Experiments on Intertemporal Choice

by

Marissa Lepper

B.A. in Economics and Psychology, Trinity University, 2015

M.A. in Economics, University of Texas at Austin, 2017

Submitted to the Graduate Faculty of

the Dietrich School of Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2023

UNIVERSITY OF PITTSBURGH

DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Marissa Lepper

It was defended on

March 27, 2023

and approved by

Lise Vesterlund, Andrew W. Mellon Professor of Economics, University of Pittsburgh

David Huffman, Professor of Economics, University of Pittsburgh

Stephanie Wang, Associate Professor of Economics, University of Pittsburgh

Kirby Nielsen, Assistant Professor of Economics, California Institute of Technology

Experiments on Intertemporal Choice

Marissa Lepper, PhD

University of Pittsburgh, 2023

This dissertation consists of three essays that contribute to the field of behavioral economics by using experimental methods to explore the causes and impacts of impatient behavior. Chapter 1 identifies excuse-based procrastination. Excuses, or justifications, are a novel environmental factor that I show can induce myopic decisions in agents who would otherwise have been patient. In a lab experiment, participants allocate work between now and later. Some decisions have uncertainty over future work that can be costlessly resolved. I show that participants remain willfully ignorant as an excuse to do less immediate work, even at the risk of increasing total work. I propose that this is due to excuses mitigating a psychic cost associated with impatient behavior that makes procrastination less attractive, which I find suggestive evidence for in a survey. Chapter 2 uses a longitudinal online experiment with real-effort tasks to explore how streaks, i.e., tracking the consecutive periods a task is performed, can serve as a psychological motivator that decreases impatient behavior for tasks that have immediate costs but delayed benefits. However, there is a tradeoff with myopic reactions to broken streaks. Allowing for flexibility, such as "cheat days" when effort costs are higher than normal, can mitigate this by mechanically preserving a streak, but can be exploited to reduce effort on those days. Chapter 3 looks at how differential time horizons for debt repayment penalties impact the choices of financially distressed borrowers who are faced with the decision of what, not if, to default. Using observational credit report data, we find that a substantial subset of such borrowers who hold a varied debt portfolio avoid defaulting on revolving credit, resulting in an eventual default on their student loans. Although defaulting on student loans has much more severe penalties than defaulting on revolving credit, they are much more delayed. We explore how this timing influences financial decision making using an online survey.

Table of Contents

Preface . . . xi

1.0 Excuse-Based Procrastination . . . 1

1.1 Introduction . . . 1

1.2 Excuse-Based Procrastination . . . 4

1.3 Experimental Design . . . 7

1.3.1 Decisions . . . 7

1.3.2 Tasks . . . 9

1.3.3 Procedures . . . 10

1.4 Results . . . 11

1.4.1 Avoiding Information as an Excuse to Procrastinate . . . 12

1.4.2 Impact on Workload . . . 14

1.4.3 Impact of Excuses on Impatience . . . 17

1.4.4 Response to Incentives . . . 18

1.4.5 Types of Impatience . . . 21

1.5 Discussion . . . 22

2.0 Streaks, Cheat Days, and the Tradeoff between Motivation and Flexibility . . . 26

2.1 Introduction . . . 26

2.2 Motivational Evidence . . . 30

2.3 Experimental Design . . . 31

2.3.1 Experiment Task . . . 32

2.3.2 Treatments . . . 33

2.3.3 Hypotheses . . . 33

2.4 Results . . . 34

2.4.1 Impact of Cost Shocks . . . 37

2.4.2 Impact on Payment . . . 41

2.5 Discussion . . . 41

3.0 Time and Punishment: Penalty Timing and Cross-Debt Default Decisions . . 43
3.1 Introduction 43
3.2 Background and Definitions 45
3.3 Evidence from Individual Data 47
3.3.1 Data 47
3.3.2 Transition into Severe Delinquency 48
3.3.3 Event Studies 51
3.3.3.1 Credit Score 51
3.3.3.2 Debt-To-Income Ratio 53
3.3.3.3 Revolving Credit Utilization Rate 54
3.4 Survey 55
3.4.1 Part One - Characteristics of Debt 55
3.4.1.1 Decisions 55
3.4.1.2 Results and Discussion 56
3.4.2 Part Two - Financial Literacy 58
3.4.2.1 Decisions 58
3.4.2.2 Results and Discussion 59
3.4.3 Part Three - Intertemporal Risk Preferences 60
3.4.3.1 Decisions 61
3.4.3.2 Results and Discussion 62
3.5 Discussion 65
Appendix A. Excuse-Based Procrastination 67
A.1 Summary Statistics 67
A.1.1 Attrition 67
A.1.2 Gender 69
A.1.3 Task Performance 71
A.2 Analysis 71
A.2.1 Order 71
A.2.2 Impact of Work Schedules on Outcomes No Excuse Treatments 72
A.2.3 Heterogeneous Classifications 72

A.3 Instructions 74

Appendix B. Time and Punishment: Penalty Timing and Cross-Debt Default Decisions 75

B.1 Data Set 75

B.1.1 Sampling and Construction of Panel 75

B.1.2 Defining and Identifying Events 76

B.2 Additional Analysis 81

B.2.1 Life Cycle Differences 81

B.2.2 Summary Statistics 82

B.2.3 Event Study: Credit Score 82

B.2.4 Event Study: Debt to Payment Ratio 83

B.2.5 Payment-to-Income Ratio 84

B.3 Survey Instrucionts 85

Bibliography 86

List of Tables

Table 1: Experimental Design: Work Requirement of each Allocation Option by State 8

Table 2: Timeline of Both Experimental Sessions 11

Table 3: OLS Estimates of Excuses on Task Allocation 16

Table 4: Task Allocations in the Conflicted State for Each Work Schedule 19

Table 5: OLS Estimate of Work Schedule on Excuse Uptake and Procrastination . . 20

Table 6: Correlation Table of Survey Measures 23

Table 7: Streaks are Motivating Unless Broken 36

Table 8: Analysis of Behavior by Cost 38

Table 9: Average payment by treatment by if they experienced a cost shock 41

Table 10: Portfolio Composition One Quarter Before Event by Debt Type 50

Table A.1.1: Summary Statistics By Attrition Status 68

Table A.1.2: Summary Statistics by Gender 70

Table A.1.3: Average performance statistics excluding practice round 71

Table A.2.1: Order of Questions for Choice-To-Reveal 72

Table A.2.2: Impact of Work Schedule Conditions in the No-Excuse Treatment 73

Table A.2.3: Heterogeneous Classification by Decisions 74

Table B.1.1: Severe Delinquency, Default, and Collections Process for Student Loans . . 78

Table B.2.1: Summary Statistics One Quarter Before Event by Debt Type 82

Table B.2.2: Portfolio Composition One Quarter Before Event by Debt Type 82

List of Figures

Figure 1: Emotion When Doing Things Late Depending on the Presence of an Excuse 5
Figure 2: Self-Image of Doing Things on Time 6
Figure 3: Decision Screen in the Excuse Treatment 9
Figure 4: A Completed Transcription Task 10
Figure 5: State Contingent Decisions Compared to Information Acquisition 13
Figure 6: Percent Minimizing Current Work By Information Set 14
Figure 7: Frequency of Impatient Decisions in the Conflicted State by Treatment . . . 17
Figure 8: Impatient Decisions Split by Type 21
Figure 9: Subjects' Decision Screen 33
Figure 10: Work Conditional on Previous Day 35
Figure 11: Intensive Margin of Effort on High Cost Days 40
Figure 12: Transition into Severe Delinquency along Extensive Margin 49
Figure 13: Portfolio Composition of Debts in Repayment (%) 50
Figure 14: Overlaid Coefficient Plot: Credit Score 52
Figure 15: Overlaid Coefficient Plot: Debt-to-Income Ratio over All Open Accounts . . 53
Figure 16: Overlaid Coefficient Plot: Revolving Utilization Rate 54
Figure 17: Ranking of Debt Characteristics in Default Decisions by Importance 57
Figure 18: Comparison of Financial Literacy Across Debt Types 59
Figure 19: Distribution of Correct Answers by Type of Debt 60
Figure 20: Histogram of Participants' Indifferent Point δ 61
Figure 21: Scatterplot of the Row Participants Switch At in Price List Two and Three 63
Figure 22: Expected Value of Lotteries which are Indifferent to \$1 Now For Sure 64
Figure B.1.1: Age Distribution by Subsample in the First Quarter of 2004 76
Figure B.1.2: Presence of Outstanding Governmental Agency Debt 79
Figure B.1.3: Amount Past Due 80
Figure B.1.4: Balance on Collections Accounts 80

Figure B.2.1: Transition into Severe Delinquency along Extensive Margin 81
Figure B.2.2: Overlaid Coefficient Plot: Payment-to-Income Ratios over All Open Accounts 84

Preface

Acknowledging everyone who has been a part of my research journey would be a thesis in itself, so I first want to thank **everyone** who has supported and encouraged me over the past few years.

I am immensely grateful to my academic mentors, particularly my advisor Lise Vesterlund, whose expertise, insights, and guidance have been instrumental in my academic success and helped me become a better economist than I ever hoped to be. Thank you, Lise, for your willingness to go above and beyond for me, for hopping on calls from all around the world, for providing me with untold opportunities and support, and for your true, unwavering belief in me. I am also incredibly thankful to Kirby Nielsen for her invaluable mentorship, feedback, and friendship, and to David Huffman, Alistair Wilson, Stephanie Wang, David Danz, the Cohen Fellows, and everyone else who has helped me with my research.

My family's unconditional love and support has been a constant source of motivation. Thank you to my mom, Danielle Lepper, my dad, John Lepper, and my sister, Jess Lepper, for celebrating my ups and helping me through my downs, for inspiring me, and for never doubting me. I also need to thank Socks and Sissi, who are no longer with us – and would not be able to read if they were – but who were a huge source of comfort for many years.

Finally, I would like to acknowledge the support of my friends throughout graduate school. Sierra Arnold, thank you for always being there for me and being such a caring person. David Agorastos, your constant companionship has made a huge difference in my life. Pun Winachukul, thank you for your support as my officemate and co-RA and continued support even after you've graduated. I also need to thank Juan Campenero, Maddie Arnsbarger, Priyoma Mustafi, Kelly Hyde, Neeraja Gupta, and Mallory Avery for all of their help and encouragement on this journey. The countless hours we've spent in the office, on Zoom, texting, and at happy hour have helped me grow as a person and a researcher, and I am thankful I have made lifelong friends. And finally, I'd like to thank my non-grad school friends, especially Berit Backes and Joanna Ahn, who have spent many hours cheering me on, listening to me vent, and reminding me that it's okay to have fun.

1.0 Excuse-Based Procrastination

Individuals are frequently tempted to delay tasks to the detriment of their well-being – even when they know they shouldn't. I show that costly procrastination increases when excuses, or justifications, are present. I propose this is driven by excuses attenuating a psychic cost associated with procrastinating, which I find suggestive evidence for using a survey. I explore the behavioral impact of having excuses available using a within-subject longitudinal laboratory experiment. Participants select bundles of work over two different periods. Some decisions have uncertainty over future work that can be costlessly resolved. I show that participants remain willfully ignorant as an excuse to select options with less immediate work, even at the risk of increasing total work. While 91 percent of participants minimize overall work when given full information, 37 percent avoid the information needed to do so when given the chance. The presence of an optional excuse increases procrastination fourfold and overall workload by 10 percent. Thus, the presence of an excuse has important implications for inference of individual time preferences. These results also help explain long-standing puzzles, such as why individuals do not learn from their past impatient mistakes.

1.1 Introduction

Tradeoffs between current and future utility often create tension between what someone *wants* to do and what they *should* do. Indeed, we've all tried, with varying levels of success, to go to bed at a reasonable time, save more for retirement, or finish a paper instead of watching television. Individuals who prioritize their current selves in these situations will act impatiently, often to the detriment of their overall well-being. One such impatient behavior is procrastination, where individuals delay tasks to a point that they themselves view as suboptimal (Ericson and Laibson, 2019). Procrastination has been argued to be costly across a breadth of labor, health, and financial outcomes, underscoring the importance of understanding how to avoid these costly delays.[1]

[1] By definition, procrastination is costly from the individual's perspective. Additionally, for papers documenting the cost of procrastination, see Ariely and Wertenbroch (2002a); Ashraf et al. (2006); Augenblick et al. (2015); Bénabou and Tirole (2004a); Carroll et al. (2009); Carvalho et al. (2016a); DellaVigna and Malmendier (2004, 2006); Duflo et al. (2011); Ericson (2017); Ericson and Laibson (2019); Giné et al. (2010); Heidhues and Kőszegi (2010); Heidhues and Strack (2021); Kaur et al. (2015); O'Donoghue and Rabin (1999).

I identify a novel environmental factor that can influence whether an individual procrastinates in a given decision: the presence of excuses. I argue that the availability of an optional excuse, or justification, for impatient behavior can lead to *excuse-based procrastination* in individuals who would otherwise have been patient. I propose this pattern of behavior, i.e., appearing patient absent an excuse and impatient otherwise, can be explained by excuses attenuating a psychic cost or feeling of guilt associated with procrastinating. The logic underlying this mechanism is intuitive: People feel bad when they make decisions they themselves believe to be suboptimal – unless they can justify it to themselves. This cost is novel and not found in existing models of intertemporal choice, but can play a role in determining whether or not someone procrastinates. Furthermore, excuses represent a novel policy intervention, as targeting excuse-based procrastination may be best addressed by removing access to excuses.

My study consists of an online survey and a longitudinal laboratory experiment. My survey provides suggestive evidence for an emotional cost of procrastination and that this cost attenuates in the presence of excuses. My experiment shows that the availability of an optional excuse increases impatient behavior. My within-subject design allows me to classify impatient behavior that only happens when an excuse is present as procrastination. The specific excuse is the ability to remain willfully ignorant of the future consequences of one's actions.[2] Despite being a direct result of their own choices, this ignorance can be used to justify decisions that reduce current costs, even at the risk of increasing future costs.

Participants in my laboratory experiment complete real-effort tasks in two periods, separated by a week, for a fixed payment. In the first session, participants choose between bundles that correspond to work now (i.e., tasks to be completed that session) and work later (i.e., tasks to be completed one week later during the second session). Participants make three types of decisions. Two are made in the *No-Excuse* treatment with full information, each of in a different state of the world. In one state, the same option minimizes work in both weeks. In contrast, in the other state, minimizing current work comes at the direct expense of increasing future work – and overall work.[3] The other decision is in the *Excuse* treatment, where I grant participants discretion over their information set about the underlying state. Participants don't know which state they are in

[2] My design adapts the Dana et al. (2007) moral wiggle room framework to the intertemporal domain.

[3] In one state, participants decide between 10 tasks now and 13 later, or 13 tasks now and 20 later. In the other state, they decide between 10 tasks now and 20 later or 13 tasks in both weeks.

and therefore face uncertainty over the mapping between current and future work. However, they can resolve this uncertainty by clicking a button before making their decision if they want to do so.

The results of my experiment show that individuals avoid information as an excuse to procrastinate. While 91 percent of decisions made with full information minimize overall workload, 37 percent of participants avoid the information needed to do so when given the chance. This leads to a 27 percentage point increase in impatient decisions, i.e., selecting bundles that minimize current work at the expense of overall workload, in the Excuse treatment. This four-fold increase in impatient behavior comes from participants only being willing to delay tasks at a cost to their future self if an excuse is available, which I classify as excuse-based procrastination. While excuse-based procrastination reduces immediate work, it substantially and significantly increases both future and total workload. Finally, I vary the work schedules in auxiliary treatments and find that participants respond predictably to changes in the incentives to procrastinate.[4]

Understanding the role of excuses in the intertemporal choice domain has both theoretical and empirical contributions. Excuse-based procrastination, i.e., acting patient only when an excuse is present, is not predicted by current intertemporal choice models. However, this pattern of behavior can be derived by adding an excuse-attenuated emotional cost of procrastination to various models, such as temptation and self-control (Gul and Pesendorfer, 2001), time-inconsistent preferences (Laibson, 1997; O'Donoghue and Rabin, 1999; Strotz, 1955), or dual-self models (Fudenberg and Levine, 2006; Shefrin and Thaler, 1981).[5]

Empirically, excuse-based procrastination has important implications for the interpretation of revealed time preferences (Andreoni et al., 2015; Augenblick et al., 2015; Carvalho et al., 2016a; Cohen et al., 2020; Falk et al., 2018; Frederick et al., 2002; Heidhues and Strack, 2021). I find that the same individual appears significantly less patient when an excuse is present. Moreover, excuse-based procrastination accounts for a substantial share of impatient behavior observed in my study, with nearly 70 percent of impatient decisions I observe being made in the presence of an excuse. This suggests that failing to take the supply of excuses available to a decision-maker into account can lead to disparate and misspecified measures of patience.

Excuse-based procrastination can help explain persistent naïveté, a long-standing puzzle (Ali,

[4]Zimmermann (2020) finds that other forms of motivated beliefs also respond to incentives.

[5]Costly self-control, like that found in Gul and Pesendorfer (2001), may exacerbate the demand for excuses as a form of emotional regulation, akin to that found in charitable giving (Andreoni et al., 2017; DellaVigna et al., 2012).

2011; Augenblick and Rabin, 2018; DellaVigna and Malmendier, 2006; Ericson and Laibson, 2019; John, 2019; LeYaouanq and Schwardmann, 2022). Justifications for impatient behavior, even self-created ones, attenuate the link between actions and consequences. This may result in agents failing to update their beliefs and continuing to make myopic choices even when faced with the same decisions. Therefore, excuse-based procrastination can inhibit the ability of individuals to learn from the past, leading to short-sightedness that is not tempered by past myopic "mistakes."

Finally, studying excuses in the intertemporal choice domain contributes to our understanding of how excuses and self-deception impact decision-making. I show that excuses trigger impatience and lead to procrastination when choices differentially impact current and future versions of oneself, akin to how they increase selfishness in choices between oneself and others (Bartling et al., 2014; Bicchieri et al., 2022; Dana et al., 2007; Exley, 2015; Exley and Kessler, 2019; Gino et al., 2016; Grossman and van der Weele, 2016; Linardi and McConnell, 2011; Serra-Garcia and Szech, 2022). However, in stark contrast to the less salient costs resulting from excuse-based selfishness, decision-makers experience the consequences of their excuse-based procrastination first-hand. This shows that self-deception is not only the basis of self-serving biases (DiTella et al., 2015; Gneezy et al., 2019; Mijović-Prelec and Prelec, 2010; Saccardo and Serra-Garcia, 2022), but can also be used to justify *present*-self-interested decisions.

The rest of the chapter is organized as follows: Section 1.2 discusses my proposed mechanism and presents the results of my survey. I describe my experimental design in Section 1.3 and the results of my laboratory experiment in Section 1.4. Section 1.5 concludes.

1.2 Excuse-Based Procrastination

This chapter explores the role of excuses in procrastination by introducing excuse-based procrastination. Importantly, excuse-based procrastination is not an ex-post justification of myopic behavior, but rather context-dependent procrastination that happens only when an excuse is available. This means that someone who procrastinates in one decision-making environment won't do so in another that is identical, minus an optional excuse.

I propose that excuse-based procrastination results from an immediately-felt emotional cost of procrastinating that is attenuated by excuses. In other words, people feel bad when they procrastinate and, absent an excuse, this psychic cost aligns current and long-run incentives. However, the ability to justify an action – for example, by remaining willfully ignorant of the future consequences – alleviates this cost and makes procrastination more appealing. Both the psychic costs of procrastinating and excuses will therefore play into the calculus of when an individual procrastinates.

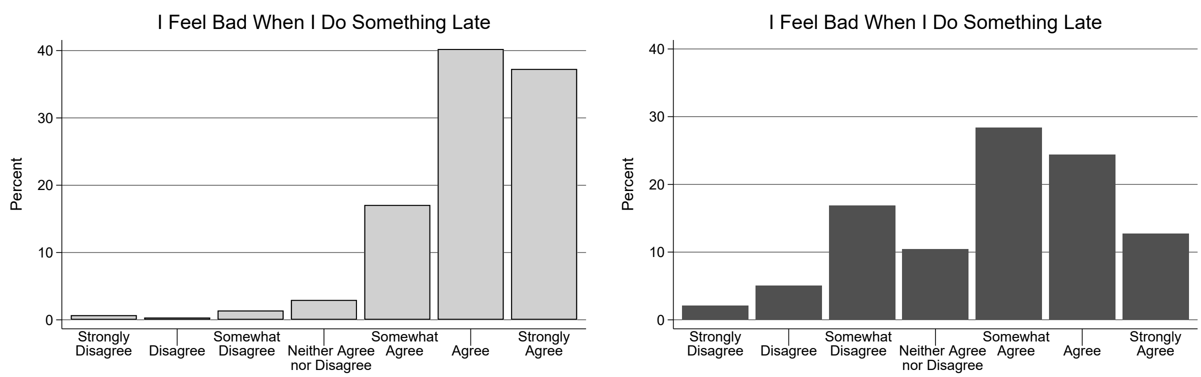

(a) Without An Excuse

(b) With An Excuse

Figure 1: Emotion When Doing Things Late Depending on the Presence of an Excuse

I use a survey ($N = 600$) run on Prolific to evaluate the emotional cost of procrastination that is attenuated by excuses. The survey takes four minutes on average and participants receive a \$2 flat payment. Participants answer a series of questions about doing things "on time/late," a proxy for procrastination.[6] I assess this cost by asking participants how much they agree with the statement "I feel bad when I do something late" on a seven-point likert scale.[7] The light gray bars in panel 1a show the right-skewed distribution of participant answers. Almost 95 percent of participants answering above "Somewhat Agree" and over two-thirds answering either "Agree" or "Strongly Agree." The dark bars in panel 1b show the histogram of agreement to the same statement except with the caveat of having an excuse.[8] Excuses appear to reduce the emotional

[6]This was intentionally vague to capture the general feeling associated with delaying tasks. However, participants were given examples like paying bills late, turning in work or school assignments late, arriving late to meetings, etc.

[7]Kaiser and Oswald (2022) show that self-report surveys are a reliable measurement of feelings.

[8]Exact wording: "I don't feel bad when I do something late if I have an excuse as to why I did not do it on time."

cost of doing something late, with those agreeing or strongly agreeing, the two strongest options, dropping to one-third. Moreover, the No-Excuse distribution first order stochastically dominates the Excuse distribution (Somer's $d = 0.53$, $p < 0.001$).

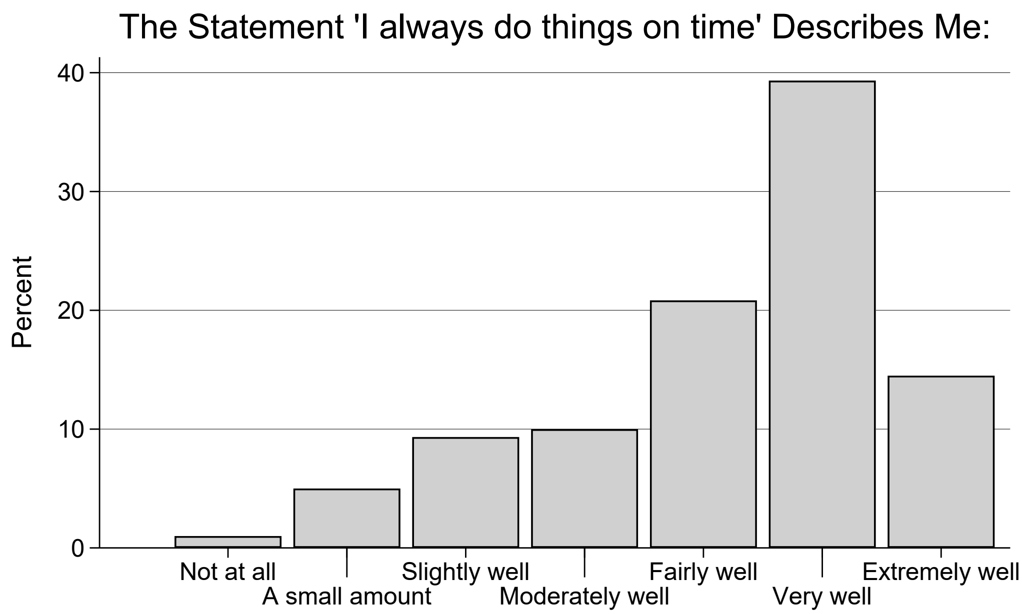

Figure 2: Self-Image of Doing Things on Time

This psychic cost resembles classes of models where non-consumption terms, such as self-image, enter into agents' utility functions Bénabou and Tirole, 2016; Köszegi, 2006.[9] Figure 2 shows a histogram of participants' responses to how well the statement "I always do things on time" describes them on a seven-point likert scale. This self-report data provides valuable insight into how participants view themselves, with the right-skewing distribution suggesting that many think of themselves as being on time. In fact, most participants select either "Very Well" or "Extremely Well." An individual who thinks of themselves as adept at not delaying tasks may then experience an emotional cost when they act in a way contrary to their positive self-image.

[9] In general, this can be thought of as an internalized social norm as there has been an overwhelmingly positive view of temperance throughout society, both historically and currently. Plato and Aristotle viewed self-restraint as one of the core ideals a person could hold (Aristotle et al., 2011 Edition; Plato et al., 2016 Edition), while Adam Smith argued that "self-command" is the foundation of all virtue (Smith, 1759). Moreover, some form of self-control is viewed as righteous in Buddhism, Christianity, Hinduism, Judaism, and many other religions. See Akerlof and Kranton (2000); Bénabou and Tirole (2011) for how norms can influence decision making.

Overall, my survey provides suggestive evidence about the relationship between excuses and procrastination. I further explore this using a laboratory experiment that manipulates the availability of an excuse to study their impact on behavior.

1.3 Experimental Design

Isolating excuse-based procrastination requires a longitudinal experiment with incentivized decisions over an intertemporally non-fungible good and treatments that alter the availability of an "excuse". I fulfill these characteristics by running a laboratory experiment where participants allocate real effort tasks across two time periods for a fixed payment.[10] I introduce the availability of an excuse by varying participants' abilities to control their knowledge of how many tasks they will complete in the future for a given decision.

1.3.1 Decisions

In a 3x4 within-subject design, participants make three work allocation decisions in each of the four work schedule conditions. I first discuss the *Baseline* work schedule condition and then present the parameters and results of the other three conditions in Section 1.4.4.[11]

Work decisions involve choices between two bundles, Option X and Option Y, which split across Week One (the current session) and Week Two (a future session). In all decisions, selecting Option X means completing 10 tasks in Week One, while Option Y means completing 13. I therefore will refer to Option X as *Less Now* and Option Y as *More Now* for the remainder of the paper. For both options, the potential Week Two allocations are 13 or 20 tasks. While *Less Now* always minimizes current effort, the mapping between the two options and future work varies.

Participants make decisions in two states of the world, both shown in Table 1. In the first state of the world (Table 2a), the same option minimizes both current and overall work. In this state, *Less Now* leads to fewer tasks in both weeks (10 in Week One and 13 in Week Two, compared to

[10]Using real-effort tasks reduces potential confounds, such as corner solutions that arise due to the ability to directly substitute across time, transaction costs and beliefs over payment reliability. For further discussion, see Andreoni and Sprenger (2012a); Augenblick et al. (2015); Augenblick and Rabin (2018).

[11]The remaining work schedule conditions differ from *Baseline* only in the number of tasks.

the 13 and 20 allocations of *More Now*) and 10 fewer tasks overall. Since working less now aligns with working less overall, I call this the Aligned state. In the second state (Table 2b), in contrast, working less now has the direct expense of working more overall. Selecting *Less Now* leads to three fewer tasks Week One (10 instead of 13), but seven additional tasks in Week Two (20 instead of 13) and four additional tasks overall (30 instead of 26). Although picking *Less Now* increases overall effort for the same payment, participants may be tempted to pick it to delay their work.[12] This creates an intertemporal tension between the current and future self and produces the opportunity for impatient behavior not found in the Aligned state. I therefore call this the Conflicted state.

Table 1: Experimental Design: Work Requirement of each Allocation Option by State

	Week One	Week Two
Less Now	10 Tasks	13 Tasks
More Now	13 Tasks	20 Tasks

(a) Aligned State

	Week One	Week Two
Less Now	10 Tasks	20 Tasks
More Now	13 Tasks	13 Tasks

(b) Conflicted State

Two treatments – The No-Excuse and the Excuse treatments – encompass the three work allocation decisions. Two decisions, one in each of the two states, are part of the No-Excuse treatment where full information is exogenously given. The third decision constitutes the Excuse treatment. In this decision, as shown in Figure 3, participants can maintain uncertainty over the state of the world, and therefore which option minimizes future – and total – work. Each state is equally likely, and the uncertainty can be costlessly resolved by pressing a button. If a participant's preferences differ by state -– such as wanting *Less Now* in the Aligned state and *More Now* in the Conflicted state – they should simply click the button, learn the state, and incorporate this knowledge into their decision. However, participants who would otherwise act patiently may avoid the information needed to do so and instead minimize Week-One work. This pattern of behavior represents an excuse-taker, or someone who is impatient only when an excuse is present.

[12]Sessions ran early in the semester to reduce differences in marginal effort costs. Over 90 percent of participants did not differ in self-reported busyness across the two weeks, and any differences would impact both treatments.

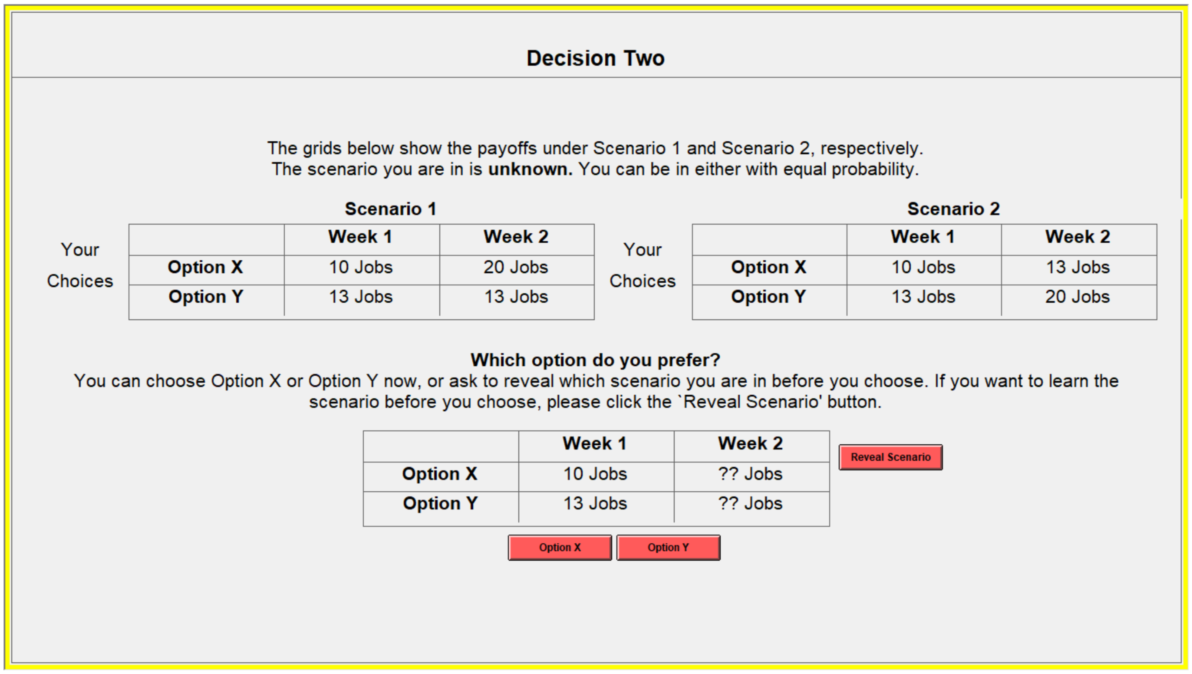

Figure 3: Decision Screen in the Excuse Treatment

The state realization in the Excuse treatment varies at the individual level and is assigned ex-ante, allowing me to know the incentives faced by each participant even if they themselves do not. Through this, I can identify not only excuse takers – acting patient in the Conflicted state in the No-Excuse treatment but picking *Less Now* in the Excuse treatment without revealing the state – but also excuse-based procrastination itself. Excuse-based procrastination will appear as an apparent preference reversal in decisions made in the same exact state, and therefore with the same effort incentives, but with differing availability of excuses. Comparing levels of observed impatience in the Conflicted state between the No-Excuse and Excuse treatments identifies a causal increase in procrastination due to the presence of an optional excuse.

1.3.2 Tasks

In both weeks, participants are given the task to transcribe rows of 35 blurry Greek letters, as shown in Figure 4. For each row, they must press buttons that corresponded with the Greek letters listed.[13] The submit button appears after participants enter 35 letters. If the entered letters are less than 80 percent accurate (i.e., fewer than 28 are correct), a warning appears and participants can correct their transcription by either deleting individual letters from the end or starting over.[14]

Figure 4: A Completed Transcription Task

1.3.3 Procedures

Four groups of sessions, each comprising of two sessions exactly one week apart, were conducted in person at the Pittsburgh Experimental Economics Laboratory (PEEL) using a gender-balanced undergraduate population. Participants received \$6 and \$30 after the two sessions, with the larger payment serving to minimize attrition. Out of the original 78 participants who began the study, 72 returned the following week.[15] I exclude those who do not return in my analysis, but all results are robust to including them.[16]

[13]I adapt Augenblick et al. (2015); Augenblick and Rabin (2018)'s task to zTree (Fischbacher, 2007)

[14]The average accuracy is 98 percent for Week One and 99 percent for Week Two.

[15]The sample size was selected to achieve 90 percent power based on ex-ante power calculations using data from Augenblick et al. (2015); Dana et al. (2007) and PEEL attrition data.

[16]Importantly, there is no difference in task performance or assigned Week-Two tasks.

In Week One, participants receive detailed instructions and answer comprehension questions. They next complete a practice round consisting of transcribing four tasks to reduce projection bias and ensure awareness of effort costs (Augenblick and Rabin, 2018) prior to making their 12 allocation decisions, one of which is implemented.[17] Participants learn their Week-One task allocation and complete them. They move through the study at their own pace and leave once their own tasks are finished, preventing boredom stemming from waiting from influencing decisions. To obfuscate social signals that may arise through the timing of departure (Bénabou and Tirole, 2006), the exact decision that is implemented is randomized at the individual level. Participants learn their Week-Two task allocation directly before leaving. This small temporal distance between choices and knowledge of future consequences allows time for participants to reap the benefits of excused myopia while still limiting the role of intertemporal information preferences or ambiguity aversion.

To avoid participants aiming to minimize their time in the lab instead of the number of tasks, both sessions are designed to be as similar as possible. In Week Two, instructions are read out loud, participants complete another practice round and make risk preference elicitation decisions, and then complete their remaining tasks.[18] Participants move through the study at their own pace and leave after finishing their own tasks.

Table 2: Timeline of Both Experimental Sessions

	Instructions	Practice	Decisions	Work	Payment
Week 1	✓	✓	Main	✓	\$6
Week 2	✓	✓	Bonus	✓	\$30

1.4 Results

I start by examining if participants avoid information as an excuse to procrastinate in Section 1.4.1. In doing so, I identify excuse-takers – those who minimize overall workload with full information but then don't reveal the state and minimize current work when given the chance. Next, I

[17] Participants make all Excuse treatment decisions first, with decision order randomized within each treatment.
[18] Each session, two subjects receive their decision fo one row of a list comparing \$10 gambles to a certain \$5.

quantify how excuses impact outcomes by taking an omniscient view of the underlying state to look at how excuses impact workload in Section 1.4.2 and impatience in Section 1.4.3. In Section 1.4.4, I explore how excuse uptake responds to the incentives to procrastinate, namely an increase in current benefits or a decrease in future costs of delaying tasks. Finally, in Section 1.4.5, I show that excuse-based procrastination accounts for most impatient behavior throughout the entire study.

1.4.1 Avoiding Information as an Excuse to Procrastinate

The first step in identifying excuse-based procrastination is looking at information acquisition decisions. Standard models of economic theory predict demand for information with instrumental value(Stigler, 1961). Following this, the share of participants who costlessly reveal the state should be at least as high as – if not higher than – the share of participants who make choices that vary between the two states in the No-Excuse treatment.[19] Figure 5 shows clear evidence of information avoidance as a significantly larger share make state contingent decisions in the No-Excuse treatment (light grey bar) than opt to learn the underlying state in the Excuse treatment (dark grey bar). Consistent with the predictions of excuse-based procrastination, we see the vast majority of participants make state-contingent decisions (95.8 percent), yet only 62.5 percent reveal the state ($t = 5.95$, $p < 0.01$).[20]

I next explore allocation decisions to see whether information avoidance increases procrastination. Figure 6 shows the percent choosing the *Less Now* work allocation option conditional on their information set.[21] Decisions in the four left-most bars are made with full information – whether chosen in the Excuse treatment (dark grey bars) or exogenously given in the No-Excuse treatment (light grey bars) – while the furthest right one shows decisions made with uncertainty pooled across the underlying state. Excuse-based procrastination will appear as participants minimizing overall work with full information, i.e., picking *Less Now* in the known Aligned states and *More Now* in the known Conflicted states, but minimizing current work when they don't know the state.

When participants know they are in the Aligned state (the two right-most bars), they overwhelmingly choose to minimize both current and overall workload, with 90.3 percent picking *Less*

[19]Individuals can have positive demand for information without instrumental value (Eliaz and Schotter, 2010).
[20]This amount is all state-contingent decisions, not just the ones that minimize overall work.
[21]In Week One, *Less Now* corresponds to 10 tasks and *More Now* to 13. *Less Now* leads to 13 Week-Two tasks in the Aligned state and 20 in the Conflicted state, the reverse of the *More Now* Week-Two allocations.

Now. Participants who know there are in the Conflicted state (the center two bars), on the other hand, face a decision of doing less overall but at the expense of doing more now. However, participants continue to prioritize minimizing total work, with only 6.2 percent picking *Less Now*. In total, 92 percent of decisions made with full information minimize overall workload. Additionally, participants are equally likely to minimize overall work in both treatments, meaning behavior does not depend on if information is exogenously or endogenously revealed.[22] The presence of an excuse, therefore, does not lead to impatient behavior unless a participant chooses to take it.

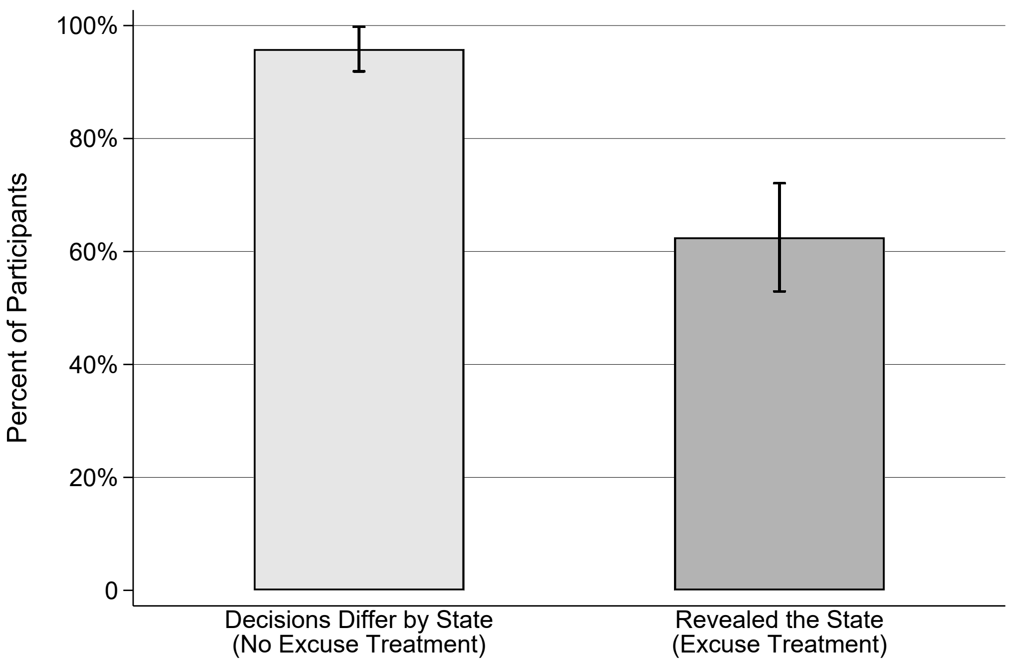

Figure 5: State Contingent Decisions Compared to Information Acquisition

While decisions made with full information demonstrate a high degree of patience in the face of a costly task, decisions made under uncertainty tell a different story. Two-thirds of participants without information pick *Less Now*, minimizing current work without knowing its impact on future work. 76.2 percent of the participants who always minimized overall workload in the No-Excuse treatment but didn't reveal the state select *Less Now*. I classify these participants – who account

[22]Participants in the Excuse treatment who reveal the state minimize overall work marginally more (95.5 percent) than participants in the No-Excuse treatment (90.9, $t = 0.990$, $p = 0.167$). However, comparing decisions across treatments made by decisions who revealed the information removes any differences.

for about 22 percent of my sample – as **excuse-takers**.[23] They appear impatient (patient) when there is (not) an excuse available, demonstrating the importance of understanding the supply of excuses within a decision-making environment when evaluating time preferences.[24]

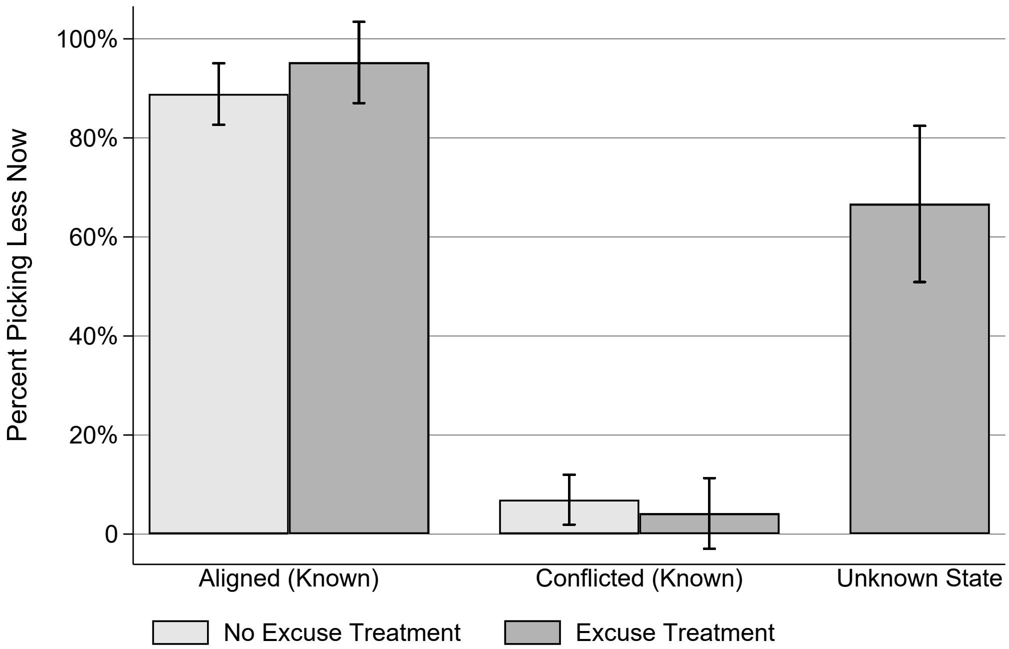

Figure 6: Percent Minimizing Current Work By Information Set

1.4.2 Impact on Workload

I now take an omniscient view of the underlying state, assigned ex-ante prior to the decision, while remaining agnostic towards information sets. This means I know the number of tasks associated with each decision, even if a participant does not. I examine the costs of excuse-based procrastination by exploiting my within-subject design to compare how allocations differ by whether or not an excuse is present. I can therefore go beyond identifying an apparent preference reversal not previously predicted by economic theory by quantifying the costs associated with it.

My analysis centers on the excuse-takers I identify in Section 1.4.1. These participants engage in excuse-based procrastination by minimizing overall work in the No-Excuse treatment – selecting

[23] Excuse-takers do not differ in terms of gender, self-reported busyness, or risk-preferences.
[24] Further typology, including classifications of patient and myopic types, is in the appendix.

Less Now in the Aligned state but *More Now* in the Conflicted state – but avoiding the information in the Excuse treatment and then picking *Less Now*, thereby potentially procrastinating.

Table 3 presents the results of the following specification

$$Y_{id} = \beta_0 + \beta_1 * ExcuseTreatment_d + \beta_2 * ExcuseTaker_i + \beta_3 * Interaction_{id} + \epsilon_i$$

where individual i's chosen workload Y_{id} in decision d is regressed on three indicator variables: *Excuse Treatment*, for decisions made in the presence of an excuse (whether taken or not), *Excuse Taker*, for decisions made by those who take excuses (whether one is present or not), and their interaction – decisions made by excuse-takers when an excuse is available, i.e., excuse-based procrastination.

Column 1 looks at the total workload, or number of tasks across both weeks, associated with a decision. The constant shows that non-excuse-taker participants in the No-Excuse treatment complete, on average, 25.24 tasks across both weeks. The null results for the *Excuse Taker* and *Excuse Treatment* coefficients confirm the results presented in Section 1.4.1. First, the presence of an excuse does not impact the behavior of those who do not take it. In the Excuse treatment, participants who are not excuse-takers – including those who revealed the state and those who picked *More Now* after avoiding the information – only allocate, on average, an additional half task compared to the No-Excuse treatment. Second, taking an excuse does not impact decisions made when an excuse is not available. On average, excuse-takers allocation 0.185 fewer tasks than those who don't take excuses in the No-Excuse treatment.

Looking at the interaction term – the task allocation amounts of excuse-takers in the Excuse treatment – allows me to quantify the cost of excuses. When No-Excuse is available, excuse-takers have an average total workload of 25.43 tasks (the sum of the constant and the *Excuse Taker* coefficient). When an optional excuse becomes available in the Excuse treatment, however, total workload for these same participants increases by 3 tasks on average (the sum of the *Excuse Treatment* coefficient and the *Interaction* term). This amounts to over a 10 percent increase in workload that could have been avoided by a simple press of a button to reveal the state.

Columns 2 and 3 split up workload by week. Similar to column 1, there is no significant difference in average workload either week when comparing non-excuse-takers to themselves when an excuse is and is not present (the *Excuse Treatment* coefficients) and when comparing them to

Table 3: OLS Estimates of Excuses on Task Allocation

	(1) Total Work	(2) Week One	(3) Week Two
Excuse Treatment	0.500	0.306	0.194
	(0.430)	(0.201)	(0.353)
Excuse Taker	-0.185	0.139	-0.324
	(0.475)	(0.116)	(0.388)
Interaction	2.500**	-1.972***	4.472***
	(0.993)	(0.230)	(0.908)
Constant	25.24***	11.53***	13.71***
	(0.291)	(0.0279)	(0.285)
N	216	216	216

Note: Standard errors in parentheses clustered at the individual level. Excuse Taker indicators decisions made by participants who minimize overall workload in the No-Excuse treatment but minimize current work after avoiding information in the Excuse treatment. Excuse Treatment indicates decisions made with an excuse present. The interaction term represents decisions made by excuse-takers in the Excuse treatment. Outcome variables are, in order of columns, the overall workload associated with a given decision and the workload split across Week One and Week Two.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

excuse-takers when there is No-Excuse present (the *Excuse Taker* coefficients). Additionally, when No-Excuse is present, there is no significant difference between the average workloads for Week One (11.53 tasks) and Week Two (13.71 tasks, $\chi^2 = 61.39$, $p < 0.001$), showing that participants smooth their effort across the two sessions.

Column 2 shows that taking an excuse reduces current workload by an average of 1.67 tasks (the sum of the *Excuse Treatment* and interaction coefficients). This is an almost 15 percent reduction from the 11.67 Week-One tasks excuse-takers allocate on average in the No-Excuse treatment (the summer of the constant + *Excuse Taker* coefficient). However, this has the steep tradeoff of increasing their average Week-Two workload by one-third. The implied intertemporal rate of substitution of excuse-takers increases from 0.14 in the No-Excuse treatment to 0.8 in the Excuse treatment – an increase of 471 percent.[25] This means that the same group of participants appear very patient when No-Excuse is present but very impatient when one becomes available, leading to disparate estimates of their time preferences depending on the decision making environment.

1.4.3 Impact of Excuses on Impatience

The previous analysis shows that a quarter of participants are classified as excuse-takers and use information avoidance as an excuse to prioritize their current self, leading to a substantial increase in future and overall workload on average. However, excuse-takers may be in the Aligned state without knowing it, where picking *Less Now* is what they would do if they knew the state. Therefore, to fully explore how the presence of excuses impacts *impatience*, I now restrict my analysis to decisions made in the Conflicted state while abstracting away from information sets.

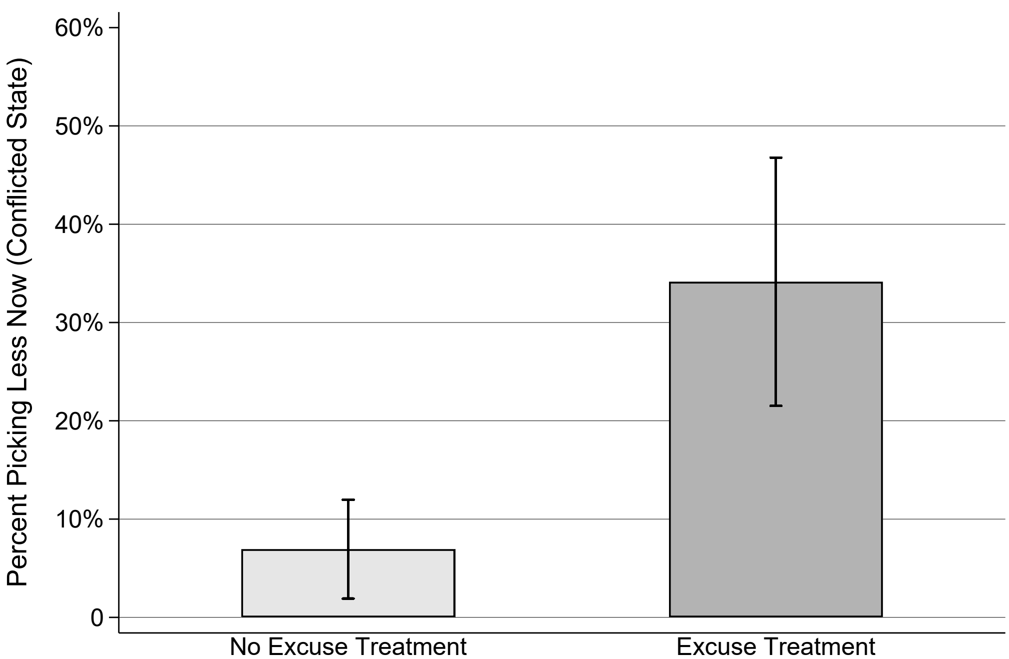

Figure 7: Frequency of Impatient Decisions in the Conflicted State by Treatment

[25]In comparison, non-excuse-takers implied intertemporal rate of substitution is 0.19 in the No-Excuse treatment and 0.17 in the Excuse treatment.

In the Conflicted state, participants, whether they know it or not, face a direct tradeoff between current and overall work. I classify selecting *Less Now* $= (10, 20)$ over *More Now* $= (13, 13)$, or picking to do three fewer tasks immediately at the expense of doing seven additional tasks the following week, as the impatient choice. Figure 7 looks at all decisions made in the Conflicted state across my experiment and compares the prevalence of selecting *Less Now* depending on if an excuse is (dark grey bar) or is not (light grey bar) available. Since I am looking at decisions made by the same participants facing identical incentives minus the optional excuse, I can not only directly identify how the presence of an excuse impacts impatience, but also how it leads to procrastination, or the delaying of tasks at a cost an individual themselves feels is suboptimal.

In the No-Excuse treatment, only 6.9 percent of participants knowingly act impatient, demonstrated by the right-most bar in Figure 7.[26] Without excuse-based procrastination, the Excuse treatment should have a similarly low level of impatience. However, impatience increases by 394.2 percent when an excuse becomes available (t=3.93, $p < 0.01$). This represents procrastination, as the individual acting impatiently would not have done so if they knew the incentives they face.

1.4.4 Response to Incentives

Three additional treatments vary the work schedules to test how different costs and benefits (i.e., different amounts of work now and/or later) impact the prevalence of excuse-based procrastination.[27] All work schedule conditions are comprised of the same three work allocation decisions as the *Baseline* condition: two in the No-Excuse treatment and one in the Excuse treatment. Table 4 provides the work schedules for each condition in the Conflicted state.[28]

I manipulate the benefit of procrastination by lowering the amount of Week-One tasks associated with *Less Now*, thus reducing the current workload. This makes procrastinating more attractive, and can increase the demand for excuses. To hold Week One *More Now* and Week-Two allocation amounts constant, I compare *Baseline* to ↓ *Week One* and ↓ *Week Two* to ↓ *Both Weeks.*

[26]Binarize the r=.75 interest rate in Augenblick et al. (2015) results in the prediction of roughly 25 percent picking Less Now in the Conflicted No-Excuse decision, a similar level of patience that is observed in Augenblick and Rabin (2018). This may be evidence that binary choices instead of convex decision sets make the patient option more salient to participants, increasing the cost of acting impatiently.

[27]This also shows addresses concerns that participants may be confused, not care about incentives, or avoid revealing for other reasons, as excuse-based procrastination would then remain constant across treatments.

[28]Switching the Week-Two allocations shows the incentives in the Aligned state.

Table 4: Task Allocations in the Conflicted State for Each Work Schedule

Baseline

	w1	w2
X	10	20
Y	13	13

↓ Week One

	w1	w2
X	8	20
Y	13	13

↓ Week Two

	w1	w2
X	10	18
Y	13	13

↓ Both Weeks

	w1	w2
X	8	18
Y	13	13

I alter the cost of procrastination through the maximum Week-Two allocation, or in other words, the future workload if the Conflicted state is realized. I hypothesize that the demand for excuses is inversely related to the severity of the consequence, i.e., the difficulty of engaging in self-deceptive behavior. To identify this, I compare ↓ *Week Two* to *Baseline* and ↓ *Both Weeks* to ↓ *Week One*, holding Week-One and the minimum Week-Two allocations constant.

Table 5 shows the results of regressing the differential work schedules on information avoidance and excuse taking, i.e., making state contingent decisions in the No-Excuse treatment but avoiding the state and picking *Less Now* in the Excuse condition. The work schedule changes do not impact behavior when decisions are made with full information.[29] This means that in the No-Excuse treatment, participants minimize total workload and mostly select *Less Now* in the Conflicted state. Any changes found in the Excuse treatment can therefore be attributed to excuse-based procrastination.

Looking at the *Lowered Cost* coefficient in Column 1, we see that there is an almost 10 percentage point increase in not revealing the state when the potential cost of excuse-based procrastination decreases (i.e., the maximum amount of Week-Two tasks is 18 instead of 20), while the benefit channel is insignificant. This suggests that information avoidance is driven by a sophisticated self-deceptive decision-making process that is mitigated by greater potential consequences, rather than a more impulsive process influenced by current desires. Columns 2, on the other hand, only has a significant interaction term. This implies that excuse-taking and, as a direct result, excuse-based procrastination, respond only to tandem changes in costs and benefits. However, taken together this shows that participants do respond to the incentives to procrastinate given in the experiment.

[29] Analysis shown in the appendix

Table 5: OLS Estimate of Work Schedule on Excuse Uptake and Procrastination

	(1) Information Avoidance	(2) Excuse Takers
Increased Benefit	-0.0278	-0.0417
	(0.053)	(0.046)
Lowered Cost	0.0972**	-0.0278
	(0.049)	(0.053)
Interaction	-0.0139	0.125*
	(0.064)	(0.052)
Constant	0.333***	0.250***
	(0.056)	(0.052)
N	288	288

Note: Standard errors in parentheses clustered at the individual level. Increased Benefit is an indicator variable for *Less Now* having fewer Week-One tasks. Lowered Cost is an indicator variable for a lower amount of maximum Week-Two tasks. Outcome variables are not revealing the state in the Excuse treatment and percent of participants who minimize current work after not revealing the state in the Excuse treatment and also make state-contingent decisions in the No-Excuse treatment. Does not include attrition.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

1.4.5 Types of Impatience

Throughout the experiment, participants are presented with multiple opportunities to make impatient decisions. In some situations, participants can knowingly benefit their current self at the expense of their future one. I demonstrate that procrastination which increases impatience can arise from excuses, where participants *unknowingly* make impatient decisions despite acting patiently when faced with the same incentives under full information and being allowed control over their information set. I quantify the full extent of excuse-based procrastination on patience by comparing the prevalence of both types. Figure 8 splits up all observed impatient behavior throughout the entire study – i.e., every instance of selecting *Less Now* in the Conflicted state – as either known impatience, broken down by if the information is exogenously or endogenously endowed, and excuse-based procrastination, i.e., decisions made with uncertainty.

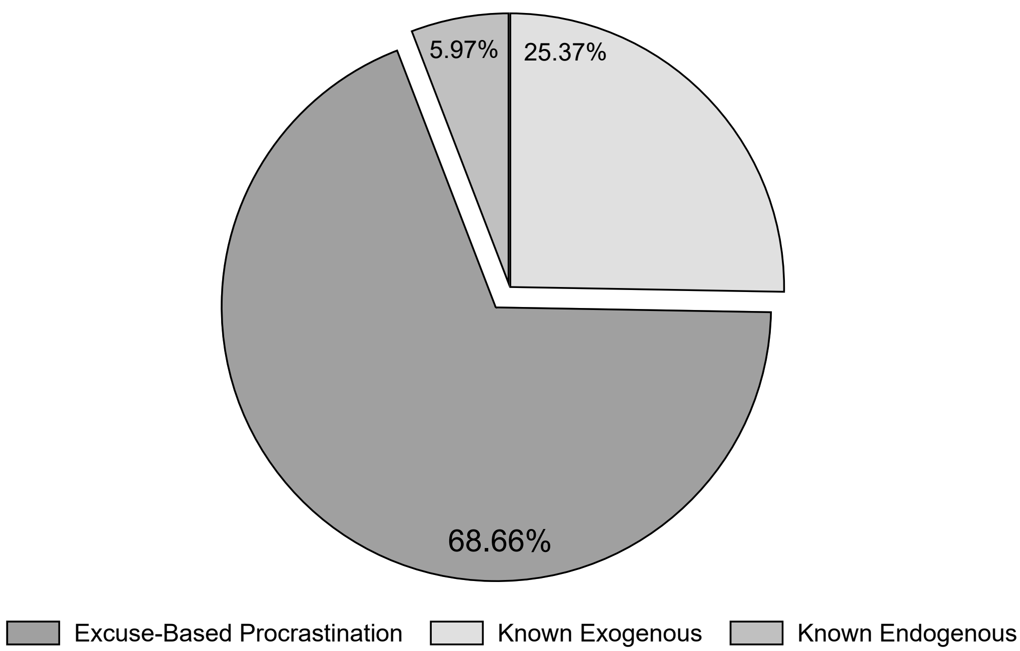

Figure 8: Impatient Decisions Split by Type

Note: Excuse-Based Procrastination is picking *Less Now* in the Excuse treatment without information. *Known Exogenous* and *Known Endogenous* are known impatience in the No-Excuse and Excuse treatment, respectively.

I document very little known impatience. Instead, excuse-based procrastination seems to be the dominant form of impatience, with almost 70 percent of all impatient behavior I observe across the study actually being excuse-based.

1.5 Discussion

This paper introduces and identifies excuse-based procrastination, where individuals use excuses in the intertemporal choice domain to justify short-sighted decisions. Excuse-based procrastination can occur when there is both the possibility for procrastination and the opportunity to take an excuse. This has costs to both the individual and to society and can lead to inefficient policy instruction and the misestimation of time preferences. Moreover, excuse-based procrastination can exacerbate persistent naïveté about an agent's own self-control problems.

I conduct a survey that shows suggestive evidence of an immediately-felt cost of doing things late that is attenuated by excuses. In addition to the measures discussed in Section 1.2, participants also indicate how much they agree with the statement "When I do something late, I usually give an excuse explaining why." I find evidence of the relationship between the act of delaying tasks, feeling bad when completing tasks late, and the use of excuses. Table 6 shows pairwise correlations between the self-reported measures in my survey. Participants who more strongly agreed with feeling bad when doing things late are also more likely to report that they perceive themselves as getting things done on time. Both directions of this relationship are plausible: individuals who feel worse when they do things late may be more likely to avoid doing so, and similarly, people who don't think of themselves as some who does things late may feel worse when they do end up doing so. Finally, both my measure of feeling an emotional cost when doing something late and self-image measures are correlated with self-reported use of excuses.

While the results of my survey are self-report measures assessing agreement with doing something late, they are consistent with my proposed mechanism of an emotional cost of procrastination. Moreover, these results are consistent with the results of my within-subjects longitudinal laboratory experiment where individuals work in two separate time periods. I examine if individuals exploit uncertainty, which can be resolved at the press of a button, over the quantity of future work as an excuse to select options that require less work today at the risk of more overall work by comparing decisions made in the Excuse treatment to decisions made when participants are exogenously given full information in the No Excuse treatment.

Table 6: Correlation Table of Survey Measures

	Use Excuses	Feel Bad (No Excuse)	Feel Bad (Excuse)	Delay
Use Excuses	1.000			
Feel Bad (No Excuse)	0.194	1.000		
	(0.000)			
Feel Bad (Excuse)	0.080	0.422	1.000	
	(0.050)	(0.000)		
Delay	0.112	0.240	0.098	1.000
	(0.006)	(0.000)	(0.016)	

Note: P values in parenthesis. *Delay* is the self-reported extent the statement "I always do things on time" describes themselves. *Use Excuses* and *Feel Bad ([No] Excuse)* are the agreement with the statements, in order, "When I do something late, I usually give an excuse explaining why", " I feel bad when I do something late", and reverse-coded "I don't feel bad when I do something late if I have an excuse as to why I did not do it on time."

I document the avoidance of information with instrumental value and how it is used as an excuse to minimize current work. Participants with full information prioritize minimizing total work, with over 91 percent of participants making the state-contingent decisions needed to do so in the No Excuse treatment. However, 37.5 percent do not acquire the information needed to ensure they make task-minimizing decisions. The 22 percent of my sample who remain willfully ignorant as an excuse to minimize their current workload are classified as excuse takers. The presence of an excuse significantly increases their overall workload and their implied intertemporal rate of substitution. Overall, excuses increase the prevalence of procrastination fourfold. Auxiliary treatments show that excuse-based procrastination is robust to different parameters and responds to incentives, notably information avoidance decreasing as consequences of the excuse increase.

Excuse-based procrastination differs from knowingly procrastinating since the excuses do not serve as an ex-post justification, but rather lead to behavioral changes. Excuse-based procrastination accounts for almost 70 percent of procrastination seen in my study. This is suggestive evidence that some short-sighted behaviors could be attributable to excuses rather than what we previously considered to be traditional impatience.

Excuse-based procrastination can lead to both individual and societal costs as agents increase short-sighted behavior in a variety of economic domains where there is an opportunity for procrastination and an excuse is available. Examples include Uber drivers who would drive if the rate is above what they think their reservation wage should be but avoid checking for surge charges as an excuse to not work guilt-free. Similarly, an employee who would open a 401(k) account if they understood their employer's generous contribution matching plan may avoid opening e-mails about it as an excuse to continue consuming guilt-free.[30] Students and workers who create excuses to procrastinate end up diminishing their human capital acquisition and failing to reach their potential, particularly in environments with minimal oversight and individual control over effort. Firms can also exploit excuse-based procrastination. One such example is casinos purposefully removing outside windows to impact gamblers' senses of time, thereby creating an excuse for patrons – who are still able to check the time themselves – to delay leaving and continue gambling into the morning hours. Excuse-based procrastination increases individual long-run sub-optimal decisions, some of which generate negative externalities. The general equilibrium impacts of this behavioral bias, therefore, reach far beyond the individual decision-maker.

Understanding the mechanisms that affect impatient decision-making can yield large economic gains. However, these gains cannot be fully realized without separating excuse-based procrastination from known procrastination, as the two patterns of behavior respond differently to information interventions. Specifically, traditional impatient decisions require strong external commitment devices while their excuse-based counterparts instead need access to excuses removed.

For example, a patient person may research the health effects of cigarettes and decide that future costs are not worth any present benefit. A traditionally impatient individual would smoke regardless of their knowledge about the health impacts and would require more severe policies such

[30]These examples highlight complementarities between excuse-based procrastination and motivated information avoidance (Andries and Haddad, 2020; Ganguly and Tasoff, 2017; Golman et al., 2017; Karlsson et al., 2009; Oster et al., 2013), where the sophisticated agents avoid information to change their behavior.

as high tax rates to prevent smoking. In contrast, an excuse taker would not smoke if they knew the health impacts but avoid doing research to allow themselves to continue smoking. Policies that put health warnings on cigarette packs could thus reduce smoking caused by excuse-based procrastination without impacting the behavior in the other two groups. Alternatively, excuse-based procrastination can influence the effectiveness of campaigns aimed at vaccination rates. People who would get vaccinated if they had full information about the long-run benefits may avoid that information as an excuse to stay unvaccinated. Increasing the availability and awareness of information would not change their behavior – instead, they must be exogenously shown it.

Understanding the impact excuses have on procrastination is important for a variety of reasons. First, the emotional cost of procrastination and its attenuation through excuses must be taken into account when modeling intertemporal choice. Second, the supply of excuses will impact the interpretation of measured time preferences. Finally, excuses provide a novel policy lever and help us understand why, despite repeat exposure, individuals continue to engage in impatient behavior that time and again proves costly to their future selves.

2.0 Streaks, Cheat Days, and the Tradeoff between Motivation and Flexibility

(joint with Kirby Nielsen)

This paper explores streaks as a novel technique to counteract impatient behavior in economic decision-making when actions have current costs but delayed benefits. Streaks are a powerful psychological motivator triggered by tracking the number of consecutive periods an action is performed, increasing the motivation for current effort. However, myopic reactions to broken streaks can lead to a tradeoff that causes streaks to backfire. We use a longitudinal real-effort experiment to study this tradeoff. We find that, conditional on working the day before, being randomly assigned to having consecutive days worked tracked, as opposed to total days, significantly increases the probability of working. However, conditional on *not* working the day prior, those in the streaks treatment work significantly less. Allowing for flexibility, such as cheat days when effort costs are high, can act as insurance against breaking streaks, thus mitigating both the motivating and demotivating aspects. We find that the Streaks treatment performs better in terms of payment when individuals do not experience an exogenous cost shock, while the most flexible treatment performs better when there is one. The Cheat Day treatment falls in between, with no difference in payment depending on cost shock.

2.1 Introduction

This paper explores streaks, or tracking the consecutive number of periods an agent performs an action, as a novel technique to increase effort provision of costly actions that have delayed benefits, such as exercising, studying, or dieting. Impatient agents often avoid such actions, leading to potentially suboptimal effort allocations and detrimental long-run outcomes. Streaks act as a *psychological* motivator; rather than directly altering incentives, such as effort costs or monetary benefits, they alter psychological incentives. This can increase effort provision in situations where costs may otherwise outweigh the discounted future benefits. The streaks paradigm is versatile and nearly costless to implement, making it an attractive solution for reducing impatient behavior.

Although streaks have the potential to motivate individuals, they come with an inherent tradeoff which demonstrates the importance of examining their role in a controlled setting. Streaks are motivating while maintained, but can backfire once broken. Inevitably, a streak will get broken in a high-cost period where the psychological incentive to continue the streak is not sufficiently large enough to motivate continued effort. This broken streak may be demotivating and decrease future participation, even on low-cost days when the agent would have otherwise worked. The streak paradigm therefore may backfire if the short-run increase in effort within the streak does not compensate for the long-run decrease in motivation.

To identify the effects streaks have on effort, we use a longitudinal online experiment with real-effort tasks. This allows us to control effort costs, avoid selection effects by randomizing participants into the streak paradigm, and measure effort on both the extensive and intensive margin. Participants earn a delayed piece-rate payment for learning vocabulary words by transcribing their definitions, a task with current costs but delayed benefits. Each day during the nine-day study, participants were given 15 minutes to learn up to K words, earning a "gold star" and an additional delayed bonus. We designed this incentive structure so that participants have a monetary incentive to learn more vocabulary words regardless of the value of K, but the gold star incentive allows us to vary the streak incentive, as explained below. Furthermore, we vary K daily; this allows us to exogenously manipulate the cost of maintaining a streak as assigning participants a prohibitively-large K exogenously breaks their streak. However, they can still complete tasks to earn payment, allowing us to measure both the intensive and extensive margins of effort.

We alter the way gold stars are tracked using between-subject treatments. In the *Baseline* treatment, participants are shown the total number of "gold star" days completed thus far. This serves as a control treatment in which we simply track the number of "full effort" days. In contrast, in the *Streaks* treatment, we only display the *consecutive* number of gold star days completed. This difference affects the gold star counter when a participant learns fewer than K words on a given day—in the *Baseline* treatment, learning fewer than K simply does not increase the counter, while in the *Streaks* treatment, it resets the counter to zero. Since we exogenously assign K, and we assume that individual effort costs are orthogonal to treatment assignment and the daily K, we can identify the psychological impact of streaks through differences in effort, on both the extrinsic and intrinsic margins, across treatments.

We find that streaks motivate continued effort: In the *Baseline* treatment, the likelihood that a participant works, conditional on working the previous day, is 78 percent, and this increases to 84 percent in the *Streaks* treatment. However, we also find that breaking a streak is demotivating: Individuals in the *Baseline* treatment return 65 percent of the time after failing to reach K, while individuals in the *Streaks* treatment return only 59 percent of the time. Although we are under-powered to detect significant differences, we find that which treatment performs better directionally depends on if there is a cost shock or not: with a cost shock, the extra flexibility in the *Baseline* treatment leads to better outcomes, otherwise, the motivating power of streaks increases payment. However, as individuals are less likely to return after breaking their streak later on, this result may not hold in the long-run.

Our results suggest that streaks are a powerful motivator, but their effectiveness can be short-lasting, raising the question of whether it is possible to harness the power of streaks while minimizing their drawbacks. We explore this further in an additional treatment that examines whether flexibility or what we refer to as "cheat days," can mitigate the demotivation that breaking a streak causes. Our *Cheat Day* treatment is the same as our *Streaks* treatment, except the gold star counter does not reset if K is not met on days when K is particularly high. This gives participants flexibility on high-cost days—effectively providing insurance against unexpected cost shocks—but also has a trade-off: While cheat days mechanically preserve the streak and thus reduce the demotivating effect of breaking the streak, they may also reduce total effort on high-cost days, since individuals no longer have the psychological motivation to increase effort to maintain their streak. The probability of working in the *Cheat Day* treatment falls directionally in between that of the *Streaks* treatment and that of the *Baseline* treatment, although not significantly different from either. Moreover, payment does not differ between participants who do or do not experience a cost shock, demonstrating how flexibility can act as insurance, while still having the benefit of not having complete flexibility.

Although streaks have not been explicitly studied in the economics literature, they have important implications for both policy-makers and firms. For instance, streaks are used in substance abuse recovery programs like Alcoholics Anonymous, where continuous sobriety is rewarded by earning chips. However, streaks can also potentially backfire in these situations, such as perfect attendance awards increasing absenteeism after the initial absence (Robinson et al., 2019). This underscores the importance of understanding streaks' tradeoffs, as encouraging lifetime patient behavior is just as important as short-run patient behavior in many policy-relevant domains.

Streaks also have significant implications for firms, as they can be used to incentivize customer retention and increase the productivity of workers. Rideshare companies such as Lyft and Uber use "Streak Bonuses" and "Consecutive Trip Promotions," respectively, to encourage drivers to stay online and to move the supply of drivers into areas with high demand.[1] Moreover, Lyft allows drivers to take one fifteen-minute break without breaking their streak, mirroring our cheat day implementation. Outside of motivating workers, firms like Snapchat, Duolingo, and The New York Times also use streaks as incentives for their consumers. The popularity of streaks for these companies shows how powerful they can be, while consumers' reactions to breaking a streak show how painful it can be, with Snapchat even having a team dedicated to dealing with requests to restore streaks. Some companies, such as Duolingo, go beyond simply using streaks as a motivational aspect, and monetize the pain associated with breaking a streak, allowing their users to purchase "Streak Freezes" or pay almost $10 to repair a missed streak. Moreover, they find that automatically granting users a "weekend pass" increases the returning number of users on Monday by 4 percent. We supplement our experiment by analyzing observational data from Duolingo.

This chapter is the first to document the motivational aspect of non-incentivized streaks and its tradeoff with flexibility, thus contributing to a variety of literatures. First, we provide a theoretical contribution to understanding intertemporal choice by identifying a framework that mitigates impatient behavior without altering the pecuniary benefits or effort costs. Allowing for streaks to enter to utility function with loss aversion with respect to the current streak, similar to cumulative prospect theory (Tversky and Kahneman, 1992), can help explain this behavior. This contributes to our understanding of how reference points can influence effort allocation (Abeler et al., 2011; Allen et al., 2017; Camerer et al., 1997; Corgnet et al., 2018; Farber, 2015; Heath, 1999).

Moreover, this implies that longer streaks introduce the potential for greater "loses," aligning with our findings in our experiment that those in the *Streaks* treatment exert more effort in the intensive margin on high-cost days that occur further along in the experiment and our findings using observation data from DuoLingo that those with longer streaks are less likely to return. Interestingly, this result is counter to the intuition behind models of habit formation which predict

[1] Both Uber and Lyft incentivize streaks using bonus payments, while we rely on a natural streak incentive. As this is a starker framework, it further shows how streaks may be underused by firms.

that if streaks motivate agents to perform an action for a longer consecutive timeframe, they should be *more*, not less, likely to re-start the streak (Acland and Levy, 2015; Dynan, 2000; Havrankey et al., 2017; Royer et al., 2015).

Another way that framing can influence intertemporal choice includes goal setting (Clark et al., 2020; Corgnet et al., 2015; Harding and Hsiaw, 2014; Hsiaw, 2013). However, streaks have a unique result: if the goal is to reach the target each day, individuals should be willing to restart their broken streak. Finally, we contribute to the growing literature on the tradeoffs between flexibility and rigidity in incentive schemes designed to increase effort (Bénabou and Tirole, 2004b; Burger et al., 2011; Falk and Kosfeld, 2006; Himmler et al., 2019; Hsiaw, 2018; Koch and Nafziger, 2016, 2020) and understanding the role of costly commitment devices (Ariely and Wertenbroch, 2002b; Bisin and Hyndman, 2020; Bryan et al., 2010).

The rest of the chapter is organized as follows: Section 2.2 discusses results from an analysis of Duolingo language learning data. We describe the experimental design of our online experiment in Section 2.3 and its results in Section 2.4. Section 2.5 concludes.

2.2 Motivational Evidence

As preliminary motivational evidence, we collect data from individuals who use Duolingo, a popular language-learning program that heavily emphasizes streaks. Maintaining a streak requires completing one lesson per day, which only takes a few minutes. Duolingo learners who use the phone app receive daily reminders which make their current streak length very salient.

We pay 127 Spanish learners on MTurk $2 to join a Duolingo classroom.[2] This gives us access to their Duolingo history, including the time-stamp of all lessons completed. There was heterogeneity in usage levels, with participation rates varying from two days to over one thousand days. Most importantly for our purposes, the number of consecutive days also worked varied a great deal. We use this data to calculate streaks, which also vary in length, with our longest recorded streak lasting 656 days. The average streak length is 49 days.

[2] Participants know what data we receive. We remove each participant immediately after collecting their data.

There is significant path dependency in the use of Duolingo. 84 percent of individuals who use Duolingo on day $t-1$ return on day t. In stark contrast, only 3 percent of individuals who *don't* use Duolingo in time $t-1$ return the following day (Fisher's exact, $p < 0.001$). We find a positive correlation between the length of consecutive days worked thus far and the likelihood of using Duolingo the following day; that is, individuals with longer streaks are more likely to continue the streak tomorrow (correlation=0.223, $p < 0.001$). Additionally, in a multiple choice question, 80 percent of subjects report having broken a past streak due to an unexpected time constraint rather than changes in preferences.

Despite having our participants' full Duolingo use history, our analysis still has some limitations. First, Duolingo offers a "streak freeze" which essentially "glosses over" a missed day, meaning that a participant who has used Duolingo for 27 consecutive days but misses day 28 can still maintain their streak. These can be purchased in the app and are automatically enacted upon missing a day. Since we cannot see streak freezes in our data, we cannot say for sure which streaks are actually broken. However, this means the streaks that we calculate are actually a *lower* bound of streaks.

Second, this data is plagued by endogeneity concerns. Individuals with shorter streaks are likely to be less serious about learning the language or have other unobserved traits that predict they would be less likely to continue using Duolingo. Additionally, there can be a selection into using Duolingo and apps that are known for keeping track of streaks. Finally, costs can be correlated across days, meaning that when someone breaks a streak due to a high cost in day t, the high cost continues into day $t+1$. This means that while this evidence provides suggestive evidence of our hypothesis, we cannot identify the *causal* impact of breaking a streak in this data. Thus, we turn to a controlled experiment.

2.3 Experimental Design

We ran a 9-day longitudinal study online. Our experiment mimics how streaks are used by firms like Duolingo while still affording us control over the setting. We recruited 180 participants from Prolific and invited them to participate in a longitudinal study that could last up to two weeks, without telling them the exact horizon.[3] We randomly assigned each subject to one of three treatments: the *Baseline* treatment, the *Streak* treatment, and the *Cheat Day* treatment. We describe the main task and treatments below.

2.3.1 Experiment Task

On each day of the experiment, participants learn vocabulary words for 15 minutes by transcribing each word's definition. We chose this task for a few reasons. First, it is very simple and easy for participants to understand. Second, there is a clear metric of effort: the number of words learned on a given day. Finally, compared to other real-effort tasks, learning vocabulary may carry additional benefits due to feelings of personal improvement and motivation.

We pseudo-randomly assign subjects a number that is their "gold star amount," (hereafter called the "target") every day. Subjects receive $0.01 for each word they learned up to the target amount, but earn $1.50 total if they reach the target. All payments are made on the final day of the study. By varying the target, we effectively manipulate the cost of maintaining a streak.

We chose three potential targets: 15 words, 50 words, and 125 words. We ran a pilot session before our main experiment in order to collect data on how many words participants could complete in 15 minutes. We chose the target of 125 as a level that almost no one could reach in 15 minutes. Thus, a target of 125 words serves as an exogenous cost shock to maintaining a streak. We chose the targets of 15 and 50 to be easy and moderate targets, respectively, while ensuring most participants would be able to reach the target in 15 minutes. Breaking a streak due to unexpected cost shocks aligns with the responses from our survey of why individuals ended up breaking their DuoLingo streak. Exogenous cost shocks that are orthogonal to individual effort costs allow us to disentangle the impact of breaking the streak from other explanations for a decrease in effort post-broken streak, such as the correlation between effort costs in consecutive days.

Most days, we randomly assign participants either a 15- or 30-word target. Additionally, each participant is assigned to receive the 125-word target on either day three, day eight, or never. This allows us to compare the effect of breaking streaks of different lengths (i.e., two or seven-day streaks) while maintaining enough power for each comparison.

[3]This reduces backwards unraveling that may occur with full knowledge of the end date.

2.3.2 Treatments

In all treatments, subjects have a "tracker" that keeps track of the number of days they reach their target; treatments vary exactly what this tracker counts. We display the tracker prominently in the corner of each screen (shown below in Figure 9), and remind subjects of the current value of their tracker at the beginning of each day.

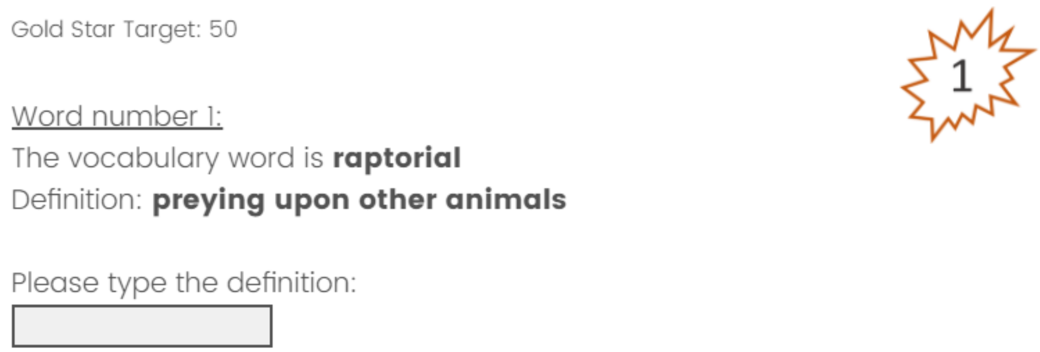

Figure 9: Subjects' Decision Screen

In the *Baseline* treatment, the tracker increases each day a subject reaches their target and never resets. In the *Streak* treatment, the tracker increases for each *consecutive* day a subject reaches their target, and it resets to zero if a subject fails to reach their target on any given day. In the *Cheat Days* treatment, the tracker works the same as in the *Streak* treatment but does not reset to zero if the subject fails to reach the 125-word target (though does reset if they fail to reach a 15- or 50-word target). In other words, streaks are not broken by failing to meet the target on high-cost days in the *Cheat Day* treatment, but they are broken in the *Streak* treatment.

2.3.3 Hypotheses

We have three main hypotheses of how subjects' effort will differ by treatment. First, we hypothesize that subjects are incentivized to maintain their streaks. We test this by comparing effort in the *Streak* treatment to effort in the *Baseline* treatment, for those subjects who successfully reached the target on the previous day. These subjects have the same earnings, same experience with the task, etc., and the only difference is that subjects in the *Streak* treatment will have their trackers reset to zero if they don't meet the target.

Hypothesis 1. ***Streaks are motivating****: Conditional on meeting the target on day $t-1$, subjects in the Streak treatment will be more likely to work on day t than subjects in the Baseline treatment.*

Second, we hypothesize that this motivational power of a streak can backfire once individuals break the streak. We test this by again comparing effort in the *Streak* treatment to effort in the *Baseline* treatment, this time for subjects who did *not* reach the target on the previous day. Individuals in the *Streak* treatment had their tracker reset to zero after failing to reach the target, while individuals in the *Baseline* treatment did not.

Hypothesis 2. ***Breaking a streak is demotivating****: Conditional on not meeting the target on day $t-1$, subjects in the Streak treatment will be less likely to work on day t than subjects in the Baseline treatment.*

Finally, we look at the effect of "cheat days." In the *Cheat Day* treatment, subjects' trackers are not reset on high cost days after failing to reach the target.

Hypothesis 3. ***Cheat days reduce motivation and demotivation****: The effects of hypotheses 1 and 2 will be reduced in the Cheat Day treatment compared to the Streak treatment. Overall, this leads to an ambiguous prediction between the Cheat Day treatment and the Baseline treatment.*

2.4 Results

We find that streaks are, indeed, motivating. The left side of Figure 10 shows the percentage of subjects who choose to exert any effort on day t, conditional on meeting the target on day $t-1$. Comparing the *Streaks* treatment to the *Baseline* treatment shows that while 84% of subjects in the *Streak* treatment work the following day, only 78% of subjects in the *Baseline* treatment do.

In Table 7, we confirm our first result using a Probit regression. In column (1), we regress the binary decision to work at time t on an indicator variable for being in the *Streak* treatment relative to the *Baseline* treatment, an indicator for meeting the target on date $t-1$, and the interaction of these two. We find that meeting the target on date $t-1$ is positively and significantly correlated with working on date t. This is not too surprising, since individuals with higher intrinsic motivation will be more likely to meet the target and more likely to work the following day. What we find, though, as seen through the interaction term, is that this effect is significantly larger in the *Streak* treatment compared to the *Baseline* treatment. That is, while individuals in the *Streak* treatment are not more likely to work per se, they are more likely to work to maintain their streak. Thus, we find support for Hypothesis 1.

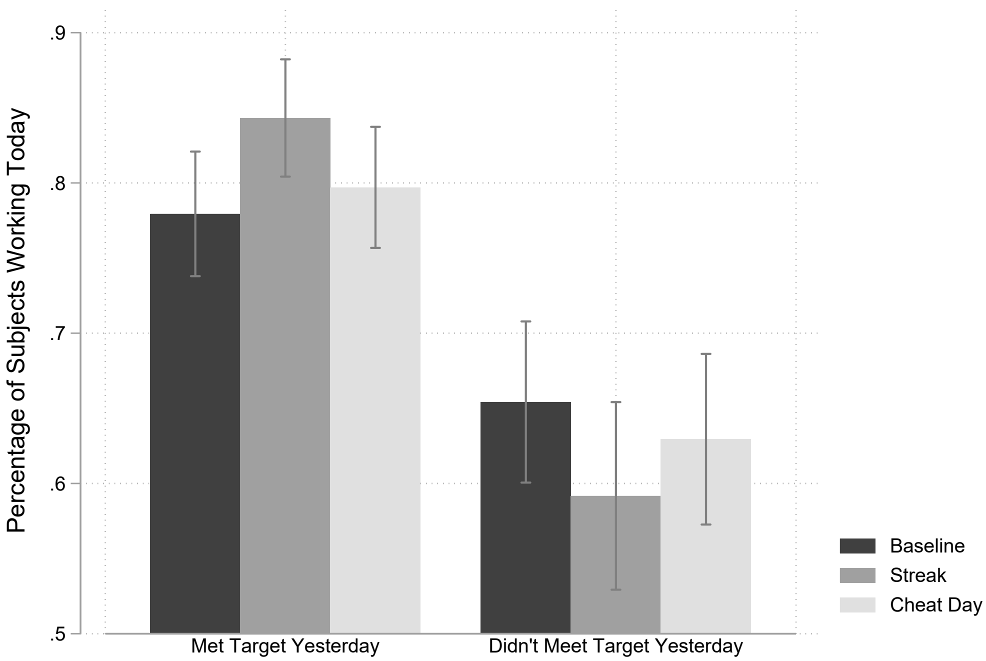

Figure 10: Work Conditional on Previous Day

Result 1. ***Streaks are motivating***: *Conditional on meeting the target on day $t-1$, subjects in the Streak treatment are more likely to work on day t than subjects in the Baseline treatment.*

The right-hand side of Figure 10 shows how streaks are also demotivating. We again compare the two left-most bars. Conditional on *not* meeting the target on date $t-1$, individuals in the *Streak* treatment are less likely than individuals in the *Baseline* treatment to work on date t. Only 59% of individuals in the *Streak* treatment work, compared to 65% in the *Baseline* treatment. We confirm our results again using a regression, shown in column (2) of Table 7. Similar to the results above,

Table 7: Streaks are Motivating Unless Broken

	(1)	(2)	(3)	(4)
Completed any tasks that day				
Streak Treatment	-0.236	0.238	-0.236	0.238
	(0.229)	(0.166)	(0.229)	(0.166)
Cheat Day Treatment			-0.103	0.0609
			(0.203)	(0.154)
Met Target on $t-1$	0.676***		0.676***	
	(0.149)		(0.149)	
Streak × Met Target on $t-1$	0.474*		0.474*	
	(0.271)		(0.271)	
Cheat Day × Met Target on $t-1$			0.164	
			(0.237)	
Did Not Meet Target on $t-1$		-0.374***		-0.374***
		(0.132)		(0.132)
Streak × Did Not Meet Target on $t-1$		-0.402*		-0.402*
		(0.231)		(0.231)
Cheat Day × Did Not Meet Target on $t-1$				-0.127
				(0.209)
Constant	0.0941	0.770***	0.0941	0.770***
	(0.136)	(0.109)	(0.136)	(0.108)
Observations	792	891	1208	1359

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

we find that while individuals in the *Streak* treatment are no less likely to work overall than those in the *Baseline* treatment, they *are* less likely to work on days following a broken streak.

Result 2. ***Breaking a streak is demotivating****: Conditional on not meeting the target on day $t-1$, subjects in the Streak treatment are less likely to work on day t than subjects in the Baseline treatment.*

Finally, we turn to our *Cheat Day* treatment to see the extent to which increasing flexibility in the streak affects individuals' motivation. On both sides of Figure 10, the right-most bar shows the percentage of subjects who work at all on date t in the *Cheat Day* treatment, separated by whether they met the target on date $t-1$ or not. As expected, we find that cheat days mitigate both streak effects: Streaks are less motivating when individuals have cheat days, but high-cost days do not lead to discouragement when they do not break the streak. Thus, we find that individuals are *less* likely to work in the *Cheat Day* treatment compared to the *Streak* treatment on days when they met the target on the previous day, but they are *more* likely to work on days when they did not meet the target.

Columns (3) and (4) of Table 7 confirm this result. While the interaction of being in the *Cheat Days* treatment with (not) meeting the target on day $t-1$ is closer to zero in both conditions when compared to that of the *Streaks* condition, it is not significantly different from either the *Streaks* treatment ($\chi^2 = 3.07$, p = 0.215) or the *Baseline* treatment.

Hypothesis 4. ***Cheat days reduce motivation and demotivation****: The effects of results 1 and 2 are reduced in the Cheat Day treatment compared to the Streak treatment.*

2.4.1 Impact of Cost Shocks

One unique aspect of our design is our control over the cost of meeting the target and therefore continuing the streak. As a quick motivating example for this analysis, imagine that someone runs for an hour each day. Part of the cost of running for an hour includes the opportunity cost of their time. Therefore, days when they are busy in general can be thought of as "high cost" days where it is more difficult to meet their target of running for an hour. However, even if they are not able to run for a full hour, they can still run a shorter amount of time—not reaching their target and experiencing less total benefit than if they had, but running any amount would still be beneficial.

Table 8: Analysis of Behavior by Cost

	(1)	(2)
	Worked Any	Words Learned
Streak Treatment	0.0401	-0.126
	(0.165)	(1.579)
Cheat Day Treatment	0.0111	-0.983
	(0.149)	(1.447)
High Cost Day	-0.179	43.42***
	(0.217)	(5.831)
Streak × High Cost Day	0.0240	13.97
	(0.313)	(9.017)
Cheat Day × High Cost Day	-0.0351	3.124
	(0.293)	(8.415)
Constant	0.609***	30.04***
	(0.106)	(1.125)
Observations	1359	991

Worked Any is a probit over the binary variable of learning any words

Words Learned conditions on learning any words

Standard errors clustered at the individual level in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

However, the high cost of reaching the target can serve as an excuse and allow someone to instead not work at all.

Similarly, the number of words assigned as a target changes the difficulty of meeting the target in our study.[4] On days with the highest target, where it is unlikely that someone would be able to finish the tasks, they are still able to learn vocabulary words for a piece-rate bonus that day. We therefore look at two different outcomes that can be impacted by the cost: whether someone works at all, and conditional on working, how many words they learn.

First, we focus on whether the *Streaks* treatment can motivate individuals to work on high-cost days in order to preserve their streak. The first column of Table 8 uses a Probit regression to explore the extensive margin decision to work any amount on a given day. Unlike the analysis in the previous section, this does not condition on an individual's history. Across the entire study, there does not seem to be a treatment difference in the number of days worked. This may be due to the motivational aspect of streaks helping mainly at the beginning – at the study of the study, a marginally larger percentage of participants in the *Streaks* treatment work each day compared to the *Baseline* treatment. However, this switches on day four, coinciding with the first cost shock. It is therefore unknown which treatment would be more motivational if continued longer. Moreover, we are underpowered to detect any significant differences in the work decisions by day.

The second column of Table 8 shows the results of a linear regression of the number of words learned, conditional on working at all. Although there is very little difference in the number of words learned when the cost is not high, participants in the *Streaks* treatment who do work on a high-cost day learn marginally more words compared to the *Baseline* treatment ($F = 2.24, p = 0.138$). This does have an impact on outcomes, though: while only one individual in the *Baseline* treatment completed 125 words and met the high-cost target, five people in the *Streaks* treatment did. This aligns with the hypothesis that individuals work harder to maintain their streak on high-cost days.

This change in motivation is evident when looking at Figure 11, which shows the number of words learned for individuals who have a target of 125 words. Although we don't see any difference between the *Streak* treatment and *Baseline* treatment in the number of words learned on Day 3, we do see a significant difference on Day 8. Individuals in the *Baseline* treatment learn 75 words on average, while individuals in the Streak treatment learn 102 words ($p = 0.052$). This is consistent

[4]Any additional cost shocks would be orthogonal to the exogenous treatment assignment and cost shocks.

with individuals in the *Streak* treatment who decide to try to keep their streak by exerting more effort to try to maintain their 8-day streak. The fact that this difference is significant on Day 8 and not on Day 3 also suggests that individuals place more value on longer streaks, and worker even harder to maintain their streak as the streak length grows.

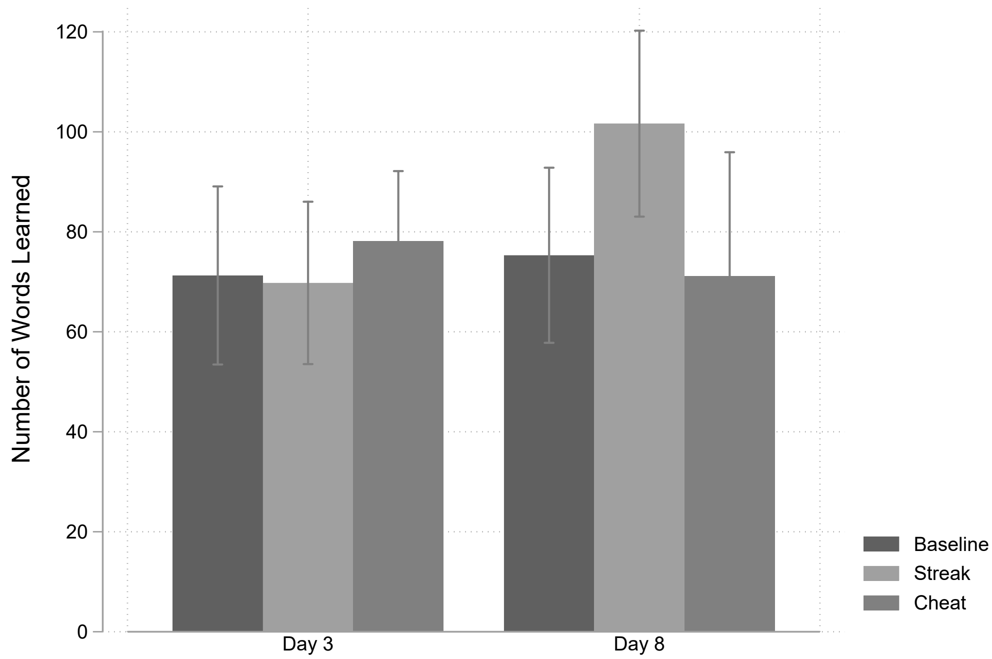

Figure 11: Intensive Margin of Effort on High Cost Days

To explore how flexibility impacts behavior on high-cost days, we now look at the *Cheat Days* treatment and see that, compared to the *Streak* treatment, having a cheat day reduces motivation on both the extensive and intensive margin. This is most evident on day 8: only 50 percent of workers in the *Cheat Day* treatment worked at all, compared to 64 percent in the *Streaks* treatment and 81 percent in the *Baseline* treatment. The high motivation of those working seen in the *Streaks* treatment is reduced as well, with participants in the *Cheat Day* treatment learning on average 71 words, even fewer than those in the *Baseline* treatment. Thus, while we find evidence that cheat days do help maintain continued motivation, they do attenuate effort on days when costs are high.

2.4.2 Impact on Payment

To compare performance across treatments, we look at the average payment for each treatment in Table 9. Looking at total payments, we find that the demotivation that arises from breaking a streak does not counteract how motivating they are, with participants in the *Streaks* treatment earning \$0.38 more on average compared to the *Baseline* treatment ($t = -0.519, p = 0.302$).

Table 9: Average payment by treatment by if they experienced a cost shock

	No High Cost	High Cost	Total
Baseline Treatment	\$8.05	\$9.23	\$8.84
Streaks Treatment	\$9.57	\$9.02	\$9.22
Cheat Day Treatment	\$9.15	\$9.16	\$9.16
Total	\$ 8.88	\$9.15	\$9.06

However, we can see both the motivational and demotivation aspects of streaks when comparing outcomes between those who experienced a cost shock and those who didn't. Of the participants who never had a high target, those in the *Streaks* treatment earned, on average, \$1.52 more than those in the *Baseline* treatment ($t = 0.992, p = 0.164$). In comparison, when looking at participants who *did* have a high-cost day, the *Baseline* treatment earned more, on average, by \$0.20 ($t = 0.261, p = 0.397$). This shows how random cost shocks can greatly influence the efficacy of streaks. Unsurprisingly, those in the *Cheat Day* treatment earned in between the other two treatments. Interestingly, they earned the same amount regardless of if they experienced a cost shock or not, showing how flexibility truly acts as insurance against cost shocks.

2.5 Discussion

We study the motivational power of streaks in an experiment that allows us to explicitly control cost shocks, economic benefits, and the framing of streaks, in a setting that allows us to measure effort at both an intensive and extensive margin. We complement our experiment with suggestive evidence from the language-learning app, Duolingo. We find that streaks increase the likelihood of effort in order to maintain the streak, but decrease effort once the streak is broken. Keeping a streak can also increase the motivation to work harder on high-cost days in order to prevent breaking it. Finally, we find that allowing for flexibility on high-cost days reduces the motivation on those days, but increases the probability of working in the future by not breaking the streak. These results can help inform the costs and benefits of implementing these non-monetary psychological incentive mechanisms, though we note that our results are limited to a short time window (less than two weeks). Future work should explore the tradeoff between streaks and flexibility on a longer horizon to see the long-run impacts of this tradeoff.

The streak and cheat day frameworks are costless to implement, acting as cheap interventions to help individuals make patient decisions, potentially improving their health and financial standing. The streak framework is versatile and can be applied to a wide variety of situations. Understanding the best way to implement the streak framework can be useful for educators in incentivizing attendance, dietitians in helping people stick to a healthy diet, or understanding how to best help people with substance dependency remain sober. Moreover, streaks have important implications for firms that want to increase worker productivity and consumer retention. Our results contribute to the understanding of how to impact effort without changing the pecuniary benefits of doing so.

3.0 Time and Punishment: Penalty Timing and Cross-Debt Default Decisions

(joint with David Agorastos)

Using observational credit report data, we show that a substantial subset of financially distressed borrowers who hold a varied debt portfolio avoid defaulting on revolving credit, resulting in an eventual default on their student loans. The severity of student loan default compared to credit card default makes this a suboptimal financial decision that has negative consequences at both the individual and aggregate levels. We explore how the different time horizon of default and therefore penalties of these debts influence default decisions using a survey and find that participants rank the timing of default as the most important factor in influencing default decisions. Additionally, we find a higher level of financial literacy for credit cards than for student loans.

3.1 Introduction

Debt is pervasive in our financial landscape, with consumer debt in the United States alone reaching \$17 trillion in 2022.[1] High levels of indebtedness can mean that many borrowers are unable to repay their debts. In fact, 16.2 percent of borrowers are delinquent, meaning that 2.5 percent of debt has not been paid in over 90 days. Financially-distressed borrowers with multiple delinquent debts must make the difficult choice of which debt to pay and which to default on. These decisions impact not only their own financial well-being, but also the prices of debts and the solvency of federally subsidized programs, affecting financial markets and the broader economy.

Despite the fact that borrowers usually hold multiple types of debt, consumer finance research typically examines decisions over one type of debt. However, contractual terms (including interest rates), repayment plans, collateralization, bankruptcy, and penalties associated with default vary between types of debt and can impact repayment decisions. This paper examines the impact of the time horizon of the contractual terms on the default decisions of financially-distressed borrowers who hold two types of debt: student loans and revolving credit, such as credit card debt.

[1]See Federal Reserve Bank of New York Household Debt and Credit.

Student loans and revolving credit account for a significant portion of consumer debt in the United States, with balances of $1.60 trillion and $986 billion respectively. Both of these types of debts do not have collateral, however, they differ in terms of the severity and timeline of default penalties. Defaulting on student loans carries more severe, albeit delayed, consequences than defaulting on revolving credit. In this paper, we explore the trade-offs that borrowers face when deciding which types of debt to prioritize when they are unable to cover the minimum monthly payments on all of their debt.

Our analysis uses data from observational credit reports in the United States and focuses on how default rates differ between student loans and revolving credit. The borrowers who default on either loan are similar in terms of observables, both being young adults with low income and bad credit scores. We look at the first incident of default and find that the default rate for student loans is higher than that of revolving credit. Despite the more severe penalties associated with student loan default, around one third of these defaults occur among borrowers who have both types of debt. To better understand the impacts of default decisions, we use event studies of default on financial variables of interest, such as credit scores.

Our results imply that borrowers often opt to default on debt with later but harsher penalties, rather than the debt with immediate consequences. For example, when a borrower is delinquent on both their credit card debt and their student loans, they face the choice of paying their credit card bill and avoiding default, or paying their student loans to reduce the probability of default in the future. They may choose to pay the former to avoid default in the short run, despite the more severe penalties of student loan default that include wage garnishment and withholding of tax returns and social security payments.

We complement our empirical analysis with a survey, which yields three main results that provide insight into the mechanisms behind financially-distressed borrowers' default decisions. First, most of the participants report that the timing of default is the most influential aspect in this decision-making process, even more so than the overall severity of the penalties. This further emphasizes the importance of understanding how differences in time horizons for punishments of different debts influence financial decision making.

Second, our survey reveals that individuals know more about the consequences of defaulting on credit card debt than on student loans. This holds true even for those who hold both types of debt, for whom this lack of knowledge may be driven by a lack of motivation to acquire even easily available information on the debt with the later default date. Borrowers may therefore have a misspecified model as they do not fully understand the severity of the penalties, increasing the default rate on student loans. Finally, our survey does not provide evidence for differential risk preferences caused by the presence of an intertemporal tradeoff in the timing of the lottery's resolution. We discuss this result and how the lottery decision in our survey differs from the risk present in the default decisions.

Taken together, our empirical analysis and survey results provide insights into the complex trade-offs and decision-making processes that borrowers face when managing their debt obligations. We specifically highlight how the time horizon of default is more important to borrowers' default decisions than the severity of the penalty, which provides insight into ways to design contractual terms of debt to increase repayment. Reducing default can greatly improve the financial well-being of borrowers and help reduce poverty traps.

The rest of the chapter is organized as follows: Section 3.2 defines the terms used in the paper and provides institutional knowledge. Section 3.3 presents and discusses the results of our analysis of the credit report data, followed by the results and discussion of our survey in Section 3.4. Finally, Section 3.5 concludes.

3.2 Background and Definitions

Our main analysis focuses on the first transition into *severe delinquency* for revolving credit, i.e., credit card debt, and student loans. We define severe delinquency as paying under the minimum payment for 90 or more days, which allows us to compare decisions across debts with different time frames for default and penalties. For example, delinquency on federal student loans is not reported to the credit bureaus until 90 days have elapsed since the last payment, while the first missed payment is typically reported for revolving credit.

Borrowers who continue to miss minimum payments then transition from delinquency to default and begin incurring severe penalties. Student loan default happens after 420 days without a payment, while revolving credit default only requires missing payment for 90 days. However, this means that the 90-day delinquent period for student loans coincides with the time horizon for default on revolving credit, allowing us to compare the transition of revolving credit from delinquent to default to the transition of student loans to delinquency. This situation is one where a borrower must decide between an immediate cost by defaulting on revolving credit versus risking future default on student loans. Overall, we find that approximately 30 percent of student loan borrowers who enter severe delinquency, i.e., do not make a payment for over 90 days, incur the penalties associated with student loan default.[2]

We restrict our analysis to revolving credit and student loans only. These two debts are unique given that neither is collateralized, meaning there is no underlying asset that can be repossessed in the event of default. This lack of collateral leads to these debts having otherwise severe penalties for defaulting or being delinquent for an extended period of time. Comparing decisions made about different types of uncollateralized debt can provide unique insights into how borrowers take the time-horizons of default into account.[3]

Despite both being uncollateralized, student loans and revolving credit differ in several key ways that make repayment decisions between the two interesting to examine. Both debts incur interest on their balance, but student loans typically have a significantly larger balance while revolving credit typically has a higher interest rate. Moreover, interest on revolving credit compounds daily after 30 days, while student loans accrue interest at a fixed monthly rate unless an individual defaults. For repayment decisions above the minimum payment, paying off revolving credit may be an attractive option due to the higher interest rate and lower balance. However, when a borrower cannot afford both minimum decisions, additional attributes of the debt should be taken into account.

While both debts can be sent to a private collection agency after default – lowering credit scores and reducing access to future credit by a similar amount, as shown in Section 3.3 – other penalties differ across the two. Defaulting on revolving credit closes that line of credit, reducing access to current money, while the federal government can garnish wages and withhold both tax refunds and social security benefits indefinitely and without a court order after student loan default.[4] Moreover, except under rare circumstances, student loan debt is not dischargeable through bankruptcy.[5]

[2]The appendix contains further discussion of why we choose the definitions we did.

[3]We show in the appendix that borrowers overwhelmingly prioritize repayments on debts that can result in repossession, such as auto loans or mortgages

[4]See Understanding Delinquency and Default. States can also suspend driver's licenses and professional licenses.

[5]See Discharge in Bankruptcy.

The differences in terms of student loans and revolving credit present trade-offs for borrowers who cannot pay all of their minimum monthly payments. Remaining current on revolving credit allows borrowers to smooth their consumption (up until they hit their credit limit) but remaining current on student loans serves only to stave off the severe penalties for defaulting on them later on. In the short term for young, low-credit-score borrowers, the pecuniary penalties will be larger for revolving credit; however, the pecuniary penalties for eventually defaulting on student loans will be much higher.[6] In other words, the terms of these debts cause financially-distressed borrowers to choose between immediate pain in the form of losing credit access and lasting pain, such as the garnishment of wages and withholding of both tax returns and social security benefits.

3.3 Evidence from Individual Data

We use proprietary data from the Experian Credit Reporting Bureau to gain insights on consumer credit decisions made by borrowers who hold a diverse debt portfolio and are unable to meet all their payment obligations. More specifically, we look at the rate at which borrowers transition into delinquency for revolving credit compared to student loans, which we define as not making a minimum payment for at least 90 days. However, for ease of understanding, we will refer to this as *default decisions* made by student loan defaulters or revolving credit defaulters.

3.3.1 Data

Our data set is a large nationally-representative administrative panel consisting of anonymized credit files that track 1,000,000 individuals from the 2004 Q1 to 2012 Q2 and an additional 100,000 from 2013 Q1 to 2018 Q2.[7] For everyone in the panel, we see information about their debt portfolio, delinquencies and defaults, credit score, estimated individual and household income, and some limited demographic characteristics such as birthdate, ZIP code, state, and public records such as those related to bankruptcy.[8]

[6] Yannelis (2020) use a parsimonious quantitative model to explore strategic student loan default, where borrowers exploit the timing of the penalties similar to cash-flow managers/strategic defaulters in the mortgage industry.

[7] A comprehensive discussion of the panel's sampling and construction can be found in the Appendix.

[8] The credit score is the Vantage Score 3.0 produced by the Experian Credit Bureau. Individual and household income is estimated by the Experian Credit Bureau and validated against IRS income data.

In terms of debt portfolios, we observe the number of trades, balance, original credit amount or available credit, the minimum monthly required payment, amount past due or in collections, delinquencies and defaults, and information on debt inquiries for all major debt types – auto loans and leases, mortgage-type trades, revolving credit excluding and including home equity lines of credit, and deferred and non-deferred student loans.[9] In total, we observe more than 200 variables describing the debt portfolios of borrowers. In our analysis, we restrict the sample to borrowers 35 years or younger who hold student loans in the panel and do not originate a mortgage-type trade within eight quarters of their last observation.

This dataset affords us the unique ability to document how the trade-offs presented by holding different debt portfolios change borrowers' default choices and to analyze how a borrower's financial characteristics (e.g. debt portfolio composition, balances, minimum monthly payments, credit score, and income) change before and after the event of default on a given debt or set of debts.

3.3.2 Transition into Severe Delinquency

Figure 12 shows the default rate by age and the debt portfolio held by borrowers.[10] There are several striking results. First, borrowers with multiple types of debt typically only default on one, and borrowers with only student loans default at a higher rate than borrowers with only revolving credit. Next, regardless of the debt portfolio, the rate of defaulting on student loans is at least as high as that of revolving credit. Finally, both student loan defaulters and revolving credit defaulters tend to be young adults. Although the young age of student loan defaulters is unsurprising given that Looney and Yannelis (2015) show that most delinquency and default on student loans happens shortly after a borrower enters repayment, the similarity in the ages of revolving credit defaulters is both interesting and puzzling.

Student loan defaulters and revolving credit defaulters appear similar on observables beyond just age.[11] On average, in the quarter they default both have only three years of credit history, subprime credit scores, low incomes, and high debt-to-income and payment-to-income ratios. However, the two types of defaulters differ in their credit scores and amount of deferred student loan debt.

[9] The data does not allow us to distinguish between federal and private student loans, however, since 92 percent of all student loans are federal, we treat all student loans as if they are federal.

[10] The appendix includes analysis with auto loans and a discussion about the life cycle.

[11] The Appendix shows full summary statistics for each type of debt one quarter before defaulting.

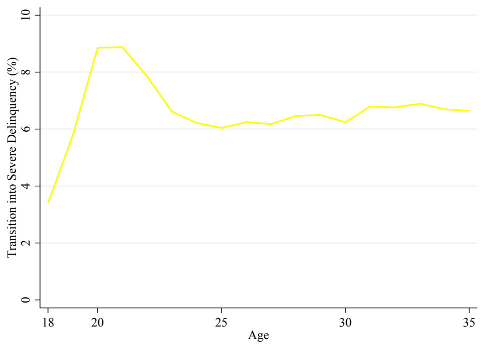

(a) Student

(b) Revolving

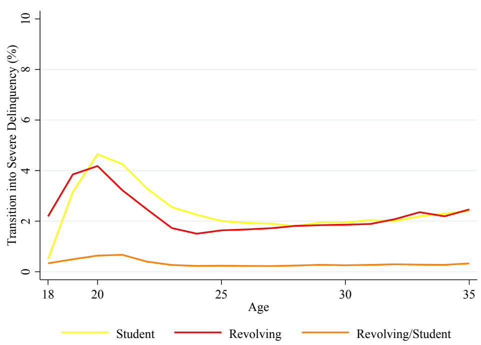

(c) Revolving/Student

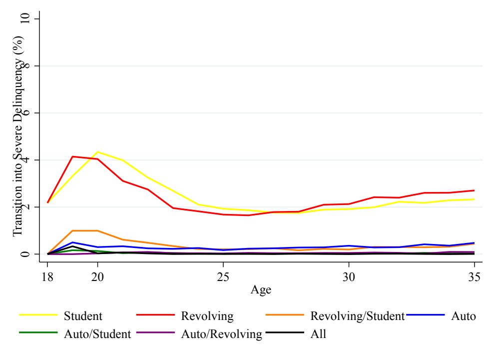

(d) All

Figure 12: Transition into Severe Delinquency along Extensive Margin

Note: Panels represent the types of debts in repayment. Lines show the percentage of borrowers at each age who transition into delinquency or default on a given debt by not making the minimum payment for over 90 days.

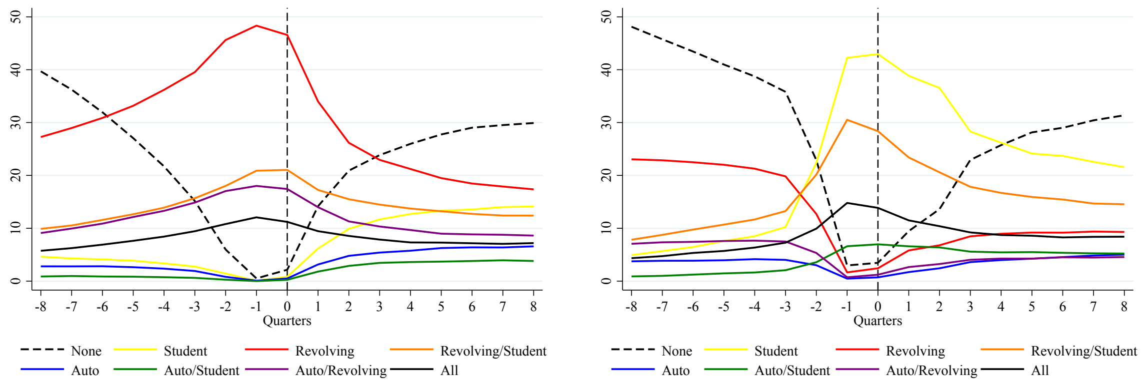

(a) Revolving

(b) Student

Figure 13: Portfolio Composition of Debts in Repayment (%)

Note: Panels represent the debt that transitions into severe delinquency. Lines represent the percentage of borrowers that hold a given portfolio in which each debt is in repayment.

Figure 13 shows further evidence that borrowers only default on one debt. For a given portfolio of debts in repayment, it shows the proportion of agents who default on each type of debt. Interestingly, a substantial number of borrowers no longer hold any debt eight-quarters post-default. This may be due to default causing young borrowers to lose access to (or be priced out from) credit, or enter into a nonrepayment status, such as deferment or forbearance on their student loans.[12]

Table 10: Portfolio Composition One Quarter Before Event by Debt Type

Debt	None	STU	REV	REV_STU	AUT_STU	AUT_REV	ALL
Revolving	0	0	23	46	0	10	21
Student	2	43	1	31	7	0	15

Note: Rows show the percentage of borrowers who transition have made under the minimum payment for over 90 days for each type of debt, split by the overall debt portfolio held. STU is a portfolio with student debt, REV is a portfolio with revolving credit, and AUT is a portfolio with auto loans.

Table 10 shows, for each debt type, the percentage of borrowers holding a given portfolio one quarter before default. It shows that almost half of student loan defaulters only have student loans and no other types of debt in the pre-event quarter, while almost half of revolving credit defaulters

[12]See the U.S. Department of Education's Office of Federal Student Aid web page Get Temporary Relief.

have both student loans and revolving credit. In other words, most borrowers who default on student loans only have student loans, while most borrowers who have both student loans and revolving credit opt to default on revolving credit. However, one-third of borrowers with both types of loans instead opt to default on student loans, which is the behavior we are interested in.

3.3.3 Event Studies

To examine the dynamics of borrowers' debt portfolio and financial characteristics around the event of default, we use the following model:

$$Y_{i,j,t} = \alpha + \sum_{q=-8,q\neq-1}^{Q=8} \beta_q \mathbb{1}_{(q=t-\bar{t}_{i,j})} + \mathbf{X\Gamma} + \iota_i + \tau_t + \varepsilon_{i,t}.$$

$Y_{i,j,t}$ represents borrower i's type of debt j in quarter t and $\bar{t}_{i,j}$ denotes borrower i entering severe delinquency (or default) on debt-type j in quarter t. $\mathbf{X}$ is a vector of controls, including demographic characteristics like age and income and portfolio characteristics like the types of debt held. Borrower ι and time τ fixed effects control for time-invariant, unobserved heterogeneity. We cluster standard errors at the Zip code to adjust for unobserved correlation common to borrowers in the same space such as prevailing local economic conditions.

Our analysis focuses on the impact of default on credit score, total debt balance across all open accounts, debt-to-income ratio, payment-to-income ratio, and utilization rate across all open revolving accounts. The coefficients of interest β_q are the difference in the dependent variable compared to the baseline of one quarter before the event. This baseline is very similar for both student loan defaulters and revolving credit defaulters. Additional analysis and discussion can be found in the appendix for each event study.

3.3.3.1 Credit Score

Figure 14 shows the evolution of credit scores around the first transition into severe delinquency, or default, on auto loans, revolving credit, and student loans.

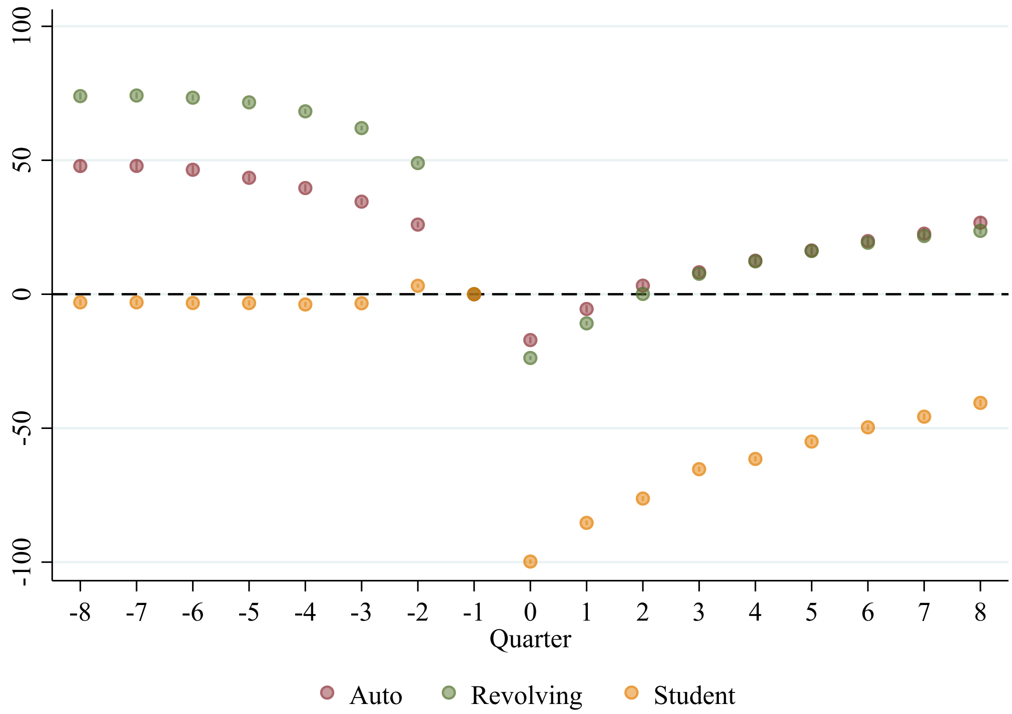

Figure 14: Overlaid Coefficient Plot: Credit Score

Note: The pre-event quarter coefficient is zero for all debt types. We control for age, log income, the debt portfolio held, and whether the debt is in repayment status. The y-axis shows the point difference from the pre-event quarter.

The downward sloping pre-tend for revolving credit shows that credit score is a good predictor of default risk for this type of debt. However, the credit score of student loan defaulters does not predict their risk. However, both types end up with a similar credit score at the event time. Defaulting on student loans has a harsher immediate hit to the credit score, decreasing their credit score on average by 100 points compared to the average decrease of 24 points seen for revolving credit. However, revolving credit defaulters already lost approximately 100 points from their 8-quarter-before-event credit scores of approximately 575, the full penalty is similar overall.

Over the full post-event period, revolving credit defaulters' credit scores increase by approximately 45 points compared to a higher increase of 60 points for student loan defaulters' credit scores. Many factors can potentially explain why student loan defaulters are able to recover their credit score quicker, such as forbearance or alternative repayments plans and other programs available for federal student loans that do not exist for revolving credit, or differences in borrower behavior and financial characteristics after the event, such as inquiring and opening new accounts, repayment, or revolving credit utilization.

3.3.3.2 Debt-To-Income Ratio

Figure 15 shows how the (total) debt-to-income ratio, which helps determine the credit risk of borrowers, evolves around their first default. The appendix presents a related analysis using the payment-to-income ratio, a positively correlated but distinct measure of credit risk.

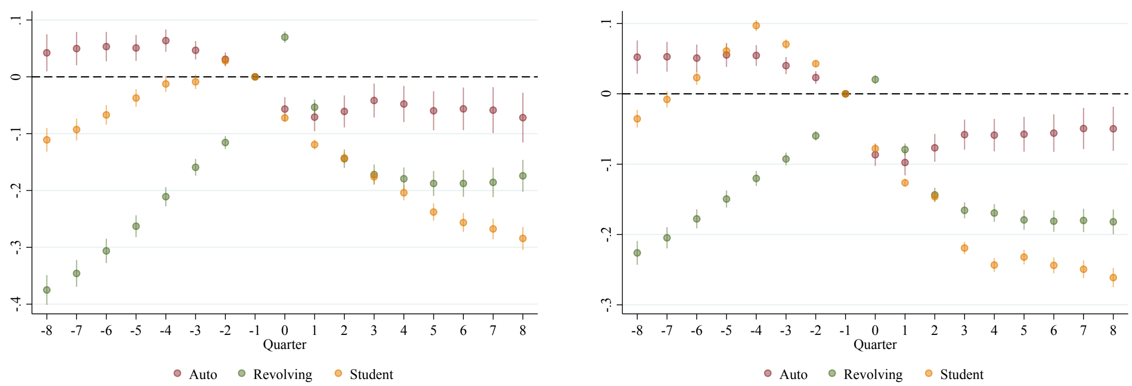

(a) Excluding Deferred Student Loans

(b) Including Deferred Student Loans

Figure 15: Overlaid Coefficient Plot: Debt-to-Income Ratio over All Open Accounts

Note: The pre-event quarter coefficient is zero for all debt types. We control for age, the debt portfolio held, and whether the debt is in repayment status. The ratio is transformed using the natural logarithm so that the y-axis is interpreted as approximately $(\exp(x) - 1) \times 100$ percent change from the pre-event quarter.

Comparing the two graphs suggests that the large increase in revolving credit defaulters' debt-to-income ratios is due to increased utilization of revolving credit, rather than deferred student loans entering repayment.[13] In contrast, the wave patterns for the debt-to-income ratio that include deferred student loans for student loan defaulters is primarily driven by a continuous increase of income starting four quarters before the event.

[13]We analyze this and the impact of wage changes further in the appendix.

3.3.3.3 Revolving Credit Utilization Rate

Figure 16 shows the evolution of the utilization rate on revolving credit. Revolving credit is unique in its ability to be fungible across accounts, meaning borrowers can pay off debt in one account by borrowing up to their credit limit in another account.[14]

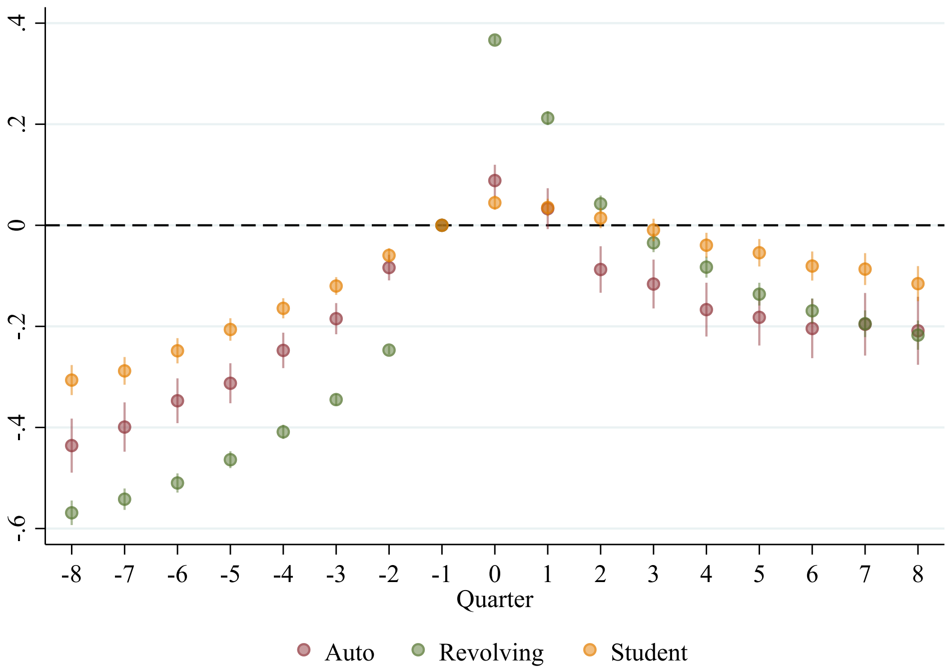

Figure 16: Overlaid Coefficient Plot: Revolving Utilization Rate

Note: The pre-event quarter coefficient is zero for all debt types. We control for age, log-income, the debt portfolio held, and whether the debt is in repayment status. The ratio is transformed using the natural logarithm so that the y-axis is interpreted as approximately $(\exp(x) - 1) \times 100$ percent change from the pre-event quarter.

The utilization of revolving credit increases by approximately 57 and 31 percent for revolving credit defaulters and student loan defaulters, respectively, during the eight-quarter-pre-event period. The discontinuity at the event for revolving credit is likely due to closure of the account they defaulted on, thereby reducing the credit limit associated with any given balance. After default, the utilization rate for student loan defaulters drops the least, although the number of inquiries for

[14]In principle, student loans can be used to make payments on other debts like auto loans and revolving credit as both are valid living expenses; however, this is only true while the borrower remains in school.

revolving credit remains lower than that for revolving credit defaulters. Further analysis is required to determine if this is caused by reduced access to credit after default or by selection.

3.4 Survey

We conduct a survey consisting of three parts and a demographic survey to gain a deeper understanding of how borrowers choose which debt to default on. We aim to help explain the decision to default on revolving credit over student loans made by one-third of financially constrained borrowers in Section 3.3 who hold both types of debt by documenting that intertemporal choice plays a role in default decisions and exploring how they are influenced by financial literacy and the inherent intertemporal tradeoff.

Our sample consists of 98 participants from two online survey platforms, MTurk and Prolific. Participants earn a fixed-rate payment plus a bonus based on one randomly selected part, for an average payment of \$3.64. The survey takes 13 minutes to complete on average. Our sample is 57 percent male[15]. The average age of participants is 39 years old, and the majority are employed full time with a four-year degree, earning a median average annual salary of around \$60,000. Furthermore, 87 percent of participants have credit card debt, and 39 percent have student loans. It is important to note, however, that the analysis in Section 3.3 focuses on a significantly younger sample who are consequently more financially constrained and hold a different portfolio of debt.

3.4.1 Part One - Characteristics of Debt

Part One of our survey comprises of two questions that aim to identify characteristics of different types of debt that can influence borrowers' default decisions using two questions.

3.4.1.1 Decisions

Question One is a free-response question in which participants identify the two most important characteristics of debt to consider if a borrower cannot make the minimum payment on all of

[15]The MTurk sample is 65 percent male, while the Prolific sample is gender-balanced.

their debt and must decide which uncollateralized debt to default on.[16] While coming up with characteristics themselves relies on participants' financial literacy, the same is true for borrowers who actually face the decision. In Question Two, participants rank five factors based on how influential they are in borrowers' default decisions. The five factors are listed below in random order:

- Size of Penalties: Prioritizing the debt with more severe overall penalties, regardless of when they happen
- Timing of Default: Prioritizing the debt that is closest to default
- Size: Prioritizing the debt with the larger balance
- Interest Rates: Prioritizing the debt with the higher interest rate
- Timing of Penalties: Prioritizing the debt with more immediate penalties

The two questions provide insights into different thought processes. The free response asks participants how *they* would make the decision, while Question Two elicits second-order beliefs on what others find important by paying $0.50 for each ranking that matches the modal response.

3.4.1.2 Results and Discussion

Our analysis focuses on the rankings from Question Two. Figure 17 shows the full distribution of ranks for each factor, which significantly differs by term (Kruskal-Wallis equality of populations rank test ($\chi^2 = 18.686, p < 0.01$). Almost half of participants rank timing of default as the most important factor in repayment decisions, clearly demonstrating our hypothesis that default time horizon impacts intertemporal choice.[17] Interestingly, the time horizon of penalties is ranked much lower, with an average ranking of third important and only 17 percent rating it as the most important. This shows that participants view default as its own distinct event.

Compared to those who have not defaulted on student loans, the 15 individuals in our sample who have directionally rank the size of the debt as mattering more ($diff = 0.516, t = 1.381, p = 0.170$), the interest rate as mattering more ($diff = 0.170, t = 0.420, p = 0.338$), the size of the penalties as mattering less ($diff = 0.338, t = 1.023, p = 0.154$), the timing of the penalties as

[16] We discovered during the data collection that some participants were using search engines to generate their answer, so we replaced all text in the survey with screenshots to mitigate this issue.

[17] The is no impact of the time preferences measured in Part Three (Spearman's $rho = -0.039, p = 0.742$)

mattering more ($diff = 0.128, t = 0.328, p = 0.372$), and the timing of default as mattering less ($diff = 0.475, t = 1.060, p = 0.291$). The results of this test may be biased by the small sample of student loan defaulters and the differences in comparison groups compared to Section 3.3.[18] However, it still provides insight into the priorities of those who have experienced our decision of interest firsthand. Our findings suggest that student loan defaulters prioritize reducing interest payments and other current penalties rather than reducing the severity of default risk and eventual penalties, which aligns with the observed behavior.

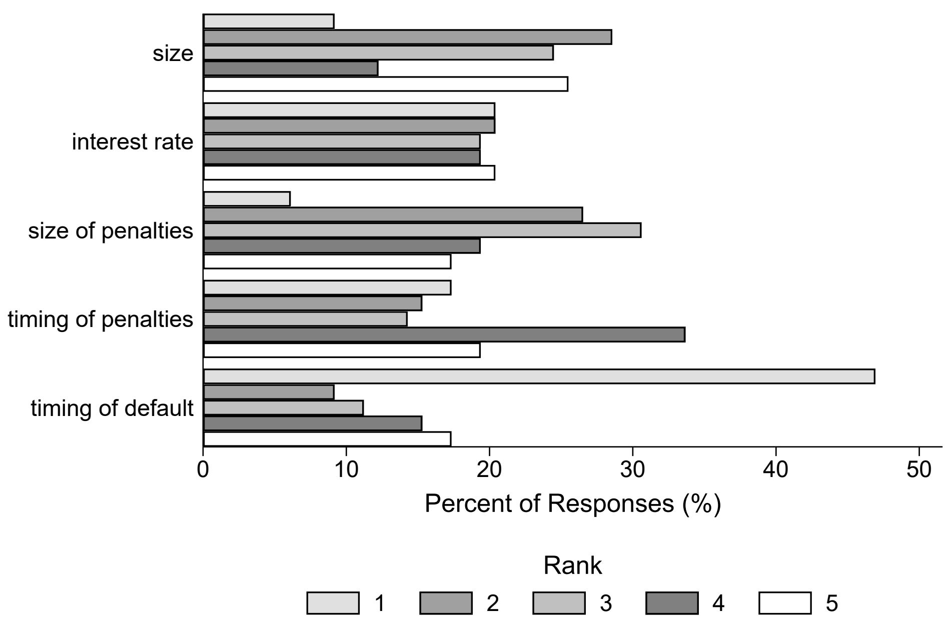

Figure 17: Ranking of Debt Characteristics in Default Decisions by Importance

These differences may represent an underlying difference in preferences without necessarily being a result of behavioral biases. However, we elicit beliefs over what is *influential* to default decisions, rather than on normative views of what *should* impact default decisions. This means that student debt defaulters may be sophisticated or have self-reflected, leading them to better understand the myopic prioritization that led to student loan default.

[18]Here they are older and without financial constraints, there they are other young, financially distressed borrowers.

3.4.2 Part Two - Financial Literacy

Part Two is a financial literacy test. Participants have two minutes to classify ten items as a potential consequence of student loan default, credit card default, both, or neither. They earn \$0.25 per correct answer. This section aims to assess participants' understanding of default consequences, providing insight into the relationship between financial literacy and default decisions.

Financial literacy can play a crucial role in default decisions by influencing borrowers' evaluation of the costs and benefits associated with defaulting on different types of debt. A misspecified model of the disutility associated with defaulting on student loans versus defaulting on revolving credit can lead to suboptimal decision-making. For example, if someone does not understand the severity of defaulting on student loans, they may prioritize preserving their current access to credit, even if they would avoid default on student loans if they knew the true consequences.

While we don't measure the cause of differences in financial literacy, we theorize that information acquisition decisions may be influenced by the differences in timelines between the two types of debt. Despite similar availability of information, individuals may be more motivated to seek out information regarding the earlier default (i.e., revolving credit). Niativety about the negative long-run consequences of student loan default can therefore serve as an excuse for borrowers to prioritize revolving credit, such as is shown in my first chapter.

3.4.2.1 Decisions

The ten items to classify,, sorted by their correct answer, are:

- **Student Loan Default Only**: Money taken from your paycheck without a court order; the IRS does not give you your tax refund
- **Revolving Credit Default Only**: Debt can be discharged through bankruptcy; Increased interest rate on this debt; Debt gains interest daily
- **Both**: Debt sent to a private collection agency; Decreased credit score
- **Neither**: Repossession of your degree and/or the items you purchased; Loss of access to social services such as food stamps; Imprisonment until the debt is repaid

3.4.2.2 Results and Discussion

Figure 18 compares participants' financial literacy of student loans to revolving credit. We define a correct match if the individual identifies if an item is or is not a potential consequence of the debt correctly, without differentiating between that category/both or the other category/neither.

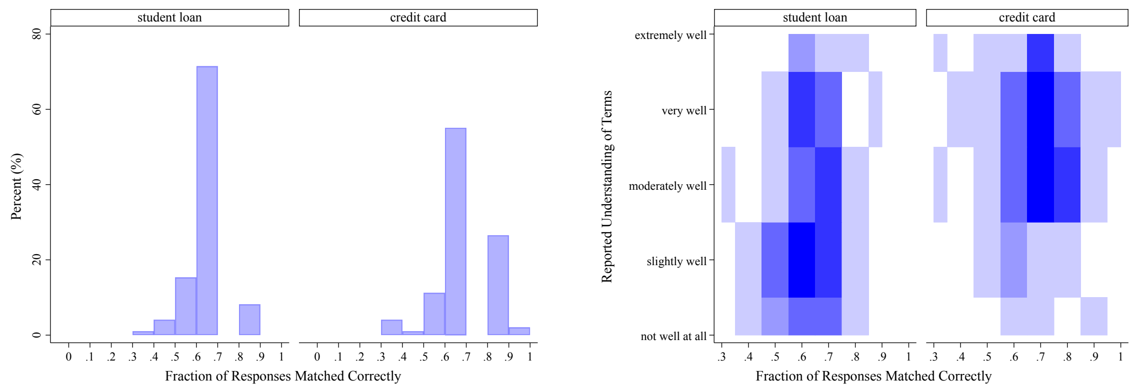

(a) Distribution of Financial Literacy by Debt

(b) Self-Reported vs Actual Financial Literacy

Figure 18: Comparison of Financial Literacy Across Debt Types

Note: Actual financial literacy is determined by correctly categorizing items as potential consequences of defaulting on that debt term. Self-reported financial literacy is reported on a five-point Leikert scale.

Our analysis reveals that, on the whole, our sample displays greater financial literacy for revolving credit than for student loans, with individuals themselves being aware of this superior knowledge. However, this outcome is not necessarily surprising as a larger proportion of individuals in our sample hold revolving credit in their portfolio than student loans. Given this, we know analyze our findings while conditioning on portfolio type.

Figure 19 shows financial literacy split by debt portfolio. Our results demonstrate that all portfolio types are more knowledgeable about the consequences of defaulting on revolving credit than student loans, even amongst those who have both student loans and revolving credit and therefore have a clear incentive to understand both types of debt.Our findings open up questions for future exploration into the mechanism behind the information acquisition decisions of agents.

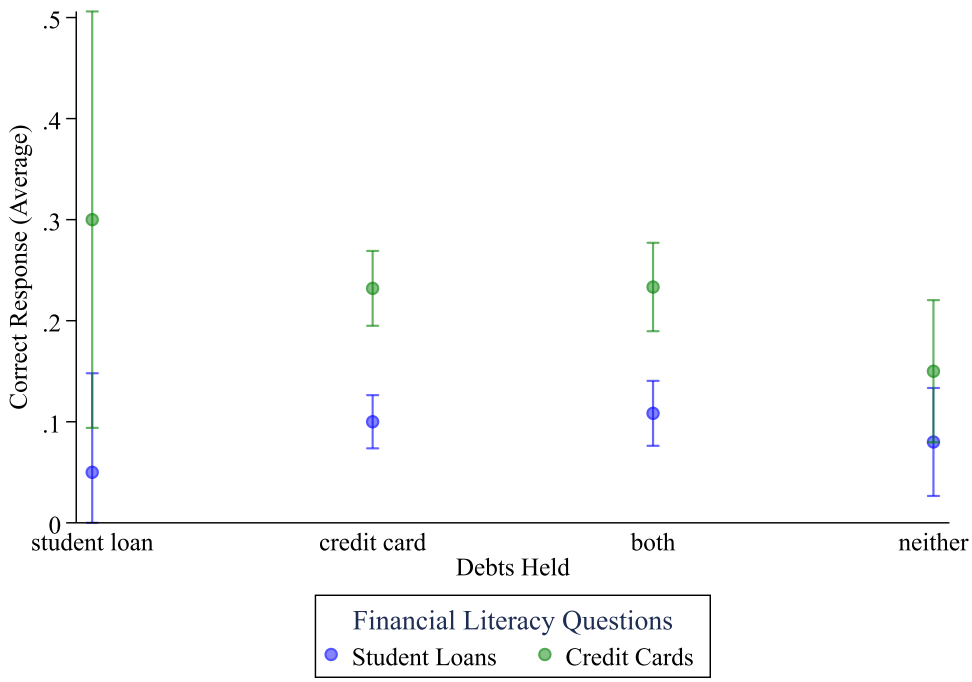

Figure 19: Distribution of Correct Answers by Type of Debt

Note: Financial literacy is assessed through the number of correct classification of potential consequences for defaulting on that debt type.

3.4.3 Part Three - Intertemporal Risk Preferences

Part Three elicits time and risk preferences using three price lists to study the intertemporal risk financially-constrained borrowers face.[19] While our findings show an ex-post preference for defaulting on student loans over revolving credit, it is more accurate to say borrowers prioritize avoiding immediate penalties at the expense of increasing their future default risk.

Given the time-frame and the severity of the penalties, this would suggest that these agents are either extremely risk-seeking or have a very large discount factor – unless the presence of future risk is serving as an excuse to enable borrowers to make *potentially* myopic decisions. This extends the framework of Exley (2015), where the presence of a tradeoff between oneself and others creates differential risk preferences that increase selfishness, to the intertemporal choice domain.[20]

[19]See Ahlbrecht and Weber (1997); Andreoni and Sprenger (2012b) for work on risk/time preferences interactions.

[20]Drucker and Kaufmann (2022) explore the role of risk in excuse-driven present bias by using real-effort tasks and adding the probability of having an easier task.

3.4.3.1 Decisions

Bonuses for this part are paid in two time periods: t_0, the same day, and t_1, 15 days after. This delay was chosen to ensure that $t = 1$ fell into a different payment period for online workers who receive their payments at different frequencies.[21] There was a guaranteed bonus of $0.10 on both payment dates, thus reducing corner solutions that may arise from participants attempting to minimize transaction costs.

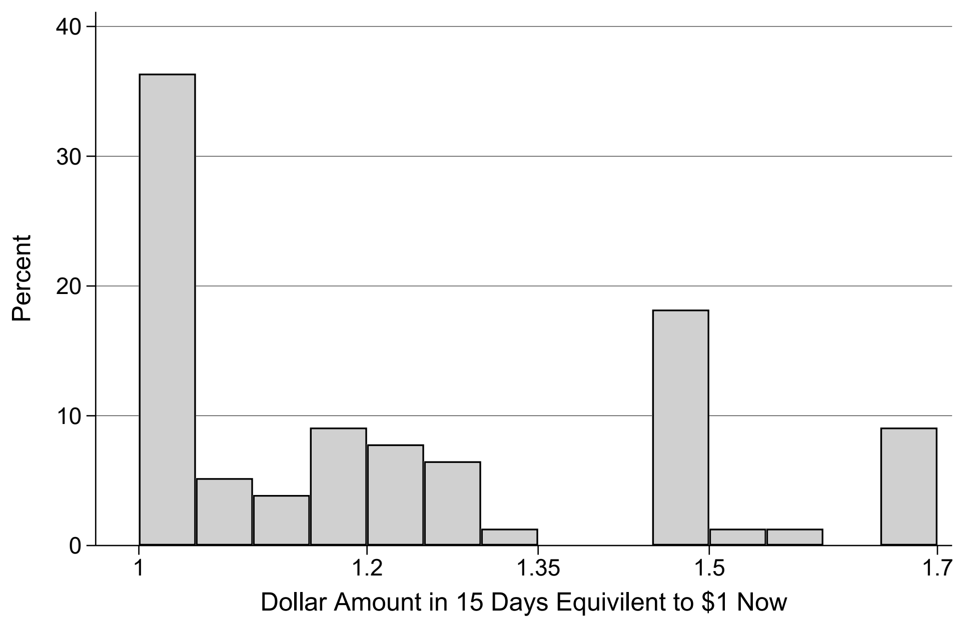

Figure 20: Histogram of Participants' Indifferent Point δ

Note: δ is one cent above the highest amount 15 days in the future that an immediate payment of $1 is preferred to.

Price List One elicits time preferences by presenting participants with a choice between a certain payment at the present time ($t = 0$) or a delayed payment of $1 and some cents in 15 days ($t = 1$). The first row is a choice between $t_0 = \$1$ and $t_1 = \$0.95$, which increases by $0.05 in each of the subsequent 15 rows.[22] We use this list to define $\delta = \$1 + \$\frac{Row}{100}$, where *row* refers to the last $t = 1$ payment that $t_0 = \$1$ was preferred to. To account for the inability to determine the exact indifference point, one cent is added to δ. The full distribution of δ is shown in Figure 20, with $t_0 = \$1$ being equivalent to $t_1 = \$1.17$ on average.

[21] While possible to cash out every day, only 3 percent of our sample reports doing so. 74 percent cash out every 3, 7, or 14 days, and the remainder cash out every 30 days.

[22] Fourteen participants made choices with row differences of $0.10 and are therefore excluded.

We use Price List Two to assess the risk preferences of our participants.[23] This list requires them to choose between receiving a certain payment of \$1 or a lottery for \$2 with varying probabilities, both paid in $t = 0$. The probability of winning is zero percent in the first row and increases by ten percent in each of the subsequent 11 rows. Price List Three is also used to elicit risk preferences, but with an added intertemporal tradeoff. In this list, the certain payment remains fixed as \$1 in $t = 0$, while the risky payment is a lottery for $2 * \delta$ payable in t_1. In theory, this value is equivalent to $t_0 = \$2$, i.e., the lottery amount in Price List Two. Therefore, without the presence of excuse-drive differential risk preferences, we would expect their risk preferences to be similar across both lists.[24] This choice maps into default decisions by measuring the amount of additional risk of *losing* the lottery subjects will take on to avoid losing the certain payment now. However, since our price lists are not in the loss domain, these will excuses will increase risk-aversion rather than risk-seeking behavior.

3.4.3.2 Results and Discussion

Our sample displays high levels of risk-aversion, in line with the results of other online studies. Specifically, 78 percent of participants opted for the certain payment of \$1 over a fifty percent chance of winning \$2, indicating risk aversion.[25] We calculate the expected value of the lottery they are indifferent to using the midpoint between the highest percentage for which the certain payment of \$1 was preferred and the lowest percentage that was preferred to the certain payment of \$1. This represents the price-differential needed to compensate participants for the risk of the lottery. On average, the expected value of the lottery is \$1.22.

Figure 21 displays a scatterplot comparing the row in which each participant switches to preferring the lottery to the certain payment in both Price List Two and Three. The dashed line represents the 45-degree line, representing having identical risk preferences in both lists. Overall, 63 percent of participants switch to the same row in both price lists. This increases to 95 percent

[23] We exclude anyone who makes a dominated decision or switches multiple times in any price list. Seven participants are excluded as they never switch in Price List One, making it impossible to determine an indifference point for them. This brings the effective sample size to 84 participants.

[24] Abdellaoui et al. (2011) find that individuals have the same discounted risk preferences across time, however, all lotteries they evaluate have the certain payment happening immediately.

[25] 40 percent switch either directly before or after the 50 percent lottery, which may better be classified as risk-neutrality. In that case, 51.52 percent of participants are risk averse.

when considering those who switch within one row of each other. However, it is important to note that switching within one row encompasses a 20 percentage point range. Only 4.55 percent of participants switch to a later row in Price List Three by more than one row, while 3.56 percent switch by more than one row earlier.

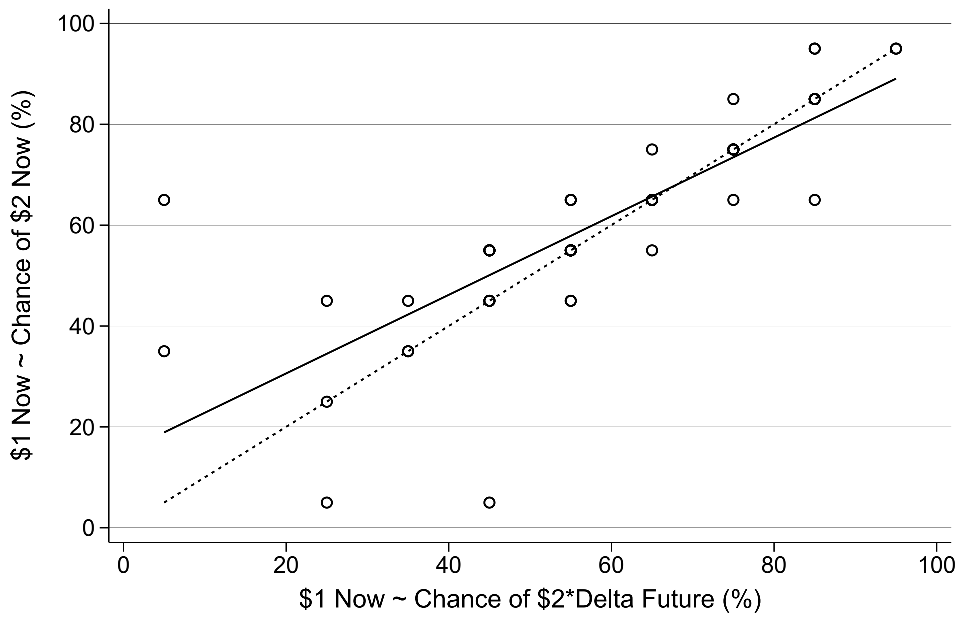

Figure 21: Scatterplot of the Row Participants Switch At in Price List Two and Three

Note: Dots show the row that participants switch from preferring \$1 now to preferring a lottery with a $(row * 10)$ percent chance of \$2 now or $\$2 * \delta$ in the future, with $\$\delta$ being the future amount they are indifferent to \$1 now.

While our analysis does not find evidence to support our hypothesis of present-biased differential risk preferences, we acknowledge several assumptions we make that could potentially influence our results.[26] First, the difference in the absolute size of the lotteries may lead participants to appear more risk-seeking. We explore this in Figure 22 by comparing the distribution of the expected values implied by decisions in each price list. When normalized to $t = 0$ dollars, the expected value from Price List Three is nearly identical to that of Price List Two, as shown in the left graph. In contrast, the nominal expected values, shown in the right graph, show *more* risk-seeking behavior.

[26] Some assumptions seem unlikely to impact our results, such as measurement error caused by the course measurement of δ, which would bias participants away from having the same switchpoint.

Additionally, 26 percent of participants have a $\delta = \$1.01$. As they seem to discount very little and are facing almost identical lotteries, it is unsurprising that almost all of these participants switch on the same row for both risk price lists, making up 52 percent of participants with the same risk presences. For other participants with larger δ, the 8-percentage-point range of the true probability may reduce our ability to detect differences. While we use the mid-point for both price lists, it is possible that some participants have a higher indifference point in Price List Three that just isn't different enough to merit switching to another row.

Our analysis assumes that *intra*-temporal risk preferences do not differ across time. If individuals are generally more risk-seeking in terms of future money, then switching at the same row *does* show an increase in risk aversion caused by the introduction of an intertemporal tradeoff. It is also important to consider the limitations of our analysis not related to measurement error as potential alternative explanations for the observed default decisions, such as difference of loss-aversion due to differential framing of the gain-loss domain or the differences in time-horizons and scale of the decisions. Moreover, the differing characteristics between our sample and that in Section 3.3 may influence both time and risk preferences. The student loan defaulters are much younger and more credit constrained, potentially increasing risk-seeking and leading to more present-biased preferences in monetary domains (Carvalho et al., 2016b; Schildberg-Hörisch, 2018).

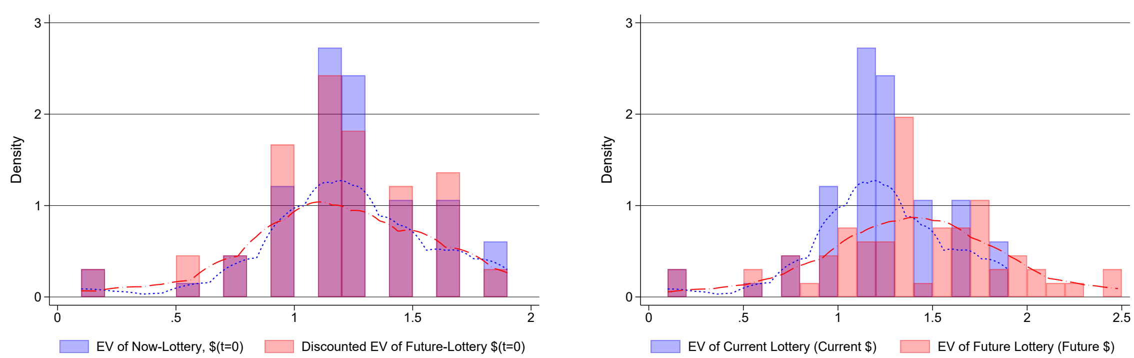

(a) Real Expected Value

(b) Nominal Expected Value

Figure 22: Expected Value of Lotteries which are Indifferent to $1 Now For Sure

Note: Expected values use the midpoint between the highest probability lottery to which $1 now is preferred and the lowest probability lottery preferred over $1 now. Real expected value is normalized using elicited time preferences.

There are additional aspects of real-world decision making that our survey does not capture. For example, the large and important stakes associated with default decisions, such as the immediate loss of access to a line of credit, may increase motivated reasoning, increasing the demand for excuses or general over-optimism (Breig et al., 2021). Moreover, one's actions can reduce the probability of future default, giving another potential avenue for borrowers to bias their beliefs not present in a simple lottery. Therefore, while finding present-biased differential risk preferences in our survey would indicate its influence on default decisions, not finding them simply shows that further research is needed to better understand the complex interplay between risk preferences and intertemporal decision-making in the context of debt default.

3.5 Discussion

In this chapter, we explore consumer finance decisions made by financially constrained borrowers with multiple types of debt who cannot afford all of their minimum payments. We are among the first to look at cross-debt decisions and explore the influence of different terms of debt. We focus our analysis on student loans and revolving credit, two types of uncollateralized debts.[27] Specifically, we look at how the different time horizons for defaulting and therefore incurring penalties impacts the decision of which debt to default on.

While revolving credit goes into default after fewer missed payments, resulting in the closure of that credit line and therefore an immediate loss in access to money, student loans have much more severe, albeit delayed, penalties. These include garnishment of wages and the loss of all future tax returns and social security payments. Despite these harsh penalties, we find that a substantial proportion of borrowers in our observation credit report data set who have both student loans and revolving credit end up defaulting on the former.

We explore the mechanisms behind these costly financial decisions using an online survey. We find that individuals overwhelmingly indicate that the timing of default is the most influential factor for default decisions made by financially constrained borrowers; that financial literacy is much higher for revolving credit than for student loans, even amongst those who hold both; and that there are no differential risk preferences depending on the presence of an intertemporal tradeoff.

[27] We also show that borrowers prioritize avoiding default on loans with collateral.

This chapter opens up an avenue for future research to continue exploring how behavioral biases can influence these types of decisions. While default decisions may result from extreme time discounting or preferences for minimizing interest and closing out accounts with lower balances, there is no doubt that the intertemporal aspect also plays a role. Behavioral biases such as present bias, over-optimism and over-confidence in one's ability to change their financial situation in the future, and motivated reasoning in the form of differential risk preferences or motivated information avoidance may influence these default decisions as well.

Understanding the mechanisms behind these default decisions is important as debt account for a large portion of the American economy. Suboptimal repayment and default decisions carry harsh financial penalties that can negatively impact financial well-being and credit poverty traps. Understanding the root cause of these decisions can inform policy about how to improve financial decision making on an individual level as well as government policy about how to structure the terms of debt to improve welfare.

Appendix A Excuse-Based Procrastination

A.1 Summary Statistics

A.1.1 Attrition

Table A.1.1 lists summary statistics for the sample divided by if they returned for the second session (first column), did not (second column), and for the overall sample. For each variable, the mean is listed with standard errors in parenthesis below. Two-sample t-tests were run to test if the subjects returning the second week were significantly different than those who did not return.

Age is measured in years. *Gender* can be interpreted as the percentage female.[1] The next group of five questions were answered using a five-point Leikert scale, with 0 coded as “Strongly Disagree” and 4 coded as “Strongly Agree”. *Careful Decision* was the statement “I made each decision in this study carefully.”. *Random Decisions* was the statement “I made decisions in this study randomly”. *Understood* was the statement “I understood what my decisions meant for how many jobs I’d have to complete this week and next week.” *Busyness Today* and *Busyness Next Week* were questions asking, respectively, “How busy (are you this week)/(do you expect to be next week)?”. The following five averages are of Week One performance. Time is the number of seconds from first seeing the transcription to submitting an acceptable answer. Attempts in the number of times the participants submitted an answer, which would be 1 if their first answer was acceptable. Resets is the number of times participants hit the “Reset” button to fully clear the Greek letters they had submitted. Back is the number of times the participant hit the “Back” button to clear the last Greek letter they had submitted. Accuracy was the percent of Greek letters submitted which were in the correct spot when an acceptable answer was submitted.

Gender is the main characteristic that differs between those who do and do not return the second week, as all 6 participants who did not return were male. Table A.1.2 shows how men and women differ in similar summary statistics. On the surface, those who returned appear older. However, this is due to an outlier and a small sample: one participants who did not return reported

[1] The demographic question had a third option, “Other or Non-Binary” in addition to male or female, but no participants identified as such.

Table A.1.1: Summary Statistics By Attrition Status

	Returned	Attrition	Total
Age	19.90	23.17**	20.15
	(0.144)	(3.177)	(0.279)
Gender	0.556	0***	0.513
	(0.058)	(0)	(0.058)
Careful Decision	3.097	3.333	3.115
	(0.095)	(0.494)	(0.094)
Random Decisions	0.500	1.333*	0.564
	(0.106)	(0.843)	(0.117)
Understood	3.681	3.333	3.654
	(0.083)	(0.666)	(0.090)
Busyness Today	2.861	2.833	2.859
	(0.108)	(0.477)	(0.105)
Busyness Next Week	2.708	2.500	2.692
	(0.114)	(0.562)	(0.112)
Avg Time (Sec)	52.54	52.24	52.51
	(1.198)	(2.654)	(1.121)
Avg Attempts	1.059	1.033	1.057
	(0.016)	(0.033)	(0.015)
Avg Resets	0.102	0.0500	0.0982
	(0.016)	(0.034)	(0.015)
Avg Back	1.083	1.272	1.097
	(0.141)	(0.471)	(0.135)
Avg Accuracy	98.58	99.41	98.64
	(0.317)	(0.274)	(0.294)
Num Tasks (W1)	11.166	11.166	11.166
	(0.243)	(0.872)	(0.233)
N	72	6	78

Mean of each variable, SE in parenthesis

* p<0.10 ** p<0.05 *** p < 0.001

being 39 years old, while the other 5 reported being between 19-21. Only undergraduates between the ages of 18-22 were invited to participant and this participant reported being in their Sophomore year, so it is seems likely that they mis-entered 19. Without this outlier, the two samples do not differ by age.

Those who did not return reported making their decisions more randomly, with their answer to that statement increasing from “Strongly Disagree” to “Disagree”. Two participants who did not return selected “Strongly Agree”, while only one participant who did return selected that option. The remaining 2/3rds of those who did not return all selected “Strongly Disagree”. Of those who returned, 68% selected “Strongly Disagree”, and an additional 23% selected “Disagree”. This is some suggestive evidence that those who did not return put less thought into the study in general. This difference in decision making did not correspond with differing responses to if they made careful decisions, and did not seem to stem from misunderstanding the questions.

Task performance and self-reported busyness do not differ between those who do and do not return for week two, reducing doubt that attrition was caused by a decision made prior to take choices. Since the difference between those who did and did not return are minimal, the results presented in the paper do not include those who did not return. However, all results are robust to including all individuals.

A.1.2 Gender

Table A.1.2 shows Table A.1.1, but split by gender rather than attrition status. Women are, on average, about one year older than the men in the study, and complete the tasks almost four seconds slower. The difference in speed does not correspond to an increase in accuracy, resets, backs, or attempts. Overall, this would add under a minute to the first session and between 1 - 1.5 minutes to the second session, so although the difference is significant it does not lead to a large, in magnitude, difference in time spent in the lab. Women, however, do self-report making their decisions more carefully than their male counterpoints. However, they don’t report acting significantly different in terms of self-reported randomness of the decisions, unlike the attritors. This seems to suggest that randomness was more associated with not returning, rather than gender.

Table A.1.2: Summary Statistics by Gender

	Female	Male	Difference	se
Age	20.631	19.7	.931*	.551
Careful Decision	3.421	2.825	.596**	.178
Random Decisions	.3684	.75	-.381	.232
Understood	3.631	3.675	-.043	.183
Busy Today	2.973	2.75	.223	.211
Busy Next Week	2.578	2.8	-.221	.226
Avg Time	50.597	54.333	-3.736*	2.217
Avg Attempt	1.046	1.067	-.021	.030
Avg Reset	.085	.110	-.025	.030
Avg Back	1.107	1.088	.019	.272
Avg Accuracy	98.735	98.550	.184	.593
N	40	38	78	

A.1.3 Task Performance

All participants saw the exact same tasks, i.e., the same row of blurry greek letters, in the same order. Performance did not improve throughout the session, which most tasks taking a little over 50 seconds on average. Table A.1.3 compares participation statistics between the two weeks, of which there is very little difference. Taken together, this shows that there is not learning happening within the task. Therefore it is unlikely that participants would push the task off to the next week expecting to perform better at it.

Table A.1.3: Average performance statistics excluding practice round

Variable	W1	W1	W2
Time (Sec)	52.5	52.3	46.74
Attempts	1.05	1.05	1.03
Resets	.09	.10	.06
Backs	1.09	1.08	.91
Accuracy	98.64	98.57	99.12
With Attrition	Yes	No	-

A.2 Analysis

A.2.1 Order

Participants first see all four Excuse treatment decisions in a random order, then the eight No-Excuse treatment decisions in a random order. The order was pre-randomized at the session level. However, two sessions used the same randomization. Although the assignment was still random, this does make it so that the order of the questions is not fully balanced. Table A.2.1 shows how many participants saw each parameter in which order (i.e. 16 participants made the Baseline decision first, 15 made it second, 27 made it third, and 14 made it fourth).[2]

[2]Table A.2.1 shows the order of all subjects, including the six who did not return for the second week.

Table A.2.1: Order of Questions for Choice-To-Reveal

Order	Baseline	↓ Week One	↓ Week Two	↓ Both
1	18	12	16	32
2	16	22	28	12
3	28	20	12	18
4	16	24	22	16

A.2.2 Impact of Work Schedules on Outcomes No Excuse Treatments

A.2.3 Heterogeneous Classifications

I split my sample into three types: Patient, those who always minimize overall work; Impatient, those who always minimize current work; and Excuse Takes, those who engage in excuse-based procrastination as described in the main text. Table A.2.3 shows how participant decisions map into classification.

Sixty percent of the sample are patient and always minimize total work, getting the information needed to do so. Interestingly, there are no impatient participants in my sample.The five subjects who selected *Less Now* in the No-Excuse Conflicted state then pick *More Now* in the No-Excuse Aligned state. These were participants who mostly had preferences which did not depend on effort, i.e. preferences for the number 13 or odd numbers, or participants who enjoyed the task.

Excuse-takers are only identifiable through my within-subject design as they are identical to the patient participants when there is no excuse, and identical to the impatient ones when there is. About seven percent of my sample made state-contingent decisions to minimize overall work in the No-Excuse treatment, did not reveal the information, but then selected *More Now.* This may suggest the presence of another type: a self-signaler. These participants may be so averse to procrastinating they want to opt out of deciding at all and just pick the potentially patient option.

Table A.2.2: Impact of Work Schedule Conditions in the No-Excuse Treatment

	(1) Minimize Current Work	(2) Select *Less Now*
↓ Week One	0.007	-0.007
	(0.021)	(0.021)
↓ Week Two	0.014	0.000
	(0.017)	(0.010)
↓ Both Weeks	0.007	-0.007
	(0.029)	(0.025)
Constant	0.910^{***}	0.479^{***}
	(0.032)	(0.012)
N	576	576

Note: Standard errors in parentheses clustered at the individual level. OLS estimates from regressing work schedule conditions on an indicator for (in order of columns) minimizing total workload and selecting *Less Now* when the state of the world is exogenously given. Omitted group is *Baseline*. Includes controls for demographic variables. Does not include attrition.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.2.3: Heterogeneous Classification by Decisions

	No Excuse Aligned	No Excuse Conflicted	Revealed The State	Excuse Treatment	Percent
Patient	Less Now	More Now	Yes	*Depends*	59.72%
Impatient	Less Now	Less Now	-	Less Now	0%
Excuse Taker	Less Now	More Now	No	Less Now	22.22%

A.3 Instructions

Click here for experiment and survey instructions.

Appendix B Time and Punishment: Penalty Timing and Cross-Debt Default Decisions

B.1 Data Set

B.1.1 Sampling and Construction of Panel

The sample has two parts: 70 percent random and 30 percent oversampled on ever 90 or more days delinquent (i.e. severe delinquency) or derogatory borrowers. Within each subsample, 80 percent of the borrowers exist in the first quarter of 2004 and 20 percent of the borrowers enter after the first quarter of 2004. An additional 100,000 borrowers, with the same sampling are merged in to the dataset in the first quarter of 2013.

The random sample was constructed to be representative of the population of adults in the United States of America that have a credit file in the first quarter of 2004; thus, it is representative of the population along the dimensions of age, location, and severe delinquency or derogatory history in the first quarter of 2004. Borrowers then enter the panel at approximately the growth rate of the population of the United States of America to maintain representativeness throughout the time-series.

For our analysis, we use the 769,989 borrowers comprising the 70 percent random sample. Wrangling the data, we drop 1,260 borrowers due to global identifier inconsistencies, 38,522 due to deceased information present in the borrower's credit file, 5,019 borrowers due to not being residents of the 50 states, Washington D.C., and Puerto Rico; we then drop 520,501 borrowers that never held student trades, 916 borrowers due to holding trades before the given birthdate, and 106 borrowers due to holding delinquent trades before the age of 14.

The remaining 203,665 borrowers we use in our analysis; approximately 26.5 percent of the original random sample. Figure B.1.1 shows the age distribution of the 203,665 borrowers we use in our analysis by subsample as above described and in total in the first quarter of 2004. Of these 203,665 borrowers, 145 are missing Zip code, 278 are missing State, 10,678 are missing a credit score, and 71,357 are missing income after existing or entering the panel. In parts of the graphical analysis and the entire regression analysis, we further restrict the sample to borrowers

35 years old or younger, that do not originate a mortgage-type trade within eight quarters of their last observation, with a dense panel for +/- eight quarters around the event of transitioning into severe delinquency or derogatory status for the first time, thereby resulting in 179,698 borrowers in 10,885 Zip codes.

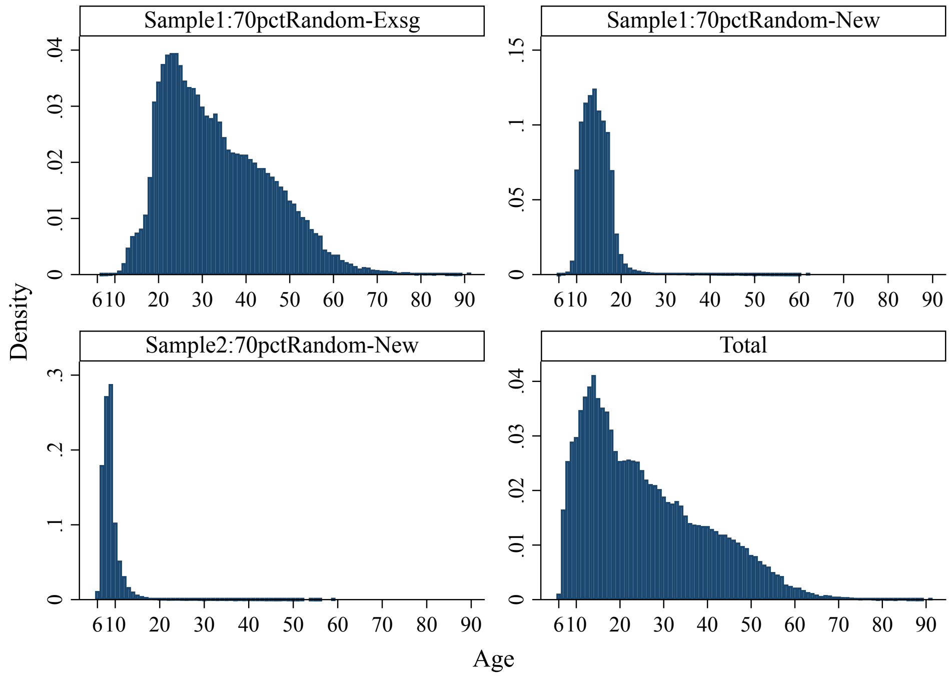

Figure B.1.1: Age Distribution by Subsample in the First Quarter of 2004

B.1.2 Defining and Identifying Events

The events of interest in this analysis are (i) the first origination of any trade, (ii) the first origination of each loan type, (iii) the first repayment on each loan type, and (iv) the first transition into severe delinquency or derogatory status on each loan type. For borrowers that exist at the beginning of the panel, variables for the number of months since the oldest trade was opened and number of months since the most recent 90 or more days delinquency or derogatory is used to calculate the age of the borrower at the first origination of any trade and the first transition into severe delinquency or derogatory status on any trade. The first origination of each trade and the first repayment on each trade cannot be calculated for borrowers that exist at the beginning of the panel if a given trade has already been originated or entered repayment with this data. For borrowers that enter the panel, the age of the borrower at each of these events is directly observed. Only observed events are used in this analysis; and "first" events are known to be the first event.

The event of most interest in this analysis is the first transition into severe delinquency or derogatory status on each trade. we define severe delinquency as 90 or more days delinquent for each trade. Federal student loan borrowers in delinquency are not reported to the credit bureaus until being 90 or more days delinquent. Borrowers in delinquency on auto, revolving, and mortgage-type trades are reported after a required payment has been missed by one or more payment billing cycle, so typically after one month.

However, most non-pecuniary penalties for delinquency or derogatory status do not occur until being 90 or more days delinquent. For these reasons and to provide a consistent measure across trades the event of transitioning into severe delinquency or derogatory status is defined as being 90 days or more days delinquent on at least one trade in a given loan type. Thus, we refer to this as transition into 90 or more days past due along the extensive margin. It is possible to be both current and delinquent on some trades within a given loan type, but we are not directly able to measured this consistently across loan types.

For auto loans or leases, mortgage-type trades (i.e. first mortgages, second mortgages (or home equity loans), and home equity lines of credit), and revolving trades we measure transition into severe delinquency using variables for worst present status. These variables are not available for home equity lines of credit (separately) and student loans, so we use variables for the number of trades ever 90 or more days delinquent or derogatory to measure transition into severe delinquency.

One problem that arises from directly comparing across these measures is that the former does not allow for measurement along the intensive margin but is defined at a higher frequency of within the last 6 months, whereas the latter does allow for measurement along the intensive margin but is defined at a lower frequency of within the last 24 months. Thus, we restrict my analysis to the first transition into severe delinquency or derogatory status along the extensive margin on each loan type. This makes practical sense as well because the first event of this type has persistent effects on the financial life of borrowers, remaining on a credit file for up to seven years.

The non-pecuniary penalties, such as repossession, foreclosure, and loss of credit access, for loan types other than federal student loans typically begin at 90 or more days delinquency. However, delinquent federal student loans enter technical default at 271 days of delinquency and after 420 days of delinquency the loan is reported as defaulted, or in derogatory status, to the credit bureaus and is placed with a collection agency. Penalties on student loans include but are not limited to wage garnishment and treasury offset.[1]

A variable for presence of outstanding governmental agency debts allows us to measure these delayed penalties associated with federal student loan default. Table B.1.1 shows how federal student loan delinquency, default, and placement with a collection agency maps into the quarterly frequency of the sample. Without delays in reporting, a federal student loan that has transitioned into severe delinquency will enter default and be placed with a collector (and transferred to a collections trade) three to four quarters after the event.

Table B.1.1: Severe Delinquency, Default, and Collections Process for Student Loans

Timeline of Status Mapped into Quarterly Reporting				
Quarter After Event	90+ DPD or Derogatory Flag	Balance on Open Accounts	DPD Range	Status
-1	0	x	0-90	Current/Not Reported
0	1	x	91-181	Severe Delinquency
1	1	x	182-272	Severe Delinquency or in Technical Default
2	1	x	273-363	in Technical Default
3	1	$< x$	364-454	in Technical Default or Default and in Collections
4	1	$< x$	455-545	Default and in Collections

The variables we use to measure transition into severe delinquency for student loans does not distinguish between delinquency and default. However, the most severe penalties for non-payment

[1] See the U.S. Department of Education's Office of Federal Student Aid web page Collections.

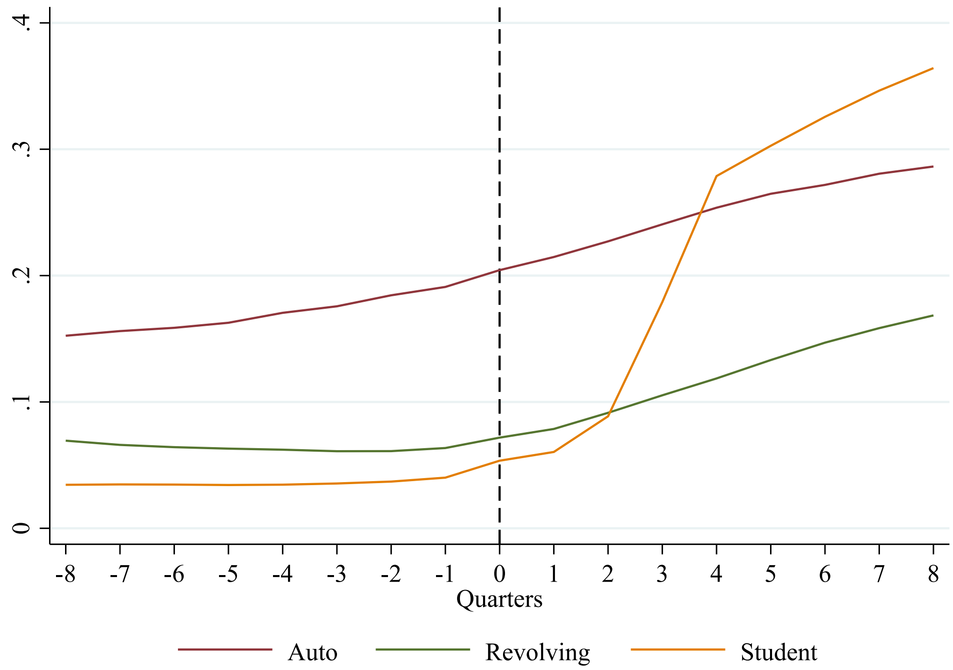

Figure B.1.2: Presence of Outstanding Governmental Agency Debt

on federal student loans are associated with default. Figure B.1.2 shows the fraction of borrowers that continue to miss payments and likely transition from severe delinquency to default. It shows that approximately 30 percent of student loan borrowers that transition into severe delinquency face the severe penalties associated with federal student loan default. Although this measure is a proxy since it also includes borrowers with outstanding governmental agency debts other than defaulted student loan, such as tax lien and unpaid child/family support, it is likely a strong proxy based on the subsample being used.

Figure B.1.3 shows that the timing of the amount of debt past due on severely delinquent student loans matches up well with the timeline in Table B.1.1. It suggests that within five quarters of student loan borrowers transitioning into severe delinquency, that borrowers either arrange an alternate repayment plan to cure the delinquency or do not and are transferred to the collections system as depicted by the partial decrease of the amount past due not in collections.

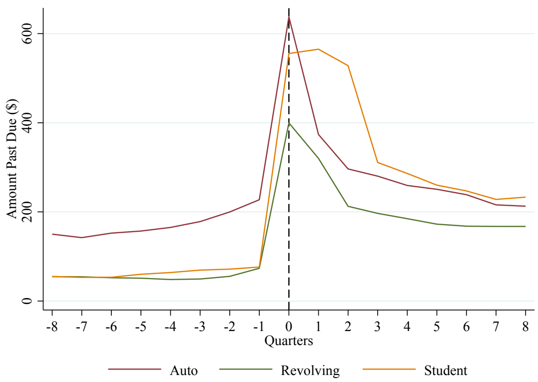

(a) 90 to 180 DPD

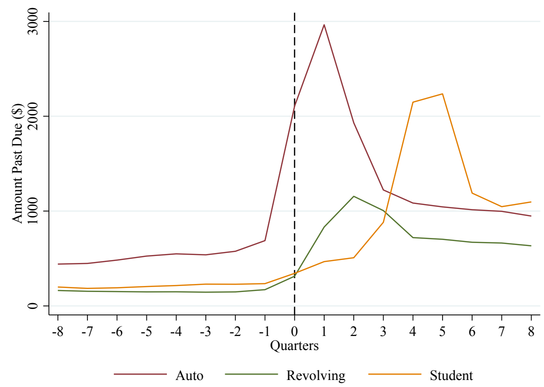

(b) Derogatory Excluding Collections

Figure B.1.3: Amount Past Due

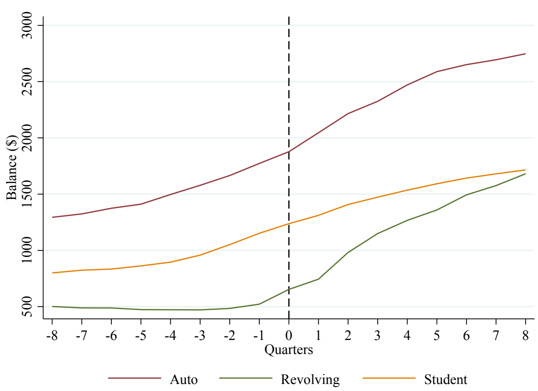

(a) In Collections

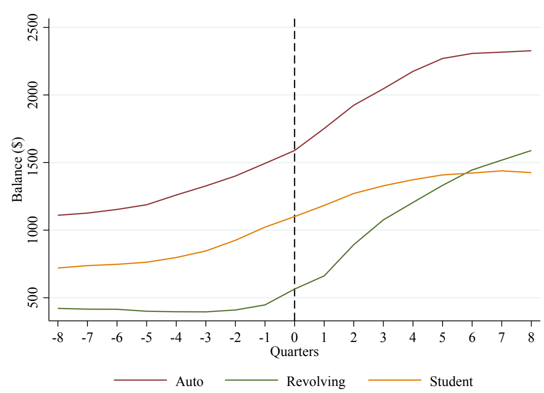

(b) Placed with Collectors

Figure B.1.4: Balance on Collections Accounts

B.2 Additional Analysis

B.2.1 Life Cycle Differences

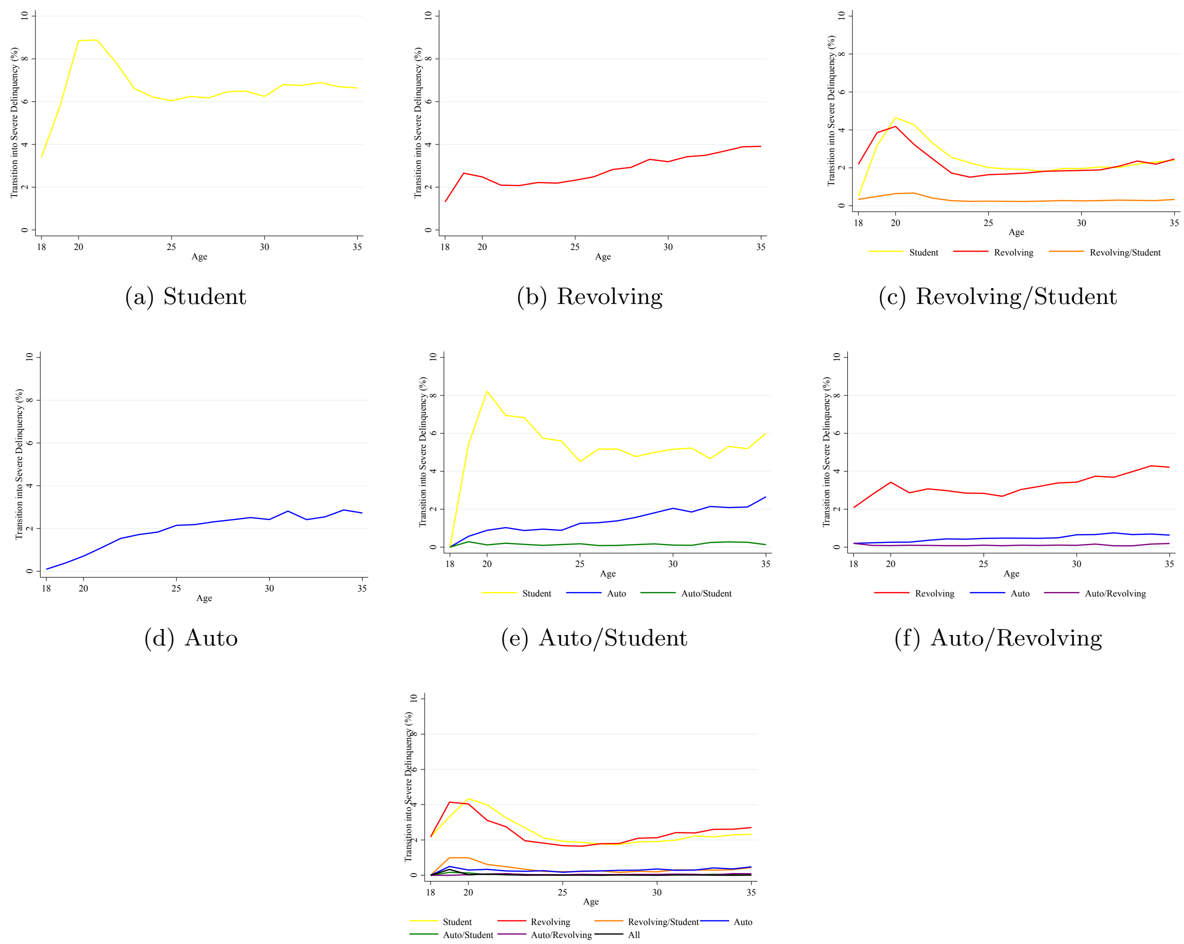

(a) Student (b) Revolving (c) Revolving/Student

(d) Auto (e) Auto/Student (f) Auto/Revolving

(g) All

Figure B.2.1: Transition into Severe Delinquency along Extensive Margin

Note: Panels represent the debt portfolio held and in repayment for borrowers (the lines) that do not make the minimum payment for over 90 days.

The age profiles for transitioning into severe delinquency suggest that there is unobserved selection into holding different debt portfolios at different ages. Except for the youngest student loan borrowers in repayment, borrowers that hold only one debt type transition into severe delinquency at higher rates than the youngest borrowers that hold the same debt portfolio.[2]

[2]The peaks of student loan default are reassuring as these borrowers are negatively selected, they are likely no

B.2.2 Summary Statistics

Table B.2.1: Summary Statistics One Quarter Before Event by Debt Type

Debt	Count	Age	Qs_Since_Orig	Bal	BID	Pymt	Cred	Inc	DTI	DIDTI	PTI
Auto	11885	27	18.90	22806	28678	592	503	37	0.650	0.870	0.210
Auto		4.200	12.70	22500	27690	432	66	16	0.510	0.850	0.110
Revolving	36255	24.60	11.40	14473	22044	352	525	34	0.470	0.870	0.140
Revolving		4.400	10.60	21024	27957	655	79	15	0.580	1.130	0.120
Student	42102	25.10	12.10	20884	24435	334	574	30	0.770	0.930	0.140
Student		4.100	11.20	26152	31410	495	77	15	0.840	1.130	0.130

Table B.2.2: Portfolio Composition One Quarter Before Event by Debt Type

Debt	None	STU	REV	REV_STU	AUT	AUT_STU	AUT_REV	ALL
Auto	1	1	1	2	24	28	15	29
Revolving	0	0	23	46	0	0	10	21
Student	2	43	1	31	0	7	0	15

B.2.3 Event Study: Credit Score

Credit scores for revolving credit severe delinquents begin to trend downwards approximately 6 quarters before the event, indicating that the credit score is correctly predicting that these borrowers' probabilities of incurring a severe delinquency in the next 24 months are increasing; however, for student loan severe delinquents the credit score is predicting no such risk.

We find that student loan severe delinquents' credit scores discontinuously decrease by 100 points from the pre-event quarter, whereas auto loan and revolving credit severe delinquents' credit scores decrease by 17 and 24 points from the pre-event quarter, respectively. This is likely due to a combination of reporting, since delinquent federal student loans are not reported to credit bureaus until 90 days delinquency, and differential borrower behavior such as higher revolving utilization

longer enrolled in a program, enrolled in a for-profit program, or are private student loan borrowers that must make payments while enrolled in a program.

rates. Over the full pre-event period, revolving credit and student loan severe delinquents lose approximately 100 points from their 8-quarter-before-event credit scores of approximately 575, so they are penalized similarly overall. This suggests that revolving credit and student loan severe delinquents are even more similar before the event than they first appear based on the summary statistics in Appendix Table B.2.1.

Over the full post-event period, revolving credit severe delinquents' credit scores increase by approximately 45 points compared to a higher increase of 60 points for student loan severe delinquents' credit scores. There are many factors that may cause this difference in credit score recovery, such as programs that are available for federal student loans, such as forbearance or alternative repayments plans, that are not available for auto loans and revolving credit or differences in borrower behavior and financial characteristics after the event, such as inquiring and opening new accounts, repayment, or revolving credit utilization.

B.2.4 Event Study: Debt to Payment Ratio

In particular, we show that approximately 70 percent of borrowers that transition into severe delinquency have all student loans in deferral 4 quarters before the event and approximately 60 percent of borrowers transition from having some or all student loans in deferral to none in deferral 4 quarters before the event. In short, a debt-to-income ratio that does not include deferred student loan debt likely underweights some student loan borrowers' true credit risk.[3] Additionally, we show that wage income steadily increases for borrowers that transition into severe delinquency on auto loans and revolving credit in the eight-quarter-before-event period and adjusts afterwards, but increases sharply for borrowers that transition into severe delinquency on student loans in the four-quarter-before-event period and does not adjust afterwards. The coevolution of balance and income are then reflected in the evolution of the change in the log-debt-to-income ratios.

Figure 15 suggests that the large increases in the debt-to-income ratios relative to the pre-event quarter for borrowers that transition into severe delinquency on revolving credit is primarily driven

[3] Although we are not privy to the exact details of the VantageScore 3.0 credit scoring model, we impute the total balance variable, including deferred student loans, to compute the debt-to-income ratio including deferred student loans. Combined with the fact that credit score is constant on average before the first transition into severe delinquency, it appears that the model does not include a predicted value for the amount of time until a borrower enters repayment. If true, the model would likely fare better at predicting student loan default by including a predicted value for the amount of time until a borrower enters repayment.

by a sharp increase in the utilization rate instead of deferred student loans entering repayment, which we show in the Appendix. In contrast, the wave patterns for the debt-to-income ratio including deferred student loans for borrowers that transition into severe delinquency on students loans is primarily driven by a continuous increase of income starting four quarters before the event and for borrowers that transition into severe delinquency on auto loans is driven by the combination of the income estimate adjusting downwards and uncured delinquencies transferring to collections after the event.

B.2.5 Payment-to-Income Ratio

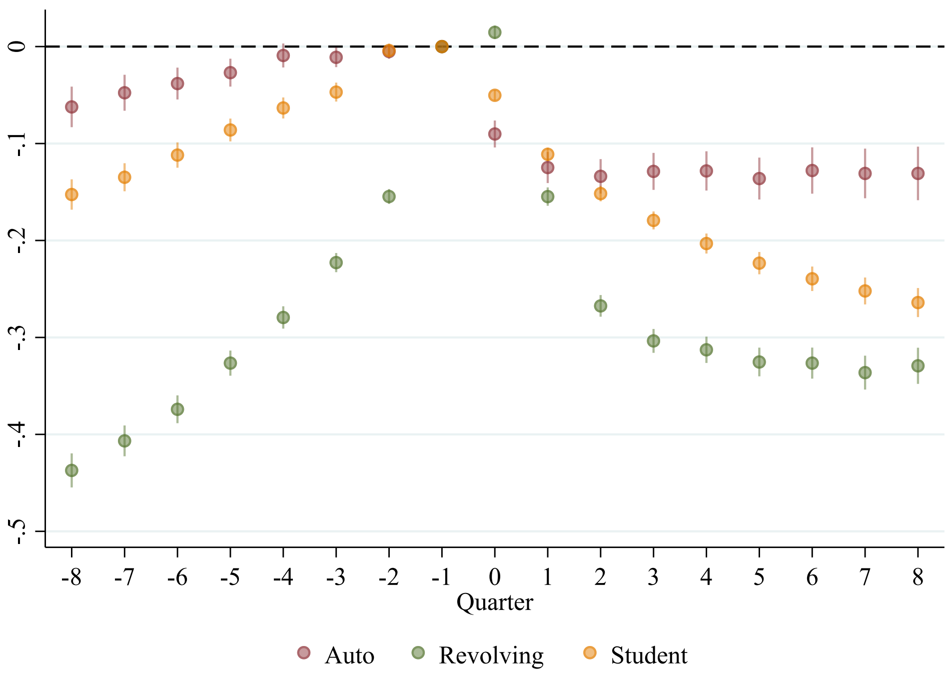

Figure B.2.2: Overlaid Coefficient Plot: Payment-to-Income Ratios over All Open Accounts

Note: The pre-event quarter coefficient is zero for all debt types. Regressions include controls for age, the debt portfolio held, and whether the debt is in repayment status. The ratio is transformed using the natural logarithm so that the y-axis is interpreted as approximately $(\exp(x) - 1) \times 100$ percent change from the pre-event quarter.

Although on average the minimum required monthly payment to remain current on debts scale proportionally with balance, non-standard repayment plans such as income-based repayment options available to federal student loan borrowers tie payment to income instead of debt balance. In the data, we do not observe the repayment plan of the borrower, so it is difficult to systematically identify student loan borrowers with and without standard repayment plans.[4]

Figure B.2.2 shows the evolution of the (total) payment-to-income ratio around the first transition into severe delinquency on auto loans, revolving credit, and student loans. The payment-to-income ratios increase in the eight-quarter-pre-event period for borrowers transitioning into severe delinquency on all three debt types; however, it increases most significantly for revolving credit due to the sharp increase in the utilization rate. In the pre-event quarter, payment-to-income ratios are relatively high at levels of 0.21, 0.14, and 0.14 for auto loan, revolving credit, and student loans, respectively. Although the sharp increase in the payment-to-income ratio appears to be reflected in the credit score of borrowers transitioning into severe delinquency on auto loans and revolving credit before the event, it is not reflected in the credit score of borrowers transitioning into severe delinquency on student loans.[5]

B.3 Survey Instrucionts

Click here for survey instructions.

[4] There are seven repayment plans: standard, graduated, income-contingent, income-based, pay as you earn (PAYE), revised pay as you earn (REPAYE), and alternative. Income-contingent repayment plans existed for Direct Stafford Loans since the 1990's, but were not widely available or taken-up. The income-based, PAYE, and REPAYE repayment plans were introduced in July 2009, July 2013, and December 2015, respectively. Current take-up for these newer, more widely available repayment plans is significant.

[5] A predicted value for the payment-to-income ratio for student loan borrowers can be estimated using amortization formulas and the information available in the credit file.

Bibliography

Abdellaoui, Mohammed, Enrico Diecidue, and Ayse Öncüler. 2011. "Risk preferences at different time periods: An experimental investigation." *Management Science* 57 (5): 975–987.

Abeler, Johannes, Armin Falk, Lorenz Goette, and David Huffman. 2011. "Reference Points and Effort Provision." *American Economic Review* 101 (2): 470–92. 10.1257/aer.101.2.470.

Acland, Dan, and Matthew R Levy. 2015. "Naiveté, projection bias, and habit formation in gym attendance." *Management Science* 61 (1): 146–160.

Ahlbrecht, Martin, and Martin Weber. 1997. "An empirical study on intertemporal decision making under risk." *Management Science* 43 (6): 813–826.

Akerlof, George, and Rachel Kranton. 2000. "Economics and Identity." *The Quarterly Journal of Economics* 115 715–753.

Ali, S Nageeb. 2011. "Learning self-control." *The Quarterly Journal of Economics* 126 (2): 857–893.

Allen, Eric J, Patricia M Dechow, Devin G Pope, and George Wu. 2017. "Reference-dependent preferences: Evidence from marathon runners." *Management Science* 63 (6): 1657–1672.

Andreoni, James, Michael A Kuhn, and Charles Sprenger. 2015. "Measuring time preferences: A comparison of experimental methods." *Journal of Economic Behavior & Organization* 116 451–464.

Andreoni, James, Justin M Rao, and Hannah Trachtman. 2017. "Avoiding the Ask: A Field Experiment on Altruism, Empathy, and Charitable Giving." *Journal of Political Economy* 125 (3): 625–653.

Andreoni, James, and Charles Sprenger. 2012a. "Risk Preferences Are Not Time Preferences." *American Economic Review* 102 (7): 3357–3376.

Andreoni, James, and Charles Sprenger. 2012b. "Risk Preferences Are Not Time Preferences." *American Economic Review* 102 (7): 3357–76. 10.1257/aer.102.7.3357.

Andries, Marianne, and Valentin Haddad. 2020. "Information aversion." *Journal of Political Economy* 128 (5): 1901–1939.

Ariely, Dan, and Klaus Wertenbroch. 2002a. "Procrastination, Deadlines, and Performance: Self-Control by Precommitment." *Psychological Science* 13 (3): 219–224.

Ariely, Dan, and Klaus Wertenbroch. 2002b. "Procrastination, deadlines, and performance: Self-control by precommitment." *Psychological science* 13 (3): 219–224.

Aristotle, Robert Bartlett, and Susan Collins. 2011 Edition. *Aristotle's Nicomachean Ethics.* University of Chicago's Press.

Ashraf, Nava, Dean Karlan, and Wesley Yin. 2006. "Tying Odysseus to the Mast: Evidence From a Commitment Savings Product in the Philippines." *The Quarterly Journal of Economics* 121 (2): 635–672.

Augenblick, Ned, Muriel Niederle, and Charles Sprenger. 2015. "Working over Time: Dynamic Inconsistency in Real Effort Tasks." *The Quarterly Journal of Economics* 130 (3): 1067–1115.

Augenblick, Ned, and Matthew Rabin. 2018. "An Experiment on Time Preference and Misprediction in Unpleasant Tasks." *The Review of Economic Studies* 86 (3): 941–975.

Bartling, Björn, Florian Engl, and Roberto Weber. 2014. "Does Willful Ignorance Deflect Punishment? - An Experimental Study." *European Economic Review* 70 512–524. https://doi.org/10.1016/j.euroecorev.2014.06.016.

Bénabou, Roland, and Jean Tirole. 2004a. "Willpower and Personal Rules." *Journal of Political Economy* 112 (4): 848–886.

Bénabou, Roland, and Jean Tirole. 2004b. "Willpower and personal rules." *Journal of Political Economy* 112 (4): 848–886.

Bénabou, Roland, and Jean Tirole. 2006. "Incentives and prosocial behavior." *American economic review* 96 (5): 1652–1678.

Bénabou, Roland, and Jean Tirole. 2011. "Identity, Morals, and Taboos: Beliefs as Assets." *The Quarterly Journal of Economics* 126 (2): 805–855.

Bénabou, Roland, and Jean Tirole. 2016. "Mindful Economics: The Production, Consumption, and Value of Beliefs." *Journal of Economic Perspectives* 30 (3): 141–164.

Bicchieri, Cristina, Eugen Dimant, and Silvia Sonderegger. 2022. "It's Not A Lie if You Believe the Norm Does Not Apply: Conditional Norm-Following and Belief Distortion."Technical report, Working Paper. 10.2139/ssrn.3326146.

Bisin, Alberto, and Kyle Hyndman. 2020. "Present-bias, procrastination and deadlines in a field experiment." *Games and economic behavior* 119 339–357.

Breig, Zachary, Matthew Gibson, and Jeffrey Shrader. 2021. "Why do we procrastinate? present bias and optimism."Technical report, R&R The Economics Journal.

Bryan, Gharad, Dean Karlan, and Scott Nelson. 2010. "Commitment devices." *Annu. Rev. Econ.* 2 (1): 671–698.

Burger, Nicholas, Gary Charness, and John Lynham. 2011. “Field and online experiments on self-control.” *Journal of Economic Behavior & Organization* 77 (3): 393–404.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. “Labor supply of New York City cabdrivers: One day at a time.” *The Quarterly Journal of Economics* 112 (2): 407–441.

Carroll, Gabriel D, James J Choi, David Laibson, Brigitte C Madrian, and Andrew Metrick. 2009. “Optimal Defaults and Active Decisions.” *The Quarterly Journal of Economics* 124 (4): 1639–1674.

Carvalho, Leandro S, Stephan Meier, and Stephanie W Wang. 2016a. “Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday.” *American Economic Review* 106 (2): 260 – 284.

Carvalho, Leandro S., Stephan Meier, and Stephanie W. Wang. 2016b. “Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday.” *Am. Econ. Rev.* 106 (2): 260–84. 10.1257/aer.20140481.

Clark, Damon, David Gill, Victoria Prowse, and Mark Rush. 2020. “Using goals to motivate college students: Theory and evidence from field experiments.” *Review of Economics and Statistics* 102 (4): 648–663.

Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White. 2020. “Measuring Time Preferences.” *Journal of Economic Literature* 58 (2): 299–347.

Corgnet, Brice, Joaquín Gómez-Miñambres, and Roberto Hernán-Gonzalez. 2015. “Goal setting and monetary incentives: When large stakes are not enough.” *Management Science* 61 (12): 2926–2944.

Corgnet, Brice, Joaquín Gómez-Miñambres, and Roberto Hernán-Gonzalez. 2018. “Goal setting in the principal–agent model: Weak incentives for strong performance.” *Games and Economic Behavior* 109 311–326.

Dana, Jason, Roberto Weber, and Jason Xi Kuang. 2007. “Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness.” *Economic Theory* 33 (1): 67–80.

DellaVigna, Stefano, John A List, and Ulrike Malmendier. 2012. “Testing for Altruism and Social Pressure in Charitable Giving.” *The Quarterly Journal of Economics* 127 (1): 1–56.

DellaVigna, Stefano, and Ulrike Malmendier. 2004. “Contract design and self-control: Theory and evidence.” *The Quarterly Journal of Economics* 119 (2): 353–402.

DellaVigna, Stefano, and Ulrike Malmendier. 2006. “Paying Not to Go to the Gym.” *American Economic Review* 96 (3): 694–719. 10.1257/aer.96.3.694.

DiTella, Rafael, Ricardo Perez-Truglia, Andres Babino, and Mariano Sigman. 2015. "Conveniently Upset: Avoiding Altruism by Distorting Beliefs about Others' Altruism." *American Economic Review* 105 (11): 3416–42.

Drucker, Luca, and M Kaufmann. 2022. "Excuse-Driven Present Bias." working paper, Centeral European University.

Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2011. "Nudging Farmers to use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review* 101 (6): 2350–90.

Dynan, Karen E. 2000. "Habit Formation in Consumer Preferences: Evidence from Panel Data." *American Economic Review* 90 (3): 391–406. 10.1257/aer.90.3.391.

Eliaz, Kfir, and Andrew Schotter. 2010. "Paying for confidence: An experimental study of the demand for non-instrumental information." *Games and Economic Behavior* 70 (2): 304–324. https://doi.org/10.1016/j.geb.2010.01.006.

Ericson, Keith. 2017. "On the Interaction of Memory and Procrastination: Implications for Reminders, Deadlines, and Empirical Estimation." *Journal of the European Economic Association* 15 (3): 692–719.

Ericson, Keith, and David Laibson. 2019. "Intertemporal Choice." In *Handbook of Behavioral Economics: Applications and Foundations 1*, Volume 2. 1–67, Elsevier.

Exley, Christine. 2015. "Excusing Selfishness in Charitable Giving: The Role of Risk." *The Review of Economic Studies* 83 (2): 587–628.

Exley, Christine, and Judd Kessler. 2019. "Motivated Errors."Technical report, Working Paper. 10.3386/w26595.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman, and Uwe Sunde. 2018. "Global Evidence on Economic Preferences." *The Quarterly Journal of Economics* 133 (4): 1645–1692.

Falk, Armin, and Michael Kosfeld. 2006. "The Hidden Costs of Control." *American Economic Review* 96 (5): 1611–1630. 10.1257/aer.96.5.1611.

Farber, Henry 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, Urls. 2007. "z-Tree: Zurich toolbox for ready-made economic experiments." *Experimental Economics* 10 171–178.

Frederick, Shane, George Loewenstein, and Ted O'donoghue. 2002. "Time discounting and time preference: A critical review." *Journal of economic literature* 40 (2): 351–401.

Fudenberg, Drew, and David K Levine. 2006. “A Dual-Self Model of Impulse Control.” *American Economic Review* 96 (5): 1449–1476.

Ganguly, Ananda, and Joshua Tasoff. 2017. “Fantasy and Dread: The Demand for Information and the Consumption Utility of the Future.” *Management Science* 63 (12): 4037–4060.

Giné, Xavier, Dean Karlan, and Jonathan Zinman. 2010. “Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation.” *American Economic Journal: Applied Economics* 2 (4): 213 – 235.

Gino, Francesca, Michael Norton, and Roberto Weber. 2016. “Motivated Bayesians: Feeling Moral While Acting Egoistically.” *Journal of Economic Perspective* 30 (3): 189–212. 10.1257/jep.30.3.189.

Gneezy, Uri, Silvia Saccardo, and Roel van Veldhuizen. 2019. “Bribery: Behavioral Drivers of Distorted Decisions.” *Journal of the European Economic Association* 17 (3): 917–946. 10.1093/jeea/jvy043.

Golman, Russell, David Hagmann, and George Loewenstein. 2017. “Information Avoidance.” *Journal of Economic Literature* 55 (1): 96–135.

Grossman, Zachary, and Joël van der Weele. 2016. “Self-Image and Willful Ignorance in Social Decisions.” *Journal of the European Economic Association* 15 (1): 173–217.

Gul, Faruk, and Wolfgang Pesendorfer. 2001. “Temptation and Self-Control.” *Econometrica* 69 (6): 1403–1435.

Harding, Matthew, and Alice Hsiaw. 2014. “Goal setting and energy conservation.” *Journal of Economic Behavior & Organization* 107 209–227.

Havrankey, T, M Rusnak, and A Sokolava. 2017. “Habit formation in consumption: A meta-analysis.” *European Economic Review* 95 142–167.

Heath, C. 1999. “Goals as Reference Points.” *Cognitive Psychology* 38 (1): 79–109.

Heidhues, Paul, and Botond Köszegi. 2010. “Expoiting Naïvet/’e about Self-Control in the Credit Market.” *American Economic Review* 100 2279–2303.

Heidhues, Paul, and Philipp Strack. 2021. “Identifying Present Bias from the Timing of Choices.” *American Economic Review* 111 (8): 2594–2622. 10.1257/aer.20191258.

Himmler, Oliver, Robert Jäckle, and Philipp Weinschenk. 2019. “Soft Commitments, Reminders, and Academic Performance.” *American Economic Journal: Applied Economics* 11 (2): 114–42. 10.1257/app.20170288.

Hsiaw, A. 2018. “Goal bracketing and self-control.” *Games and Economic Behavior* 111 100–121.

Hsiaw, Alice. 2013. “Goal-setting and self-control.” *Journal of Economic Theory* 148 (2): 601–626.

John, Anett. 2019. "When Commitment Fails: Evidence from a Field Experiment." *Management Science* 66 (2): 503–529, https://pubsonline.informs.org/doi/abs/10.1287/mnsc.2018.3236.

Kaiser, Caspar, and Andrew J Oswald. 2022. "The Scientific Value of Numerical Measures of Human Feelings." *Proceedings of the National Academy of Sciences* 119 (42): e2210412119.

Karlsson, Niklas, George Loewenstein, and Duane Seppi. 2009. "The Ostrich Effect: Selective Attention to Information." *Journal of Risk and Uncertainty* 38 (2): 95–115.

Kaur, Supreet, Michael Kremer, and Mullainathan Sendhil. 2015. "Self-Control at Work." *Journal of Political Economy* 6 1227 – 1277.

Koch, Alexander K., and Julia Nafziger. 2016. "Goals and bracketing under mental accounting." *Journal of Economic Theory* 162 305–351. https://doi.org/10.1016/j.jet.2016.01.001.

Koch, Alexander K., and Julia Nafziger. 2020. "Motivational goal bracketing: An experiment." *Journal of Economic Theory* 185 104949. https://doi.org/10.1016/j.jet.2019.104949.

Köszegi, Botond. 2006. "Ego Utility, Overconfidence, and Task Choice." *Journal of the European Economic Association* 4 (4): 673–707.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2): 443–477.

LeYaouanq, Yves, and Peter Schwardmann. 2022. "Learning about One's Self." *Journal of the European Economic Association* 20 (5): 1791–1828.

Linardi, Sera, and Margaret McConnell. 2011. "No Excuses for Good Behavior: Volunteering and the social environment." *Journal of Public Economics* 95 445 – 454.

Looney, Adam, and Constantine Yannelis. 2015. "A crisis in student loans?: How changes in the characteristics of borrowers and in the institutions they attended contributed to rising loan defaults." *Brookings Papers on Economic Activity* 2015 (2): 1–89.

Mijović-Prelec, Danica, and Drazen Prelec. 2010. "Self-Deception as Self-Signalling: a Model and Experimental Evidence." *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 365 (1538): 227. 10.1098/rstb.2009.0218.

O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later." *The American Economic Review* 89 (1): 103–124.

Oster, Emily, Ira Shoulson, and E Dorsey. 2013. "Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease." *American Economic Review* 103 (2): 804–30.

Plato, Allan Bloom, and Adam Kirsch. 2016 Edition. *Plato's The Republic.* New York Books.

Robinson, Carly D., Jana Gallus, Monica H. Lee, and Todd Rogers. 2019. "The demotivating effect (and unintended message) of awards." August, https://scholar.harvard.edu/

files/todd_rogers/files/the_demotivating_effect_and_unintended_message_of_awards_vf.pdf.

Royer, Heather, Mark Stehr, and Justin Sydnor. 2015. "Incentives, Commitments, and Habit Formation in Exercise: Evidence from a Field Experiment with Workers at a Fortune-500 Company." *American Economic Journal: Applied Economics* 7 (3): 51–84. 10.1257/app.20130327.

Saccardo, Silvia, and Marta Serra-Garcia. 2022. "Cognitive Flexibility or Moral Commitment? Evidence of Anticipated Belief Distortion."Technical report, Working Paper. 10.2139/ssrn.3676711.

Schildberg-Hörisch, Hannah. 2018. "Are risk preferences stable?" *Journal of Economic Perspectives* 32 (2): 135–54.

Serra-Garcia, Marta, and Nora Szech. 2022. "The (In)Elasticity of Moral Ignorance." *Management Science* 68 (7): 4815–4834.

Shefrin, H. M., and Richard Thaler. 1981. "An Economic Theory of Self-Control." *Journal of Political Economy* 89 392–406. 10.1086/260971.

Smith, Adam. 1759. *The Theory of Moral Sentiments.* Liberty Classics.

Stigler, George J. 1961. "The Economics of Information." *Journal of Political Economy* 69 (3): 213–225.

Strotz, Robert. 1955. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review Economic Studies* 23 (3): 165–180. 10.2307/2295722.

Tversky, Amos, and Daniel Kahneman. 1992. "Advances in prospect theory: Cumulative representation of uncertainty." *Journal of Risk and uncertainty* 5 (4): 297–323.

Yannelis, Constantine. 2020. "Strategic Default on Student Loans."Technical report, R&R Journal of Finance.

Zimmermann, Florian. 2020. "The Dynamics of Motivated Beliefs." *American Economic Review* 110 (2): 337–61. 10.1257/aer.20180728.