates a profit that beats my company's profit expectation in the last period
- If my recommended investment decision generates a profit that beats my company's profit expectation in a randomly selected period
- None of the above

Understanding Quiz V

Question 1.

If you beat your company's profit expectations in a period, this is good news for

- Only your IQ rank score
- Only the project's quality
- Both your IQ rank score and the project's quality
- Neither your IQ score nor the project's quality

Question 2.

The Statistician's Report	
Your IQ Rank Score	**Project Quality**
20	75
40	71
60	69
80	66
100	65

If you think your IQ rank score is 60, what should you expect your project's quality to be according to the statistician's report?

- 75
- 71
- 69
- 66

Question 3.

If you think your IQ rank score is 80, what should you expect your project's quality to be according to the statistician's report?

- 75
- 71
- 69
- 66

Bibliography

[1] N. Bhutta, A. Fuster, and A. Hizmo, *Paying too much? price dispersion in the us mortgage market*, .

[2] S. Andersen, J. Y. Campbell, K. M. Nielsen, and T. Ramadorai, *Sources of inaction in household finance: Evidence from the danish mortgage market*, *American Economic Review* (Forthcoming).

[3] A. Ponce, E. Seira, and G. Zamarripa, *Borrowing on the wrong credit card? evidence from mexico*, *American Economic Review* **107** (2017), no. 4 1335–61.

[4] J. Gathergood, N. Mahoney, N. Stewart, and J. Weber, *How do individuals repay their debt? the balance-matching heuristic*, *American Economic Review* **109** (2019), no. 3 844–75.

[5] A. Lusardi and O. S. Mitchell, *The economic importance of financial literacy: Theory and evidence*, *Journal of economic literature* **52** (2014), no. 1 5–44.

[6] A. Lusardi and P. Tufano, *Debt literacy, financial experiences, and overindebtedness*, *Journal of Pension Economics & Finance* **14** (2015), no. 4 332–368.

[7] J. Y. Campbell, *Restoring rational choice: The challenge of consumer financial regulation*, *American Economic Review* **106** (2016), no. 5 1–30.

[8] J. Beshears, J. J. Choi, D. Laibson, and B. C. Madrian, *Behavioral household finance*, Working Paper 24854, National Bureau of Economic Research, July, 2018.

[9] R. Chetty, A. Looney, and K. Kroft, *Salience and taxation: Theory and evidence*, *American economic review* **99** (2009), no. 4 1145–77.

[10] V. Stango and J. Zinman, *Limited and varying consumer attention: Evidence from shocks to the salience of bank overdraft fees*, *The Review of Financial Studies* **27** (2014), no. 4 990–1030.

[11] D. Karlan, M. McConnell, S. Mullainathan, and J. Zinman, *Getting to the top of mind: How reminders increase saving*, *Management Science* **62** (2016), no. 12 3393–3411.

[12] P. Bordalo, N. Gennaioli, and A. Shleifer, *Memory, attention, and choice*, tech. rep., National Bureau of Economic Research, 2017.

[13] B. Handel and J. Schwartzstein, *Frictions or mental gaps: What's behind the information we (don't) use and when do we care?*, *Journal of Economic Perspectives* **32** (2018), no. 1 155–78.

[14] R. E. Nisbett and L. Ross, *Human inference: Strategies and shortcomings of social judgment*, .

[15] S. E. Taylor and S. C. Thompson, *Stalking the elusive" vividness" effect.*, *Psychological review* **89** (1982), no. 2 155.

[16] P. Bordalo, N. Gennaioli, and A. Shleifer, *Salience and consumer choice*, *Journal of Political Economy* **121** (2013), no. 5 803–843.

[17] B. Kőszegi and A. Szeidl, *A model of focusing in economic choice*, *The Quarterly journal of economics* **128** (2012), no. 1 53–104.

[18] S. Soroka, P. Fournier, and L. Nir, *Cross-national evidence of a negativity bias in psychophysiological reactions to news*, *Proceedings of the National Academy of Sciences* **116** (2019), no. 38 18888–18892.

[19] R. F. Baumeister, E. Bratslavsky, C. Finkenauer, and K. D. Vohs, *Bad is stronger than good*, *Review of general psychology* **5** (2001), no. 4 323–370.

[20] D. Kahneman, *Prospect theory: An analysis of decisions under risk*, *Econometrica* **47** (1979) 278.

[21] C. R. Sunstein, *Empirically informed regulation*, *The University of Chicago Law Review* **78** (2011), no. 4 1349–1429.

[22] M. Bertrand and A. Morse, *Information disclosure, cognitive biases, and payday borrowing*, *The Journal of Finance* **66** (2011), no. 6 1865–1893.

[23] E. Seira, A. Elizondo, and E. Laguna-Müggenburg, *Are information disclosures effective? evidence from the credit card market*, *American Economic Journal: Economic Policy* **9** (2017), no. 1 277–307.

[24] J. S. Hastings, B. C. Madrian, and W. L. Skimmyhorn, *Financial literacy, financial education, and economic outcomes*, *Annu. Rev. Econ.* **5** (2013), no. 1 347–373.

[25] X. Gabaix, *A sparsity-based model of bounded rationality*, *Quarterly Journal of Economics* **129** (2014), no. 4 1661–1710.

[26] N. Karlsson, G. Loewenstein, and D. Seppi, *The ostrich effect: Selective attention to information*, *Journal of Risk and uncertainty* **38** (2009), no. 2 95–115.

[27] S. Benartzi and R. Thaler, *Heuristics and biases in retirement savings behavior*, *Journal of Economic Perspectives* **21** (September, 2007) 81–104.

[28] B. J. Keys and J. Wang, *Minimum payments and debt paydown in consumer credit cards*, *Journal of Financial Economics* (2018).

[29] N. Stewart, *The cost of anchoring on credit-card minimum repayments*, *Psychological Science* **20** (2009), no. 1 39–41.

[30] U. Fischbacher, *z-tree: Zurich toolbox for ready-made economic experiments*, *Experimental economics* **10** (2007), no. 2 171–178.

[31] C. A. Sims, *Implications of rational inattention*, *Journal of monetary Economics* **50** (2003), no. 3 665–690.

[32] S. T. Fiske and S. E. Taylor, *Social cognition: From brains to culture.* Sage, 2013.

[33] V. A. Thompson, *Dual-process theories: A metacognitive perspective.*, .

[34] P. N. Johnson-Laird, *Mental models and human reasoning*, *Proceedings of the National Academy of Sciences* **107** (2010), no. 43 18243–18250.

[35] D. Kahneman, *Maps of bounded rationality: Psychology for behavioral economics*, *American economic review* **93** (2003), no. 5 1449–1475.

[36] J. S. B. Evans, *The heuristic-analytic theory of reasoning: Extension and evaluation*, *Psychonomic Bulletin & Review* **13** (2006), no. 3 378–395.

[37] A. Tversky and D. Kahneman, *Judgment under uncertainty: Heuristics and biases*, *Science* **185** (1974), no. 4157 1124–1131.

[38] G. Gigerenzer and W. Gaissmaier, *Heuristic decision making*, *Annual review of psychology* **62** (2011) 451–482.

[39] J. Schwartzstein, *Selective attention and learning*, *Journal of the European Economic Association* **12** (2014), no. 6 1423–1452.

[40] X. Gabaix, *Behavioral inattention*, tech. rep., National Bureau of Economic Research, 2017.

[41] N. H. Frijda, *The emotions.* Cambridge University Press, 1986.

[42] I. P. Levin, S. L. Schneider, and G. J. Gaeth, *All frames are not created equal: A typology and critical analysis of framing effects*, *Organizational behavior and human decision processes* **76** (1998), no. 2 149–188.

[43] N. Herscovics and L. Linchevski, *A cognitive gap between arithmetic and algebra*, *Educational studies in mathematics* **27** (1994), no. 1 59–78.

[44] K. Stacey and M. MacGregor, *Learning the algebraic method of solving problems*, *The Journal of Mathematical Behavior* **18** (1999), no. 2 149–167.

[45] S. Mullainathan and E. Shafir, *Scarcity: Why having too little means so much.* Macmillan, 2013.

[46] A. Lusardi, P.-C. Michaud, and O. S. Mitchell, *Optimal financial knowledge and wealth inequality*, *Journal of Political Economy* **125** (2017), no. 2 431–477.

[47] G. Loewenstein, C. R. Sunstein, and R. Golman, *Disclosure: Psychology changes everything*, *Annu. Rev. Econ.* **6** (2014), no. 1 391–419.

[48] G. Gigerenzer and U. Hoffrage, *How to improve bayesian reasoning without instruction: frequency formats.*, *Psychological review* **102** (1995), no. 4 684.

[49] M. Zaki, *Interest rates: Prices hidden in plain sight*, *Available at SSRN 3168043* (2018).

[50] U. Bhattacharya, A. Hackethal, S. Kaesler, B. Loos, and S. Meyer, *Is unbiased financial advice to retail investors sufficient? answers from a large field study*, *The Review of Financial Studies* **25** (2012), no. 4 975–1032.

[51] F. D'Acunto, N. Prabhala, and A. G. Rossi, *The promises and pitfalls of robo-advising*, *The Review of Financial Studies* **32** (2019), no. 5 1983–2020.

[52] M. Bianchi and M. Brière, *Augmenting investment decisions with robo-advice*, *Université Paris-Dauphine Research Paper* (2021), no. 3751620.

[53] M. Bertrand, D. Karlan, S. Mullainathan, E. Shafir, and J. Zinman, *What's advertising content worth? evidence from a consumer credit marketing field experiment*, *The quarterly journal of economics* **125** (2010), no. 1 263–306.

[54] J. J. Choi, D. Laibson, and B. C. Madrian, *Why does the law of one price fail? an experiment on index mutual funds*, *The Review of Financial Studies* **23** (2010), no. 4 1405–1432, [/oup/backfile/content$_public/journal/rfs/$23/4/10.1093$_rfs_hhp$097/1/hhp097.pdf].

[55] J. J. Choi, D. Laibson, and B. C. Madrian, *$100 bills on the sidewalk: Suboptimal investment in 401 (k) plans*, *Review of Economics and Statistics* **93** (2011), no. 3 748–763.

[56] P. Adams, S. Hunt, C. Palmer, and R. Zaliauskas, *Testing the effectiveness of consumer financial disclosure: Experimental evidence from savings accounts*, *Journal of Financial Economics* **141** (2021), no. 1 122–147.

[57] B. Ferman, *Reading the fine print: Information disclosure in the brazilian credit card market*, *Management Science* **62** (2016), no. 12 3534–3548.

[58] V. Stango and J. Zinman, *Fuzzy math, disclosure regulation, and market outcomes: Evidence from truth-in-lending reform*, *The Review of Financial Studies* **24** (2011), no. 2 506–534.

[59] S. Agarwal, S. Chomsisengphet, N. Mahoney, and J. Stroebel, *Regulating consumer financial products: Evidence from credit cards*, *The Quarterly Journal of Economics* **130** (2015), no. 1 111–164.

[60] S. Kulkarni, S. Truffa, and G. Iberti, *Seatbelts or statements—which loan regulations benefit whom?*, .

[61] H. Ozyilmaz and G. Zhang, *The debt payment puzzle*, .

[62] G. M. Becker, M. H. DeGroot, and J. Marschak, *Measuring utility by a single-response sequential method*, *Behavioral science* **9** (1964), no. 3 226–232.

[63] D. Danz, L. Vesterlund, and A. J. Wilson, *Belief elicitation: Limiting truth telling with information on incentives*, tech. rep., National Bureau of Economic Research, 2020.

[64] T. N. Cason and C. R. Plott, *Misconceptions and game form recognition: Challenges to theories of revealed preference and framing*, *Journal of Political Economy* **122** (2014), no. 6 1235–1270.

[65] P. Bohm, J. Lindén, and J. Sonnegård, *Eliciting reservation prices: Becker–degroot–marschak mechanisms vs. markets*, *The Economic Journal* **107** (1997), no. 443 1079–1089.

[66] A. Lusardi and O. S. Mitchell, *Financial literacy around the world: an overview*, *Journal of pension economics & finance* **10** (2011), no. 4 497–508.

[67] D. Read, G. Loewenstein, M. Rabin, G. Keren, and D. Laibson, *Choice bracketing*, in *Elicitation of preferences*, pp. 171–202. Springer, 1999.

[68] A. Ellis and D. J. Freeman, *Revealing choice bracketing*, *arXiv preprint arXiv:2006.14869* (2020).

[69] D. L. Chen, M. Schonger, and C. Wickens, *otree—an open-source platform for laboratory, online, and field experiments*, *Journal of Behavioral and Experimental Finance* **9** (2016) 88–97.

[70] B. Greiner, *Subject pool recruitment procedures: organizing experiments with orsee*, *Journal of the Economic Science Association* **1** (2015), no. 1 114–125.

[71] R. Hanna, S. Mullainathan, and J. Schwartzstein, *Learning Through Noticing: Theory and Evidence from a Field Experiment **, *The Quarterly Journal of Economics* **129** (06, 2014) 1311–1353, [https://academic.oup.com/qje/article-pdf/129/3/1311/30629812/qju015.pdf].

[72] B. Handel and J. Schwartzstein, *Frictions or mental gaps: What's behind the information we (don't) use and when do we care?*, *Journal of Economic Perspectives* **32** (February, 2018) 155–78.

[73] P. Heidhues, B. Kőszegi, and P. Strack, *Unrealistic expectations and misguided learning*, *Econometrica* **86** (2018), no. 4 1159–1214.

[74] I. Esponda, E. Vespa, and S. Yuksel, *Mental models and learning: The case of base-rate neglect*, tech. rep., 2020.

[75] N. D. Weinstein, *Unrealistic optimism about future life events.*, *Journal of personality and social psychology* **39** (1980), no. 5 806.

[76] O. Svenson, *Are we all less risky and more skillful than our fellow drivers?*, *Acta psychologica* **47** (1981), no. 2 143–148.

[77] D. A. Moore and P. J. Healy, *The trouble with overconfidence.*, *Psychological review* **115** (2008), no. 2 502.

[78] C. Camerer and D. Lovallo, *Overconfidence and excess entry: An experimental approach*, *American economic review* **89** (1999), no. 1 306–318.

[79] B. M. Barber and T. Odean, *Boys will be boys: Gender, overconfidence, and common stock investment*, *The quarterly journal of economics* **116** (2001), no. 1 261–292.

[80] U. Malmendier and G. Tate, *Ceo overconfidence and corporate investment*, *The journal of finance* **60** (2005), no. 6 2661–2700.

[81] U. Malmendier and G. Tate, *Who makes acquisitions? ceo overconfidence and the market's reaction*, *Journal of financial Economics* **89** (2008), no. 1 20–43.

[82] R. Bénabou and J. Tirole, *Self-confidence and personal motivation*, *The quarterly journal of economics* **117** (2002), no. 3 871–915.

[83] B. Köszegi, *Ego utility, overconfidence, and task choice*, *Journal of the European Economic Association* **4** (2006), no. 4 673–707.

[84] F. Zimmermann, *The dynamics of motivated beliefs*, *American Economic Review* **110** (2020), no. 2 337–61.

[85] M. M. Möbius, M. Niederle, P. Niehaus, and T. S. Rosenblat, *Managing self-confidence*, *NBER Working paper* **17014** (2014).

[86] S. Chen and H. Schildberg-Hörisch, *Looking at the bright side: The motivational value of confidence*, *European Economic Review* **120** (2019) 103302.

[87] D. Eil and J. M. Rao, *The good news-bad news effect: asymmetric processing of objective information about yourself*, *American Economic Journal: Microeconomics* **3** (2011), no. 2 114–38.

[88] S. Ertac, *Does self-relevance affect information processing? experimental evidence on the response to performance and non-performance feedback*, *Journal of Economic Behavior & Organization* **80** (2011), no. 3 532–545.

[89] A. Coutts, *Good news and bad news are still news: Experimental evidence on belief updating*, *Experimental Economics* **22** (2019), no. 2 369–395.

[90] I. Esponda and D. Pouzo, *Berk–nash equilibrium: A framework for modeling agents with misspecified models*, *Econometrica* **84** (2016), no. 3 1093–1130.

[91] R. H. Berk, *Limiting behavior of posterior distributions when the model is incorrect*, *The Annals of Mathematical Statistics* (1966) 51–58.

[92] A. Alekseev, G. Charness, and U. Gneezy, *Experimental methods: When and why contextual instructions are important*, *Journal of Economic Behavior & Organization* **134** (2017) 48–59.

[93] R. P. Larrick, K. A. Burson, and J. B. Soll, *Social comparison and confidence: When thinking you're better than average predicts overconfidence (and when it does not)*, *Organizational Behavior and Human Decision Processes* **102** (2007), no. 1 76–94.

[94] D. A. Moore and D. A. Small, *Error and bias in comparative judgment: on being both better and worse than we think we are.*, *Journal of personality and social psychology* **92** (2007), no. 6 972.

[95] T. Hossain and R. Okui, *The binarized scoring rule*, *Review of Economic Studies* **80** (2013), no. 3 984–1001.