Title
Essays in Applied Behavioral Economics

Author
Zhang, Guangli

Publication Date
2021

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Essays in Applied Behavioral Economics

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

Doctor of Philosophy
in
Economics

by

Guangli Zhang

Committee in charge:

Professor Peter Kuhn, Co-Chair
Professor Emanuel Vespa, Co-Chair
Professor Gary Charness

June 2021

The Dissertation of Guangli Zhang is approved.

Professor Gary Charness

Professor Emanuel Vespa, Committee Co-Chair

Professor Peter Kuhn, Committee Co-Chair

June 2021

Essays in Applied Behavioral Economics

Acknowledgements

Throughout the writing of this dissertation, I have received a great deal of support from many amazing people. I am heavily indebted to my dissertation committee – Peter Kuhn, Emanuel Vespa and Gary Charness – for their guidance and encouragement over the years. I want to thank Peter for teaching me how to ground my research ideas and improving my projects through his comprehensive feedback. I want to thank Manu for teaching me how to design, conduct experiments and present scientific findings. I want to thank Gary for providing practical responses to my questions and always being available to talk about research ideas.

I would like to thank Erik Eyster, Ignacio Esponda, Sevgi Yuksel, Cheng-Zhong Qin, Richard Startz, Ted Frech, Youssef Benzarti, Alisa Tazhitdinova, the UCSB Applied and Experimental Research Groups for providing valuable feedback to my research projects.

I want to thank my co-authors: Jeffrey Cross and Hakan Özyılmaz, with whom I explored many exciting projects. I look forward to many more years of collaboration. In addition, I want to thank Yongwook Kim and Ganghua Mei for the stimulating conversations about research that helped me to stay productive during the COVID-19 lockdown.

Finally, I would like to thank my family, especially my wife and my long-time friend, Gladys, for supporting me throughout this long and winding journey. I could not have completed this milestone without her.

Curriculum Vitæ
Guangli Zhang

Contact

Address	North Hall 3045, Department of Economics, University of California, Santa Barbara, Santa Barbara, CA 93106
Email	guanglizhang@ucsb.edu
Webiste	https://sites.google.com/view/guanglizhang

Education

2021	Ph.D. in Economics (Expected), University of California, Santa Barbara.
2016	M.A. in Economics, University of California, Santa Barbara.
2015	M.Sc in Economics, University College London
2013	B.S. in Economics, Purdue University

Research Interests

Labor Economics; Household Finance; Public Economics; Experimental and Behavioral Economics

Working Papers

The Effect of Unemployment Benefit Pay Frequency on UI Claimants' Job Search Behaviors
(with Hakan Özyılmaz) *The Debt Payment Puzzle: An Experimental Investigation.*
(with Jeffrey Cross) *Paying for Integers.*

Work in Progress

Intra-Monthly Time-Use Patterns: Evidence from Unemployed Workers in New Jersey
(with Hakan Özyılmaz) *The Debt Payment Puzzle II: Policy Implications*
(with Jeffrey Cross). *Cutting Fewer Corners: Taxicabs and Tipping Policies*

Teaching Experience

2020	Personnel Economics, Macroeconomics, Microeconomics
2019	Personnel Economics, Macroeconomics, Financial Management
2018	Behavioral Finance, Microeconomics
2017	Microeconomics, Macroeconomics

2016	Statistics for Economics, Macroeconomics

Invited Seminars and Conference Presentations

2020	113th Annual Conference on Taxation (discussant)
2019	North America ESA

Research Grants

2019	(co-investigator) Institute for Social, Behavioral and Economic Research (ISBER). UC Santa Barbara, *"Why do Workplace Teams Bunch at Kink Points? Evidence from the Laboratory."* ($2,850)
2018	(co-investigator) Departmental Research grant, UC Santa Barbara, *"The Debt Payment Puzzle: An Experimental Investigation."* ($2,770)

Professional Affiliation

American Economic Association
Economic Science Association
National Tax Association

Referee Service

Journal of Labor Economics

Abstract

Essays in Applied Behavioral Economics

by

Guangli Zhang

This dissertation consists of three essays that investigate how the interaction of psychological and institutional factors shape individuals' daily decisions in the fields of labor economics and household finance. The essays use a combination of quasi-experimental and experimental approaches. The common theme under this research agenda is to first investigate the potential mechanisms for household behaviors that seem to be inconsistent with neoclassical theory and then use my findings to inform behaviorally motivated policy recommendations.

The first essay presents new evidence on how UI (Unemployment Insurance) benefit pay frequencies affect the job search behaviors of UI claimants in the United States. By exploiting quasi-experimental variations in states' benefit pay schedules, the essay find that switching from biweekly to weekly pay significantly increases UI claimants' unemployment durations. This observed effect can be partly rationalized by the more frequent end-of-the-month positive benefit shocks under weekly pay schedules. The essay concludes that the previously overlooked policy parameter, *benefit pay frequency*, has important effects on the job search behaviors of UI claimants.

The second essay, based on joint work with Hakan Özyılmaz, studies the sources of suboptimal allocations observed in credit card repayments using a diagnostic laboratory experiment. The essay shows that optimization ability and limited attention are jointly insufficient to explain the puzzle. Moving beyond existing results, the essay documents that the inherent negative frame of the debt payment problem interferes with subjects' ability to optimize and hinders learning. The essays shows 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, The essay finds that the debt frame increases subjects' focus on the irrelevant balance information.

The third essay, based on a joint work with Jeffrey Cross, explores customers' preferences toward leaving integer tips in the context of taxi rides. With the advent of cashless payment systems, customers are increasingly being presented convenient, low-stake tip suggestions following purchases. Despite the rising frequency of these interactions, we still know little about the preferences underlying tipping behavior. Previous research, for example, has documented that customers tend to tip integer amounts, but has not been able to disentangle if this is due to smaller cognitive costs associated with tipping an integer amount or direct utility benefits from integer tips. Combining a theoretical model with plausibly exogenous variation in the occurrence of integer tip suggestions, the essay shows that customers' behaviors are consistent with a model where they experience direct utility benefits from giving integer tips. The essay estimates that this leads to a 0.6 percentage point increase in tip rates and an approximately 2.38 million dollars transfer from riders to drivers as a result of a 2012 rate fare change that increased the probability of integer tip suggestions.

Contents

List of Figures

List of Tables

Chapter 1

The Effect of Unemployment Benefit Pay Frequency on UI Claimants' Job Search Behaviors

1.1 Introduction

The effects of UI (Unemployment Insurance) generosity on workers' unemployment durations have been extensively studied.[1] A robust finding is that higher UI generosity (measured in benefit amount and/or potential durations) lengthens unemployment duration.[2] Consequentially, policy discussions regarding the UI program mostly center on these two 'generosity' parameters. However, non-monetary policy parameters might also have important impacts on individuals' decisions. This paper examines the effect of a previously overlooked policy parameter – *benefit pay frequency* – on UI claimants' search behaviors. Using plausible

[1]See Krueger and Meyer (2002) and Schmieder and Von Wachter (2016) for a summary of past studies.

[2]This effect combines a welfare reducing moral hazard effect and a welfare enhancing liquidity effect (Chetty, 2008). The moral hazard effect occurs when increases in UI generosity reduce UI claimants' net incentives to search. Independently, the liquidity effect occurs when increases in UI generosity enable UI claimants with limited consumption-smoothing capabilities to afford to wait for better jobs.

state-year level policy variations in benefit pay frequency, I find switching from biweekly to a more frequent weekly pay schedule increases UI claimants' unemployment durations (or equivalently, decreases reemployment hazard).[3]

Why does benefit pay frequency matter for households' labor supply decisions? First, several studies of consumption responses to anticipated income note that even unconstrained households exhibit "excess-sensitivity" to anticipated income.[4] These findings suggest that a considerable fraction of many households consume hand-to-mouth. A more frequent pay schedule could potentially reduce households' tendencies to spend excessively by imposing a smoother income flow. The improved consumption smoothing capability would therefore reduce households' urges to find a job quickly. Second, Vellekoop (2018) documents households' intra-monthly budgeting cycles are driven by the end-of-the-month rent/ mortgage payments. Therefore, fluctuations in end-of-the-month cash-on-hand that are generated by variations in benefit pay frequencies could potentially affect households' consumption and labor supply decisions.[5]

Motivated by these recent findings from the household finance literature, benefit pay frequencies could affect UI claimants' search behaviors through a combination of mechanical and behavioral channels. Mechanically, a more frequent weekly benefit pay schedule would increase the occurrences of positive (monthly) liquidity shocks during unemployment. Zhang (2017) who evaluated income on a monthly basis, finds that households with biweekly pay schedule can receive three paychecks (instead of two) once every six months, whereas households with weekly pay schedules can receive five paychecks (instead of four) once every three

[3]Throughout the paper, I define the reemployment hazard (h_t) as the likelihood of finding a job at the end of period t, conditional on entering period t unemployed. In addition, one can roughly interpret a 10% increase in reemployment hazard equivalent to a 10% decrease in expected unemployment duration.

[4]Under the Life-Cycle/Permanent Income Hypothesis (Friedman et al., 1957; Modigliani & Brumberg, 1954), the frequency or timing of benefit payments should not matter as forward-looking rational agents' expenditures do not respond to shapes or paths of anticipated inflow of income. See Browning and Lusardi (1996), Browning and Crossley (2001a) and Jappelli and Pistaferri (2010) for a summary of past studies.

[5]Note, both mechanisms are operating through the *liquidity effect* as in Chetty (2008).

months. The quasi-experimental variations in the timing and the magnitude of the extra benefit can generate different liquidity shocks to UI claimants being paid under different payment frequencies. Second, in order for these extra benefit to have impacts, households should exhibit excess sensitivities to these anticipated income shocks. Common explanations for this phenomenon include the limited ability to smooth consumption due to liquidity constraints (Browning & Crossley, 2001b), quasi-hyperbolic discounting (Ganong & Noel, 2019; Gerard & Naritomi, 2019), illiquid savings (Kaplan, Violante, & Weidner, 2014), household's tendency to hold lifetime wealth in cash (Olafsson & Pagel, 2018) and reliance on rules-of-thumb or heuristics (Zhang, 2017). In this paper, I investigate the combined effect from the mechanical and behavioral channels of the pay frequency effect on households.

The UI program in the United States provides an ideal environment to examine the pay frequency effect, as benefit pay frequency varies across states and over time. However, the impact of pay frequency has received little attention partly due to the small monetary differences between weekly and biweekly pay.[6] This paper makes two contributions to the UI literature. First, it quantifies the pay frequency effect in the context of labor supply under unemployment insurance. Second it documents UI claimants' search responses to anticipated fluctuations in the monthly benefit amount, holding the benefit amount constant. To my best knowledge, neither benefit pay frequency nor the timing of extra benefits have been extensively explored in the context of labor supply under social insurance policies.[7] Therefore, I propose two plausible dimensions of heterogeneity for evaluating social benefits.[8]

This paper uses data from the 1985-2007 Survey of Income and Program Participation

[6]Apart from Fishman, Farrell, Gardiner, Barnow, and Trutko (2003), who collected information on continued UI certification frequency for 8 states and discussed its impact on UI takeup in 2003.

[7]Note, the consumption responses to payment and expenditure timing has been well studied. For example, Castner, Henke, et al. (2011) finds food stamp recipients spend a disproportionally large fraction of their SNAP benefits at the start of their benefit month – this is known as the "SNAP cycle". In a recent study, Beatty, Bitler, Cheng, and Van der Werf (2019) finds the SNAP cycle is more pronounced for workers who are paid on a weekly or monthly basis.

[8]In a follow up study, I plan to apply the same type of analysis to Workers Compensation.

(SIPP) to provide new evidence of the effects of benefit pay frequencies. SIPP is well suited to this project because it contains information from different states from 1985-2007. The panel structure allows me to follow individuals over the course of 2.5 to 4 years and observe the transitions in and out of unemployment, and the weekly employment variable allows me to obtain the precise lengths of unemployment. The detailed measures of individuals' assets prior to unemployment in SIPP's supplemental surveys allow me to control for UI claimants' monetary constraints at the time of unemployment.

The paper is divided into three parts. I start by estimating the effects of different benefit pay frequencies on UI claimants' reemployment hazard. Next, I investigate a potential mechanism by examining the effect from receiving anticipated extra benefit checks on UI claimants' reemployment hazard under different pay frequencies. Lastly, I examine several policy implications relating to the frequency of benefit payments.

Section 1.2 uses plausible quasi-experimental changes in UI benefit pay frequency to examine its impact on UI claimants' unemployment durations (or reemployment hazards). In my benchmark analysis, I estimate Cox proportional hazard models with state fixed effects and year fixed effects. When I restrict my analysis to New York (1993), Washington (1996) and Massachusetts (2003), I find switching from biweekly to weekly pay frequency increases expected duration by 2-4 weeks (decreases reemployment hazard by 22%). To assess the robustness of the benchmark estimates, I adopt a event study framework to estimate the dynamic impact of pay frequency change on job finding hazard.[9] Interestingly, results from the event study design suggests the main effect seems to only have short run impacts on UI claimants' job finding hazard.

In the presence of households who are hand-to-mouth, I examines a possible channel that can rationalize the pay frequency effect — the frequency and the magnitude of extra benefit checks in Section 1.3. Variations in the end-of-the-month cash on hand can have

[9] States with less than 100 observations from 1985-2007 are excluded from the main analysis.

important impacts on UI claimants' monthly cash flows as most major expenditures – such as rent, credit card debt, utility bills, mortgage – occur around the end of each month. In this section, I adopt the quasi-experimental variation introduced in Zhang (2017) by estimating UI claimants' responses to the anticipated extra benefit checks under weekly or biweekly pay frequencies.[10] Overall, I find (possibly) receiving an anticipated extra benefit at the end of the month leads to a 34.0% (or 16.3%) decrease in the next month's reemployment hazards for claimants under weekly (or biweekly) pay. The finding seems to suggest that the effect of extra benefit check on unemployment durations exhibit diminishing marginal returns.[11] Given that UI claimants under weekly pay can experience *twice* as many end-of-the-month benefit shocks as those under biweekly pay, the frequency of extra benefit checks plays a more important role in affecting UI claimants' job search behaviors.

To design a cost-effective social benefit program, policy makers needs to know the potential costs and benefits from switching to a more frequent benefit pay schedule. Therefore I investigate implications from varying pay frequencies on various policy outcomes. First, using data from the Annual Survey of Government Employment and Payroll (ASGEP), I estimate the impact of switching to weekly pay on state governments' annual UI administrative costs. I find close to zero and statistically insignificant effect. Second, using samples from 1985-2007 SIPP, I find switching to weekly pay does not affect UI eligible workers' UI take-up rate. Third, I find suggestive evidence that switching to weekly pay could potentially reduce the liquidity effect (gains from consumption smoothing) from increases in UI benefits.

This paper is closely related to the literature on estimating the consumption smoothing benefit of unemployment insurance. Several papers have estimated the effect from varying the pre-unemployment asset level on unemployment durations and found a considerable liq-

[10] Note, (1) in terms of the *magnitude*: the extra benefit amount is equivalent to a 25% (or 50%) increase in the regular monthly benefit level for weekly (or biweekly) pay states; (2) in terms of the *occurrence frequency*: the extra benefit month occurs 4 times (or 2 times) under weekly (or biweekly) pay.

[11] For example, receiving two separate $500 extra checks lead to larger responses in UI claimants' durations than receiving a one-time $1000 extra check.

uidity effect (Card, Chetty, & Weber, 2007; Chetty, 2008; LaLumia, 2013). All these papers find UI claimants with larger pre-unemployment assets tend to search longer as they are more capable to smooth their consumptions during unemployment. However, studying the effect from varying cash on hand *during* unemployment spell is equally important. A recent strand of literature has examined the optimal path of benefits (Ganong & Noel, 2019; Gerard & Naritomi, 2019; Kolsrud, Landais, Nilsson, & Spinnewijn, 2018; Lindner & Reizer, 2020; Schmieder & von Wachter, 2017). In particular, most of these studies have relied on temporal variations in benefit extensions to estimate the fiscal cost and consumption smoothing benefits from a step-wised benefit path.[12] On the other hand, this is the first paper that examines the impact of variations in (i) benefit pay frequency and (ii) monthly benefit amount under a constant (weekly) benefit path. Findings from this paper highlight the importance of incorporating the frequency and the timing of income and expenditure streams when evaluating the consumption smoothing benefit of UI.

In addition to the empirical literature on unemployment insurance, this paper also relates to the recent household finance literature that examines households' consumption and borrowing responses to the anticipated timing and frequency of income (Aguila, Kapteyn, & Perez-Arce, 2017; Baugh & Correia, 2018; Berniell, 2018; Olafsson & Pagel, 2018; Zhang, 2017) or the timing of consumption commitments Vellekoop (2018). In particular, from two closely related works, Leary and Wang (2016) and Baugh, Leary, and Wang (2018) find households experience more financial shortfalls when the timing of income and expenditure streams are misaligned. The new evidence documented in this paper indicates that households' imperfect budgeting responses to the anticipated liquidity streams can have significant spillover effects on their labor supply decisions, at least in the context of unemployment insurance.

The rest of the paper is organized as follows. Section 1.2 studies the impact of pay fre-

[12] A noticeable exception is Lindner and Reizer (2020), who find that front-loading UI benefit payments leads to shorter unemployment duration and increases in reemployment wage.

quency variations on UI claimants' reemployment hazards. Section 1.3 explores one potential mechanism. Section 1.4 discusses related policy implications of switching to weekly benefit pay frequency. Section 1.5 includes a series of robustness checks for the pay frequency effect estimation and section 2.5 concludes.

1.2 Empirical Evidence: Pay Frequency and Reemployment Hazard

This section briefly discusses the UI operations and pay frequency procedures used in the United States when this paper was written. Section 1.2.2 explains the empirical strategy, Section 1.2.3 describes the dataset and Section 1.2.4 presents the first empirical finding.

1.2.1 UI Benefit Pay Frequency in the United States

The Unemployment Insurance (UI) benefit program, part of the Federal Social Security Act of 1935, is designed to provide periodic economic support for individuals who are laid off involuntarily (Price, 1985). In most states, the program ensures a weekly benefit amount (WBA) for up to 26 weeks determined by the claimant's earnings from the most recent four calendar quarters, i.e. the base period. Eligible individuals file an initial claim in the state they reside in during the first week of unemployment, and may then wait three weeks or longer before the claim is processed. Maximum WBAs, otherwise known as coverage generosity, continued claim certifications, and payment requirements and frequencies vary by state. The map in Figure 1.1 shows the weekly certification benefits in 2007 for each state.

To be eligible for receiving continued benefits, claimants are required to file certification periodically after the initial claim. The certification process asks claimants to report their earnings, job offers, and job search activities for the past benefit week(s) they are claiming benefits for. Many states impose a minimum amount of weekly job search requirement and only a small fraction of the claimants' search activities were audited.[13] The payday is usually 2-4 days after the certification day depends on the state and most UI benefits were distributed via mailed checks. Most states require weekly certification; some states allow claimants to choose to file either weekly or biweekly and others require claimants to file biweekly.[14]

Prior to the 1980s, many states use a biweekly payment frequency due to the limited capability in filing and payment technology (Blaustein, 1979). Back then, in-person claim and mail claim were the two predominant ways to file a continued UI certification. Starting from the mid 1990s, the introduction of the more advanced telephone and online filing systems as well as the direct deposit payment system induced several states to opted-in for a more frequent (weekly) filing process for continued UI claims. As of today, nine states still require claimants to file for continued certification and receive payments on a biweekly basis.

The Department of Labor (DOL) does not explicitly record the UI pay frequency in the *Significant Provisions of State Unemployment Insurance Laws*, I relied on two complementary information sources to recover the state level benefit pay frequency policies. The primary source comes from the Benefit Accuracy Measurement (BAM) program, administrated by the DOL. For the purpose of auditing, BAM samples around 10 UI claimants every week for each state. Important to this study, BAM contains information on continued benefit filing frequency starting from 1985. For each state, I used the year that the share of alternative claiming method accelerated by the greatest amount and designated it the event year.[15] To complement with the BAM survey, I manually collected the pay frequency information from archived documents on state government websites. Using *Google.com* and *Archive.org*, I was

[13]Only 10 randomly selected UI claimants are chosen to be audited by the Benefit Accuracy Measurement program every week for each state.

[14]Almost all states define a benefit week as a calendar week from Sunday through Saturday; New York is the only state that defines a benefit week as Monday through Sunday.

[15]Apart from Nevada and Ohio, all switcher states changed their pay frequency from Biweekly to Weekly pay.

able to verify all UI pay frequency information starting from the mid-1990s.[16] I find several states introduced weekly benefit payment: District of Columbia (2008), Maryland (2014), Massachusetts (2003), Minnesota (2007), Montana (2015), New Hampshire (2003), New Jersey (2014), New Mexico (1999), New York (1993), Oregon (1992), Rhode Island (1996), Utah (1994), Virginia (1998), Washington (1996) and Wyoming (2018). Most of these states fully switched to a weekly filing system after 1-5 years. I use Massachusetts, New York and Washington for my main analysis.[17]

1.2.2 Empirical Strategy

Most states changed their benefit pay frequency to encourage the use of more advanced filing technology by claimants, i.e. mail to telephone, or telephone to internet.[18] Given that the changes in the continued filing technologies are implemented without prior notice, it is unlikely that UI claimants would exhibit anticipatory responses to such policy variations.

To study the effects of switching from biweekly to weekly pay frequencies on UI claimants' unemployment exits, my empirical strategy exploits the state level variations in UI payment frequencies.[19] In particular, I estimate a series of Cox proportional hazard models with the

[16] I have attached the detailed UI pay frequency records along with their document sources in the appendix, see Table A1.1

[17] The rest of the switcher states are excluded from my analyses due to limited observations – concerns with attenuation bias.

[18] For example, the original communication sent by the New Jersey Department of Labor stated: "The New Jersey Department of Labor and Workforce Development is encouraging all of our unemployment insurance customers to claim their benefits each week by using our Internet application at www.NJUIFILE.net. That's correct! Instead of claiming your benefits every two weeks, you may now claim them each week and receive a benefit payment each week...."

[19] As documented in Anderson and Meyer (1997), UI take up decision is endogenous and can be largely affected by factors such as benefit generosity or potential compensated duration. Controlling for benefit generosity and possible UI duration extensions, I find the elasticity of take up with respect to pay frequency change is insignificant and close to zero, which suggest the endogeneity in take-up decision is unlikely to be affected by the policy change in benefit pay frequency. Similarly, Ebenstein and Stange (2010) finds no impact of UI filing technology on UI takeup.

following specification:

$$\log h_{ist} = \alpha_t + \beta_1 \mathbf{1}\{weeklypay_s\} + \mathbf{X}_{ist} \tag{1.1}$$

where h_{ist} is the hazard rate of exiting unemployment for individual i from state s at unemployment week t. α_t is the flexible non-parametric baseline hazard rate at the given week t conditional on surviving. $weeklypay_s$ is a dummy indicates the pay frequency for state s at a given year. Specifically, $weeklypay_s = 1$ if pay frequency is on a weekly basis, and $weeklypay_s = 0$ if pay frequency is on a biweekly basis. $\mathbf{X_{ist}}$ is a set of controls: (1) state level controls that include start-of-the-spell monthly unemployment rate and UI generosity; (2) Industry, occupation fixed effect and (3) individual specific controls such as 10-piece log-linear spline for the claimant's pre-unemployment wage, total wealth, age, education, marital status and being on the seam between interviews to adjust for the seam effect. Lastly, $\mathbf{X_{ist}}$ also includes (4) year fixed effects that capture changes over time that vary uniformly across states and (5) state fixed effects that capture time invariant cross state differences. Standard errors are clustered at the state level.

Because UI benefit is not well measured under the SIPP survey, I use three alternative proxies for claimants' benefits: (i) individual predicted benefit, (ii) state-year level simulated replacement rate and (iii) state-year level maximum benefit. The first proxy – *predicted benefit* – follows the two-step approach from Chetty (2008). In the first step, I predict claimants' pre-unemployment log annual wages using their observable characteristics (as included in $\mathbf{X_{ist}}$ in Eq.(1.1)). In the second step, I plug the predicted wages into a UI calculator to obtain claimants' predicted UI benefits.[20] The construction of the second proxy – *simulated replacement rate* – follows the standard two-step procedure (East & Kuka, 2015; Gruber, 1997; Kroft & Notowidigdo, 2016). The idea here is to use to policy change in state-year level UI

[20] I used UI calculator program from Kuka (2020).

generosity to proxy for average claimants' UI benefit. In step one, I predict claimants' pre-unemployment log annual wages using observable characteristics. In step two, I use a fixed 1993 national sample to compute the average weekly benefits and UI replacement rate for all state-year combination in the data set. The two-stage simulated replacement rate only depends on observable demographic characteristics and variations from state laws. Lastly, given the fact that approximately 50% of UI claimants receive the maximum benefit Chetty (2008), I also use the *maximum weekly benefit* to proxy for individual claimant's UI benefit.

1.2.3 Data and Sample

I use unemployment spell data from the Survey of Income and Program Participation (SIPP) from 1985-2007. SIPP is a panel data that contains weekly employment status so I can follow an unemployed worker over time. I closely follow Chetty (2008) and Kroft and Notowidigdo (2016) when constructing my sample for this part of the analysis: I restrict my sample to be prime-age males who (a) report searching for a job, (b) are not on temporary layoff, (c) have at least 3 months of work history in the survey (to compute pre-unemployment earnings), (d) took up UI benefits within the first month of unemployment. Furthermore, to reduce the influence of outliers and restrict my attention to search behavior in the first year after job loss, I censor unemployment duration at 50 weeks. Lastly, All monetary values are adjusted into 1990 dollars using CPI-U.

Apart from the aforementioned sample construction criteria, I make additional restrictions on individuals' wealth measures. For UI claimants in the SIPP data, wealth measures are collected through the topical module - "asset and liquidity" - which only happens 2 to 3 times in a panel. Therefore, about one-half unemployment spells does not contain wealth measures prior to the unemployment. One approach is to use ex-post (post-/ during-unemployment) wealth measures to proxy for ex-ante wealth. However, UI claimant's ex-post wealth level is endogenous to factors such as unemployment duration (Gruber, 2001) and thus is a noisy indicator to an individual's ability to consumption smooth. Since the pay frequency effect could potentially affect UI claimant's search effort through the liquidity channel, I restrict my sample to those with information on pre-unemployment total (liquid and non-liquid) wealth holding. The aforementioned restrictions leave me 3,406 unemployment spells in the pooled sample.

1.2.4 Empirical Result

I begin by providing graphical evidence on the effect of change in pay frequency on duration for UI recipients in the treated states. Then I use regression analyses to complement the graphic analysis. In this part, I split my sample into two subsamples according to the frequency of receiving UI benefits: weekly or biweekly. In particular, prior to the policy change, the switcher state is included in the biweekly group; after the policy change, the switcher state is moved to the weekly group.

Figure 1.2 is the Kaplan-Meier survival curves for UI claimants under weekly or biweekly benefit pay frequencies. The survival curve for claimants under biweekly pay frequency is slightly lower, indicating a less frequency pay schedule is associated with shorter unemployment duration. Partly due to the limited number of samples for the switcher states, the difference is statistically insignificant under a non-parametric Wilcoxon test for equality with p=0.1598.

The graphic analysis provide some preliminary evidence that suggests the impact of UI benefits on search duration could be affected by the frequency of benefit payment given a similar benefit replacement rate. However, the result from this simple comparison could potentially be driven by individual or state specific characteristics. To complement the graphic analysis, I run a set of estimations using semi-parametric Cox proportional hazard model

(Eq.(1.1)) that includes a rich set of controls. Findings from the regression analysis are consistent with the graphic analysis.

The main results are presented in Table 1.2. The reported estimates are hazard coefficients. In column 1, I estimated effect of pay frequency change without controlling for state or individual observable characteristics. In column 2, I include the full sets of controls and restrict my sample to those with pre-unemployment total wealth only. In Column 3, I replace pre-unemployment total wealth with pre-unemployment net wealth. This further reduces the number of samples. Column 3 is my preferred specification as it represents the estimation of Eq. (1.1) using the most stringent set of controls. The key coefficient of interest is the $WeeklyPay$ dummy that varies over time. Under all columns, the estimated hazard coefficient β_1 is negative and significant. In particular, $\beta_1 = -0.255$ (SE 0.063) indicates that switching from biweekly to weekly pay frequency leads to a decrease in the likelihood of exiting unemployment spell by 22% for an average UI claimant.[21]

Result 1: *For Massachusetts, New York and Washington, switching from biweekly to weekly pay frequency leads to a 10% to 22% decrease in the reemployment hazard. This roughly translates to 2 to 4 weeks of additional unemployment for average UI claimants with mean spell length equal to 18 weeks.*

This result provides a first evidence on the important role of pay frequency in designing UI policy. However, there still exists concerns with both internal and external validity of the estimation. First, states have implemented other concurrent reforms (such as filing technology change) could bias the true causal effect. Second, the treated and control states might be on different outcome trends prior to the treatment. Third, the baseline estimation relies on policy variations from three states only, the results might not be generalized to states with very different demographic or socio-economic characteristics. In response to these concerns,

[21] The percentage change in hazard rate (caused by the change from biweekly to weekly pay frequency) is computed using $exp(\beta_1) - 1$.

I further assess the validity of the baseline two-way fixed effect research design in Section 1.5. Overall, results from additional analyses are suggests the benchmark estimation is robust.

1.3 Mechanism: Pay Frequency and Monthly Benefit Shocks

This section proposes a potential mechanism that could rationalize the pay frequency effect – the frequency of the end-of-the-month extra benefit checks. Section 1.3.1 introduces the institutional settings of extra benefit checks under the UI system. Sections 1.3.2 and 1.3.3 presents the empirical strategy, data and sample. Section 1.3.4 presents empirical results.

1.3.1 Institutional Background - Extra Benefit Checks

Standard UI benefit schedule in the US is evaluated under a weekly basis. Typical benefit schedules consist of a predetermined weekly benefit amount (WBA) not exceeding a potential benefit duration (PBD), based on UI claimants' pre-unemployment earnings during the "base period". As noted by Zhang (2017), under either weekly or biweekly pay schedules, in months with five calendar weeks, UI claimants receive an extra payment check. The amount of extra benefit check is equivalent to 25% (or 50%) of monthly benefits under weekly (or biweekly) pay schedules. Under a six-month unemployment duration, claimants receive two (or one) extra benefit months under weekly (or biweekly) pay. Given that many major expenditures – rent, mortgage payment, utility bills – occurs on a monthly basis, receiving this extra benefit check towards the end of the calendar month could have a significant impact on a UI claimant's liquidity in the following month.[22]

Figures 1.3 and 1.4 show examples of the monthly benefit paths under weekly pay and biweekly pay schedules. I assume a UI claimant is entitled to receive the first UI benefit at the beginning of May 2020; the constant weekly benefit amount of $400 paid on Tuesdays

[22] These periodic expenditures are sometimes referred as "consumption commitments".

terminates at the end of November 2020 (26 weeks). Under the weekly pay schedule, a UI claimant can experience up to two extra benefit shocks, whereas claimants under the biweekly pay schedule can only experience up to one extra benefit shock. Therefore, an average UI claimant under weekly pay schedule are more likely to have extra cash-on-hand at the end of each month during their unemployment spell.

Since there exist variations in both the magnitude and the frequency of extra benefit checks, there should be differential responses by UI claimants to these positive end-of-the-month liquidity shocks among UI claimants under weekly or biweekly pay schedules. In particular, if the effect on unemployment duration exhibit diminishing marginal returns to extra benefit checks, we would expect to observe larger responses under weekly pay schedules. That is, holding the total extra benefit amount constant, the overall duration increases from receiving multiple smaller shocks is expected to be higher than the overall duration increases from receiving a single large shock. On the other hand, since such benefit shocks can be anticipated by forward-looking UI claimants, the impact might be small and insignificant as rational agents should have already internalized this anticipated volatility in their monthly benefit paths.

1.3.2 Empirical Strategy

Next, I analyze the effect of receiving an anticipated end-of-the-month extra benefit check on UI claimants' reemployment hazards in the subsequent month. I exploit quasi-experimental variations in the monthly benefit levels – the variation mainly depends on the claimant's timing of unemployment.

An ideal experiment to study this effect is to compare individuals' reemployment hazards between those who have or have not received the extra benefit checks. However, this comparison requires information on the exact timing of benefit distribution. Due to the SIPP's data limitations, I use the differential probabilities of receiving an extra benefit check in calendar year-months to proxy for UI claimants' the treatment status. In particular, there are months in calendar years where it is never possible to receive extra benefits – as illustrated in Table A1.2, these months vary from year-to-year. Under this setup, a UI claimant is *not treated* in the t^{th} month of unemployment if the probability of receiving an extra benefit check $= 0$. Similarly, a UI claimant is (possibly) *treated* in the t^{th} month of unemployment if the probability of receiving an extra benefit check > 0.[23]

To examine whether extra benefit checks affects UI claimants' search behavior in the subsequent calendar month, I estimate a series of Cox proportional hazard models with the following specification:

$$\log h_{ist} = \alpha_t + \beta_1 \mathbf{1}\{PosExtra_{is,t-1}\} + \mathbf{X}_{ist} \tag{1.2}$$

where h_{ist} is the hazard rate of exiting unemployment for individual i from state s at time t. α_t is the flexible non-parametric baseline hazard rate at the given week t conditional on surviving. $\mathbf{1}\{PosExtra_{is,t-1}\}$ is a dummy indicates the status of receiving extra benefit checks from the previous month. Specifically, $\mathbf{1}\{PosExtra_{is,t-1}\} = 1$ if the probably of having received the extra check is > 0, and $\mathbf{1}\{PosExtra_{is,t-1}\} = 0$ if the probably of having received the extra check $= 0$. $\mathbf{X_{ist}}$ is a set of controls: (1) state level controls that include start-of-the-spell monthly unemployment rate and UI generosity; (2) Industry, occupation fixed effect and (3) individual specific controls such as 10-piece log-linear spline for the claimant's pre-unemployment wage, total wealth, age, education, marital status and being on the seam between interviews to adjust for the seam effect. Lastly, $\mathbf{X_{ist}}$ also includes (4) year fixed effects that capture changes over time that vary uniformly across states, (5) state fixed effects

[23]Given a normal processing and filing time, most UI checks are likely to be distributed towards the end of each week. I use the possibility of receiving extra benefit checks on Wednesday, Thursday and Friday as an alliterative categorization. The results are qualitatively equivalent.

that capture time invariant cross state differences. Standard errors are clustered at the state level and (6) calendar month fixed effects that capture the seasonal patterns of reemployment hazard.

In addition to Eq (1.2), I also interacts $\mathbf{1}\{PosExtra_{t-1}\}$ with $\mathbf{1}\{WeeklyPay\}$ in a separate regression. Eq (1.3) allows me to examine the differential extra benefit effects for states before and after they switched from biweekly to weekly pay frequency. If the extra benefit attributes to the observed pay frequency effect, we expect to see a positive and significant estimates for the interaction term.

$$\log h_{ist} = \alpha_t + \beta_1 \mathbf{1}\{PosExtra_{is,t-1}\} + \beta_2 \mathbf{1}\{WeeklyPay_s\} \\ + \beta_3 \mathbf{1}\{WeeklyPay_s\} \times \mathbf{1}\{PosExtra_{is,t-1}\} + \mathbf{X}_{ist} \quad (1.3)$$

1.3.3 Data and Sample

The sample is identical to Section 1.2. The only difference is that I use monthly (instead of weekly) unemployment status because the variations occur at a monthly basis.

1.3.4 Empirical Result

The estimated results are presented in Table 1.3. The reported estimates are hazard coefficient. For columns (1)–(3), the key coefficient of interest β_1 is negative under both the pooled sample and the two sub-samples. In particular, possibly receiving an extra check in the previous month is estimated to reduce this claimant's reemployment hazard by 23%. Under the cross sectional comparison (columns (2) and (3)), I find the point estimate is more than two times larger for weekly pay states. In column (4), when I interact pay frequency policy change dummy with the extra benefit dummy, I find that the previously documented pay frequency effect mainly operates through the end-of-the-month extra benefit channel: the effect of receiving extra benefits is significantly stronger after states switched from biweekly to weekly pay!

I note that the magnitude of the extra benefit check is equivalent to a 25% (or 50%) increase in monthly benefit amount under the weekly (or biweekly) pay schedule. UI claimants under weekly pay schedules are twice as likely to experience a positive benefit shock during unemployment. This finding suggests that the frequency of benefit shocks plays a more important role in improving UI claimants' capability to smooth consumption during unemployment.

Result 2: *Possibly receiving an extra benefit check could reduces UI claimants' reemployment hazards for the subsequent month. The estimated effect is larger under weekly pay states.*

The documented larger responses to the anticipated end-of-the-month positive benefit shocks under weekly pay suggest the consumption smoothing gains from receiving extra benefit checks exhibit diminishing returns. That is, holding the total extra benefit amount constant, the effect of receiving multiple smaller benefit on UI claimants' reemployment hazards is estimated to be larger than the effect of receiving a single large shock.

Given that the extra check months occur more frequently under the weekly pay schedule, the results provide some support for the existence of the pay frequency effect. In particular, the weekly pay schedule mechanically leads to more occurrences of extra benefit checks in UI claimants' monthly benefit paths. Relative to the biweekly pay schedule, the greater likelihood of having income shocks aligned with end-of-the-month major expenditures under the weekly pay schedule could significantly increase their cash on hand for the subsequent months. Therefore, holding the weekly benefit and pre-unemployment wealth constant, UI claimants under the weekly pay schedule are more able to smooth consumption during unemployment. As a result, switching from biweekly to weekly pay leads to longer unemployment durations, because UI claimants can afford to wait longer after a switch in pay frequency.

Table 1.4 presents related results from heterogeneity analysis. In particular, I separately estimated the hazard coefficient for different sub-samples. Overall, I find the extra benefit effect is mainly driven by UI claimants who are not liquidity constrained. In particular, the point estimate is large and statistically significant for UI claimants: (1) with above median pre-unemployment total wealth; (2) married with working spouses and (3) who are home-owners with or without mortgage payments. The result could be interpreted in several ways. First, there might exist some threshold level of extra benefit amount for claimants to become responsive to it; Second, unconstrained households might have stronger consumption commitments and response more intensively to extra benefit shocks; Third, UI claimants who were not liquidity constrained prior to unemployment might be less capable to smooth consumption during unemployment and exhibit higher sensitivity to extra benefit shocks. Due to data limitation, I am not able to tease out these explanations.

Result 3: *Responses to extra benefit shocks seems to be driven by liquidity unconstrained UI claimants.*

One potential concern with the liquidity-effect based explanation is that UI claimants who anticipates an extra benefit check might *intentionally* delay search effort and capture this “additional" benefit. That is, the response to extra benefit checks might be a result of the incentive-distorting moral hazard effect. However, I argue that moral hazard might be less of a concern. Consider the following example: suppose that a UI claimant receives benefit check every week, and the potential job also makes salary payments every week. Whether this month has five or four paychecks makes absolutely no different to this UI claimant's job finding incentives, because this claimant would get a fifth check whether she find a job or not. This suggests that the potential distortion from moral hazard effect would be quite small.

1.4 Policy Implications

Findings from the previous sections have shown that switching from biweekly to weekly pay leads to longer unemployment duration as it decreases UI claimants' job finding hazards. In addition, I find evidence suggests that such effect is likely a result of improved consumption smoothing capabilities for UI claimants. In this section, I investigate two policy relevant implications from switching benefit pay frequencies. Specifically, Section 1.4.1 estimates the impact of varying benefit pay frequency on UI administrative costs and UI Take-up. Section 1.4.2 studies the interactions between benefit pay frequency and increases in benefit amount (WBA). Overall, I find variations in benefit pay frequency: (1) does not have significant effect on states' UI administrative cost or take-up, (2) and potentially crowds-out the consumption smoothing gains from increases in WBA.

1.4.1 Impact on UI Administrative Costs and Take-up

Switching to a more frequent certification frequency might occur additional administrative processing cost for state governments, as the weekly volume of benefit certification are likely to be doubled. On the other hand, such variation could also affect the certification cost for UI claimants. To examine these potential policy impacts, I follow Ebenstein and Stange (2010) and estimate a series of two-way FE regression:

$$y_{s,t} = \alpha_0 + \beta_1 \mathbf{1}\{WeeklyPay_{s,t}\} + \beta_2 \mathbf{1}\{PostPhone_{s,t}\} + \beta_2 \mathbf{1}\{PostNet_{s,t}\} + \mathbf{X}_{st} \quad (1.4)$$

where $y_{s,t}$ is the outcome of interest: {log of Full-Time Employment, log of Payroll, fraction of Employment Part-Time, UI Take-up rate} for state s and calendar year t. $WeeklyPay$ is a dummy that varies with state and calendar year t. $PostPhone$ and $PostNet$ are dummies that indicates whether this state implemented phone or internet claiming for continued certifications.[24] In my sample period, all switcher states changed filing frequency at the same year they adopted phone claiming. Lastly, $X_{i,t}$ includes state FEs, year FEs, max. WBA and state unemployment rate.

Data used in this section are drawn from two separate sources. The UI administrative employment and payroll information is obtained from the Annual Survey of Government Employment and Payroll (ASGEP) dataset from 1992-2007. The data contains state level annual expenditure and employment information under the "Social Insurance Administration" item. The UI take-up rate is obtained from the 1985-2007 SIPP data. I first use unemployed individuals' pre-unemployment annual wages to predict their UI eligibility. Then I compute the the UI take-up rate using the number of UI takers divided by the number of UI eligible individuals for each given state year.[25]

Table 1.5 presents results for UI administration costs and UI take-up rates. Columns (1), (2) and (3) are estimated using ASGEP data. The point estimate for β_1 are all negative and close to zero indicating switching to weekly benefits does not seem to have significant impacts on switcher states' employment costs (even after controlling for filing technology). Further, as illustrated from column (3), the insignificant response in employment cost are not driven by increases in part-time employment. Column (4) uses the 1985-2007 SIPP UI eligible unemployed sample to examine the effect on UI take-up rates. The point estimate for β_1 is small and close to zero.

Result 4: *Switching to weekly pay does not seem to affect state's UI administrative cost. In addition, the potential change in the continued certification costs due to pay frequency variation does not seem to affect UI take-up rate.*

[24] The policy event time for technology adoptions are obtained from the BAM survey.
[25] I use UI calculator from Kuka (2020) to estimate unemployed workers' UI eligibility.

1.4.2 Interaction with Variations in WBA

Does switch to week pay *crowd-out* the consumption smoothing gains from increases in UI benefits? This section investigates the interaction between benefit pay frequency and increases in UI benefit amount. In particular, I separately estimates UI claimants' responses to UI benefit increase under different benefit pay frequencies.

The empirical strategy follows earlier literature (Chetty, 2008; Krueger & Meyer, 2002; Meyer, 1990) that exploits state and year variation in the maximum weekly benefit amount. The treatment group is UI claimants with higher earnings that are likely to be affected by the increase in the Max WBA. The control group is UI claimants with lower earnings that are not going to be affected by the change in the Max WBA. The identification assumption requires the two groups to follow parallel trends over time in absence of the Max WBA changes within sub-samples.

The baseline Cox proportional hazard model closely follows Chetty (2008) and Kroft and Notowidigdo (2016):

$$\log h_{it} = \alpha_t + \beta_1 \log b_i + \beta_2 (t \times \log b_i) + \mathbf{X_{it}} \tag{1.5}$$

where h_{it} is the hazard rate of exiting unemployment for individual i at time t. α_t is the flexible non-parametric baseline hazard rate at the given week t conditional on surviving. b_i is the weekly benefit amount that this individual receives. The coefficient β_1 is the elasticity of the hazard rate with respect to UI benefits at $t = 0$. The inclusion of $(t \times \log b_i)$ allows the effect of benefit varying with duration. $\mathbf{X_{it}}$ controls for state and year fixed effect for the purpose of the difference-in-difference design. In addition, $\mathbf{X_{it}}$ also includes for occupation and industry dummies; 10 piece log wage spline for claimant's pre-unemployment wage; log total wealth and other individual specific linear controls (education, age, marital status and being on the seam week between interviews).

I use the identical sample as in Section 1.2. Table 1.6 provides a descriptive summary for the subsamples divided into weekly and biweekly pay frequencies. Although UI claimants from biweekly states have longer unemployment duration, receive higher WBA on average and have higher pre-unemployment annual wage, the State-Year level UI generosity measured by simulated replacement rate or predicted benefit amount appears to be similar across the two subsamples.

The main results are reported in Table 1.7. I report estimations of duration elasticity from the pooled sample, as well as two sub-samples. The variable of interest is duration elasticity – UI claimants' likelihood of find a job at the first week of unemployment in response to increases in UI benefit. I control for state fixed effects and year fixed effects, industry and occupation fixed effects, 10-piece linear spline of pre-unemployment annual wage earnings, pre-unemployment total wealth and other individual specific demographics.

The results columns (2) to (3) suggests that UI claimants' behavioral responses are slightly stronger under biweekly states compared to the point estimate under the pooled sample. For the subsamples, a 10% increase in benefit is estimated to reduce reemployment hazard by 4% (or 3%) under a Biweekly (or Weekly) pay frequency. Although potentially due to smaller sample size, the point estimate for the Weekly pay sub-sample is statistically insignificant.

Result 5: *Given an identical % increase in UI benefit, UI claimants under the Biweekly pay frequency seems to be more responsive to it, though the difference in estimated hazard elasticity are not statistically significant.*

According to the traditional view (Meyer, 1990; Moffitt, 1985), the difference in the estimated duration elasticities are driven purely by the differences in moral hazard, i.e. the degree of incentive distortion is larger under weekly pay frequencies. Chetty (2008), however, might interpret the observed heterogeneous behavioral responses to changes in UI benefits as differences in the liquidity effect, i.e. UI claimants under weekly pay are more able to smooth consumption during unemployment. Given that the present discounted benefit levels are almost identical under the two pay frequencies, moral hazard is unlikely to be the main driver of this observed difference in duration elasticities.[26] Therefore, I conclude that there exists some degree of *substitutability* between the weekly pay schedule and increases in UI benefits. In particular, given that weekly payments potentially improves UI claimants capability to consumption-smoothing, the additional liquidity gains from increases in WBA are less helpful for UI claimants under the weekly pay schemes.

1.5 Robustness

In this section, I include a series of robustness checks for the estimation of pay frequency effect. First, to eliminate the potential contamination effects from the variations in UI filing technologies, I use a two-way fixed effect framework to examine the impact of filing technology on UI take-up and UI claimants' reemployment hazards independent from changes in pay frequency. Second, I use a event-study framework to visually examine the validity of the parallel-trends assumption for the two-way fixed effect design. Third, I ran a series of permutation tests to compare the baseline estimation to 1,000 randomly generated benchmarks.[27] The main result survives under all these tests.

1.5.1 Variations in the Continued Filing Method

One concern with the baseline empirical strategy presented in section 1.2 is that the effect can be confounded by other concurrent policy changes. For all three switcher states in

[26]Formally, since the government cannot observe agents' search effort (e), moral hazard occurs when agents only consider the private marginal product of work – wage minus benefit (w-b) – and their private costs when choosing search effort. Note that the private marginal product of work is lower than the social marginal product of work (w). Thus, increases in benefit (b) distorts search incentive.

[27]In addition to these checks, I adopt a Synthetic Control Method (SCM) to improve the comparability between switcher and non-switcher states prior to the policy change. The SCM results are discussed in detail in Appendix A.2.

my sample, the pay frequency variation is accompanied by the adoption of telephone filing technology for continuing claims. In this subsection, I restrict my sample to states that only varied continued filing methods to estimate its impact on UI take up and UI claimants' reemployment hazards. The results suggest that the estimated pay frequency effect is likely not driven by variations in continued filing method.

I use the UI claims filing method data collected by the US Department of Labor from 1985 to 2007. The claim filing data is originally collected from the BAM (Benefit Accuracy Measurement) program. The BAM samples around surveys 400 UI claimants in each state-year level. The data contains claiming method information for both initial and continued claims. For the purpose of this project, I restrict my attention to variations in continued claim methods.[28] To date policy changes, I follow Ebenstein and Stange (2010) to look for the sharp changes in claim method usage. In particular, for each state, I infer the time of filing technology change to be the year that the share of claims filed via telephone or internet increased by the greatest amount.

The empirical approach exploit the state-year level variations in UI continued filing methods. I examine the pay frequency effect after the inclusion of technology adoption dummies: $\mathbf{1}\{PostPhone_s\}$ and $\mathbf{1}\{PostNet_s\}$. I use Cox hazard models to estimate the pay frequency effect after accounting for technology adoption changes (Eq. (1.6)). The Cox hazard estimation uses the identical sample as in Section 1.2. Note that the left-hand side variable $\log h_{ist}$ represents the reemployment hazard for individual i from state s at t^{th} week of unemployment.

$$\log h_{ist} = \alpha_t + \beta_1 \mathbf{1}\{WeeklyPays\} + \beta_2 \mathbf{1}\{PostPhone_s\} + \beta_3 \mathbf{1}\{PostNet_s\} + \mathbf{X}_{ist} \quad (1.6)$$

[28] Ebenstein and Stange (2010) use the same BAM dataset to examine the impact of initial claim methods on UI take-up and find no effect.

Results are presented in Table 1.8. Due to small sample size, the standard errors are quite large. In both cases, I cannot reject the effect of technology adoption is different from zero for states that did not change pay frequencies. This result complements Ashenfelter, Ashmore, and Deschênes (2005), who finds increases in continued filing costs – increases in monitoring of job searches – has no effect on the duration of continued claims.

Result 6: *The pay frequency effect is not likely to be driven by technology adoption.*

1.5.2 A event study framework

To examine the validity of parallel assumption, I estimate event-study models with leading and lagging treatment dummies, so we can assess the pre-treatment time trends in the hazard coefficient in the following specification:

$$\log h_{ist} = \alpha_t + \sum_{k=-4, k\neq -1}^{4} \delta_k D_s^k + \mathbf{X}_{ist} \tag{1.7}$$

where α_t and $\mathbf{X}_{ist}$ are defined as they were in Eq. (1.1). D_s^k is a dummy variable that equals to 1 after state s changed from biweekly to weekly pay. The endpoints are set to address imbalances in the sample. The endpoints are binned so that $D_s^4 = 1\{event\ time \geq 4\}$ and $D_s^{-4} = 1\{event\ time \leq -4\}$.

Figure 1.6 presents the event study plot on the dynamic effects of weekly pay frequency on UI claimants' job finding hazards. I fail to reject the null of having pre-treatment trends. Interestingly, the treatment effect seems to be concentrated in the first two years after the pay frequency switch, suggesting the baseline two-way fixed effect estimates is main driven by short term responses. Overall, the 95% confidence intervals are quite large, potentially due to the small sample size in the SIPP data.

Result 7: *There seems to be no presence of pre-treatment trends. The pay frequency effect seems* *to be driven by short term responses.*

1.5.3 Permutation Test

In this section, I implement a non-parametric permutation tests for the purpose of randomization inference: Comparing to a large number of possible random assignments, is my baseline result significantly different from them? How different is it?[29]

I randomly assign the 3 switcher states in the sample with the event years following the actual pay frequency policy implantation timetable, i.e. 1 state in 1993, 1 state in 1996 and 1 state in 2003.[30] Following random treatment assignments, I re-estimate the placebo pay frequency effect following the baseline specification (Table 1.2, Column (3)). Then I repeat this process for 1,000 times to obtain a distribution of estimated coefficients. The p-value in this context is defined as the probability that the baseline estimate is obtained purely by chance and is computed by the following expression:

$$p\text{-}value = \frac{\sum_{i=1}^{1000} \mathbb{1}|\beta^i_{baseline} \geq \beta_{placebo}|}{1000}$$

Figure A3.1 plots the empirical distribution of the placebo estimates using 1,000 random treatment assignments. The dashed line is the point estimate from the baseline estimation ($\beta = -0.255$). Comparing to the estimated placebo treatment effects, the actual effect is not significant (p-value = 0.121). This is potentially driven by the small sample size or the unbalanced dataset. The result from permutation test provides a conservative p-value for the baseline estimation.

[29]I also implement a similar permutation test for the estimated extra benefit effect, see Appendix A.3 for more detail.

[30]Note that the assignment only applies to states that started with biweekly pay frequency at the beginning of the study. See Table A1.1 for the list of qualified states.

1.6 Conclusion

There is a large literature that studies the impact from receiving unemployment insurance on job search behavior (Krueger & Meyer, 2002; Schmieder & Von Wachter, 2016). Many of these paper evaluated the effects of benefit generosity. However, there has not been many research conducted on the non-monetary aspects of UI policy – in the case of benefit pay frequency and the timing of benefits, there was none.[31]

This paper uses data from the 1985-2007 SIPP to investigate the effect of benefit pay frequency on job search behavior by presenting three pieces of empirical evidence. First, utilizing quasi-experimental changes in benefit pay frequency, this paper finds switching to a more frequent weekly pay schedule reduces UI claimants' job finding hazard. Second, using variations in the timing of the extra benefit checks, the paper finds suggestive evidence that the pay frequency effect is partly due to the more frequent occurrences of the end-of-the-month extra benefit checks under the weekly pay schedule. Third, switching from biweekly to weekly pay does not seem to increase states' UI administrative costs, nor UI eligible workers' UI take-up rate.

Furthermore, I investigate the interactions between benefit pay frequency and benefit amount increases using a standard Difference-in-Difference research design (Chetty, 2008; Kroft & Notowidigdo, 2016; Krueger & Meyer, 2002; Meyer, 1990). I separately estimate the effects of benefit increase on unemployment durations for states under weekly or biweekly pay schedules. Results from the additional analyses imply that the magnitude and the significance of the liquidity effect due to increases in UI benefit amount may vary with benefit pay frequency. Overall, findings from this paper highlight the importance of benefit pay frequency and pay timing when evaluating the consumption smoothing benefit from social insurance policies.

[31] See O'Leary (2004) for the recent a summary about the effects of changing continued certification requirements.

There are several limitations to this paper that future research could address. First, although this paper finds suggestive evidence on the linkage between pay frequency and the liquidity effect, the paper does not directly test this due to data limitations. Future research could make use of the high frequency transaction data to investigate potential differences in consumption, saving and borrowing behavior across weekly and biweekly pay frequencies. Such a study would provide more concrete evidence on the causal relationship between pay frequency and consumption smoothing. For example, Baugh and Correia (2018) use account aggregator data to investigate the borrowing pattern for employed workers across different pay frequencies. Ganong and Noel (2019) uses JPMCI data to study the consumption patterns for UI claimants.

Second, the idea of evaluating the impact of pay frequency and pay timing variation on consumption and other household behaviors can be easily applied to evaluating different social benefit programs. Future studies could expand this research agenda and investigate related questions. For example, (1) Do we observe a similar pay frequency effect in other settings? (2) How big is the welfare gain if the timing and the frequency of social benefits align with individuals' expenditure streams?

Figures and Tables

Figure 1.1: UI Benefit Filing and Pay Frequency Policies

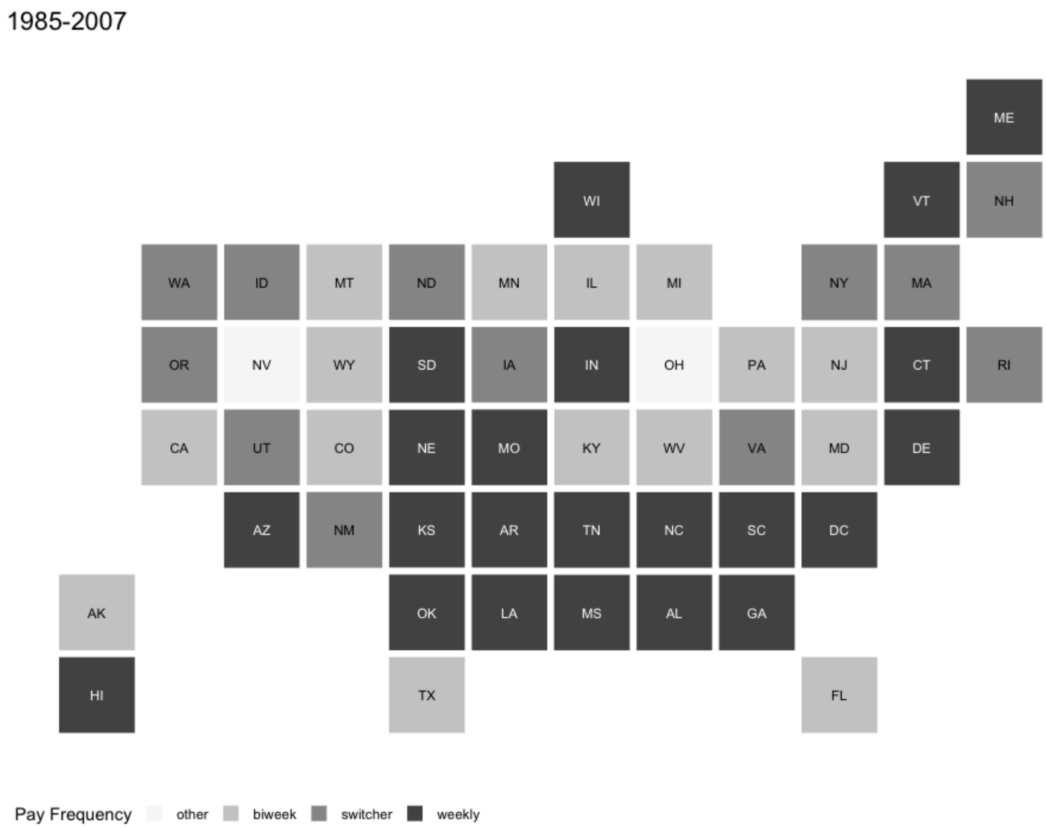

Notes: The US map shows the UI benefit pay frequency by state at the time of 2007: States that pays biweekly are in light gray; states that switched pay frequency from biweekly to weekly pay are in darker gray; states that pays weekly are in black. Nevada switched pay frequency two times; Ohio allowed for either weekly or biweekly filing. The pay frequency policy information are collected from a combination of the archived state websites via *archive.org* and survey results from the Benefit Accuracy Measurement Audit.

Figure 1.2: Survival Curves - Comparing biweekly and weekly Pay Frequency

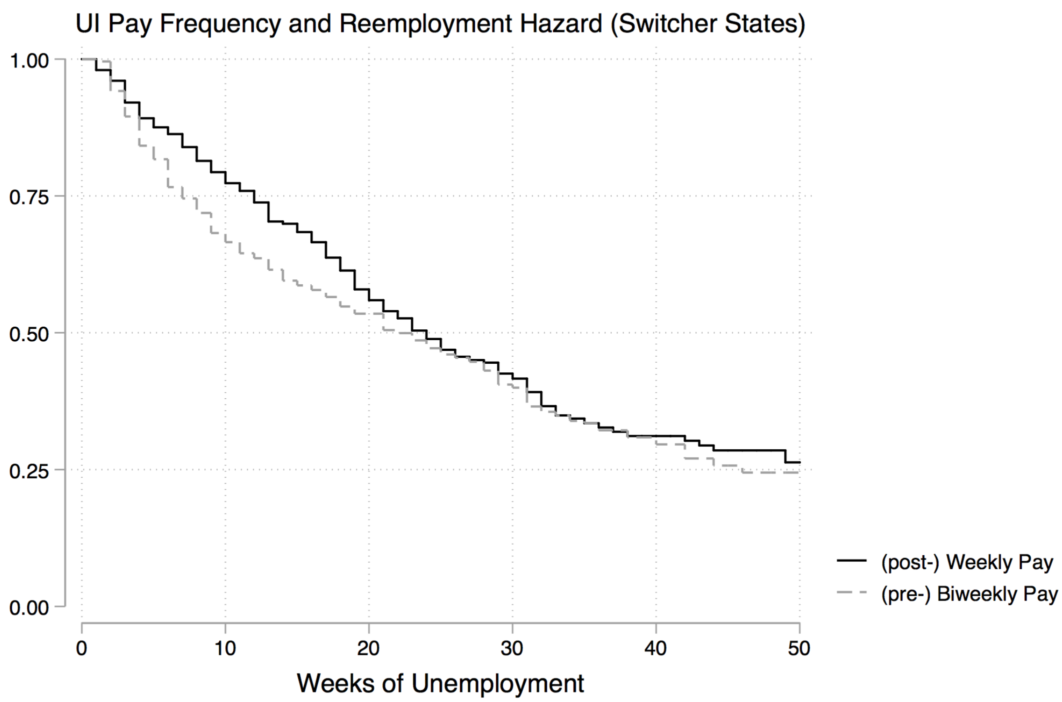

Notes: Figure shows individual level unemployment duration from SIPP 1985-2007 for Massachusetts, New York and Washington (the switcher states). The vertical axis indicates the fraction of unemployed sample. The dashed line represents the probability of exiting unemployment for UI claimants from the switcher states prior to the change in pay frequency (distribute benefit payment on a weekly basis); the solid line represents the probability of exiting unemployment for UI claimants from the switcher states post the change (distribute benefit payment on a biweekly basis). Following Chetty (2008), these two Kaplan-Meier survival curves adjusts for the seam effect. The unemployment duration is censored at 50 weeks.

Figure 1.3: Weekly Pay - Extra Benefit Shocks under Monthly Benefit Path

Extra Benefits under Weekly Pay

June, 2020

1	**2**	3	4	5	6	7
8	**9**	10	11	12	13	14
15	**16**	17	18	19	20	21
22	**23**	24	25	26	27	28
29	**30**					

Sept, 2020

	1	2	3	4	5	6
7	**8**	9	10	11	12	13
14	**15**	16	17	18	19	20
21	**22**	23	24	25	26	27
28	**29**	30				

Monthly Benefit Schedule (May to Nov)

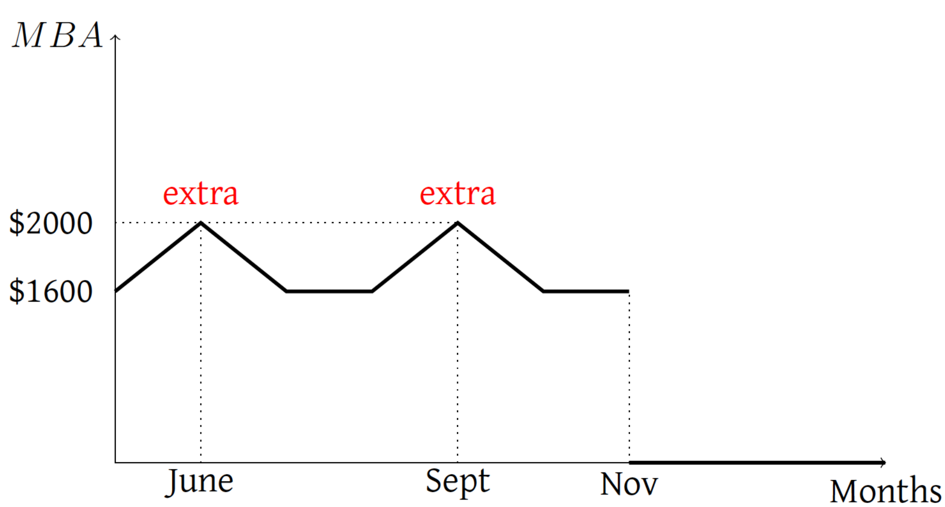

Figure 1.4: Biweekly Pay - Extra Benefit Shocks under Monthly Benefit Path

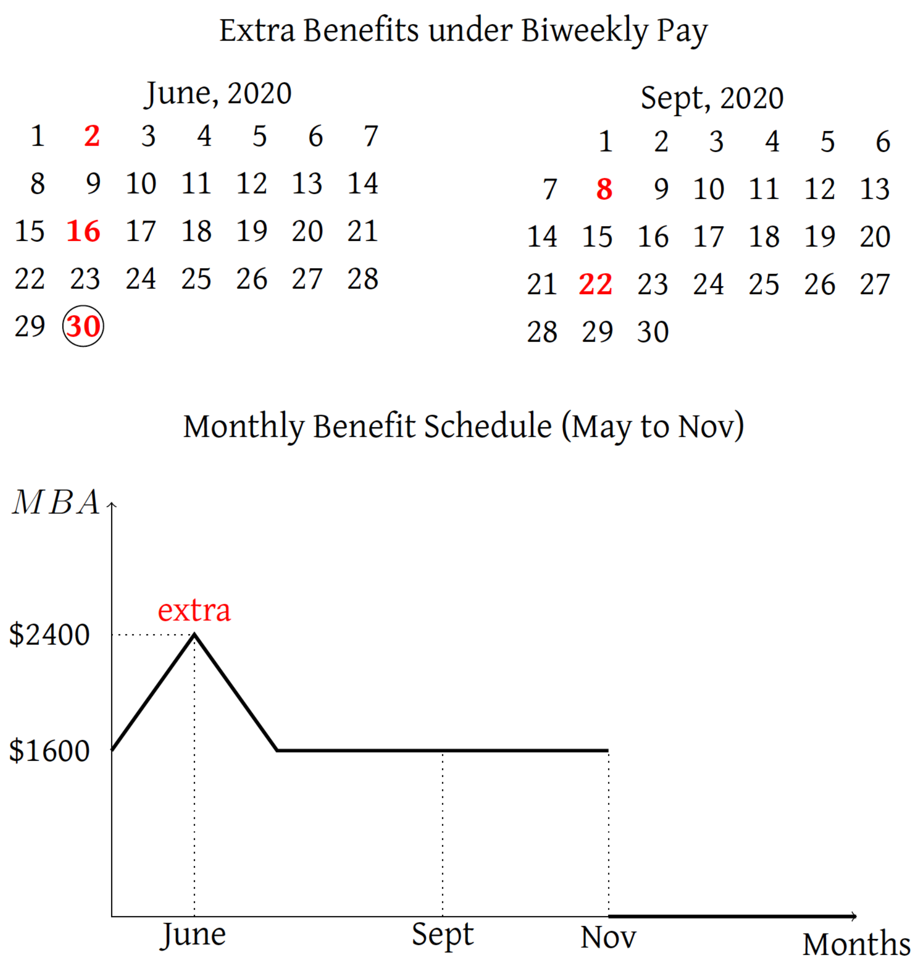

Figure 1.5: Permutation test for inference of baseline estimation: pay frequency effect

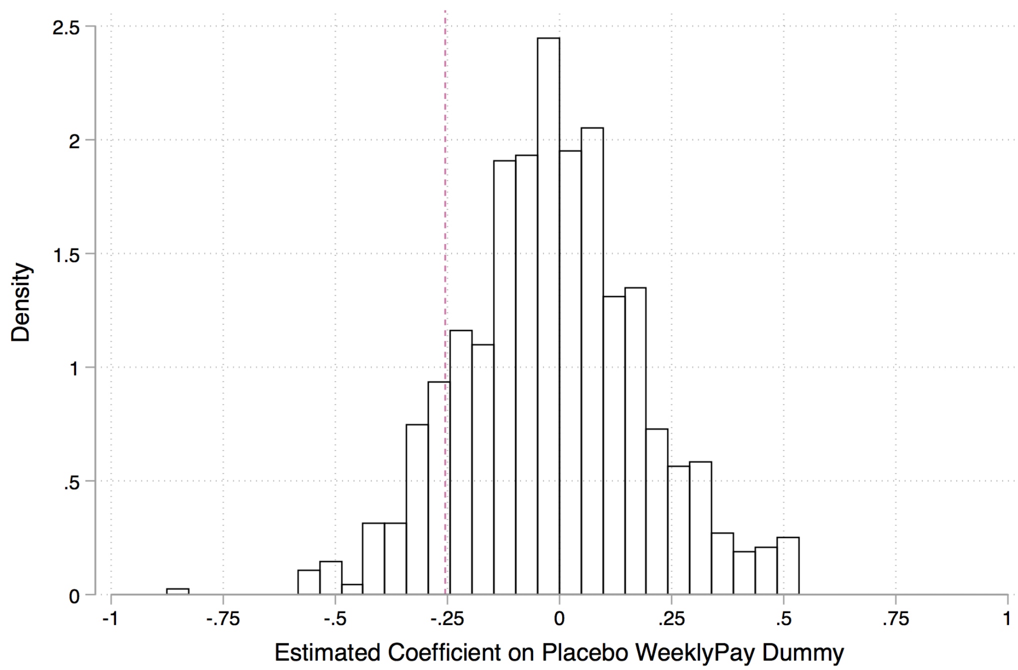

Notes: Figures shows the empirical distribution of estimated placebo treatment effects from 1,000 random assignments. Dashed line is the actual treatment effect estimated from Table 1.2 Column (3). p-value under the permutation test is 0.121

Figure 1.6: Dynamic effects of weekly pay on UI claimants' job finding hazards

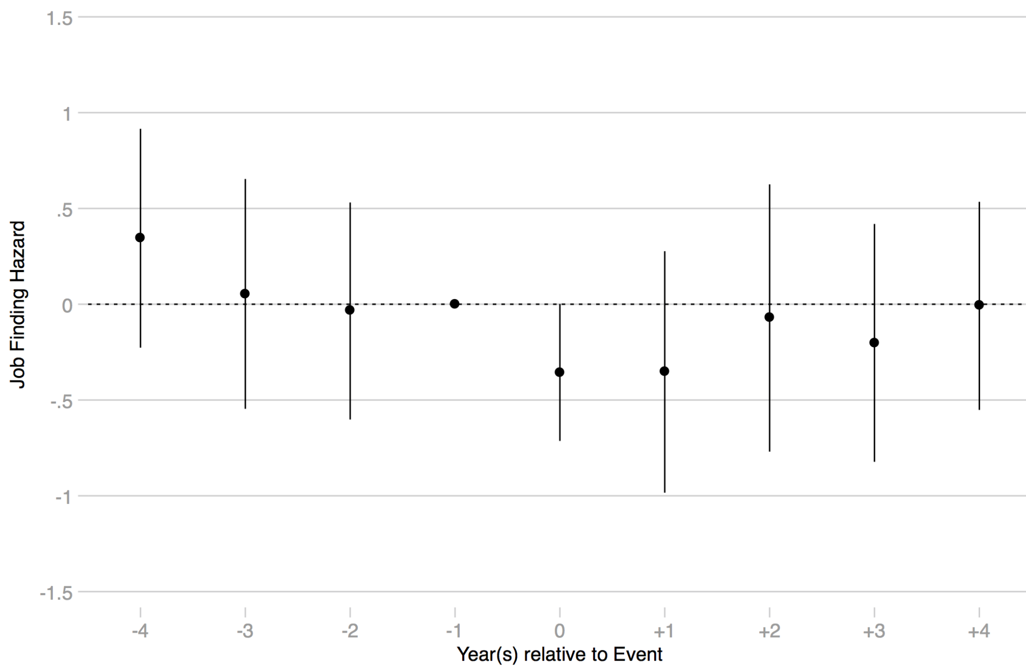

Table 1.1: Descriptive Statistics for Switcher and Control States, SIPP 1985-2007

	Pooled		Switcher		Non-Switcher	
Variable	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Unemployment Duration (weeks)	18.13	13.70	18.69	14.27	18.03	13.60
Average UI weekly benefit ($)	170.59	29.67	194.21	25.27	166.46	28.43
Maximum UI weekly benefit ($)	235.67	54.51	273.37	37.14	229.07	54.38
Simulated replacement rate	0.51	0.05	0.50	0.02	0.51	0.05
Age	38.75	11.33	39.28	11.58	38.66	11.28
Years of Education	12.32	2.84	12.73	2.78	12.25	2.85
Married	0.63	0.48	0.58	0.49	0.63	0.48
Pre-ue annual wage ($)	21183.49	16122.33	22468.60	15967.98	20958.74	16141.39
Pre-ue liquid wealth ($)	32749.87	96605.55	40208.85	111824.70	31194.76	94078.79
Pre-ue unsecured debt ($)	4889.97	18534.49	6026.81	31496.22	4691.15	15170.78
Pre-ue home equity ($)	35268.72	57448.00	50209.10	74291.17	32655.82	53554.67
# Spells	3,646		507		2,919	

Notes: The data presented are individual level unemployment spells from 1985-2007 SIPP data. The average and maximum UI weekly benefit amount are obtained from the US Department of Labor. All dollar values are converted into 1990 values. The sample is restricted to male UI claimants only. The pay frequency policy information are collected from archived state websites via *archive.org* and the BAM survey.

Table 1.2: Impact of switching from biweekly to weekly pay frequency on UI claimants' reemployment hazard

	(1)	(2)	(3)
$\mathbb{1}\{\text{WeeklyPay}\}$ (β_1)	-0.105	-0.116	-0.255
	(0.048)	(0.075)	(0.063)
log(WBA)	-	-0.621	-0.493
		(0.104)	(0.150)
State FE, year FE	×	×	×
Industry FE, occupation FE and seam dummy	×	×	×
Education, marriage and age		×	×
10-piece pre-ue annual wage spline		×	×
State log unemployment rate (at layoff time)		×	×
Pre-unemployment log total wealth		×	
Pre-unemployment log net wealth			×
# Spells	3,383	3,176	1,904

Notes: All columns report result from semi-parametric Cox proportional hazard model from estimating equation (1.1). The key coefficient (β_1) is the change in hazard rate with respect to pay frequency policy changes. Data are individual-level unemployment spells from 1985-2007 SIPP. I include state fixed effects, year fixed effects, industry and occupation fixed effects, a 10-piece linear spline of the pre-unemployment annual wage, pre-/ post-unemployment total wealth, onseam indicator and other individual specific controls such as education and marital status. Standard errors clustered by state are in parentheses.

Table 1.3: Impact of receiving extra benefit checks on UI claimants' reemployment hazard

	(1) Pooled	(2) Weekly Pay	(3) Biweekly Pay	(4) Interaction
$\mathbb{1}\{\text{PosExtra}_{t-1}\}$ (β_1)	-0.277	-0.428	-0.179	-0.178
	(0.095)	(0.152)	(0.144)	(0.119)
$\mathbb{1}\{\text{WeeklyPay}\}$ (β_2)	-	-	-	-0.112
				(0.079)
$\mathbb{1}\{\text{WeeklyPay}\}$ x $\mathbb{1}\{\text{PosExtra}_{t-1}\}$ (β_3)	-	-	-	-0.182
				(0.094)
log(WBA)	x	x	x	x
State FE, year FE, month FE	x	x	x	x
Industry FE, occupation FE and seam dummy	x	x	x	x
Education, marriage and Age	x	x	x	x
10 piece pre-ue annual wage spline	x	x	x	x
pre-ue net wealth	x	x	x	x
State log unemployment rate (at layoff time)	x	x	x	x
# Spells	1,800	650	1,135	1,680

Notes: All columns report result from semi-parametric Cox proportional hazard model from estimating Eq. (1.2) and (1.3). Data are individual-level unemployment spells from 1985-2007 SIPP. I include state fixed effects, year fixed effects, calendar month fixed effects, industry and occupation fixed effects, a 10-piece linear spline of the pre-unemployment annual wage, pre-unemployment net wealth, onseam indicator and other individual specific controls such as education and marital status. Standard errors clustered by state are in parentheses. I restrict my sample to those who stay unemployed for at least 1 month.

Table 1.4: Extra benefit checks and reemployment hazard, heterogeneity analysis

	(1) Below Median pre-ue Wealth	(2) Above Median pre-ue Wealth	(3) Spouse Working	(4) Spouse Not Working	(5) Single	(6) Renter	(7) Homeowner w/ Mortgage	(8) Homeowner w/o Mortgage
$\mathbb{1}\{\text{PosExtra}_{t-1}\}(\beta_1)$	-0.055	-0.313	-0.345	-0.044	-0.011	-0.111	-0.225	-0.235
	(0.132)	(0.112)	(0.134)	(0.109)	(0.106)	(0.136)	(0.109)	(0.219)
log(WBA)	x	x	x	x	x	x	x	x
State FE, year FE, month FE	x	x	x	x	x	x	x	x
Education, marriage and Age	x	x	x	x	x	x	x	x
Industry FE, occupation FE and seam dummy	x	x	x	x	x	x	x	x
10 piece pre-ue annual wage spline	x	x	x	x	x	x	x	x
pre-ue total wealth	x	x	x	x	x	x	x	x
State log unemployment rate (at layoff time)	x	x	x	x	x	x	x	x
# Spells	1,601	1,823	1,357	2,067	1,283	1,172	1,766	486

Notes: All columns report result from semi-parametric Cox proportional hazard model from estimating variants of Eq. (1.2). Data are individual-level unemployment spells from 1985-2007 SIPP. Standard errors clustered by state are in parentheses. I restrict my sample to those who stay unemployed for at least 1 month.

Table 1.5: Benefit pay frequency, UI administrative cost and UI take-up

	(1) Log(FTE)	(2) log(Payroll)	(3) Frac. PTE	(4) UI Take-Up
$\mathbb{1}$\{WeeklyPay\} (β_1)	-0.063	-0.039	0.004	-0.004
	(0.078)	(0.085)	(0.015)	(0.023)
$\mathbb{1}$\{PostPhone\} (β_2)	-0.027	-0.044	-0.012	-0.010
	(0.059)	(0.058)	(0.013)	(0.015)
$\mathbb{1}$\{PostNet\} (β_3)	0.019	-0.002	-0.007	-0.004
	(0.072)	(0.075)	(0.013)	(0.021)
State FE, year FE	x	x	x	x
State unemployment rate	x	x	x	x
State Max WBA	x	x	x	x
Pre-ue wage, wealth, education, age, marriage	-	-	-	x
Industry FE, Occupation FE	-	-	-	x
Data source		ASGEP 92-07		SIPP 85-07, UI eligible
Observation level		(state x year) 610		(individual) 15,580

Notes: All columns report result from linear regression model from estimating Eq. (1.4). Columns (1)-(3) are estimated using state x year observations from 1992-2007 ASGEP dataset. Column (4) is estimated using individual-level unemployment spells from 1985-2007 SIPP. I excluded NV and OH due to their non-standard certification frequencies. Standard errors clustered by state are in parentheses.

Table 1.6: Descriptive Statistics for UI recipients, SIPP 1985-2007

	Weekly Pay		Biweekly Pay	
Variable	Mean	Std. Dev.	Mean	Std. Dev.
Duration	17.64	13.56	18.16	13.80
Average UI benefit amount ($)	163.36	32.11	174.74	26.40
Maximum UI benefit amount ($)	219.00	59.05	245.91	48.04
Predicted UI benefit ($)	183.34	75.61	186.60	83.78
Simulated replacement rate	0.52	0.04	0.51	0.06
Age	38.59	11.39	38.76	11.31
Years of Education	12.36	2.64	12.28	2.96
Married	0.60	0.49	0.63	0.48
Pre-ue annual wage ($)	20226.46	14017.69	21823.23	17111.16
Pre-ue liquid wealth ($)	32216.64	97894.54	31791.09	88986.34
Pre-ue unsecured debt ($)	5171.25	20627.17	4817.604	16594.31
Pre-ue home equity ($)	31129.15	55169.38	37149.59	57904.9
# Spells	1,334		2,279	

Notes: The data presented are individual level unemployment spells from 1985-2007 SIPP data. The average and maximum UI weekly benefit amount are obtained from the US Department of Labor. All dollar values are converted into 1990 values. The sample is restricted to male UI claimants only. The pay frequency policy information are collected from archived state websites via *archive.org* and the BAM survey.

Table 1.7: Duration Elasticity, by Pay Frequency

	(1) Pooled	(2) Weekly Pay	(3) Biweekly Pay
log(WBA)	-0.399	-0.365	-0.509
	(0.180)	(0.323)	(0.025)
State FE, year FE	x	x	x
Industry FE, occupation FE and seam dummy	x	x	x
Education, Marriage, Age	x	x	x
10 piece pre-ue annual wage spline	x	x	x
Pre-ue net wealth	x	x	x
State unemployment rate (at layoff time)	x	x	x
# Spells	1,904	742	1,294

Notes: All columns report semi-parametric Cox proportional hazard model results from estimating equation (1.5). The reported coefficients are elasticities of hazard rate with respect to UI benefits. Data are individual-level unemployment spells from 1985-2007 SIPP. Standard errors clustered by state are in parentheses.

Table 1.8: Impact of continued filing technology on UI claimants' reemployment hazard

	(1) Baseline	(2) w/ Tech.
$\mathbb{1}\{WeeklyPay\}$ (β_1)	-0.255	-0.232
	(0.063)	(0.082)
$\mathbb{1}\{PostPhone\}$ (β_2)	-	-0.056
		(0.161)
$\mathbb{1}\{PostNet\}$ (β_3)	-	0.070
		(0.126)
log(WBA)	×	×
State FE, year FE	×	×
Education, marriage and age	×	×
Industry FE, occupation FE and seam dummy	×	×
10-piece pre-ue annual wage spline	×	×
pre-ue net wealth	×	×
State log unemployment rate (at layoff time)	×	×
# Spells	1,904	1,904

Notes: Columns (1) presents the baseline estimation of the pay frequency effect. Column (2) report result from the Cox proportional hazard model (1.6). Data are from 1985-2007 SIPP. All switcher states in my sample changed pay frequency at the same year as they adopted telephone filing. Standard errors clustered by state are in parentheses.

Chapter 2

The Debt Payment Puzzle: An Experimental Investigation

with Hakan Özyılmaz

2.1 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 (Bhutta, Fuster, and Hizmo (2020), Andersen, Campbell, Nielsen, and Ramadorai (in press), Ponce, Seira, and Zamarripa (2017)). 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, Ponce et al. (2017) and Gathergood, Mahoney, Stewart, and Weber (2019), 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

[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) Gathergood et al. (2019) 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, Gathergood et al. (2019) 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] Ponce et al. (2017) 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. Gathergood et al. (2019) 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.

potential explanations for this puzzling behavior. Two immediate explanations are financial literacy and limited attention. Researchers in household finance have long emphasized the role of financial literacy (Lusardi and Mitchell (2014), Lusardi and Tufano (2015)). 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 (Campbell (2016), Beshears, Choi, Laibson, and Madrian (2018)). The behavioral economics literature has emphasized the role of limited attention in consumer choice (Chetty, Looney, and Kroft (2009), Stango and Zinman (2014), Karlan, McConnell, Mullainathan, and Zinman (2016), Bordalo, Gennaioli, and Shleifer (2017)). 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.

[4]Handel and Schwartzstein (2018) 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.

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 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 Ponce et al. (2017) and Gathergood et al. (2019). 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 Gathergood et al. (2019).

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

[6] Ponce et al. (2017) find the share of optimal allocations to be approximately 15% among people who hold two comparable credit cards using observational data. Gathergood et al. (2019) find this rate to be 11.8% .

around salience is that it is an established cognitive mechanism that guides choice behavior in various contexts (Nisbett and Ross (1980), Taylor and Thompson (1982)). Its applications in behavioral economics have been particularly fruitful in capturing deviations from rational choice in simple environments (Bordalo, Gennaioli, and Shleifer (2013), Kőszegi and Szeidl (2012)).

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 (Soroka, Fournier, and Nir (2019), Baumeister, Bratslavsky, Finkenauer, and Vohs (2001), Kahneman (1979)). 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

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

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.

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 (Ponce et al. (2017), Gathergood et al. (2019)). 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 (Sunstein (2011)). 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 (Bertrand and Morse (2011), Seira, Elizondo, and Laguna-Müggenburg (2017)). 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 (Hastings, Madrian, and Skimmyhorn (2013), Lusardi and Mitchell (2014)). 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

[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 Campbell (2016) for a discussion of these trade-offs.

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 problem. 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 (Gabaix (2014)), models of salience (Bordalo et al. (2013)), focusing (Kőszegi and Szeidl (2012)) and selective attention (Karlsson, Loewenstein, and Seppi (2009)). 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 (Benartzi and Thaler (2007), Gathergood et al. (2019)), we present systematic evidence on how an algebraically identical allocation problem under different frames induces different distributions of heuristic use over subjects.

2.2 Evidence for Suboptimal Repayments

The purpose of the baseline experiment is two-fold. First, it helps us documenting the severity and persistence of suboptimal credit card repayments in a controlled environment where many institutional factors that confound consumers' incentives to fully pay off the more expensive credit card are removed. Second, it documents that the suboptimal repayments extend beyond the standard explanations for consumer mistakes such as limited attention and optimization ability.

2.2.1 Baseline Design

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, we *vividly* 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.[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. The vivid display of balance information in our design mirrors the vivid display of balance information on actual credit card statements and online accounts, and is an aspect of the decision environment we manipulate in further treatments.[10]

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

[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.

[10] We will revisit the vivid display of balance information and its potential role in driving the puzzle in Section 2.3 of the paper.

stage consists of five decision periods.[11] 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.

Some Merits of 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 (Ellis & Freeman, 2020; Read, Loewenstein, Rabin, Keren, & Laibson, 1999). 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

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

component of credit score calculations.[12] 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 (Keys & Wang, 2018; Stewart, 2009). 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.[13]

Parameter Choices and Balance Reallocation

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 distinguish potential balance-

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

[13]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.

[14]Gathergood et al. (2019) document that the observed annual interest rate difference is 15% at the 90th percentile corresponding to a monthly interest rate difference of 1.25%. Ponce et al. (2017) find the average monthly interest rate gap to be 1.1% in their data.

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 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.

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 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.

[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.

[16] See Figure B5.7 for a screenshot of these periods.

[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 B.7.

Subjects are advanced to the next stage if they complete a stage or if they exceed the maximum allotted time.[19]

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 (Fischbacher (2007)). 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.

[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.

[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 $x+y=300, \quad x,y\geq 0$

2.2.2 Baseline Results

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 2.3 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.

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 2.5, only 22.4% of the 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. Ponce et al. (2017) find the share of optimal allocations to be approximately 15% among people who hold two comparable credit cards. Gathergood et al. (2019) find this rate to be 11.8%.

Attending to the interest rate information is a necessary and sufficient step for optimal decision-making from the perspective of rational and a large class of attention-based behavioral decision models. Since we track the information buttons that a subject clicks, we have a concrete measure that we can use as a proxy for the attentiveness of our subjects to their interest rate information at the time of decision making. We find that 93.2% of our subjects'

$first$ period decisions are made after actively clicking a button to acquire interest rate information.[22] [23] While the optimality rate of payments that are made $without$ paying attention to interest rates is 5.5% (statistically equivalent to 0%, $p = 0.30$), this rate significantly increases to 21.55% (statistically unequal to 100%, $p < 0.0001$) among payments that are made after acquiring interest rate information ($p = 0.024$). These results suggest that while paying attention to interest rates is indeed necessary for our subjects to optimize and a significant predictor of the optimality rate, the puzzle remains present to a large extent even among the "attentively made" payments.

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 repayments is misallocated.[24] Ponce et al. (2017) report that consumers misallocate 50% of their repayments to the low interest rate card and Gathergood et al. (2019) report a misallocation level of 48.5%.[25] 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 of allocations made to the high interest rate card by stage. Figure 2.6 provides some suggestive evidence on subjects' tendency to allocate more towards the card with higher balances. In

[22] 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.

[23] We exclude non-first periods for the purpose of this subsection as it is not necessary for a subject to learn her interest information in those periods. The reason for this exclusion is that the interest rate is fixed within a stage and hence a subject who learns her interest rate in the first period within a stage has no incentive to click on the interest rate button in later periods within the same stage.

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

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

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.[26] 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.[27] 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.[28]

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

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

[27] 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%.

[28] 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%.

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 greater than the assigned current balance on the right card and we denote this treatment with the dummy variable *Higher Balance*.[29]

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 2.2 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 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. Gathergood et al. (2019) 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 alloca-

[29] 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.

tion behavior. Ponce et al. (2017), 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.

We find that subjects do not learn to make better decisions within a stage or between stages. Figure 2.7 shows the average fraction of correctly made allocations and the share of optimal allocations within and between stages.[30] 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 ev-

[30] 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.

idence 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).[31]

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.[32]

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 (Gabaix, 2014; Sims, 2003). 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' allocation 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 2.8, the drastic increase in incentives to optimize only leads to small improvement in the share of optimal allocations. In fact, the share of optimal balance reallocations is 27.3% - which is slightly higher than the share of optimal allocations observed in the main part of the experiment. Our findings from balance reallocation is consistent with previous findings (Gathergood et al., 2019; Ponce et al., 2017), which have documented the degree of misallocation is virtually invariant to the economic stakes.

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

[32]Both Gathergood et al. (2019) and Ponce et al. (2017) find that the fraction of correctly made allocations do not increase with the length of account tenure.

2.3 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.[33] Moreover, there has been significant advances in behavioral economics literature that incorporates these insights from psychology to develop descriptive theories of financial decision making (Bordalo et al. (2013), Kőszegi and Szeidl (2012), Gabaix (2014), Schwartzstein (2014), Handel and Schwartzstein (2018)).

2.3.1 Optimization Ability

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 (Nisbett and Ross (1980)) and they think such vividness effects to be generated through differential attention to one portion of the environment (Taylor and Thompson (1982)).[34] Comparing to interest rate information, the vividly displayed balance information might therefore attract greater attention and influence the subsequent decisions more heavily.

[33]These mechanisms include selective attention (Nisbett and Ross (1980), Fiske and Taylor (2013)), mental models (Thompson (2009), Johnson-Laird (2010)), dual process theories (Kahneman (2003), Evans (2006)) and heuristics (Tversky and Kahneman (1974), Gigerenzer and Gaissmaier (2011)).

[34]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 Gabaix (2017) for a detailed discussion.

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 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 (Frijda (1986)).[35] 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 (Soroka et al. (2019), Baumeister et al. (2001), Kahneman (1979)).

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 B.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.

[35] Levin, Schneider, and Gaeth (1998) 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. Kahneman (1979) 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.

2.3.2 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. Table 2.3 presents an overview of our treatments.[36] 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*.[37] 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.

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

[36] See Appendix B.7 for the screenshots of the interface of these new treatments.

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

tively 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.*

2.3.3 Results from Mechanism Treatments

Role of Information Vividness

In Figure 2.9, 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 B1.1 and B1.2 in Appendix).

Figure 2.10 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 B1.3 and B1.4 in Appendix).

Result 8: *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 **D**R do not seem to learn to make better decisions within or between stages, similar to the subjects in **D**B. These results suggest that subjects in **D**R also struggle with learning how to make their allocations correctly.

Role of Framing

In Figure 2.11, 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** ($p = 0.02$). The average allocation to the high interest rate account increases by 46.14 ECU - from 332.4 ECU to 378.54 ECU ($p = 0.04$). The results are qualitatively similar when we relax our sample restrictions and control for demographic information (See Tables B1.5 and B1.6 in Appendix).

Figure 2.12 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 B1.7 and B1.8 in Appendix).

Result 9: *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 **I**B exhibit small yet significant learning which stands in contrast to the subjects' behavior in **D**B. This suggests that the debt frame of the problem do not only interfere with subjects' ability to optimize but also hinders learning.

2.3.4 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 B.2.

2.3.5 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.[38]

In Figure 2.13, Panel A shows the average click rates on current balance and interest rate buttons in each period by stage for the subjects in **D**N. 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 **I**N. In sharp contrast with the click patterns in **D**N, we find that subjects in **I**N 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 **I**N clicks, on average, 0.6 times less on the current balance compared to a subject who is assigned to **D**N ($p = 0.005$) while the click rates on interest rate information is similar across treatments ($p = 0.87$). See Table B4.2. 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.[39]

[38]See Figures B5.11 and B5.12 for the screenshots of the interface.

[39]When we compare the time spent on information buttons across treatments, we find that subjects in **D**N 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 **I**N ($p = 0.07$). Moreover, we find that subjects in **I**N treatment spend significantly less time on the current balance information compared to the subjects in **D**N ($p = 0.001$) although there is no difference in the time spent on interest rate information across these two treatments ($p = 0.62$). See Table B4.3.

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 **D**N and interest rate button if the subject is assigned to **I**N. Figures B4.1 and B4.2 presents the click order data.

Result 10: *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 B.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.[40]
2. **Balance Matching (BM)**: Allocate more into the account with higher balances.[41]
3. **Interest Matching (IM)**: Allocate more into the account with higher interest rates.[42]

Panel A of Table 2.5 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.

[40] 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.

[41] Our definition of the balance matching heuristic is less strict than Gathergood et al. (2019) although it still captures the same intuition that greater balances on an account lead to greater allocations on that account.

[42] 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.

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 2.5 we show the heuristic distribution under each frame when we weaken the classification requirement.[43] This approach allows us to classify a significantly higher portion of the subjects and the results remain qualitatively similar.

Result 11: *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.*

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 B.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

[43]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.

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.

2.4 Discussion

2.4.1 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 (Campbell (2016), Beshears et al. (2018)). 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.[44] 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 (Bertrand and Morse (2011), Seira et al. (2017)). 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

[44] 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.

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.[45]

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 (Hastings et al. (2013), Lusardi and Mitchell (2014)). While confirming that optimization ability is associated with improved decision making, we find that a significant majority 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.[46] 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

[45]Both Bertrand and Morse (2011), Seira et al. (2017) explore psychology-guided disclosures in similar borrowing situations and find them to have modest effects.

[46]There is a substantial educational psychology literature that discusses mechanisms that underlie errors in algebraic thinking and methods to overcome these errors (Herscovics and Linchevski (1994), Stacey and MacGregor (1999)).

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.[47] 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 Lusardi, Michaud, and Mitchell (2017)) should enhance the descriptive power of these models and the accuracy of policy evaluations obtained under these models.

2.4.2 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 Bordalo et al. (2013), a salient thinker allocates strictly greater attention to balance information compared to interest rate information since the balance information shows greater variability.[48] Similar to salience theory, both Kőszegi and Szeidl (2012)'s model of focusing and Gabaix (2014)'s model of sparsity predict greater attention to balance information as the range of outcome utilities differ more in that

[47] A related psychology and economics literature investigates how scarcity might affect various cognitive functions and lead to suboptimal behavior in many domains (e.g. Mullainathan and Shafir (2013)).

[48] 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.

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. Bordalo et al. (2013) 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 Kőszegi and Szeidl (2012) and Gabaix (2014)'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 selective attention where people derive direct utility from attending to information (e.g. Karlsson et al. (2009)). 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.

2.5 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 intrinsic 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.

Figures and Tables

Figure 2.1: Experiment Interface

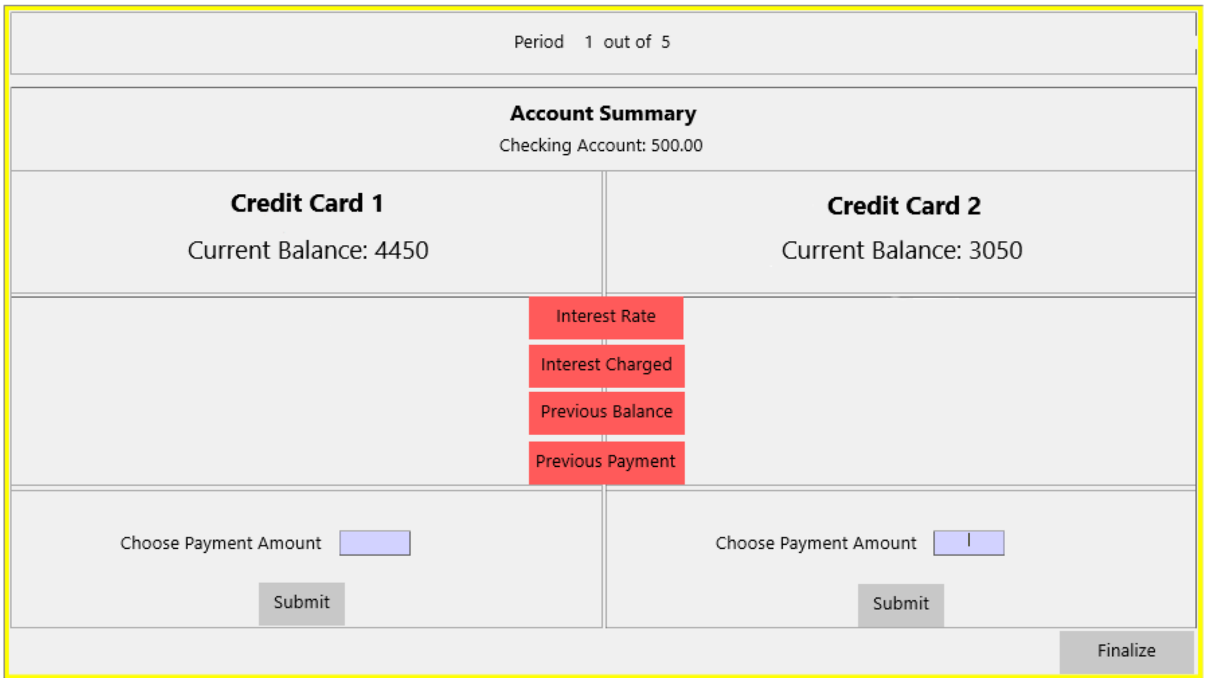

Figure 2.2: Experiment Timeline

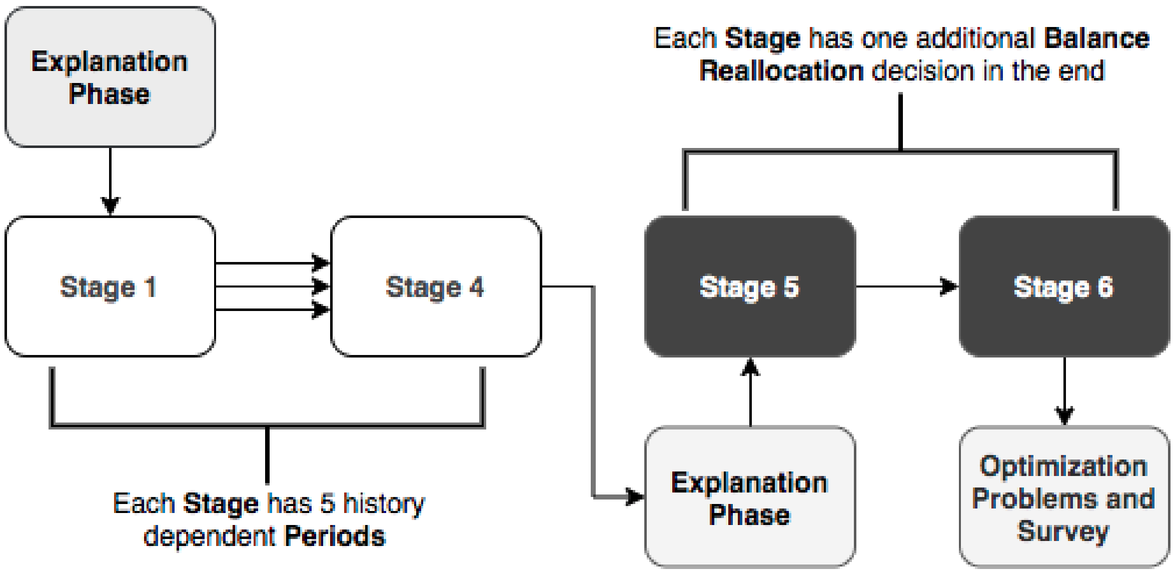

Figure 2.3: Proportion of Subjects Acquiring Interest Rate Information

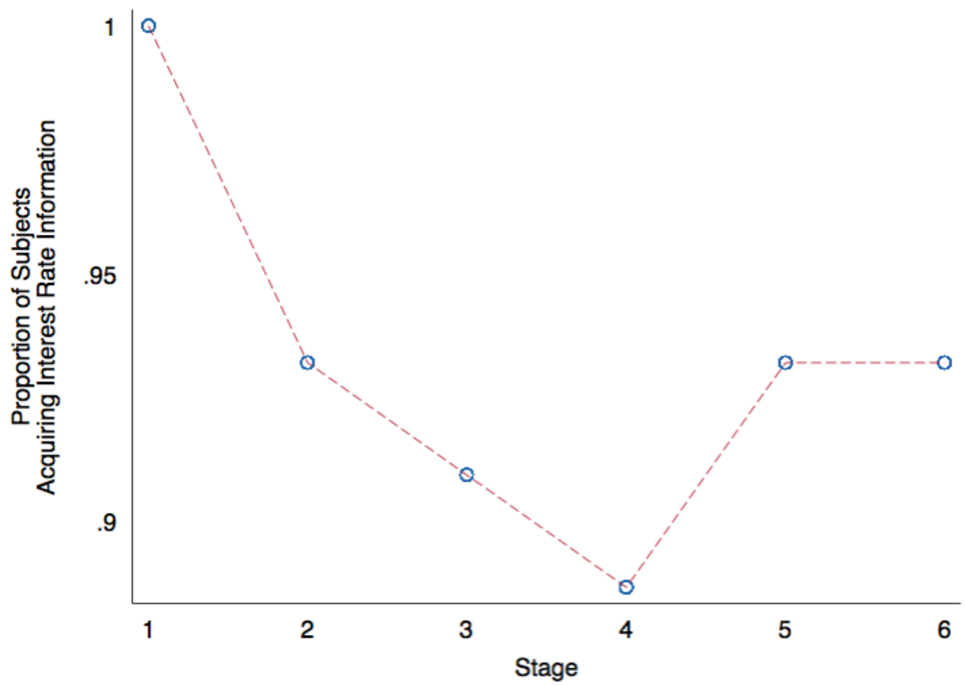

Notes: Figure shows the proportion of subjects acquiring interest rate information by the *first period* of each stage.

Figure 2.4: Distribution of Subjects' Optimization Abilities

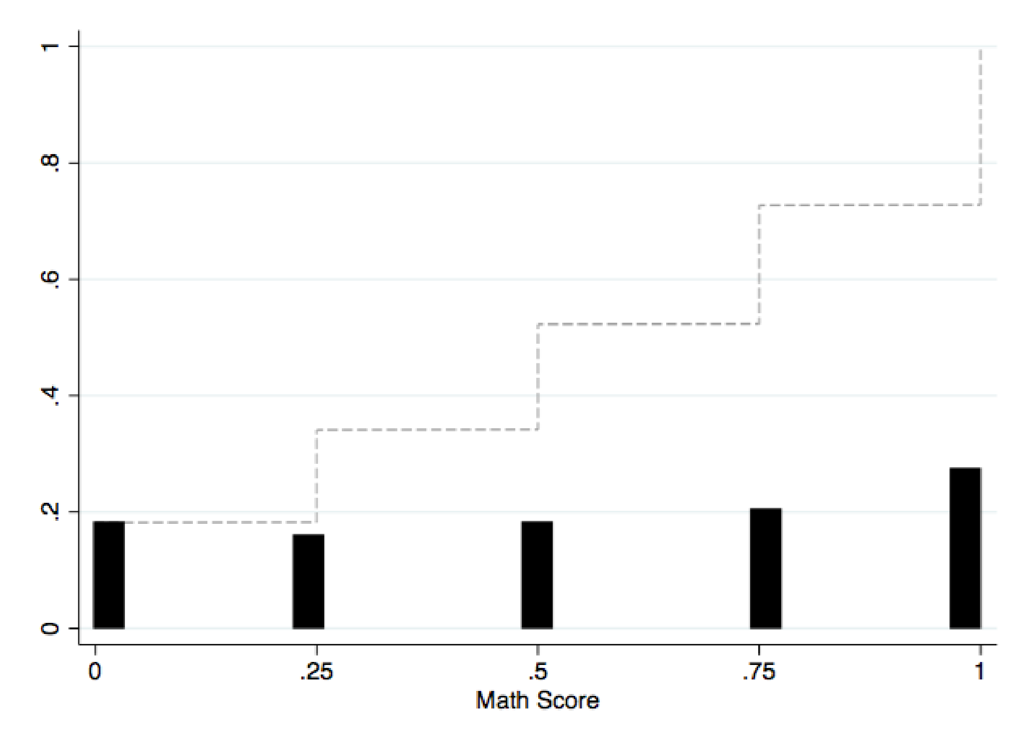

Notes: Figure shows the distribution of subjects' optimization abilities. *Math Score* represents the fraction of correctly answered optimization problems. Each bar represents the fraction of subjects achieving a certain score. The dotted line represents the empirical cumulative distribution function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

Figure 2.5: Distribution of Allocations

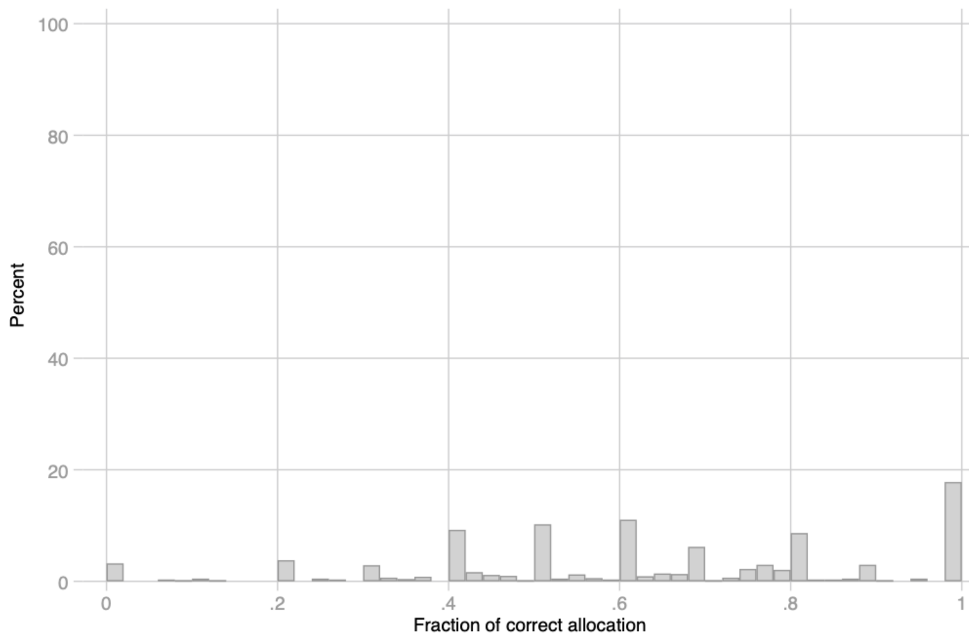

Notes: Figure shows the distribution of fraction of payments subjects make toward the high interest rate card in each period from the Debt Balance Treatment. We have 264 observations at the $subject \times stage$ level. The histogram contains 50 equally sized bins. The rational choice theory predicts a distribution with full mass located at 1.

Figure 2.6: Allocation Patterns Across Stages - Period 1 Decisions

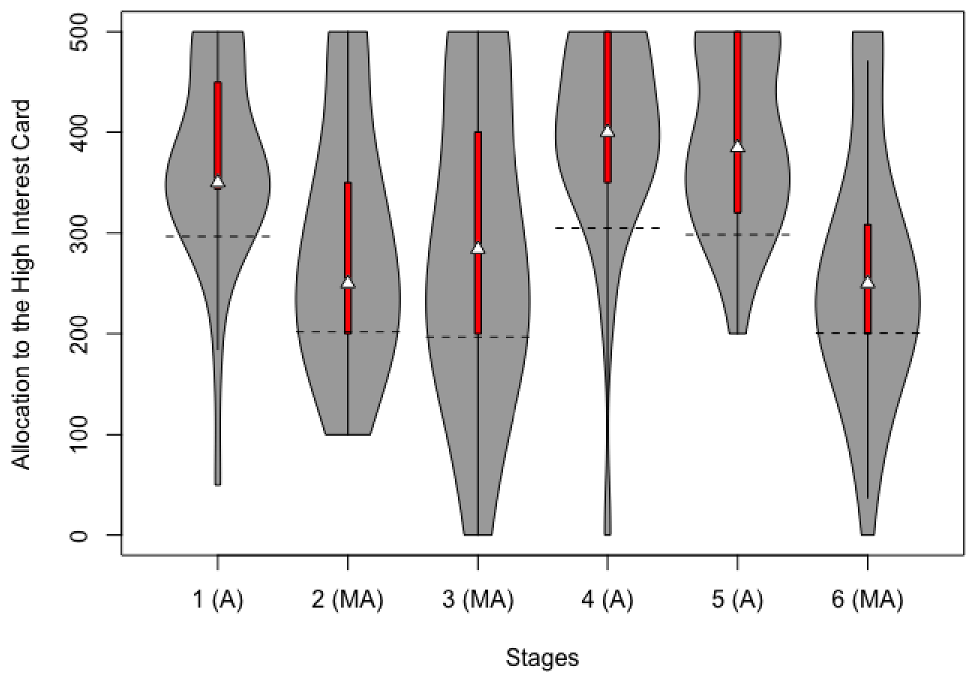

Notes: The violin plot shows the distribution of repayments subjects make toward high interest rate card in the first period of each stage. The center white dot represents the median allocation towards the higher interest rate card in a given stage. The thick bars around the median represents allocations 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.

Figure 2.7: Measures of Optimality Within and Between Stages

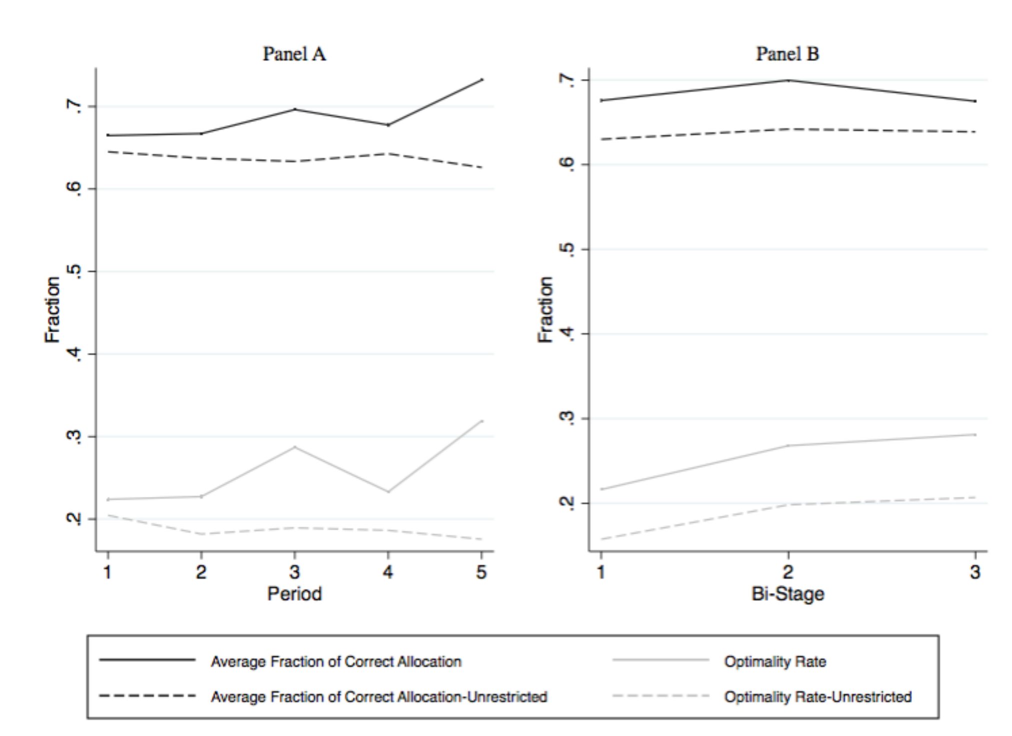

Notes: Panel A shows both the average fraction of correctly made allocations and the share of optimal allocations by periods within a stage. Panel B shows the same optimality measures by bi-stages. A bi-stage consists of two consecutive stages with one aligned and one misaligned stage. The solid lines 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.

Figure 2.8: Distribution of Balance Reallocation Decisions

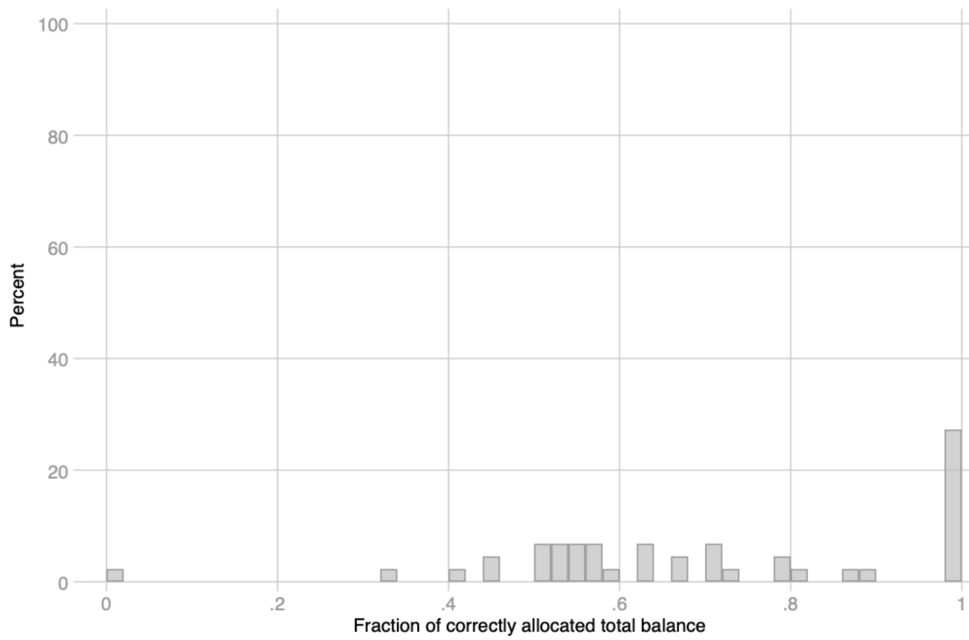

Notes: Figure shows the distribution of fraction of total balances subjects reallocate toward the high interest rate card in balance reallocation periods. The distribution is represented with 50 equally sized bins. We have 88 subject × stage observations. The rational choice theory predicts a distribution with full mass located at 1.

Figure 2.9: Optimality Measures Across Debt Treatments - Period 1 Decisions

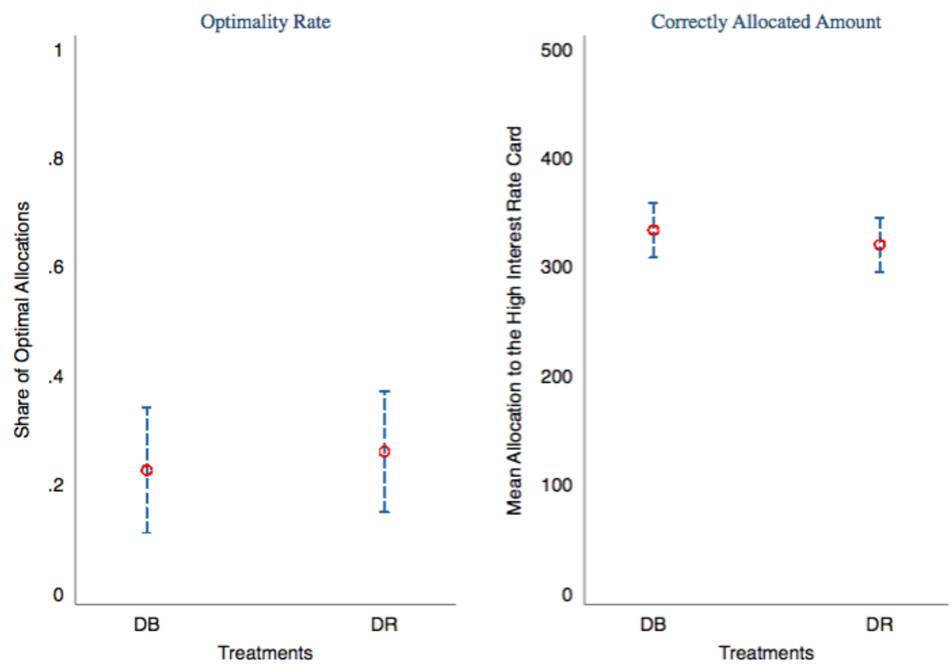

Notes: Panel A shows the share of optimal allocations made under **DB** and **DR**. The whiskers indicate 95% confidence interval calculated using subject-level clusters. Panel B shows the average allocation made to the high interest rate card under **DB** and **DR**.

Figure 2.10: Allocation Patterns Across Debt Treatments - Period 1 Decisions

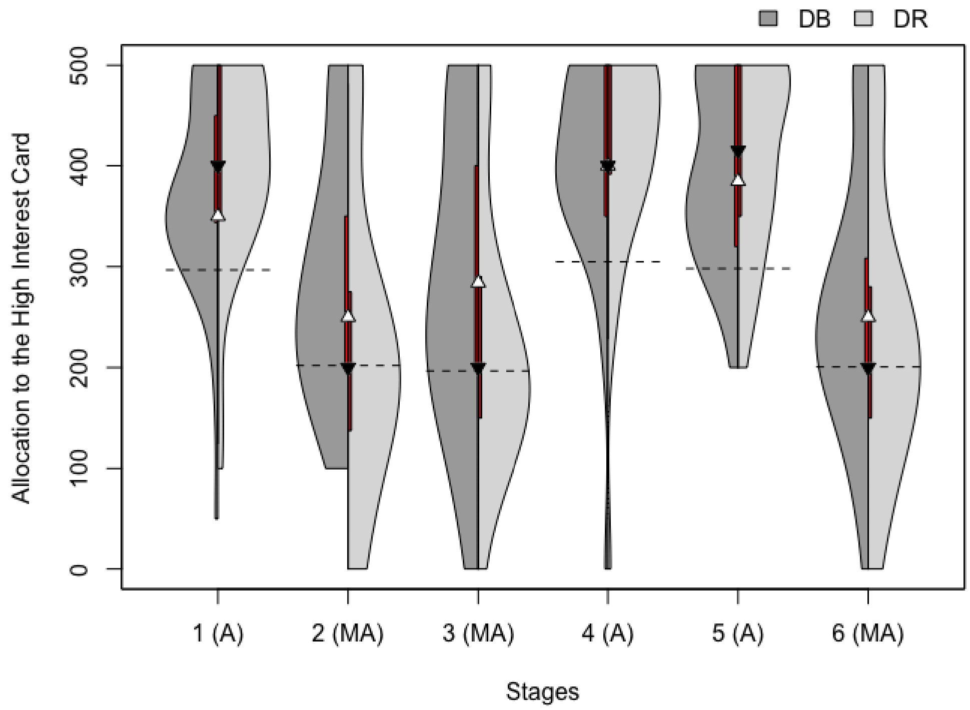

Notes: The violin plots show the distribution of repayments subjects make toward the high interest rate card in the first period of each stage. The upward white triangle and the downward black triangle represent the median allocation towards the higher interest rate card in a given stage for **DB** and **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 2.11: Comparison of Balance Treatments

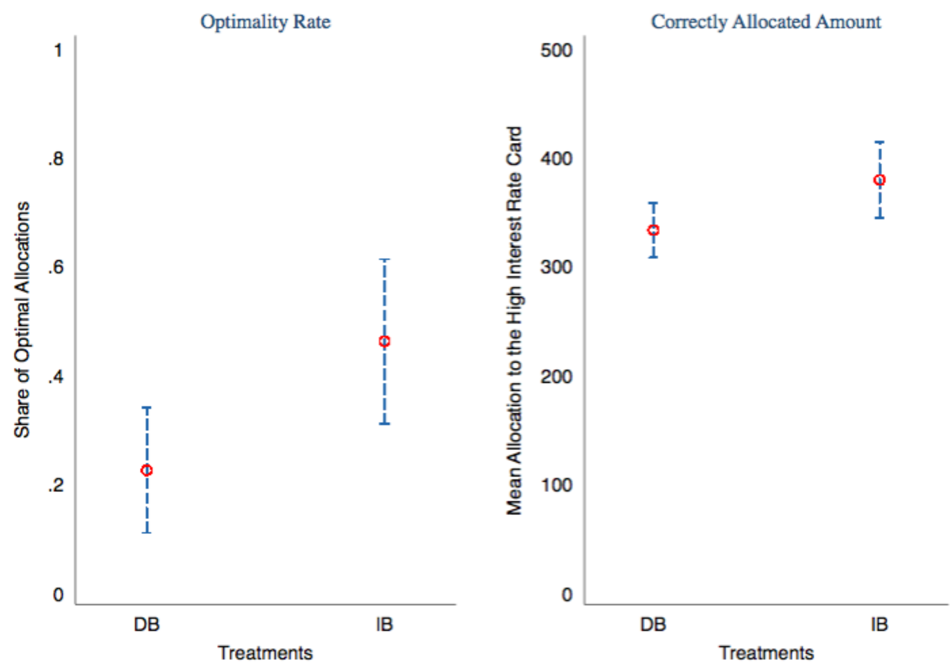

Notes: Panel A shows the share of optimal allocations made in **DB** and **IB**. The whiskers indicate 95% confidence interval calculated using subject-level clusters. Panel B shows the average allocation made to the high interest rate card.

Figure 2.12: Allocation Patterns Across Vivid Balance Treatments - Period 1 Decisions

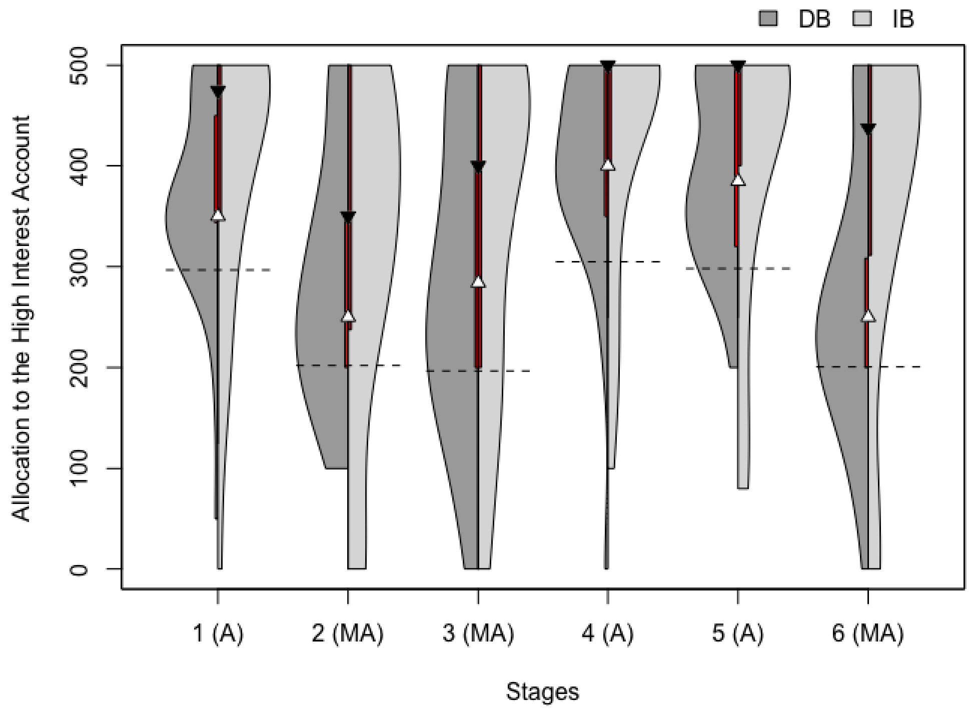

Notes: The violin plots show the distribution of repayments subjects make toward the high interest rate card in the first period of each stage. The upward white triangle and the downward black triangle represent the median allocation towards the higher interest rate card in a given stage for **DB** and **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 2.13: Average Click Rates Across No-Vivid Treatments

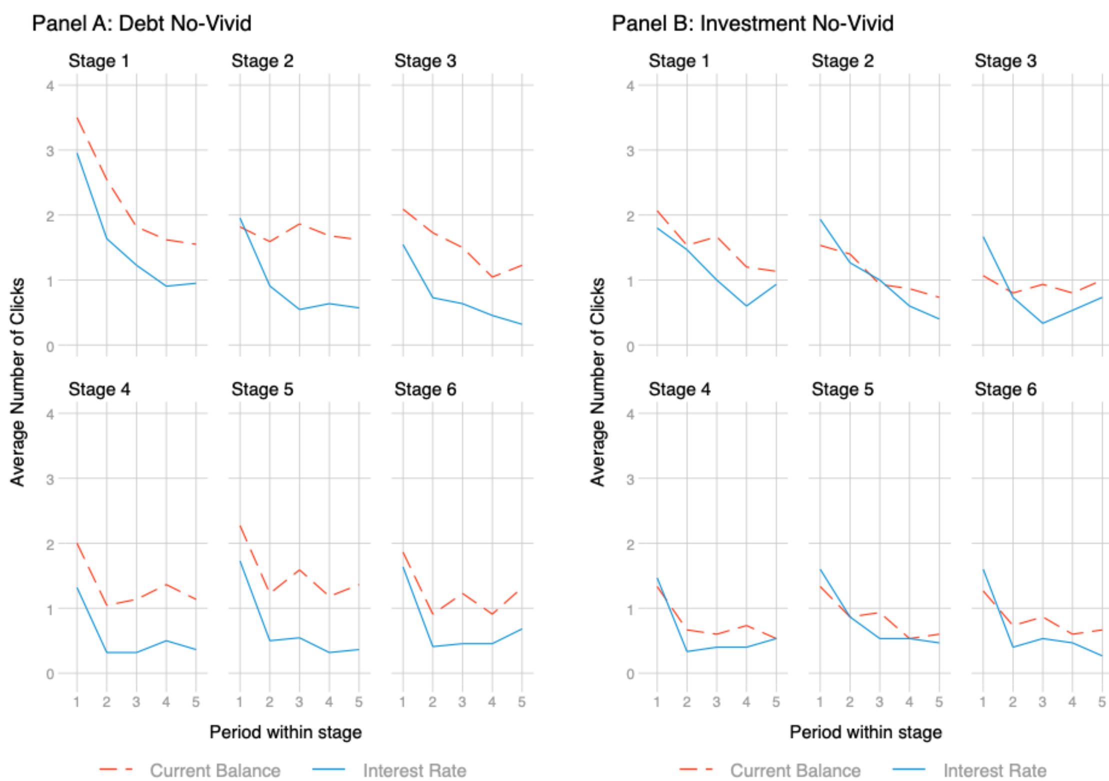

Notes: Panel A documents the difference in average click rates on interest rate and current balance button for each period in each stage under DN treatment. Panel B presents the same measures for IN treatment.

Table 2.1: Parameter Choices and Balance Reallocation

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

Table 2.2: 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

Notes: 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.

Table 2.3: 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

Table 2.4: 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

Table 2.5: 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

Notes: 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 **D**B, DR, DN and across the investment treatments **I**B, IR, IN.

Chapter 3

Paying for Integers

with Jeffrey Cross

3.1 Introduction

In many markets, consumers are faced with situations where they are expected to voluntarily pay extra in the form of a tip for no additional good or service. Historically, tipping is prevalent in particular markets within the United States, such as the restaurant industry, where tips have accounted for more than $40 billion of revenue (Azar, 2008). In the early 2010s, however, cloud-based point-of-sale systems like Square, Inc. were introduced. These systems allow firms to present and customize suggested tip functions in their payment interface. As a result, consumers are increasingly encountering formal prompts and suggestions for tips in settings like coffee houses, where previously there were none. Despite the increasing prevalence in new markets and large revenue in traditional 'tipping markets', economists still understand little about the determinants of consumer tipping behavior.

In this paper, we exploit the unique setting of New York City taxi rides, where we observe high frequency, trip-level responses to preset tip suggestions. Similar to previous work on tips, we document that customers do respond to default tip suggestions. Despite the fact that default tip suggestions do not cluster at integer tip amounts, however, we find that customers have a tendency to tip integer amounts. Furthermore, customers exhibit this behavior despite the fact that tips in this setting are automatically incorporated into final prices by the credit card machine. We thus ask: do customers respond differently to tip suggestions based on whether or not the suggested tip amount is an integer and, if so, what does this reveal about human behavior?

To theoretically explain the tendency for passengers to give integer tips, we use the model of Donkor (2020) as a starting point. In this model, a passenger's preferred tip rate absent a menu (i.e., custom tip) is where the marginal costs associated with tipping more is equal to the marginal gains from smaller norm-deviation costs. When presented with a menu, she then decides if it is worth paying the cognitive costs to tip her preferred tip rate or if she would instead like to select an option from the menu, which has no cognitive costs associated with it. In this model there is no reason for clustering at integer tip amounts, so we extend the decision that passengers make by incorporating lower cognitive costs when giving custom tips that are integers and lump-sum utility gains when giving an integer tip. Both mechanisms, differential cognitive costs and lump-sum utility gains for integer tips, lead to increases in the frequency of integer custom tips. Only in the presence of lump-sum utility gains when giving integer tips, however, are customers more likely to give the suggested tip amount if it is an integer. We leverage this implication of our model, in combination with plausibly random variation in whether a customer is presented an integer tip suggestion, to provide evidence on whether the pattern of clustering at integers is driven, in part, by customers experiencing lump-sum utility gains from tipping an integer amount.

Endogeneity in prices, tipping schemes, and consumer purchasing decisions can make studying consumers' tipping behavior challenging. It could be the case, for example, that integer tip suggestions only occur when customers purchase a certain combination of goods.

Alternatively, it is possible that customers are more likely to select integer tip suggestions because it is easier to calculate the total amount that they must pay. If these customers differ from those that purchase other combinations of goods, then this would bias estimates of the relationship between integer tip suggestions and tipping behavior.

In the context of New York City taxi trips, however, we are able to overcome many of the endogeneity concerns due to the fact that 1) tips are automatically added to the fare amount when selected off the menu and 2) plausibly random variation in tip suggestions depending on the credit card payment machine and surcharges throughout the day. Every passenger that pays with a credit card during the time period that we study is faced with a menu of three tip suggestions: 20, 25, and 30 percent. The total that is used to calculate these suggestions, however, varies between the two credit card payment machines as one does not include the Metropolitan Transportation Authority (MTA) tax of $0.50 while the other does. Importantly, this means that two customers with identical trips, i.e., same date and distance, will receive slightly different tip suggestions depending on the credit card payment machine the taxi is using, which is not evident from the exterior. One of these customers could thus be "treated" with an integer tip suggestion, while the other is presented a nearly identical non-integer tip suggestion. Since surcharges change by day of the week and time, the credit card payment machine that presents integer tips to customers changes thereby allowing us to isolate the effect of an integer tip suggestion from potential confounds like differences in tipping behavior throughout the day or by credit card machine.

We leverage this variation in the occurrence of integer tip suggestions to examine whether customers' behaviors are consistent with a model where they experience direct utility benefits from giving integer tips. Across a variety of specifications and estimation strategies we find consistent support for this hypothesis in the form of increases in take-up of default tip suggestions and tip rates. In our preferred specification where we control for average differences in tipping behavior by driver, hour of that date, and pickup and drop-off census blocks, we find that the probability a passenger selects the default option increases by more than 21 percentage points and tip rates increase by more than 0.6 percentage points. Intuitively, the increase in tip rates is due to the fact that, in our context, passengers tend to give custom tips smaller than all menu options. As passengers switch from custom tips to selecting an option from the menu, this leads to an increase in average tip rates.

The likelihood that a customer faces an integer tip suggestion is jointly determined by the interaction between the fare rate and the tip rate used for the tip suggestions. Given that customers' tipping behaviors respond to integer tip suggestions, a change in prices can indirectly impact the likelihood of integer tip suggestions and with this, revenue. We explore the magnitude of this effect using an increase in the fare rate from 40 to 50 cents in September 2012 that increased the probability of integer tip suggestions from 3% to 21%. When we decompose the effect of the fare rate change on revenue, our estimates suggest that the increase in integer tip suggestions after the fare change led to an increase in revenue of approximately 1.4 cent per trip. With over 170 million taxi trips and 41,000 unique drivers this leads to a transfer of 2.38 million dollars from riders to drivers in the year following the policy change.

Our paper is closely related to the literature that documents clustering around integers or round numbers in other domains of individual decision making (Allen, Dechow, Pope, & Wu, 2017; Lynn, Flynn, & Helion, 2013).[1] People's tendency to use integer or round numbers is commonly associated with lower cognitive cost (Isaac, Wang, & Schindler, 2020; Schindler & Wiman, 1989) or lower trading negotiation cost (Harris, 1991). Although several studies have suggested that this clustering pattern can be rationalized with people's direct preference towards integers or round numbers, there is limited causal evidence. Our paper contributes to this literature in two ways: first, we document a similar pattern of clustering at integer values in the context of taxi tipping; second, we provide theoretical underpinning and causal evidence for this behavior. Specifically, our finding is consistent with a previously under-

[1]Round numbers refer to integers that end with 5 or 0.

explored mechanism that suggests people derive direct utility gains from giving integer tips.

Our paper also relates to the strand of literature that examines the potential drivers of tipping behaviors. The literature has offered causal evidence for a number of mechanisms: for example, customers' tipping decisions can be affected by the default suggestions (Alexander, Boone, & Lynn, 2021; Haggag & Paci, 2014; Hoover, 2019), their compliance to social norms (Donkor, 2020; Thakral & Tô, 2019) and their degrees of social preferences (Azar, 2007; Chandar, Gneezy, List, & Muir, 2019). Similar to this literature, we offer causal evidence for an underlying determinant of consumer tip behavior, utility gains from tipping integers. The modeling approach of this paper, however, is mostly related to Donkor (2020) who focus on estimating parameters for the optimal tipping menu in the presence of customers who conforms to social norms. We extend his model by introducing utility costs/ benefits that are associated with integer tips. Our model enables us to further decompose the mechanism that affects tipping behavior, particularly as it relates to the tendency to tip integers.[2] Our paper thus contributes to our understanding of the determinants of tipping behavior by providing additional causal evidence that passengers respond to integer tip amount suggestions in a manner consistent with direct utility gains from integer tips.

The rest of our paper is structured as follows. Section 3.2 describes the institutional setting, our dataset and the sampling restrictions. Section 3.3 presents descriptive evidence and two models of tipping behavior. Section 3.4 describes the variations tip suggestions and our main econometric specification. Section 3.5 presents customers' estimated responses to inte-

[2] In a related note, our paper offers a new insight that could potentially reconcile the conflicting findings on the magnitude of the default effect in the context of taxi tipping. Using quasi-experimental variation in tip suggestion, Haggag and Paci (2014) documents increase in default raises average tip rate significantly. On the other hand, Chandar et al. (2019) ran a field experiment that manipulates default tip suggestions via Uber and only finds a moderate effect from increasing default options. Chandar et al. (2019) attributes this difference to the reduction in norm compliance under no monitoring (in the case of Uber tipping). The integer effect uncovered in our paper offers an additional explanation to this seemingly conflicting evidence. Specifically, the default change in Haggag and Paci (2014) is accompanied by the increase in the occurrences of integer tip suggestions, whereas all default options offered in Chandar et al. (2019)'s study are in integer terms already. Therefore, the estimated default effect from Haggag and Paci (2014) is a combination of integer and default effect which is indeed greater than Chandar et al. (2019)'s estimated 'net' default effect.

ger tip suggestions. Section 3.6 discusses the implications of varying fare rate and tip suggestions on revenue. Section 3.7 concludes.

3.2 Context and Data

We use data provided by the Taxi and Limousine Commission (TLC) of New York City to estimate the effect of integer tip suggestions on tipping behavior and driver revenue. As of 2008, the entire taxi fleet was outfitted with new equipment that allowed customers to pay using credit cards and also the electronic collection of trip data. Nearly the entire fleet used equipment provided by either Creative Mobile Technologies (CMT) or VeriFone Incorporation (VTS).[3] Taxi cabs equipped by both of these vendors had a Passenger Information Monitor (PIM) which, at the end of a trip, displayed a payment screen. At this point, the devices show a tip menu to passengers who pay with credit cards. Passengers can then choose to give a tip based off the menu options, manually enter in an amount, or provide a separate cash tip.

3.2.1 Context

For standard rate fares, passengers are charged $\$2.50$ and a $\$0.50$ Metropolitan Transportation Authority (MTA) tax upon entering the cab. The fare increases by an additional $\$0.40$, or $\$0.50$ after September 4, 2012, for every fifth of a mile or for every minute where the vehicle travels less than 12 miles per hour. Throughout the period of our analysis, there is a night surcharge of $\$0.50$ for trips between 8 PM and 6 AM and a $\$1.00$ surcharge for trips between 4 and 8 PM on weekdays.

At the end of each trip, passengers are shown trip expenses through the touch-screen payment device. Passengers that pay with a credit card are then presented with a tip menu

[3] We do not use data from a third vendor, Digital Dispatch Systems, which accounted for less than 5% of electronic transmission devices in use in 2010.

that varied by vendor over time. An example of this screen for a CMT outfitted vehicle in 2012 is shown in Appendix C3.1. Based on the selection of the passenger for the tip, the total is calculated and the passenger proceeds with payment. If the taxi uses a CMT device, the tip menu calculates tips on the total fare, which includes the base fare, MTA tax, tolls, and any surcharges. Alternatively, for a VTS device, tips are calculated using the base fare and the surcharge, but does not include tolls or MTA tax. In Figure 3.1, we show the menu of tip suggestions for CMT and VTS devices over time. Prior to February 9, 2011, customers in taxicabs with CMT devices were presented with tip suggestions that were 15, 20, and 25 percent. From February 9, 2011, onward all options on the CMT menu went up to higher tip percentages of 20, 25, and 30. For VTS devices, tip suggestions changed in January of 2012. Prior to that month, they offered a tip menu of dollar amounts (\$2, \$3, and \$4) if the base fare and surcharge was under \$15, and suggestions of 20, 25, and 30 percent for larger fares. After that month, VTS offered only the percentage choices (20, 25, and 30), regardless of the trip fare.

3.2.2 Data

Our data consists of trip (ride) level data on all tax rides in New York City and surrounding counties from 2010 to 2013. For each trip, our data records the date, time, and geographic location of the pickup and drop-off. Each observation is recorded with a unique medallion number and a taxi driver license number. These numbers identify a unique cab and driver for any given year, but cannot be used to identify drivers or cabs across years. In addition, the equipment records information on trip time, trip distance, fare amount, tolls, tax, surcharge, rate code, and payment method. For all customers that pay digitally when using a credit card, we observe the tip entered into the credit card machine. Importantly, however, we do not observe tips for trips paid with cash, and we cannot interpret manually entered tips of 0 when paying with a credit card as a tip of 0.

To account for potential differences in customer characteristics, we use data from the American Community Survey's 5 year estimates (2006-2010), which consists of census tract level summary statistics. We leverage the GPS coordinates for each pickup and drop-off location to assign each trip pickup and drop-off census tracts. We then merge this with the ACS census tract variables so that we can characterize the median income of where a customer is picked up and dropped off.

We take many of the same steps to cleaning the data that have been used in the previous literature, see Haggag and Paci (2014). Since we do not observe tip information for trips or tips paid by cash, we drop these and focus on trips paid with credit cards that have positive tips in our analysis.[4] In addition, our primary analysis focuses on all trips that use standard rate fares. We do this in large part, since our primary results leverage plausibly exogenous variation in tip suggestions present in the standard rate fare, which is not present with all other rate fares.[5] The conclusions from our analysis, however, do not change when including trips with all rate fares. To ensure that our results are not influenced by drivers changing between vendors, we drop all drivers that change vendors within the same year. Similar to Farber (2015), for simplicity we focus on a random sample of drivers in all of the analysis that follows. Specifically, since we cannot track drivers across years, we use a sample of 2,000 random taxi driver and car pairs for each year.[6]

In our primary analysis, we utilize variation in the decimal places of a constant menu of tip suggestions, 20, 25, and 30 percent. Our preferred subsample focuses on the time window

[4]In our sample, about 55% of the payments were made by cash. The differences between trips with cash and credit payments are shown in Table C1.1. The table does not highlight the fraction of zero tip trips that are paid with credit cards. To highlight that this is a small fraction, we include these trips in our primary descriptive figures but we will exclude them in the regression analysis.

[5]The rate for trips between JFK to Manhattan, for example, is fixed and would introduce non-random variation in tip suggestions.

[6]We include the detailed data refinement procedure in Appendix C.1. To ensure that our results are robust to the larger dataset, we drew a second random sample. The conclusions are identical regardless of the sample we use.

from February to August of 2012, where all standard fare rides were subject to the same rate fare and menu of tip suggestions, regardless of vendor. This offers the key advantage of a single distribution relating rate fare to tip suggestions that all customers are subject to for a significant length of time.

3.3 Tipping Behavior

Previous research (e.g., Haggag and Paci 2014) has documented that customers respond to tipping menus by changing tipping behavior. Little attention, however, has been paid towards patterns of customer tipping behavior and the implications for tipping rates. In this section, we will first document descriptive evidence on patterns of customer tipping behavior. We will then introduce a theoretical model to guide our empirical analyses.

3.3.1 Descriptive Evidence of Tipping Behavior

A typical taxi ride experience ends with tip payments. On the payment screen, taxi passengers are often prompted with three tip suggestions and a number pad that allows them to enter any non-negative custom tip amounts. It is well documented that passengers' tip decisions are influenced by defaults and menus. Before analyzing the distribution of these choices, we first want to analyze whether there are any differences between VTS and CMT trips. Table 3.1 presents the summary statistics at the trip level for our preferred subsample, split by CMT and VTS. Although there do not appear to be differences in the details of the trips (e.g., distance or time), there are differences in average tipping behavior. Passengers that ride in vehicles with VTS equipment tend to give a higher tip rate in part due to selecting options from the menu at a higher rate. This is likely due to differences in equipment or presentation of the tip suggestions, which has been highlighted by previous research (Hoover, 2019) and will not be the focus of our analysis.

As illustrated in Figure 3.2a, over 40% of the tips were made at the default options under a menu of tip suggestions at 20, 25, and 30 percent. However, the distribution of tip rate does not provide a holistic perspective of customers' tipping behaviors as it ignores potential patterns that might exist in the nominal tip values. Indeed, when we plot the raw distribution of tip amounts in Figure 3.2b, we observe clustering of tips at integer values. This is unlikely since the second decimal place of tip suggestions are equally likely to be an even number as we show in Figure 3.3a.[7] Although the second decimal places are equally likely, it is evident in Figure 3.3b that passengers are more likely to tip the suggested amount when the low suggestion is an integer. The clustering of tip amounts at integers appears to be driven, in part, by this increased tendency for passengers to select default integers when they are integers in combination with custom tips at integer values, as is evident in Figure C3.3.

Overall, a visual inspection of aggregate tipping behavior suggests that (1) customers respond to default suggestions, (2) customers tend to tip at integer values, and (3) the tendency to give integer tips is evident in custom and, to a lesser extent, default tips.

3.3.2 Models of Tipping Behavior

Given the large number of taxi drivers, we will model tipping behavior as being primarily influenced by the pressure of social norms (Azar, 2007) instead of strategic incentives (Azar, 2008).[8] Following Donkor (2020), consider a passenger i that gives a tip of $t_i\%$ at the end of her taxi ride that costs F_i. She believes that the socially accepted tipping rate to give based on the ride is $T_i\%$, which can vary by passenger. If her chosen tip rate is different than what she believes is the socially accepted tipping rate, then she incurs a norm-deviation cost of $v(T_i, t_i)$. Assume that for any fixed T_i, the norm-deviation cost of $v(T_i, t_i)$ is convex with respect to t_i with a minimum at $t_i = T_i$. When making her tipping decision she is presented a menu

[7] The distribution of tip suggestions for each of the options is shown in Figure C3.2

[8] There are over 10,000 Yellow taxis in New York City, which minimizes the potential for repeated passenger and driver interactions.

of tipping options D, which consists of a variety of suggested tipping percentages. Without loss of generality, denote the preferred option out of the menu for customer i as t_i^D. In order to choose an option that is not on the menu, she incurs a cost c_i that reflects the cognitive costs associated with finding her ideal tip percentage. The utility maximization problem for passenger i can be written as:

$$Max_{t_i} U = -t_i F_i - v(T_i, t_i) - c_i \cdot \mathbb{1}\{t_i \neq t_i^D\} \tag{3.1}$$

The first term represents the passenger's expenditure. The second term represents the cost of deviating from her perceived socially accepted tipping rate, T_i, and the last term captures the cost of computing a tip not presented in the tip menu.

Since the cognitive cost for all custom tip options are the same, the utility-maximizing custom tip rate is the tip rate that maximizes the first two terms. Given the assumption on the functional form of $v(T_i, t_i)$ the utility-maximizing custom tip satisfies:

$$\frac{\partial v}{\partial t_i} = -F_i \tag{3.2}$$

Intuitively, this shows that she will increase her tip rate until the marginal return of reducing the norm deviation cost, $\frac{\partial v}{\partial t_i}$, is equal to the marginal cost of increasing the tipping rate, $-F_i$. For now, denote the custom tip that solves equation (3.2) as t_i^C. As the left panel from Figure 3.5 shows, this means that, even absent cognitive costs, she will not give the socially accepted tipping rate, but will instead "shade" downwards and give a custom tip rate less than T_i.

Working backwards, the passenger then decides if she will give the custom tip or instead choose a default option from the menu. It is only worth the cognitive cost of calculating and manually entering the custom tip if:

$$\underbrace{[-t_i^C F_i - v(T_i, t_i^C)]}_{U(t_i^C) \text{ if } c_i = 0} - \underbrace{[-t_i^D F_i - v(T_i, t_i^D)]}_{U(t_i^D)} > c_i \tag{3.3}$$

The left side of the equation captures the utility gains from manually entering a custom tip relative to selecting a default option if cognitive costs were 0. A passenger compares this to the cognitive costs on the right side, and then decides if the custom tip is worth calculating. All else equal, passengers are more likely to select custom tips if their cognitive costs are low or, alternatively, if they strongly prefer the custom tip to the default options.

Given the utility problem presented in equation (3.1), it is difficult to explain the pattern of tips at integer values. Tipping integer values represent different tipping rates across trips so t_i^D does not naturally cluster at integer values. In addition, it is unlikely that utility-maximizing custom tips that satisfy equation (3.2) would lead to disproportionately more integer custom tips relative to non-integer custom tips. To better explain the concentration of tips at integer values, we will now propose two extensions to the model.

First, it is possible that passengers give integer tips because it decreases the cognitive costs associated with computing the ideal tip. In other words, if a passenger believes that the suggested options are too high, she might choose a lower tip that is close to the ideal tip percentage, but is an integer and is thus less cognitively costly. To incorporate this into the passenger's problem, let there be a lower cognitive cost $c_i^{int} < c_i^{non}$ when the selected tipping choice is an integer. Define the difference in cognitive costs as $\alpha_i = c_i^{non} - c_i^{int} > 0$, which is passenger specific. The second potential mechanism behind customers tipping integer amounts is that passengers, in general, feel more comfortable giving integer tips. We model this as a lump-sum utility gain, b_i, whenever a passenger tips an integer, regardless of whether it is on the menu or not. We can then write the utility maximization problem for passenger i as:

$$Max_{t_i} U = -t_i F_i - v(T_i, t_i) - \mathbb{1}\{t_i \neq t_i^D\}[c_i^{non} - \underbrace{\alpha_i \cdot \mathbb{1}\{t_i F_i \in \mathbb{Z}\}]}_{\text{Reduced Cognitive Costs}} + \underbrace{b_i \cdot \mathbb{1}\{t_i F_i \in \mathbb{Z}\}}_{\text{Integer Utility Gain}} \quad (3.4)$$

which nests the previous model of passenger utility, shown in equation (3.1), but we now allow for integer tips to directly impact utility as a lump-sum utility gain and lower cognitive costs when giving a custom tip. Figure 3.6 illustrates an passenger's decision process under the extended model.

The inclusion of differential cognitive costs and lump-sum utility gains when giving an integer tips impact the utility-maximization problem in two key ways. First, in the model presented in equation (3.1), the choice of custom tip rate is where the marginal return of reducing the norm deviation cost, $\frac{\partial v}{\partial t_i}$, is equal to the marginal cost of increasing the tipping rate, $-F_i$. As the right panel from Figure 3.5 shows, however, this need not be the case in the extended model. The tip rate that satisfies equation (3.2) might not be utility-maximizing if the tip amount is not an integer. Intuitively, this is because the benefits from an integer tip suggestion, $b_i + \alpha_i$, can outweigh the lower utility from not equating the marginal return of reducing the norm deviation cost to the marginal cost of increasing the tipping rate.

To show this more formally, define the custom tip rate that satisfies equation (3.2) as t_i^{non} and the preferred custom integer tip rate of t_i^{int}, which has a lower cognitive cost.[9] For arbitrary benefits and additional cognitive costs b_i and α_i, she will choose to give a non-integer custom tip that satisfies equation (3.2) if:

$$\underbrace{[-t_i^{non} F_i - v(T_i, t_i^{non})]}_{U(t_i^{non}) \text{ if } c_i^{non}=0} - \underbrace{[-t_i^{int} F_i - v(T_i, t_i^{int})]}_{U(t_i^{int}) \text{ if } c_i^{int}=0} > \underbrace{c_i^{non} - c_i^{int}}_{\alpha_i} + b_i \quad (3.5)$$

[9]Intuitively, given the functional form assumptions on $v(T_i, t_i)$ the customer will have a preferred integer custom tip rate that rounds up or down from $F_i t_i^{non}$.

The left-hand side represents utility gains from giving the preferred non-integer tip, which satisfies equation (3.2), relative to the integer tip. If this outweighs the benefit of giving an integer tip, shown on the right-hand side, then she will give the non-integer tip. As the benefits of the integer tip increase, she is increasingly likely to prefer an integer custom tip. Alternatively, as the benefits from an integer tip approach 0, she is more likely to give the custom tip rate that satisfies equation (3.2), t_i^{non}.

The second way that α_i and b_i impact how a passenger tips is through the decision between custom and default tips. Denote the preferred custom tip as t_i^C, which need not satisfy equation (3.2), and define I^C and I^D as indicator variables equal to one if the custom and default tip rates lead to integer tip amounts. When choosing between the custom and default tip rate, she will give the custom tip rate if:

$$\underbrace{[-t_i^C F_i - v(T_i, t_i^C) + I^C \cdot b_i]}_{U(t_i^C) - c_i^{non}} - \underbrace{[-t_i^D F_i - v(T_i, t_i^D) + I^D \cdot b_i]}_{U(t_i^D)} > \underbrace{c_i^{non} - \alpha_i \cdot I^C}_{\text{Cognitive Costs}} \tag{3.6}$$

where the left-hand side represents gains from giving custom tips without considering the cognitive costs. If this is larger than the cognitive costs on the right-hand side, then she will choose to “pay" the cognitive cost for the custom tip rate, t_i^C.

The gains and cognitive costs associated with the custom tip rate now depend on whether the preferred custom tip rate leads to an integer tip, but also, importantly, on whether the default tip suggestion is an integer. For a given default tip rate, a small change in the fare that leads to an integer default tip suggestion sharply increases the utility of the default tip option based on the magnitude of b_i. If b_i is small, then whether or not the default tip is an integer will likely have no effect on the choice between a custom or default tip rate. If it is large enough, however, then this could switch customers away from a custom tip rate towards the default tip option.[10] Importantly, this highlights that the impact of integer default tip suggestions depend on the magnitude of b_i instead of α_i.

In summary, assuming b_i is larger than 0 on average, we have the following hypothesis:

Hypothesis 1 (H1): *When percent-based tip suggestions lead to integer tip amounts, passengers are more likely to choose the default option.*

The intuition for this hypothesis is evident in equation (3.6). If $b_i > 0$, then when $I^D = 1$ (integer tip amount suggestion) the left-side decreases thereby lowering the likelihood the passenger gives a custom tip and, in turn, increasing the likelihood a default option is chosen. A natural implication of our first hypothesis in this setting where custom tip rates, t_i^C, tend to be lower than the default tip rates, t_i^D:

Hypothesis 2 (H2): *Average tip rates of passengers are higher when percent-based tip suggestions lead to integer tip amounts.*

This hypothesis falls naturally from Hypothesis 1 combined with the fact that custom tip rates tend to be lower than the default tip suggestions. Passengers that normally would give a lower custom tip rate switch to the default tip suggestion, which leads to a higher average tip rate. In the following section, we detail how we leverage our empirical setting to test these hypotheses.

3.4 Identification Strategy

To identify what drives the clustering of integer tips, we need to combine the predictions obtained from the theoretical models with random variation in the frequency of integer tip

[10]There is no obvious reason that utility from custom tips would have a similarly sharp change to integer default tip suggestions. We explore this, however, in Appendix C.2 where we parameterize the utility function. We find no response in custom tip utility, including cognitive costs, when there are integer default tip suggestions.

suggestions that we observe in the raw data. In this section, we first describe the quasi-random variation in the frequency of integer tip suggestions across the two vendors. We then present the econometric specifications for our main analysis.

3.4.1 Variation in Tip Suggestions

From February to August of 2012 all standard fare rides where customers paid with a credit card, regardless of the vendor or fare, were presented with tip suggestions of 20, 25, and 30 percent. This means that for every fifth of a mile or for every minute where the vehicle travels less than 12 miles per hour, the three options on the menu of suggested tips increase by $\$0.08$, $\$0.10$, and $\$0.12$, respectively. With a base fare of $\$2.50$, this means that the lowest tip suggestion on the VTS menu, $\gamma_{i,j}^{VTS} = 0.2 \cdot F_i$, between pickup location i and drop-off location j can be defined as:[11]

$$\gamma_{i,j}^{VTS} = 0.50 + 0.08x(d, mph) + 0.2s \tag{3.7}$$

where x is a function of distance (d) and speed (mph), which implicitly depend on the locations i and j. In addition, the tip suggestion depends on s, which is a categorical variable equal to $\$0.00$ if there is no surcharge, $\$0.50$ if it is a night surcharge, and $\$1.00$ if there is a peak weekday surcharge. Alternatively, since CMT trips include the MTA tax ($\$0.50$) and tolls when calculating the tip percentage, the lowest tip suggestion, $\gamma_{i,j}^{CMT}$, can be defined as:

$$\gamma_{i,j}^{CMT} = 0.60 + 0.08x(d, mph) + 0.2(s + \tau) \tag{3.8}$$

where the base tip suggestion increases by $\$0.10$, reflecting the MTA tax, and tip suggestions now depend on the cost of tolls, τ.

A key implication from these two formulas is that, at any given point in time the probabil-

[11] We focus on the lowest option on the menus here since it is the option that passengers select most frequently.

ity of an integer is different depending on the vendor and whether or not there is a surcharge. This is shown in Figure 3.7, which plots the low tip suggestions by $x(d, mph)$, surcharges, and vendor. Points shown in black represent integer tip suggestions. For taxicabs using VTS, regardless of the pickup and drop-off locations, the probability that γ_{ij}^{VTS} is an integer is 0, except when there is a night surcharge, $s = \$0.50$. For VTS trips that travel at least a fifth of a mile, it is only in the case where $s = \$0.50$ and $x = 25y + 5$ where y is an integer greater than or equal to 0. On the other hand, trips in CMT taxicabs have an integer γ_{ij}^{CMT} when $x = 25y + 5$, and there is no surcharge, $s = \$0.00$. If there is a night surcharge, then the probability of a low integer tip suggestion is 0. Alternatively, if it there is a peak weekday surcharge, $s = \$1.00$, then a CMT trip will give an integer tip suggestion for the lowest option if $x = 25y + 15$. In summary, for the majority of the day CMT trips have a positive probability of an integer low tip suggestion, but this changes at night when only VTS trips have a positive probability of an integer low tip suggestion.

Assume that for any given passenger, $x(d, mph)$, is fixed and known. This raises two concerns when estimating the impact of integer tips on tipping behavior based on the tipping suggestion formulas shown in equations (3.7) and (3.8). First, if customers can sort on vendors then this could lead to non-random variation in probability of an integer tip. This concern of selection by the rider is mitigated by the fact that, prior to entry, taxicabs with VTS or CMT credit card machines appear essentially identical. Second, it is possible that, absent an integer tip suggestion, customers (trips) that are likely to have an integer tip suggestion are different in tipping behavior compared to those that are unlikely to have an integer tip suggestion. For example, in this case where $x(d, mph)$ is deterministic, even though there is not sorting by vendor, all customers except those with $x = 25y + 5$ or $x = 25y + 15$ will never be presented with a low integer tip suggestion. Although there is no reason to believe ex-ante that these customers should differ from those that have a low, or zero, probability of receiving an integer tip suggestion, we mitigate this concern by controlling for average tipping behavior in pickup and drop-off locations.

In reality, even when taking a trip between the same locations, $x(d, mph)$ is likely to vary due to exogenous factors like rainfall. There is also, however, the concern that drivers could manipulate $x(d, mph)$ in a way that leads to an unobserved correlation between tipping behavior and the probability of an integer tip suggestion. If there was manipulation by drivers, to induce more frequent integer occurrences then this should be apparent in more frequent tip suggestions ending with a 0. One would expect that this would show up as a higher frequency of 0 in the second decimal place, however, which is not what we see in Figure 3.3a.

3.4.2 Econometric Specifications

Each passenger that pays with a credit card faces a set of tip suggestions at the end of the trip. None of these options have an integer tip suggestion in the vast majority of the time, but, as we detailed in the previous section, the treatment and control groups change throughout the day depending on the vendor and distance of the trip. To isolate the effect of the integer tip suggestion, we attempt to control for each of these factors, such as time and distance, which could influence tipping behavior.

Let D_{ijcdhm} denote whether a trip from location i to location j in taxi c on date d, pickup hour h, and pickup minutes m has a nominal tip suggestion in the menu that is an integer. We estimate the effect of D on the probability a customer selects an option from the tipping menu and the tipping rate, defined as the fraction of the fare rate tipped, as:

$$y_{ijcdhm} = \alpha + \beta D_{ijcdhm} + \delta x_{ijcdhm} + \gamma I_{ijcdh} + \epsilon_{ijcdhm} \tag{3.9}$$

where y is either an indicator for whether a passenger gives a suggested tip or the tipping rate, which we define as the tip amount divided by the rate fare used when determining the tip suggestion. Our coefficient of interest is β, which estimates the effect of a an integer tip suggestion on selecting a default tip option or the tipping rate. According to our hypotheses based on the behavioral model, if customers experience increased utility from tipping integers ($b_i > 0$) then we will find that integer tip suggestions increase the probability of selecting the default tip option and the average tipping rate ($\beta > 0$). Our model also shows that tipping behavior varies by the fare, so we linearly control for $x(d, mph)$. In addition, we control for average differences in tipping by driver, location, and over time with driver, date by hour, pickup census block, and drop-off census block fixed effects, I_{ijcdh}, in our preferred specification. Although this is our preferred specification, we will vary the controls to ensure the robustness of our results to alternative specifications. In all specifications standard errors are two-way clustered at the driver level and the pick-up date level.

To test for heterogeneous effects of integer tip suggestion, we define D^1_{ijcdhm} if the lowest of the three tip suggestions is an integer nominal tip suggestion, and D^2_{ijcdhm} if either of the other two options are an integer nominal tip suggestion. We then estimate:

$$y_{ijcdhm} = \alpha + \beta_1 D^1_{ijcdhm} + \beta_2 D^2_{ijcdhm} + \delta x_{ijcdhm} + \gamma I_{ijcdh} + \epsilon_{ijcdhm} \tag{3.10}$$

where β_1 estimates the impact of an integer tip suggestion if it is the lowest option on the menu, and β_2 shows the effect if it is either of the other options. If the effect of an integer tip suggestion is identical regardless of its place on the menu, then our estimate for β_1 and β_2 should be the same. We include the same controls in our preferred specification and cluster the standard errors at the driver level.

3.5 Results

We present our empirical findings in this section. In Section 3.5.1, we show baseline results from regression analyses that utilize quasi-random variation in the occurrences of integer tip suggestions. In Section 3.5.2, we complement the baseline findings with a series of robustness checks. In Section 3.5.3, we provide additional results on the tip-rounding behaviors for individuals who opted-in for custom tips.

3.5.1 Regression Analysis

Table 3.2 presents our primary results for the effect of integer tip suggestions on tipping an option from the menu. We find that, when presented with an integer tip suggestion, passengers are approximately 21 percentage points more likely to give a tip equal to one of the suggested options. The last three columns highlight that this response can be largely attributed to when the lowest suggestion is an integer. Although this provides some evidence that clustering around integer values, we cannot rule out that a part of the effect we are estimating is due to some customers always tipping an integer value, regardless of the suggestion or percentage. If this behavior is at lower integer values, then this could also potentially explain that the effect is largely driven by the low option.

If the mechanism behind the result shown in Table 3.2 is customers always tipping integer values, regardless of tip suggestions, there should be no difference in the average tipping rate when the suggestions are an integer. This strategy would be, by definition, independent of the tipping suggestion, and thus should not lead to a larger (or smaller) tipping rate based on integer tip suggestions on the menu. Table 3.3 presents our main results for the effect of integer tip suggestions on passenger tipping rates. In our preferred specification, we find that passengers increase their tipping rates by 0.006 (0.6 percentage points) when they are presented with an integer tip suggestion. Column 6 shows that this is due to integer tip suggestions for the lowest option, and not the other two options.

The results from Table 3.2 and 3.3 show that customers are, in fact, responding differently to the presentation of integer tip suggestions, particularly if it is the lowest option. From the visual evidence shown in Figure 3.2b, it was evident that customers tend to give integer tips. These results match our hypotheses, which suggests that based on our behavioral model passengers do experience additional utility from giving an integer tip. Specifically, only when $b_i > 0$ would our model suggest this pattern of behavior. Even when there are cognitive costs associated with tipping, there is no reason that cognitive costs associated with giving custom tips change discontinuously when the tip suggestion is an integer.

3.5.2 Robustness Checks

Leveraging plausibly exogenous variation in the decimals of tip suggestions in 2012, we found that customers increase the their tipping rate by 0.6 percentage points when presented a tip suggestion that is an integer. This is driven almost entirely by integer tip suggestions in the lowest menu option. To ensure that our findings are not a result of how we defined the tipping rate, the sample restrictions, or the identification strategy used, we will take alternative approaches to each of these aspects of our analysis.

First, in our primary results we defined the tipping rate as the tip amount divided by the rate fare used when determining the tip suggestion. This varies by vendor since CMT includes tolls and the MTA tax, while VTS does not. To make the denominator comparable, we calculate the total cost of the trip, except for the tip, and use this as the denominator when defining the tip rate. The results from estimating equations (3.9) and (3.10) with this alternative method of calculating the tip rate are shown in Table 3.4. Our results with this method are similar to our primary results, albeit slightly smaller with our preferred specification showing a change in the tip rate of approximately 0.56 percentage points.

Second, to check that our results are not driven by our sample restrictions, we conduct two additional robustness checks. We first expand the sample to include all trips, not just standard rate trips during the time period of February to August 2012.[12] As we show in Table 3.5,

[12]The restriction to standard rate trips is because other trips, such as from JFK Airport to Manhattan, are

including these trips increases our estimates of the impact on customer tip rates. Instead of increasing the trips during this period into our analysis, we could also expand the time period that we analyze. Specifically, we can use all CMT trips starting in February 2011 until August 2012, which faced the same tip suggestions of 20, 25, and 30%. If customers do not select into a specific vendors, then we should find similar estimates focusing on CMT over this period compared to our primary results. Table 3.6 shows that this is the case. We find a very similar pattern of results with an even higher estimated effect of integer tip suggestions on the tip rate.

Lastly, we verify that our results are not driven by our identification strategy in a couple ways. First, we estimate our preferred specifications with placebo treatment effects. To do this, we compare our preferred estimates with a series of placebo treatment effects in which we randomly assign treatment status to each trip for the main sampling period.[13] The placebo estimation follows the baseline specification (3.9). We repeat the process for 1,000 times to obtain a distribution of estimated placebo effects. We define the p-value in this context as the probability that the baseline estimate is obtained purely by chance. The results from placebo tests are shown in Figure C3.4. We find our baseline estimations are significantly different from their respective random benchmarks (p-value< 0.001).

In addition, we utilize variation over time in VTS tip suggestions. From January 22 to January 26 2012, VTS updated the tip suggestions for fares under $\$15$ from a fixed menu of $\$2$, $\$3$, and $\$4$ to 20%, 25%, and 30%. The difference between the tip suggestions following the policy change varied based on how far the fare was from $\$10$. For a fare of exactly $\$10$, two of the suggestions were identical ($\$2$ and $\$3$) with the only change being replacing the rarely used $\$4$ option with a $\$2.50$ option.[14] Outside of this option, this means that if a fare is

charged at a fixed rate leading to non-random spikes in our data at certain tip suggestions. JFK to Manhattan, for example, is charged at a fixed rate of $\$45$, which leads to, absent any surcharges or tolls, a 9 tip suggestion.

[13]We maintain the proportion of trips with integer suggestions when assign our placebo treatments.

[14]Prior to the change, customers used the option to give a $\$4$ tip for a fare ranging from 9 to 11 dollars less than 4 percent of the time.

slightly above or below $\$10$, the two primary tip suggestions are thus nearly identical before and after the menu change, except they are no longer integers. In addition, customers are no longer presented an option at the right-tail of the distribution shown in Figure 3.2a, but instead have a 25% tip rate option.

We use the change in the VTS menu and two alternative empirical strategies to analyze the impact of this menu change on selecting an option from the menu and customer tip rates. Limiting our sample to trips with total fares in the range of $\$9$ to $\$11$ to ensure that the percent and dollar values are similar for two of the options, we estimate a regression discontinuity in time and an event study design. Our regression discontinuity estimates are local linear regressions using a triangular kernel and a bandwidth of 30 days, while our event study estimates control for changes in rate fares in September 2012, tipping percentages for CMT in February 2011, vendor and date fixed effects, and controlling linearly for $x(d, mph)$. When we estimate a regression discontinuity in time, we find a statistically significant decrease in the tip rate of 0.7 percentage points, which is shown in Figure C3.5.[15] This is very similar to the results from our event study in Figure C3.6. We find that, consistent with our primary results, switching away from integer tip suggestions decreased the probability that a passenger chose to tip an option from the menu and it decreased the average tip rate. In terms of magnitude, the estimated impact on tip rate is similar, but the estimated effect on selecting default tip options is much smaller than our primary results, likely due to the replacement of the $4 option with a $2.50 option.

[15] Although we do estimate a decrease in the selection of options from the tipping menu using a regression discontinuity design, we do not present them here since we are not able to confidently classify the tipping menu around the cutoff. It took more than half a week for changes to the tip menu to occur for all vehicles, which is a process we do not observe. For this reason, we do not present these results here.

3.5.3 Who are the Switchers?

Our results highlight that tip options that are integers lead to an increased probability of selecting an option from the tip menu and an overall increase in the tip rate. One question that we have not addressed, however, is what passengers that switch to a menu option do if there are no integer options on the menu. From the theory, we would expect these to be customers with high b_i and/or $T_i \approx 20\%$. This matches what we find in the data, which shows that over half of customers that choose to give a tip not on the menu give either a tip of one dollar, or they round (up or down) from the lowest tip suggestion to the nearest dollar or fifty cents.

To better understand how customers are giving custom tips, we create bins based on the decimals of the lowest tip suggestion. We then calculate what fraction of custom tips within each bin exhibit rounding up or down to the nearest integer or rounding up or down to the nearest 50 cents. We plot the pattern of passenger behavior by decimal bin in Figure 3.8. Regardless of the decimal place of the suggestions, customers tend to round down to the nearest dollar or 50 cents far more than they round up. In fact, only at the point where the decimal places of the tip suggestion are in the range of 91-99 do more customers round up to the nearest integer more than they round down. Empirically, this pattern highlights the mechanism behind our results showing an increase in the tip percentage at the integer values. Customers, when choosing to give a custom tip, tend to choose a lower integer rather than one larger than the lowest option on the tip menu. When presented with an integer tip, however, customers that are rounding for an integer value are less likely to do so since they can select the tip suggestion and avoid the cognitive cost. Since most of these customers tend to round down, this leads to an increase in the tip rate.

3.6 Implications

In the previous sections, we first documented that customers tend to tip integer values. We then presented a model of tipping behavior that can explain this observed pattern in the data. On the one hand, when customers choose to give a tip that is not on the menu the cognitive cost could be lower when choosing an integer. On the other hand, it could be the case that customers actually experience a lump-sum increase in utility when tipping an integer, regardless of if the tip is from the menu or a custom tip. Importantly, both of these mechanisms can explain clustering of tip amounts at integer values, but they differ in one key dimension – the response when presented an integer tip on the menu.

To understand the mechanism underlying the tendency for customers to tip integers, we leverage plausibly exogenous variation in the occurrence of integer tip suggestions. We find consistent evidence that customers do experience a lump-sum increase in utility, which is shown through increases in the tip rate and the probability that a customer chooses a tip from the menu when presented with an integer tip suggestion. In the data, this is driven by customers that, instead of tipping an integer below or above the tip suggestion, are now choosing the option from the menu.

In this section, we consider the implications of customer's preference for integers on the impact that price changes (rate fares) and tip suggestions have on revenue.

3.6.1 Implications for the Impact of Price Changes on Revenue

Let a taxi-drivers revenue for a trip be given by $R = F + t \cdot F$, where F is the total fare and t is the optimal tip rate from the passenger that is maximizing the utility function in equation (3.4). Denote the average tip rate if customers do not face an integer suggestion as $\bar{t}^{non}$. Similarly, define the average tip rate of those that do face an integer suggestion as $\bar{t}^{int}$, where $\bar{t}^{int} - \bar{t}^{non} = \eta$. Assume that the fraction of trips that face an integer suggestion is z

and the fraction that do not is $1 - z$. The average taxi-drivers revenue can then be rewritten:

$$\bar{R} = \bar{F} + \bar{F}[z\bar{t}^{int} + (1-z)\bar{t}^{non}] = \bar{F} + \bar{F}[\bar{t}^{non} + z\eta] \tag{3.11}$$

the revenue from taxi-drivers now depends on how often the tip suggestion presented is an integer and the impact this has on average tipping rates.

To examine the importance of this channel, we will now consider the impact of increasing the rate fare from 0.40 to 0.50, which occurred on September 4, 2012. Assuming that η is unchanged, the change in revenue is:

$$\Delta\bar{R} = \bar{R}_1 - \bar{R}_0 = \Delta F + \underbrace{\bar{F}_1\bar{t}_1^{non} - \bar{F}_0\bar{t}_0^{non}}_{\text{Non-integer Tip Change}} + \underbrace{\eta[\bar{F}_1 z_1 - \bar{F}_0 z_0]}_{\text{Integer Tip Change}} \tag{3.12}$$

where the last term, representing the integer tip change, can simplify further if the likelihood of an integer suggestion after the fare change, z_1, is the same as before the fare change, z_0. In the case of this fare change, however, z_1 increased significantly as we show in Figure 3.9. Average total fares increase from approximately 10.36 ($\bar{F}_0$) to 12 ($\bar{F}_1$) along with a large increase in the probability of an integer tip suggestion from 3.11% (z_0) to 21.27% (z_1). By combining these parameters with our preferred estimate for η from column (3) of Table 3.3, we are able to calculate the last component of the change in revenue. In other words, we can calculate how much the average revenue of a trip increased as a result of the tendency for customers to tip a higher percentage when presented with integer tip suggestions. Plugging in all of the aforementioned parameters into the last component of equation (3.12) we find that this led to an approximately 1.4 cent increase in revenue per trip. With over 170 million taxi trips and 41,000 unique drivers this leads to a transfer of 2.38 million dollars from riders to drivers.

The previous example highlights that when the tip menu is based on percentages, changes in prices (rate fares) can significantly change the nominal values of the options presented to customers. Switching from a rate fare of 40 cents to 50 cents increased the likelihood of an integer tip suggestion by approximately 18 percentage points. Given the differential response of customers to integer tip suggestions, this led to an estimated transfer of 2.38 million dollars from riders to drivers in the year following the policy change. This result emphasizes the key role that the interaction between prices and tip suggestions can have on revenue.

3.6.2 Implications for Default Tip Suggestions

The differential response of passengers based on the tip suggestion decimal places is evident in Figure 3.8. Even when the decimals are in the .91 to .99 range, passengers round down to the 50 cent or dollar at a higher rate than they round up. This leads to the natural question: should drivers round up tip suggestions to the nearest dollar? Intuitively, this seems like it would increase revenue as customers that would have rounded down will now tip the suggested option.

In order to analyze whether "rounding upwards" of tip suggestions would lead to higher tips, on average, we return to our theoretical framework shown in equation (3.4). An intuitive way to think about the impact that this would have on tips is to consider the removal of the lowest option on the tip menu, which is replaced by a new option, denoted as $t_i^{D'}$. Assume that this option is a larger tip rate, $t_i^{D'} \geq t_i^D$ that amounts to rounding up to the nearest integer, $t_i^{D'} F_i \in \mathbb{Z}$.[16]

Unfortunately, our estimates cannot definitively answer how such a menu change would impact average customer tips. We are limited by our identification strategy, which leverages plausibly random tip suggestions that are integer tips at the current menu of suggested tip rates: 20, 25, and 30%. This means that although we find a higher tip rate when 20% tip suggestions are integers, we cannot conclusively say that this will be the case if they are

[16] It is worth noting that $t_i^{D'}$ is not a single value, but changes based on the total fare.

presented with, for example, a 21% suggested tip rate that is an integer. It is possible that increasing the tip rate leads to a large increase in the fraction of customers that choose a custom tip instead of the option from the menu. We can provide some suggestive evidence that customers are unlikely to switch away from the menu option to a custom tip by plotting the selection of the $\$2$ dollar tip suggestion for VTS trips under $\$15$ dollars prior to the menu change in January of 2012. As Figure 3.10 illustrates, customers are unlikely to take-up this suggested tip option when it is a higher percentage of the total fare. However, in the region of the 20% tip it is relatively flat, which suggests that a small change in the tip rate in order to round up the tip suggestion is unlikely to decrease selection of that option significantly.

In summary, our results point towards an alternative tip menu where tip suggestions are rounded up to the nearest dollar. Although we cannot definitively answer whether this would increase revenue, we use the case of 2, 3, and 4 dollar VTS tip suggestions for all trips prior to the January 2012 to show that the down-side from such a policy depends on how large of a shift in the suggested tip rate is. Figure 3.10 shows a small decrease in selecting an integer tip suggestion when it is approximately 20%, but as the tip rate approaches 25 to 30% customers give custom tips at a much higher rate. Importantly, this non-linear relationship between selecting the menu option and tip rates highlight that it may be the case that rounding up tip suggestions can increase revenue, but only when it does not increase the tip rate significantly as this can push passengers away from selecting the tip suggestion.

3.7 Conclusions

Previous research has highlighted that the menu of tip suggestions presented to customers impact the amount that they tip. Using detailed data on millions of trips in New York taxicabs, however, we document that customers tend to tip at integer values even when the menu of tip suggestions are rarely integer values. By extending the model of Donkor (2020), we show that this behavior can be explained by decreasing cognitive costs associated with computing a custom tip that is an integer or from additional utility gains when giving an integer tip. Despite the fact that both model extensions can lead to clustering at integer tips, only in the presence of utility gains at integer values do customers tip differently when presented with an integer suggestion on the tipping menu.

To estimate if customers respond differentially to integer tip suggestions, we leverage plausibly exogenous variation in when integer tip suggestions are given to customers. Across a variety of sample restrictions, specifications, and estimation strategies, we find that when customers are presented with an integer they are more likely to give a tip equal to the suggested value and they tip at a higher rate. That customers respond differentially to the presentation of an integer tip suggestion provides evidence that the tendency to tip integer values is driven, at least in part, by a preference to give integer tips.

Customers' differential responses to integer tip suggestions has natural implications for how prices and tip menus impact revenue. Specifically, our estimates of how customers respond to integer tip suggestions imply that the rate fare change in September 2012 increased average annual revenue for the NYC taxi industry by approximately 2.38 million dollars per year due to the fact that it increased the probability of integer tip suggestions by 18 percentage points. In addition, our finding that customers experience utility gains from integer suggestions could have implications for revenue-maximizing tip menus. Incorporating a model of consumer utility that accounts for the differential response of consumers to integer prices and tip suggestions is an area that deserves the attention of future work.

Figures and Tables

Figure 3.1: Timeline of Fare and Tip Suggestion Changes

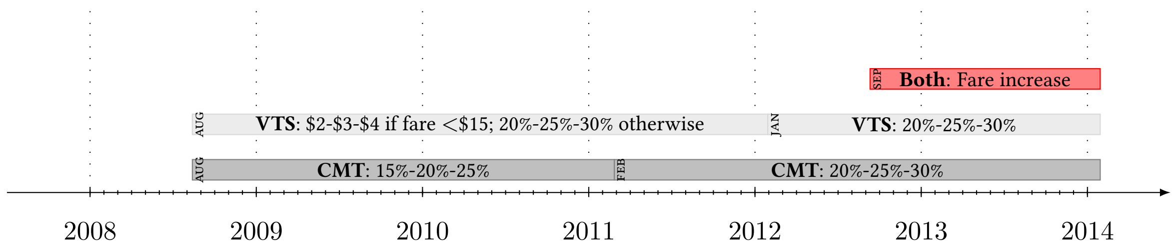

Notes: This figure shows the timing of tip and fare rate changes. NYC taxi cabs were equipped with electronic payment systems around August, 2008. At the beginning, VTS implemented a $-% hybrid tip suggestion menu: the tip prompt is programmed to display 2, 3, and 4 dollars of suggestions if the rate fare (surcharge + fare) is less than $15, and 20, 25, 30 percent if otherwise. On the other hand, the default menu for CMT was 15, 20, and 25 percent. On February 9, 2011, CMT increased their default suggestion to 20, 25, and 30 percent. On the week of January 22, 2012, VTS removed the $ tip suggestions for rate fare below $15 and set their tip suggestion to 20, 25, and 30 percent. On September, 2012, fare rate increased from 40 cents to 50 cents per one fifth of a mile.

Figure 3.2: Distribution of Tip Rate and Tip Amounts: Feb – Aug 2012

(a) Distribution of Tip Rate

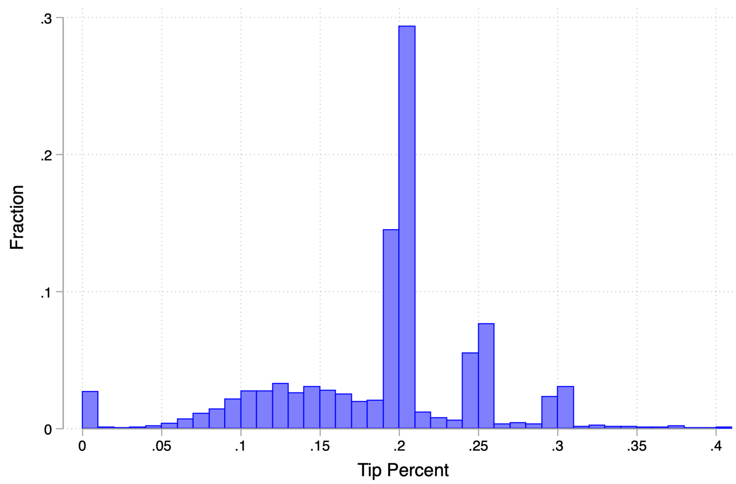

(b) Distribution of Tip Amount

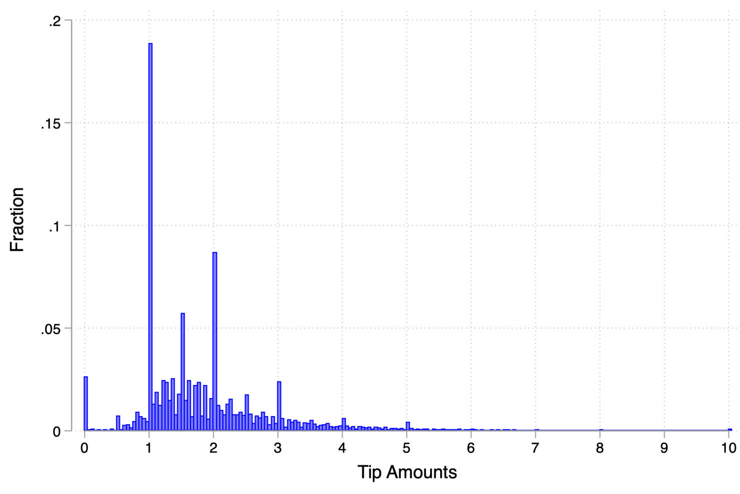

Notes: Panel (a) shows the distribution of tip % for all non-airport trips that were paid by credit card. Panel (b) figure shows the distribution of tip amounts for all non-airport trips that were paid by credit card. Extreme tip rate ($> 99^{th}$ percentile) are excluded from the figure. Tip rate is defined as the tip amount divided by the total rate fare.

Figure 3.3: Distribution of Second Decimal Places

(a) Low Tip Suggestion

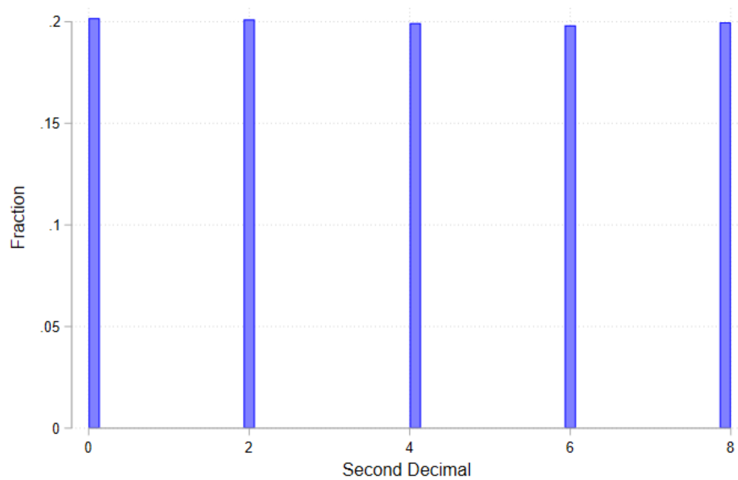

(b) Fraction Tipping the Suggested Amount

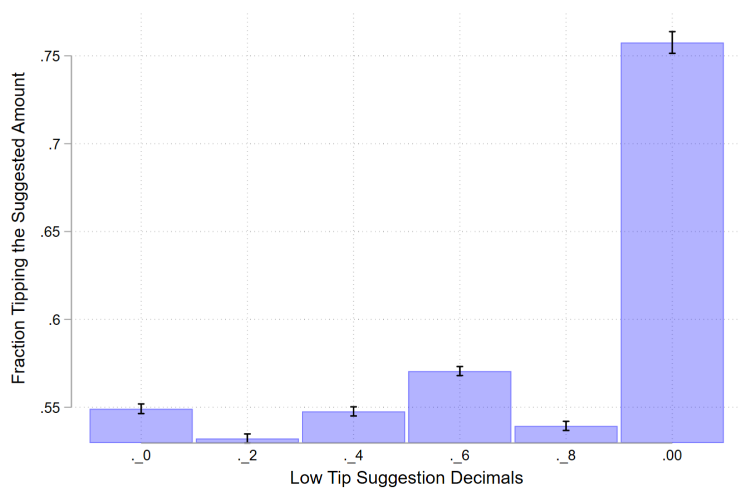

Notes: Panel (a) presents the distribution of second decimal places for tip amounts provided by the default suggestions. The pattern indicates that 0, 2, 4, 6, and 8 are approximately equally likely to appear in the tip suggestions. Panel (b) shows the fraction of customers that tip the suggested amount for each second decimal place of the low tip suggestion. The pattern indicates that the tip rate is significantly higher when the low tip suggestion ends with ".00", i.e., is an integer, compared to all other tip suggestions.

Figure 3.4: Distribution of Selected Suggested Tip Amounts: Feb – Aug 2012

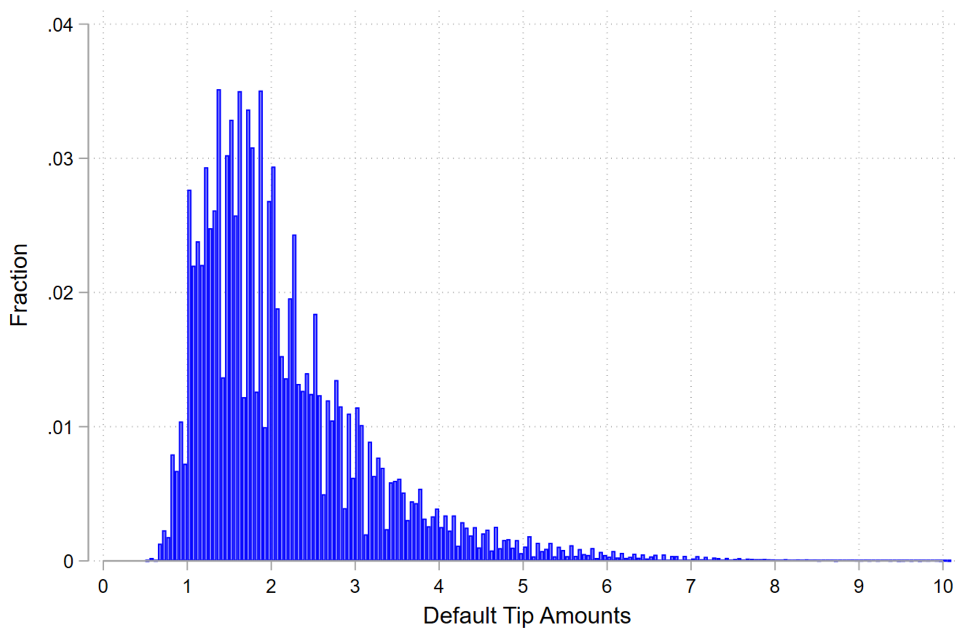

Notes: This figure shows the distribution of selected, by the passenger, default tip amounts for all non-airport trips that were paid by credit card. Default tips includes all non-zero tips that are equal to one of the tip suggestions. Extreme default tip amounts ($> 99^{th}$ percentile) are excluded from the figure. All tip amounts are in nominal dollar value.

Figure 3.5: Individual's Utility Maximization under Baseline and Extended Models

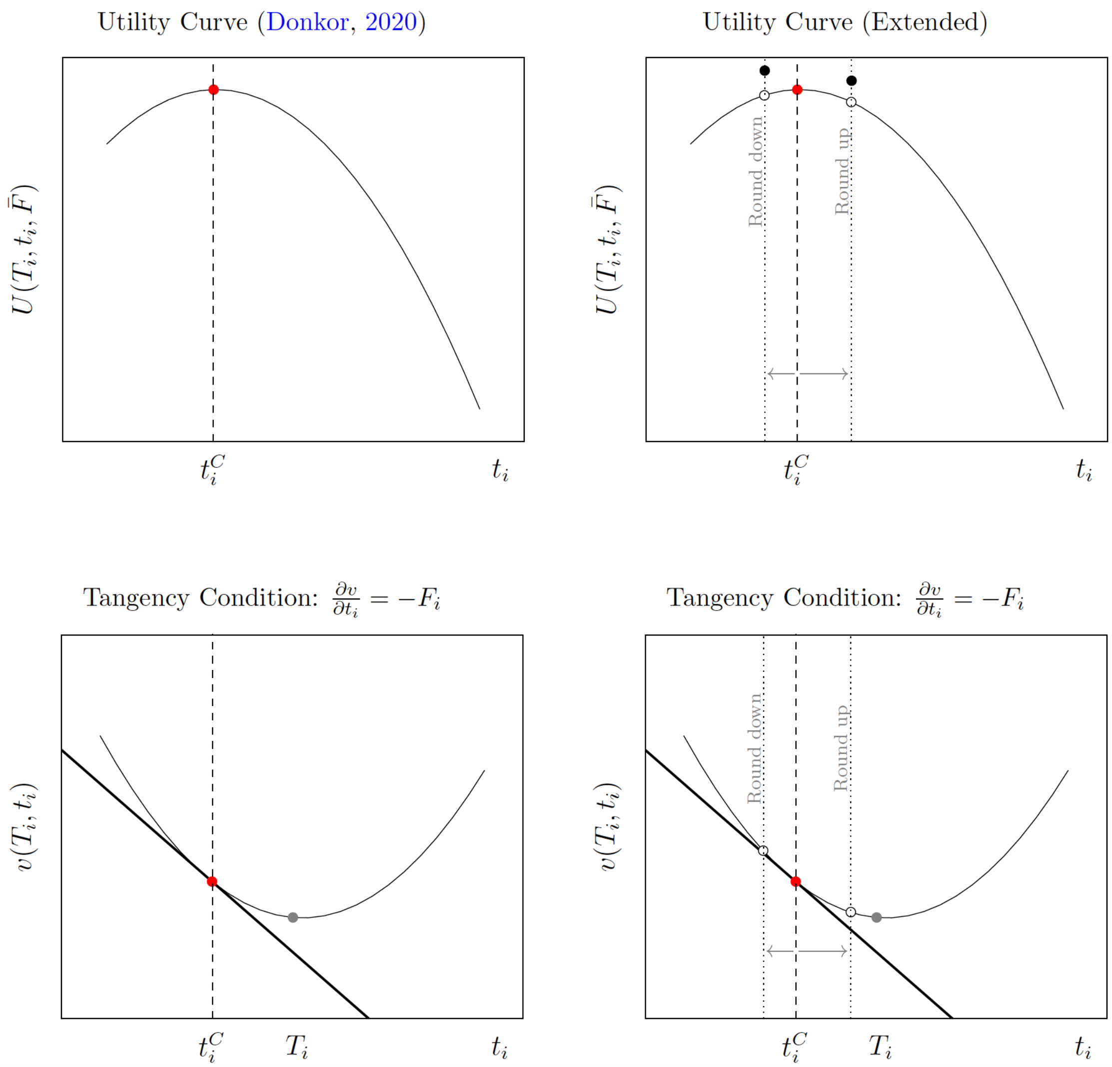

Notes: The left and right panels compare and contrast a typical passenger's utility maximization decision under the baseline model (Donkor, 2020) and the extended model. The top panel presents utility curves. The bottom panel presents the corresponding tangency condition. In particular, the convex function represents norm deviation cost ($v(T_i, t_i)$) and the downward sloping line represents the cost of increasing tipping rate. Under the baseline model, the passenger solves the utility maximization by choosing $t_i = t_i^C$ at the tangent point. Under the extended model, given the utility function is not continuous at integer tip amounts, the tangency condition might not lead to a global maximum. In this specific example, it is optimal for the passenger to *Round* t_i^C *down* to the nearest integer.

Figure 3.6: The Decision Process for Agents under the Extended Model

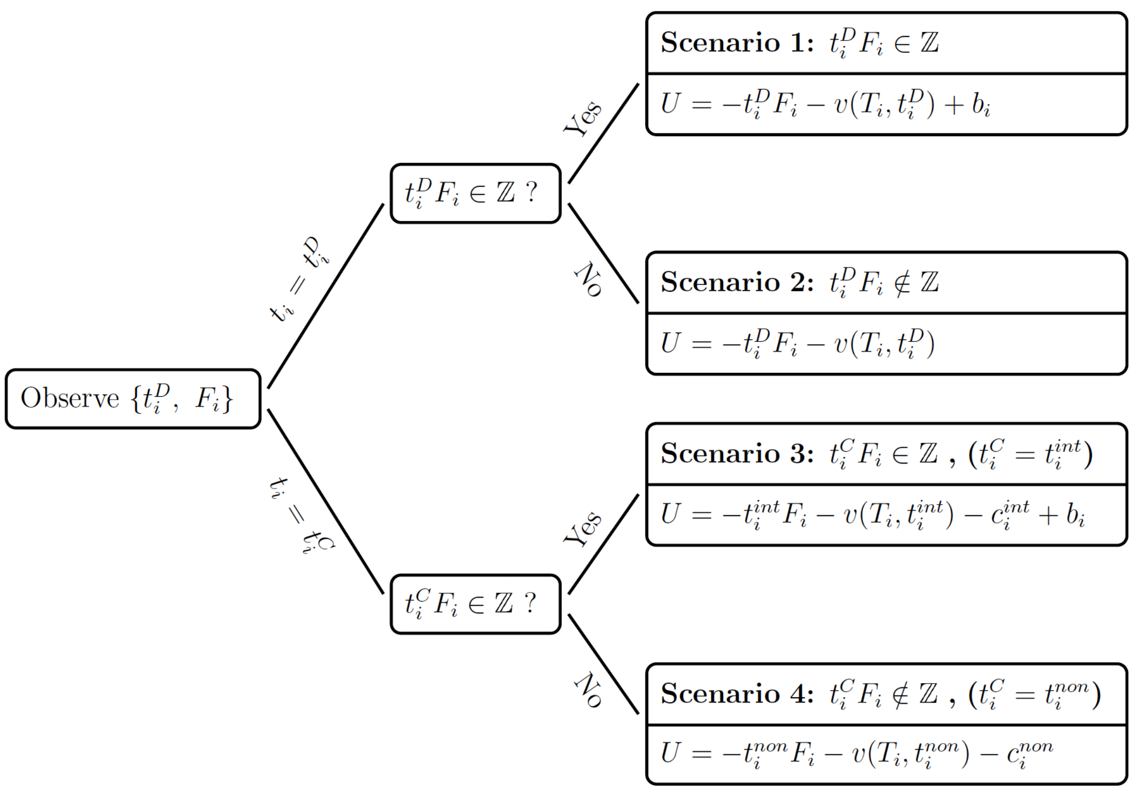

Notes: This figure presents the decision process that is described by the extended model (equation. 3.4). After observing t_i^D and F_i, a customer decide whether or not to the choose custom tipping option. If she chooses to enter a custom tip, she will incur a cognitive cost of either c_i^{non} or c_i^{int} depending on whether the chosen tip amount is an integer. We assume the cognitive cost is lower if she chooses to tip an integer amount $c_i^{int} < c_i^{non}$. In addition, she experiences utility gains (b_i) from choosing an integer tip, i.e, $t_i \in \mathbb{Z}, t_i \in \{t_i^C, t_i^D\}$.

Figure 3.7: Low Tip Suggestion by $x(d, mph)$, Surcharges and Vendor

(a) VTS

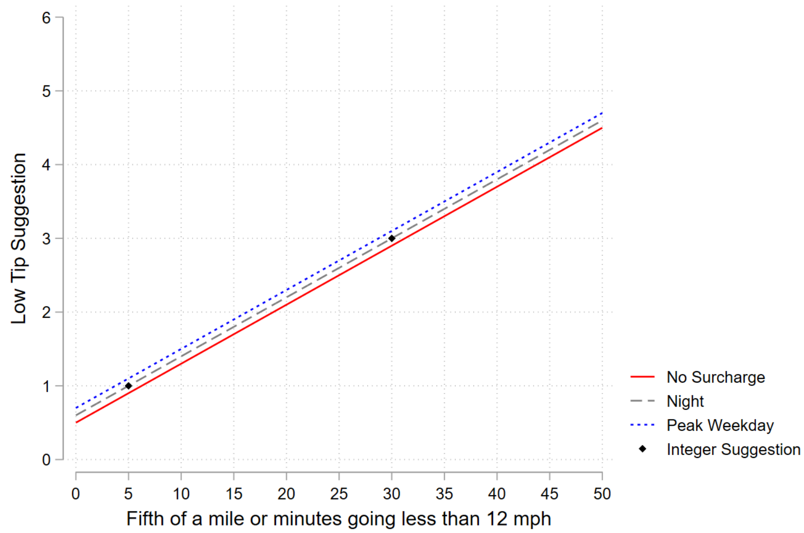

(b) CMT

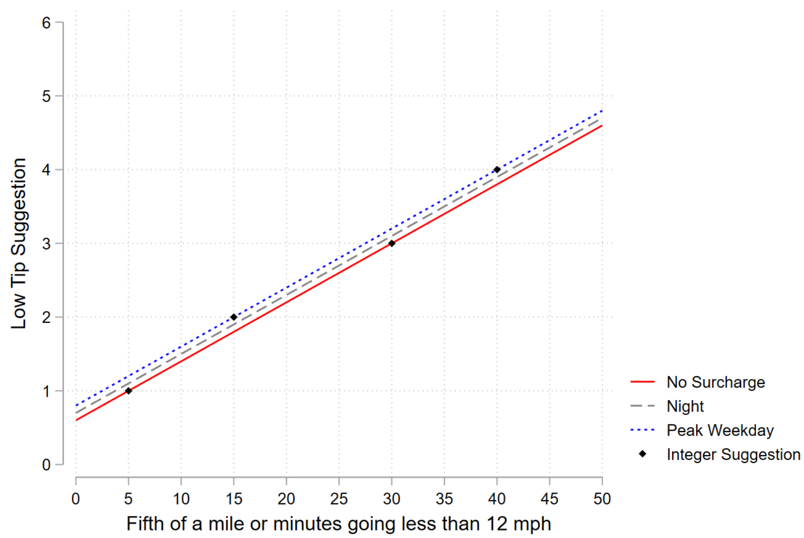

Notes: Figures show the mapping from $x(d, mph)$ and surcharges to Low tip suggestions for each vendors from Feb – Aug 2012. For VTS, integer suggestions only appears in weekdays during peak hours. For CMT, integer suggestions could appear when there's no surcharge, or during the peak weekdays (surcharge = \$1.00 from 4pm-8pm on weekdays).

Figure 3.8: Fraction of Custom Tips that Round to Nearby Values

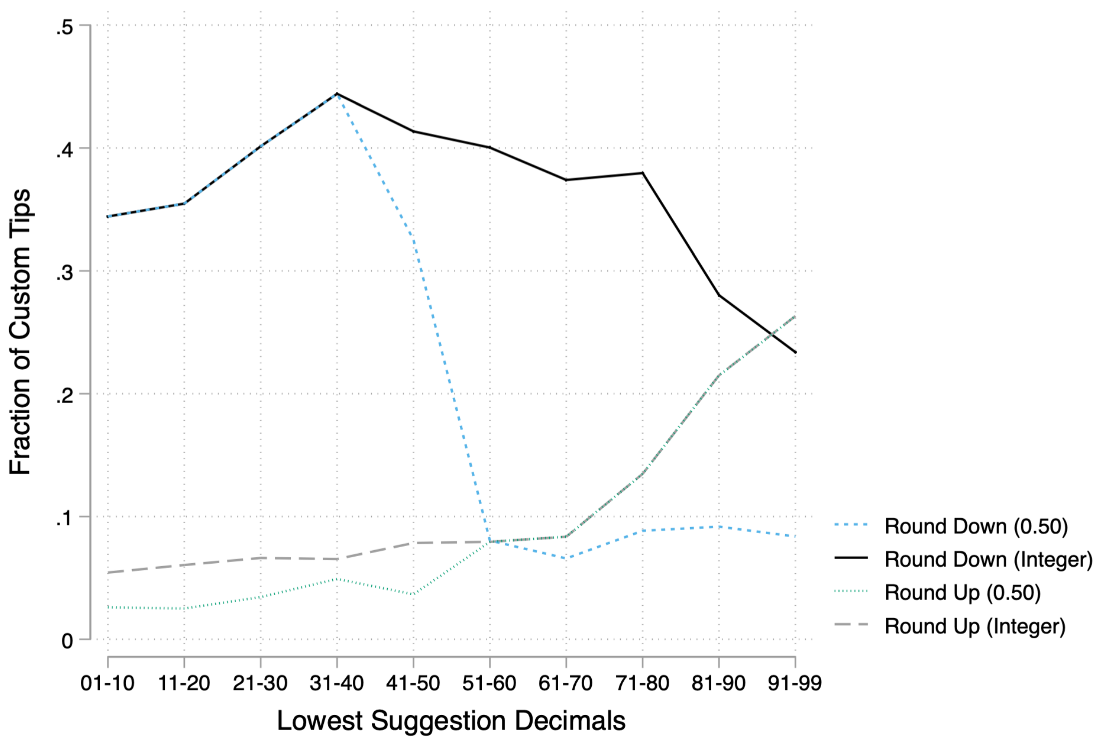

Notes: This figure shows rounding heuristics used by customers when they opt-in for custom tips. The horizontal axis presents decimal places for the lowest tip suggestion amount. The decimal places are divided into 10 equally spaced bins. The vertical axis represents the fraction of custom tips that either rounds the lowest default suggestion up or down to the nearest dollar or 50 cents. Overall, we observe that customers are more likely to round down regardless of the decimal places of the suggestions.

Figure 3.9: The Probability of Integer Tip Suggestions by Month-Year

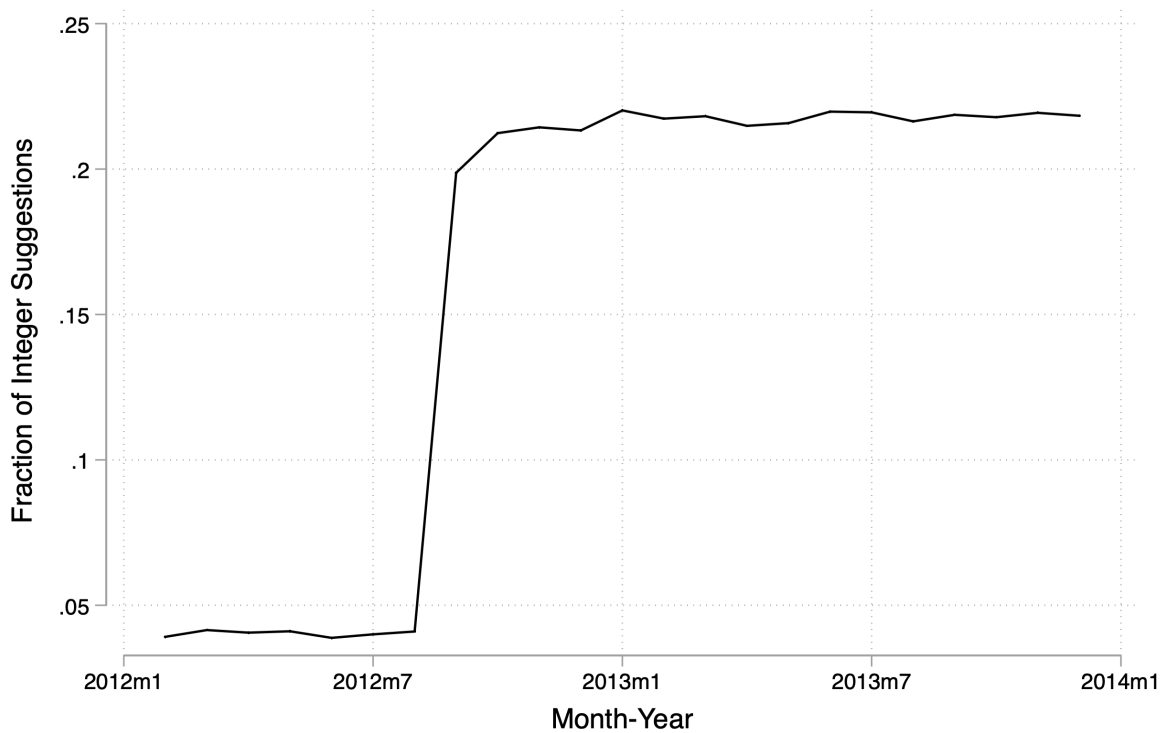

Notes: This figure shows the average probability of having an integer tip suggestion overtime from 2012.Feb onward. Around 2012.Sept, per-unit fare rate increased from \$0.40 to \$0.50. Given a tip suggestion menu: 20, 25, and 30 percent, this fare increase has significantly raised the probability of having an integer tip suggestion.

Figure 3.10: Fraction of Customers Choose a $2 Default, by Fare+Surcharge: VTS (pre-2012)

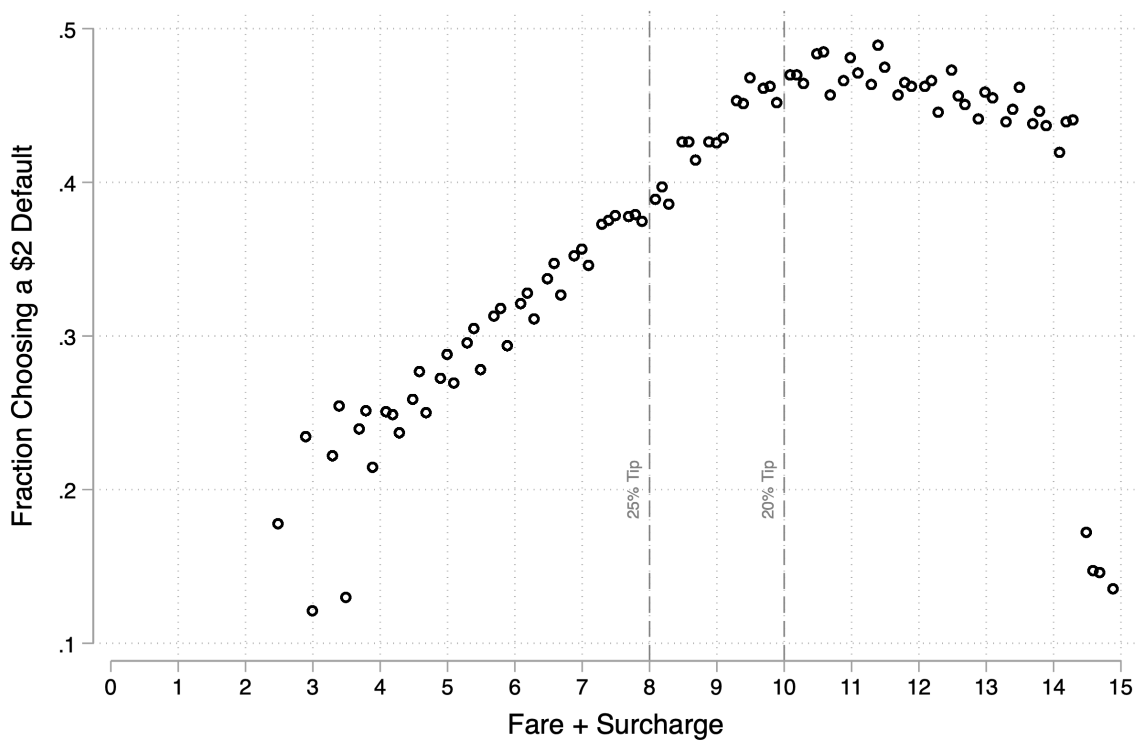

Notes: This figure shows the fraction of VTS customers choose a $2 default before 2012.Jan. we divide the *rate fare* (Fare + Surcharge) into 150 equally sized bins, then we compute the average $2 tip take-up rate for each bin. The 20% and 25% marks represents *rate fare* values that are equivalent to a $2 tip.

Table 3.1: Summary Statistics by Trip (Ride): Feb–Aug 2012

	(1) VTS	(2) CMT	(3) Difference
Fare Amount	9.48	9.48	-0.00
	(5.00)	(4.87)	(0.01)
Tip Amount	1.86	1.93	0.07***
	(1.35)	(1.27)	(0.00)
Tip Rate	0.20	0.19	-0.00***
	(0.12)	(0.09)	(0.00)
Trip Length (in minutes)	12.08	12.13	0.05**
	(7.27)	(7.69)	(0.02)
Trip Distance (in miles)	2.54	2.54	-0.01
	(2.10)	(2.12)	(0.01)
Zero Tip	0.03	0.02	-0.02***
	(0.18)	(0.13)	(0.00)
Pr(Select 'low' default)	0.42	0.37	-0.05***
	(0.49)	(0.48)	(0.00)
Pr(Select 'middle' default)	0.13	0.11	-0.02***
	(0.33)	(0.31)	(0.00)
Pr(Select 'high' default)	0.05	0.04	-0.02***
	(0.23)	(0.19)	(0.00)
Observations	358,416	351,643	710,059

Notes: This table presents the summary statistics for the random sample of 2,000 taxi drivers during the time period of our main study: February to August 2012. During this period of time, the % tip suggestions are identical to CMT and VTS in all trips: 20, 25 and 30 percent. Tip rate is defined as the tip amount divided by the total fare excluding the tipped amount. *rate fare* (the value F_i used for tip computation) is defined differently for CMT and VTS. For CMT: Rate Fare = fare + surcharge + mta tax + tolls; For VTS: Rate Fare = fare + surcharge. Standard deviations are in parenthesis.

Table 3.2: Impact of Integer Tip Suggestions on Selecting a Default Suggestions

	(1)	(2)	(3)	(4)	(5)	(6)
Any Default Integer	0.19506***	0.21361***	0.21462***			
	[0.00590]	[0.00739]	[0.00746]			
Low Option Integer				0.23487***	0.23239***	0.23388***
				[0.00569]	[0.00642]	[0.00642]
Mid or High Option Integer				-0.00530	0.05978***	0.05806***
				[0.00670]	[0.01357]	[0.01357]
Outcome Mean	.554	.554	.554	.554	.554	.554
x(d,mph) Control	No	Yes	Yes	No	Yes	Yes
Date FE	No	Yes	No	No	Yes	No
Pick-up Hour FE	No	Yes	No	No	Yes	No
Driver FE	No	Yes	Yes	No	Yes	Yes
Drop-off Block FE	No	No	Yes	No	No	Yes
Pickup Block FE	No	No	Yes	No	No	Yes
Date by Hour FE	No	No	Yes	No	No	Yes
Clusters (Driver)	1,658	1,656	1,655	1,658	1,656	1,655
Clusters (Date)	213	213	213	213	213	213

Notes: This table shows the estimated impact of having an integer tip suggestion option on the probability that a custom tips a suggested amount. The results shown here are for all standard rate fare trips from February to August of 2012, which did not involve a pickup or drop-off at an airport. All estimates are from a linear probability model with specifications varying by column. The specifications of the first three columns are repeated in the next 3 columns. The first three columns show the effect if any of the options on the menu are an integer, while the last three columns show the effect if the lowest option is an integer or the other two options are integers. The first (and fourth) column has no controls or fixed effects. The second (and fifth) column includes date, driver, and hour fixed effects. The third (and sixth) column includes date by hour, driver, pickup census block, drop-off census block, and nearest integer to the low suggestion fixed effects. Apart from the first and the fourth column, we control for the linear relationship between $x(d, mph)$ and the outcome of interest. Standard errors in brackets are two-way clustered at the driver level and the pick-up date level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.3: Impact of Integer Tip Suggestions on Tip Rates

	(1)	(2)	(3)	(4)	(5)	(6)
Any Default Integer	0.00796***	0.00716***	0.00631***			
	[0.00080]	[0.00067]	[0.00064]			
Low Option Integer				0.01334***	0.01013***	0.00826***
				[0.00074]	[0.00073]	[0.00073]
Mid or High Option Integer				-0.01215***	-0.00568***	-0.00300**
				[0.00121]	[0.00124]	[0.00118]
Outcome Mean	.193	.193	.193	.193	.193	.193
x(d,mph) Control	No	Yes	Yes	No	Yes	Yes
Date FE	No	Yes	No	No	Yes	No
Pick-up Hour FE	No	Yes	No	No	Yes	No
Driver FE	No	Yes	Yes	No	Yes	Yes
Drop-off Block FE	No	No	Yes	No	No	Yes
Pickup Block FE	No	No	Yes	No	No	Yes
Date by Hour FE	No	No	Yes	No	No	Yes
Clusters (Driver)	1,658	1,656	1,655	1,658	1,656	1,655
Clusters (Date)	213	213	213	213	213	213

Notes: This table shows the estimated impact of having an integer tip suggestion option on the tipping rate, defined as the tip amount divided by the rate fare used when determining the tip suggestion. The results shown here are for all standard rate fare trips from February to August of 2012, which did not involve a pickup or drop-off at an airport. All estimates are from an ordinary least squares regression with specifications varying by column. The specifications of the first three columns are repeated in the next 3 columns. The first three columns show the effect if any of the options on the menu are an integer, while the last three columns show the effect if the lowest option is an integer or the other two options are integers. The first (and fourth) column has no controls or fixed effects. The second (and fifth) column includes date, driver, and hour fixed effects. The third (and sixth) column includes date by hour, driver, pickup census block, drop-off census block, and nearest integer to the low suggestion fixed effects. Apart from the first and the fourth column, we control for the linear relationship between $x(d, mph)$ and the outcome of interest. Standard errors in brackets are two-way clustered at the driver level and the pick-up date level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.4: Impact of Integer Tip Suggestions on Tip Rates (Alternative Definition)

	(1)	(2)	(3)	(4)	(5)	(6)
Any Default Integer	0.00935***	0.00624***	0.00556***			
	[0.00075]	[0.00066]	[0.00063]			
Low Option Integer				0.01258***	0.00901***	0.00747***
				[0.00077]	[0.00072]	[0.00073]
Mid or High Option Integer				-0.00559***	-0.00556***	-0.00330***
				[0.00122]	[0.00124]	[0.00118]
Outcome Mean	.187	.187	.187	.187	.187	.187
x(d,mph) Control	No	Yes	Yes	No	Yes	Yes
Date FE	No	Yes	No	No	Yes	No
Pick-up Hour FE	No	Yes	No	No	Yes	No
Driver FE	No	Yes	Yes	No	Yes	Yes
Drop-off Block FE	No	No	Yes	No	No	Yes
Pickup Block FE	No	No	Yes	No	No	Yes
Date by Hour FE	No	No	Yes	No	No	Yes
Clusters (Driver)	1,658	1,656	1,655	1,658	1,656	1,655
Clusters (Date)	213	213	213	213	213	213

Notes: This table shows the estimated impact of having an integer tip suggestion option on the tipping rate, defined as the tip amount divided by the total fare excluding the tipped amount. The results shown here are for all standard rate fare trips from February to August of 2012, which did not involve a pickup or drop-off at an airport. All estimates are from an ordinary least squares regression with specifications varying by column. The specifications of the first three columns are repeated in the next 3 columns. The first three columns show the effect if any of the options on the menu are an integer, while the last three columns show the effect if the lowest option is an integer or the other two options are integers. The first (and fourth) column has no controls or fixed effects. The second (and fifth) column includes date, driver, and hour fixed effects. The third (and sixth) column includes date by hour, driver, pickup census block, drop-off census block, and nearest integer to the low suggestion fixed effects. Apart from the first and the fourth column, we control for the linear relationship between $x(d, mph)$ and the outcome of interest. Standard errors in brackets are two-way clustered at the driver level and the pick-up date level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.5: Impact of Integer Tip Suggestions on Tip Rates: Including non-Standard Rate Trips

	(1)	(2)	(3)	(4)	(5)	(6)
Any Default Integer	0.00506***	0.01106***	0.00900***			
	[0.00062]	[0.00065]	[0.00066]			
Low Option Integer				0.00845***	0.01420***	0.01146***
				[0.00064]	[0.00061]	[0.00073]
Mid or High Option Integer				-0.01080***	-0.00640***	-0.00445***
				[0.00114]	[0.00114]	[0.00111]
Outcome Mean	.193	.193	.193	.193	.193	.193
x(d,mph) Control	No	Yes	Yes	No	Yes	Yes
Date FE	No	Yes	No	No	Yes	No
Pick-up Hour FE	No	Yes	No	No	Yes	No
Driver FE	No	Yes	Yes	No	Yes	Yes
Drop-off Block FE	No	No	Yes	No	No	Yes
Pickup Block FE	No	No	Yes	No	No	Yes
Date by Hour FE	No	No	Yes	No	No	Yes
Clusters (Driver)	1,663	1,662	1,661	1,663	1,662	1,661
Clusters (Date)	213	213	213	213	213	213

Notes: This table shows the estimated impact of having an integer tip suggestion option on the tipping rate, defined as the tip amount divided by the rate fare used when determining the tip suggestion. The results shown here are for all trips from February to August of 2012, even if the drop-off or pickup location is an airport. All estimates are from an ordinary least squares regression with specifications varying by column. The specifications of the first three columns are repeated in the next 3 columns. The first three columns show the effect if any of the options on the menu are an integer, while the last three columns show the effect if the lowest option is an integer or the other two options are integers. The first (and fourth) column has no controls or fixed effects. The second (and fifth) column includes date, driver, and hour fixed effects. The third (and sixth) column includes date by hour, driver, pickup census block, drop-off census block, and nearest integer to the low suggestion fixed effects. Apart from the first and the fourth column, we control for the linear relationship between $x(d, mph)$ and the outcome of interest. Standard errors in brackets are two-way clustered at the driver level and the pick-up date level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.6: Impact of Integer Tip Suggestions on Tip Rates: CMT (2011 Feb – 2012 Aug)

	(1)	(2)	(3)	(4)	(5)	(6)
Any Default Integer	0.00804***	0.00642***	0.00591***			
	[0.00068]	[0.00056]	[0.00053]			
Low Option Integer				0.01595***	0.01198***	0.00988***
				[0.00064]	[0.00058]	[0.00059]
Mid or High Option Integer				-0.01100***	-0.00792***	-0.00522***
				[0.00077]	[0.00079]	[0.00079]
Outcome Mean	.189	.189	.189	.189	.189	.189
x(d,mph) Control	No	Yes	Yes	No	Yes	Yes
Date FE	No	Yes	No	No	Yes	No
Pick-up Hour FE	No	Yes	No	No	Yes	No
Driver FE	No	Yes	Yes	No	Yes	Yes
Drop-off Block FE	No	No	Yes	No	No	Yes
Pickup Block FE	No	No	Yes	No	No	Yes
Date by Hour FE	No	No	Yes	No	No	Yes
Clusters (Driver)	2,317	2,311	2,310	2,317	2,311	2,310
Clusters (Date)	573	573	573	573	573	573

Notes: This table shows the estimated impact of having an integer tip suggestion option on the tipping rate, defined as the tip amount divided by the rate fare used when determining the tip suggestion. The results shown here are for all standard rate fare trips from February 2011 to August of 2012 that use CMT equipment and did not involve a pickup or drop-off at an airport. All estimates are from an ordinary least squares regression with specifications varying by column. The specifications of the first three columns are repeated in the next 3 columns. The first three columns show the effect if any of the options on the menu are an integer, while the last three columns show the effect if the lowest option is an integer or the other two options are integers. The first (and fourth) column has no controls or fixed effects. The second (and fifth) column includes date, driver, and hour fixed effects. The third (and sixth) column includes date by hour, driver, pickup census block, drop-off census block, and nearest integer to the low suggestion fixed effects. Apart from the first and the fourth column, we control for the linear relationship between $x(d, mph)$ and the outcome of interest. Standard errors in brackets are two-way clustered at the driver level and the pick-up date level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Appendix A

Appendix for “The Effect of Unemployment Benefit Pay Frequency on UI Claimants’ Job Search Behaviors”

A.1 Additional Figures and Tables

Figure A1.1: Screenshot of Archived UI website

(a) New Jersey UI website: 2013, Biweekly Filing

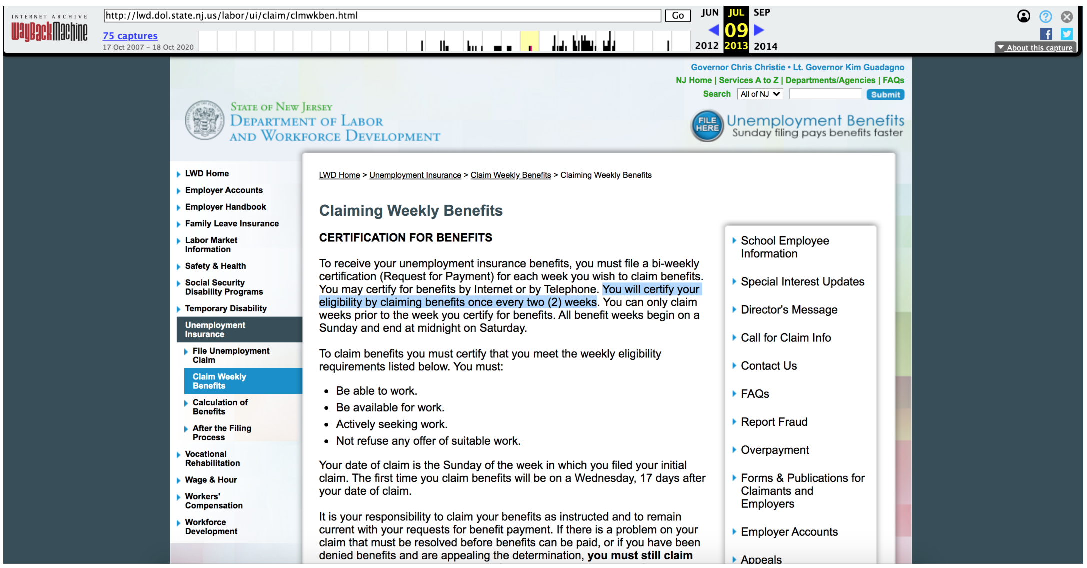

(b) New Jersey UI website: 2014, Weekly Filing

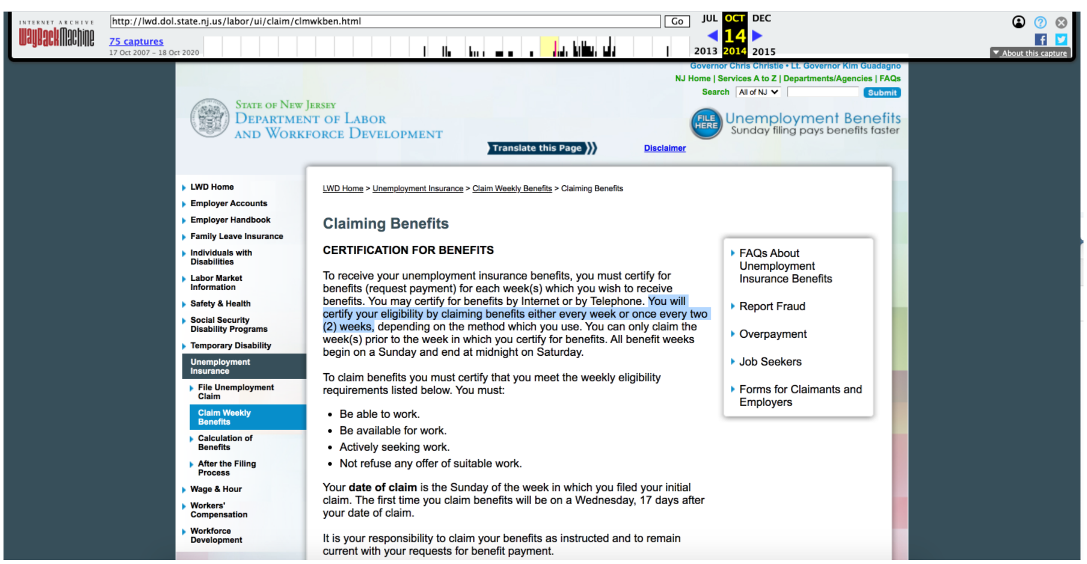

Notes: The pictures show screenshot taken from *Archive.org* for New Jersey's historical UI websites. From change in filing instructions, I can back-out the changing date for UI filing frequency policy.

Table A1.1: State policies on pay frequency, 1985-2016

States	Start w/	Switch (year)	States	Start w/	Switch (year)
Alabama	weekly	-	Missouri	weekly	-
Arizona	weekly	-	Nebraska	weekly	-
Arkansas	weekly	-	Nevada*	weekly	*
California	biweekly	-	New Hampshire	biweekly	weekly (2003)
Colorado	biweekly	-	New Jersey	biweekly	weekly (2013)
Connecticut	weekly	-	New Mexico	biweekly	weekly (1999)
Delaware	weekly	-	New York	biweekly	weekly (1993)
District of Columbia	biweekly	weekly (2007)	North Carolina	biweekly	weekly (1997)
Florida	biweekly	-	Ohio	weekly	either (2003)
Georgia	weekly	-	Oklahoma	biweekly	either (2004)
Hawaii	weekly	either (2008)	Oregon	biweekly	weekly (1992)
Illinois	weekly	-	Pennsylvania	biweekly	-
Indiana	weekly	-	Rhode Island	biweekly	weekly (1996)
Kansas	weekly	-	South Carolina	weekly	-
Kentucky	biweekly	-	Tennessee	weekly	-
Louisiana	weekly	-	Texas	biweekly	-
Maryland	biweekly	weekly (2013)	Utah	biweekly	weekly (1994)
Massachusetts	biweekly	weekly (2003)	Virginia	biweekly	weekly (1996)
Michigan	biweekly	-	Washington	biweekly	weekly (1996)
Minnesota	biweekly	weekly (2008)	West Virginia	biweekly	weekly (2014)
Mississippi	weekly	-	Wisconsin	weekly	-

Notes: Data collected from the BAM survey and the archived state government’s website via *archive.org*. *Nevada switched from biweekly to weekly in 1994, and switched back to biweekly after 1999.

Table A1.2: The Timing of No-Extra Check Month From 1985-2007

Year	Month 1	Month 2
1985	Feb	Jun
1986	Feb	Nov
1987	Feb	-
1988	-	-
1989	Feb	Apr
1990	Feb	Sep
1991	Feb	Jun
1992	Feb	-
1993	Feb	-
1994	Feb	-
1995	Feb	-
1996	Jun	-
1997	Feb	Nov
1998	Feb	-
1999	Feb	-
2000	Apr	-
2001	Feb	Sep
2002	Feb	Jun
2003	Feb	Nov
2004	Feb	-
2005	Feb	-
2006	Feb	Apr
2007	Feb	Sep

Notes: The table displays the months which it is never possible to have extra benefit checks. From 1985-2007: 74.33% of such months are in February, 10.9% in June, 5.67% in September, 5.43% in November and 3.67% in April. In addition, such months vary from year-to-year.

Table A1.3: State policies on the adoption of initial claiming technology

States	Tel. (year)	Int. (year)	States	Tel. (year)	Int. (year)
Alabama	2002	2007	Missouri	1997	2003
Arizona	2001	2006	Nebraska	2001	2007
Arkansas	2004	2005	Nevada	1999	2003
California	1994	2003	New Hampshire	2000	2003
Colorado	1991	2003	New Jersey	1999	2002
Connecticut	2002	2006	New Mexico	-	2003
Delaware	2002	2010	New York	1999	2003
District of Columbia	2001	2004	North Carolina	2001	2003
Florida	1996	2003	Ohio	1997	2005
Georgia	1990	2011	Oklahoma	2000	2003
Hawaii	2000	2010	Oregon	1994	2004
Illinois	1998	2003	Pennsylvania	1998	2003
Indiana	-	2005	Rhode Island	1997	2003
Kansas	1999	2003	South Carolina	1999	-
Kentucky	2004	2004	Tennessee	2002	2004
Louisiana	2003	2006	Texas	1998	2003
Maryland	1996	2003	Utah	1998	2006
Massachusetts	1996	2013	Virginia	2002	2004
Michigan	2003	2004	Washington	1999	2003
Minnesota	2000	2003	West Virginia	2004	2010
Mississippi	2005	2010	Wisconsin	1994	2003

Notes: Data collected from the BAM survey. Tel. stands for Telephone filing and Int. stands for Internet filing.

Table A1.4: State policies on the adoption of continued claiming technology

States	Tel. (year)	Int. (year)	States	Tel. (year)	Int. (year)
Alabama	1998	2009	Missouri	1995	2003
Arizona	1996	2007	Nebraska	1996	2006
Arkansas	2000	2005	Nevada	1999	2003
California	2005	2012	New Hampshire	2000	2004
Colorado	1998	2006	New Jersey	1997	2009
Connecticut	1996	2006	New Mexico	1999	2006
Delaware	2005	2013	New York	1994	2003
District of Columbia	2004	2005	North Carolina	1997	2003
Florida	1995	2004	Ohio	2003	2005
Georgia	1994	2003	Oklahoma	1996	2008
Hawaii	2001	2011	Oregon	1994	2003
Illinois	1994	2003	Pennsylvania	1996	2003
Indiana	-	2006	Rhode Island	1997	2010
Kansas	1997	2003	South Carolina	1995	-
Kentucky	1997	2004	Tennessee	1994	2004
Louisiana	1996	2003	Texas	1996	2007
Maryland	1996	2003	Utah	1995	2004
Massachusetts	2003	2004	Virginia	1998	2004
Michigan	1996	2010	Washington	1996	2007
Minnesota	1997	2005	West Virginia	2001	2007
Mississippi	2004	2010	Wisconsin	1994	2004

Notes: Data collected from the BAM survey. Tel. stands for Telephone filing and Int. stands for Internet filing. Indiana is a special case where it changed the filing technology from mail filing to internet filing directly. I have therefore excluded Indiana from the analysis.

A.2 The Effect of Severance Pay under Different UI Pay Frequencies

Severance pay is a lump-sum cash transfer from employers to their employees at the time of layoff. Since such payment does not count against the size of UI benefits, we can interpret the impact from receiving severance pay as a form of liquidity effect. However, as noted in Chetty (2008), severance pay status is not determined at random - its eligibility highly relates to one's job tenure. In fact, most firms have minimum job tenure threshold and the size of severance pay usually increase in job tenure (as a step-function).

To obtain a reasonable estimation of the liquidity effect using severance pay status, it is important to control for UI claimant's job tenure in my sample. Unfortunately, due to the short-panel nature of the SIPP data, over 72% of my sample are left-censored (no information on the exact job starting date). To overcome this limitation, I predict job tenure for each individual using OLS on the Mathematica sample from Chetty (2008). The predictors consists of linear form of pre-unemployment wage, age, education and martial status. I find the results using predicted tenure and the actual tenure are similar under the Mathematica sample. This indicates the predicted tenure is a plausible proxy for the actual tenure in my SIPP 1996-2007 sample.

In the following, I estimate the Cox proportional hazard model regression from Chetty (2008) using sample from SIPP 1996-2007:

$$\log h_{it} = \alpha_t + \beta_1 sev_i + \beta_2 (t \times sev_i) + \mathbf{X_{it}} \tag{A.1}$$

where h_{it} is the hazard rate of exiting unemployment for individual i at time t. α_t is the flexible non-parametric baseline hazard rate at the given week t conditional on surviving. sev_i is an indicator for receiving severance payment at the time of unemployment. $(t \times sev_t)$ allows the effect of severance pay to interact with duration. $\mathbf{X_{it}}$ includes state fixed effects, year fixed effects, industry and occupation fixed effects, state unemployment rate and individual predicted weekly benefit amount and individual demographic controls identical to the previous parts. In addition, since I've argued that tenure plays an important role in determining severance pay eligibility, I control for this by using a 10 piece (predicted) job tenure spline. Assuming that severance pay status is "random" conditional on job tenure, β_1 will identify the liquidity effect on reemployment hazard at the beginning of the unemployment spell.

A.2.1 Data and Sample

Table A2.1: Descriptive Statistics (Mean) by Pay Frequency and Severance Pay Status, SIPP 1996-2007

	WEEKLY		BIWEEKLY	
	Severance	No Severance	Severance	No Severance
Unemployment Duration (weeks)	21.99	18.69	24.68	20.33
Age	42.26	39.22	44.18	39.24
Years of education	13.73	12.44	13.70	12.42
$\mathbb{1}$(Married)	0.68	0.57	0.74	0.60
Simulated replacement rate	0.50	0.50	0.50	0.49
State unemployment rate	5.17	5.27	5.30	5.34
Predicted Tenure (weeks)	317.44	289.74	337.93	290.46
Pre-unemployment Annual wage ($)	31,283.78	20,372.10	41,271.39	22,906.31
Liquid Wealth ($)	62,888.79	36,130.25	74,250.97	37,458.53
Unsecured debt ($)	3,954.91	6,211.85	12,541.08	5,301.08
Home Equity ($)	55,725.25	30,362.11	56,873.46	37,487.31
# Spells	78	1,258	104	1,372

Notes: The data presented are individual level unemployment spells from 1996-2007 SIPP data. Tenure is predicted using the Mathematica sample from Chetty (2008). All dollar values are converted to 1990 values. The sample restricted to prime age male UI claimants only. The pay frequency policy information are collected from archived state websites via *archive.org*.

I use unemployment spell data from SIPP 1996-2007. Starting from the 1996 Panel, SIPP included questions on severance pay recipient status and the amount of severance pay. To facilitate interpretation, I apply the identical sampling restriction as the parts. In particular, I focus on prime age male unemployed worker who (a) report searching for a job, (b) are not on temporary layoff, (c) have at least 3 months of work history in the survey (to compute pre-unemployment earnings), (d) took up UI benefits within the first month of unemployment. Unemployment duration is censored at 50 weeks and all monetary values are in 1990 dollars.

Table A2.1 provides a descriptive summary for my sample. In this sample, less than 10% of the UI claimants received severance pay - the number is slightly lower compare to the Mathematica sample. Among both “weekly" and “biweekly" states, UI claimants with severance pay at the time of unemployment look significantly different from those without severance pay. In particular, severance pay recipients tend to be older, more educated, more likely to be married, have longer predicted job tenure and higher net wealth.

A.2.2 Empirical Result

Figure A2.1 shows the effect of receiving severance pay on unemployment duration for UI claimants under each pay frequency controlling for job tenure. From visual inspection, I find results from both panels resemble the previous finding from Chetty (2008) – receiving severance pay leads to significantly lower reemployment hazard.

Table A2.2 displays a series of regression results. The coefficient of interest is the percent change in hazard rate with respect to severance pay status (β_1). Columns (1) (3) and (5) display results from pooled regression for Combined, weekly and biweekly sample. The estimated hazard coefficient of -0.232 (s.e=0.086) from the combined sample is very close to Chetty (2008)’s estimation of -0.233 (s.e=0.071), who used data from the Mathematica. My estimation from row 1 implies receiving severance pay leads to 20.7% reduction in UI claimant’s reemployment hazard. For the “weekly" and “biweekly" sub-samples, the coefficients have similar magnitudes, but is only statistically significant under “weekly" states. In the second row, I further restrict my sample to UI claimants with pre-unemployment asset information - this eliminates about 40% of the sample. Interestingly, the magnitude of the estimated β_1

Figure A2.1: Survival Curve - Effect of Severance Pay on Duration - by Pay Frequency

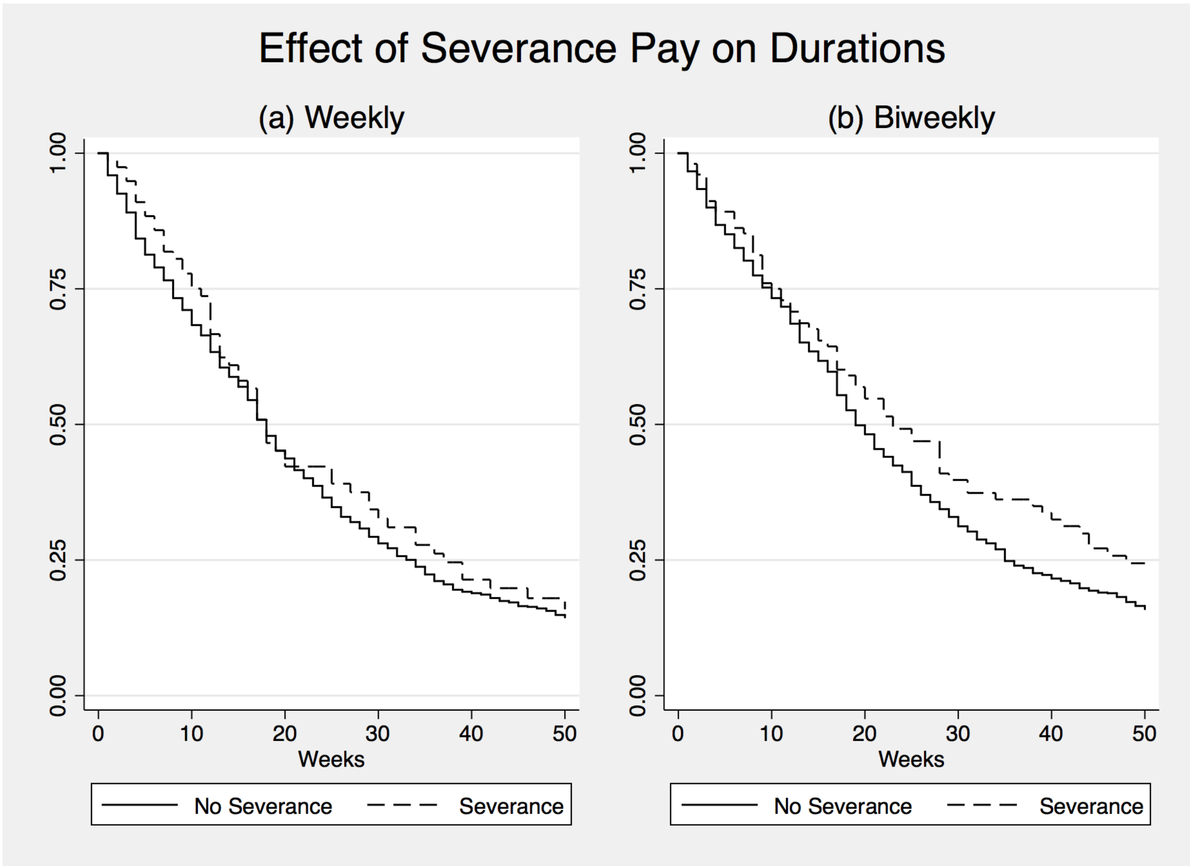

Note: Figure shows individual level unemployment duration from SIPP 1996-2007 controlling for job tenure. The vertical axes indicates the fraction of unemployed sample. The figure is divided into two panels according to UI benefit pay frequency. For each panel, the solid line represents the hazard of exiting unemployment for UI claimants without severance pay; the dashed line represents the probability of exiting unemployment for UI claimants with severance pay. Following Chetty (2008), the unemployment duration is censored at 50 weeks.

increases, so does the gap in the hazard coefficient between the two sub-samples.

To further investigate the causal interpretation of the estimated β_1, I run a series of regressions stratified by the relative length of predicted job tenure. Since the size of severance pay increases in job tenure, this regression allow me to test whether an increase in severance pay generosity leads to a bigger liquidity effect (lower reemployment hazard). As shown in Table A2.2 row 3 and 4, I find the liquidity effect is indeed stronger for severance pay recipients

Table A2.2: Effects of Severance Pay, by Pay Frequency

	COMBINED		WEEKLY		BIWEEKLY	
	Pooled	Stratified	Pooled	Stratified	Pooled	Stratified
	(1)	(2)	(3)	(4)	(5)	(6)
Severance Pay	-0.232		-0.254		-0.224	
	(0.086)		(0.095)		(0.157)	
Sev., *Ex-ante assets*	**-0.294**		**-0.340**		-0.280	
	(0.123)		(0.141)		(0.209)	
(Tenure$<$Median)$\times$ Sev.		-0.048		-0.091		-0.011
		(0.122)		(0.156)		(0.122)
(Tenure$>$Median)$\times$ Sev.		-0.357		-0.426		**-0.421**
		(0.118)		(0.155)		(0.202)
Equality p-value		0.057		0.166		0.123
# Spells	2,790	2,790	1,306	1,306	1,451	1,451
# Spells, *Ex-ante assets*	1,741		785		932	

Notes: All columns report semi-parametric Cox proportional hazard model results from estimating equation (A.1). The reported coefficients are the percent change in hazard rate with respect to severance pay status. Data are individual-level unemployment spells from 1996-2007 SIPP. For Pooled regression (columns (1) (3) and (5)), I include state fixed effects, year fixed effects, industry and occupation fixed effects, a 10-piece linear spline of the pre-unemployment annual wage, onseam indicator and other individual specific controls - education, age, marital status and total wealth. For stratified regression (columns (2) (4) and (6)), I allow controls to interact with tenure quantiles. The second row controls for household total wealth and restricts the sample to have pre-unemployment assets. The final row display the F-test result comparing coefficients for UI claimants from long or short job tenure quantiles. Standard errors clustered by state are in parentheses.

with longer job tenure. The pattern is a consistent across all sub-samples.

Under mean spell length, severance pay amount is equivalent to a 69% increase in UI benefit level under the biweekly pay states and 70% increase in UI benefit level under the weekly pay states. Using my estimated β_1 from Table A2.2 row 2, a 10% increase in UI benefit level would reduce reemployment hazard through the liquidity channel by 3.5% (insignificant) under the "biweekly" states and 4.1% under the "weekly" states.

Result A2.1: *Using cross-sectional variation in severance pay status, I find UI claimants respond to the effect of severance pay similarly under the two pay frequencies.*

A.3 Permutation Test for Extra Benefit Effect

Figure A3.1: Permutation test for inference of baseline estimation: extra benefit effect

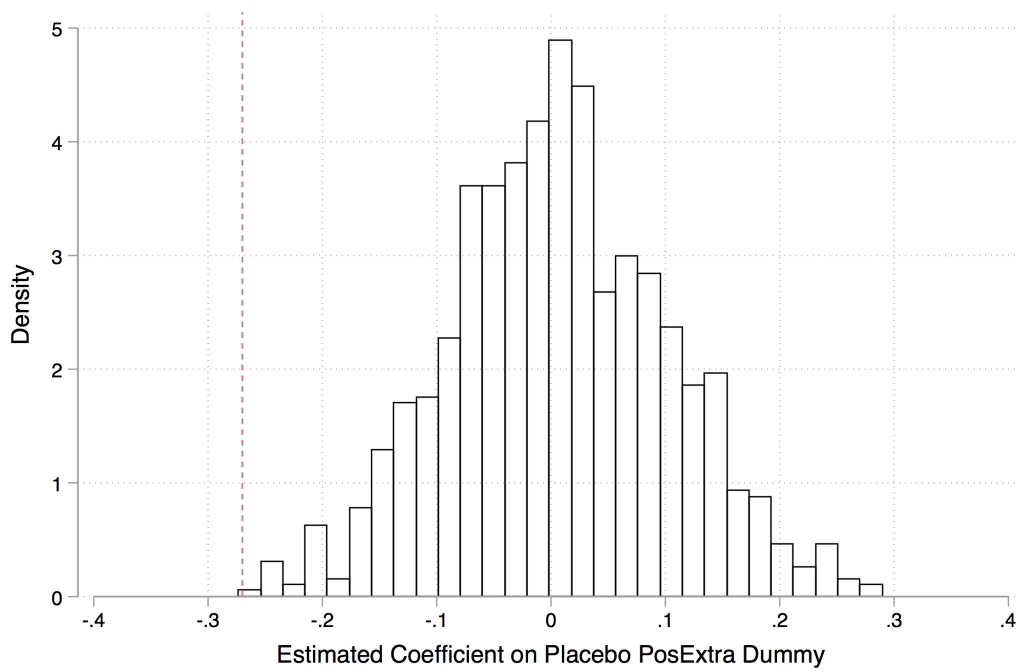

Note: Figures shows the empirical distribution of estimated placebo treatment effects from 1,000 random assignments. Dashed line is the actual treatment effect estimated from Table 1.3 Column (1). p-value under the permutation test is 0.001

I randomly assign treatment status (*PosExtra*) calendar months following the actual treatment timetable (Table A1.2).For example, given that 1985 has two *NoExtra* months, I randomly assign two calendar months within 1985 to be *NoExtra* and assign the rest as *PosExtra*. Following random treatment assignments, I re-estimate the placebo pay frequency effect following the specification (Table 1.3, Column (1)). Then I repeat this process for 1,000 times to obtain a distribution of estimated coefficients. The p-value in this context is defined as the probability that the baseline estimate is obtained purely by chance and is computed by the following expression:

$$p\text{-}value = \frac{\sum_{i=1}^{1000} \mathbb{1}|\beta^{i}_{baseline} \geq \beta_{placebo}|}{1000}$$

Figure A3.1 plots the empirical distribution of the placebo estimates using 1,000 random treatment assignments. The dashed line is the point estimate from the baseline estimation ($\beta = -0.277$). Comparing to the estimated placebo treatment effects, the actual effect is statistically significant (p-value = 0.001).

Appendix B

Appendix for “The Debt Payment Puzzle: An Experimental Investigation”

B.1 Additional Results

Table B1.1: 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

Notes: 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 B1.2: 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

Notes: 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 B1.3: 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
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.27$	$p = 0.90$	$p = 0.94$			
$\beta_{\mathbf{DR}xHigherInterestRate} = 0$				$p = 0.47$	$p = 0.39$	$p = 0.84$
$\beta_{\mathbf{DR}xHigherBalance} = 0$				$p = 0.02$	$p = 0.0044$	$p = 0.02$

Notes: 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 B1.4: 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
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.27$	$p = 0.92$	$p = 0.93$			
$\beta_{\mathbf{DR}xHigherInterestRate} = 0$				$p = 0.47$	$p = 0.38$	$p = 0.84$
$\beta_{\mathbf{DR}xHigherBalance} = 0$				$p = 0.02$	$p = 0.0044$	$p = 0.02$

Notes: The table executes the analysis in Table B1.3 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 B1.5: 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

Notes: 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 B1.6: 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

Notes: 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 B1.7: 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
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.01$	$p = 0.0001$	$p = 0.0027$			
$\beta_{\mathbf{IB}xHigherInterestRate} = 0$				$p = 0.04$	$p = 0.07$	$p = 0.03$
$\beta_{\mathbf{IB}xHigherBalance} = 0$				$p = 0.21$	$p = 0.77$	$p = 0.95$

Notes: 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 B1.8: 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
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.001$	$p = 0.0001$	$p = 0.003$			
$\beta_{\mathbf{IB}xHigherInterestRate} = 0$				$p = 0.04$	$p = 0.07$	$p = 0.03$
$\beta_{\mathbf{IB}xHigherBalance} = 0$				$p = 0.20$	$p = 0.78$	$p = 0.95$

Notes: The table executes the analysis in Table B1.3 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 B1.9: 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

Notes: 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 B1.10: 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
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.005$	$p = 0.0006$	$p = 0.0001$			
$\beta_{\mathbf{IR}xHigherInterestRate} = 0$				$p = 0.72$	$p = 0.76$	$p = 0.7$
$\beta_{\mathbf{IR}xHigherBalance} = 0$				$p = 0.53$	$p = 0.54$	$p = 0.26$

Notes: 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 B1.11: 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

Notes: 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 B1.12: 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
$\beta_{HigherInterestRate} = \beta_{HigherBalance}$	$p = 0.005$	$p = 0.0006$	$p = 0.0001$			
$\beta_{\mathbf{IR}xHigherInterestRate} = 0$				$p = 0.71$	$p = 0.78$	$p = 0.69$
$\beta_{\mathbf{IR}xHigherBalance} = 0$				$p = 0.53$	$p = 0.53$	$p = 0.26$

Notes: The table executes the analysis in Table B1.10 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.

B.2 Role of Vividness under the Investment Frame

In Figure B2.1, 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 B1.9 and B1.11 in Appendix).

Figure B2.1: Comparison of Investment Treatments

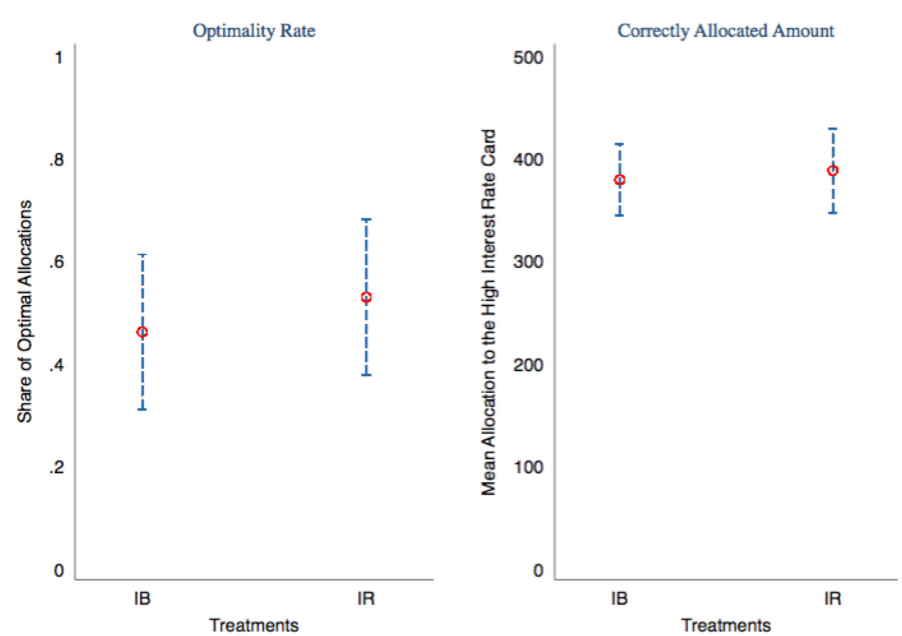

Notes: Panel A shows the share of optimal allocations made in **IB** and **IR**. The whiskers indicate 95% confidence interval calculated using subject-level clusters. Panel B shows the average allocation made to the high interest rate card.

Figure B2.2 documents further evidence that allows us to compare the allocation patterns

Figure B2.2: Allocation Patterns Across Investment Treatments - Period 1 Decisions

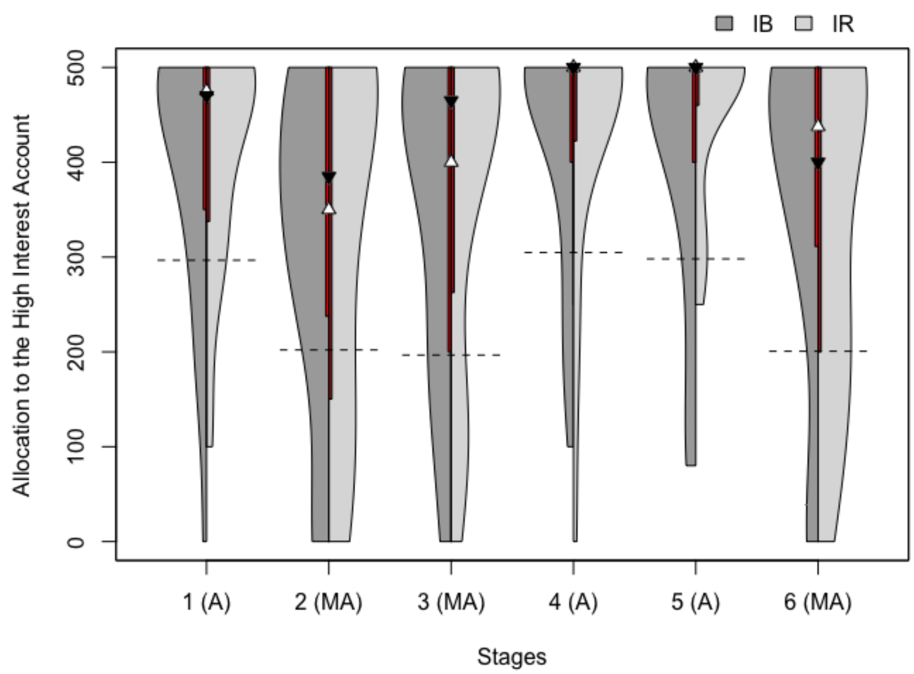

Notes: The violin plots show the distribution of repayments subjects make toward the high interest rate fund in the first period of each stage. The upward white triangle and the downward black triangle represent the median allocation towards the higher interest rate card in a given stage for **IB** and **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.

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 B1.10 and B1.12).

Result B2.1: *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.*

B.3 Learning

Table B3.1: 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

Notes: 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 B3.2: 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

Notes: 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 B3.3: 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

Notes: 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 B3.4: 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

Notes: 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 B3.5: 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

Notes: 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 B3.6: 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

Notes: 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 B3.7: 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

Notes: 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 B3.8: 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

Notes: 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.

Figure B3.1: Within and Between Stage Learning

(a) Within Stage Learning

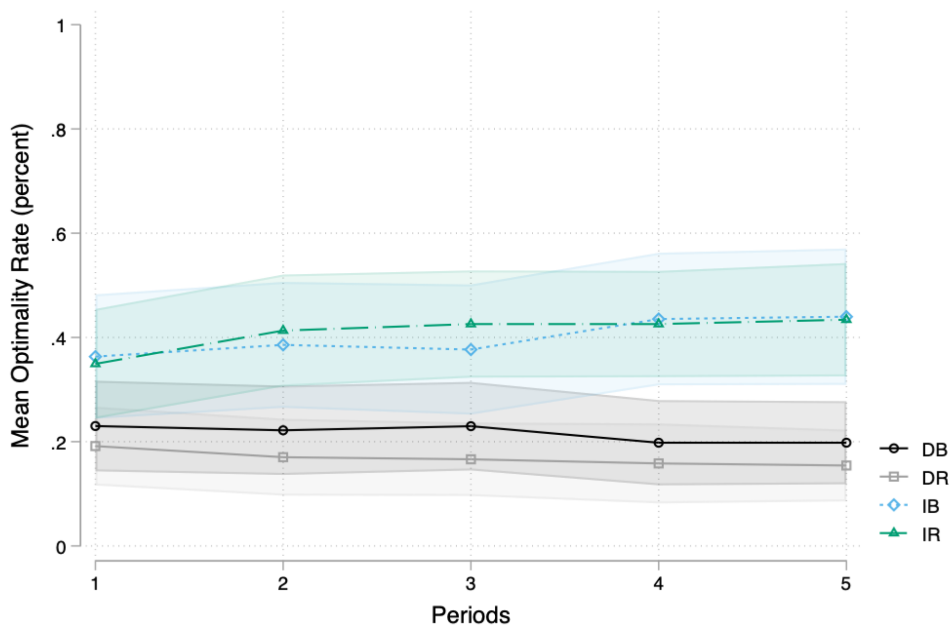

(b) Between Bi-Stage Learning

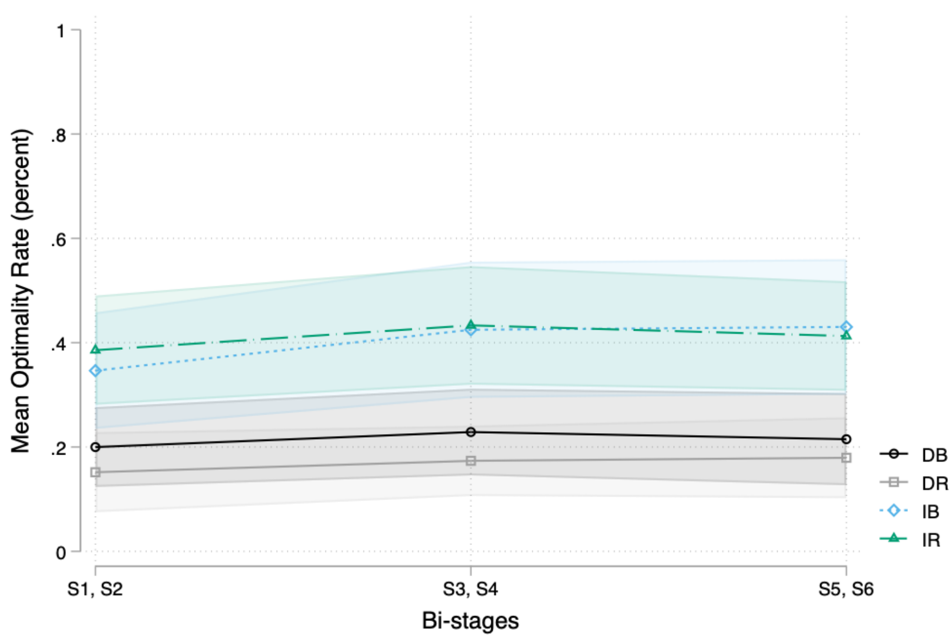

Notes: Figures control for individual fixed effects. Standard errors are clustered at the individual level.

B.4 Information Acquisition and the Measures of Optimality

In Section 10, 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 B4.1 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 ($p = 0.06$) 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 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 ($p = 0.70$). 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 B4.1: *Each click to interest rate button is correlated with an increase in the correctly*

Table B4.1: 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

Notes: 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.

allocated amount whereas each click to current balance button correlates with 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.

Table B4.2: 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

Notes: 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.

Table B4.3: 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

Notes: 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.

Figure B4.1: Click Order for All Periods: Debt No-Vivid

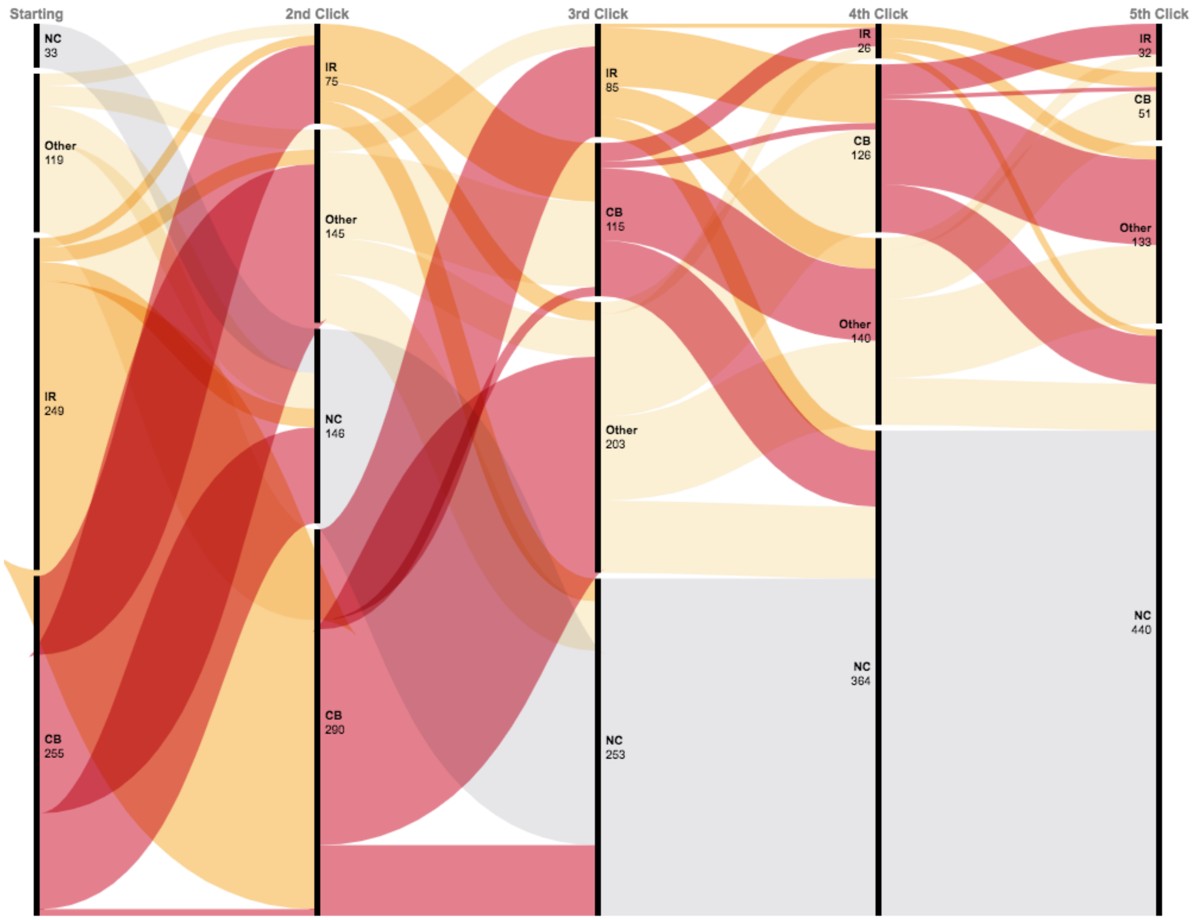

Notes: The river chart shows the click order information for all decision period from Debt No-Vivid Treatment. On average, majority of the subjects click Current Balance (CB) in their first click in each period.

Figure B4.2: Click Order for All Periods: Investment No-Vivid

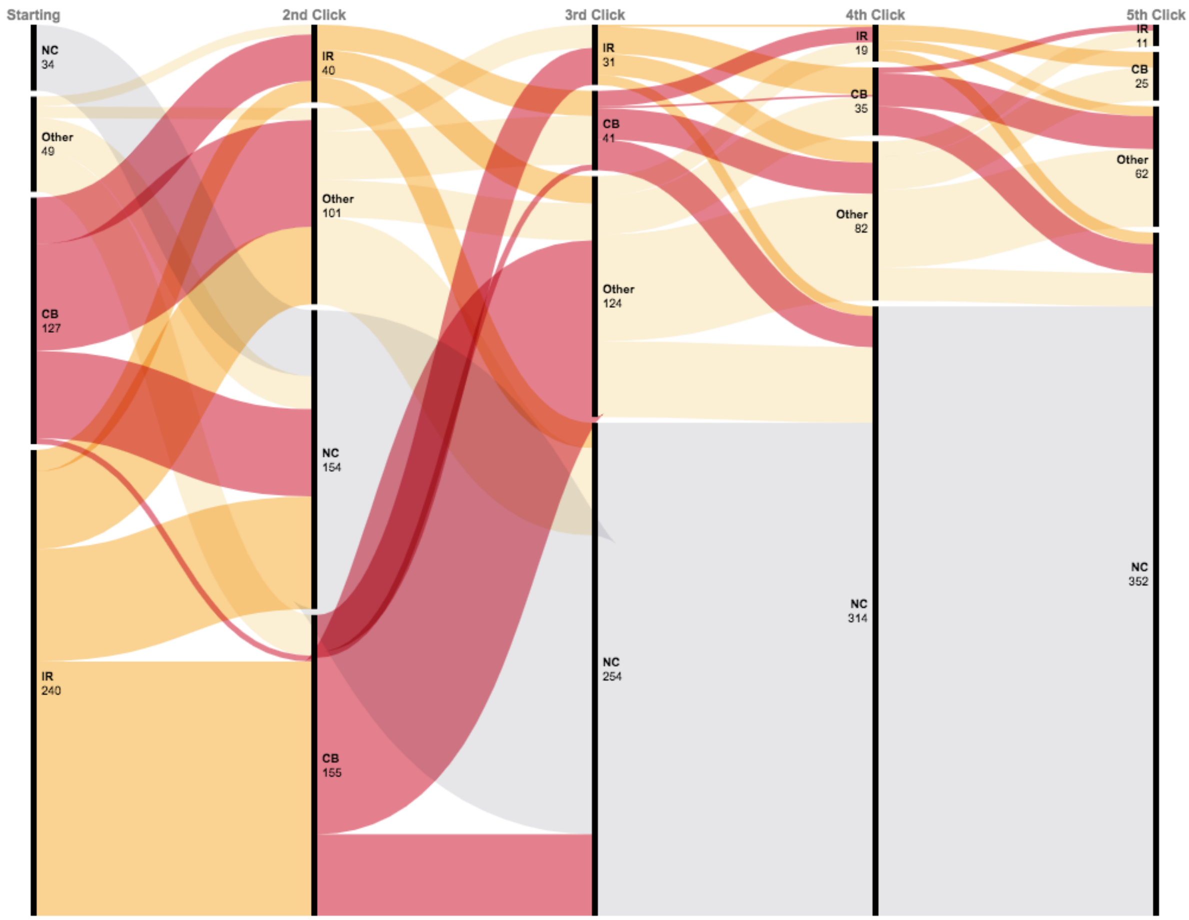

Notes: The river chart shows the click order information for all decision period from Investment No-Vivid Treatment. On average majority of the subjects click Interest Rate (IR) in their first click in each period.

B.5 Use of Heuristics - Heuristic Transition Matrices

Figure B5.1: Allocation Heuristics Transition Matrix

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

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

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

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

Notes: 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.

Figure B5.2: Typical Allocation Heuristics

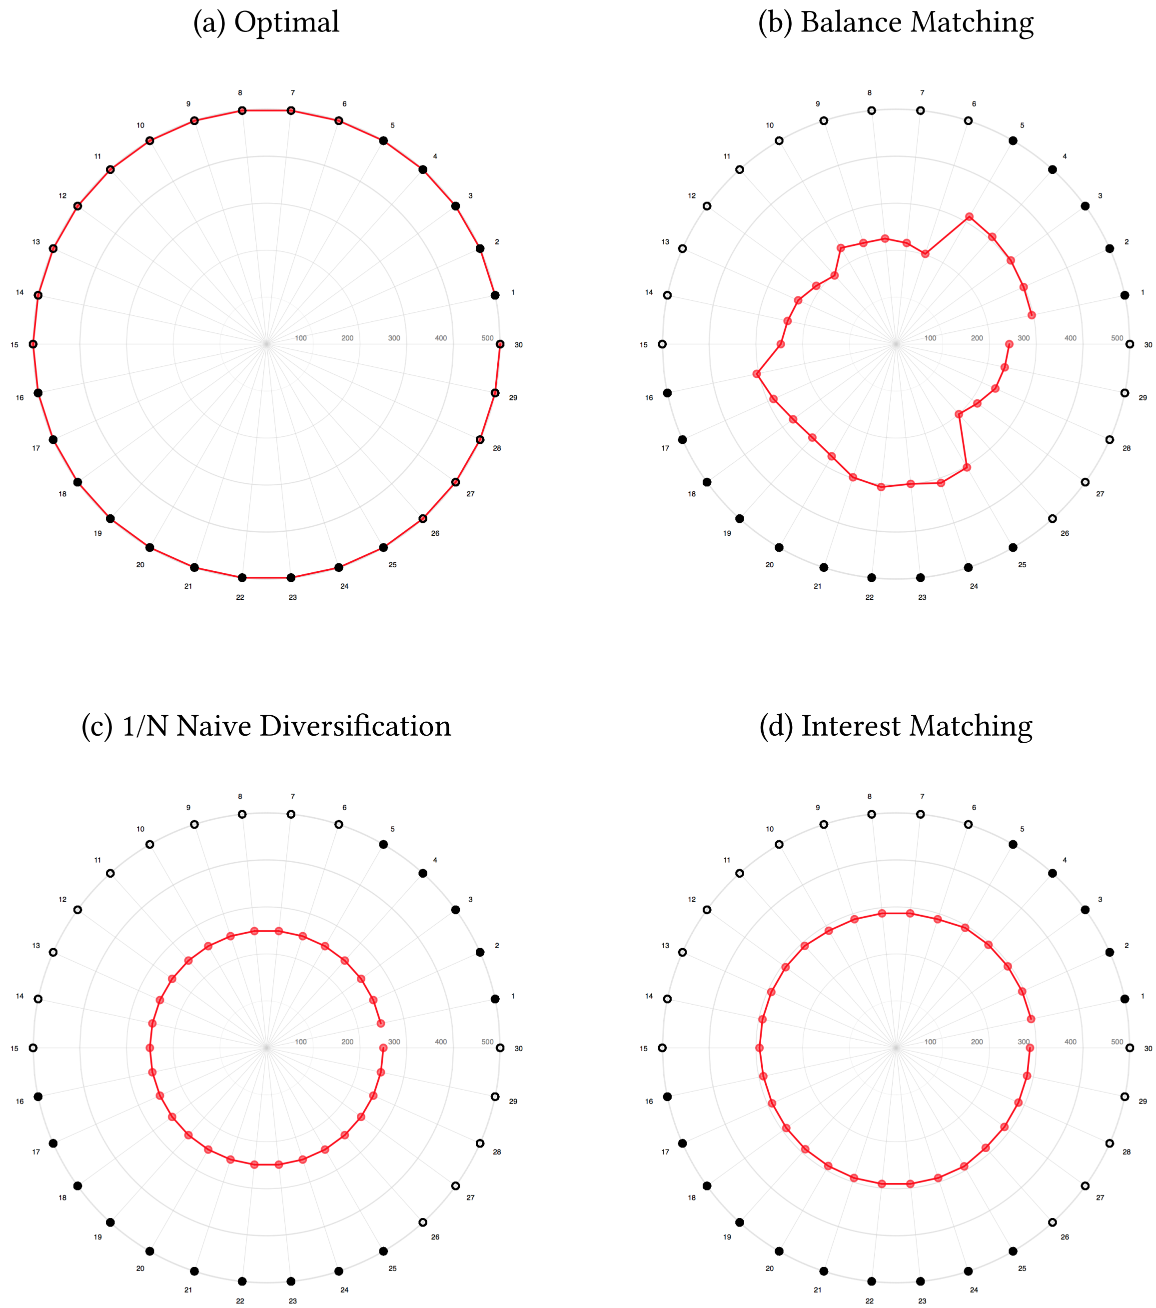

Notes: The polar figures show allocation patterns for the four common allocation heuristics. The red connected dot represents allocation towards higher interest account. Around the perimeter, the black dots indicate stages where higher interest comes with higher initial balances whereas the hollowed dots indicate stages where higher interest comes with lower initial balances. As illustrated in Figure (a), the optimal allocation rule requires subject to allocate 100% of their per period endowment (=500 ECU) for 30 payments/ investments toward the higher interest account.

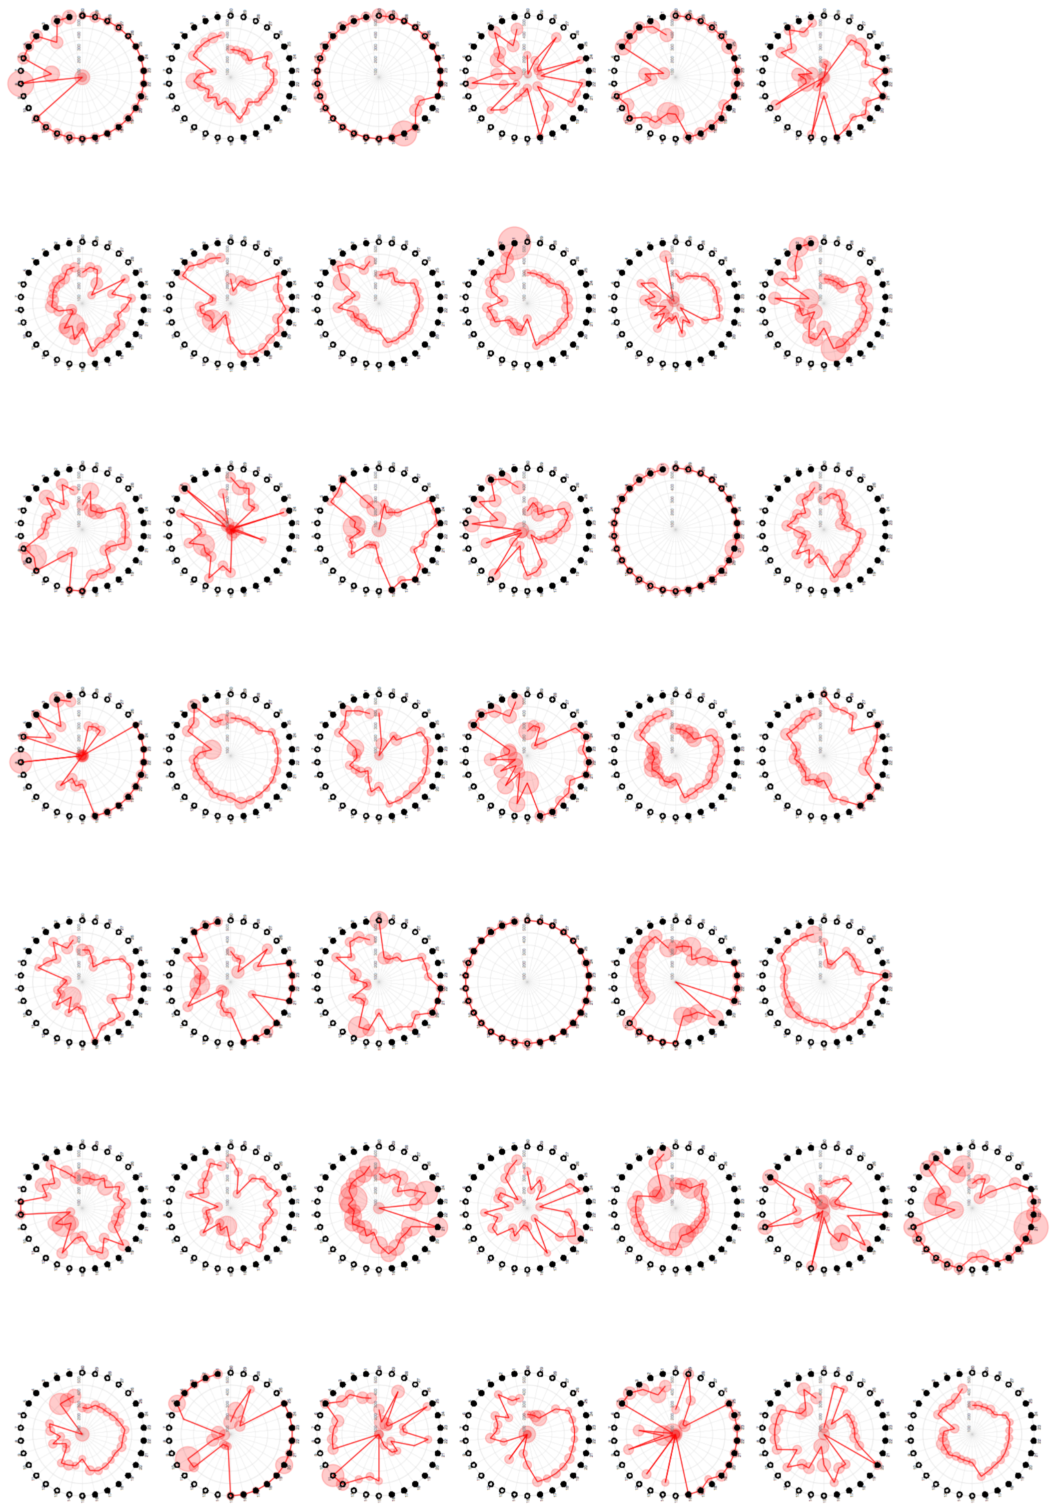

Figure B5.3: Raw Allocation Weighted by Time Spent per Period: **DB**

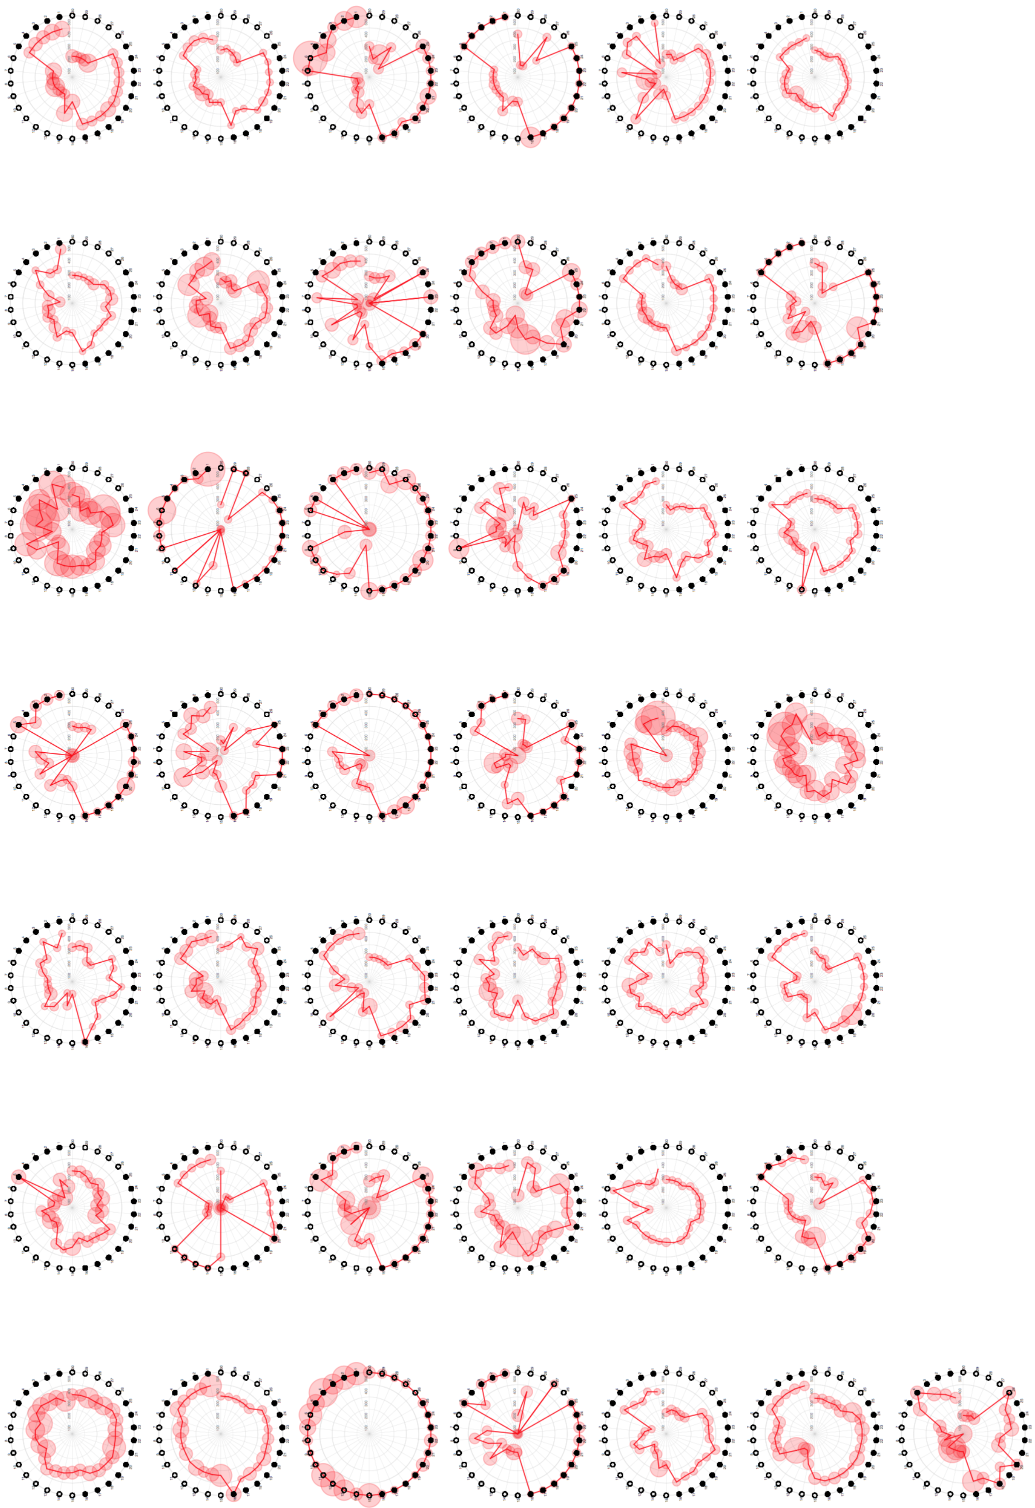

Figure B5.4: Raw Allocation Weighted by Time Spent per Period: **DR**

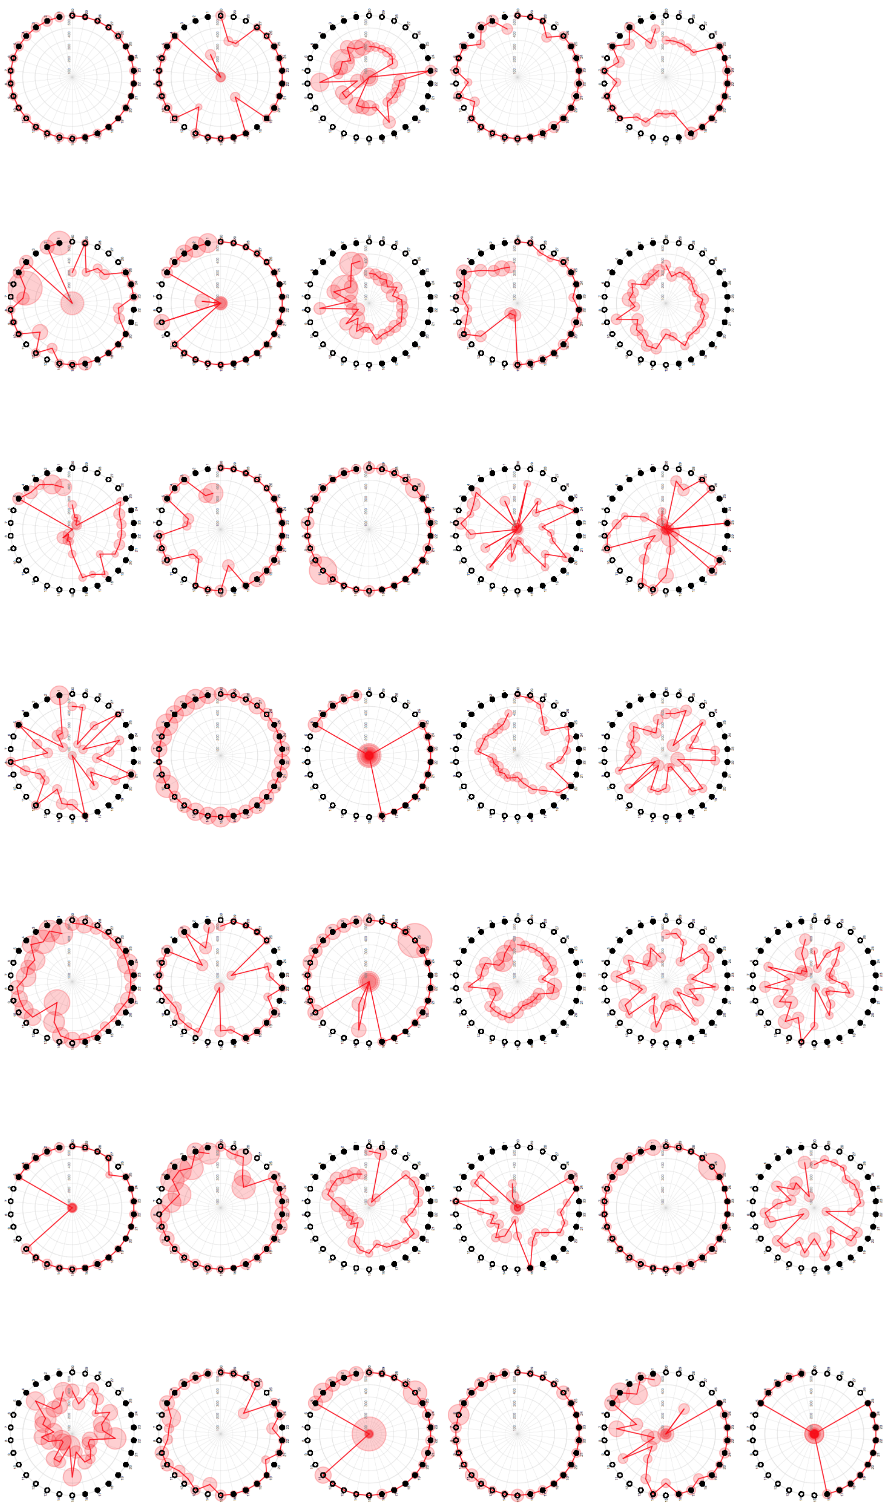

Figure B5.5: Raw Allocation Weighted by Time Spent per Period: **IB**

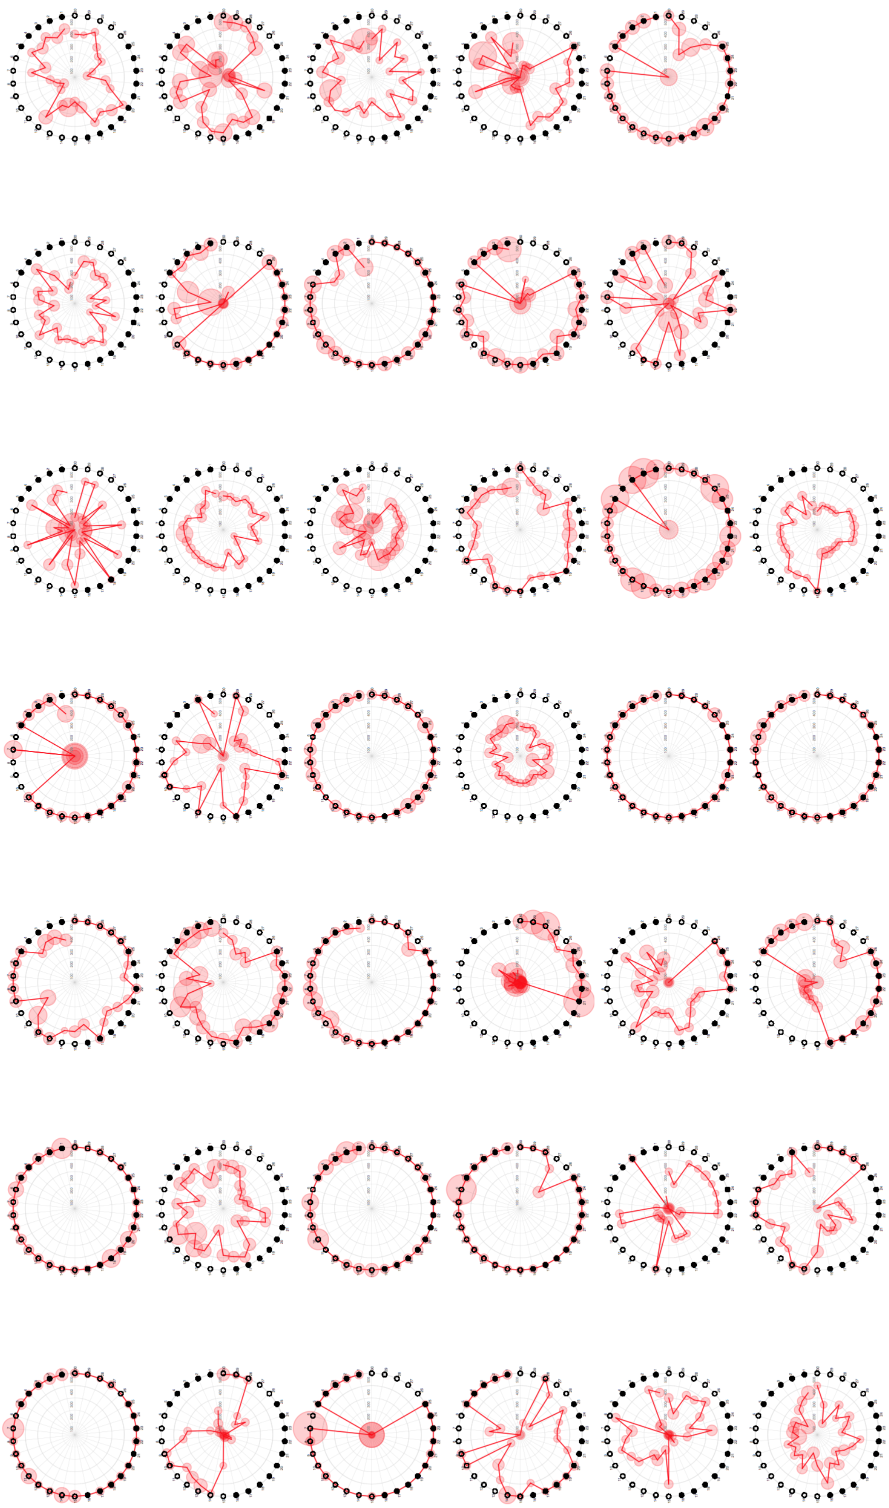

Figure B5.6: Raw Allocation Weighted by Time Spent per Period: **IR**

B.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 Taylor and Thompson (1982), 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 (Bordalo et al. (2013), Kőszegi and Szeidl (2012), Gabaix (2014)). 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.

B.7 Experiment Interface and Instructions

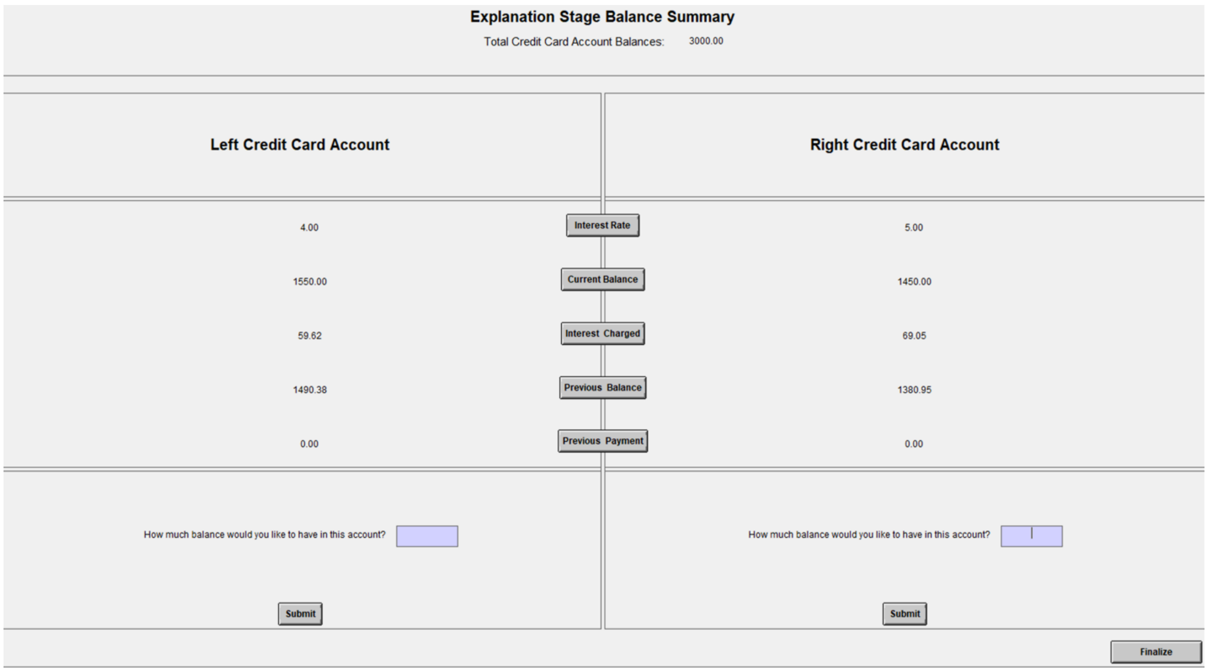

Figure B5.7: Experiment Interface for the treatment **DB** in Balance Reallocation Periods

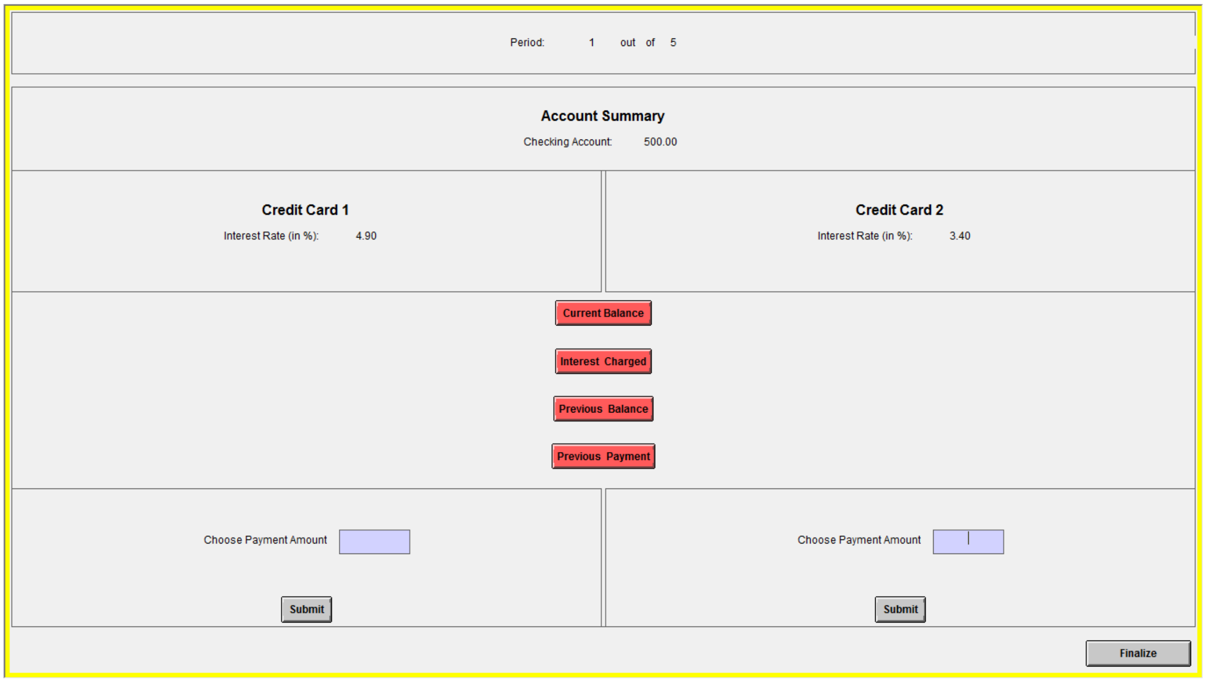

Figure B5.8: Experiment Interface for the treatment **DR**

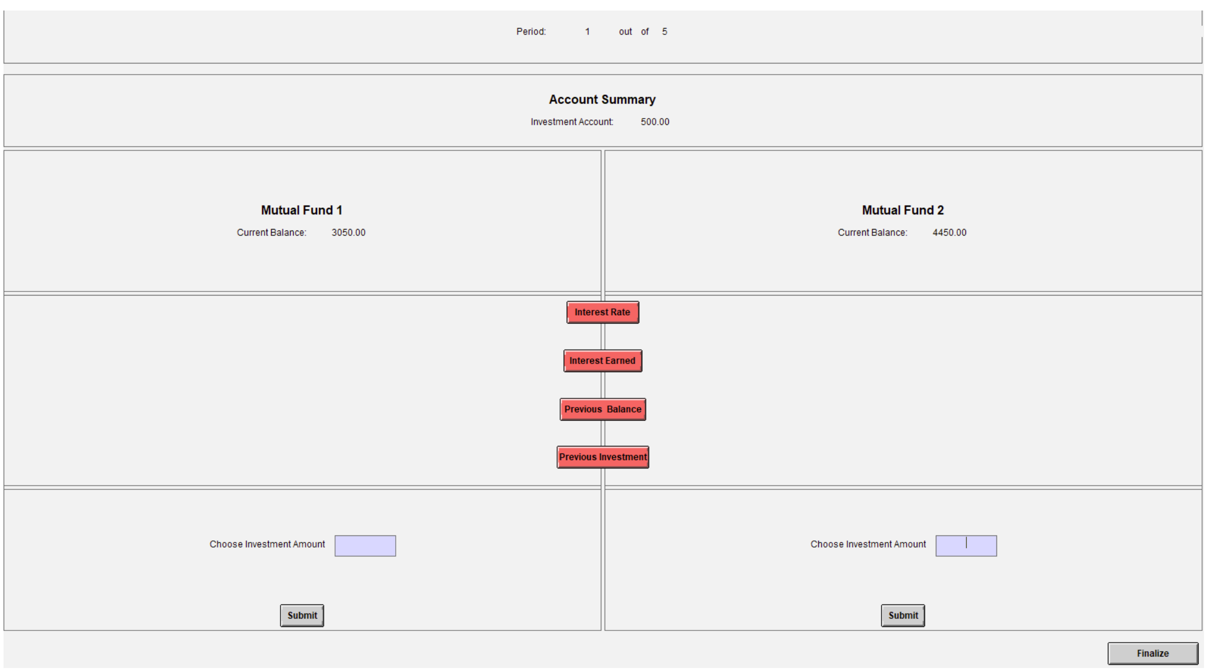

Figure B5.9: Experiment Interface for the treatment **IB**

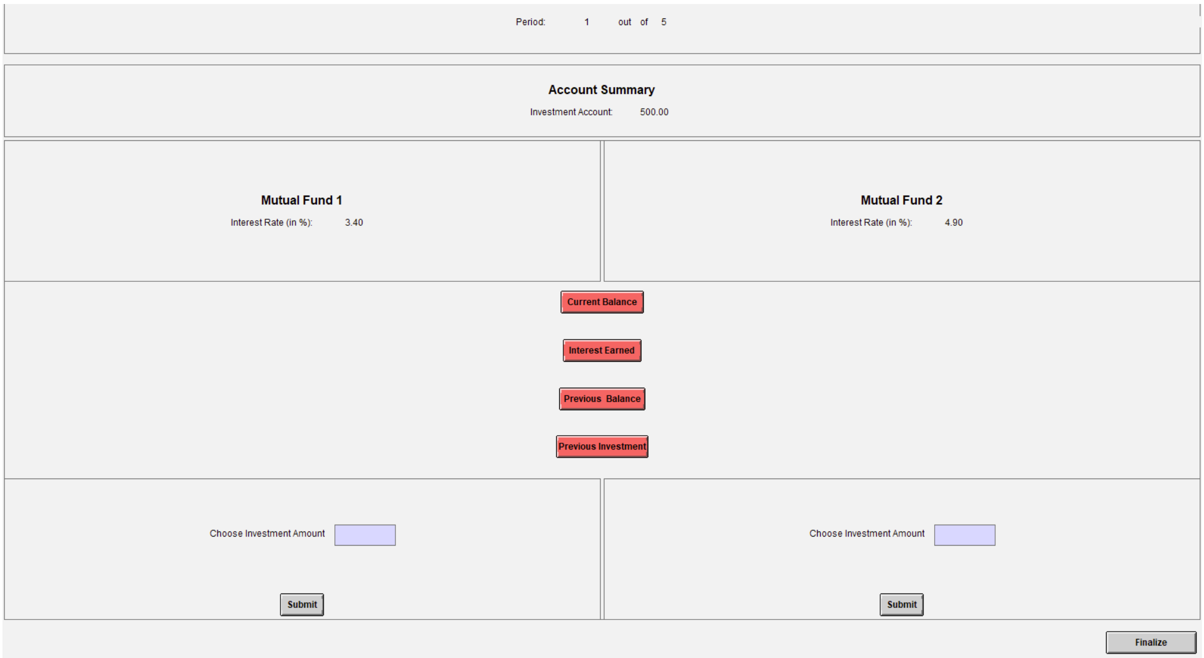

Figure B5.10: Experiment Interface for the treatment **IR**

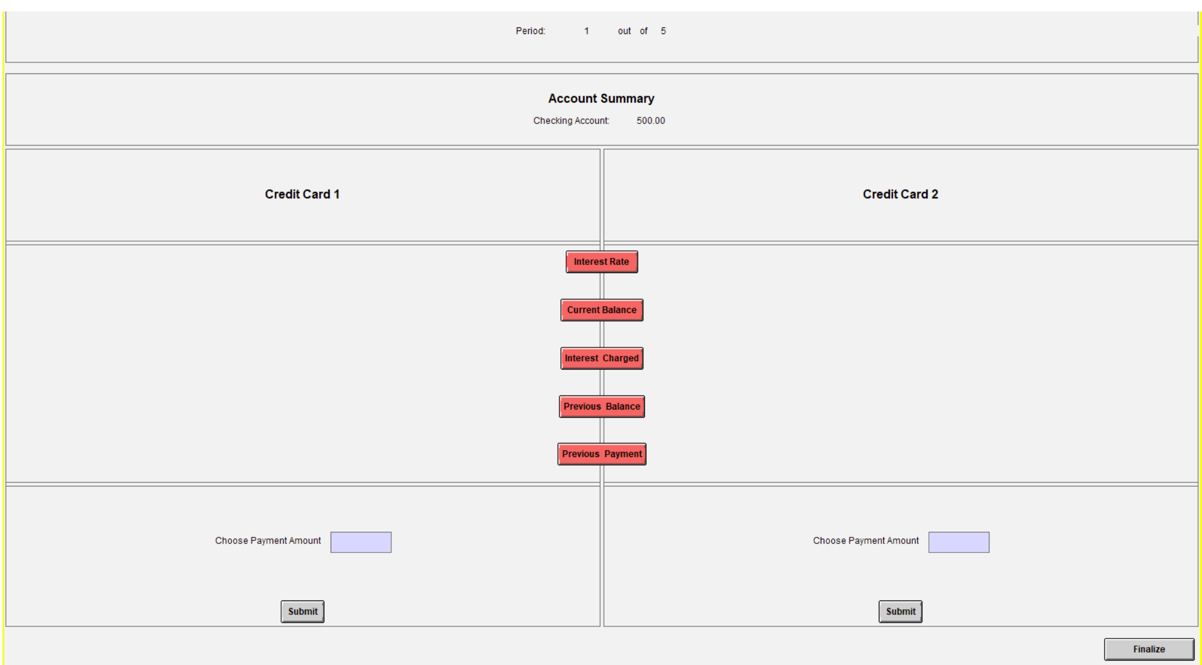

Figure B5.11: Experiment Interface for the treatment DN

Figure B5.12: 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

$$(1+4\%)\,(2{,}000 - \text{Payment to Left Card}) + (1+5\%)\,(1{,}000 - \text{Payment to Right Card})$$

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 C

Appendix for “Paying for Integers”

C.1 Data Refinement Procedure

The refinement procedure follows Haggag and Paci (2014).

1. Drew 2,000 random taxi driver and car pairs for each year
2. Dropped duplicate observations.
3. Drop-off time occurs before pick-up time.
4. Drop-off time occurs after subsequent trip pick-up time.
5. Ride duration was zero or longer than 3 hours.
6. Trip distance was zero or greater than 100 miles.
7. Surcharge amount was greater than $1.00.
8. Fare was less than $2.50 or negative fare amounts.
9. MTA tax was larger than $0.50.
10. Driver drove fewer than 100 rides for a given year.
11. Multiple cars were associated with the same driver during the same shift.
12. Driver’s shift was longer than 20 hours.
13. Driver’s shift was shorter than 30 minutes.
14. Either the pickup or drop-off location could not be mapped to census tract in New York, New Jersey, Connecticut or Pennsylvania
15. Dropped fares were categorized as “Dispute" or “No Charge"
16. Switched variable names between “Tip Amount" and “Tolls Amount" for Dec2011 fare.

17. Dropped rides with cash transactions.[1]

Table C1.1: Summary of Cash and Credit Differences: Feb-Aug 2012

	(1) Cash	(2) Credit	(3) Difference
Fare Amount	8.58	9.48	-0.90***
	(4.76)	(4.93)	(0.01)
Trip Length (in minutes)	10.78	12.10	-1.33***
	(7.25)	(7.48)	(0.01)
Trip Distance (in miles)	2.20	2.54	-0.34***
	(2.03)	(2.11)	(0.00)
Fraction VTS	0.50	0.50	-0.00***
	(0.50)	(0.50)	(0.00)
Pickup Location Median Income	95,948.18	95,919.98	28.20
	(38,495.11)	(36,906.74)	(62.21)
Fraction Low Option Integer	0.03	0.02	0.01***
	(0.17)	(0.16)	(0.00)
Fraction Mid or High Option Integer	0.01	0.01	-0.00
	(0.10)	(0.10)	(0.00)
Observations	785,300	710,059	1,495,359

Notes: This table presents the summary statistics for the random sample of 2,000 taxi drivers during the time period of our main study: February to August 2012. Standard deviations are in parenthesis.

[1] See Table C1.1 for a comparison between cash and credit transactions.

C.2 Simulating Impact of Integer Default Tip Suggestions

The impact of integer default tip suggestions on the utility from the menu option is relatively straightforward. A small change in the fare that leads to an integer tip suggestion can impact the utility of this option depending on the value of b_i, as is shown in Figure C2.1. When $b_i = 0$ there is no change in the utility of the default tip option based on whether or not the tip suggestion, $t_i^D F_i$, is an integer. However, when $b_i > 0$ a passenger's utility from the default option exhibits discontinuously higher utility when the tip suggestion is an integer.

Figure C2.1: Individual's Utility from Default Tip Suggestion in Response to Fare amount, by Different b_i

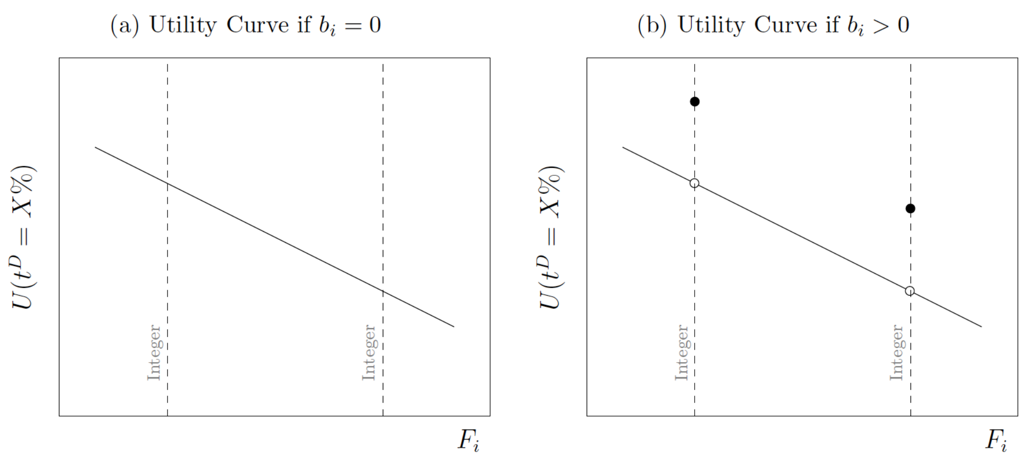

Notes: Figures presents the relationship between utility of taking default tip suggestion and fare amounts under the extended model by different b_i, for a given default tip rate $t_i^D = X\%$. Panel (a) presents the relationship when we set $b_i = 0$. Under this case, increases in F_i *smoothly* decreases one's utility. Panel (b) presents the same relationship when we set $b_i > 0$. Under this case, we observe discontinuous sharp increases in $U(t_i^D)$ when $t_i^D F_i \in \mathbb{Z}$.

The impact of integer default tip suggestions on the utility from custom tips is less clear as the preferred custom tip depends on comparing the tip rate that satisfies equation (3.2) with alternative tip rates that lead to integer tips, as shown in the right panel from Figure 3.5. Although it is unlikely, one could imagine that tip rates that satisfy equation (3.2) tend to lead to integer tip suggestions when default tip rates are integers. To explore this, we parameterize the utility function and plot the utility of the preferred custom and default tips based on distance to the integer tip suggestion.

The primary piece of customer's utility that we need to put structure to in order to simulate utility is the norm-deviation cost $v(T_i, t_i)$. Following (Donkor, 2020), we define the norm-deviation cost as $\theta(T_i - t_i)^2$. We can then write a generic passenger's utility function as:

$$Max_{t_i} U = -t_i F_i - \underbrace{\theta(T_i - t_i)^2}_{v(T_i,t_i)} - \mathbb{1}\{t_i \neq t_i^D\}[c_i^{non} - \alpha_i \cdot \mathbb{1}\{t_i F_i \in \mathbb{Z}\}] + b_i \cdot \mathbb{1}\{t_i F_i \in \mathbb{Z}\} \quad \text{(C2.1)}$$

where θ "scales" the impact of deviating from what the passenger perceives as the socially accepted tip.

We are primarily interested in investigating whether, under reasonable parameters, utility from custom tips exhibit a discontinuity when default tip suggestions are integers. For this exercise, we will thus make the following parameter assumptions:

- $T_i = 0.15$ or $T_i = 0.18$
- $t_i^D = 0.2$
- $\theta = 1000$
- $c_i^{non} = 0.6$
- $\alpha = 0.1$
- $b_i = 0.1$

For fares ranging from 0 to 100, we then calculate the utility for the tip rate that satisfies equation (3.2) and the closest tip rates that lead to integer tips. We then calculate $U(t_i^C)$ as the custom tip, integer or not, that gives the highest utility to the passenger for that fare.

Alternatively, for default tips, we calculate the utility for a single default tip rate of 0.2 for all the fares from 0 to 100.

Given the default tip rate of 0.2, integer default tip suggestions will occur at fares of 5, 10, 15, etc. To highlight any discontinuities in utility around these values, we calculate the average utility for default and custom tip rates at values around the integer default tip suggestion. Specifically, we calculate the distance between the fare and the closest fare that leads to an integer default tip suggestion. In practice, this means that fares of 4.5 and 9.5 would be treated similarly since their decimal distance is -0.5 (-50 cents), while fares of 5.5 and 10.5 would have decimal distance equal +0.5. We then calculate the average default tip option and custom tip option utility based on the decimal distance. If there is a discontinuity, on average, then this would be shown in a spike at the value of 0. Figure C2.2 shows that this is evident for default tip suggestions, but not custom tips. Importantly, the lack of a discontinuity for custom tip rates does not appear to be a result of the choice of T_i as the results are robust to alternative T_i besides those shown here. In addition, in all alternative specifications for the other parameters (θ, c_i^{non}, α, and b_i) that we have simulated, the conclusions are similar although the utility levels and magnitudes of the spikes for the default option can vary.

In summary, the simulation shown in Figure C2.2 highlights that custom tip utility appears to be continuous when presented with default tip suggestions. Intuitively, this is because the primary concern was that custom tip rates that satisfy equation (3.2) lead to integer tips more frequently when the default tip suggestion is also an integer. There is no reason ex-ante to think that this would be the case, which is supported by Figure C2.2.[2]

[2]Intuitively, one could think that customer's prefer a tip rate of 0.1, which would also frequently have integer tip suggestions when $t_i^D = 0.2$. Our theory, however, would suggest that even if passenger's believe the socially accepted tip rate is 0.1, they would "shade downwards" their preferred custom tip.

Figure C2.2: Custom and Default Tip Utility by Distance to Integer Default Tip Suggestion

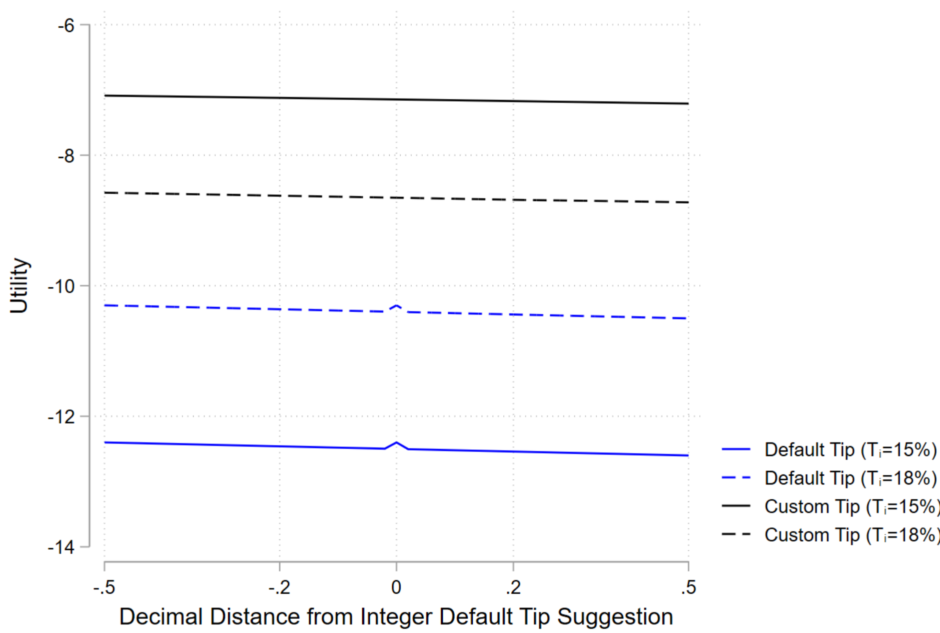

Notes: This figure plots the utility of choosing a custom tip compared to a default tip option based on the distance of the fare from the closest fare that leads to a default tip suggestion that is an integer. The range of fares used to create this figure is from 0 to 100. For the default tip rate of 0.2 used here, this means that the utility shown at 0 corresponds to the average utility at fares of 5, 10, 15, etc, while -0.5 represents 4.5, 9.5, 14.5, etc. The utility function used for this figure is:

$$U = -t_i F_i - \underbrace{\theta(T_i - t_i)^2}_{v(T_i,t_i)} - \mathbb{1}\{t_i \neq t_i^D\}[c_i^{non} - \alpha_i \cdot \mathbb{1}\{t_i F_i \in \mathbb{Z}\}] + b_i \cdot \mathbb{1}\{t_i F_i \in \mathbb{Z}\}$$

where we set $\theta = 1000$, $c_i^{non} = 0.6$, $\alpha = 0.1$, and $b_i = 0.1$. Solid lines show when $T_i = 0.15$ and dashed lines show when we set $T_i = 0.18$. To calculate the default tip utility for each fare, we change F_i leaving all else constant. To calculate the custom tip utility for each fare, we change F_i and find the custom tip rate that maximizes utility, ignoring the default option, at that point.

C.3 Additional Figures and Tables

Figure C3.1: Passenger Display for CMT in 2012

Notes: This figure shows the screen for a CMT outfitted vehicle in 2012. The source is the online appendix to Haggag and Paci (2014), Figure A.1, which was a photo taken by the authors.

Figure C3.2: Distribution of Tip Suggestion: Feb – Aug 2012

(a) Low Suggestion (20%)

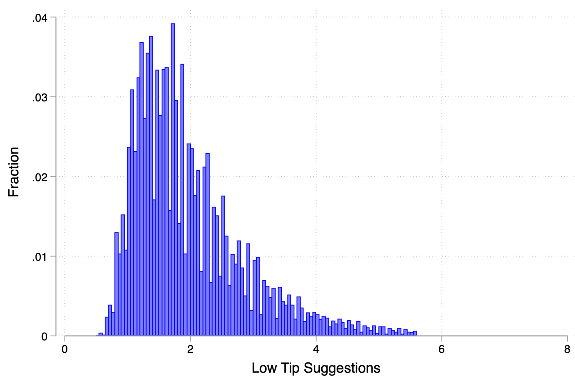

(b) Middle Suggestion (25%)

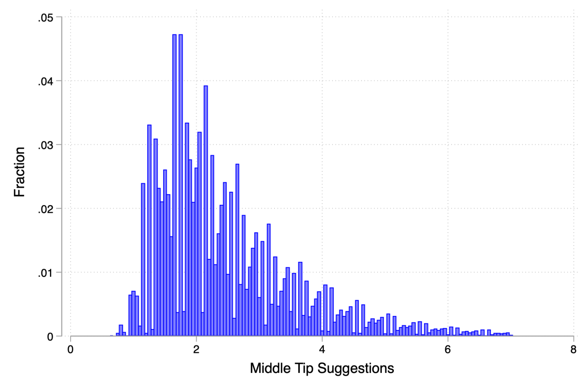

(c) High Suggestion (30%)

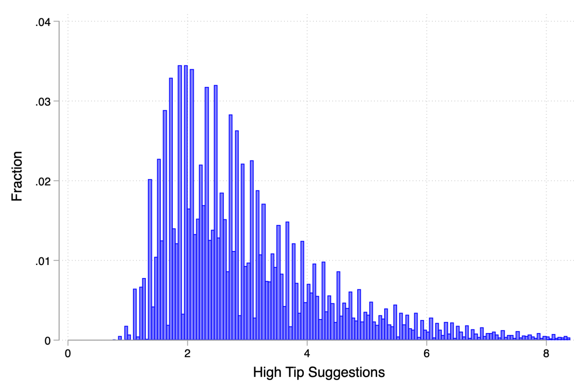

Notes: Panels (a) (b) and (c) shows the distributions of tip suggestions for the low, middle and high options. Extreme tip suggestion ($> 99^{th}$ percentile) are excluded from the figure. During Feb–Aug 2012, the % tip suggestion options (20-25-30) were identical to CMT and VTS taxis.

Figure C3.3: Distribution of Custom Tip Amount: Feb – Aug 2012

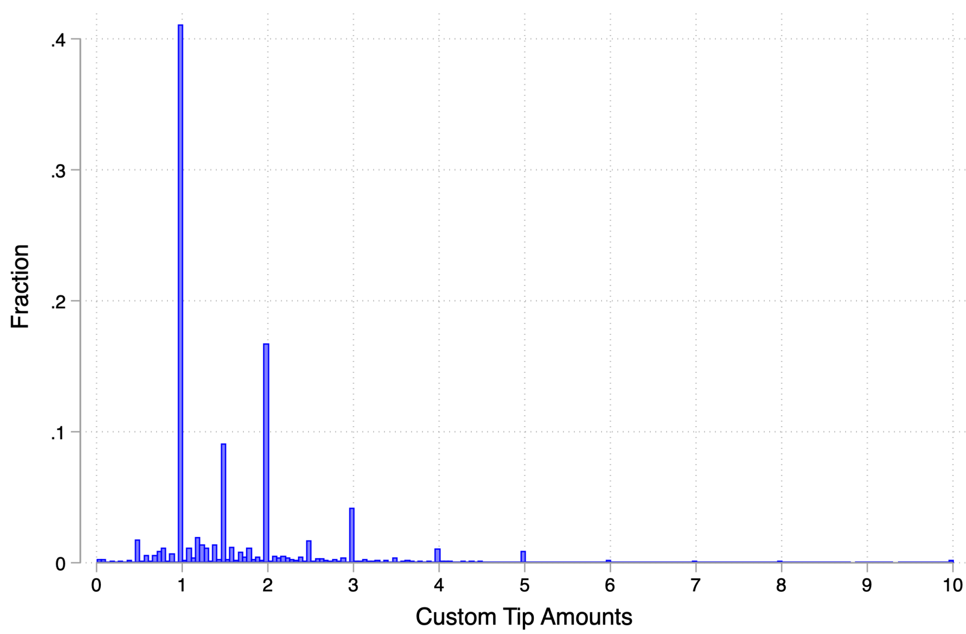

Notes: This figure shows the distribution of Custom tip amounts for all non-airport trips that were paid by credit card. Custom tips includes all non-zero tips that are not equal to any of the tip suggestions. Extreme custom tip amounts ($> 99^{th}$ percentile) are excluded from the figure. All tip amounts are in nominal dollar value.

Figure C3.4: Placebo Effects

(a) Default Take-up

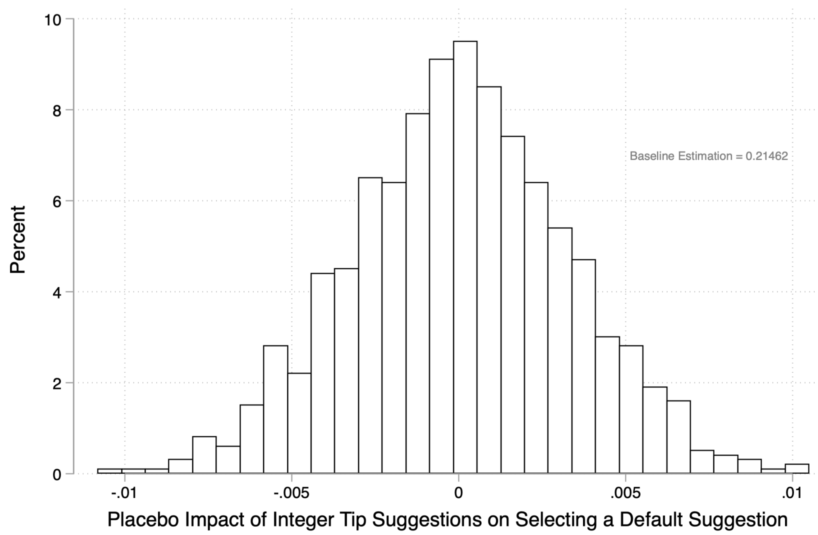

(b) Tip Rate

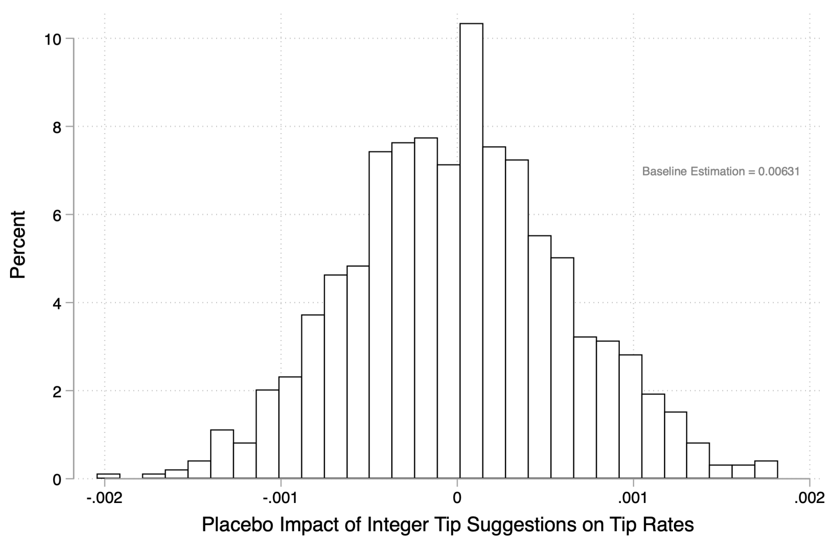

Notes: Figures shows the empirical distribution of estimated placebo treatment effects from 1,000 random treatment (trip with integer tip suggestion) assignments. The actual treatment effects are estimated from Table 3.2 Column (3) and Table 3.3 Column (3). p-values under the placebo tests are both < 0.001.

Figure C3.5: Impact of the VTS Menu Change on Tip Rates: RD in Time

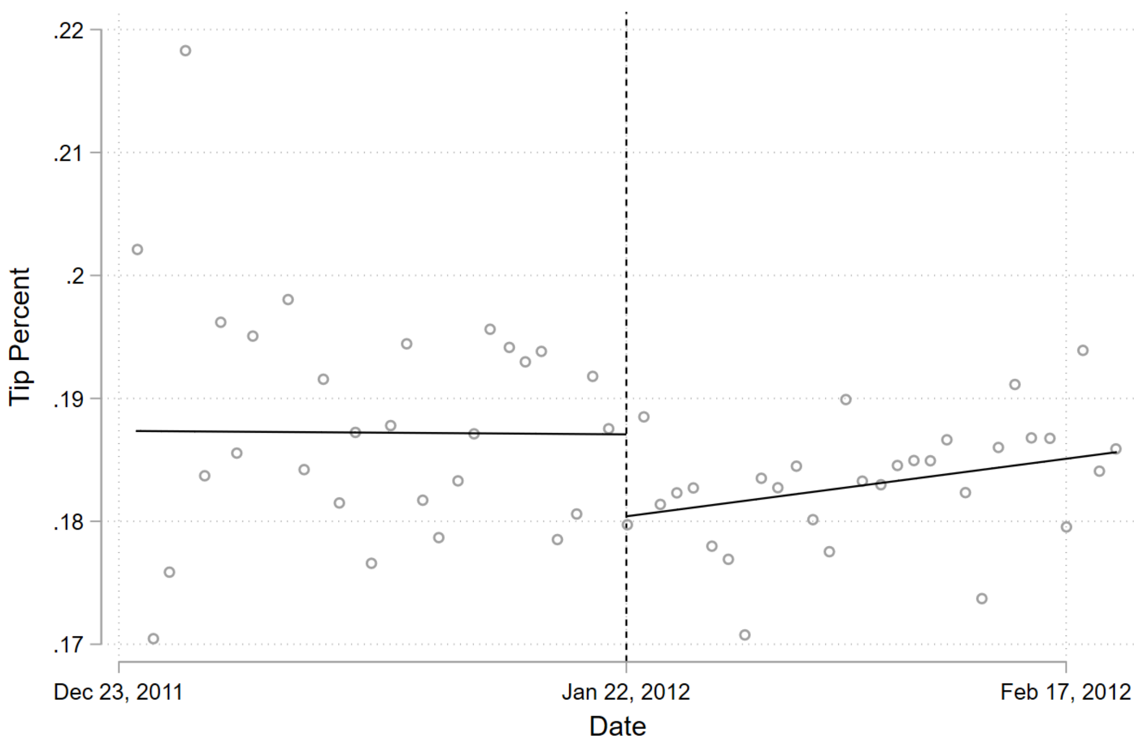

Notes: This figure presents the plot of regression discontinuity in time. Prior to the event, the tip suggestion menu offered by VTS had 2, 3, and 4 dollars suggestions for rate fare (fare + surcharge) below $15. After the week of January 22, 2012, VTS removed the dollar tip suggestion and replaced it with the 20, 25, 30 percent tip suggestions.

Figure C3.6: Effect of VTS Menu Change in 2012

(a) Selecting Options from the Menu

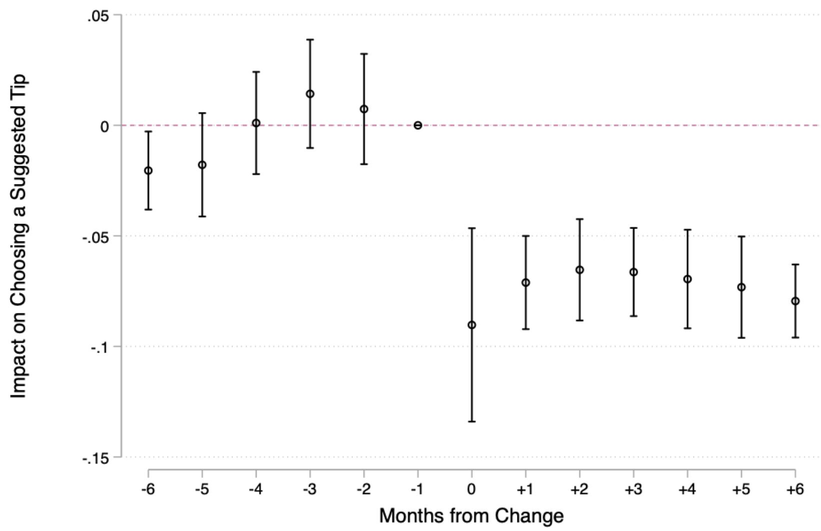

(b) Tip Rate

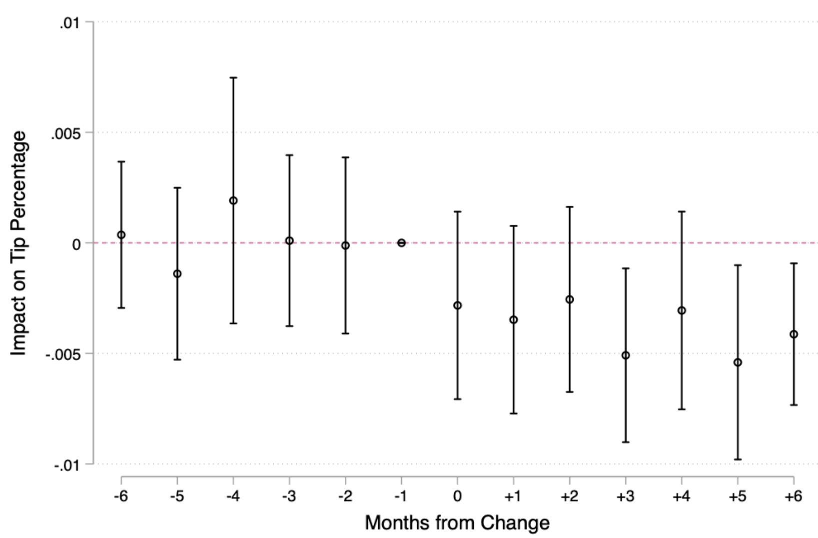

Notes: Figures show the event study plots for VTS menu change in Jan.2012. We control for tip or fare policy variations, and we include pick-up date fixed effects and vendor fixed effects. In addition, we include $x(d, mph)$ control. We include samples with non-zero tips. We cluster the standard errors at the pick-up date level.

References

Aguila, E., Kapteyn, A., & Perez-Arce, F. (2017). Consumption smoothing and frequency of benefit payments of cash transfer programs. *American Economic Review*, *107*(5), 430–35.

Alexander, D., Boone, C., & Lynn, M. (2021). The effects of tip recommendations on customer tipping, satisfaction, repatronage, and spending. *Management Science*, *67*(1), 146–165.

Allen, E. J., Dechow, P. M., Pope, D. G., & Wu, G. (2017). Reference-dependent preferences: Evidence from marathon runners. *Management Science*, *63*(6), 1657–1672.

Andersen, S., Campbell, J. Y., Nielsen, K. M., & Ramadorai, T. (in press). Sources of inaction in household finance: Evidence from the danish mortgage market. *American Economic Review*.

Anderson, P. M., & Meyer, B. D. (1997). Unemployment insurance takeup rates and the after-tax value of benefits. *The quarterly journal of economics*, *112*(3), 913–937.

Ashenfelter, O., Ashmore, D., & Deschênes, O. (2005). Do unemployment insurance recipients actively seek work? evidence from randomized trials in four us states. *Journal of econometrics*, *125*(1-2), 53–75.

Azar, O. H. (2007). The social norm of tipping: A review. *Journal of Applied Social Psychology*, *37*(2), 380–402.

Azar, O. H. (2008). Strategic behavior and social norms in tipped service industries. *The BE Journal of Economic Analysis & Policy*, *8*(1), 0000102202193516821778.

Baugh, B., & Correia, F. (2018). Does paycheck frequency matter for households' decisions? evidence from financial account data. *working paper*.

Baugh, B., Leary, J. B., & Wang, J. (2018). When is it hard to make ends meet? *RRC Paper No. NB1l7-05. Cambridge, MA: National Bureau of Economic Research*.

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of general psychology*, *5*(4), 323–370.

Beatty, T. K., Bitler, M. P., Cheng, X. H., & Van der Werf, C. (2019). *Snap and paycheck cycles* (Tech. Rep.). National Bureau of Economic Research.

Benartzi, S., & Thaler, R. (2007, September). Heuristics and biases in retirement savings behavior. *Journal of Economic Perspectives*, *21*(3), 81-104. Retrieved from http://www.aeaweb.org/articles?id=10.1257/jep.21.3.81 doi: 10.1257/jep.21.3.81

Berniell, I. (2018). *Pay cycles: Individual and aggregate effects of paycheck frequency* (Tech. Rep.). Documento de Trabajo.

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

Beshears, J., Choi, J. J., Laibson, D., & Madrian, B. C. (2018, July). *Behavioral household finance* (Working Paper No. 24854). National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w24854 doi: 10.3386/w24854

Bhutta, N., Fuster, A., & Hizmo, A. (2020). Paying too much? price dispersion in the us mortgage market.

Blaustein, S. J. (1979). *Insured unemployment data* (No. 24). National Commission on Employment and Unemployment Statistics.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2013). Salience and consumer choice. *Journal of Political Economy*, *121*(5), 803-843. doi: 10.1086/673885

Bordalo, P., Gennaioli, N., & Shleifer, A. (2017). *Memory, attention, and choice* (Tech. Rep.). National Bureau of Economic Research.

Browning, M., & Crossley, T. F. (2001a). The life-cycle model of consumption and saving. *Journal of Economic Perspectives*, *15*(3), 3–22.

Browning, M., & Crossley, T. F. (2001b). Unemployment insurance benefit levels and consumption changes. *Journal of public Economics*, *80*(1), 1–23.

Browning, M., & Lusardi, A. (1996). Household saving: Micro theories and micro facts. *Journal of Economic literature*, *34*(4), 1797–1855.

Campbell, J. Y. (2016). Restoring rational choice: The challenge of consumer financial regulation. *American Economic Review*, *106*(5), 1–30.

Card, D., Chetty, R., & Weber, A. (2007). Cash-on-hand and competing models of intertemporal behavior: New evidence from the labor market. *The Quarterly journal of economics*, *122*(4), 1511–1560.

Castner, L., Henke, J., et al. (2011). *Benefit redemption patterns in the supplemental nutrition assistance program* (Tech. Rep.). Mathematica Policy Research.

Chandar, B., Gneezy, U., List, J. A., & Muir, I. (2019). *The drivers of social preferences: Evidence from a nationwide tipping field experiment* (Tech. Rep.). National Bureau of Economic Research.

Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of political Economy*, *116*(2), 173–234.

Chetty, R., Looney, A., & Kroft, K. (2009). Salience and taxation: Theory and evidence. *American economic review*, *99*(4), 1145–77.

Donkor, K. B. (2020). How social norms and menus affect choices: Evidence from tipping.

East, C. N., & Kuka, E. (2015). Reexamining the consumption smoothing benefits of unemployment insurance. *Journal of Public Economics*, *132*, 32–50.

Ebenstein, A., & Stange, K. (2010). Does inconvenience explain low take-up? evidence from unemployment insurance. *Journal of Policy Analysis and management*, *29*(1), 111–136.

Ellis, A., & Freeman, D. J. (2020). Revealing choice bracketing. *arXiv preprint arXiv:2006.14869*.

Evans, J. S. B. (2006). The heuristic-analytic theory of reasoning: Extension and evaluation. *Psychonomic Bulletin & Review*, *13*(3), 378–395.

Farber, H. S. (2015). Why you can't find a taxi in the rain and other labor supply lessons from cab drivers. *The Quarterly Journal of Economics*, *130*(4), 1975–2026.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, *10*(2), 171–178.
Fishman, M. E., Farrell, M., Gardiner, K. N., Barnow, B., & Trutko, J. (2003). *Unemployment insurance non-monetary policies and practices: How do they affect program participation? a study of 8 states* (Tech. Rep.). U.S. Department of Labor Employment and Training Administration.
Fiske, S. T., & Taylor, S. E. (2013). *Social cognition: From brains to culture.* Sage.
Friedman, M., et al. (1957). theory of the consumption function.
Frijda, N. H. (1986). *The emotions.* Cambridge University Press.
Gabaix, X. (2014). A sparsity-based model of bounded rationality. *Quarterly Journal of Economics*, *129*(4), 1661-1710.
Gabaix, X. (2017). *Behavioral inattention* (Tech. Rep.). National Bureau of Economic Research.
Ganong, P., & Noel, P. J. (2019). *Consumer spending during unemployment: Positive and normative implications* (Tech. Rep.). National Bureau of Economic Research.
Gathergood, J., Mahoney, N., Stewart, N., & Weber, J. (2019). How do individuals repay their debt? the balance-matching heuristic. *American Economic Review*, *109*(3), 844–75.
Gerard, F., & Naritomi, J. (2019). *Job displacement insurance and (the lack of) consumption-smoothing* (Tech. Rep.). National Bureau of Economic Research.
Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, *62*, 451–482.
Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *The American Economic Review*, *87*(1), 192.
Gruber, J. (2001). The wealth of the unemployed. *ILR Review*, *55*(1), 79–94.
Haggag, K., & Paci, G. (2014). Default tips. *American Economic Journal: Applied Economics*, *6*(3), 1–19.
Handel, B., & Schwartzstein, J. (2018). Frictions or mental gaps: What's behind the information we (don't) use and when do we care? *Journal of Economic Perspectives*, *32*(1), 155–78.
Harris, L. (1991). Stock price clustering and discreteness. *The Review of Financial Studies*, *4*(3), 389–415.
Hastings, J. S., Madrian, B. C., & Skimmyhorn, W. L. (2013). Financial literacy, financial education, and economic outcomes. *Annu. Rev. Econ.*, *5*(1), 347–373.
Herscovics, N., & Linchevski, L. (1994). A cognitive gap between arithmetic and algebra. *Educational studies in mathematics*, *27*(1), 59–78.
Hoover, H. (2019). Default tip suggestions in nyc taxi cabs. *Available at SSRN 3333460.*
Isaac, M. S., Wang, Y., & Schindler, R. M. (2020). The round-number advantage in consumer debt payoff. *Journal of Consumer Psychology.*
Jappelli, T., & Pistaferri, L. (2010). The consumption response to income changes.
Johnson-Laird, P. N. (2010). Mental models and human reasoning. *Proceedings of the National Academy of Sciences*, *107*(43), 18243–18250.
Kahneman, D. (1979). Prospect theory: An analysis of decisions under risk. *Econometrica*, *47*, 278.

Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *American economic review*, *93*(5), 1449–1475.

Kaplan, G., Violante, G. L., & Weidner, J. (2014). *The wealthy hand-to-mouth* (Tech. Rep.). National Bureau of Economic Research.

Karlan, D., McConnell, M., Mullainathan, S., & Zinman, J. (2016). Getting to the top of mind: How reminders increase saving. *Management Science*, *62*(12), 3393–3411.

Karlsson, N., Loewenstein, G., & Seppi, D. (2009). The ostrich effect: Selective attention to information. *Journal of Risk and uncertainty*, *38*(2), 95–115.

Keys, B. J., & Wang, J. (2018). Minimum payments and debt paydown in consumer credit cards. *Journal of Financial Economics*.

Kolsrud, J., Landais, C., Nilsson, P., & Spinnewijn, J. (2018). The optimal timing of unemployment benefits: Theory and evidence from sweden. *American Economic Review*, *108*(4-5), 985–1033.

Kőszegi, B., & Szeidl, A. (2012). A model of focusing in economic choice. *The Quarterly journal of economics*, *128*(1), 53–104.

Kroft, K., & Notowidigdo, M. J. (2016). Should unemployment insurance vary with the unemployment rate? theory and evidence. *The Review of Economic Studies*, *83*(3), 1092–1124.

Krueger, A. B., & Meyer, B. D. (2002). Labor supply effects of social insurance. *Handbook of public economics*, *4*, 2327–2392.

Kuka, E. (2020). Quantifying the benefits of social insurance: Unemployment insurance and health. *Review of Economics and Statistics*, *102*(3), 490–505.

LaLumia, S. (2013). The eitc, tax refunds, and unemployment spells. *American Economic Journal: Economic Policy*, *5*(2), 188–221.

Leary, J. B., & Wang, J. (2016). *Liquidity constraints and budgeting mistakes: Evidence from social security recipients* (Tech. Rep.). Working Paper.

Levin, I. P., Schneider, S. L., & Gaeth, G. J. (1998). All frames are not created equal: A typology and critical analysis of framing effects. *Organizational behavior and human decision processes*, *76*(2), 149–188.

Lindner, A., & Reizer, B. (2020). Front-loading the unemployment benefit: An empirical assessment. *American Economic Journal: Applied Economics*, *12*(3), 140–74.

Lusardi, A., Michaud, P.-C., & Mitchell, O. S. (2017). Optimal financial knowledge and wealth inequality. *Journal of Political Economy*, *125*(2), 431–477.

Lusardi, A., & Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *Journal of economic literature*, *52*(1), 5–44.

Lusardi, A., & Tufano, P. (2015). Debt literacy, financial experiences, and overindebtedness. *Journal of Pension Economics & Finance*, *14*(4), 332–368.

Lynn, M., Flynn, S. M., & Helion, C. (2013). Do consumers prefer round prices? evidence from pay-what-you-want decisions and self-pumped gasoline purchases. *Journal of Economic Psychology*, *36*, 96–102.

Meyer, B. (1990). Unemployment insurance and unemployment spells. *Econometrica*, *58*(4), 757–782.

Modigliani, F., & Brumberg, R. (1954). Utility analysis and the consumption function: An

interpretation of cross-section data. *Franco Modigliani*, *1*(1), 388–436.
Moffitt, R. (1985). Unemployment insurance and the distribution of unemployment spells. *Journal of Econometrics*, *28*(1), 85–101.
Mullainathan, S., & Shafir, E. (2013). *Scarcity: Why having too little means so much.* Macmillan.
Nisbett, R. E., & Ross, L. (1980). Human inference: Strategies and shortcomings of social judgment.
Olafsson, A., & Pagel, M. (2018). The liquid hand-to-mouth: Evidence from personal finance management software. *The Review of Financial Studies*, *31*(11), 4398–4446.
O'Leary, C. J. (2004). Ui work search rules and their effects on employment.
Ponce, A., Seira, E., & Zamarripa, G. (2017). Borrowing on the wrong credit card? evidence from mexico. *American Economic Review*, *107*(4), 1335–61.
Price, D. N. (1985). Unemployment insurance, then and now, 1935-85. *Soc. Sec. Bull.*, *48*, 22.
Read, D., Loewenstein, G., Rabin, M., Keren, G., & Laibson, D. (1999). Choice bracketing. In *Elicitation of preferences* (pp. 171–202). Springer.
Schindler, R. M., & Wiman, A. R. (1989). Effects of odd pricing on price recall. *Journal of Business Research*, *19*(3), 165–177.
Schmieder, J. F., & Von Wachter, T. (2016). The effects of unemployment insurance benefits: New evidence and interpretation. *Annual Review of Economics*, *8*, 547–581.
Schmieder, J. F., & von Wachter, T. (2017). A context-robust measure of the disincentive cost of unemployment insurance. *American Economic Review*, *107*(5), 343–48.
Schwartzstein, J. (2014). Selective attention and learning. *Journal of the European Economic Association*, *12*(6), 1423–1452.
Seira, E., Elizondo, A., & Laguna-Müggenburg, E. (2017). Are information disclosures effective? evidence from the credit card market. *American Economic Journal: Economic Policy*, *9*(1), 277–307.
Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, *50*(3), 665–690.
Soroka, S., Fournier, P., & Nir, L. (2019). Cross-national evidence of a negativity bias in psychophysiological reactions to news. *Proceedings of the National Academy of Sciences*, *116*(38), 18888–18892.
Stacey, K., & MacGregor, M. (1999). Learning the algebraic method of solving problems. *The Journal of Mathematical Behavior*, *18*(2), 149–167.
Stango, V., & Zinman, J. (2014). Limited and varying consumer attention: Evidence from shocks to the salience of bank overdraft fees. *The Review of Financial Studies*, *27*(4), 990–1030.
Stewart, N. (2009). The cost of anchoring on credit-card minimum repayments. *Psychological Science*, *20*(1), 39–41.
Sunstein, C. R. (2011). Empirically informed regulation. *The University of Chicago Law Review*, *78*(4), 1349–1429.
Taylor, S. E., & Thompson, S. C. (1982). Stalking the elusive" vividness" effect. *Psychological review*, *89*(2), 155.
Thakral, N., & Tô, L. (2019). *Tipping and the dynamics of social norms* (Tech. Rep.). Mimeo.

Thompson, V. A. (2009). Dual-process theories: A metacognitive perspective.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131. Retrieved from http://www.jstor.org/stable/1738360

Vellekoop, N. (2018). Explaining intra-monthly consumption patterns: The timing of income or the timing of consumption commitments?

Zhang, Y. (2017). Consumption responses to pay frequency: Evidence from "extra" paychecks. *ACR North American Advances.*